Data-Driven Network Resource Allocation for Controlling Spreading Processes

Shuo Han, *Member, IEEE*, Victor M. Preciado, Cameron Nowzari, and George J. Pappas, *Fellow, IEEE*

Abstract—We propose a mathematical framework, based on conic geometric programming, to control a susceptible-infected-susceptible viral spreading process taking place in a directed contact network with unknown contact rates. We assume that we have access to time series data describing the evolution of the spreading process observed by a collection of sensor nodes over a finite time interval. We propose a data-driven robust optimization framework to find the optimal allocation of protection resources (e.g., vaccines and/or antidotes) to eradicate the viral spread at the fastest possible rate. In contrast to current network identification heuristics, in which a single network is identified to explain the observed data, we use available data to define an uncertainty set containing all networks that are coherent with empirical observations. Through Lagrange duality and convexification of the uncertainty set, we are able to relax the robust optimization problem into a conic geometric program, recently proposed by Chandrasekaran and Shah [1], which allows us to efficiently find the optimal allocation of resources to control the worst-case spread that can take place in the uncertainty set of networks. We illustrate our approach in a transportation network from which we collect partial data about the dynamics of a hypothetical epidemic outbreak over a finite period of time.

Index Terms—Spreading processes, resource allocation, networked dynamics, robust optimization

1 Introduction

Modeling and analysis of spreading processes in complex networks is a rich and interdisciplinary research field with a wide range of applications. Examples include disease propagation in human populations [2], [3], [4], [5], [6] or information spreading in social networks [7], [8], [9], [10]. A classical model of disease spreading is the susceptible-infected-susceptible (SIS) epidemic model [2], [3]. This model was originally proposed in the context of 'unstructured' populations [6]. Due to the current availability of accurate datasets describing complex patterns of network connectivity, the classical SIS model has been extended to model spreading processes in 'networked' populations using a variety of Markov models [4], [5], [7], [11], [12], [13], [14], [15], [16], [17].

There is a fast-growing body of literature on containing epidemic outbreaks given limited control resources. In the context of epidemiology, these resources can be pharmaceutical (e.g., vaccines and antidotes) as well as non-pharmaceutical actions (e.g., traffic control and quarantines). Since these resources are costly, it is of relevance to develop computational tools to optimize the allocation of resources throughout a population to control an outbreak. This problem has attracted the attention of the network science community, resulting in a variety of vaccination heuristics. For example, Cohen et al. [18] proposed a vaccination strategy, called *acquaintance immunization policy*, and proved it to be much more efficient than random vaccine allocation. Borgs et al. [19] studied theoretical limits in the control of spreads in undirected network by distributing antidotes. Chung et al. [20] proposed an immunization strategy based on PageRank centrality. Similar problems have also been studied recently in the communication and control community [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [33].

We base our work on [34], [35], where Preciado et al. developed a convex optimization framework to find the cost-optimal distribution of vaccines and antidotes in both directed and undirected networks when the network structure and spread rates are fully specified. In comparison, our current work does not assume full information about the network structure and spreading rates, since this information is only partially known in most practical applications. To elaborate on this point, let us consider the following setup. Assume that each node in a network represents subpopulations (e.g., city districts) connected by edges that are determined by commuting patterns between districts. In practice, one can use traffic information and geographical proximity to infer the existence of an edge connecting districts; however, it is very challenging to use this information to estimate the contact rates between subpopulations. Inspired by this practical realization, we consider a networked SIS model taking place in a contact network with unknown contact rates. To extract information about these unknown rates, we assume that we have access to time series data describing the evolution of the spreading process observed by a collection of sensor nodes over a finite time interval. Such time series data can be obtained from web services such as Google Flu Trends [36], or public health agencies such as the Center for Disease Control in the US [37].

A possible approach to recover the spreading rates is the use of network identification techniques [38], [39], [40], [41], [42], [43], [44], [45], [46], [47]. However, these techniques are designed to find only one of the many networks that are

• *The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104. E-mail: {hanshuo, preciado, cnowzari, pappasg}@seas.upenn.edu.*

Manuscript received 12 Nov. 2014; revised 17 July 2015; accepted 3 Nov. 2015. Date of publication 10 Nov. 2015; date of current version 8 Jan. 2016. Recommended for acceptance by S. Roy
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TNSE.2015.2500158

coherent with empirical observations [48]. Furthermore, as illustrated in [49], [50], these techniques can lead to unsuccessful network identification. In contrast to network identification techniques, we propose a data-driven robust optimization framework to find the optimal allocation of protection resources (e.g., vaccines and/or antidotes) over a set of control nodes to eradicate the viral spread at the fastest possible rate. In contrast to current network identification heuristics, in which a single network is identified to explain the observed data, we define an uncertainty set containing all networks that are consistent with the observed data. Through Lagrange duality and convexification of the uncertainty set of networks, we can relax the robust optimization problem into a *conic geometric program*, which has recently been proposed by Chandrasekaran and Shah [1]. In this context, we are able to efficiently find the optimal allocation of resources to control the worst-case spread that can take place in the uncertainty set of networks. We illustrate our approach in a transportation network from which we collect partial data about the dynamics of a hypothetical epidemic outbreak over a finite period of time. We discover that incorporating observations into the uncertainty set of networks significantly helps reduce the worst-case bound on the spreading rate. As we increase either the length of time over which observations are taken or the number of sensor nodes, the bound on the spreading rate decreases monotonically and converges after a relatively small number of observations (either in time or in the number of nodes). Furthermore, even though our allocation algorithm does not have access to the true underlying contact network, the resulting allocation performs surprisingly close to the full-knowledge optimal allocation.

The rest of the paper is organized as follows. In Section 2, we provide some preliminaries and formulate the problem under consideration. In Section 3, we introduce the conic geometric programming framework and provide the details about how to cast our problem into this framework. In Section 4, we illustrate our approach with numerical simulations using data from the air transportation network.

2 Preliminaries and Problem Definition

We begin by introducing the notation and preliminary results needed in our derivations. In the rest of the paper, we denote by $\mathbb{R}^n_+$ (respectively, $\mathbb{R}^n_{++}$) the set of n-dimensional vectors with nonnegative (respectively, positive) entries. For $d \in \mathbb{N}$, we define $[d]$ as the set of integers $\{1, \ldots, d\}$. We denote vectors using boldface and matrices using capital letters. We denote by $\mathbf{0}$ the vector of all zeros. Given two vectors $\mathbf{a}$ and $\mathbf{b}$ of equal dimension, $\mathbf{a} \succeq \mathbf{b}$ indicates component-wise inequality.

2.1 Graph-Theoretic Nomenclature

A *weighted*, *directed* graph is defined as the triad $\mathcal{G} \triangleq (\mathcal{V}, \mathcal{E}, \mathcal{W})$, where $\mathcal{V} \triangleq \{v_1, \ldots, v_n\}$ is a set of n nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of ordered pairs of nodes called directed edges, and the weight function $\mathcal{W} : \mathcal{E} \to \mathbb{R}_{++}$ associates *positive* real weights to the edges in $\mathcal{E}$. Throughout the paper, we may use v_i and i interchangeably for all $i \in [n]$. By convention, we say that (v_j, v_i) is an edge from v_j pointing towards v_i. We define the in-neighborhood of node v_i as $\mathcal{N}_i \triangleq \{j \in [n] : (v_j, v_i) \in \mathcal{E}\}$. We define the weighted *in-degree* of node v_i as $d_i \triangleq \sum_{j \in \mathcal{N}_i} \mathcal{W}\big((v_j, v_i)\big)$. A directed path from v_{i_1} to v_{i_l} in $\mathcal{G}$ is an ordered set of vertices $(v_{i_1}, v_{i_2}, \ldots, v_{i_{l-1}}, v_{i_l})$ such that $(v_{i_s}, v_{i_{s+1}}) \in \mathcal{E}$ for $s = 1, \ldots, l-1$. A directed graph $\mathcal{G}$ is *strongly connected* if, for every pair of nodes $v_i, v_j \in \mathcal{V}$, there is a directed path from v_i to v_j. The *adjacency matrix* of a weighted, directed graph $\mathcal{G}$, denoted by $A_{\mathcal{G}}$, is an $n \times n$ matrix with entries $a_{ij} = \mathcal{W}\big((v_j, v_i)\big)$ if edge $(v_j, v_i) \in \mathcal{E}$, and $a_{ij} = 0$ otherwise. In this paper, we only consider graphs with positively weighted edges; hence, adjacency matrices are always nonnegative. Given an $n \times n$ nonnegative matrix A, we can always associate a directed graph $\mathcal{G}_A$ such that A is the adjacency matrix of $\mathcal{G}_A$. Finally, a nonnegative matrix A is *irreducible* if and only if its associated graph $\mathcal{G}_A$ is strongly connected.

Given an $n \times n$ matrix M, we denote by $\lambda_1(M), \ldots, \lambda_n(M)$ the eigenvalues of M, where we order them according to their magnitudes, i.e., $|\lambda_1| \geq |\lambda_2| \geq \ldots \geq |\lambda_n|$. We denote the corresponding eigenvectors by $\mathbf{v}_1(M), \ldots, \mathbf{v}_n(M)$. We call $\lambda_1(M)$ the spectral radius (or dominant eigenvalue) of M, which we also denote by $\rho(M)$.

