2015 American Control Conference
Palmer House Hilton
July 1-3, 2015. Chicago, IL, USA

Optimal resource allocation for competing epidemics over arbitrary networks

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

***Abstract*— This paper studies an SI_1SI_2S spreading model of two competing behaviors over a bilayer network. In particular, we address the problem of determining resource allocation strategies that ensure the extinction of one behavior while *not necessarily* ensuring the extinction of the other, and pose a marketing problem in which such a model can be of use. Our discussion begins by extending the SI_1SI_2S model to node-dependent infection and recovery parameters and generalized graph topologies, contrasting prior work. We then find conditions under which a chosen epidemic becomes extinct. We show that a distribution of resources which realizes this goal always exists for some budget under mild assumptions. We address the case in which the available budget is not sufficient for extinction by establishing analytic means for mitigating the spreading rate of the unwanted behavior. We demonstrate a method for tractably computing solutions to each problem via geometric programming. Our results are validated through simulation.**

I. INTRODUCTION

Modeling, analysis, and control of spreading processes in complex networks has recently garnered significant attention from the research community. The applications for such methods are diverse: epidemiology, social modeling, cyber security, and product adoption serve as suitable examples. However, prior work has focused primarily on the case of a single spreading network. It is clear that such an abstraction is limited in modeling capacity; many real world networks transmit phenomena through markedly different channels, motivating the need to study multi-layer models.

This paper studies such a model (heterogeneous SI_1SI_2S) with which the spread of competing beliefs and behaviors through social interaction can be modeled. We direct our attention to a set of problems focused on a single theme: controlling a spreading process so as to eliminate a chosen behavior while allowing the possibility another survives in steady state. This is a natural choice of equilibrium concept for several socially relevant frames; a few readily come to mind: the effect of political strategies on the opinions of the populace, the ramifications of gossip in professional networks, and the influence of marketing strategies on consumer behavior. We expand upon the latter in Section II.

We consider three problems in detail: (i) stabilization of a chosen equilibrium, (ii) allocation of resources to attain the chosen equilibrium, and (iii) mitigation of spread given a fixed budget. The solutions of these problems form a strong base for future work in the field. They not only serve to generalize prior efforts in the single-layer (SIS) domain, but also to strengthen the results of earlier works in the multilayer domain, both of which we now review.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Pennsylvania, PA 19104, USA, {nwatk,cnowzari,preciado,pappasg}@upenn.edu

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], [2] 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 [3]–[5], 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 [6], and the development of a model for the spread of competing ideas using the 'Susceptible-Infected-Susceptible' (SIS) model on scale-free networks [7].

A more recent trend is the investigation into systems with multiple pathogens *and* multiple spreading layers. An overview of this research area can be found in [8]. Particular examples of interest include an investigation into the effects of pathogen interaction on overlay networks with SIR dynamics [9], the development of a model in which disease awareness and infection spread on separate layers of SIS dynamics [10], [11], the development of a model (SI_1SI_2S) that generalizes the classic SIS model to a competitive multilayer framework [12], and work to find conditions under which processes in the SI_1SI_2S model can coexist [13].

We concern ourselves with finding resource allocations which control the system at optimal cost. This approach was used in controlling the mean-field approximation for the single layer SIS model in [14], and a multilayer model in [15]; we accomplish this for the SI_1SI_2S model here.

Statement of contributions: We begin by extending the work of [12] and [13] by generalizing the homogeneous parameter model to a heterogeneous setting and extending the allowable graph layer topologies to a larger class. Extending the model to the heterogeneous case - one in which each node (or agent) is affected differently by the spread of the behavior - is of importance for determining a method of controlling the processes at hand. In particular, such an extension allows for the design of a resource allocation strategy to effectively protect against an outbreak of an undesired phenomenon. Extending the allowable layer topologies from undirected graphs with identical node sets to the case of directed graphs with arbitrary node sets allows for more realistic modeling. In particular, we can capture the effect of agents having asymmetric influence over each other, and can analyze the case in which sets of agents are immune to a particular contagion: the work presented in [13] does not allow for this possibility.

We determine conditions for creating an exponentially stable equilibrium in the generalized model under which one pro-

978-1-4799-8686-6/$31.00 ©2015 AACC

cess becomes extinct while allowing the possibility that the other survives, a technique which we motivate by framing in the context of a marketing problem. We develop a method for computing a minimal-cost set of resource allocations for which the desired equilibrium is attained. In consideration of the case for which budget constraints prevent stabilization of the desired equilibrium, we develop a method for minimizing the rate of spread of the unwanted behavior. We demonstrate that each of these problems can be formulated as an equivalent geometric program, and hence solved tractably. Note that the proofs have been omitted for brevity, but will be made available in future publications.

A. Notation and Mathematical Review

Let $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$ denote the set of real and nonnegative 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 $x_i \in \mathbb{R}$. We use the notation $||\vec{x}||_1$ to denote the 1-norm of $\vec{x}$, i.e. $\sum_i |x_i|$. We use $|S|$ to denote the cardinality of a finite set.

