

Optimal Resource Allocation for Competitive Spreading Processes on Bilayer Networks

Nicholas J. Watkins [✉], Cameron Nowzari [✉], Victor M. Preciado [✉], and George J. Pappas [✉]

Abstract—We consider a competitive epidemic process in a bilayer network, and develop a framework to find an optimal allocation of control resources to eliminate one of the epidemics. We consider the SI_1SI_2S model, a recent generalization of the popular SIS model to the case of two competitive epidemics. We start our analysis by extending the standard SI_1SI_2S formulation with homogeneous parameters to a heterogeneous setting with edge-dependent infection rates, and node-dependent recovery rates. We then find necessary and sufficient conditions under which the mean-field approximation of a chosen epidemic process stabilizes to extinction exponentially quickly. Leveraging this result, we develop a framework for the solution of two optimization problems. In the first, we find an optimal allocation of control resources in order to eradicate the chosen epidemic at a minimum cost. In the second, we are given a fixed budget and propose a method which provably attains the extinction condition when sufficient capital is available, and otherwise mitigates the spread of the unwanted epidemic as much as possible. We explore the efficacy of our methods through extensive simulation.

Index Terms—Behavioral science, optimization, stochastic systems.

I. INTRODUCTION

MODELING, analysis, and control of spreading processes in complex networks has recently garnered significant attention from the research community. The potential applications for such methods are diverse: the spread of biological epidemics, social behaviors, and cybersecurity threats can all be formalized within this framework. Prior efforts have focused primarily on the case of single-layer spreading networks; however, such an abstraction is limited in modeling capacity. In principle, spreading over networks can take place through markedly different channels, which motivates the study of multilayer models.

This paper studies a multilayer, heterogeneous compartmental epidemic model, in which the spread of competing epidemics, such as beliefs and behaviors, can be modeled. We direct our

Manuscript received January 16, 2016; revised July 19, 2016; accepted August 23, 2016. Date of publication September 9, 2016; date of current version March 16, 2018. This work was supported in part by the National Science Foundation grant CNS-1302222, “NeTS: Medium: Collaborative Research: Optimal Communication for Faster Sensor Network Coordination,” IIS-1447470 “BIGDATA: Spectral Analysis and Control of Evolving Large Scale Networks,” and in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP), a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Recommended by Associate Editor Randy Freeman.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: nwatk@upenn.edu; cnowzari@upenn.edu; preciado@upenn.edu; pappasg@upenn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCNS.2016.2607838

attention to the problem of controlling a spreading process in order to quickly eliminate a chosen epidemic in a competitive environment. This is a natural concept in modeling several socially relevant problems. For example, we may use this model to study the effects of political strategies on the opinions of the populace, predict the ramifications of gossip in professional networks, and understand the influence of marketing strategies on consumer behavior.

Literature Review: Many well-known models of spreading processes in networks are developed for the case of a single contagion spreading over a single network layer; we refer the reader to [1]–[3] for an overview. Recent efforts have been made in extending this body of work to account for the possibility of competitive and/or coexistent processes on single-layer networks. Particular examples include investigations into the effects of multiple pathogens in a single-layer “Susceptible-Infected-Removed” (*SIR*) model [4]–[6], a study of an extension to the *SIR* model (*SICR*) for assessing the effects of competition and cooperation between pathogens spreading on a single network [7], and the development of a model for the spread of competing ideas using the “Susceptible-Infected-Susceptible” (*SIS*) model on scale-free networks [8].

A more recent trend is the investigation into systems with multiple pathogens *and* multiple spreading layers, in which each contagion spreads over a specified layer. An overview of this research area can be found in [9]. Particular examples of interest include an investigation into the effects of pathogen interaction on overlay networks with *SIR* dynamics [10], the development of a model in which disease awareness and infection spread on separate layers of *SIS* dynamics [11], [12], the development of a model (SI_1SI_2S) that generalizes the classic *SIS* model to a competitive multilayer framework [13], and work to find conditions under which processes in the SI_1SI_2S model can coexist [14].

We concern ourselves with the design of an optimization framework for allocating resources to achieve an optimal cost network design. Similar problems have been studied for controlling the single layer *SIS* model in [15], and a noncompetitive multilayer model in [16]. The work we present here is the first to consider an allocation problem which leverages interprocess competition, which we incorporate by studying a variant of the SI_1SI_2S process.

Statement of Contributions: We develop a computational framework for determining resource allocations which realize an optimal-cost network which controls the SI_1SI_2S process presented in [13] and [14] to a desired equilibrium. More

specifically, we begin by introducing a heterogeneous version of the model with edge-dependent spreading rates and node-dependent recovery rates in order to enable us to capture the effects of asymmetric influence among agents. We leverage this added flexibility to design networks which exploit interprocess competition in eliminating a chosen epidemic. This equilibrium concept is useful in modeling various situations in which competitive epidemics may occur, such as a marketing firm wanting to influence their customer base in order to eliminate competitors. Our technical contributions evolve from addressing this task, and address several control-theoretic facets of the mean-field SI_1SI_2S model left previously unexplored.

Organization: The remainder is organized as follows. In Section III, we determine necessary and sufficient conditions for exponentially stabilizing the desired equilibrium of the mean-field model. In Section IV, we formulate an optimal resource allocation problem in which costs can be paid to change the parameters of the model. We then develop two tractable optimization methods. Our first computes a minimal-cost resource allocation which attains the desired equilibrium. Our second addresses the situation in which a budget is specified, and we aim to mitigate the prevalence of the unwanted epidemic process when the available budget is not sufficient for realizing the desired equilibrium. In Section V, we explore the efficacy of the mean-field control policies developed against the stochastic process behavior through extensive Monte Carlo simulations. With respect to the preliminary work presented in [17], this paper extends the results pertaining to the effects of competition, provides proofs of our main results, and adds significant simulations comparing the mean-field model to the stochastic SI_1SI_2S process.

A. NOTATION AND BACKGROUND

Let \mathbb{R} , $\mathbb{R}_{\geq 0}$, and $\mathbb{R}_{>0}$ denote the set of real, non-negative real numbers, and positive real numbers, respectively. We use the notation $\vec{x} \in \mathbb{R}^n$ to denote an n -dimensional column vector, and \vec{x}^T to denote its transpose, both with components $x_i \in \mathbb{R}$. Fix a probability space (Ω, \mathcal{F}, P) , and let $X : \Omega \mapsto \mathbb{R}$ be a random variable; we denote its expectation by $E[X]$. We use $|S|$ to denote the cardinality of a finite set.

We say a matrix A is irreducible if no similarity transformation exists which places A into block upper-triangular form. We denote by $\text{diag}(\vec{a})$ a matrix with entries $\text{diag}(\vec{a})_{ii} = a_i$ for all i and 0 elsewhere. We use the notation $\lambda_{\max}(A)$ to denote the maximum taken over the real parts of the eigenvalues of a matrix A , i.e., $\lambda_{\max}(A) = \max_i \{\Re(\lambda_i(A))\}$. We call a matrix A such that $\lambda_{\max}(A) < 0$ Hurwitz, or stable. We will make repeated use of the Perron-Frobenius Lemma, stated as follows:

Proposition 1 (Perron-Frobenius [18]): Let A be a non-negative, irreducible matrix. Then, there exists a vector \vec{u} such that $u_i > 0$ for all i , and $Au = \lambda^*u$, where $\lambda^* > 0$ is the eigenvalue of A with the maximum absolute value, that is, the leading eigenvalue.

Graph Theory: A *directed graph* (digraph) is given by a triplet $G = (V, E, A)$ in which V is the set of vertices, $E \subseteq V \times V$ the ordered set of edges, and $A \in \{0, 1\}^{|V| \times |V|}$ the ad-

jacency matrix, that is, $a_{ij} = 1$ if and only if there exists an edge $(i, j) \in E$ connecting node i to node j . We define the set of in-neighbors of node i given the adjacency matrix A as $\mathcal{N}_i^{\text{in}} = \{j \in V \mid a_{ji} = 1\}$.

A path p is given by an ordered set of vertices $p = (v_1, v_2, \dots, v_m)$ such that (v_k, v_{k+1}) is an edge in E for all $k \in \{1, 2, \dots, m-1\}$. We say that some path p connects node v_i and v_j if the path starts at node v_i and ends at node v_j . We say a digraph is strongly connected if there exists a path connecting node v_i to node v_j for all $v_i, v_j \in V$. The adjacency matrix of a strongly connected digraph is irreducible.