2.2 SIS Model in Directed Networks

In our work, we model the spread of a disease using an extension of the networked discrete-time SIS model proposed by Wang et al. in [17]. In contrast to Wang's model, we consider directed networks (instead of undirected) with non-homogeneous transmission and recovery rates (instead of homogeneous) as described below. In all SIS models, each node can be in one out of two possible states: *susceptible* or *infected*. Over time, nodes switch their states according to a stochastic process parameterized by (*i*) a set of infection rates $\{\beta_{ij} \in (0,1)\}_{(v_j, v_i) \in \mathcal{E}}$ representing the rates at which an infection can be transmitted through the edges in the network, and (*ii*) a set of recovery rates $\{\delta_i \in (0,1)\}_{v_i \in \mathcal{V}}$ representing the rates at which nodes recover from an infection. We define $p_i(t)$ to be the probability of node v_i being infected at a particular time slot $t \in \mathbb{N}$. In the epidemiological problem considered herein, it is convenient to associate each node not to an individual, but a subpopulation living in a particular district.[1] In this context, the variable $p_i(t)$ represents the fraction of the population being infected at time t. In the original model proposed by Wang et al. [17], the infection and recovery rates were assumed to be homogeneous, i.e., $\beta_{ij} = \beta$ and $\delta_i = \delta$, and the evolution of $p_i(t)$ was described by a set of difference equations obtained from a mean-field approximation (see [17], eq. (5)-(6)). In our work, we consider the case of non-homogeneous contact and recovery rates. This case has been studied in the work of Wan et al. [21], which shows that the set of difference equations that describe the system dynamics can be written as follows:

$$p_i(t+1) = (1 - p_i(t)) \left\{ 1 - \prod_{j \in \mathcal{N}_i} \big[1 - \beta_{ij} p_j(t)\big] \right\} + (1 - \delta_i) p_i(t) \tag{1}$$

1. Although in the original networked SIS model [17], nodes represented individuals in a social network, we find the interpretation of nodes as districts better-suited for epidemiological applications.

for $i \in [n]$. This is a system of nonlinear difference equations for which one can derive sufficient conditions for global stability as follows. First, notice the following upper bound of (1)

$$\begin{aligned} p_i(t+1) &\leq 1 - \prod_{j \in \mathcal{N}_i} \left[1 - \beta_{ij} p_j(t)\right] + (1-\delta_i) p_i(t) \\ &\leq \sum_{j \in \mathcal{N}_i} \beta_{ij} p_j(t) + (1-\delta_i) p_i(t), \end{aligned} \tag{2}$$

where the last upper bound is a close approximation of (1) for $p_i(t) \ll 1$ and/or $\beta_i \ll 1$. For convenience, we define the complementary recovery rate of node v_i as $\delta_i^c \triangleq 1 - \delta_i$, and the vector $\mathbf{d}^c := \left(\delta_1^c, \ldots, \delta_n^c\right)^T$. We also define the matrix of infection rates $B_{\mathcal{G}} \triangleq \left[\beta_{ij}\right]$, where we assume $\beta_{ij} = 0$ for all pairs (i,j) such that $\left(v_j, v_i\right) \notin \mathcal{E}$. Notice that $B_{\mathcal{G}}$ maintains the same sparsity pattern as $A_{\mathcal{G}}$. Using the upper bound in (2), we define the following linear discrete-time system $\widehat{p}_i(t+1) = \sum_{j \in \mathcal{N}_i} \beta_{ij} \widehat{p}_j(t) + \delta_i^c \widehat{p}_i(t),\ i \in [n]$, which can be written in matrix-vector form as $\widehat{\mathbf{p}}(t+1) = M(B_{\mathcal{G}}, \mathbf{d}^c)\widehat{\mathbf{p}}(t)$, where $\widehat{\mathbf{p}}(t) \triangleq (\widehat{p}_1(t), \ldots, \widehat{p}_n(t))^T$ and the state matrix is given by $M(B_{\mathcal{G}}, \mathbf{d}^c) \triangleq B_{\mathcal{G}} + \operatorname{diag}(\mathbf{d}^c)$. Hence, the linear system is asymptotically stable if

$$\rho(M(B_{\mathcal{G}}, \mathbf{d}^c)) \ < \ 1, \tag{3}$$

where the spectral radius ρ of the state matrix M determines the exponential decay rate of the infection probabilities, i.e., $\|\widehat{\mathbf{p}}(t)\| \leq c\|\widehat{\mathbf{p}}(0)\|\rho^t$ for some $c > 0$. Since (2) upper bounds (1), we have that $\widehat{\mathbf{p}}(t) \succeq \mathbf{p}(t)$ for all $t \in \mathbb{N}$ when $\widehat{\mathbf{p}}(0) = \mathbf{p}(0)$. Therefore, the spectral condition in (3) is sufficient for global asymptotic stability of the nonlinear model in (1). Furthermore, the smaller the magnitude of $\rho(M)$, the faster the disease dies out.

2.3 Problem Formulation

Our main objective is to find the optimal allocation of control resources to eradicate a disease at the fastest rate possible. In order to formulate our problem, we first need to describe what pieces of information are available and what control actions we are considering. In what follows, we first describe the information available. In most real epidemiological problems, researchers do not have access to the spreading rates associated to the links connecting different districts. Therefore, the exact state matrix $M(B_{\mathcal{G}}, \mathbf{d}^c)$ is usually unknown. In order to extract information about the state matrix, we consider two different sources of information that are generally available in epidemiological problems. We classify these sources as (*i*) *prior information* about the network topology and parameters of the disease, and (*ii*) *empirical observations* about the spreading dynamics. In particular, we consider the following pieces of prior information:

P1. We assume that the sparsity pattern of the contact matrix $B_{\mathcal{G}}$ is given, although its entries are unknown. This piece of information may be inferred from geographical proximity, commuting patterns, or the presence of transportation links connecting subpopulations.

P2. We assume that we know the upper and lower bounds on the spreading rates associated to each edge, i.e., $\beta_{ij} \in \left[\underline{\beta}_{ij}, \overline{\beta}_{ij}\right]$, for all $(i,j) \in \mathcal{E}$, which may be inferred from traffic densities and subpopulation sizes.

P3. In practice, each district contains a large number of individuals. Therefore, we can use the average recovery rate in the absence of vaccination as an estimation of the nodal recovery rate. We denote this 'natural' recovery rate by δ_i^0, and assume it to be known.

Apart from these pieces of prior information, we also assume that we have access to partial observations about the evolution of the spread over a finite time interval. In particular, we assume that we observe the dynamics of the disease for $t \in [0,T]$ from a collection of sensor nodes $\mathcal{V}_S \subseteq \mathcal{V}$. In other words, we have access to the following data set:

$$\mathcal{D} \triangleq \{p_i(t) : \text{for all } i \in \mathcal{V}_S,\ t \in [T]\}. \tag{4}$$

We assume that the data are collected before any control action is taken; therefore, the evolution of $p_i(t)$ follows the dynamics in (1) with $\delta_i = \delta_i^0$ (which we assume to be known).

In what follows, we define an uncertainty set that contains all contact matrices $B_{\mathcal{G}}$ consistent with both empirical observations and prior knowledge. Based on our prior knowledge described in items P1–P3 above, we define the following uncertainty set:

$$\begin{aligned} \Delta_{B_{\mathcal{G}}}^P \triangleq \{B_{\mathcal{G}} \in \mathbb{R}^{n \times n} : \underline{\beta}_{ij} \leq \beta_{ij} \leq \overline{\beta}_{ij},\ \forall (i,j) \in \mathcal{E}; \\ \beta_{ij} = 0,\ \forall (i,j) \notin \mathcal{E}\}. \end{aligned}$$

We also define $\Delta_{B_{\mathcal{G}}}^D$ to be the set of contact matrices that are coherent with the empirical observations $\mathcal{D}$:

$$\begin{aligned} \Delta_{B_{\mathcal{G}}}^D \triangleq \{B_{\mathcal{G}} \in \mathbb{R}^{n \times n} : \{\beta_{ij}\}_{(i,j) \in \mathcal{E}} \text{ satisfy (1)} \\ \text{for } \delta_i = \delta_i^0 \text{ and } p_i(t) \in \mathcal{D},\ \forall i \in \mathcal{V}_S,\ t \in [T]\}. \end{aligned}$$

The set contains those contact matrices $B_{\mathcal{G}}$ such that the transmission rates $\{\beta_{ij}\}$ are consistent with the 'natural' disease dynamics in (1) with $\delta_i = \delta_i^0$. Notice that $\Delta_{B_{\mathcal{G}}}^D$ is defined as a collection of polynomial equality constraints on the contact rates $\{\beta_{ij}\}$ given by (1). The uncertainty set that combines information from both prior knowledge and empirical observations is defined as

$$\Delta_{B_{\mathcal{G}}} \triangleq \Delta_{B_{\mathcal{G}}}^P \cap \Delta_{B_{\mathcal{G}}}^D. \tag{5}$$

Having introduced the pieces of available information, we now describe the set of control actions under consideration. In order to eradicate the disease at the fastest rate possible, we assume that we can use pharmaceutical resources to tune the recovery rates in a collection of control nodes, i.e., δ_i for $v_i \in \mathcal{V}_C \subseteq \mathcal{V}$. In practice, these resources might be implemented by, for example, distributing vaccines and/or antidotes throughout the subpopulations located at those control districts. We assume that distributing vaccines in a district has an associated cost, which we represent as a node-dependent vaccine cost function. It is convenient to describe the vaccine cost function of a district in terms of its

complementary recovery rate δ_i^c. We denote the vaccine cost function of node i by $g_i(\delta_i^c)$. This function represents the cost of tuning the complementary recovery rate of the subpopulation at node $i \in \mathcal{V}_C$ towards the value δ_i^c. We assume that we can control the complementary recovery rate δ_i^c within a given feasible interval $[\underline{\delta}_i^c, \overline{\delta}_i^c]$, where $0 < \underline{\delta}_i^c < \overline{\delta}_i^c = 1 - \delta_i^0$. We assume that the cost of achieving $\overline{\delta}_i^c$ is zero, since it is equivalent to maintaining the natural recovery rate. We also assume that the maximum of g_i in $[\underline{\delta}_i^c, \overline{\delta}_i^c]$ is achieved at $\underline{\delta}_i^c$. Furthermore, we also assume that g_i is monotonically decreasing in the range $[\underline{\delta}_i^c, \overline{\delta}_i^c]$. In other words, as we increase the level of investment to protect a given subpopulation, we also increase the recovery rate of that subpopulation.