We say a matrix A is nonnegative if every entry a_{ij} is nonnegative. A is irreducible if no similarity transformation exists which places A into block upper-triangular form. We denote the ith column of a matrix A as A_i. We denote by $\odot$ the Hadamard product, i.e. the component-wise product of two identically-shaped arrays. We denote by $\mathrm{diag}(\vec{a})$ a matrix with entries $\mathrm{diag}(\vec{a})_{ii} = a_i$ for all i and 0 elsewhere.

A *weighted 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 set of edges, and $A \in \mathbb{R}_{\geq 0}^{|V| \times |V|}$ the weighted adjacency matrix. In such a graph, the weight $a_{ij} > 0$ if and only if there exists an edge $(i, j) \in E$ connecting node i to node j, and 0 otherwise. We define the set of in-neighbors of node i given the adjacency matrix A as $\mathcal{N}_i^{A\text{in}} = \{j \in V \mid a_{ji} > 0\}$.

A path p is given by an ordered set of vertices $p = \{v_1, v_2, \ldots, v_m\}$ such that for each pair of consecutive nodes (v_k, v_{k+1}) is an edge in E. We say that some path p connects node i and j if both i and j are listed as nodes in the path. We say a digraph is strongly connected if there exists some path p connecting node v_i to node v_j for all $i, j \in V$. 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. We define the complement V^{A^c} of the vertex set of a layer G_A as $V^{A^c} = V \setminus V^A$. We say a bilayer graph is strongly connected if each layer G_A and G_B is strongly connected with respect to its node set.

We use the technique of *geometric programming* [16]. Geometric programs form a class of quasiconvex optimization problems which have *posynomial* objective functions and inequality constraints, and *monomial* equality constraints.

In the language of geometric programming, a function $f : \mathbb{R}_{\geq 0}^n \to \mathbb{R}$ is called a *monomial* if it can be written in the form $f(\vec{x}) = c\, x_1^{r_1} x_2^{r_2} \ldots x_n^{r_n}$, 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 can be made into convex optimization problems by performing a logarithmic change of variables and a logarithmic transformation of the objective and constraint functions. For more details on geometric programs and their solution, we refer the reader to [17].

II. Problem Statement

In this section, we construct an example application, formalize our model, and present the problems we seek to solve.

A. An Application in Marketing

We consider a problem faced by two firms, $\mathfrak{A}$ and $\mathfrak{B}$, in a competitive market: how might marketing resources be allocated so as to maximize effect? We model the relations between consumers by a multilayer graph G, with layers G_A and G_B. In accordance with this choice, we let nodes in $\mathbb{A} \triangleq V^A \cap V^{B^c}$ be the customers in the base of firm $\mathfrak{A}$. We define $\mathbb{B}$ likewise. We let $\mathbb{U} \triangleq V^A \cap V^B$ be the set of customers in the base neither firm.

We define $I^{\mathcal{A}}$ to be an agent-state such that the agent is considering the purchase of $\mathfrak{B}$'s product; we define $I^{\mathcal{B}}$ similarly. In what follows, we present a model in which $I^{\mathcal{A}}$ will spread through the agents populating $\mathbb{A}$ and $\mathbb{U}$ (i.e. in layer G_A). This represents the share of $\mathfrak{A}$'s potential market considering the purchase of a competitor's device. It is clear that the managers of $\mathfrak{A}$ would like to suppress this behavior, insofar as it is possible. This is the problem we formalize and solve in the remainder.

B. Competitive Bilayer Spreading Model

We begin our discussion by extending the SI_1SI_2S model proposed in [12] and analyzed further in [13]. Our primary contributions in extending this model are allowing the processes to be influenced by heterogeneous parameters, and allowing for the graph layers to be strongly connected digraphs with arbitrary node sets. This contrasts with the work in [13], which assumes homogeneous spreading parameters and undirected layers with identical node sets. Our extension allows the possibility of asymmetric influence and nodal immunity.

Recall that we denote a bilayer graph as $G = \{G_A, G_B\}$. We consider the spreading of viruses $\mathcal{A}$ and $\mathcal{B}$ over digraphs G_A and G_B, respectively. We describe the state of node i by $X_i \in \{I^{\mathcal{A}}, I^{\mathcal{B}}, S\}$, denoting whether it is affected by virus $\mathcal{A}$ ($I^{\mathcal{A}}$), virus $\mathcal{B}$ ($I^{\mathcal{B}}$), or neither (S).