A *bilayer graph* is a collection of two graphs $G = (G_A, G_B)$ which satisfy the following property: the vertex set V and edge set E of G are such that $V = V^A \cup V^B$, and $E = E^A \cup E^B$, where V^A and V^B are the vertex sets of G_A and G_B , respectively, and E^A and E^B are the edge sets of G_A and G_B , respectively. Note that the components G_A and G_B of G define separate layers, and so allow for the specification of spreading topologies for different phenomenon in a precise, compact notation.

Geometric Programming: A function $f : \mathbb{R}_{>0}^n \rightarrow \mathbb{R}$ is called a *monomial* if it can be written in the form $f(\vec{x}) = c x_1^{r_1} x_2^{r_2} \dots x_n^{r_n}$, where $c > 0$ is used to denote a leading constant, the r_i terms represent constant powers to which the arguments are raised, and the x_i terms represent f 's arguments. A function is said to be a *posynomial* if it can be written as a sum of monomials. Geometric programs form a class of quasiconvex optimization problems which have posynomial objective functions, posynomial inequality constraints, and monomial equality constraints.

Geometric programs can be transformed into convex optimization problems by performing a logarithmic change of variables and a logarithmic transformation of the objective and constraint functions. For further details on geometric programs and their solution, we refer the reader to [19] and [20].

To ease the formal statement of some of our results, we will introduce the notion of a posynomial transformation:

Lemma 1 (Posynomial Transformations): Any function $f(x)$ of the form $f(x) = \sum_k c_k(\hat{x} - x)^{p_k}$ with domain $(0, \hat{x})$ with $\hat{x} > 0$, $c_k > 0$, and $p_k \in \mathbb{R}$ can be written as a posynomial function of a new variable $z = \hat{x} - x$ defined on the domain $(0, \hat{x})$.

Proof: Consider the variable substitution $z = \hat{x} - x$. Then, we may write the posynomial transformation $\hat{f}(\hat{x} - x) = \sum_k c_k(z)^{p_k}$, where we see that a value $z = 0 \mapsto x = \hat{x}$ and $z = \hat{x} \mapsto x = 0$. Since the transformation is continuous, the domain of \hat{f} is $(0, \hat{x})$, as specified by the hypothesis. ■

We will denote the class of functions with domain $(0, d)$ which admit a posynomial transformation in the sense of Lemma 1 by $\mathcal{P}(0, d)$. This class of functions will appear repeatedly in the remaining sections.

II. MODEL AND PROBLEM STATEMENT

We begin our technical discussion by extending the SI_1SI_2S model proposed in [13] and analyzed further in [14]. Our primary contribution in extending this model is to allow the processes to be influenced by heterogeneous parameters, and

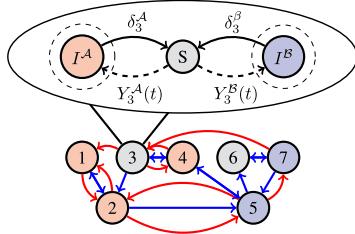


Fig. 1. A diagram of the SI_1SI_2S process, with the spreading graph for \mathcal{A} given by red edges, and the spreading graph for \mathcal{B} given by blue edges. The transition process for node 3 is explicitly illustrated, where we note that node 3 is a member of both spreading graphs, and so may have transitions to I^A and I^B .

allowing for the graph layers to be strongly connected digraphs with arbitrary node sets. This extends the work in [14], which assumes homogeneous spreading parameters and undirected layers with identical node sets. The extension allows the possibility of modeling asymmetric influence and nodal immunity, and is required to formulate the resource allocation problem.

We consider the spread of epidemics \mathcal{A} and \mathcal{B} over a bi-layer graph $G = (G_A, G_B)$, where \mathcal{A} spreads over $G_A = (V^A, E^A, A)$, \mathcal{B} spreads over $G_B = (V^B, E^B, B)$, and $|V| = n$. At any time t , we assume that each node can belong to one of three *compartments*: I^A if the node is infected by epidemic \mathcal{A} , I^B if the node is infected by epidemic \mathcal{B} , and S if the node is infected by neither. We let X_i^A , X_i^B , and X_i^S denote indicator functions corresponding to the compartments I^A , I^B , and S , respectively. We define $X_i^A(t) = 1$ if node i is in compartment I^A at time t and $X_i^A(t) = 0$ otherwise. We define X_i^B and X_i^S similarly.

We model the spread of \mathcal{A} and \mathcal{B} as a Markovian contact process in which a node i in compartment S transitions to I^A whenever it is a contacted by a node j in compartment I^A , with similar considerations holding for transitions from S to I^B . We assume all of the contact processes are stochastically independent, and occur at rates β_{ji}^A for the transitions from S to I^A and β_{ji}^B for the transitions from S to I^B , which we refer to as *spreading rates*. From this description, it then follows that the process which transitions node i from compartment S to compartment I^A is a Poisson process with rate $Y_i^A(t) = \sum_{j \in \mathcal{N}_i^A \text{ in}} \beta_{ji}^A X_j^A(t)$, and the process which transitions node i from compartment S to compartment I^B is a Poisson process with rate $Y_i^B(t) = \sum_{j \in \mathcal{N}_i^B \text{ in}} \beta_{ji}^B X_j^B(t)$. The processes which transition a node i from I^A to S and from I^B to S are Poisson processes with rates δ_i^A and δ_i^B , which we refer to as *healing rates*. A compartmental diagram of the process model is illustrated in Fig. 1.

For a general instance of the SI_1SI_2S process, studying the exact dynamics would require the enumeration of a Markov process with $O(3^n)$ states, arising from the need to explicitly account for all permissible combinations of compartmental memberships allowed by the instance of the problem. There are at least two methods of dealing with this complexity: 1) restricting considerations to simple graph topologies and 2) approximating the dynamics by a lower-dimensional system. Since our goal is to design resource allocations on graphs with arbitrary graph structures, we consider here a mean-field approximation of the

process, which reduces the dimension of the system's state space to $O(2n)$.

To clearly demonstrate how we arrive at the mean-field dynamics and give insight as to what effects the enacted approximations make, we first consider the exact equations of the expectation of the process dynamics

$$\begin{aligned} \frac{dE[X_i^A]}{dt} &= E \left[(1 - X_i^A - X_i^B) \sum_{j \in \mathcal{N}_i^A \text{ in}} \beta_{ji}^A X_j^A - \delta_i^A X_i^A \right], \\ \frac{dE[X_i^B]}{dt} &= E \left[(1 - X_i^A - X_i^B) \sum_{j \in \mathcal{N}_i^B \text{ in}} \beta_{ji}^B X_j^B - \delta_i^B X_i^B \right], \end{aligned} \quad (1)$$

where we have used the substitution $X_i^S = (1 - X_i^A - X_i^B)$ in order to reduce dimension. Note that the equations described by the system (1) are not closed: they contain terms of the form $E[X_i^A X_j^A]$ and $E[X_i^B X_j^B]$, which cannot be represented in terms of the dynamics of $E[X_i^A]$ and $E[X_i^B]$ without incurring error. However, without a closed set of equations, we cannot perform analysis, and so we make the approximations $E[X_i^A X_j^A] \approx \Phi_i^A \Phi_j^A$ and $E[X_i^B X_j^B] \approx \Phi_i^B \Phi_j^B$, where we have introduced the symbols Φ_i^A and Φ_i^B to denote the mean-field states approximating the probability that node i is in I^A , and the probability that node i is in I^B , respectively.

Using this substitution, we arrive at a mean-field approximation of SI_1SI_2S in the style of [21]:

$$\dot{\Phi}_i^A = (1 - \Phi_i^A - \Phi_i^B) \sum_{j \in \mathcal{N}_i^A \text{ in}} \beta_{ji}^A \Phi_j^A - \delta_i^A \Phi_i^A, \quad (2)$$

$$\dot{\Phi}_i^B = (1 - \Phi_i^A - \Phi_i^B) \sum_{j \in \mathcal{N}_i^B \text{ in}} \beta_{ji}^B \Phi_j^B - \delta_i^B \Phi_i^B, \quad (3)$$

where we note that the summation for the evolution of epidemic \mathcal{A} is indexed over the set of neighbors in G_A , and the evolution of epidemic \mathcal{B} is indexed over the set of neighbors in G_B in order to reflect the fact that \mathcal{A} spreads through layer G_A , and \mathcal{B} spreads through layer G_B .

