Dynamic Management of Network Risk from Epidemic Phenomena

Aman Sinha, John Duchi, and Nicholas Bambos

***Abstract*— Despite the recent popularity of analyzing epidemic phenomena over networks, the budgeted control and protection of networks from viral propagations is not widely understood. In this paper, we blend methods from control theory and robust optimization to create a framework for network protection from epidemic environments using a limited control budget. As our emphasis is on the application of these methods to realistic contexts, our approach is designed to work with arbitrary networks, and it incorporates both decentralization as well as robustness to uncertainties in network topology. We illustrate tradeoffs between efficiency, robustness, and decentralization with respect to network protection, and we discuss methods that could build upon our framework to mitigate these tradeoffs.**

I. Introduction

The analysis of epidemic spreads in networks has gained recent interest in a variety of application domains. In addition to the obvious relevance to disease modeling and prevention, other phenomena that can be analyzed under the same theoretical framework include the spread of digital viruses, backbone router faults, viral marketing campaigns, influence in social media, and enterprise risk management.

Much of the existing literature on epidemic spreads focuses either on developing models for epidemic progression or analyzing existing epidemic models in "prototype" networks. In the former domain, the classic Susceptible-Infected-Recovered (SIR) and Susceptible-Infected-Susceptible (SIS) models, originally employed for aggregate populations without regard to any concept of a network topology (e.g. [1]), have been generalized to probabilistic models of infection over networks [2]–[4]. The latter domain includes analyses over a variety of topologies including infinite scale-free graphs [5], power-law graphs [6], and Erdős-Rényi random graphs [7], [8]. Unfortunately, the ability to rigorously analyze these structures often comes at the expense of their validity to realistic networks.

While the analysis of epidemic spread models and associated "prototype" networks serves important pedagogical purposes, what is often more important in many contexts is the ability to control epidemic spreads and/or protect networks from these spreads. The aforementioned literature on epidemic spreads, however, does not largely consider control. Numerous studies have explored heuristic methods [9], [10], and game theoretic analyses of strategies have also been applied to prototype networks [11]. However, recent optimization-oriented approaches have had the most success in controlling arbitrary directed networks composed of heterogeneous individuals with respect to a limited control budget [12]–[14]. Our approach identifies a framework that can be solved using techniques of convex optimization and, unlike previous approaches, systematically tackles issues of scalability and robustness.

We consider the problem of controlling the dynamics of interaction between a network and its environment in order to protect the network from the propagation of an epidemic. Similar to the optimization approaches in [12] and [13], our model consists of an SIS probabilistic model of infection for each agent within a network as well as an arbitrary directed communication topology between agents. In addition to the notion of a budget for our control policy, we require our method to incorporate decentralization (for scalability) and robustness to uncertainties in network topology, blending ideas from control theory with those of robust optimization. Together, these three characteristics of budget constraints, decentralization, and robustness to uncertainty make our proposed approach applicable to realistic scenarios.

This paper is organized as follows. In Sections II and III we present our model and define the problem under consideration. Subsequently, we outline our proposed approach in Section IV and illustrate its performance experimentally in Section V. We conclude with a review of our results, their implications, and open questions for further investigation.

II. Model Framework

A. System Dynamics

To capture the heterogeneity of realistic networks, we consider a generalization of the SIS epidemic model presented in [6] to include directed forms of interaction. Namely, we consider a connected system of N nodes with state given by $\mathbf{s} = [s_1, s_2, ..., s_N]^T \in \{0,1\}^N$. The boolean state s_i of node i is 0 when i is uninfected and 1 when i is infected with the contaminating virus. The structure of interaction between nodes is governed by a directed graph $\mathscr{G} = \{\mathscr{N}, \mathscr{E}, A\}$, where $\mathscr{N} = \{1,2,...,N\}$ is the set of nodes, $\mathscr{E} \subseteq \mathscr{N} \times \mathscr{N}$ is the set of edges, and $A \in \mathbb{R}_+^{N\times N}$ is the adjacency matrix governing the strength of interactions. In the context of epidemics, $[A]_{ij}$ governs the spreading rate from node j to node i, so it is reasonable to consider heterogeneous graphs in which $A \neq A^T$. A node cannot infect itself, so $[A]_{ii} = 0 \, \forall \, i$.

We model the health of the system as a continuous-time Markov process: each healthy node becomes infected at a rate proportional to the health of its neighbors and the strength of its pairwise interactions with them. Furthermore,

A. Sinha is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305. amans@stanford.edu

J. Duchi is with the Departments of Statistics and Electrical Engineering, Stanford University, Stanford, CA 94305. jduchi@stanford.edu

N. Bambos is with the Departments of Electrical Engineering and Management Science & Engineering, Stanford University, Stanford, CA 94305. bambos@stanford.edu

each node i recovers at a rate $r_i > 0$. Thus, the process is modeled as follows:

$$s_i(t) : \begin{cases} 0 \to 1 \text{ at rate } \mathbf{e}_i^T A\mathbf{s}(t) \\ 1 \to 0 \text{ at rate } r_i \end{cases} \tag{1}$$

where $\mathbf{e}_i$ denotes the i-th unit vector. Because $r_i > 0$, the process has a unique absorbing state $\mathbf{s} = \mathbf{0}$ that can be reached with positive probability from any other state, and $P(\mathbf{s}(t) \neq \mathbf{0}) \propto \exp(-\alpha t)$ for some $\alpha > 0$ [15]. Our goal is to protect this system from external threats. Thus, we now consider an external environmental force.

