

# Control of Generalized Discrete-time SIS Epidemics via Submodular Function Minimization

Nicholas J. Watkins and George J. Pappas

**Abstract**—In this paper, we study a novel control method for a generalized *SIS* epidemic process. In particular, we use predictive control to design optimal protective resource distribution strategies which balance the need to eliminate the epidemic quickly against the need to limit the rate at which protective resources are used. We expect that such a controller may be useful in mitigating the spread of biological diseases which do not confer immunity to those who have been infected previously, with sexually transmitted infections being a prominent example of such. Technically, this paper provides a novel contribution in demonstrating that the particular combinatorial optimal control problem used to design resource allocations has an objective function which is submodular, and so can be solved in polynomial time despite its combinatorial nature. We test the performance of the proposed controller with numerical simulations, and provide some comments on directions for future work.

**Index Terms**—Biological systems; Stochastic systems; Network analysis and control

## I. INTRODUCTION

**D**UE to increased population mobility and a rapidly changing climate, the likelihood of a major biological epidemic occurring in modern times is significant [1]. It is therefore prudent to invest effort into understanding the control of epidemics. Epidemic control is not a new field of research. Indeed, the modern study of epidemic processes dates back to at least the 1920s [2], with work focusing specifically on the optimal control of spreading processes dating back to at least the 1970s [3]–[6]. Notably, nearly all such work focuses on the analysis and control of mean-field approximations to epidemic processes (see, e.g., [7] for a recent review). While known to be accurate for homogeneous populations asymptotically (i.e. in the limit of large populations) [8], it is unclear how well mean-field approximations work in practical settings, as they ignore the underlying stochastic nature of the contacts which drive disease dissemination. This has given rise to recent work studying the control of epidemics without using mean-field approximation.

It was shown in [9] that varying the healing rate of an *SIS* process sufficiently aggressively enables a controller to drive the contagion out of the network quickly. It was shown in [10] that for any priority order strategy, there exists a sufficiently large budget so as to guarantee that the epidemic is driven out of the network quickly. It was shown in [11] that it is possible to construct a system of ordinary differential equations which gives provable upper- and lower- bounds on

infection statistics, and to use these approximations to enforce a stability constraint that guarantees fast extinction of the epidemic. However, in all of these works, the optimization problems which fundamentally characterize the controllers are  $\mathcal{NP}$ -hard. As such, only suboptimal solutions may be obtained, where the performance attained may vary.

It is interesting to consider if there is any way to design predictive controllers for stochastic spreading process models in such a way that the underlying optimization is tractable. Indeed, it has been known for some time that optimizing seed selection under general threshold spreading models is a submodular maximization problem [12], [13], and so can be well approximated by a greedy algorithm. Similar sorts of structural properties have been identified in many variations of problems on threshold-type spreading models, (see, e.g., [14]–[16]), but such models are only appropriate for modeling non-recurrent epidemics. That is, they study processes in which each agent in the network can only actively spread the phenomenon for one interval of time, and afterwards remains inactive. While such a feature seems appropriate in many contexts, e.g. when a disease confers immunity to survivors, this is not always the case. Indeed, many sexually transmitted infections (e.g. chlamydia, gonorrhea) do not confer such immunity (see, e.g., [17]).

The study of discrete-time models with recurrent compartmental memberships seems to have originated in [18], where the stability of a deterministic approximation to a discrete-time Susceptible-Infected-Susceptible (*SIS*) model was studied. Since then, deterministic approximations to discrete-time models have been studied in a few contexts [19]–[22]. While it is known that in some cases the open-loop stability of the deterministic approximation implies a similar notion of open-loop stability of the stochastic model [19]–[21], it appears that [23] was the first to consider closed-loop control of such processes. Whereas [23] studied the control of the discrete-time *SIS* process by direct control of the process parameters, here we study the control of a generalized *SIS* model by way of direct allocation of protective resources.