We will more thoroughly examine the interrelation of the mean-field model and the stochastic process in Section V. However, the majority of our work will be guided by seeking answers to the following questions with respect to the mean-field model:

- Extinction: what conditions are sufficient to extinct a chosen process quickly?
- Optimal extinction: can we compute an optimal allocation of resources to attain a desired extinction quickly?
- Fixed budget mitigation: given a fixed budget, can we limit the prevalence of a desired process effectively?

Answers to these questions are of interest to the community of researchers currently engaged in the study of competitive epidemic spreading processes. As a particular example, we may consider a situation in which a firm would like to quell smear campaigns occurring on its network of customers in the most expedient and cost-effective manner possible. We may represent this within the framework of our model as a problem of finding conditions under which an unwanted process is driven out of existence as quickly and efficiently as possible. Our work

shows that computing an optimal cost network to realize this goal is feasible in the mean-field regime, and provides a step forward from the earlier works considering single-layer spreading processes.

III. EXTINCTION CONDITIONS

This section addresses the first of our stated problems, that is, finding conditions under which an unwanted epidemic extincts, or more concretely:

Problem 1 (Extinction): For a specified SI_1SI_2S spreading process on a bilayer graph $G = (G_A, G_B)$, determine conditions for the parameters of the subgraph G_A under which a chosen behavior \mathcal{A} extincts quickly.

In particular, we are concerned with stabilizing a mean-field equilibrium $\bar{\Phi} = [(\bar{\Phi}^A)^T, (\bar{\Phi}^B)^T]^T$ where $\bar{\Phi}_i^A$ and $\bar{\Phi}_i^B$ are the steady states of Φ_i^A and Φ_i^B , $\bar{\Phi}_i^A = 0$ for all i , and the values of $\bar{\Phi}_i^B$ are given by the solutions of the system

$$\frac{\bar{\Phi}_i^B}{(1 - \bar{\Phi}_i^B)} = \frac{1}{\delta_i^B} \sum_{j \in \mathcal{N}_i^B \text{ in}} \beta_{ji}^B \bar{\Phi}_j^B. \quad (4)$$

Note that the solution of (4) may be computed numerically by methods similar to those used in [22] for the SIS steady-state equations, and is unique due to the uniqueness of the SIS endemic equilibrium [23]. With the ability to claim knowledge of the values $\{\bar{\Phi}_i^B|_{\Phi_i^A=0}\}_{i \in V}$, we may now construct a result to Problem 1. In fact, we find necessary and sufficient conditions for the desired equilibrium to be exponentially stable:

Theorem 1 (Mean-Field Exponential Stability): For any SI_1SI_2S spreading process on a strongly connected bilayer graph G with mean field dynamics given by (2) and (3), the equilibrium $\bar{\Phi} = [(\bar{\Phi}^A)^T, (\bar{\Phi}^B)^T]^T$ with $\bar{\Phi}_i^A = 0$ for all i and $\bar{\Phi}^B$ given by the solutions of (4) is (locally) exponentially stable if and only if

$$J_{11} = \text{diag}(1 - \bar{\Phi}^B)(\beta^A)^T - \text{diag}(\vec{\delta}^A), \quad (5)$$

is Hurwitz, where $\vec{\delta}^A$ is the vector of \mathcal{A} 's recovery rates, and β^A is the matrix of \mathcal{A} 's spreading rates, which we assume to inherit \mathcal{A} 's sparsity pattern. ■

Proof: See Appendix A. ■

Remark 1 (Homogeneous Threshold): Note that this is similar to, but more general than, the stability results presented in [14]. In particular, the condition in [14] requires that all infection rates β_{ij}^A and recovery rates δ_i^A take on homogeneous values β and δ such that $\frac{\beta}{\delta} < \frac{1}{\lambda_{\max}(\text{diag}(1 - \bar{\Phi}^B)A)}$. By inspection, our result permits parameter choices which are excluded by this condition.

The form of the matrix we need to stabilize to guarantee the extinction of epidemic \mathcal{A} is similar to the matrix needed to guarantee extinction when we ignore competition. In particular, we note that a simple consequence of prior work on the SIS process (see, for example, [15], [24]) is that a sufficient condition for the exponentially fast elimination of the process spreading \mathcal{A} is that $\lambda_{\max}((\beta^A)^T - \text{diag}(\vec{\delta}^A)) < 0$ holds. By accounting for persistent competition among the epidemic processes, we might expect that our condition allows for more aggressive parameter selections. We will show that this is true in a rigorous sense

with our next result, which we will develop by first considering a technical lemma, and then specializing to our setting.

Lemma 2 (Row Compression Inequality): Let $M \in \mathbb{R}_{\geq 0}^{n \times n}$, $\vec{\kappa} \in [0, 1]^n$, and $\vec{\gamma} \in \mathbb{R}^n$. Then, the following inequality holds:

$$\lambda_{\max}(\text{diag}(\vec{\kappa})M - \text{diag}(\vec{\gamma})) \leq \lambda_{\max}(M - \text{diag}(\vec{\gamma})). \quad (6)$$

Proof: See Appendix B. ■

An immediate consequence of Lemma 2 is that competition in any particular node helps to prevent the persistence of an unwanted behavior. We make this formal as follows:

Proposition 2 (Benefits of Competition): Take any set of values $(\beta^A, \vec{\delta}^A)$ such that $\delta_i^A > 0$ for all i , and epidemic \mathcal{A} meets the SIS extinction condition, i.e., $\lambda_{\max}((\beta^A)^T - \text{diag}(\vec{\delta}^A)) < 0$. Then, for any realization of the SI_1SI_2S process, we must also have exponential elimination of \mathcal{A} . Moreover, if $\bar{\Phi}_i^B > 0$ for some i , then there exists some $\hat{\beta}^A$ with $\hat{\beta}_{ij}^A \geq \beta_{ij}^A$ for all i and j , and $\hat{\beta}_{ij}^A > \beta_{ij}^A$ for some i and j such that the set of parameters $(\hat{\beta}^A, \vec{\delta}^A)$ exponentially eliminates \mathcal{A} .

Proof: Exponential elimination of \mathcal{A} is a direct consequence of Theorem 1 when we apply Lemma 2 with $M = (\beta^A)^T$, $\vec{\kappa} = (1 - \bar{\Phi}^B)$ and $\vec{\gamma} = \vec{\delta}^A$. To prove the existence of a pair $(\hat{\beta}^A, \vec{\delta}^A)$ satisfying our claim, consider the matrix $\hat{\beta}^A$ with entries $\hat{\beta}_{ji}^A = \frac{1}{(1 - \bar{\Phi}_i^B)} \beta_{ji}^A$. Then

$$\text{diag}(1 - \bar{\Phi}^B)(\hat{\beta}^A)^T = (\beta^A)^T.$$

Since we only consider $\delta_i^A > 0$ for all i , we have that $\bar{\Phi}_i^B \in [0, 1)$ for all i . Hence, $\hat{\beta}_{ij}^A \geq \beta_{ij}^A$ for all i and j , where the inequality is strict for the case where $\bar{\Phi}_i^B > 0$. ■

Remark 2 (Benefits of Competition): Proposition 2 admits an explicit characterization of *how much* competition helps: for all agents $i \in V^A$, we can guarantee extinction even when the spreading rates associated with the incoming edges of node i are increased by a factor of up to $\frac{1}{(1 - \bar{\Phi}_i^B)}$ compared to the SIS case. While the quantitative utility of this observation depends on the particular cost functions in a given problem instance, this result qualitatively shows that the existence of a persistent spreading process \mathcal{B} that is competing with \mathcal{A} is *guaranteed* to make it easier to make process \mathcal{A} extinct quickly.

IV. OPTIMAL RESOURCE ALLOCATION

Having established conditions for the exponential stability of the desired equilibrium, we now focus our attention on establishing means for designing resource allocations which create networks with desirable control properties. We first consider the problem of designing a set of resource allocations in order to eliminate a chosen process at optimal cost when we are given functions which relate the chosen process parameter values to resource expenditures.