B. Environmental Forcing

We add a forcing term to the homogeneous system above in the form of an additional node that models the bulk environmental behavior surrounding the local system.[1] Defining an $(N+1)$-st state $s_{N+1} := s_{env}$, the total state is now $\hat{\mathbf{s}} = [\mathbf{s}^T, s_{env}]^T \in \{0,1\}^{N+1}$, and the dynamics can now be written as follows (for $i = \{1,2,...,N+1\}$):

$$s_i(t) : \begin{cases} 0 \to 1 \text{ at rate } \mathbf{e}_i^T \hat{A}\hat{\mathbf{s}}(t) \\ 1 \to 0 \text{ at rate } r_i \end{cases} \tag{2a}$$

$$\hat{A} = \left(\begin{array}{c|c} A & \mathbf{b} \\ \hline \mathbf{0} & 0 \end{array} \right). \tag{2b}$$

where $\mathbf{b} \in \mathbb{R}_+^N$ encodes the effects of the environment on the system, and we assume that the environment is large enough that s_{env} is not affected by $\mathbf{s}$. Importantly, we set $r_{i+1} := r_{env} \in \mathbb{R}_+$. In particular, we allow $r_{env} = 0$ since the environment need not be able to heal, in which case the local system $\mathbf{s}$ has an absorbing state when $\mathbf{b} = \mathbf{0}$ or $s_{env}(0) = 0$.

We denote $P(\mathbf{1}^T\mathbf{s}(t) > 0)$ as the instantaneous "energy of infection" at time t. We now show that we can upper bound this quantity by the 2-norm of a variable with linear dynamics.

Proposition 1: The Markov process dynamics (2) can be bounded in the following sense:

$$P(\mathbf{1}^T\mathbf{s}(t) > 0) \leq \sqrt{N}\|\mathbf{z}(t)\|_2, \tag{3}$$

where $\mathbf{z}(t)$ satisfies the following dynamics:

$$\frac{d}{dt}\mathbf{z}(t) = D\mathbf{z}(t) + \mathbf{b}e^{-r_{env}t}s_{env}(0), \tag{4a}$$

$$\mathbf{z}(0) = \mathbf{s}(0),\ D = A - \text{diag}(\mathbf{r}). \tag{4b}$$

Sketch of Proof: Following a similar style to the proof of *Theorem 3.1* in [6], we introduce a new variable $\hat{\mathbf{s}}^d = [(\mathbf{s}^d)^T, s_{env}^d]^T \in \mathbb{N}^{N+1}$ that evolves according to the dynamics:

$$s_i^d(t) : \begin{cases} j \to j+1 \text{ at rate } \mathbf{e}_i^T \hat{A}\hat{\mathbf{s}}^d(t) \\ j \to j-1 \text{ at rate } r_i s_i^d(t). \end{cases} \tag{5}$$

Define the following variables:

$$\mathbf{p}(t) := E[\mathbf{s}(t)],\ \hat{\mathbf{p}}(t) := E[\hat{\mathbf{s}}(t)] \tag{6a}$$

$$\mathbf{z}(t) := E[\mathbf{s}^d(t)],\ \hat{\mathbf{z}}(t) := E[\hat{\mathbf{s}}^d(t)] \tag{6b}$$

[1] Our approach can easily be generalized to multiple environmental nodes.

Then, $\hat{\mathbf{p}}(t)$ and $\hat{\mathbf{z}}(t)$ evolve as follows:

$$\frac{d}{dt}\hat{\mathbf{p}}(t) = \hat{D}\hat{\mathbf{p}}(t) - E\left[\text{diag}(\hat{\mathbf{s}}(t))\hat{A}\hat{\mathbf{s}}(t)\right] \tag{7a}$$

$$\frac{d}{dt}\hat{\mathbf{z}}(t) = \hat{D}\hat{\mathbf{z}}(t) \tag{7b}$$

$$\hat{D} = \left(\begin{array}{c|c} D & \mathbf{b} \\ \hline \mathbf{0} & -r_{env} \end{array} \right). \tag{7c}$$