**Statement of Contributions:** *To the best of the authors' knowledge, this paper is the first to develop a feedback controller for discrete-time, recurrent epidemic processes using node removal as a means of actuation.* With respect to prior work, the present article builds on a preliminary paper [23], which studied the control of a discrete-time *SIS* process by direct control of the process parameters. Principally, our work here differs in that we focus on control via node removal. This change in focus allows us to model control actions more realistically: protective devices (e.g. latex gloves, barriers,

N.J. Watkins and G.J. Pappas are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA, {nwatck, pappasg}@seas.upenn.edu.

condoms) are discrete objects by nature, and so should be treated as such in the model.

The primary technical contribution of this work is in demonstrating that a particular combinatorial optimal control problem can be posed as a submodular minimization problem with respect to a ground set that is a subset of the set of nodes, and so can be solved efficiently (i.e., in polynomial time with respect to the size of the graph). Additionally, we investigate the performance of the controller by way of numerical simulations, and provide some comments on directions for future work which we believe to be promising.

## II. PROBLEM STATEMENT

In this section, we formally develop the problem studied in the paper. Section II-A details the generalized *SIS* model studied. Section II-B details the actuation model studied. Section II-C details the optimal control problem studied.

### A. Epidemic Model

We study an epidemic process on an  $n$  node graph  $\mathcal{G}$ , with node set  $\mathcal{V}$ , and edge set  $\mathcal{E}$ . Epidemic processes are mathematical objects which evolve on graphs, in which each node is considered to be an agent, and each edge a relationship between agents. The current health status of an agent is modeled by a collection of compartments, where at each time, every agent belongs to exactly one compartment. The process we study has one susceptible (i.e. healthy) compartment (denoted by  $S$ ), and  $p$  infected spreading compartments (denoted by  $\{I_k\}_{k \in [p]}$ ). A transition from susceptibility to infection comes from contact between a susceptible node and an infected node. Other transitions are due to random events internal to the agent (e.g. progressing through the stages of the disease). Note that the multiplicity of infected compartments here is used to shape the amount of time an agent spends in infection.

We denote by  $X_i^\ell(t)$  an indicator random variable which takes the value 1 if node  $i$  is in compartment  $\ell$  at time  $t$ , and takes the value 0 otherwise. We take compartmental memberships to evolve in a way that generalizes the standard discrete-time *SIS* process [18]. A node  $i$  transitions from susceptible to the first infected compartment (i.e.  $I_1$ ) on the increment from time  $t$  to time  $t + 1$  through contact with nodes which are in an infected compartment at time  $t$ . Given that a node  $i$  is infected at time  $t$ , it transitions to other model compartments independently of all external phenomena. We then have that the indicator  $X_i^S$  evolves as

$$X_i^{S+} = X_i^S(1 - Z_{i(X)}^{S \rightarrow I_1}) + \sum_{k=1}^p X_i^{I_k} Z_{i(X)}^{I_k \rightarrow S}, \quad (1)$$

where the random variable  $Z_{i(X)}^{S \rightarrow I_1}$  is an indicator that at least one infection event (indicated by random variables  $Y_{ij}$ ) has occurred on an edge between node  $i$  and an infected node  $j$  in the set of neighbors of  $i$  (denoted  $\mathcal{N}_i$ ), and is defined as

$$Z_{i(X)}^{S \rightarrow I_1} \triangleq \min\{1, \sum_{j \in \mathcal{N}_i} Y_{ij} \sum_{k=1}^p X_j^{I_k}\}, \quad (2)$$

and the random variables  $Z_{i(X)}^{I_k \rightarrow S}$  are indicators denoting a transition from the  $k$ 'th infected compartment to susceptibility. Likewise, we have that the indicator  $X_i^{I_1}$  evolves as

$$X_i^{I_1+} = X_i^S Z_{i(X)}^{S \rightarrow I_1} + \sum_{k=1}^p X_i^{I_k} Z_{i(X)}^{I_k \rightarrow I_1} - X_i^{I_1} \sum_{\ell \in \mathcal{L}} Z_{i(X)}^{I_1 \rightarrow \ell}, \quad (3)$$