In the context of a marketing problem, we may think of spending on resources such as product giveaways, consumer incentive programs, advertisement campaigns, etc. designed to affect the perception of a company within a given market. To model this effect, we assume that for every designable parameter β_{ij}^A and δ_i^A , we are given cost functions f_{ij} and g_i which relate a desired parameter value to a capital expenditure, the particular charac-

teristics of which we assume to be application specific. With this notion developed, we may state our problem more formally as follows:

Problem 2 (Optimal Extinction): Consider an SI_1SI_2S spreading process on a bilayer graph $G = (G_A, G_B)$. Given sets of cost functions $\{f_{ij}\}_{(i,j) \in E^A}, \{g_i\}_{i \in V^A}$, determine a minimum cost allocation of resources to enforce the extinction conditions for the equilibrium of Problem 1.

From the discussion in Section III, we may formally cast Problem 2 as the following optimization program:

$$\begin{aligned} \text{minimize}_{\{\beta^A, \delta^A\}} \quad & \sum_{(i,j) \in E^A} f_{ij}(\beta_{ij}^A) + \sum_{i \in V^A} g_i(\delta_i^A) \\ \text{subject to} \quad & \lambda_{\max}(J_{11}(\beta^A, \delta^A)) < 0, \end{aligned} \quad (7)$$

where J_{11} is defined by (5). Note that (7) is nonconvex in general; it is an eigenvalue problem. However, if we allow ourselves to restrict considerations to a reasonable class of cost functions, we may develop a computational method for arriving at a solution tractably. Our work follows a similar line of development to that which was studied in [15] for single-layer, susceptible-infected-susceptible contagions, in which a convexification scheme was developed for an epidemic control resource allocation problem.

In particular, we will consider a method for transforming (7) into a convex problem when the cost functions are structured to make aggressive processes (i.e., those with higher spreading rates) cost more.

Theorem 2: Consider a realization of the dynamics (2)–(3) with an equilibrium point of the form $\bar{\Phi} = [(\bar{\Phi}^A)^T, (\bar{\Phi}^B)^T]^T$ with $\bar{\Phi}_i^A = 0$ for all i , and $\bar{\Phi}^B$ given by the solutions of (4). Define $z_i = (1 - \bar{\Phi}_i^B)$ for all $i \in V^B$, and consider any set of monotonically decreasing posynomial cost functions $\{f_{ij}\}_{\{(i,j) \in E^A\}}$, any set of functions $\{g_i \in \mathcal{P}(0, \hat{\delta}_i^A)\}_{i=1}^{|V^A|}$, and any $\epsilon \in (0, \min_i \{\hat{\delta}_i^A\})$. Then, an optimal solution of (7) can be computed by the solution of the following geometric program:

$$\begin{aligned} \text{minimize}_{\{\beta^A, \vec{t}, \lambda, \vec{u}\}} \quad & \sum_{(i,j) \in E^A} f_{ij}(\beta_{ij}^A) + \sum_{i \in V^A} \hat{g}_i(t_i) \\ \text{subject to} \quad & \frac{\sum_{j \in \mathcal{N}_i^A} \beta_{ji}^A z_j u_j + t_i u_i + \epsilon u_i}{\lambda u_i} \leq 1, \quad \forall i \in V^A, \\ & \frac{t_i}{\bar{\delta}} \leq 1, \quad \forall i \in V^A, \\ & \frac{(\bar{\delta} - \hat{\delta}_i^A)}{t_i} \leq 1, \quad \forall i \in V^A, \\ & \beta_{ij}^A, u_i \geq 0, \quad \forall i, j \in V^A, \\ & 0 \leq \lambda \leq \bar{\delta}, \end{aligned} \quad (8)$$

where $\bar{\delta} > \max_i \{\hat{\delta}_i^A\}$, \hat{g}_i denotes the posynomial transformation of g_i and we set $\delta_i^{A*} = \bar{\delta} - t_i^*$, where t_i^* is given by the optimal solution to (8).

Proof: Recall that the condition that we need to attain to guarantee local exponential stability is that J_{11} is Hurwitz.

Noting that the only negative values of J_{11} are from the term $-\text{diag}(\vec{\delta}^A)$, we can assert that the matrix $J_{11} + \bar{\delta}I + \epsilon I$ is a non-negative matrix, since each $\delta_i^A \leq \bar{\delta}$ by definition. Moreover, since J_{11} is irreducible, $J_{11} + \bar{\delta}I + \epsilon I$ must be so as well.

Proposition 1 then gives the existence of $\lambda > 0$ and \vec{u} with $u_i > 0$ for all i such that the equation

$$(J_{11} + \bar{\delta}I + \epsilon I)\vec{u} = \lambda \vec{u}$$

is satisfied. If we relax the equation and make the substitution $t_i = \bar{\delta} - \delta_i^A$ for all i , we can see that the inequalities

$$\frac{\sum_{j \in \mathcal{N}_i^A} \beta_{ji}^A z_j u_j + t_i u_i + \epsilon u_i}{\lambda u_i} \leq 1 \quad \forall i \in V^A, \quad (9)$$

compose eigenvalue equations when met with equality. It remains to show that any optimal solution to the geometric program is such that the constraints defined by (9) are met with equality.

For purposes of identifying a contradiction, assume that there exists an optimal solution in which β_{ij}^{A*} is the computed optimal value of β_{ij}^A for some constraint i for which (9) was not met with equality. Noting that β_{ij}^A affects no other constraint, we may increase β_{ij}^{A*} to some other value $\tilde{\beta}_{ij}^A > \beta_{ij}^{A*}$ such that (9) is met with equality. In doing so, we improve the value of the objective function, since f_{ij} was specified as monotonically decreasing. It must then be that our assumed solution was not optimal, and we have proven that (9) is met with equality at any optimal solution. By noting that the constraint $0 \leq \lambda \leq \bar{\delta}$ holds, we see that the leading eigenvalue of J_{11} is negative, and the extinction condition required by Theorem 2 is realized. By applying our use of the posynomial transformation, we may set $\delta_i^{A*} = \bar{\delta} - t_i^*$.

Proving the existence of a feasible solution for any permissible choice of program data remains. We proceed by construction. Select $\beta^A = \alpha A$ and $\vec{\delta}^A = \gamma \vec{1}$. Then, we can write the eigenvalue constraint as

$$\lambda_{\max}(\text{diag}(1 - \bar{\Phi}^B)(\alpha A)^T - \gamma I + \bar{\delta}I + \epsilon I) < \bar{\delta},$$

where if we choose $\gamma = \min_i \{\hat{\delta}_i^A\}$, we can reduce this to

$$\lambda_{\max}(\text{diag}(1 - \bar{\Phi}^B)(A)^T) < \frac{(\gamma - \epsilon)}{\alpha}.$$

Since we can choose any $\alpha > 0$, our proof is complete. \blacksquare

Remark 3 (Cost Function Restrictions): We have found a convex formulation for Problem 2 for the specified class of cost functions. However, the restriction is slight within the context of the problem. Given that the parameter β_{ij}^A is a rate of spread, it is natural to associate it with a monotonically decreasing cost function; this captures the intuition that enforcing a phenomenon to be less aggressive is costly when attempting to extinct it. Since we may choose any $g_i \in \mathcal{P}(0, \hat{\delta}_i^A)$, possible choices for g_i are many. To make the extent of this flexibility concrete, we note that $\mathcal{P}(0, \hat{\delta}_i^A)$ includes the class of shifted finite-order polynomials with positive coefficients.

We now shift our focus to a setting in which exponential extinction may not be possible. In particular, we consider a situation in which we are given a fixed operating budget $\mathcal{C} > 0$, and we are tasked with mitigating the spread of the unwanted

behavior insofar as is possible. In the interest of making best use of the resources available, we concern ourselves with solving the following problem:

Problem 3 (Fixed Budget Mitigation): Consider an SI_1SI_2S spreading process on a bilayer graph $G = (G_A, G_B)$. Given sets of cost functions $\{f_{ij}\}_{(i,j) \in E^A}$, $\{g_i\}_{i \in V^A}$, determine an allocation of resources which conforms to a budget $\mathfrak{C} > 0$ such that the chosen behavior \mathcal{A} extincts if possible, and mitigates the extent of its spread otherwise.