Note that $\text{diag}(\hat{\mathbf{s}}(t))\hat{A}\hat{\mathbf{s}}(t) \succeq \mathbf{0}$. Given equal deterministic initial conditions, $\hat{\mathbf{s}}^d(0) = \hat{\mathbf{s}}(0)$, we claim that

$$\hat{\mathbf{p}}(t) \preceq \hat{\mathbf{z}}(t), \tag{8}$$

but we postpone its full proof (by analyzing the dynamics of the difference between $\hat{\mathbf{z}}$ and $\hat{\mathbf{p}}$) for a more comprehensive study. The intuition behind this argument is that s_i^d leaves 0 at least as fast as s_i, and both variables fall back to 0 at the same rate. By using this claim and applying the Markov and Cauchy-Schwartz inequalities to the energy of infection, we have:

$$\begin{aligned} P(\mathbf{1}^T\mathbf{s}(t) > 0) &= P(\mathbf{1}^T\mathbf{s}(t) \geq 1) \\ &\leq \mathbf{1}^T\mathbf{p}(t) \\ &\leq \mathbf{1}^T\mathbf{z}(t) \\ &\leq \sqrt{N}\,\|\mathbf{z}(t)\|_2 . \end{aligned} \tag{9}$$

Using the equal deterministic initial conditions, we can easily expand (7b) and (7c) into the final form (4). ■

III. Problem Formulation

We aim to control the effects of the environment on the local system with a limited set of resources, i.e. a budget constraint. The strength of node i's interaction with the environment is given by b_i of $\mathbf{b}$. We consider modifying these values in a linear manner through a control vector $\mathbf{v}$, which results in a nontrivial problem if our budget does not allow us to eliminate environmental interaction completely. Including $\mathbf{v}$, the dynamics are

$$\frac{d}{dt}\mathbf{z}(t) = D\mathbf{z}(t) + \mathbf{n}(t) \tag{10a}$$

$$\mathbf{n}(t) = (\mathbf{b} - \mathbf{v}(t))e^{-r_{env}t}s_{env}(0). \tag{10b}$$

Of course, in many realistic scenarios, $\mathbf{v}$ cannot be changed continuously in time. Instead, we consider the scenario where resources are reallocated at regular intervals h (i.e. $\mathbf{v}$ is piecewise-constant). Defining the discrete-time variables $\mathbf{x}(k) := \mathbf{z}(kh)$ and $\mathbf{w}(k) := \mathbf{v}(t), t \in [kh,(k+1)h)$, we have:

$$\mathbf{x}(k+1) = F\mathbf{x}(k) + G\mathbf{u}(k) \tag{11a}$$

$$F := e^{hD},\ G := \int_0^h e^{D(h-\tau)}e^{-r_{env}\tau}d\tau \tag{11b}$$

$$\mathbf{u}(k) = (\mathbf{b} - \mathbf{w}(k))e^{-r_{env}kh}s_{env}(0) \tag{11c}$$

$$\mathbf{0} \preceq \mathbf{w}(k) \preceq \mathbf{b},\ \|\mathbf{w}(k)\|_1 \leq c, \tag{11d}$$

where (11d) denotes the constraints on our control input. Control resources are nonnegative ($\mathbf{0} \preceq \mathbf{w}$), the total budget per time period h is c, and the best that can ever be done is to isolate the system from the environment ($\mathbf{w} \preceq \mathbf{b}$). Our objective is to minimize J, an upper bound on the total integrated energy of infection of the system:

$$J = \sqrt{N} \int_0^\infty \|\mathbf{z}(t)\|_2 dt \geq \int_0^\infty P(\mathbf{1}^T \mathbf{s}(t) > 0) dt, \tag{12}$$

which can be approximated by the following:

$$J \approx \sqrt{N} \sum_{k=0}^{T} \|\mathbf{x}(k)\|_2 \tag{13}$$

for small h and large T. As stated, this is a convex optimization problem that can be solved using the receding-horizon or model predictive control framework (MPC) [16] . In other words, at time m, we solve the problem

$$\begin{aligned} &\text{minimize } J_m := \sqrt{N} \sum_{k=m+1}^{T+m} \|\mathbf{x}(k)\|_2 \\ &\text{subject to } (11), \end{aligned} \tag{14}$$

apply the policy values $\mathbf{u}(m)$, and then repeat for time $m+1$. The MPC framework allows us to easily account for exogenous changes in model parameters such as r_{env} by updating values before solving the optimization problem for successive time steps.

A. Decentralization and Robustness Requirements

For large networks (large N), the centralized approach above can become highly inefficient and unrealistic, so we seek a decentralized solution method that approximates the centralized MPC method. Specifically, we consider the case where the system is divided into M groups with N_i nodes in the i-th group, such that $\sum_i N_i = N$. We denote the state of the i-th group as $\mathbf{x}^i$, so $\mathbf{x} = [(\mathbf{x}^1)^T, ..., (\mathbf{x}^M)^T]^T$. Similarly, each group i sets the corresponding elements of the control vector $\mathbf{u}^i$ via $\mathbf{w}^i$.[2]

Furthermore, we consider the situation where each group i knows its intrinsic local dynamics but has uncertainty about its interaction with neighboring groups. In other words, we assume that the diagonal blocks D_{ii} are known to each group i, but the off-diagonal blocks D_{ij} are known only within a certain uncertainty region to groups i and j ($i \neq j$). In addition to the decentralization requirement stipulated above, our control policy should also be robust to these uncertainties in group-wise interactions.