Transitions from S to $I^{\mathcal{A}}$ occur at a rate determined by the product of the infection rate of $\mathcal{A}$ with respect to node i (denoted $\beta_i^{\mathcal{A}}$) with the weighted sum of i's infected neighbors $Y_i^{\mathcal{A}} = \sum_{j \in \mathcal{N}_i^{A\text{in}}} a_{ji}\, \mathbf{1}_{I^{\mathcal{A}}}(X_j)$, where $\mathbf{1}_{I^{\mathcal{A}}}(X_j) = 1$ if $X_j = I^{\mathcal{A}}$ and 0 otherwise. Similar considerations hold for transitions from S to $I^{\mathcal{B}}$. Transitions from $I^{\mathcal{A}}$ to S occur at a rate $\delta_i^{\mathcal{A}}$; likewise for transitions from $I^{\mathcal{B}}$ to S. A compartmental illustration of the process at node i is given in Figure 1. It can now be seen that the dynamics of the infinitesimal generator [19] can be written as

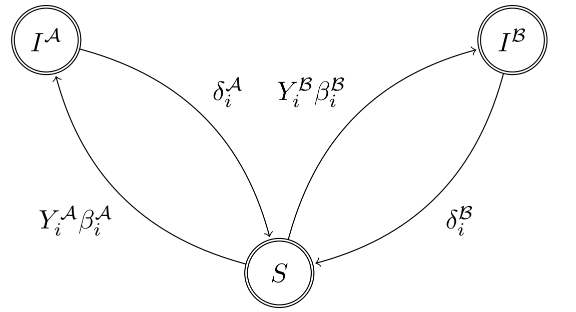

Fig. 1: Stochastic compartmental model for node i.

$$\begin{aligned}
&\Pr\left(X_i(t+\triangle t)=S \middle| X_i(t)=I^{\mathcal{A}}, X(t)\right) \approxeq \delta_i^{\mathcal{A}}\triangle t,\\
&\Pr\left(X_i(t+\triangle t)=S \middle| X_i(t)=I^{\mathcal{B}}, X(t)\right) \approxeq \delta_i^{\mathcal{B}}\triangle t, \qquad (1)\\
&\Pr\left(X_i(t+\triangle t)=I^{\mathcal{A}} \middle| X_i(t)=S, X(t)\right) \approxeq \beta_i^{\mathcal{A}}\, Y_i^{\mathcal{A}}(X(t))\triangle t,\\
&\Pr\left(X_i(t+\triangle t)=I^{\mathcal{B}} \middle| X_i(t)=S, X(t)\right) \approxeq \beta_i^{\mathcal{B}}\, Y_i^{\mathcal{B}}(X(t))\triangle t,
\end{aligned}$$

where $X(t)$ denotes the state of the Markov process at time t, and $\triangle t$ is an arbitrarily small time increment.

We note that the number of possible process states X grows exponentially with the number of nodes in the model: analysis of the exact dynamics is not tractable for large problems. Hence, we do not attempt to analyze the dynamics (1) directly, but rather their mean-field approximation:

$$\begin{aligned}
\dot{\Phi}_i^{\mathcal{A}} &= s_i\beta_i^{\mathcal{A}} \sum_{j\in\mathcal{N}_i^{A\text{in}}} a_{ji}\Phi_j^{\mathcal{A}} - \delta_i^{\mathcal{A}}\Phi_i^{\mathcal{A}},\\
\dot{\Phi}_i^{\mathcal{B}} &= s_i\beta_i^{\mathcal{B}} \sum_{j\in\mathcal{N}_i^{B\text{in}}} b_{ji}\Phi_j^{\mathcal{B}} - \delta_i^{\mathcal{B}}\Phi_i^{\mathcal{B}}, \qquad (2)\\
\dot{s}_i &= \delta_i^{\mathcal{A}}\Phi_i^{\mathcal{A}} + \delta_i^{\mathcal{B}}\Phi_i^{\mathcal{B}} - s_i\beta_i^{\mathcal{B}}\sum_{j\in\mathcal{N}_i^{B\text{in}}} b_{ji}\Phi_j^{\mathcal{B}} - s_i\beta_i^{\mathcal{A}}\sum_{j\in\mathcal{N}_i^{A\text{in}}} a_{ji}\Phi_j^{\mathcal{A}}
\end{aligned}$$

where we allow $\Phi_i^{\mathcal{A}}$ to represent the probability that node i is in state $I^{\mathcal{A}}$, with $\Phi_i^{\mathcal{B}}$ and s_i similarly defined. Note that this approximation can be arrived at via elementary Markov chain theory, but can be seen as a special case of the generalized epidemic mean-field (GEMF) model presented in [20].

Note that the additivity of probability must be satisfied, which implies $\Phi_i^{\mathcal{A}} + \Phi_i^{\mathcal{B}} + s_i = 1$ for all i. We use this to further simplify (2) to the dynamics:

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

$$\dot{\Phi}_i^{\mathcal{B}} = (1-\Phi_i^{\mathcal{A}}-\Phi_i^{\mathcal{B}}) \sum_{j\in\mathcal{N}_i^{B\text{in}}} b_{ji}\beta_i^{\mathcal{B}}\Phi_j^{\mathcal{B}} - \delta_i^{\mathcal{B}}\Phi_i^{\mathcal{B}}, \qquad (4)$$

In the coming sections, we use this model to address three central problems:

1) Stabilization (Section III): what conditions are needed for asymptotic stability of the desired equilibrium?
2) Optimal Stabilization (Section IV): can we effectively allocate resources to attain a desired equilibrium?
3) Rate Control (Section V): given a fixed budget, can we limit the spreading rate of a particular behavior?