and the indicators  $X_i^{I_k}$  for  $k \neq 1$  evolve as

$$X_i^{I_k+} = \sum_{\ell \in \mathcal{L} \setminus \{S\}} X_i^\ell Z_{i(X)}^{\ell \rightarrow I_k} - X_i^{I_k} \sum_{\ell \in \mathcal{L}} Z_{i(X)}^{I_k \rightarrow \ell}, \quad (4)$$

where  $\mathcal{L}$  is the set of all compartmental labels. Note that for all nodes to belong to a unique compartment at all times, we must have that  $Z_{i(X)}^{\ell \rightarrow \ell'}$  each take values on  $\{0, 1\}$  and satisfy  $\sum_{\ell'} Z_{i(X)}^{\ell \rightarrow \ell'} = 1$  with probability one.

To make matters concrete, we fix one particular way of enforcing this. We assume all  $Y_{ij}$  in (2) are independently distributed Bernoulli random variables with known success probabilities, and that when a node is infected, the compartmental membership random variables evolve independently as a discrete-time Markov chain with  $p + 1$  states, structured so that  $S$  is its unique absorbing state, and that  $S$  is reached in finite time with probability one. This level of generality allows us to model the amount of time taken to recover from an infection with more precision than a standard *SIS* epidemic, in which there is only one infected compartment, and  $Z_{i(X)}^{I_1 \rightarrow S}$  is an independent Bernoulli random variable at all times. Such an assumption forces the time taken to recover to be distributed as a geometric random variable. Under the model specified here, the time taken to recover follows a discrete phase-type distribution, and so is quite general [24].

### B. Actuation Model

We consider allocating protective barriers in order to preventing the spread of an infection. Formally, our controller actuates the process detailed in Section II-A by selecting a subset of nodes  $\mathcal{P}$  to protect against infection. Because protective devices (e.g. latex gowns, gloves, condoms) are often single-use, it is appropriate to model the economic cost of protecting the set of nodes  $\mathcal{P}$  as being the sum over all edges in which one adjacent node is a susceptible node that is protected, and the other adjacent node is infected.

That is, we only pay to protect a particular node if it interacts with an infected person, and we pay in proportion to the extent of interaction between the protected node and its infected neighbors over the discretized time period. For example, if we are allocating condoms to mitigate the spread of chlamydia and we want to protect a particular person  $i$ , we will provide  $i$  with one condom for each contact between it and all infected partners. Mathematically, we have

$$\mathcal{C}(\mathcal{P}|X) \triangleq \sum_{i \in \mathcal{S}(X)} \mathbb{1}_{\{i \in \mathcal{P}\}} \sum_{j \in \mathcal{N}_i \cap \mathcal{I}(X)} c_{ij}, \quad (5)$$

where  $\mathbb{1}_{\{ \cdot \}}$  is a  $\{0, 1\}$  indicator function, the non-negative constants  $c_{ij}$  model the cost of providing a protective barrier for all interactions between  $i$  and  $j$  over one time period,  $\mathcal{S}(X)$  denotes the set of susceptible nodes at state  $X$ , and

$\mathcal{I}(X)$  denotes the set of infected nodes at state  $X$ . We model protecting a node by removing it from the graph. That is, a node  $i$  which is in the set of protected nodes  $\mathcal{P}$  and is susceptible at time  $t$  is susceptible at time  $t + 1$  with probability one. Mathematically, we then have that the controlled dynamics for  $X_i^S$  follow

$$X_i^{S+} = X_i^S(1 - \mathbb{1}_{\{i \notin \mathcal{P}\}} Z_{i(X)}^{S \rightarrow I_1}) + \sum_{k=1}^p X_i^{I_k} Z_i^{I_k \rightarrow S}, \quad (6)$$

and the controlled dynamics for  $X_i^{I_1}$  are changed similarly.

### C. Optimal Control Problem

If we were strictly concerned with minimizing the total accumulated cost for our controller, it is feasible to take the action  $\mathcal{P} = \emptyset$  at all times. Doing so incurs zero cost. However, this may result in the infection persisting in the population for a long time. This is a socially undesirable outcome. In general, we may wish to consider the presence of infection as a sort of soft cost imposed on the controller.