IV. Proposed Approach

We tackle the issues of decentralization and robustness to uncertainty through the use of reduced-order models and techniques from robust optimization.

A. Reduced-Order Models

A naive approach to ensuring decentralization is to allocate a fraction of the total budget to each group and have it ignore interactions with other groups when choosing a control policy. This can be highly suboptimal. On the other hand, making each group find a solution with respect to the entire matrix D defeats the purpose of decentralization, i.e. that each group solves a smaller, more manageable problem. Thus, we allow each group to account for its interactions with others via reduced-order models, with inherent tradeoffs between the degree of decentralization (corresponding to the size of the reduced-order models) and the accuracy of each group's corresponding model.

Although this approach can work for arbitrary M, we describe the case with $M = 3$ for simplicity. We postpone analysis of uncertainties in D_{ij} until Section IV-B. For all subsequent analysis, we assume that D_{ii} is Hurwitz $\forall\, i$, i.e. that the local system needs protection or isolation from the environment and not from any of its own nodes. Without loss of generality, any unstable components of the local system can simply be modeled as part of the environment.

Consider the local continuous dynamics for the first group. The state-space system can be summarized by the following:

$$S^1 = \left[\begin{array}{c|c} D_{11} & (D_{12},\, D_{13},\, I) \\ \hline I & \end{array}\right] \tag{15}$$

with state $\mathbf{z}^1$ and input $[(\mathbf{z}^2)^T, (\mathbf{z}^3)^T, (\mathbf{n}^1)^T]^T$. A variety of methods can be used to perform model reduction, reducing the state of system i to size $k_i \ll N_i$, such as using balanced truncation or via a Hankel-norm approximation [17]–[19]. Either of these methods involves finding the controllability and observability gramians (W_C^1 and W_O^1 respectively). Note that both W_C^i and W_O^i are positive definite, since all D_{ii} are Hurwitz and the input/output matrices of S^i are full rank. Namely, W_C^1 and W_O^1 solve the following Lyapunov equations (with similar expressions for the other groups):

$$D_{11} W_C^1 + W_C^1 D_{11}^T + (D_{12},\, D_{13},\, I)\,(D_{12},\, D_{13},\, I)^T = 0 \tag{16a}$$

$$D_{11}^T W_O^1 + W_O^1 D_{11} + I = 0. \tag{16b}$$

We omit a derivation of the well-known model reduction procedure (see, e.g., [17]–[19]). However, we note that the reduced system is developed through a transformation analogous to a standard similarity transformation. Specifically, the reduced system for the first group is

$$\begin{aligned} \tilde{S}^1 &= \left[\begin{array}{c|c} T_L^T D_{11} T_R & T_L^T\,(D_{12},\, D_{13},\, I) \\ \hline T_R & \end{array}\right] \\ &:= \left[\begin{array}{c|c} \tilde{D}_{11} & \left(\tilde{D}_{12},\, \tilde{D}_{13},\, \tilde{B}^1\right) \\ \hline \tilde{C}^1 & \end{array}\right] \end{aligned} \tag{17}$$

corresponding to the reduced state $\tilde{\mathbf{z}}^1$, where T_L and T_R are developed via manipulations of the singular value decomposition for $W_C^1 W_O^1$ [17]. The initial condition is given by $\tilde{\mathbf{z}}^1(0) = \tilde{B}^1 \mathbf{z}^1(0)$.[3] It is helpful to think of $\tilde{B}^1$ and $\tilde{C}^1$ behaving as compression and expansion operators respectively. When reduced to order k_i, the relative error of the associated transfer function, denoted as $\tilde{S}^i(j\omega)$, is proportional to the sum of the $N_i - k_i$ smallest Hankel singular values:

$$\left\| S^i - \tilde{S}^i \right\|_\infty \leq 2 \sum_{p=k_i+1}^{N_i} \sigma_p^i \tag{18a}$$

[2]We can write equivalent expressions $\mathbf{z}^i$, $\mathbf{n}^i$, and $\mathbf{v}^i$ for the continuous-time dynamics.

[3]Note that $\tilde{B} = (\tilde{C})^\dagger$, where $(\cdot)^\dagger$ is the Moore-Penrose pseudoinverse.

$$\sigma_p^i = \sqrt{\lambda_p(W_C^i W_O^i)}, \tag{18b}$$

where $\lambda_p(A)$ is the p-th largest eigenvalue of positive semidefinite matrix A. Intuitively, the reduced-order system captures the modes of behavior that interact the most with other groups and the environment (i.e. the most controllable/observable modes). In another sense, (18) captures the tradeoff between decentralization and the accuracy of the models used for developing control policies.