We are now in a position to state the control problem under consideration:

Problem 1 (Data-driven optimal allocation). *Assume we are given the following pieces of information about a viral spread:*

- *(i) prior information about the state matrix (as described in P1–P3);*
- *(ii) a finite (and possibly sparse) data series representing partial evolution of the spread over a set of sensor nodes $\mathcal{V}_S \subseteq \mathcal{V}$ during the time interval $t \in [T]$ (i.e., $\mathcal{D}$ in (4));*
- *(iii) a set of vaccine cost functions g_i for all $i \in \mathcal{V}_C$, and a range of feasible recovery rates $[\underline{\delta}_i^c, \overline{\delta}_i^c]$ such that $1 - \delta_i^0 = \overline{\delta}_i^c \geq \delta_i^c \geq \underline{\delta}_i^c > 0$;*
- *(iv) a fixed budget $C > 0$ to be allocated throughout a set of control nodes in $\mathcal{V}_C \subseteq \mathcal{V}$, so that $\sum_{i \in \mathcal{V}_C} g_i(\delta_i^c) \leq C$.*

Find the cost-constrained allocation of control resources to eradicate the disease at the fastest possible exponential rate, measured as $\rho(M(B_\mathcal{G}, \mathbf{d}^c))$, over the uncertainty set $\Delta_{B_\mathcal{G}}$ of contact matrices coherent with prior knowledge and the observations in $\mathcal{D}$.

From the perspective of optimization, Problem 1 is equivalent to finding the optimal allocation of resources to minimize the worst-case (i.e., maximum possible) decay rate $\rho(M(B_\mathcal{G}, \mathbf{d}^c))$ for all $B_\mathcal{G} \in \Delta_{B_\mathcal{G}}$. This can be cast as a robust optimization problem in the following:

$$\begin{aligned} \underset{\mathbf{d}^c}{\text{minimize}} \quad & \sup_{B_\mathcal{G} \in \Delta_{B_\mathcal{G}}} \rho(M(B_\mathcal{G}, \mathbf{d}^c)) \\ \text{subject to} \quad & \sum_{i \in \mathcal{V}_C} g_i(\delta_i^c) \leq C, \\ & \delta_i^c \leq \delta_i^c \leq \overline{\delta}_i^c, \quad i \in \mathcal{V}_C, \end{aligned} \tag{6}$$

where the first constraint accounts for our budget limit C. In general, the set $\Delta_{B_\mathcal{G}}$ is nonconvex due to the observation-based uncertainty set $\Delta_{B_\mathcal{G}}^D$. In Section 3.3, we will define a convex superset $\widehat{\Delta}_{B_\mathcal{G}}^D \supset \Delta_{B_\mathcal{G}}^D$, such that problem (6) can be relaxed into a conic geometric program. In our numerical simulations, we verify that this relaxation provides a good approximation based on real network data. From here on, we will refer to problem (6) as the *robust allocation problem*.

By solving problem (6), one is able check whether the disease can be eradicated in the worst case for the given amount of budget. For a given total budget C, if the optimal value of problem is greater than or equal to 1, then it is possible that the disease cannot be eradicated for some network that is consistent with available observations. Similarly, one can also choose to use problem (6) as a way to find the minimum budget for eradicating the disease; for instance, one may vary C using the bisection procedure until the optimal value of problem (6) falls below 1.

For practical implementation, when the spread of the disease has just started at $t = 0$, one does not have any observations available and may choose an allocation based only on prior knowledge about the network. After the disease has spread across the network for some time, one can use available observations to revise the allocation by solving the robust allocation problem (6). As we demonstrate via simulations in Section 4, one major benefit of the data-driven formulation is that problem (6) only requires a small number (relative to the number of nodes in the network) of observations to converge while avoiding identification of the underlying network $B_\mathcal{G}$.

3 Data-Driven Resource Allocation

In this section, we develop a mathematical framework to solve the robust allocation problem described above. Our solution is based on geometric programming [51] and its conic extension recently proposed by Chandrasekaran and Shah in [1]. We start our exposition by briefly reviewing some concepts used in our formulation.

3.1 Robust Geometric Programming

Geometric programs (GPs) are a type of quasiconvex optimization problem that can be easily transformed into a convex program and solved in polynomial time. Let $x_1, \ldots, x_n > 0$ denote n decision variables and define $\mathbf{x} \triangleq (x_1, \ldots, x_n) \in \mathbb{R}_{++}^n$. In the context of GP, a *monomial* $m(\mathbf{x})$ is defined as a real-valued function of the form $m(\mathbf{x}) \triangleq d x_1^{a_1} x_2^{a_2} \ldots x_n^{a_n}$ with $d > 0$ and $a_i \in \mathbb{R}$. A *posynomial* function $f(\mathbf{x})$ is defined as a sum of monomials, i.e., $f(\mathbf{x}) \triangleq \sum_{k=1}^K c_k x_1^{a_{1k}} x_2^{a_{2k}} \ldots x_n^{a_{nk}}$, where $c_k > 0$ and $a_{ik} \in \mathbb{R}$. It is convenient to write down a posynomial as the product of a vector of nonnegative coefficients $\mathbf{c} \triangleq (c_1, \ldots, c_K)$ and a vector of monomials $\mathbf{m}(\mathbf{x}) \triangleq (m_1(\mathbf{x}), \ldots, m_K(\mathbf{x}))^T$, such that $f(\mathbf{x}) = \mathbf{c}^T \mathbf{m}(\mathbf{x})$. Notice that $\{m_k(\mathbf{x})\}_{k=1}^K$ is the set of all K monomials involved in our posynomial. Posynomials are closed under addition, multiplication, and nonnegative scaling. A posynomial can be divided by a monomial, with the result a posynomial.

A GP is an optimization problem of the form (see [51] for a comprehensive treatment):

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}_{++}^n}{\text{minimize}} \quad & f_0(\mathbf{x}) \\ \text{subject to} \quad & f_i(\mathbf{x}) \leq 1, \quad i \in [m], \\ & h_j(\mathbf{x}) = 1, \quad j \in [p], \end{aligned} \tag{7}$$

where f_i are posynomial functions and $h_j(\mathbf{x}) \triangleq d_j x_1^{b_{1,j}} x_2^{b_{2,j}} \ldots x_n^{b_{n,j}}$ are monomials. To write f_i in vector-product form, we can define a vector $\mathbf{c}_i$ of positive coefficients such that $f_i(\mathbf{x}) = \mathbf{c}_i^T \mathbf{m}(\mathbf{x})$, so that the posynomial constraints in (7) can be written as $\mathbf{c}_i^T \mathbf{m}(\mathbf{x}) \leq 1$.

A GP is a quasiconvex optimization problem [52] that can be convexified using the logarithmic change of variables $y_i = \log x_i$ (see [51] for more details on this transformation). After this transformation, the GP in (7) takes the form

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^n}{\text{minimize}} \quad & \widetilde{f}_0(\mathbf{y}) \\ \text{subject to} \quad & \widetilde{f}_i(\mathbf{y}) \le 0, \quad i \in [m], \\ & \mathbf{b}_j^T\mathbf{y} + \log d_j = 0, \quad j \in [p], \end{aligned} \tag{8}$$

where $\widetilde{f}_i(\mathbf{y}) \triangleq \log f_i(e^{\mathbf{y}})$ for $i \in \{0, 1, \ldots, m\}$ and $\mathbf{b}_j \triangleq \left(b_{1,j}, \ldots, b_{n,j}\right)^T$ (i.e., the exponents of the monomial h_j) for $i \in [m]$. As a result of this transformation, the optimization problem (8) is convex and can be efficiently solved in polynomial time (see [52, Chapter 4.5] for more details).

In this paper, we shall use conic GP, which is a conic extension of GP, to solve the following *robust GP with coefficient uncertainties*:

$$\underset{\mathbf{x}\in\mathbb{R}^n_{++}}{\text{minimize}} \quad f_0(\mathbf{x}) \tag{9}$$

$$\text{subject to} \quad \sup_{\mathbf{c}_i\in\mathcal{C}_i} \mathbf{c}_i^T\mathbf{m}(\mathbf{x}) \le 1, \quad i \in [m], \tag{10}$$

$$h_j(\mathbf{x}) = 1, \quad j \in [p],$$

where $\mathbf{c}_i \in \mathbb{R}_+^K$ is a vector of coefficients contained in an uncertainty set $\mathcal{C}_i \subseteq \mathbb{R}_+^K$. The robust GP in (9) extends the formulation of the standard GP in (7) to account for uncertainties in the coefficients of the posynomial functions f_i for $i \in [m]$.