Previous works studying optimal control of epidemics (e.g., [3]–[6]) often do so by posing a problem of the form

$$\min_{\pi \in \Pi} \mathbb{E}_{\pi} \left[ \sum_{\tau=0}^{\infty} J(X(\tau), \theta(\tau)) | X \right], \quad (7a)$$

$$J(X, \theta) \triangleq \mu c(\theta) + (1 - \mu) q(\mathbb{I}(X)) \quad (7b)$$

where  $\Pi$  is the set of all non-anticipating control policies which map observations of the process as it evolves to changes in the spreading model's parameters  $\theta$ ,  $c(\theta)$  is the instantaneous cost of setting the processes parameters to  $\theta$ ,  $\mathbb{I}(X)$  is a function which maps the state  $X$  to the number of infected nodes in state  $X$ , and the function  $q$  is some non-negative valued function which determines the extent to which we should care about the existence of infection. In our case, if it were so that the process  $\{X(t)\}$  was a Markov chain on a small state space and the set of possible control actions were small, (7a) could be solved by treating it as a Markov decision process and applying a standard solution technique, e.g. value iteration. However, even for standard *SIS* processes,  $\{X(t)\}$  evolves as a Markov chain with  $2^n$  states, and there are  $2^n$  possible choices of the set  $\mathcal{P}$ . That is, the complexity of using a standard Markov decision process algorithm here is exponential, due to the curse of dimensionality. As such, it is an interesting task to construct a controller which allows for the same qualitative tradeoff as (7a), but is computationally tractable to implement.

To accomplish this, we consider applying controls which solve the infinite-horizon optimal control problem

$$\min_{\mathcal{P} \subseteq \mathcal{V}} h_X(\mathcal{P}) \quad (8a)$$

$$h_X(\mathcal{P}) \triangleq \mu \mathcal{C}(\mathcal{P}|X) + (1 - \mu) Q(\mathcal{P}|X) \quad (8b)$$

where  $\mathcal{C}(\mathcal{P}|X)$  is the cost function (5), and  $Q(\mathcal{P}|X)$  is the cost of applying the rollout policy which protects all susceptible nodes for all future times, given that the set of nodes  $\mathcal{P}$  is protected at the current time (see, e.g., [25, Section 6.4] for

background on the use of rollout policies in infinite horizon optimal control). Intuitively, our controller designs the protection set  $\mathcal{P}$  at the current time while *anticipating* that at all future times, every reasonable action will be implemented to eradicate infection as quickly as possible. Thus, our controller anticipates that a cost  $c_{ij}$  occurs at every time where exactly one of the node  $i$  or  $j$  is susceptible. That is,  $Q$  is defined mathematically as

$$Q(\mathcal{P}|X) \triangleq \mathbb{E}_{\Theta(\mathcal{P})} \left[ \sum_{\tau=1}^{\infty} \sum_{\{i,j\} \in \mathcal{E}} (X_i^S(\tau) \oplus X_j^S(\tau)) c_{ij} | X \right], \quad (9)$$

where  $\Theta(\mathcal{P})$  is the measure induced by protecting the set of nodes  $\mathcal{P}$  at the current time, and by  $\oplus$  we denote the exclusive or operator, i.e. for two variables  $Y, Z \in \{0, 1\}$ ,

$$Y \oplus Z \triangleq \begin{cases} 0, & Y = 0, Z = 0; \\ 1, & Y = 1, Z = 0; \\ 1, & Y = 0, Z = 1; \\ 0, & Y = 1, Z = 1. \end{cases} \quad (10)$$

Note that because we update our decision  $\mathcal{P}$  at every time, this total protection strategy is not actually implemented. Rather,  $Q$  plays the same role in (8a) as  $q(\mathbb{I})$  plays in (7a). It penalizes the existence of infection in the network, and so allows a control designer a means for trading off between an immediate resource expenditure and the rate of decay in the population of infected individuals by appropriately selecting  $\mu$ . However, it is not immediately clear that actions can be efficiently computed as in (8a), as the set over which the optimum must be computed has  $2^n$  elements, and it is not clear that the objective  $h_X$  is sufficiently well structured so as to enable efficient optimization.