Now we employ these reduced models to incorporate the dynamics of groups 2 and 3 into group 1's optimization. Specifically, the first group considers the reduced state $\mathbf{z}_r^1 = [(\mathbf{z}^1)^T, (\tilde{\mathbf{z}}^2)^T, (\tilde{\mathbf{z}}^3)^T]^T$ as well as a reduced input $\mathbf{n}_r^1$ evolving according to the following state-space dynamics:

$$\frac{d}{dt}\mathbf{z}_r^1(t) = D_r^1 \mathbf{z}_r^1(t) + \mathbf{n}_r^1(t) \tag{19a}$$

$$D_r^1 = B_r^1 D C_r^1 \tag{19b}$$

$$\mathbf{n}_r^1(t) = (B_r^1 \mathbf{b} - \mathbf{v}_r^1(t)) e^{-r_{env} t} s_{env}(0) \tag{19c}$$

$$B_r^1 = \text{diag}\left(I,\, \tilde{B}^2,\, \tilde{B}^3\right),\ C_r^1 = \text{diag}\left(I,\, \tilde{C}^2,\, \tilde{C}^3\right), \tag{19d}$$

where $\text{diag}(\cdot)$ forms a block diagonal matrix from its matrix arguments, and $\mathbf{v}_r^1$ is the reduced control input corresponding to $\mathbf{n}_r^1$. Defining appropriate discrete-time counterparts, group 1 solves the following optimization problem at time step m:

$$\begin{aligned}
&\text{minimize } J_m^1 := \sqrt{N} \sum_{k=m+1}^{T+m} \|\mathbf{x}_r^1(k)\|_2 \\
&\text{subject to} \\
&\mathbf{x}_r^1(k+1) = F_r^1 \mathbf{x}_r^1(k) + G_r^1 \mathbf{u}_r^1(k), \\
&F_r^1 := e^{hD_r^1},\ G_r^1 := \int_0^h e^{D_r^1(h-\tau)} e^{-r_{env}\tau} d\tau, \\
&\mathbf{u}_r^1(k) = (B_r^1 \mathbf{b} - \mathbf{w}_r^1(k)) e^{-r_{env} kh} s_{env}(0) \\
&\mathbf{0} \preceq C_r^1 \mathbf{w}_r^1(k) \preceq \mathbf{b},\ \left\|C_r^1 \mathbf{w}_r^1(k)\right\|_1 \leq c,
\end{aligned} \tag{20}$$

The main difference with the formulation in Section III is that the dynamics of (10) are replaced by those of (19), which includes a reduction of the environmental and control inputs for other groups.[4] Groups 2 and 3 use similar reduced-system dynamics to minimize J_m^2 and J_m^3, and during each step of the decentralized MPC scheme, each group i only enacts its local portion of $\mathbf{u}_r^i$.

The level of communication required between groups depends on the precision with which we wish to carry out this decentralized scheme. The initial communication of reduced models as well as initial conditions for reduced states is required, after which no communication is necessary. To improve precision, however, we can allow the communication of reduced states between each time step to "reset" each local receding-horizon scheme with the proper initialization.

Finally, we observe that the collective budget constraint $\|\mathbf{w}(k)\|_1 \leq c$ is not necessarily satisfied under this scheme, since we combine policies from each group to form the collective policy. In this study, we consider the simple approach of scaling the collective policy $\mathbf{w}$ by $c/\|\mathbf{w}\|_1$ when necessary, which requires a sufficiently small amount of communication between groups at each time step. More rigorous approaches to optimally partitioning group budgets in a decentralized manner are under further investigation.

B. Incorporating Robustness Against Uncertainties

Now we consider the robust counterpart to the MPC scheme above that accounts for uncertainties in D_{ij}. The tractability of such a robust optimization problem rests on the uncertainty set D_{ij}. We employ the so-called "scenario" or polytopic uncertainty set, wherein each D_{ij} is a convex combination of a finite set of matrices:

$$\mathscr{D}_{ij} = \{D_{ij} | D_{ij} = \sum_{k=1}^{L_{ij}} \mu_k D_{ij}(k),\ \mu_k \geq 0,\ \sum_{k=1}^{L_{ij}} \mu_k = 1\}, \tag{21}$$

where L_{ij} is the number of matrices making up the vertices or "corners" of the uncertainty set, and $D_{ij}(k)$ is the k-th vertex or "corner" case of this set. We need to create robust counterparts for both model reduction as well as the optimization problem (20). For model reduction, we employ generalized balanced truncation, which replaces the Lyapunov equation for the controllability gramian with a linear matrix inequality (LMI) [20]. For example, using $M = 3$ as above, we find a (non-unique) V_C^1 satisfying:

$$D_{11} V_C^1 + V_C^1 D_{11}^T + (D_{12},\, D_{13},\, I)\,(D_{12},\, D_{13},\, I)^T \preceq 0, \tag{22}$$

and the corresponding errors for the reduced models are [20]:

$$\|S^i - \tilde{S}^i\|_\infty \leq 2 \sum_{p=k_i+1}^{N_i} \gamma_p^i \tag{23a}$$

$$\gamma_p^i = \sqrt{\lambda_p(V_C^i W_O^i)} \geq \sigma_p^i. \tag{23b}$$

Proposition 2: Define the matrix $B(i,j)$ as the following:

$$B(i,j) := (D_{12}(i), D_{13}(j)). \tag{24}$$

where $i \in \{1,...,L_{12}\}$ and $j \in \{1,...,L_{13}\}$. Furthermore, define $L_1 := L_{12}L_{13}$ as the total number of such matrices. Then given a matrix $V_C^1 \succeq 0$, feasibility of (22) for all $(D_{12}, D_{13}) \in \mathscr{D}_{12} \times \mathscr{D}_{13}$ is implied by the feasibility of the following:

$$D_{11} V_C^1 + V_C^1 D_{11}^T + I + L_1 B(i,j) B(i,j)^T \preceq 0\, \forall\, (i,j) \tag{25}$$