Our approach to this problem may be formalized as choosing the real component of the leading eigenvalue as the objective to our program, and adjusting the feasible set accordingly. This recovers the exponential extinction condition of Theorem 1 whenever possible, and otherwise uses the eigenvalue as a proxy for the aggregate spread of the unwanted epidemic. We formalize this as follows.

Theorem 3: Consider a realization of the dynamics (2)–(3) with an equilibrium point of the form $\bar{\Phi} = [(\bar{\Phi}^A)^T, (\bar{\Phi}^B)^T]^T$ with $\bar{\Phi}_i^A = 0$ for all i , and $\bar{\Phi}^B$ given by the solutions of (4). Define $z_i = (1 - \bar{\Phi}_i^B)$ for all $i \in V^A$, and consider any set of monotonically decreasing posynomial cost functions $\{f_{ij}\}_{\{(i,j) \in E^A\}}$, and any set of functions $\{g_i \in \mathcal{P}(0, \hat{\delta}_i^A)\}_{i=1}^{|V^A|}$. Then, Problem 3 can be solved by the following geometric program:

$$\begin{aligned} & \underset{\{\beta^A, \vec{t}, \lambda, \vec{u}\}}{\text{minimize}} \quad \lambda \\ & \text{subject to} \quad \frac{\sum_{j \in \mathcal{N}_i^A \text{ in}} \beta_{ji}^A z_i u_j + t_i u_i}{\lambda u_i} \leq 1, \quad \forall i \in V^A, \\ & \quad \frac{\sum_{(i,j) \in E^A} f_{ij} (\beta_{ij}^A) + \hat{g}_i(t_i)}{\mathfrak{C}} \leq 1, \quad \forall i \in V^A, \\ & \quad \frac{t_i}{\bar{\delta}} \leq 1, \quad \forall i \in V^A, \\ & \quad \frac{(\bar{\delta} - \hat{\delta}_i^A)}{t_i} \leq 1, \quad \forall i \in V^A, \\ & \quad \beta_{ij}^A, u_i \geq 0, \quad \forall i, j \in V^A \end{aligned} \quad (10)$$

where $\bar{\delta} > \max_i \{\hat{\delta}_i^A\}$ and \hat{g}_i denotes the posynomial transformation of g_i , and we set $\delta_i^{A*} = \bar{\delta} - t_i^*$, where t_i^* is given by the optimal solution to (10).

Proof: We will show that the stated geometric program is an equivalent problem to minimizing the eigenvalue of J_{11} . This will ensure that when the specified cost is above the optimal cost threshold, we recover the desired extinction condition; we otherwise minimize the eigenvalue as a heuristic. Noting that J_{11} is irreducible by construction and that $\bar{\delta}$ is an upper bound for all terms δ_i^A , it must be that $J_{11} + \bar{\delta}I$ is non-negative and irreducible. Hence, Proposition 1 applies and we must have the existence of some \vec{u} such that $u_i > 0$ for all i such that $(J_{11} + \bar{\delta}I)\vec{u} = \lambda\vec{u}$.

As in the proof for Theorem 2, we can relax the eigenvalue equations with the substitution $t_i = \bar{\delta} - \delta_i^A$ to obtain the inequalities

$$\frac{\sum_{j \in \mathcal{N}_i^A \text{ in}} \beta_{ji}^A z_i u_j + t_i u_i}{\lambda u_i} \leq 1, \quad \forall i \in V^A. \quad (11)$$

To show how we may attain equality of (11) at an optimal solution of (10), we may make a similar argument as to the proof of Theorem 2. However, in this case, we will show that there always exists an optimal solution which meets the constraint with equality, and construct it.

Suppose that there exists some optimal solution

$$\left\{ \{\beta_{ij}^{A*}\}_{(i,j) \in E^A}, \{\delta_i^{A*}, u_i^*, t_i^*\}_{i=1}^{|V^A|}, \lambda^* \right\},$$

at which (11) is not met with equality for some i . Since the f_{ij} functions are monotonically decreasing, we may increase the value of β_{ij}^A for some edge (i, j) until equality is attained without violating the budget constraint. Since this increase neither changes the value of λ nor makes the solution infeasible, it must be that the new solution is again optimal. Hence, given any optimal solution of (10), we may compute an optimal solution with equality in (11) by increasing values of β_{ij}^A . Given an optimal solution in which (11) is met with equality, we may then set $\delta_i^{A*} = \bar{\delta} - t_i^*$ to recover the values necessary to solve Problem 3. \blacksquare

Remark 4: Note that the program is convex for any specified $\{g_i \in \mathcal{P}(0, \hat{\delta}_i^A)\}_{i=1}^{|V^A|}$; however, particular choices of g_i may have strictly positive minimum values. Hence, there exists the possibility that (10) is infeasible. This difficulty is avoided if we restrict our choices of g_i further, for example, to functions which satisfy $\lim_{\{z \rightarrow 0^+\}} g_i(z) = 0$.

Remark 5: Formal proof that the eigenvalue minimization specified is a good proxy for optimizing the attained steady state of the chosen behavior is unavailable; however, we show in Section V that the approach works well in simulation.

We close this section by noting that the optimization programs (8) and (10) may be specialized to particular applications by the addition of further parameter constraints. Of particular interest may be the inclusion of box constraints, such that we have $\beta_{ij}^A \in [\underline{\beta}_{ij}^A, \bar{\beta}_{ij}^A]$ for all i and j , and $\delta_i^A \in [\underline{\delta}_i^A, \bar{\delta}_i^A]$ for all i , which models a scenario in which some parameter values are only partially designable. In addition, we may add constraints which enforce equality between various parameters in order to reflect a situation in which control of each spreading or healing rate cannot happen in isolation. However, since these extensions occasion no further mathematical difficulties, we will not explicitly consider them here.

V. SIMULATIONS AND DISCUSSION

Our simulations accomplish two tasks. In Section V-A, we consider the performance of the optimization methods designed with respect to the intended goals of the procedures, and find that in the mean-field regime, both methods work well. In Section V-B, we consider the accuracy of the mean-field model studied by comparison to a simulation of the exact Markov process. We find that the extinction problem studied works well for the exact process, but the mean-field model for SI_1SI_2S suffers from inaccuracy otherwise.

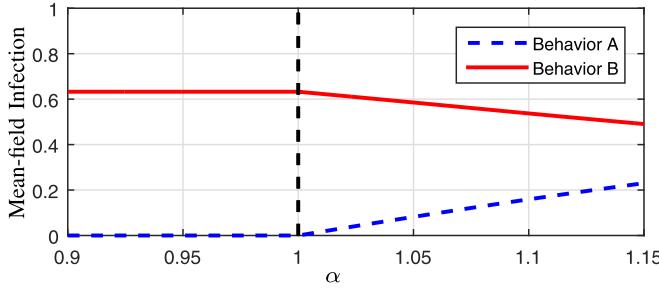


Fig. 2. A plot studying the sensitivity of the steady states of the heterogeneous mean-field SI_1SI_2S model to scaling of solutions generated by (7), denoted here by $\beta^{\mathcal{A}*}$. The average mean-field steady-state values for $\beta^{\mathcal{A}} = \alpha\beta^{\mathcal{A}*}$ are plotted on the y -axis, with the scale factor α plotted on the x -axis.

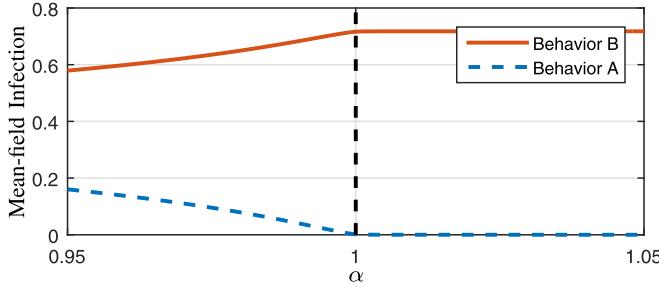


Fig. 3. Plot of the mean steady-state mean-field values of \mathcal{A} and \mathcal{B} of the solution of the optimization of Theorem 3 with the budget given by $\alpha\mathfrak{C}^*$ against α , where \mathfrak{C}^* is the optimal budget of Theorem 2.