In the forthcoming sections, we will formalize each problem and provide tractable solutions.

III. Stabilization

The goal of this section is solving the following:

Problem 1 (Stabilization) *For some specified SI_1SI_2S spreading process on a strongly connected bilayer graph G, determine conditions under which a chosen behavior $\mathcal{A}$ extincts in an asymptotically stable steady state.*

For this reason, we solve for the steady states of (3)-(4)

$$\frac{\bar{\Phi}_i^{\mathcal{A}}}{(1-\bar{\Phi}_i^{\mathcal{A}}-\bar{\Phi}_i^{\mathcal{B}})} = \frac{\beta_i^{\mathcal{A}}}{\delta_i^{\mathcal{A}}} \sum_{j\in\mathcal{N}_i^{A\text{in}}} a_{ji}\bar{\Phi}_j^{\mathcal{A}}, \qquad (5)$$

$$\frac{\bar{\Phi}_i^{\mathcal{B}}}{(1-\bar{\Phi}_i^{\mathcal{A}}-\bar{\Phi}_i^{\mathcal{B}})} = \frac{\beta_i^{\mathcal{B}}}{\delta_i^{\mathcal{B}}} \sum_{j\in\mathcal{N}_i^{B\text{in}}} b_{ji}\bar{\Phi}_j^{\mathcal{B}}. \qquad (6)$$

where $\bar{\Phi}_i^{\mathcal{A}}$ and $\bar{\Phi}_i^{\mathcal{B}}$ are the components of the steady state solutions $\bar{\Phi}^{\mathcal{A}}$ and $\bar{\Phi}^{\mathcal{B}}$ of the system. We are able to leverage the work in [18] to construct a tractable method for numerically computing the nontrivial solution of $\mathcal{B}$'s steady state equations (6) given the assumption $\bar{\Phi}_i^{\mathcal{A}} = 0$ for all i.

With the ability to claim knowledge of the steady state values $\{\bar{\Phi}_i^{\mathcal{A}}, \bar{\Phi}_i^{\mathcal{B}}\}_{i=1}^{|V^{A}|}$, we may now construct a result to Problem 1. In fact, we prove a stronger result; we find necessary and sufficient conditions for the desired equilibrium to be exponentially stable:

Theorem 1 *For any SI_1SI_2S spreading process on a strongly connected bilayer graph G with mean field dynamics given by* (3) *and* (4)*, the equilibrium $\bar{\Phi} = \{\bar{\Phi}_i^{\mathcal{A}}, \bar{\Phi}_i^{\mathcal{B}}\}_{i\in V}$ such that $\bar{\Phi}_i^{\mathcal{A}} = 0$ and $\bar{\Phi}_i^{\mathcal{B}}|_{\bar{\Phi}_i^{\mathcal{A}}=0}$ for all i is (locally) exponentially stable if and only if*

$$J = \operatorname{diag}\left((1-\bar{\Phi}^{\mathcal{B}})\odot\vec{\beta}^{\mathcal{A}}\right)A^T - \operatorname{diag}\left(\vec{\delta}^{\mathcal{A}}\right)$$

is Hurwitz.

This result is quite natural; note that the J matrix is very nearly the linearizion of a single layer SIS process about the origin. The only differences come from the competition term $(1-\vec{\bar{\Phi}}^{\mathcal{B}})$, and serve to compress the values of A^T. Intuitively, this should allow for more aggressive parameter selections. In Section VI, we will see that it is the case that the competitive terms allow for more aggressive parameter selections in simulation.

Note that this is similar to - but more general than - the stability results presented in [13]. In particular, the condition in [13] provides a scalar threshold inequality which serves to provide a sufficient condition for stability in our generalized model, provided we enforce that each node obey the threshold. Our result presents an eigenvalue condition which is necessary and sufficient for exponential stability. In Section VI, we will see that stability can be attained when the spreading parameters violate the threshold derived in [13].

IV. Optimal Stabilization

Having established conditions for stability of the desired equilibrium, we now focus our attention on establishing a solution to the following:

Problem 2 (Optimal Stabilization) *For some specified SI_1SI_2S spreading processes on a strongly connected bilayer graph G, and cost functions $\left\{f_i\left(\beta_i^{\mathcal{A}}\right)\right\}_{i=1}^{|V^A|}$ and $\left\{g_i\left(\delta_i^{\mathcal{A}}\right)\right\}_{i=1}^{|V^A|}$, determine a minimum cost allocation of resources to enforce the stability conditions for the equilibrium of Problem 1*

Note that the stabilization condition does not lead to a problem that is inherently convex; it is an eigenvector problem. However, if we allow ourselves to restrict considerations to a reasonable class of cost functions, we may extend earlier work [14] to arrive at our next result, which readily produces the solution:

Theorem 2 *Consider an equilibrium $\bar{\Phi} = \{\bar{\Phi}_i^{\mathcal{A}}, \bar{\Phi}_i^{\mathcal{B}}\}_{i\in V}$ such that $\bar{\Phi}_i^{\mathcal{A}} = 0$ and $\bar{\Phi}_i^{\mathcal{B}}|_{\bar{\Phi}_i^{\mathcal{A}}=0}$ for all i. Let $z_i = \left(1 - \bar{\Phi}_i^{\mathcal{B}}\right)$ for all i. Then, for any SI_1SI_2S spreading process on a strongly connected bilayer graph G, any set of monotonically decreasing posynomial cost functions $\{f_i\}_{i=1}^{|V^A|}$, any set of functions $\left\{g_i : \delta_i^{\mathcal{A}} \in (0, \hat{\delta}_i^{\mathcal{A}}) \mapsto \sum_k c_k \left(\hat{\delta}_i^{\mathcal{A}} - \delta_i^{\mathcal{A}}\right)^{p_k}\right\}_{i=1}^{|V^A|}$, and any $\epsilon \in \left(0, \min_i \left\{\hat{\delta}_i^{\mathcal{A}}\right\}\right)$, Problem 2 can be solved by the following geometric program:*

$$
\begin{aligned}
\underset{\vec{\beta}^{\mathcal{A}},\vec{t},\lambda,\vec{u}}{\textit{minimize}} \quad & \sum_{i\in V^A} f_i\left(\beta_i^{\mathcal{A}}\right) + \hat{g}_i\left(t_i\right) \\
\textit{subject to} \quad & \frac{\sum_{j\in\mathcal{N}_i^A} a_{ji} z_i \beta_i^{\mathcal{A}} u_j + t_i u_i + \epsilon u_i}{\lambda u_i} \le 1 \; \forall i, \\
& \frac{t_i}{\hat{\delta}_i^{\mathcal{A}}} \le 1 \; \forall i, \\
& \frac{\left(\hat{\delta} - \hat{\delta}_i^{\mathcal{A}}\right)}{t_i} \le 1 \; \forall i, \\
& \beta_i^{\mathcal{A}},\, \delta_i^{\mathcal{A}},\, u_i,\, t_i \ge 0 \; \forall i, \\
& 0 \le \lambda \le \hat{\delta}
\end{aligned}
\tag{7}
$$

where $\hat{\delta} \triangleq \max_i \left\{\hat{\delta}_i^{\mathcal{A}}\right\}$, and each $\hat{g}_i$ function is the posynomial transformation of the corresponding g_i function. Furthermore, the program is always feasible.

It is now clear that Problem 2 is solved for the specified class of cost functions. However, this restriction is slight. Given that the parameter $\beta_i^{\mathcal{A}}$ is a rate of spread, it is natural to associate it with a monotonically decreasing cost function - it captures the intuition that enforcing a phenomenon to be less aggressive is costly when attempting to extinct it. In the case of the g_i functions, the restriction is less severe. In fact, the g_i functions are structured so as to allow the $\hat{g}_i$ functions to be arbitrary posynomials - a quite flexible class.

V. RATE CONTROL

When some budget $\mathfrak{C} > 0$ is specified, we are interested in the following:

Problem 3 (Rate Control) *For some specified SI_1SI_2S spreading processes on a strongly connected bilayer graph G, and cost functions $\left\{f_i\left(\beta_i^{\mathcal{A}}\right)\right\}_{i=1}^{|V^A|}$ and $\left\{g_i\left(\delta_i^{\mathcal{A}}\right)\right\}_{i=1}^{|V^A|}$, determine an allocation of resources which conforms to a given budget $\mathfrak{C} > 0$ and limits the rate of spread of a chosen behavior $\mathcal{A}$.*

We begin formalizing this approach by providing a computable upper bound for the system:

Lemma 1 *Consider the following linear system:*

$$\dot{\Psi}^{\mathcal{A}} = \left(\mathrm{diag}\left(\vec{\beta}^{\mathcal{A}}\right) A^T - \mathrm{diag}\left(\vec{\delta}^{\mathcal{A}}\right)\right) \vec{\Psi}^{\mathcal{A}} = W \vec{\Psi}^{\mathcal{A}}$$

and let $\lambda = \lambda_{\max}(W)$. Then, for any SI_1SI_2S spreading process on a strongly connected bilayer graph G and any initial conditions $\left\{\vec{\Phi}_0^{\mathcal{A}}, \vec{\Phi}_0^{\mathcal{B}}\right\}$, the following inequality holds for all $t \ge t_0$:

$$||\vec{\Phi}^{\mathcal{A}}(t)||_1 \le ||\vec{\Psi}^{\mathcal{A}}(t)||_1 \le ||\vec{\Psi}^{\mathcal{A}}(t_0)||_1 e^{\lambda(t-t_0)}$$

with initial conditions $\vec{\Phi}_0^{\mathcal{A}} = \vec{\Psi}_0^{\mathcal{A}}$, where the dynamics of $\vec{\Phi}^{\mathcal{A}}$ and $\vec{\Phi}^{\mathcal{B}}$ are governed by (3) and (4), respectively.