Sketch of Proof: Note that $\mathscr{D}_{12} \times \mathscr{D}_{13}$ is the convex hull of all $B(i,j)$ matrices. That is, $\forall\, (D_{12}, D_{13}) \in \mathscr{D}_{12} \times \mathscr{D}_{13}$:

$$(D_{12}, D_{13}) = \sum \mu_{ij} B(i,j),\ \mu_{ij} \geq 0,\ \sum \mu_{ij} = 1. \tag{26}$$

Now consider any general matrix $X \in \mathbb{R}^{m \times n}$ which can be written as a convex combination of K matrices:

$$X = \sum_{i=1}^K \nu_i X_i,\ \ \nu_i \geq 0,\ \sum_{i=1}^K \nu_i = 1. \tag{27}$$

Since $(\nu_i X_i - \nu_j X_j)(\nu_i X_i - \nu_j X_j)^T \succeq 0$, we have:

$$\nu_i^2 X_i X_i^T + \nu_j^2 X_j X_j^T \succeq \nu_i \nu_j \left(X_i X_j^T + X_j X_i^T\right). \tag{28}$$

[4] If the local system is still too large, we can set $\mathbf{z}_r^1 = [(\tilde{\mathbf{z}}^1)^T, (\tilde{\mathbf{z}}^2)^T, (\tilde{\mathbf{z}}^3)^T]^T$, i.e. use a reduced form of group 1's own dynamics.

Then,

$$XX^T \preceq K\sum_{i=1}^{K} \nu_i^2 X_i X_i^T \preceq K\sum_{i=1}^{K} \nu_i X_i X_i^T, \tag{29}$$

where the latter inequality follows from the fact that $\nu_i \in [0,1]$. Now, assume that we have the following:

$$Y + KX_iX_i^T \preceq 0 \;\forall\, i. \tag{30}$$

Taking a convex combination of these inequalities yields:

$$Y + K\sum_{i=1}^{K} \nu_i X_i X_i^T \preceq 0 \implies Y + XX^T \preceq 0. \tag{31}$$

Substituting $D_{11}V_C^1 + V_C^1 D_{11}^T + I$ for Y, L_1 for K, $B(i,j)$ for X_i, and μ_{ij} for ν_i, we readily observe that (25) $\implies$ (22). ■

Although (25) is the robust counterpart of the specific LMI (22) for group 1 with $M = 3$, the result is completely general. In other words, to find V_C^i satisfying

$$D_{ii}V_C^i + V_C^i D_{ii}^T + (B,\, I)\,(B,\, I)^T \preceq 0, \tag{32}$$

$\forall\, B \in \prod_{j\neq i} \mathscr{D}_{ij}$, it is sufficient to simultaneously satisfy L_i "corner case" LMI's equivalent to (25), where L_i is the number of "corners" of the set $\prod_{j\neq i} \mathscr{D}_{ij}$.

To generate a robust counterpart of the optimization problem (20), we first recognize that, since all D_{ii} are fixed and $\tilde{D}_{ij}$ are simply linear functions of D_{ij}, we can parametrize D_r^i by $D_{ij}\;\forall\, i \neq j$. Thus the uncertainty set $\mathscr{D}_r^i$ is parametrized by the Cartesian product $\prod_{i\neq j} \mathscr{D}_{ij}$, which has $L := \prod_i L_i$ vertices or corner cases, denoted by $D_r^i(p)$, $p \in \{1,...,L\}$. We denote $F_r^i(p)$ as F_r^i evaluated with the corner case $D_r^i(p)$.

The robust counterpart to (20) can be written as a second-order cone program (SOCP) if we can write $\mathbf{x}_r^i(k)$ as an affine function of D_r^i. This can be approximated by Taylor-expanding the objective. Specifically, we assume that h is small enough that all terms greater than first-order in (hD_r^i) and second-order in $(r_{env}h)$ are negligible. Then we can write the following approximations (which are valid for any r_{env}):

$$(F_r^i)^k \approx I + khD_r^i \tag{33a}$$

$$(F_r^i)^k G_r^i \approx aI + bD_r^i \tag{33b}$$

$$a := h\left(1 - \frac{r_{env}h}{2}\right),\; b := h^2\left(\frac{1}{2} + k\left(1 - \frac{r_{env}h}{2}\right)\right) \tag{33c}$$