However, the constraints (10) cannot be handled naturally by numerical optimization solvers. In what follows, we propose a methodology to rewrite these constraints in a more numerically favorable manner when the uncertainty sets $\mathcal{C}_i$ in (10) can be expressed in terms of an m_i-dimensional convex cone $\mathcal{K}_i \subset \mathbb{R}^{m_i}$ as follows:

$$\mathcal{C}_i \triangleq \{\mathbf{c}_i \in \mathbb{R}_+^K : F_i\mathbf{c}_i + \mathbf{g}_i \in \mathcal{K}_i\} \tag{11}$$

for some fixed $F_i \in \mathbb{R}^{m_i\times K}$ and $\mathbf{g}_i \in \mathbb{R}^{m_i}$. The form (11) is a quite general representation for convex sets; for instance, when $\mathcal{K}_i$ is the nonnegative orthant $\mathbb{R}_+^{m_i}$, the set $\mathcal{C}_i$ becomes a convex polytope. Based on the representation (11) of $\mathcal{C}_i$, we can use duality theory to derive a more numerically favorable representation of the constraint in (10) as follows. Assuming $\mathcal{C}_i$ can be represented as (11), we have that for each $i \in [m]$, constraint (10) is equivalent to the optimal value P_i^* of the following optimization problem satisfying $P_i^* \le 1$:

$$\begin{aligned} P_i^* \triangleq \ \underset{\mathbf{c}_i}{\text{maximize}} \quad & \mathbf{c}_i^T\mathbf{m} \\ \text{subject to} \quad & F_i\mathbf{c}_i + \mathbf{g}_i \in \mathcal{K}_i, \\ & \mathbf{c}_i \succeq \mathbf{0}. \end{aligned}$$

The dual problem of the above is given by

$$\begin{aligned} \underset{\boldsymbol{\nu}_i}{\text{minimize}} \quad & \mathbf{g}_i^T\boldsymbol{\nu}_i \\ \text{subject to} \quad & F_i^T\boldsymbol{\nu}_i + \mathbf{m} \preceq \mathbf{0}, \\ & \boldsymbol{\nu}_i \in \mathcal{K}^*, \end{aligned}$$

where $\mathcal{K}^*$ is the dual cone of $\mathcal{K}$ [52]. Assume that a constraint qualification (e.g., Slater's condition [52]) holds, so that strong duality holds in this case. Then the optimal value of the dual problem is also given by P_i^*. Namely, there exists a dual feasible $\boldsymbol{\nu}_i$ such that $\mathbf{g}_i^T\boldsymbol{\nu}_i = P_i^*$. Therefore, the constraint in (10) is equivalent to:

$$\exists \boldsymbol{\nu}_i \in \mathcal{K}^* \text{ s.t. } F_i^T\boldsymbol{\nu}_i + \mathbf{m}(\mathbf{x}) \preceq \mathbf{0}, \ \mathbf{g}_i^T\boldsymbol{\nu}_i \le 1 \tag{12}$$

for each $i \in [m]$. For the uncertainty set used in our robust allocation problem, both $\mathcal{K}$ and $\mathcal{K}^*$ are the nonnegative orthant. Therefore, we can use the new constraints in (12) to replace those in (10) and rewrite the robust GP in (9) as

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^n_{++},\{\boldsymbol{\nu}_i\}_{i=1}^m}{\text{minimize}} \quad & f_0(\mathbf{x}) \\ \text{subject to} \quad & \boldsymbol{\nu}_i \succeq \mathbf{0}, \\ & F_i^T\boldsymbol{\nu}_i + \mathbf{m}(\mathbf{x}) \preceq \mathbf{0}, \quad \mathbf{g}_i^T\boldsymbol{\nu}_i \le 1, \\ & h_j(\mathbf{x}) = 1, \\ & \text{for all } i \in [m], \ j \in [p]. \end{aligned} \tag{13}$$

By applying the logarithmic transformation $y_i = \log x_i$ to (13) for all $i \in [m]$, we obtain

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^n,\{\boldsymbol{\nu}_i\}_{i=1}^m}{\text{minimize}} \quad & \widetilde{f}_0(y) \\ \text{subject to} \quad & \boldsymbol{\nu}_i \succeq \mathbf{0}, \\ & F_i^T\boldsymbol{\nu}_i + \widetilde{\mathbf{m}}(\mathbf{y}) \preceq \mathbf{0}, \quad \mathbf{g}_i^T\boldsymbol{\nu}_i \le 1, \\ & \mathbf{b}_j^T\mathbf{y} + \log d_j = 0, \\ & \text{for all } i \in [m], \ j \in [p], \end{aligned} \tag{14}$$

where $\widetilde{f}_0(\mathbf{y}) = f_0(\exp\{\mathbf{x}\})$ and $\widetilde{\mathbf{m}}(\mathbf{y}) = \mathbf{m}(\exp\{\mathbf{x}\})$ ($\exp\{\mathbf{x}\}$ is component-wise exponential). It can be shown that both $\widetilde{f}_0$ and the entries of $\widetilde{\mathbf{m}}$ are convex in $\mathbf{y}$, since they are nonnegative sums of exponentials of affine functions in $\mathbf{y}$ [52]. In fact, problem (14) is a convex problem and is a particular instance of a *conic geometric program* [1]. Problems in the form of (14) can be solved efficiently using off-the-shelf software such as CVX [53].

3.2 Robust Optimal Resource Allocation

In the following, we show how to formulate the optimization problem (6) as a conic GP using the methodology proposed in Section 3.1. In our derivations, we use the theory of nonnegative matrices and the Perron-Frobenius lemma [54]:

Lemma 2 (Perron-Frobenius). *Suppose M is an irreducible nonnegative matrix. Then, the spectral radius $\rho(M)$ of M satisfies:*

(a) $\rho(M) = \lambda_1(M) > 0$ is a simple eigenvalue of M;
(b) $M\mathbf{u} = \rho(M)\mathbf{u}$ for some $\mathbf{u} \in \mathbb{R}_{++}^n$;
(c) $\rho(M) = \inf\{\lambda \in \mathbb{R} : M\mathbf{u} \preceq \lambda\mathbf{u} \text{ for some } \mathbf{u} \succ \mathbf{0}\}$.

Remark 3. Note that the state matrix $B_{\mathcal{G}} + \mathrm{diag}(\mathbf{d}^c)$ of the linear system (2) is irreducible if the graph $\mathcal{G}$ is strongly connected. In what follows, we shall assume that the contact network $\mathcal{G}$ is strongly connected. This assumption is reasonable in the context of epidemic control, since the transportation network connecting different districts or subpopulations is strongly connected in most cases. Notice also that, as a consequence of this assumption, all the matrices in the uncertainty set $\Delta_{B_{\mathcal{G}}}$ are irreducible.

Using item (*c*) in the Perron-Frobenius lemma, the spectral radius $\rho(M)$ can be written as follows:

$$\begin{aligned}\rho(M) &= \inf\{\lambda : \exists \mathbf{u} \succ \mathbf{0} \text{ s.t. } M\mathbf{u} \preceq \lambda \mathbf{u}\} \\ &= \inf\left\{\lambda : \exists \mathbf{u} \succ \mathbf{0} \text{ s.t. } \max_{i\in[n]}\left\{\sum_{j=1}^n M_{ij}\frac{u_j}{u_i}\right\} \le \lambda\right\} \\ &= \inf\left\{\lambda : \inf_{\mathbf{u}\succ\mathbf{0}} \max_{i\in[n]}\left\{\sum_{j=1}^n M_{ij}\frac{u_j}{u_i}\right\} \le \lambda\right\} \\ &= \inf_{\mathbf{u}\succ\mathbf{0}} \max_{i\in[n]}\left\{\sum_{j=1}^n M_{ij}\frac{u_j}{u_i}\right\},\end{aligned} \tag{15}$$

where in the last equality we use the fact that $\inf\{\lambda : a \le \lambda\} = a$ for any a. Using (15), we rewrite the optimization problem (6) as

$$\begin{aligned}\min_{\mathbf{d}^c} \quad & \sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \inf_{\mathbf{u}\succ\mathbf{0}} \max_{i\in[n]}\left\{\sum_{j=1}^n M_{ij}(B_{\mathcal{G}}, \mathbf{d}^c)\frac{u_j}{u_i}\right\} \\ \text{s.t.} \quad & \sum_{i\in\mathcal{V}_C} g_i\big(\delta_i^c\big) \le C; \qquad \underline{\delta}_i^c \le \delta_i^c \le \overline{\delta}_i^c, \quad \forall i \in \mathcal{V}_C.\end{aligned} \tag{16}$$

In what follows, we will first approximate problem (16) as a robust GP with coefficient uncertainties using Lagrange duality. Observe that by exchanging the order of the sup and inf operators in problem (16), we can obtain a new problem

$$\begin{aligned}\min_{\mathbf{d}^c} \quad & \inf_{\mathbf{u}\succ\mathbf{0}} \sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \max_{i\in[n]}\left\{\sum_{j=1}^n M_{ij}(B_{\mathcal{G}}, \mathbf{d}^c)\frac{u_j}{u_i}\right\} \\ \text{s.t.} \quad & \sum_{i\in\mathcal{V}_C} g_i\big(\delta_i^c\big) \le C; \qquad \underline{\delta}_i^c \le \delta_i^c \le \overline{\delta}_i^c, \quad \forall i \in \mathcal{V}_C,\end{aligned}$$

whose optimal value yields an upper bound of the optimal value of problem (16) according to weak duality. Note that the above problem is equivalent to

$$\begin{aligned}\min_{\mathbf{d}^c,\mathbf{u}\succ\mathbf{0}} \quad & \max_{i\in[n]} \sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \left\{\sum_{j=1}^n M_{ij}(B_{\mathcal{G}}, \mathbf{d}^c)\frac{u_j}{u_i}\right\} \\ \text{s.t.} \quad & \sum_{i\in\mathcal{V}_C} g_i\big(\delta_i^c\big) \le C; \qquad \underline{\delta}_i^c \le \delta_i^c \le \overline{\delta}_i^c, \quad \forall i \in \mathcal{V}_C.\end{aligned}$$

If we introduce a slack variable

$$\lambda \triangleq \max_{i\in[n]} \sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \left\{\sum_{j=1}^n M_{ij}(B_{\mathcal{G}}, \mathbf{d}^c)\frac{u_j}{u_i}\right\},$$

we obtain the optimization problem described in the following proposition.

Proposition 4. *The optimal value of the robust allocation problem (6) is upper bounded by the following optimization problem*:

$$\begin{aligned}\underset{\mathbf{d}^c,\mathbf{u},\lambda}{\text{minimize}} \quad & \lambda \qquad (17)\\ \text{subject to} \quad & \sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \sum_{j=1}^n \beta_{ij}\frac{u_j}{u_i} + \delta_i^c \le \lambda, \quad i \in [n], \\ & \sum_{i\in\mathcal{V}_C} g_i\big(\delta_i^c\big) \le C; \qquad \underline{\delta}_i^c \le \delta_i^c \le \overline{\delta}_i^c, \quad \forall i \in \mathcal{V}_C, \qquad (18)\\ & \mathbf{u} \succ \mathbf{0}, \qquad \prod_{i=1}^n u_i = 1.\end{aligned}$$

The normalization constraint $\prod_{i=1}^n u_i = 1$ in (18) is included only for numerical purposes and does not affect optimality of the problem. We find it difficult to quantify the exact duality gap between the robust allocation problem (16) and the approximate problem (17). Furthermore, even if the approximation is exact, problem (17) is still nonconvex in general due to the non-convexity of the uncertainty set $\Delta_{B_{\mathcal{G}}}$, which will be discussed shortly. As a result, we choose to illustrate the goodness of approximation using numerical simulations presented in Section 4.