We now use this result to construct a solution to Problem 3:

Theorem 3 *For any SI_1SI_2S spreading process on a strongly connected bilayer graph G, any set of monotonically decreasing posynomial cost functions $\{f_i\}_{i=1}^{|V^A|}$, any set of functions $\left\{g_i : \delta_i^{\mathcal{A}} \in (0, \hat{\delta}_i^{\mathcal{A}}) \mapsto \sum_k c_k \left(\hat{\delta}_i^{\mathcal{A}} - \delta_i^{\mathcal{A}}\right)^{p_k}\right\}_{i=1}^{|V^{\mathcal{A}}|}$, and any budget $\mathfrak{C} > 0$, Problem 3 can be solved by the following geometric program:*

$$
\begin{aligned}
\underset{\vec{\beta}^{\mathcal{A}},\vec{t},\lambda,\vec{u}}{\textit{minimize}} \quad & \lambda \\
\textit{subject to} \quad & \frac{\sum_{j\in\mathcal{N}_i^A} a_{ji} \beta_i^{\mathcal{A}} u_j + t_i u_i}{\lambda u_i} \le 1 \; \forall i, \\
& \frac{\sum_{i=1}^{|V^{\mathcal{A}}|} f_i\left(\beta_i^{\mathcal{A}}\right) + \hat{g}_i\left(t_i\right)}{\mathfrak{C}} \le 1 \; \forall i, \\
& \frac{t_i}{\hat{\delta}} \le 1 \; \forall i, \\
& \frac{\left(\hat{\delta} - \hat{\delta}_i^{\mathcal{A}}\right)}{t_i} \le 1 \; \forall i, \\
& \beta_i^{\mathcal{A}},\, \delta_i^{\mathcal{A}},\, u_i,\, t_i \ge 0,
\end{aligned}
$$

where $\hat{\delta} \triangleq \max_i \left\{\hat{\delta}_i\right\}$, and each $\hat{g}_i$ function is the posynomial transformation of the corresponding g_i function.

Here again, we see that the specified family of cost functions for Theorem 3 are quite general, for the same reasons given in Section IV for Theorem 1. It is also of interest to note that the programs given in Theorem 1 and Theorem 3 are quite similar in form and function. For this reason, one might believe that the rate-minimization problem is a reasonable heuristic for controlling the chosen behavior when extinction is not possible. This intuition is investigated in the simulations of Section VI, and found to be sensible.

VI. Numerical Simulation

In our simulations, we consider a graph of 60 nodes: 40 in G_A, 40 in G_B, and 20 in $G_A \cap G_B$. We specify $f_i = \frac{1}{\beta_i^{\mathcal{A}}}$ for each i to guarantee the necessary monotonicity for the f_i functions. We specify $\hat{g}_i = t_i^2 + t_i + \frac{1}{t_i}$ to test the claim that the $\hat{g}_i$ functions can be arbitrary posynomials.

As a test for correctness to our solution of Problem 2, we study an Euler simulation of the dynamics engendered by the optimization problem of Theorem 2 for the system with $\tau_i^{\mathcal{B}} = \frac{\beta_i^{\mathcal{A}}}{\delta_i^{\mathcal{A}}}$ generated at random from the interval $[0, 2]$, and initial conditions for each virus and each node chosen at random on the unit interval. It should be noted that for nodes in $G_A \cap G_B$, the initial condition generated is such that a value for $\Phi_i^{\mathcal{A}}$ is chosen first, with $\Phi_i^{\mathcal{B}}$ then selected with uniform probability from the interval $\left[0, 1 - \Phi_i^{\mathcal{A}}\right]$ thereafter. Figure 2 demonstrates that the system's dynamics behave as expected. In particular, the mean infection probability for $\mathcal{A}$ equilibrates to 0, as predicted by the design of the allocation.

It is interesting to note another behavior of the simulation: Figure 3 demonstrates that the optimal solutions to the stability problem act as a sort of "threshold," with respect to the disease transmissivity $\vec{\tau}^{\mathcal{A}}$. At and below the optimal solution, the desired equilibrium is attained; above it, the steady-state differs. This behavior is analogous to the concept of a "survival threshold," developed in [13]. Note again, however, that the threshold here is more general. We plotted a perturbation of the optimal vector τ_*^{A} of the form $\alpha\vec{\tau}_*^{A}$ for simplicity, but the behavior is more complex than a scalar inequality. In particular, it is possible to find other vectors on the threshold of survival by increasing the transmissivity at some nodes, with the expense of decreasing it elsewhere - an effect not present in the homogeneous model.