With this expansion, we minimize $\sup_{D_r^i} J_m^i$ by solving the following SOCP with auxiliary variable $\mathbf{g} \in \mathbb{R}_+^T$ [21]:

$$\begin{aligned}
&\text{minimize } \;\|\mathbf{g}\|_1\\
&\text{subject to }\; (33),\\
&\left\| F_r^i(p)^{k-m}\mathbf{x}_r^i(m) + \sum_{\tau=m}^{k-1} F_r^i(p)^{k-1-\tau} G_r^i \mathbf{u}_r^i(\tau) \right\|_2 \le g_{k-m}\\
&\mathbf{u}_r^i(k) = (B_r^i\mathbf{b} - \mathbf{w}_r^i(k))e^{-r_{env}kh}s_{env}(0),\\
&\mathbf{0} \preceq C_r^i\mathbf{w}_r^i(k) \preceq \mathbf{b},\;\; \left\|C_r^i\mathbf{w}_r^i(k)\right\|_1 \le c,\\
&\forall\; k \in \{m+1,...,T+m\},\;\; p \in \{1,...,L\},
\end{aligned} \tag{34}$$

where the expressions inside $\|\cdot\|_2$ are the discrete convolution expansions for $\mathbf{x}_r^i(k)$. Thus, using polytopic uncertainties requires the one-time communication of uncertainty sets between nodes (along with reduced-order models) at the beginning of the optimization procedure.

Polytopic sets certainly appear to be the simplest form of characterizing uncertainties in D_{ij}. However, we observe that this simplicity of expression for the robust counterparts to the LMI (16a) and MPC (20) problems comes at the expense of their computation: the number of LMIs needed to robustly reduce models and the number of constraints needed to robustly minimize J_m^i grow exponentially in M and M^2 respectively (assuming the constraint sets are uncorrelated). This fact underlies the necessity of considering other uncertainty sets, such as norm-bounded uncertainties. More tractable uncertainty sets have the potential to assuage the tradeoff we observe between scalability and robustness.

V. Experimental Evaluations

We consider a model system of $N = 24$ nodes and $M = 3$ groups of equal size $N_i = N/M$. We generate adjacency matrices A and recovery rates $\mathbf{r}$ uniformly at random and scale them such that $\lambda_i(D) \in [-1, -0.33]$. Furthermore, we choose $s_{env}(0) = 1$ and $r_{env} = 0.2$ such that the environment heals (i.e. it is asymptotically stable), but it dies at a slower rate than the local system. Thus, control in the form of protection from the environment can accomplish something meaningful for this system, particularly at early times. We set $h = 0.05$, $T = 20$. Finally, the environmental forcing vector $\mathbf{b}$ is chosen with random elements, but the magnitudes of environmental interaction with half of the nodes in groups 1 and 2 are set substantially higher than for all other nodes. This makes some degree of cooperation between groups imperative for system-wide success. We assume that the groups have no a priori knowledge of this heterogeneity.

Each D_{ij} is assumed to be the convex combination of three possible "corner" matrices. For simplicity, we introduce correlations between the uncertainty sets. Specifically, we assume that all uncertainty sets $\mathscr{D}_{ij}\;\forall\; i \neq j$ are perfectly correlated, so that the matrix D is a convex combination of three "corner" cases. This reduces the effective number of vertices in $\mathscr{D}_r^i$ from 729 to 3.

We evaluate our approach with numerous sizes of reduced models $k_i = \{0,2,4,6,8\}$. Note that $k_i = 8$ corresponds to the global centralized solution. The case of $k_i = 0$ corresponds to almost complete anarchy; each group ignores the others and assumes it has access to the entire budget, but the collective policy is scaled if it is too large (as mentioned in Section IV-B). We compare these policies with two baseline cases: no control input and complete anarchy. In complete anarchy, each group assumes it is isolated from the others and has access to a fixed share c/M of the budget. This corresponds to the aforementioned situation where there is no prior knowledge for the extent to which the environment interacts with each group.

In all models, the actual state dynamics between MPC iterations are updated through a specific D matrix that is unknown to any of the controllers. Figure (1) shows (the upper bound on) instantaneous energy of infection over time

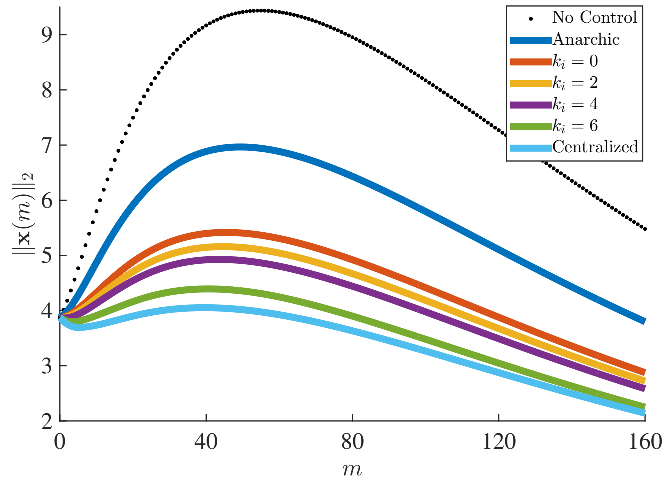

Fig. 1. Comparison of the effects of policies on instantaneous energy of infection. The major benefits of communication occur at early times since the environment and network are both asymptotically stable. Communication and dynamic budget allocation between groups assuage the severity of "overshoot" before asymptotically approaching a state of zero energy. Larger k_i results in better performance.