In the following, we discuss the conditions under which problem (17) can be solved efficiently using convex optimization. Notice that the decision variables $\mathbf{u}$ and λ of problem (17) are all strictly positive. When $\mathbf{d}^c$ is given, problem (17) is a robust GP in $\mathbf{u}$ and λ with coefficient uncertainties (in β_{ij}). In order to solve problem (17) using convex optimization, we need to make a few additional assumptions and relaxations. First of all, we need to relax $\Delta_{B_{\mathcal{G}}}$ into a convex set $\widehat{\Delta}_{B_{\mathcal{G}}}$ (i.e., $\Delta_{B_{\mathcal{G}}} \subseteq \widehat{\Delta}_{B_{\mathcal{G}}}$), so that the term

$$\sup_{B_{\mathcal{G}}\in\Delta_{B_{\mathcal{G}}}} \sum_{j=1}^n \beta_{ij}\frac{u_j}{u_i}$$

in the first constraint of problem (17) can be rewritten into a tractable form using the technique presented in Section 3.1, after we replace $\Delta_{B_{\mathcal{G}}}$ with $\widehat{\Delta}_{B_{\mathcal{G}}}$. Ideally, we would like to replace $\Delta_{B_{\mathcal{G}}}$ with its convex hull, so that the relaxation becomes exact. However, find the convex hull of the set (5) (which is a semialgebraic set) is a difficult task. Instead of attempting to find the convex hull, we give in Section 3.3 a transformation that relaxes $\Delta_{B_{\mathcal{G}}}$ into a convex polytope $\widehat{\Delta}_{B_{\mathcal{G}}}$ using the AM-GM inequality. Although such a relaxation is in general not exact, a polytopic $\widehat{\Delta}_{B_{\mathcal{G}}}$ satisfies the form of (11), so that the techniques in Section 3.1 can be applied.

We will assume that g_i is either a convex function or a posynomial for all $i \in \mathcal{V}_C$. When g_i is convex, the constraints involving $\mathbf{d}^c$ become convex, so that problem becomes a convex optimization problem after the relaxation of $\Delta_{B_{\mathcal{G}}}$ into a convex set $\widehat{\Delta}_{B_{\mathcal{G}}}$. When g_i is a posynomial, since the decision variable $\mathbf{d}^c$ is strictly positive, one can apply the logarithmic transformation on $\mathbf{d}^c$ (see the transformation between problems (7) and (8)) to convert constraints involving $\mathbf{d}^c$ into convex ones and obtain a convex optimization problem. We expect that the set of convex functions and posynomials is rich enough to model a wide variety of vaccine costs. For example,

posynomial functions can be used to fit with arbitrary accuracy any function that is convex in log-log scale. Such a feature of log-convexity is present in many cost functions representing the probability of failure versus investment of systems reliability [55]. Furthermore, there are well-developed numerical methods to fit posynomials to real data (see [51], Section 8).

3.3 Convex Set of Data-Coherent Networks

As we mentioned in Section 2.3, the uncertainty set $\Delta^D_{B_\mathcal{G}}$ is nonconvex since it is defined by a collection of polynomial equalities. In this section, we define a convex superset $\widehat{\Delta}^D_{B_\mathcal{G}} \supset \Delta^D_{B_\mathcal{G}}$, so that problem (6) becomes a conic geometric program after we substitute $\Delta^D_{B_\mathcal{G}}$ by $\widehat{\Delta}^D_{B_\mathcal{G}}$ (which changes the combined uncertainty set $\Delta_{B_\mathcal{G}}$). Moreover, since $\widehat{\Delta}^D_{B_\mathcal{G}}$ is a superset of $\Delta^D_{B_\mathcal{G}}$, if we replace $\Delta_{B_\mathcal{G}}$ in problem (17) with $\Delta_{B_\mathcal{G}} \triangleq \Delta^P_{B_\mathcal{G}} \cap \widehat{\Delta}^D_{B_\mathcal{G}}$, the optimal value of the new problem obtained still yields an upper bound of the optimal value of the original robust allocation problem (6).

We define the convex superset $\widehat{\Delta}^D_{B_\mathcal{G}}$ as follows:

$$\widehat{\Delta}^D_{B_\mathcal{G}} \triangleq \Bigg\{ B_\mathcal{G} \in \mathbb{R}^{n\times n} : \frac{1}{n}\sum_{j\in\mathcal{V}_S} \beta_{ij} p_j(t) \leq 1 - \left(1 - \frac{p_i(t+1) - p_i(t)\big(1-\delta_i^0\big)}{1-p_i(t)}\right)^{1/n} \quad \text{for all } i \in \mathcal{V}_S,\ t \in [T] \text{ s.t. } p_i(t) < 1 \Bigg\}. \tag{19}$$

Lemma 5. *The set $\widehat{\Delta}^D_{B_\mathcal{G}}$ is a superset of $\Delta^D_{B_\mathcal{G}}$.*

Proof. Consider any $i \in \mathcal{V}_S$. Recall that it always holds that $1 - \beta_{ij} p_j(t) \geq 0$. Then, from the AM-GM inequality, we have that

$$\prod_{j=1}^n [1 - \beta_{ij} p_j(t)] \leq \left(\frac{\sum_{j=1}^n [1 - \beta_{ij} p_j(t)]}{n}\right)^n = \left(1 - \frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t)\right)^n. \tag{20}$$

We can use (20) to yield a constraint on the transmission rates β_{ij} from the empirical dataset $\mathcal{D}$ defined in (4). Applying (20) to the nonlinear dynamics (1) results in the following inequality:

$$p_i(t+1) \geq p_i(t)(1-\delta_i^0) + (1 - p_i(t))\left\{1 - \left(1 - \frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t)\right)^n\right\},$$

where we have used the fact that the recovery rate δ_i is equal to the natural recovery rate δ_i^0 during empirical observations. Since $p_i(t) < 1$, we can rearrange the above inequality to obtain

$$\left(1 - \frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t)\right)^n \geq 1 - \frac{p_i(t+1) - p_i(t)(1-\delta_i^0)}{1-p_i(t)},$$

which is equivalent to

$$1 - \frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t) \geq \left(1 - \frac{p_i(t+1) - p_i(t)(1-\delta_i^0)}{1-p_i(t)}\right)^{1/n}. \tag{21}$$

Here we have used the fact

$$1 - \frac{p_i(t+1) - p_i(t)(1-\delta_i^0)}{1-p_i(t)} = \prod_{j=1}^n [1 - \beta_{ij} p_j(t)] > 0$$

according to the system dynamics (1). We can rearrange (21) to obtain

$$\frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t) \leq 1 - \left(1 - \frac{p_i(t+1) - p_i(t)(1-\delta_i^0)}{1-p_i(t)}\right)^{1/n}.$$

Finally, we use the fact

$$\frac{1}{n}\sum_{j\in\mathcal{V}_S} \beta_{ij} p_j(t) \leq \frac{1}{n}\sum_{j=1}^n \beta_{ij} p_j(t)$$

to complete the proof. □

The following comments are in order. First, the uncertainty set $\widehat{\Delta}^D_{B_\mathcal{G}}$ in (19) is defined by a collection of affine inequalities; therefore, it is a convex polytope and can be represented in the form of (11). Second, from Lemma 5, we have that the superset $\widehat{\Delta}^D_{B_\mathcal{G}}$ contains all the contact matrices $B_\mathcal{G}$ that are coherent with both prior information and empirical observations. In general, it is difficult to obtain analytical characterizations of the suboptimality induced by relaxing the original set $\Delta^D_{B_\mathcal{G}}$ to the new convex set $\widehat{\Delta}^D_{B_\mathcal{G}}$. In the following section, we illustrate our relaxation with numerical simulations and verify that the robust allocation is not overly conservative. In fact, for some realistic cases, the robust allocation achieves similar performance as the optimal allocation solved under a known contact matrix $B_\mathcal{G}$.

4 Simulations

In this section, we illustrate the robust data-driven allocation framework developed in Section 3. We consider the problem of controlling an epidemic outbreak propagating through the worldwide air transportation network [56]. The nodes in the network represent airports, whereas edges are flight connections for which we know passenger flows. Through our simulations, we demonstrate the following facts about the data-driven allocation framework. First, incorporating observations into the uncertainty set Δ_{B_G} significantly helps reduce the worst-case spreading rate bound. Second, the robust allocation framework does not need many observations to converge; in particular, the length of the observation period that we need is only a fraction of the number of nodes in the network. Finally, even though the robust allocation algorithm does not have access to the true underlying contact network $B_\mathcal{G}$, the resulting allocation achieves very similar performance compared to the optimal allocation solved using the actual $B_\mathcal{G}$.

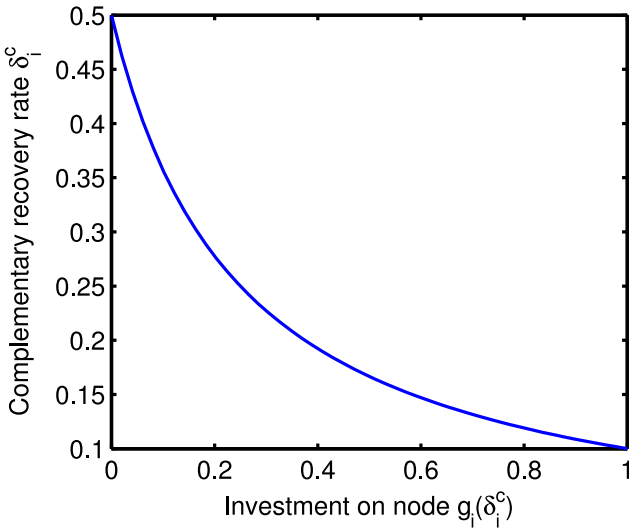

Fig. 1. Plot of the inverse of the vaccination cost function g_i^{-1}. This function represents the complementary recovery rate δ_i^c as a function of the investment on that node.