*In the body of this paper, we are concerned with determining whether solutions of (8a) can be computed efficiently (i.e. in polynomial time with respect to the size of the graph), and whether the engendered controller provides reasonable behavior.* We address computational efficiency in Section III-B, in which we show that the objective of the minimization in (8a) is submodular, which allows us to minimize  $h_X$  using polynomial time algorithms. We assess the controller's behavior in Section IV with a numerical example.

## III. PREDICTIVE CONTROL OF GENERALIZED *SIS*

In this section, we develop the mathematical foundations required to efficiently implement a controller as described in Section II. Section III-A provides some technical preliminaries which are needed in our analysis. Section III-B demonstrates that the optimization problem (8a) has sufficient structure so as to allow for the use of efficient optimization algorithms, and provides some references to software packages that can be used to solve the problem.

### A. Mathematical Preliminaries

The key mathematical concept which allows for the efficient solution of (8a) is submodularity. Submodularity is

a mathematical formalization of the concept of diminishing returns: adding an object to a larger set has less of an impact than adding the same object to a smaller set. Formally, a submodular function satisfies the following definition.

**Definition 1** (Submodular Functions). *Let  $\Omega$  be a finite ground set of objects, and suppose  $f : 2^\Omega \rightarrow \mathbb{R}$ , where  $2^\Omega$  denotes the power set of  $\Omega$ , i.e. the set of all subsets of  $\Omega$ . The function  $f$  is said to be submodular if and only if*

$$f(Z \cup \{z\}) - f(Z) \leq f(Y \cup \{z\}) - f(Y) \quad (11)$$

holds for all  $Y \subset Z \subseteq \Omega$ , and  $z \in \Omega \setminus Z$ .

It is frequently the case that the submodularity of a complicated function is verified by reducing the proof to checking the submodularity of a simpler function. We use such an argument later (Section III-B), using the exclusive-or function as the simple function, which is submodular:

**Lemma 1** (Submodularity of Restricted Exclusive Or). *Let  $\Omega$  be a finite ground set, take  $a, b \in \Omega$ ,  $A, B \in \{0, 1\}$ , and let  $\oplus$  denote the operator defined by (10). The function*

$$f(W) = A\mathbb{1}_{\{a \notin W\}} \oplus B\mathbb{1}_{\{b \notin W\}},$$

is submodular.

*Proof.* Note that if  $A = B = 0$ , then  $f(W) = 0$ , and is trivially submodular. Note also that if  $A \neq B$ , then  $f(W) = A\mathbb{1}_{\{a \notin W\}} + B\mathbb{1}_{\{b \notin W\}}$ , which is non-negative sum of submodular functions, so is submodular [26]. It remains to consider  $A = B = 1$ . In this case, if  $a = b$ , we have  $f(W) = 0$ , which is trivially submodular. It remains to consider the case in which  $a \neq b$ .

Take  $Z \supseteq Y$ ,  $z \in \Omega \setminus Z$ , and suppose neither  $a$  nor  $b$  are in  $Z$ . By  $Z \supseteq Y$ , we have that neither  $a$  nor  $b$  are in  $Y$ . If  $z \in \{a, b\}$ , then  $f(Z \cup \{z\}) - f(Z) = 1 = f(Y \cup \{z\}) - f(Y)$ . If  $z \notin \{a, b\}$ , then  $f(Z \cup \{z\}) - f(Z) = 0 = f(Y \cup \{z\}) - f(Y)$ .

Now, take  $Z \supseteq Y$ ,  $z \in \Omega \setminus Z$ , and suppose both  $a$  and  $b$  are in  $Z$ . The condition  $z \in \Omega \setminus Z$  implies that  $z \notin \{a, b\}$ , and so  $f(Z \cup \{z\}) - f(Z) = 0 = f(Y \cup \{z\}) - f(Y)$ .