A. Optimization Simulations

We consider the mean-field dynamics of a 100-node multi-layer graph $G = (G_A, G_B)$, with 80 nodes in G_A , 80 nodes in G_B , and 60 nodes in $G_A \cap G_B$. Each graph layer studied is a randomly generated strongly connected digraph with connection probability 0.5, with the set of nodes in the intersection selected uniformly at random from each subgraph. As cost functions, we study $f_{ij}(\beta_{ij}^{\mathcal{A}}) = \frac{1}{\beta_{ij}^{\mathcal{A}}}$ for each edge, and $g_i(\delta_i^{\mathcal{A}}) = (\hat{\delta}_i^{\mathcal{A}} - \delta_i^{\mathcal{A}})^2 + (\hat{\delta}_i^{\mathcal{A}} - \delta_i^{\mathcal{A}})$ for each node.

In our analysis, we have proven that any solution generated by (8) will have tight spectral constraints. Hence, we expect that if the contagion were made any more aggressive, it would survive. This is exactly what happens in Fig. 2, in which we consider the results of a study of the sensitivity of the solutions generated by (8). Here, we plot the attained mean-field steady states of a process with parameters $\alpha\beta^{\mathcal{A}*}$ as a function of α , where α is a scaling factor and $\beta^{\mathcal{A}*}$ is a solution computed by (8). It is precisely when $\alpha > 1$ that the behavior survives in an endemic state, as expected.

We study the efficacy of the fixed budget network design in Fig. 3, where we perform a similar sensitivity analysis as the one presented in Fig. 2. We plot the average mean-field steady-state values attained by a graph designed by (10) with a budget given by $\alpha\mathfrak{C}^*$, where \mathfrak{C}^* is the optimum value of the budget of the given problem instance, as computed by (8). We see that for all values of $\alpha \geq 1$, extinction is attained, as predicted by construction. For values of $\alpha < 1$, we see that the behavior survives in an endemic state. Moreover, the simulations indicate that the

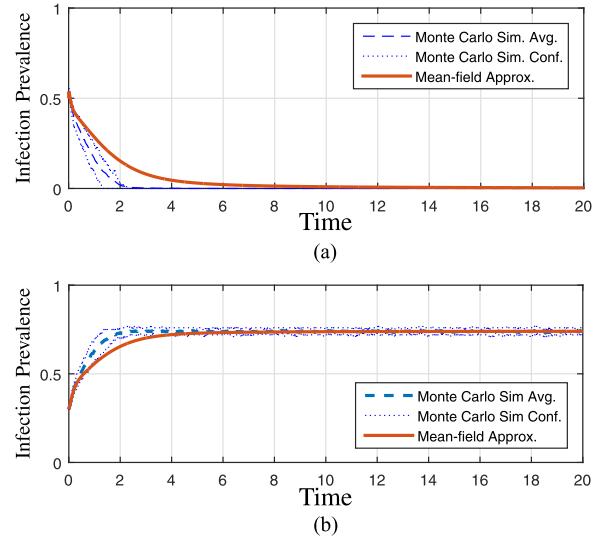


Fig. 4. The spreading behavior of the SI_1SI_2S process is compared to the mean-field approximation. These simulations were performed on the graphs used for the simulations of Section V-A. The confidence bounds plotted contain 60% of the sample paths of the simulations: a) epidemic \mathcal{A} and b) epidemic \mathcal{B} .

attained steady state grows continuously and monotonically with a decreasing budget, which suggests that the eigenvalue minimization performed serves as a suitable approach to network design when the given budget is fixed.

B. Mean-Field Simulations

To give an honest account of the utility of our results, it is necessary to investigate the relation between the behavior of the mean-field approximation we study and the SI_1SI_2S process itself. We first study networks generated as solutions to (8) with data generated as in Section V-A. A typical result of this simulation is given in Fig. 4. Here, we note that the transient response of the mean-field approximation is not tight, but the steady-state values appear to be accurate. This may be an effect of the equilibrium considered in our analysis: since epidemic \mathcal{A} 灭绝, epidemic \mathcal{B} 's dynamics eventually recover the standard SIS dynamics, which has been reported as a good approximation for sufficiently large graphs [22]. The reader should note, however, that there is a lack of scientific consensus as to when mean-field approximations are accurate, and how exactly “accuracy,” should be assessed. In particular, it is known that for many epidemic spreading models, including the SI_1SI_2S model studied here, all nodes will be susceptible with probability one in finite time. Hence, any approximation for which the approximated probability of infection does not also tend to zero as time tends to infinity will be significantly in error, asymptotically. The underdeveloped understanding of mean-field approximation in the literature is a principle motivation for our inclusion of detailed simulations in this paper.

To assess the limits of the mean-field model’s accuracy in greater generality, we consider the behavior of the model when epidemic \mathcal{A} and epidemic \mathcal{B} survive in an endemic state in a 100-node bilayer random graph. We display this in Fig. 5, where the results were taken from a 150-trial simulation of the SI_1SI_2S

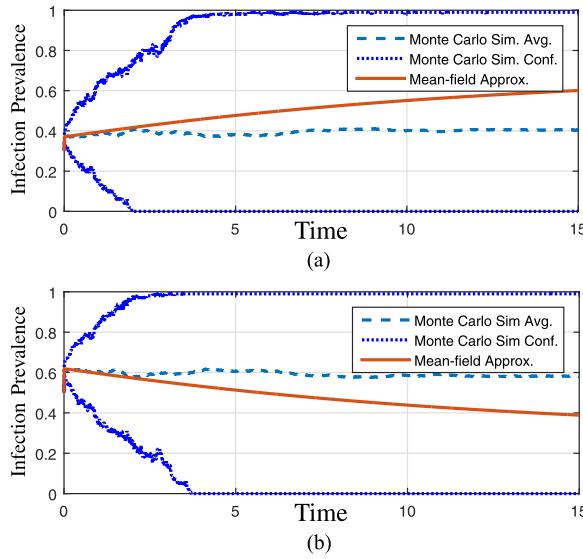


Fig. 5. Spreading behavior of the SI_1SI_2S process compared to the mean-field approximation. These simulations were performed on 100-node random graphs with 50% sparsity. The confidence bounds plotted contain 60% of the sample paths of the simulations. (a) Behavior \mathcal{A} and (b) behavior \mathcal{B} .

process, and represent a typical case. We find that the mean-field model is not necessarily accurate in this circumstance. For the particular instance displayed in Fig. 5, the error incurred between the ensemble average of the simulated SI_1SI_2S process and the mean-field approximation is approximately 0.2. Moreover, by considering the spread of the observed 60% confidence bounds of the process, we conclude that the sample paths are not well concentrated about their expectation and, hence, there is a non-negligible amount of stochasticity which is left unaccounted for by the mean-field model. This points to a need to exercise caution with respect to the analysis and control of epidemic processes, as standard heterogeneous mean-field methods used commonly in the literature may not be accurate in general.

With respect to the SI_1SI_2S model we consider here, this simulation result suggests the difficulty of pursuing more sophisticated competition models in the problem, as equilibria in which both process survive in an endemic state appear to give rise to mean-field approximations which are not necessarily accurate. It appears that it is necessary to first find a meaningful and provably accurate approximation of the process dynamics which occur when neither phenomenon die out before a rigorous analysis of a more sophisticated competitive model can be developed. As such, we hope that researchers in the field of epidemic control take heed of this result, and verify their analytical findings of approximated models against simulations of their stochastic counterparts.

VI. SUMMARY AND FUTURE WORK

The class of multilayer spreading processes is one with much potential. We have defined a framework in which the earlier work on competitive multilayer processes can be extended to a class of heterogeneously parametrized processes on generalized graph layers. Moreover, we have provided an early step in an-

alyzing competitive multilayer spreading processes by finding necessary and sufficient conditions for the exponential stability of any equilibrium of the system in which one process extincts exponentially quickly and the other survives in a prolonged state of infection.