In the interest of verifying that the results of our optimization are nontrivial, we compare the generated solution to a homogeneous condition which can be shown to stabilize the desired equilibrium: $\tau_i^{\mathcal{A}} < \frac{1}{\lambda_{\max}\left(\operatorname{diag}\left(1-\bar{\Phi}^{\mathcal{B}}\right)A^T\right)}$ for all i. The results are given in Figure 4. We see that our conditions are, in fact, more lenient: the nodes in $G_A \cap G_B$ all take transmissivity values above the homogeneous threshold. This is intuitive - the presence of competition from $\mathcal{B}$ allows for "more aggressive," selections of parameters for $\mathcal{A}$.

In order to illustrate our solution to Problem 3, we consider the optimization problem of Theorem 3 for the same conditions as the simulation for Problem 2. The results are plotted in Figure 5. Note that the apparent "noise," in the plot is a testament to the complexity of the optimization at hand - i.e. for slightly differing budgets, the optimal parameter distributions may be markedly different, and hence cause considerable changes in steady state behavior. It is interesting to note that spreading rate minimization appears to be a useful heuristic for controlling steady state behavior. In particular, the extinction of $\mathcal{A}$ is attained at a cost very near to the optimal-cost stabilization budget, and for funds well below the threshold, the attained steady state infection rate of $\mathcal{A}$ scales well.

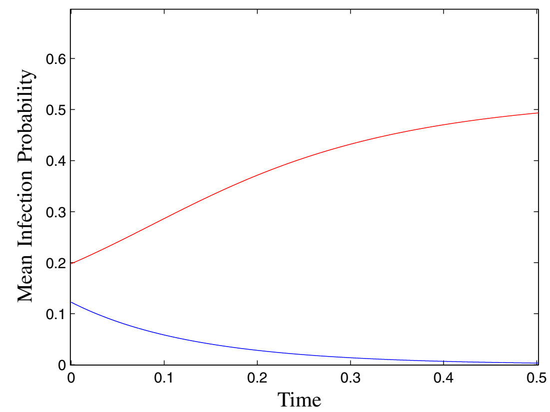

Fig. 2: A plot of mean infection probability for behavior $\mathcal{A}$ (the lower curve, in blue), and behavior $\mathcal{B}$ (the upper curve, in red) as a function of time, with solutions generated from the optimization formulation of Theorem 2.

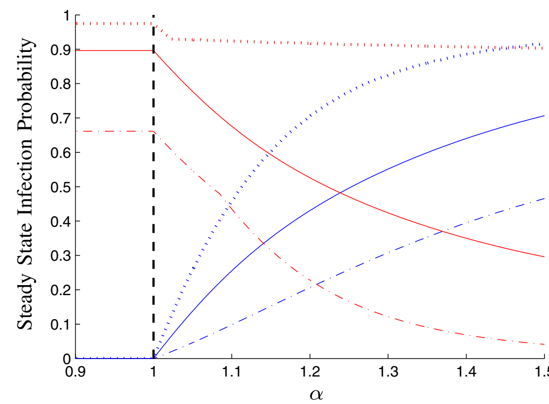

Fig. 3: A plot of steady-state infection probabilities as a function of $\alpha\vec{\tau}_*^{\mathcal{A}}$, where $\vec{\tau}_*^{\mathcal{A}}$ is a solution of the optimization problem and α is a scale factor. $\mathcal{A}$ is rendered in blue (lower on left); $\mathcal{B}$ is rendered in red (higher on left). The dotted curve represents the maximum steady-state infection probability; the solid curve represents the mean steady-state infection probability; the dashed curve represents the minimum (non-zero) steady-state infection probability.

VII. Summary and Future Work

The class of multilayer spreading processes is one with great potential. We have managed to define 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 important first step in analyzing such systems by finding necessary and sufficient conditions for the exponential stability for any equilibrium of the system in which one process dies and the other survives. This can be thought of as a characterization in which one process "dominates," the other on a sufficiently long time horizon.

Furthermore, we have developed an optimization program for determining optimal-cost parameter distributions such

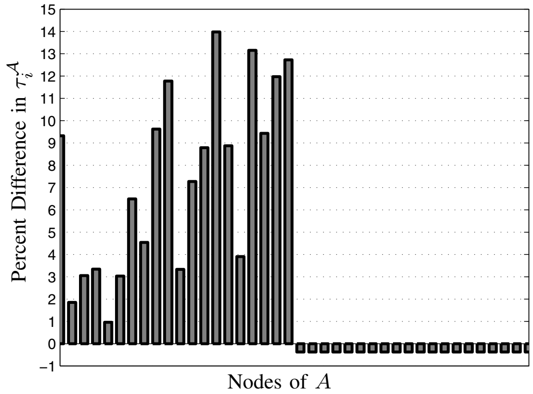

Fig. 4: A bar chart plotting the percent difference between the computed optimal solution transmissivity and the homogeneous threshold transmissivity.