Finally, take  $Z \supseteq Y$ ,  $z \in \Omega \setminus Z$ , and suppose exactly one of  $a$  or  $b$  is in  $Z$ . For concreteness, suppose  $a \in Z$ . There are two cases to consider. First, suppose  $a \notin Y$ . Then, if  $z \notin \{b\}$ , we have  $f(Z \cup \{z\}) - f(Z) = 0 = f(Y \cup \{z\}) - f(Y)$ . If  $z \in \{b\}$ , then  $f(Z \cup \{z\}) - f(Z) = -1 \leq f(Y \cup \{z\}) - f(Y) = 1$ . Now, suppose  $a \in Y$ . If  $z \notin \{b\}$ , then  $f(Z \cup \{z\}) - f(Z) = 0 = f(Y \cup \{z\}) - f(Y)$ . If  $z \in \{b\}$ , then  $f(Z \cup \{z\}) - f(Z) = -1 = f(Y \cup \{z\}) - f(Y)$ .  $\square$

### B. Efficiently Computing Optimal Sets of Protected Nodes

A principal reason that submodularity is an important concept is that it allows for a variety of combinatorial optimization problems to be solved (or approximately solved) in polynomial time. In this subsection, we demonstrate that (8a) is one such problem. As it is well-known that the minimum of a submodular function over a finite ground set can be computed in time which grows polynomially with respect to the size of the ground set (see, e.g., [27], [28]), it suffices for us to

demonstrate that the objective of (8a) is a submodular function. We accomplish this in the following result:

**Theorem 1** (Submodularity of Objective Function). *Fix a particular generalized SIS process (as detailed in Section II-A), a number  $\mu \in [0, 1]$ , and a state  $X$ . Let  $\mathcal{C}$  be defined as in (5), and  $Q$  as in (9). Then, the function  $h_X(\mathcal{P})$  as defined by (8b) is a submodular set function, where the ground set is taken to be the set of susceptible nodes at state  $X$ ,  $\mathcal{S}(X)$ .*

*Proof.* Since a function defined by a non-negative weighted sum of submodular functions is itself a submodular function [26], it suffices to show that  $\mathcal{C}$  and  $Q$  are both submodular functions. We handle each separately.

*Submodularity of  $\mathcal{C}$ :* Because the constants  $c_{ij}$  are non-negative, it suffices to note that indicators of the form  $f(W) = \mathbb{1}_{\{r \in W\}}$  are submodular. This is straightforward: if  $x \in \Omega \setminus W$ , then  $f(W \cup \{x\}) - f(W) = 1$  when  $x = r$ , and 0 otherwise, regardless of the choice of  $W$ . This implies that (11) is met with equality, satisfying Definition 1.

*Submodularity of  $Q$ :* Since the generalized SIS process defined in Section II-A is defined on a finite state space in finite time, we may construct a finite sample space for it. Call this sample space  $\Xi$ . Because having exactly one node in a pair of nodes  $i$  and  $j$  be susceptible is equivalent to having exactly one node in a pair of nodes belong to some infected compartment, we have

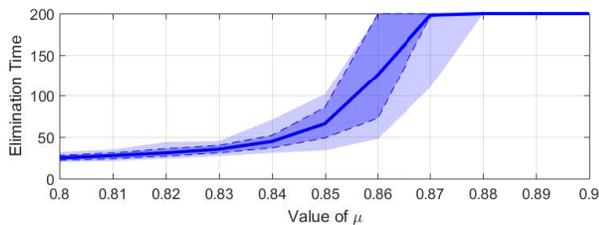
$$\begin{aligned} Q(\mathcal{P}|X) &= \mathbb{E}_{\Theta(\mathcal{P})} \left[ \sum_{\tau=1}^{\infty} \sum_{\{i,j\} \in \mathcal{E}} (X_i^I(\tau) \oplus X_j^I(\tau)) c_{ij} | X \right], \\ &= \sum_{\xi \in \Xi} \sum_{\tau=1}^{\infty} \sum_{\{i,j\} \in \mathcal{E}} (X_i^I(\tau; \mathcal{P}, \xi) \oplus X_j^I(\tau; \mathcal{P}, \xi)) c_{ij} \Pr(\xi|X), \end{aligned}$$