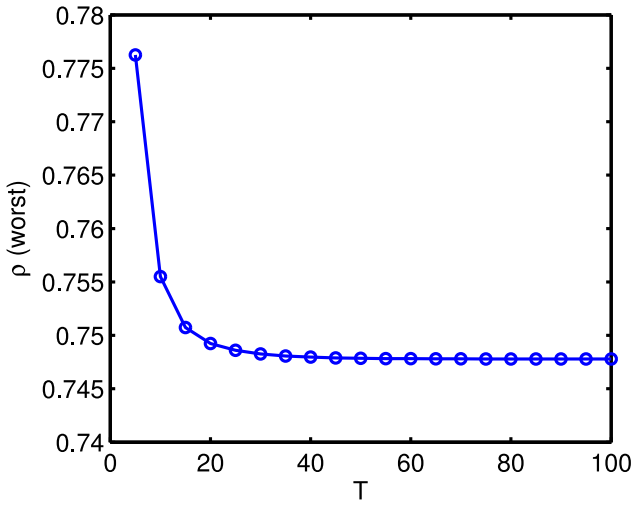

Fig. 2. Evolution of the worst-case spectral radius ρ_{wor} as a function of the number of observations T.

4.1 Numerical Setup

In our simulations, we consider the problem of controlling an epidemic outbreak propagating through a flight network comprised by the top 100 airports (based on yearly total traffic), so that $n = 100$. To illustrate the robust data-driven approach, we first generate a time series representing the dynamics of a hypothetical outbreak using the nonlinear dynamics (1). We run our simulation assuming a homogeneous value for the natural recovery rate, $\overline{\delta}_i^c = 0.5$ for all $i \in [n]$, and a link-dependent contact rate β_{ij} that is proportional to the traffic through that edge. Assuming an initial infection $p_i(0) = 0.5$ for all $i \in [n]$, we generate a time series $\{\mathbf{p}(t)\}_{t=1}^T$ representing the evolution of the infection over time.

In our data-driven framework, we assume that we do not have direct access to the matrix of infection rates $B_{\mathcal{G}}$. Instead, the data-driven allocation algorithm only has access to the observations $\{\mathbf{p}(t)\}_{t=1}^T$ for some period $t \in [T]$. Using this data, our algorithm generates an uncertainty set $\Delta_{B_{\mathcal{G}}}$ of data-coherent contact matrices. The parameters that define the uncertainty set $\Delta_{B_{\mathcal{G}}}$ are chosen as follows. For all $(i,j) \in \mathcal{E}$, unless otherwise noted, we assume an *a priori* upper bound $\overline{\beta}_{ij} = 1.5\beta_{ij}$ (i.e., the contact rate of an edge is at most 50 percent above its nominal contact rate), whereas the lower bound is $\underline{\beta}_{ij} = 0.5\beta_{ij}$ (i.e., the contact rate is at least half the nominal value). The natural recovery rate is chosen as $\overline{\delta}_i^c = 0.5$ for all $i \in [n]$, while the lower bound is chose to be $\underline{\delta}_i^c = 0.1$ for all $i \in [n]$ (i.e., the recovery rate δ_i is at most $1 - \underline{\delta}_i^c = 0.9$). We assume that the set of control nodes $\mathcal{V}_C = [n]$ and vary the set of sensing nodes.

To find the optimal allocation of vaccines, we consider the following vaccination cost function g_i given by

$$g_i(\delta_i^c) = \frac{1/\delta_i^c - 1/\overline{\delta}_i^c}{1/\underline{\delta}_i^c - 1/\overline{\delta}_i^c}$$

for all $i \in [n]$. The choice of g_i is inspired by the shape of the prototypical cost functions representing probability of failure versus investment in systems reliability [55]. These cost functions are usually quasiconvex and present diminishing returns. In our context, the infection rate plays a role similar to the probability of failure in systems reliability. It can be seen that g_i is a posynomial in δ_i^c, which makes the robust allocation problem (17) convex. The function g_i satisfies $g_i(\overline{\delta}_i^c) = 0$; namely, there is no cost by keeping δ_i^c as the natural complementary recovery rate $\overline{\delta}_i^c$. This function also satisfies $g_i(\underline{\delta}_i^c) = 1$; namely, the maximum allocation per node is one unit. Furthermore, the cost function is monotonically decreasing and exhibits diminishing returns (see Fig. 1). In this setup, our problem is to find the optimal allocation of vaccines throughout the airports in the air transportation network assuming we have a total budget equal to $C = 0.5n = 50$.

4.2 Results and Discussions

For any given uncertainty set $\Delta_{B_{\mathcal{G}}}$, we define the worst-case spectral radius $\rho_{\text{wor}}(\Delta_{B_{\mathcal{G}}})$ as the optimal value of the robust allocation problem (6). In other words, $\rho_{\text{wor}}(\Delta_{B_{\mathcal{G}}})$ represents the slowest exponential rate of disease eradication that can be achieved for those contact matrices that are coherent with our observations. In our first experiment, we illustrate the dependency of ρ_{wor} with respect to T, i.e., the number of observations available. In Fig. 2, we show the value of ρ_{wor} as we increase the observation period in the range $T = 1, \ldots, 100$. Notice how, as T grows, the amount of available information about the contact network increases and, as a result, ρ_{wor} decreases (i.e., we are able to guarantee a faster disease eradication). Notice also that the value of ρ_{wor} remains approximately unchanged after $T = 30$ observations. This result may seem surprising at first glance, since from the perspective of system observability, one would normally need as many time steps as the dimension of the system (in this case, $n = 100$) in order to identify the system. This demonstrates one of the benefits of using the robust allocation framework; namely, *it allows us to find an allocation without performing a previous system identification*.

In practice, the fraction of infected population $\mathbf{p}(t)$ is difficult to obtain directly. A common approach is to estimate $\mathbf{p}(t)$ from a random subpopulation, which leads to

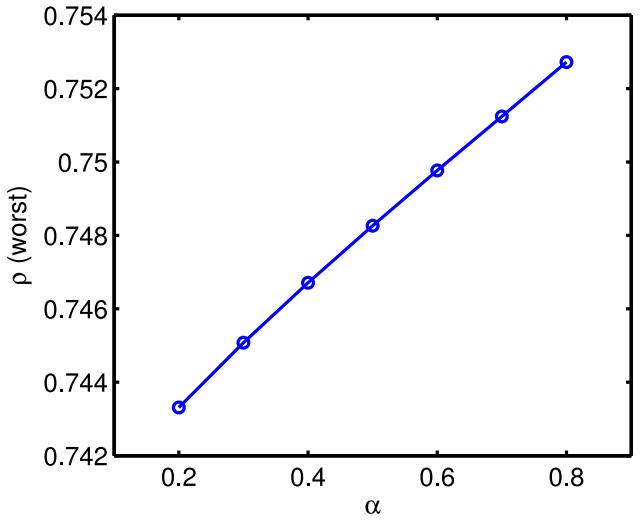

Fig. 3. Evolution of the worst-case spectral radius ρ_{wor} as a function of the *a priori* bounds on β_{ij}. For a given α, the upper and lower bounds on β_{ij} are given by $\overline{\beta}_{ij} = (1+\alpha)\beta_{ij}$ and $\underline{\beta}_{ij} = (1-\alpha)\beta_{ij}$, where β_{ij} is the nominal contact rate. The number of observations is chosen as $T = 30$.

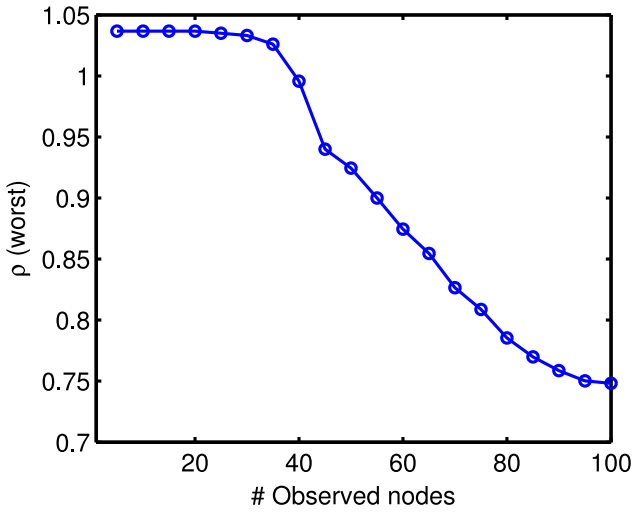

Fig. 4. Evolution of the worst-case spectral radius ρ_{wor} as a function of the number of observed nodes $|\mathcal{V}_S|$. The number of observations is chosen as $T = 30$.

random errors in $\mathbf{p}(t)$ that defines the uncertainty set $\Delta^D_{B_\mathcal{G}}$. While modeling of the errors and a rigorous treatment in minimizing the effect of errors are beyond the scope of this paper, we conduct a numerical test on the sensitivity of our optimization formulation to the errors in $\mathbf{p}(t)$. To this end, we choose $T = 30$ and introduce 5 percent relative errors in $\{\mathbf{p}(t)\}_{t=1}^T$ (i.e., each entry of $\mathbf{p}(t)$ is multiplied by a factor drawn uniformly from the interval $[0.95, 1.05]$), and the new observations are used in the robust allocation problem (6). The new worst-case spectral radius is obtained as $\rho'_{\text{wor}} = 0.6854$, which corresponds to 8.4 percent relative error compare to $\rho_{\text{wor}} = 0.7483$ obtained using the exact observation. The change in ρ_{wor} is expected, since $\mathbf{p}(t)$ changes the set $\Delta^D_{B_\mathcal{G}}$ and hence the set $\Delta_{B_\mathcal{G}}$ in the robust allocation problem (6). Nevertheless, this preliminary study of sensitivity indicates that the result is not extremely sensitive to errors in $\mathbf{p}(t)$.