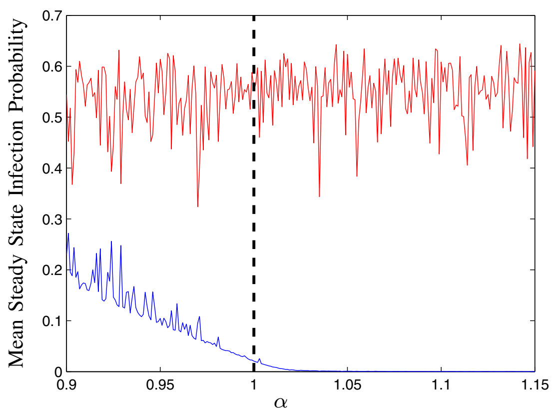

Fig. 5: A plot of mean steady-state infection probability of $\mathcal{A}$ (blue, lower curve), and $\mathcal{B}$ (red, upper curve) of the solution of the optimization of Theorem 3 with budget given by $\alpha\mathfrak{C}^{\star}$ against α, where $\mathfrak{C}^{\star}$ is the optimal budget of Theorem 2.

that the desired equilibrium is stabilized, and another which minimizes an upper bound of the spreading rate of the chosen behavior. The marketing problem we posed as motivation is just one example of the many which can be posed for such equilibria. By redefining the meaning of the variable states, we can apply our model to diverse settings: optimizing political strategies and protecting against virus spread come to mind. We leave it to the reader to conjure further examples.

This work opens many possible avenues for future research. It is obvious that a generalization to a k-layer, k-process framework would be desirable - such an extension could greatly improve the modeling capacity of the tools developed. Additionally, different assumptions can be made about 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 distribution which attains it. It may also be of interest to further define tractable methods for computing conditions of coexistence and extinction - both of which have a useful interpretation in certain contexts.

ACKNOWLEDGEMENTS

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 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.

REFERENCES

[1] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM review*, vol. 42, no. 4, pp. 599–653, 2000.

[2] M. J. Keeling and K. T. Eames, "Networks and epidemic models," *Journal of the Royal Society Interface*, vol. 2, no. 4, pp. 295–307, 2005.

[3] M. E. Newman, "Threshold effects for two pathogens spreading on a network," *Physical review letters*, vol. 95, no. 10, p. 108701, 2005.

[4] B. Karrer and M. E. J. Newman, "Competing epidemics on complex networks," *Phys. Rev. E*, vol. 84, p. 036106, Sep 2011.

[5] M. Newman and C. R. Ferrario, "Interacting epidemics and coinfection on contact networks," *PloS one*, vol. 8, no. 8, p. e71321, 2013.

[6] S. Shrestha, A. A. King, and P. Rohani, "Statistical inference for multi-pathogen systems," *PLoS computational biology*, vol. 7, no. 8, p. e1002135, 2011.

[7] Y. Wang, G. Xiao, and J. Liu, "Dynamics of competing ideas in complex social systems," *New Journal of Physics*, vol. 14, no. 1, p. 013015, 2012.

[8] M. Salehi, R. Sharma, M. Marzolla, D. Montesi, P. Siyari, and M. Magnani, "Diffusion processes on multilayer networks," *arXiv preprint arXiv:1405.4329*, 2014.

[9] S. Funk and V. A. Jansen, "Interacting epidemics on overlay networks," *Physical Review E*, vol. 81, no. 3, p. 036118, 2010.

[10] C. Granell, S. Gómez, and A. Arenas, "Dynamical interplay between awareness and epidemic spreading in multiplex networks," *Physical review letters*, vol. 111, no. 12, p. 128701, 2013.

[11] C. Granell, S. Gómez, and A. Arenas, "Competing spreading processes on multiplex networks: awareness and epidemics," *Physical Review E*, vol. 90, no. 1, p. 012808, 2014.

[12] X. Wei, N. C. Valler, B. A. Prakash, I. Neamtiu, M. Faloutsos, and C. Faloutsos, "Competing memes propagation on networks: A network science perspective," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 6, pp. 1049–1060, 2013.

[13] F. Darabi Sahneh and C. Scoglio, "Competitive epidemic spreading over arbitrary multilayer networks," *Phys. Rev. E*, vol. 89, p. 062817, Jun 2014.

[14] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, "Optimal Resource Allocation for Network Protection Against Spreading Processes," *Control of Network Systems, IEEE Transactions on*, vol. 1, pp. 99–108, Mar. 2014.

[15] X. Chen and V. Preciado, "Co-infection control in multilayer networks," in *IEEE Conference on Decision and Control (CDC)*, December 2014.

[16] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[17] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," *Optimization and Engineering*, vol. 8, no. 1, pp. 67–127, 2007.

[18] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 1, pp. 1–14, 2009.

[19] P. Mieghem, *Performance Analysis of Communications Networks and Systems*. Cambridge University Press, 2009.

[20] F. D. Sahneh, C. Scoglio, and P. V. Mieghem, "Generalized epidemic mean-field model for spreading processes over multilayer complex networks," *IEEE/ACM Transactions on Networking*, vol. 21, October 2013.