where we denote the dependence on the choice of protected nodes and the underlying sample  $\xi$  explicitly, use the notation  $\Pr(\xi|X)$  for the probability that the sample  $\xi$  is drawn from  $\Xi$ , given that the process is currently in state  $X$ , and define the shorthand notation  $X_j^I = \sum_{k=1}^p X_j^{I_k}$ . Since the terms  $c_{ij}$  and  $\Pr(\xi|X)$  are nonnegative, it suffices to show that the terms  $(X_i^I(\tau; \mathcal{P}, \xi) \oplus X_j^I(\tau; \mathcal{P}, \xi))$  are submodular functions.

Because the rollout policy removes all susceptible nodes which are adjacent to an infected node, we have that all nodes which are infected at some time  $\tau \geq t+1$  were either infected at time  $t$ , or became infected on the transition from time  $t$  to time  $t+1$ . In the case where the node  $i$  was infected at time  $t$ , the definition of the process has that all indicators  $X_i^\ell(\tau; \mathcal{P}, \xi)$  are independent of the choice  $\mathcal{P}$ , and are determined only by the particular choice of  $\xi$ . In particular, we have that

$$\begin{aligned} X_i^I(\tau; \mathcal{P}, \xi) &= \\ &\begin{cases} \mathbb{1}_{\{\tau \leq \mathcal{H}_i(\xi, X_i(t))\}}, & X_i^I = 1; \\ \mathbb{1}_{\{i \notin \mathcal{P}\}} \mathbb{1}_{\{\forall j \in \mathcal{N}_i Y_{ij}(\xi) X_j^I\}} \mathbb{1}_{\{\tau \leq \mathcal{H}_i(\xi, X_i(t+1))\}}, & X_i^S = 1, \end{cases} \end{aligned}$$

where  $\mathcal{H}_i(\xi, X_i(t))$  is a random variable denoting the next time at which an event which transitions the node  $i$  from infected to susceptible happens, given that the node is in its current compartment, and  $\vee$  denotes the logical or operation, used here because all infection events have the same effect on an unprotected susceptible node.



**Fig. 1:** The effect of varying  $\mu$  to tune controller. The region shaded in light blue contains the middle 98% of the simulated trajectories; the dark blue contains the middle 80% of the simulated trajectories. The solid blue line is the sample expectation. As expected, the controller's behavior experiences a phase transition from a domain in which future cost is prioritized (here, below  $\mu = 0.8$ ), and a domain in which instantaneous spending is prioritized (here, above  $\mu = 0.88$ ). Note: the simulation parameters reported in the example of Section IV were used to generate this figure.

Note that in the case that  $X_i^I = 1$ , submodularity follows immediately from noting that  $X_i^I(\tau; \mathcal{P}, \xi)$  is unaffected by the choice  $\mathcal{P}$ . In the case that  $X_i^S = 1$ , we have  $X_i^I(\tau; \mathcal{P}, \xi)$  is the indicator  $\mathbb{1}_{\{i \notin \mathcal{P}\}}$  multiplied by a constant in  $\{0, 1\}$ , which is determined by  $\xi$ . It follows that  $Q(\mathcal{P}|X)$  is a non-negative weighted sum of functions of the form

$$f(\mathcal{P}) = A\mathbb{1}_{\{a \notin \mathcal{P}\}} \oplus B\mathbb{1}_{\{b \notin \mathcal{P}\}},$$

with  $A, B \in \{0, 1\}$ , which are submodular by Lemma 1.  $\square$

The submodularity proven in Theorem 1 allows us to use a variety of polynomial-time algorithms [27], [29] to solve (8a) to optimality. In our simulations (Section IV), we have used an implementation of Fujishige's minimum-norm-point algorithm [30], [31] from the submodular function optimization package [29]. Alternatively, one may consider using faster *approximate* submodular minimization algorithms [32], [33] if the problem instance considered is too large to be handled efficiently in practice by other methods. It is also worth noting that (8a) may possess sufficient structure to enable the development of faster algorithms for computing the exact optimum, as it is the sum of a monotone increasing modular function  $\mathcal{C}(\mathcal{P}|X)$ , and a submodular function  $Q(\mathcal{P}|X)$  which is a weighted sum of indicators and exclusive-or operations.