Furthermore, we have developed two optimization frameworks for determining optimal resource allocations to attain the desired equilibrium within a given network. Our first method computes an optimal cost distribution of resources. Our second method studies a situation in which we are given a fixed budget, and wish to realize the extinction condition whenever possible, while we mitigate the presence of the unwanted epidemic otherwise. We have found that the designed optimization routines work well in simulation for both cases, with the eigenvalue minimization heuristic used in the fixed budget case appearing to be a good proxy for the attained average steady state of the mean-field model.

This work opens many possible avenues for future research. By redefining the meaning of the variable states, we can apply our model to diverse settings; potential examples include optimizing political strategies, protecting against viral spread and investigating the effect of marketing campaigns. A useful generalization would be an extension to a k -layer, k -process framework, as such an extension could greatly improve the modeling capacity of the tools developed. Additionally, we can place further assumptions on the set of controllable parameters and the objective of our resources allocations. For example, it may be reasonable to have control over both the spreading parameters of \mathcal{A} and \mathcal{B} , in which case it may be desirable to specify a steady state and compute an optimal allocation which attains it. Another interesting problem would be to study the effect of pricing the resource allocations in the network competitively, so as to incorporate the effect of competition among the network designers, in addition to the contagions.

Perhaps the most important question left open pertaining to this work is the relation of the mean-field approximation to the underlying stochastic process. While some form of approximation is necessary in order to avoid the exponential state-space of the exact representation of the system, other approximation schemes can be considered, and it is currently unclear which methods are most effective in which contexts. While the approximation technique applied here works well in simulation for the extinction problem considered, it is clear that a more precise understanding of the interrelation between the mean-field dynamics and the exact process dynamics is a substantial requirement for future work, for both the process studied here and epidemic spreading models, in general.

APPENDIX

A. Proof of Theorem 1

In order to provide a rigorous proof of Theorem 1, we will need to introduce two prior results, given by the following propositions:

Proposition 3 (Exp. Stability [25]): Let x_0 be an equilibrium point of the nonlinear system $x = f(x)$, where $f : D \rightarrow \mathbb{R}^n$ is continuously differentiable and the Jacobian matrix is

bounded and Lipschitz on D . Let

$$M = \frac{\partial f}{\partial x} \Big|_{x=x_0}.$$

Then, x_0 is an exponentially stable equilibrium point for the nonlinear system if and only if it is an exponentially stable equilibrium point of the linear system $\dot{x} = Mx$.

Proposition 4 (SIS Exp. Stability [23]): Consider the dynamics of the n-Intertwined SIS model with spreading rate matrix W and healing rate vector \vec{d} :

$$\dot{p} = \text{diag}(1-p)W^T p - \text{diag}(\vec{d})p \quad (12)$$

Suppose W and \vec{d} are chosen such that a non-zero equilibrium of (12) exists, and denote it by p^* . Then p^* is globally asymptotically stable, and locally exponentially stable. Moreover, if no non-zero equilibrium exists, then 0 is globally asymptotically stable and locally exponentially stable.

We may now proceed with the proof of Theorem 1.

Proof: We begin by computing the linearization of the mean-field dynamics given by (2) and (3) about $\bar{\Phi}$, which we can show to be:

$$\begin{bmatrix} \dot{\Psi}^A \\ \dot{\Psi}^B \end{bmatrix} = \begin{pmatrix} J_{11} & 0 \\ J_{21} & J_{22} \end{pmatrix} \begin{bmatrix} \vec{\Psi}^A \\ \vec{\Psi}^B \end{bmatrix} = J \begin{bmatrix} \vec{\Psi}^A \\ \vec{\Psi}^B \end{bmatrix}, \quad (13)$$

where

$$\begin{aligned} J_{11} &= \text{diag}(1 - \bar{\Phi}^B)(\beta^A)^T - \text{diag}(\vec{\delta}^A), \\ J_{21} &= -\text{diag}((\beta^B)^T \bar{\Phi}^B), \\ J_{22} &= \text{diag}(1 - \bar{\Phi}^B)(\beta^B)^T \\ &\quad - \text{diag}((\beta^B)^T \bar{\Phi}^B + \vec{\delta}^B), \end{aligned}$$

with $\vec{\delta}^B$ and β^B defined analogously to $\vec{\delta}^A$ and β^A , and $\vec{\Psi}^A$ and $\vec{\Psi}^B$ are introduced as dummy variables for the linearization of the mean-field dynamics.

Since the Jacobian matrix J of the system is component-wise bounded and Lipschitz, Proposition 3 gives us that the nonlinear dynamics given by (2) and (3) are (locally) exponentially stable if and only if the linearized system (13) is exponentially stable. It remains to show that the hypothesis ensures that J is Hurwitz.

We note that, due to the block 0 in the upper-right entry of J , the eigenvalues of J are given by the eigenvalues of J_{11} and J_{22} . Noting that J_{11} is indeed the matrix the hypothesis claims to be Hurwitz, we may turn our attention to J_{22} .

The J_{22} matrix is exactly the Jacobian of the dynamics of an n-Intertwined SIS system evaluated at its endemic equilibrium. We may now use Proposition 4 to claim that $\bar{\Phi}^B$ is a locally exponentially stable equilibrium point of single-layer model. By Proposition 3 it must be that J_{22} is Hurwitz, as it is componentwise bounded and Lipschitz.

Since both J_{11} and J_{22} are Hurwitz, it must be that J is Hurwitz. Hence, it must be that $\bar{\Phi}$ is an exponentially stable equilibrium point of the nonlinear system described by (2)-(3). Since all of the relations used in the proof are equivalences, no further considerations are necessary. \blacksquare

B. Proof of Lemma 2:

Proof: By considering the variational characterization of eigenvalues, it will suffice to show that

$$\begin{aligned} &\sup_{\vec{v} \neq 0} \Re \left[\frac{(\vec{v})^* (\text{diag}(\vec{\kappa}) M - \text{diag}(\vec{\gamma})) \vec{v}}{(\vec{v})^* \vec{v}} \right] \\ &\leq \sup_{\vec{v} \neq 0} \Re \left[\frac{(\vec{v})^* (M - \text{diag}(\vec{\gamma})) \vec{v}}{(\vec{v})^* \vec{v}} \right], \end{aligned} \quad (14)$$

holds, where we use \Re to denote an operator which returns the real part of its argument and $(\cdot)^*$ denotes the conjugate transpose of a vector. Our demonstration of this fact requires that we establish two pieces: (i) we may evaluate the supremums over $\vec{v} \in \mathbb{R}_{\geq 0}^n$ without affecting their attained value, and (ii) for each $\vec{v} \in \mathbb{R}_{\geq 0}^n$, the desired inequality follows immediately.

To argue (i), we will only explicitly consider the left hand side of (14), the right hand side follows from similar arguments. Fix some $\tilde{v} \in \mathbb{C}^n$ with components $\tilde{v}_r = \tilde{x}_r + i\tilde{y}_r$; we will show that we may always construct some vector $\hat{v} \in \mathbb{R}_{\geq 0}^n$ which increases the value of the function evaluated by the supremum. For our choice of \tilde{v} , we may write the argument of the the left hand side of (14) as

$$\Re \left[\frac{\sum_{k \neq \ell} (\tilde{v}_k)^* \tilde{v}_\ell \kappa_k m_{k\ell} + \sum_k (\tilde{v}_k)^* \tilde{v}_k \gamma_k}{\tilde{v}^* \tilde{v}} \right].$$

Now, consider the vector \hat{v} defined as $\hat{v}_r = \sqrt{x_r^2 + y_r^2}$ for all r . By construction, we have $(\tilde{v}_k)^* \tilde{v}_\ell = \hat{v}_k^2$ holds for all k , so it will suffice to show that

$$\Re \left[\sum_{k \neq \ell} \tilde{v}_k^* \tilde{v}_\ell \kappa_k m_{k\ell} \right] \leq \Re \left[\sum_{k \neq \ell} \hat{v}_k \hat{v}_\ell \kappa_k m_{k\ell} \right]$$

holds. This follows immediately once we establish that the inequality $\Re((\tilde{v}_k)^* \tilde{v}_\ell) \leq \hat{v}_k \hat{v}_\ell$ holds for all choices of k and ℓ . This can be verified by direct computation of the corresponding inequality for the squares of these values:

$$\begin{aligned} [\Re(\tilde{v}_k^* \tilde{v}_\ell)]^2 &= x_k^2 x_\ell^2 + y_k^2 y_\ell^2 + 2x_k x_\ell y_k y_\ell \\ &\leq x_k^2 x_\ell^2 + y_k^2 y_\ell^2 + x_k^2 y_\ell^2 + x_\ell^2 y_k^2 \\ &= \left((x_k^2 + y_k^2)^{\frac{1}{2}} (x_\ell^2 + y_\ell^2)^{\frac{1}{2}} \right)^2 \\ &= \hat{v}_k^2 \hat{v}_\ell^2 \end{aligned}$$

where the inequality follows from noting that

$$y_k^2 x_\ell^2 + x_k^2 y_\ell^2 - 2x_k x_\ell y_k y_\ell = (x_k y_\ell - x_\ell y_k)^2 \geq 0.$$