for all MPC policies. Despite the fact that the environment is asymptotically stable, there is a dramatically unfavorable interaction with the system at early times in the case of no control. This highlights the necessity of protection from the environment, even in situations without instability. Increasing the order of reduced models improves performance and reduces the magnitude of "overshoot" before approaching a zero-energy state. Nevertheless, larger k_i implies less decentralization and larger optimization problems.

As expected, dynamically allocating resources between groups has a major impact on performance: the complete anarchy case with an equal partitioning of resources is significantly worse than all of the cases with $k_i \geq 0$. This motivates the need for further analysis into dynamically allocating budgets heterogeneously without sacrificing decentralization. Doing so will help push the performance of decentralized models closer to the centralized limit.

VI. Conclusions and Future Work

We have developed a framework to protect an arbitrary network from viral propagations, subject to limited resources, varying degrees of decentralization, and uncertainties in network topology. Specifically, our approach combines techniques of receding-horizon control with those of model reduction and robust optimization. Our theoretical and empirical analyses reveal tradeoffs between efficiency and robustness as well as decentralization and optimality. Mitigating or assuaging these tradeoffs is the subject of further investigation. In particular, future work aims to provide upper bounds on errors of group dynamics with multiple reduced-order models, consider more scalable types of uncertainty sets, analyze dynamic schemes for optimally partitioning budgets between network groups, and incorporate dynamic network topologies. Doing so will extend the applicability of our framework to realistic systems.

VII. Acknowledgements

A. Sinha is supported by a Fannie & John Hertz Foundation Fellowship and a Stanford Graduate Fellowship.

References

[1] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 115, no. 772. The Royal Society, 1927, pp. 700–721.
[2] D. Easley and J. Kleinberg, *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, 2010.
[3] S. Banerjee, A. Chatterjee, and S. Shakkottai, "Epidemic thresholds with external agents," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 2202–2210.
[4] A. Goyal, F. Bonchi, and L. V. Lakshmanan, "Learning influence probabilities in social networks," in *Proceedings of the third ACM international conference on Web search and data mining*. ACM, 2010, pp. 241–250.
[5] R. Pastor-Satorras and A. Vespignani, "Epidemic spreading in scale-free networks," *Physical review letters*, vol. 86, no. 14, p. 3200, 2001.
[6] A. Ganesh, L. Massoulié, and D. Towsley, "The effect of network topology on the spread of epidemics," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 2. IEEE, 2005, pp. 1455–1466.
[7] J. O. Kephart and S. R. White, "Directed-graph epidemiological models of computer viruses," in *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*. IEEE, 1991, pp. 343–359.
[8] D. J. Watts, "A simple model of global cascades on random networks," *Proceedings of the National Academy of Sciences*, vol. 99, no. 9, pp. 5766–5771, 2002.
[9] R. Cohen, S. Havlin, and D. Ben-Avraham, "Efficient immunization strategies for computer networks and populations," *Physical review letters*, vol. 91, no. 24, p. 247901, 2003.
[10] F. Chung, P. Horn, and A. Tsiatas, "Distributing antidote using pagerank vectors," *Internet Mathematics*, vol. 6, no. 2, pp. 237–254, 2009.
[11] Y. Hayel, S. Trajanovski, E. Altman, H. Wang, and P. Van Mieghem, "Complete game-theoretic characterization of sis epidemics protection strategies," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 1179–1184.
[12] 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, no. 1, pp. 99–108, 2014.
[13] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. Pappas, "Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013, pp. 7486–7491.
[14] E. Gourdin, J. Omic, and P. Van Mieghem, "Optimization of network protection against virus spread," in *Design of Reliable Communication Networks (DRCN), 2011 8th International Workshop on the*. IEEE, 2011, pp. 86–93.
[15] P. Bremaud, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*. Springer Science & Business Media, 1999, vol. 31.
[16] C. E. Garcia, D. M. Prett, and M. Morari, "Model predictive control: theory and practicea survey," *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
[17] M. Safonov and R. Chiang, "A schur method for balanced model reduction," in *American Control Conference, 1988*. IEEE, 1988, pp. 1036–1040.
[18] K. Glover, "All optimal hankel-norm approximations of linear multivariable systems and their l, -error bounds," *International journal of control*, vol. 39, no. 6, pp. 1115–1193, 1984.
[19] M. Safonov, R. Chiang, and D. Limebeer, "Optimal hankel model reduction for nonminimal systems," *Automatic Control, IEEE Transactions on*, vol. 35, no. 4, pp. 496–502, 1990.
[20] G. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach*. Springer-Verlag New York, 2000.
[21] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, 2009.