In a second sets of experiments, we test how the prior knowledge on the upper bound $\overline{\beta}_{ij}$ and lower bound $\underline{\beta}_{ij}$ affect the allocation results. In particular, we choose $\overline{\beta}_{ij} = (1+\alpha)\beta_{ij}$ and $\underline{\beta}_{ij} = (1-\alpha)\beta_{ij}$ and vary α between 0 and 1; larger values of α imply that less is known about β_{ij} *a priori*. Fig. 3 shows how ρ_{wor} changes as a function of α for a given number of observations ($T = 30$). As α grows, we can see that ρ_{wor} also increases due to less knowledge about β_{ij}. It is evident from Fig. 3 that observation data indeed help restrict the size of the uncertainty set $\Delta_{B_\mathcal{G}}$; even for large values of α, the worst-case spectral radius ρ_{wor} still remains bounded with the incorporation of observations into the robust allocation problem.

Next, we numerically verify the performance of our data-driven allocation algorithm in the presence of sparse observations. In particular, we assume that we can only measure the evolution of the disease in a set of sensor nodes $\mathcal{V}_S$, which we choose to be those airports with the highest yearly total traffic. In Fig. 4, we plot the value of ρ_{wor} as we increase the number of sensor nodes from $|\mathcal{V}_S| = 1, \ldots, 100$. Notice how, as we increase the number of sensor nodes, ρ_{wor} decreases. Interestingly, for $|\mathcal{V}_S| \leq 20$ sensors, the value of ρ_{wor} hardly changes. In contrast, we observe a dramatic improvement in the value of ρ_{wor} for $|\mathcal{V}_S| \geq 40$. In fact, using only 40 sensors (out of 100 nodes), we can find an allocation that guarantees the eradication of the disease (i.e., $\rho_{\text{wor}} < 1$), even for the worst instantiation of $B_\mathcal{G}$ in $\Delta_{B_\mathcal{G}}$.

In our final simulation, we compare the allocation obtained from the data-driven framework with the allocation obtained assuming we have full access to the actual matrix of infection rates $B_\mathcal{G}$. Using the framework proposed in Preciado et al. [35], we can obtain the optimal allocation $\mathbf{d}^c_{\text{opt}}$, which is defined as the solution to the following optimization problem:

$$\begin{aligned} \underset{\mathbf{d^c}}{\text{minimize}} \quad & \rho(M(B_\mathcal{G}, \mathbf{d}^c)) \\ \text{subject to} \quad & \sum_{i \in \mathcal{V}_C} g_i(\delta^c_i) \leq C, \\ & \delta^c_i \leq \delta^c_i \leq \overline{\delta}^c_i, \quad i \in \mathcal{V}_C. \end{aligned} \tag{22}$$

The optimal value $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{opt}}))$ of problem (22) represents the fastest exponential rate at which the disease is eradicated when the contact network is completely known. As shown by Preciado et al. [35, Theorem 13], problem (22) can be converted into a geometric program and solved efficiently, as long as the cost function g_i can be expressed as a posynomial. Additionally, we denote by $\mathbf{d}^c_{\text{rob}}(T)$ the optimal solution to the robust data-driven allocation problem (6) when T time samples are available. We evaluate the spectral radius $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{rob}}(T)))$, which represents the exponential rate at which the disease is eradicated when we apply the allocation $\mathbf{d}^c_{\text{rob}}(T)$ to the actual contact network $B_\mathcal{G}$. In Fig. 5, we compare $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{rob}}(T)))$ with the optimal value $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{opt}}))$ for different values of T. Since $\mathbf{d}^c_{\text{opt}}$ is the optimal solution to problem (22), we always have $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{opt}})) \leq \rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{rob}}))$. However, Fig. 5 shows that the difference between $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{opt}}))$ and $\rho(M(B_\mathcal{G}, \mathbf{d}^c_{\text{rob}}(T)))$ is small for the particular network under investigation. The actual spectral radius is solely an empirical evaluation of the robust allocation $\mathbf{d}^c_{\text{rob}}$ and is not

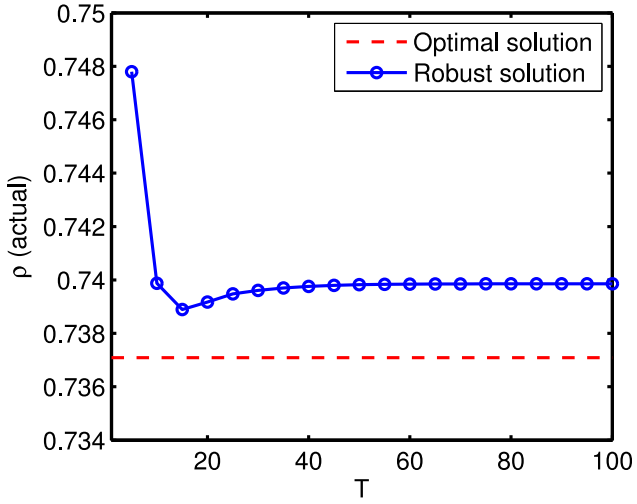

Fig. 5. The (actual) spectral radius $\rho(M(B_{\mathcal{G}}, \mathbf{d}^c))$ corresponding to both the optimal allocation $\mathbf{d}^c_{\text{opt}}$ and the robust allocation $\mathbf{d}^c_{\text{rob}}(T)$. Note that it is not guaranteed that $\rho(M(B_{\mathcal{G}}, \mathbf{d}^c_{\text{rob}}))$ decreases monotonically with T.

guaranteed to decrease as T increases. Namely, it is possible that the actual performance of the robust allocation becomes worse when more data are used. The only guarantee is that actual spectral radius is upper bounded by ρ_{wor}, which *always* monotonically decreases as T increases (as shown in Fig. 2).

Finally, it is worth mentioning that the robust data-driven allocation problem does not take significantly more time to solve than the optimal allocation problem. We have solved both allocation problems in MATLAB (R2012b) using CVX (Version 2.1, Build 1079) [53] with the Mosek solver (Version 7.0.0.106). All computations are carried out on a laptop computer equipped with a dual-core 2.5 GHz Intel Core i5 processor and 4 GB of RAM. For $n = 100$, the optimal allocation problem takes approximately 17 seconds to solve, whereas the robust allocation problem with *a priori* bounds on β_{ij} takes approximately 49 seconds for $T = 30$.

5 Conclusions

We have introduced a novel mathematical framework, based on conic geometric programming, to control a viral spreading process taking place in a contact network with unknown contact rates. We assume that we have access to time series data describing the evolution of the spreading process over a finite time period over a collection of sensor nodes. Using this data, we have developed a data-driven robust convex optimization framework to find the optimal allocation of protection resources over a set of control nodes to eradicate the viral spread at the fastest possible rate.

We have illustrated our approach using data obtained from the worldwide air transportation network. We have simulated a hypothetical epidemic outbreak over a finite time period and fed the resulting time series in our data-driven optimization algorithm. From our numerical results, we verify that (*i*) incorporating observations into the data-driven allocation algorithm significantly reduces the worst-case spreading rate bound; (*ii*) the robust allocation framework does not need many observations to converge; (*iii*) even though the robust allocation algorithm does not have access to the true underlying contact network $B_{\mathcal{G}}$, the resulting allocation achieves very similar performance compared to the optimal allocation solved under the actual $B_{\mathcal{G}}$.

Acknowledgments

This work was supported in part by the US National Science Foundation (NSF) under grants CNS-1302222, IIS-1447470, and CNS-1239224, and TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

[1] V. Chandrasekaran and P. Shah, "Conic geometric programming," in *Proc. Annu. Conf. Inf. Sci. Syst.*, 2014, pp. 1–4.
[2] N. Bailey, *The Mathematical Theory of Infectious Diseases and Its Applications*. London, U.K.: Griffin, 1975.
[3] R. M. Anderson, R. M. May, and B. Anderson, *Infectious Diseases of Humans: Dynamics and Control*, vol. 28. New York, NY, USA: Wiley, 1992.
[4] M. Newman, "Spread of epidemic disease on networks," *Phys. Rev. E*, vol. 66, no. 1, p. 016128, 2002.
[5] R. Pastor-Satorras and A. Vespignani, "Epidemic dynamics and endemic states in complex networks," *Phys. Rev. E*, vol. 63, p. 066117, May 2001.
[6] G. H. Weiss and M. Dishon, "On the asymptotic behavior of the stochastic and deterministic models of an epidemic," *Math. Biosci.*, vol. 11, no. 3, pp. 261–265, 1971.
[7] M. Draief and L. Massoulié, *Epidemics and Rumours in Complex Networks*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
[8] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
[9] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *Proc. ACM SIGKDD*, 2003, pp. 137–146.
[10] J. Leskovec, L. A. Adamic, and B. A. Huberman, "The dynamics of viral marketing," *ACM Trans. Web*, vol. 1, no. 1, p. 5, 2007.
[11] A. J. Ganesh, L. Massoulié, and D. F. Towsley, "The effect of network topology on the spread of epidemics," in *Proc. IEEE INFOCOM*, 2005, vol. 2, pp. 1455–1466.
[12] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 1–14, Feb. 2009.
[13] Y. Moreno, R. Pastor-Satorras, and A. Vespignani, "Epidemic outbreaks in complex heterogeneous networks," *The Eur. Phys. J. B*, vol. 26, pp. 521–529, 2002.
[14] V. M. Preciado and A. Jadbabaie, "Spectral analysis of virus spreading in random geometric networks," in *Proc. IEEE Conf. Decision Control*, 2009, pp. 4802–4807.
[15] V. M. Preciado and A. Jadbabaie, "Moment-based analysis of spreading processes from network structural information," *IEEE Trans. Autom. Control*, vol. 21, no. 2, pp. 373–382, Apr. 2013.
[16] R. Pastor-Satorras and A. Vespignani, "Epidemic dynamics in finite size scale-free networks," *Phys. Rev. E*, vol. 65, no. 3, p. 035108, 2002.
[17] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, "Epidemic spreading in real networks: An eigenvalue viewpoint," in *Proc. IEEE Symp. Reliable Distrib. Syst.*, 2003, pp. 25–34.
[18] R. Cohen, S. Havlin, and D. Ben-Avraham, "Efficient immunization strategies for computer networks and populations," *Phys. Rev. Lett.*, vol. 91, no. 24, p. 247901, 2003.
[19] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, "How to distribute antidote to control epidemics," *Random Struct. Algorithms*, vol. 37, no. 2, pp. 204–222, 2010.
[20] F. Chung, P. Horn, and A. Tsiatas, "Distributing antidote using pagerank vectors," *Internet Math.*, vol. 6, no. 2, pp. 237–254, 2009.
[21] Y. Wan, S. Roy, and A. Saberi, "Designing spatially heterogeneous strategies for control of virus spread," *IET Syst. Biol.*, vol. 2, no. 4, pp. 184–201, 2008.
[22] P. Van Mieghem, J. Omic, and R. Kooij, "Virus spread in networks," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 1–14, Feb. 2009.
[23] F. Sahneh and C. Scoglio, "Epidemic spread in human networks," in *Proc. IEEE Conf. Decision Control*, 2011, pp. 3008–3013.