**Remark 1** (On State and Parameter Uncertainty). It may be the case in practice that neither the state  $X$  nor the parameters of the spreading process are deterministically known. However, this situation can be readily accommodated into our controller's computations. Supposing that one has knowledge of the distribution of the state and spreading parameters, one can first sample from this, and then forward-simulate the evolution of the process with the sampled state and parameter set fixed. Verifying the submodularity of the objective function with respect to each sample path used can be done with the same argument as in the proof of Theorem 1, and so merits no further discussion here. Simulations regarding this are given in the Supplementary Material.  $\bullet$

**Remark 2** (Tuning Controller Behavior). Selecting  $\mu$  appropriately is important. For small values of  $\mu$ , the objective of (8a) is dominated by the effect of  $Q(\mathcal{A}|X)$ . In the extreme case

where  $\mu = 0$ , the controller takes an action which minimizes the future cost associated to guaranteeing the infection spreads no further. For large values of  $\mu$ , the objective of (8a) is dominated by the effect of  $\mathcal{C}(\mathcal{A}|X)$ . In the extreme case where  $\mu = 1$ , it is simple to show that the optimal action will be to protect no nodes, under any circumstance. Intuitively, intermediate values of  $\mu$  balance the immediate cost of resource expenditure against its long-term consequences. This is seen in Figure 1.  $\bullet$

#### IV. A NUMERICAL EXAMPLE

We consider a graph with 200 nodes. We assume that for each  $i, j \in \mathcal{V}$  such that  $i \neq j$ , the edge  $\{i, j\}$  is in the set of edges  $\mathcal{E}$  independently with probability 0.01. We assume that an unprotected contact between two nodes results in an infection spreading from the infected node to the susceptible node with a probability randomly generated from the unit interval. We assume that nodes come into contact with each other a maximum of three times per day, and so edge costs take values in the set  $\{0, 1, 2, 3\}$ , where the value 0 is taken if the edge  $\{i, j\} \notin \mathcal{E}$ , and the other values are chosen uniformly at random if  $\{i, j\} \in \mathcal{E}$ . Because treatments for sexually transmitted infections (e.g. chlamydia) often take in excess of a week to be effective [34], the Markov chain used to model the infectious compartments was chosen so as to make the distribution of times from infection to susceptibility uniformly distributed on the set  $\{7, 8, 9, 10\}$ .

A plot demonstrating the effect of varying  $\mu$  is given in Figure 1, where the lightest shaded region contains the middle 98% of the 100 sample paths generated, the darker shade of blue contains the middle 80% of the sample paths generated, and the central blue line gives the sample average. As anticipated in Remark 2, there is a continuous range of values  $\mu$  over which the controller's behavior markedly changes. For this particular example, this region was found to be the interval  $[\cdot 8, \cdot 9]$ . Note that in general, appropriate values of  $\mu$  must be determined for each fixed model. It may well be possible to automate a procedure for finding appropriate values of  $\mu$  online. However, developing and applying any such method would likely be quite technically involved, and is best saved as a task for future work.

The behavior of the controller with  $\mu = 0.85$  is given in Figure 2, where the lightest shaded region contains the middle 98% of the 100 generated samples, the darker shaded region contains the middle 80% of the generated samples, and the central blue line gives the sample average of the process. We see that the controller drives the epidemic out of the network quickly, while using resources at a lesser rate than one which protects all susceptible nodes at all times.

#### V. CONCLUSIONS AND FUTURE WORK

In this paper, we have developed a controller which allocates discrete protective devices in order to mitigate the spread of an epidemic process. Such a controller could be of use in fighting many forms of biological disease, which often do not confer immunity to people that have survived infections, and so are well-modeled by *SIS*-type processes.