It remains to show that

$$\begin{aligned} &\sup_{\vec{v} \in \mathbb{R}_{\geq 0}^n} \Re \left[\frac{(\vec{v})^T (\text{diag}(\vec{\kappa}) M - \text{diag}(\vec{\gamma})) \vec{v}}{(\vec{v})^T \vec{v}} \right] \\ &\leq \sup_{\vec{v} \in \mathbb{R}_{\geq 0}^n} \Re \left[\frac{(\vec{v})^T (M - \text{diag}(\vec{\gamma})) \vec{v}}{(\vec{v})^T \vec{v}} \right], \end{aligned}$$

holds. We prove this by fixing any $\vec{v} \in \mathbb{R}_{\geq 0}^n$, and noting that

$$\begin{aligned} & \frac{\sum_{k \neq \ell} v_k v_\ell \kappa_k m_{k\ell} - \sum_k v_k^2 \gamma_k}{(\vec{v})^T \vec{v}} \\ & \leq \frac{\sum_{k \neq \ell} v_k v_\ell m_{k\ell} - \sum_k v_k^2 \gamma_k}{(\vec{v})^T \vec{v}}, \end{aligned}$$

follows immediately as a consequence of $\kappa_k \in [0, 1]$ for all k and $m_{k\ell} \geq 0$ for all k and ℓ . \blacksquare

REFERENCES

- [1] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Rev.*, vol. 42, no. 4, pp. 599–653, 2000.
- [2] M. J. Keeling and K. T. Eames, "Networks and epidemic models," *J. Roy. Soc. Interface*, vol. 2, no. 4, pp. 295–307, 2005.
- [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 Syst.*, vol. 36, no. 1, pp. 26–46, Feb. 2016.
- [4] M. E. Newman, "Threshold effects for two pathogens spreading on a network," *Phys. Rev. Lett.*, vol. 95, no. 10, p. 108701, 2005.
- [5] B. Karrer and M. E. J. Newman, "Competing epidemics on complex networks," *Phys. Rev. E*, vol. 84, p. 036106, Sep. 2011.
- [6] M. Newman and C. R. Ferrario, "Interacting epidemics and coinfection on contact networks," *PLoS One*, vol. 8, no. 8, p. e71321, 2013.
- [7] S. Shrestha, A. A. King, and P. Rohani, "Statistical inference for multi-pathogen systems," *PLoS Comput. Biol.*, vol. 7, no. 8, p. e1002135, 2011.
- [8] Y. Wang, G. Xiao, and J. Liu, "Dynamics of competing ideas in complex social systems," *New J. Phys.*, vol. 14, no. 1, p. 013015, 2012.
- [9] M. Salehi, R. Sharma, M. Marzolla, M. Magnani, P. Siyari, and D. Montesi, "Spreading processes in multilayer networks," *IEEE Trans. Netw. Sci. Eng.*, vol. 2, no. 2, pp. 65–83, Apr. 2015.
- [10] S. Funk and V. A. Jansen, "Interacting epidemics on overlay networks," *Phys. Rev. E*, vol. 81, no. 3, p. 036118, 2010.
- [11] C. Granell, S. Gómez, and A. Arenas, "Dynamical interplay between awareness and epidemic spreading in multiplex networks," *Phys. Rev. Lett.*, vol. 111, no. 12, p. 128701, 2013.
- [12] C. Granell, S. Gómez, and A. Arenas, "Competing spreading processes on multiplex networks: Awareness and epidemics," *Phys. Rev. E*, vol. 90, no. 1, p. 012808, 2014.
- [13] X. Wei, N. C. Valler, B. A. Prakash, I. Neamtiu, M. Faloutsos, and C. Faloutsos, "Competing memes propagation on networks: A network science perspective," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 6, pp. 1049–1060, 2013.
- [14] F. D. Sahneh and C. Scoglio, "Competitive epidemic spreading over arbitrary multilayer networks," *Phys. Rev. E*, vol. 89, p. 062817, Jun. 2014.
- [15] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, "Optimal resource allocation for network protection against spreading processes," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 1, pp. 99–108, Mar. 2014.
- [16] X. Chen and V. Preciado, "Co-infection control in multilayer networks," in *Proc. IEEE Conf. Dec. Control*, Dec. 2014.
- [17] N. J. Watkins, C. Nowzari, V. M. Preciado, and G. J. Pappas, "Optimal resource allocation for competing epidemics over arbitrary networks," in *Proc. Amer. Control Conf.* IEEE, 2015, pp. 1381–1386.
- [18] C. R. MacCluer, "The many proofs and applications of Perron's theorem," *SIAM Rev.*, vol. 42, no. 3, pp. 487–498, 2000.
- [19] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [20] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," *Optimiz. Eng.*, vol. 8, no. 1, pp. 67–127, 2007.
- [21] F. D. Sahneh, C. Scoglio, and P. V. Mieghem, "Generalized epidemic mean-field model for spreading processes over multilayer complex networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, Oct. 2013.
- [22] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 1–14, 2009.

- [23] A. Khanafer, T. Başar, and B. Gharesifard, "Stability of epidemic models over directed graphs: A positive systems approach," *Automatica*, 2015.
- [24] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. Pappas, "Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks," in *Proc. IEEE 52nd Annu. Conf. Dec. Control.*, IEEE, 2013, pp. 7486–7491.
- [25] H. K. Khalil and J. Grizzle, *Nonlinear Systems*. Prentice-Hall, 2002, vol. 3.



Nicholas J. Watkins received the B.S. degree (Hons.) in electrical engineering from Wilkes University, Wilkes-Barre, PA, USA, in 2013 and is currently pursuing the Ph.D. degree electrical and systems engineering program at the University of Pennsylvania, University Park, PA, USA.

His current research interests include the analysis and control of epidemic spreading processes, control of cyberphysical systems subject to hard resource constraints, and applications of machine learning to control.



Cameron Nowzari received the Ph.D. degree in engineering sciences from the University of California, San Diego, CA, USA, in 2013.

Currently, he is a Postdoctoral Research Associate at the University of Pennsylvania, University Park, PA, USA. His current research interests include dynamical systems and control, sensor networks, distributed coordination algorithms, robotics, applied computational geometry, event- and self-triggered control, Markov processes, network science, and spreading processes on networks.

Dr. Nowzari was a finalist for the Best Student Paper Award at the 2011 American Control Conference and received the 2012 O. Hugo Schuck Best Paper Award from the American Automatic Control Council.



Victor M. Preciado received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2008.

Currently, he is the Raj and Neera Singh Assistant Professor of Electrical and Systems Engineering at the University of Pennsylvania, University Park, PA, USA. His research interests include network science, dynamic systems, control theory, and convex optimization with applications in sociotechnical networks, technological infrastructure, and biological

systems.

Prof. Preciado is a member of the Networked and Social Systems Engineering (NETS) Program and the Warren Center for Network and Data Sciences.



George J. Pappas received the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 1998.

Currently, he is the Joseph Moore Professor and Chair of Electrical and Systems Engineering at the University of Pennsylvania, University Park, PA, USA. He also holds secondary appointments in Computer and Information Sciences, and Mechanical Engineering and Applied Mechanics. His current research interests include hybrid systems and control, embedded control systems, cyberphysical systems, hierarchical and distributed control systems, networked control systems, with applications to robotics, unmanned aerial vehicles, biomolecular networks, and green buildings.

Prof. Pappas is a member of the GRASP Lab and the PRECISE Center. He currently serves as the Deputy Dean for Research in the School of Engineering and Applied Science.