[24] F. D. Sahneh, C. Scoglio, and P. Van Mieghem, "Generalized epidemic mean-field model for spreading processes over multilayer complex networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, pp. 1609–1620, Oct. 2013.
[25] P. Van Mieghem and J. Omic, "In-homogeneous virus spread in networks," *arXiv preprint arXiv:1306.2588*, 2013.
[26] H. J. Ahn and B. Hassibi, "Global dynamics of epidemic spread over complex networks," in *Proc. IEEE Conf. Decision Control*, 2013, pp. 4579–4585.
[27] A. Khanafer, T. Basar, and B. Gharesifard, "Stability properties of infected networks with low curing rates," in *Proc. Am. Control Conf.*, 2014, pp. 3579–3584.
[28] K. Drakopoulos, A. Ozdaglar, and J. N. Tsitsiklis, "An efficient curing policy for epidemics on graphs," *arXiv preprint arXiv:1407.2241*, 2014.
[29] X. Chen and V. M. Preciado, "Optimal coinfection control of competitive epidemics in multi-layer networks," in *Proc. IEEE Conf. Decision Control*, 2014, pp. 6209–6214.
[30] W. Mei and F. Bullo, "Modeling and analysis of competitive propagation with social conversion," in *Proc. IEEE Conf. Decision Control*, 2014, pp. 6203–6208.
[31] E. Ramirez-Llanos and S. Martinez, "A distributed algorithm for virus spread minimization," in *Proc. Am. Control Conf.*, 2014, pp. 184–189.
[32] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Stability analysis of generalized epidemic models over directed networks," in *Proc. IEEE Conf. Decision Control*, 2014, pp. 6197–6202.
[33] Y. Hayel, S. Trajanovski, E. Altman, H. Wang, and P. Van Mieghem, "Complete game-theoretic characterization of SIS epidemics protection strategies," in *Proc. IEEE Conf. Decision Control*, 2014, pp. 1179–1184.
[34] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, "Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks," in *Proc. IEEE Conf. Decision Control*, 2013, pp. 7486–7491.
[35] 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.
[36] (2015, Dec. 17). Google flu trends [Online]. Available: http://www.google.org/flutrends/about/
[37] Flu activity & surveillance—seasonal influenza (flu) [Online]. Available: http://www.cdc.gov/flu/weekly/fluactivitysurv.htm
[38] D. Materassi and G. Innocenti, "Topological identification in networks of dynamical systems," *IEEE Trans. Autom. Control*, vol. 55, no. 8, pp. 1860–1871, Aug. 2010.
[39] D. Materassi and G. Innocenti, "Unveiling the connectivity structure of financial networks via high-frequency analysis," *Physica A*, vol. 388, no. 18, pp. 3866–3878, 2009.
[40] J. Gonçalves and S. Warnick, "Necessary and sufficient conditions for dynamical structure reconstruction of LTI networks," *IEEE Trans. Autom. Control*, vol. 53, no. 7, pp. 1670–1674, Aug. 2008.
[41] Y. Yuan, G. B. Stan, S. Warnick, and J. Gonçalves, "Robust dynamical network structure reconstruction," *Automatica*, vol. 47, pp. 1230–1235, 2011.
[42] M. Timme, "Revealing network connectivity from response dynamics," *Phys. Rev. Lett.*, vol. 98, no. 22, p. 224101, 2007.
[43] M. Nabi-Abdolyousefi and M. Mesbahi, "Sieve method for consensus-type network tomography," *IET Control Theory Appl.*, vol. 6, no. 12, pp. 1926–1932, 2012.
[44] M. Nabi-Abdolyousefi and M. Mesbahi, "Network identification via node knockout," *IEEE Trans. Autom. Control*, vol. 57, no. 12, pp. 3214–3219, Dec. 2012.
[45] S. Shahrampour and V. M. Preciado, "Reconstruction of directed networks from consensus dynamics," in *Proc. Am. Control Conf.*, 2013, pp. 1685–1690.
[46] S. Shahrampuor and V. M. Preciado, "Topology identification of directed dynamical networks via power spectral analysis," *IEEE Trans. Autom. Control*, vol. 60, no. 8, pp. 2260–2265, Aug. 2015.
[47] M. Fazlyab and V. M. Preciado, "Robust topology identification and control of LTI networks," in *Proc. IEEE GlobalSIP Symp. Netw. Theory*, 2014, pp. 918–922.
[48] D. Napoletani and T. D. Sauer, "Reconstructing the topology of sparsely connected dynamical networks," *Phys. Rev. E*, vol. 77, no. 2, p. 26103, 2008.
[49] C. D. Michener and R. R. Sokal, "A quantitative approach to a problem in classification," *Evolution*, vol. 11, no. 2, pp. 130–162, 1957.
[50] D. Marinazzo, M. Pellicoro, and S. Stramaglia, "Kernel method for nonlinear granger causality," *Phys. Rev. Lett.*, vol. 100, no. 14, p. 144103, 2008.
[51] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," *Optim. Eng.*, vol. 8, no. 1, pp. 67–127, 2007.
[52] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge U.K.: Cambridge Univ. Press, 2004.
[53] CVX Research, Inc. (2012, Aug.). CVX: Matlab software for disciplined convex programming, version 2.0 [Online]. Available: http://cvxr.com/cvx
[54] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA, USA: SIAM, 2000.
[55] E. Nikolaidis, D. M. Ghiocel, and S. Singhal, eds., *Engineering Design Reliability Applications: For the Aerospace, Automotive and Ship Industries*. Boca Raton, FL, USA: CRC Press, 2007.
[56] C. M. Schneider, T. Mihaljev, S. Havlin, and H. J. Herrmann, "Suppressing epidemics with a limited amount of immunization units," *Phys. Rev. E*, vol. 84, p. 061911, Dec. 2011.

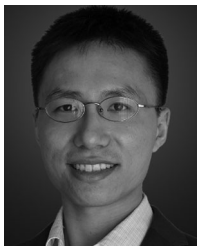

Shuo Han (S'08-M'14) received the BE and ME degrees in electronic engineering from Tsinghua University, Beijing, China in 2003 and 2006, respectively, and the PhD degree in electrical engineering from the California Institute of Technology, Pasadena in 2013. He is currently a postdoctoral researcher in the Department of Electrical and Systems Engineering at the University of Pennsylvania. His current research interests include control theory, convex optimization, applied probability, and their applications in large-scale interconnected systems. He is a member of the IEEE.

Victor M. Preciado received the PhD degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge in 2008. He is currently the Raj and Neera Singh assistant professor of electrical and systems engineering at the University of Pennsylvania. He is a member of the Networked and Social Systems Engineering (NETS) program and the Warren Center for Network and Data Sciences. His research interests include network science, dynamic systems, control theory, and convex optimization with applications in socio-technical networks, technological infrastructure, and biological systems.

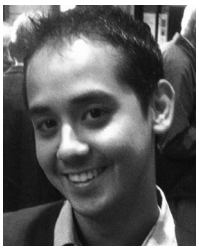

Cameron Nowzari received the BS degree in mechanical engineering from the University of California, Santa Barbara in June 2009, and the MS and PhD degrees in engineering sciences from the University of California, San Diego in December 2010 and September 2013, respectively. He is currently working as a postdoctoral research associate at the University of Pennsylvania. He 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 in the Theory category. 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, and spreading processes on networks.

George J. Pappas (S'90-M'91-SM'04-F'09) received the PhD degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, in 1998. He is currently the Joseph Moore professor and chair of the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA. He also holds a secondary appointment with the Department of Computer and Information Sciences and the Department of Mechanical Engineering and Applied Mechanics. He is a member of the GRASP Lab and the PRECISE Center. He had previously served as the deputy dean for research with the School of Engineering and Applied Science. His research interests include control theory and, in particular, hybrid systems, embedded systems, cyberphysical systems, and hierarchical and distributed control systems, with applications to unmanned aerial vehicles, distributed robotics, green buildings, and biomolecular networks. He is a fellow of the IEEE.

Dr. Pappas has received various awards, such as the Antonio Ruberti Young Researcher Prize, the George S. Axelby Award, the Hugo Schuck Best Paper Award, the George H. Heilmeier Award, the National Science Foundation PECASE award and numerous best student papers awards at ACC, CDC, and ICCPS.