

Optimal policy design to mitigate epidemics on networks using an SIS model

Carlo Cenedese, Lorenzo Zino, Michele Cucuzzella, Ming Cao

Abstract—Understanding how to effectively control an epidemic spreading on a network is a problem of paramount importance for the scientific community. The ongoing COVID-19 pandemic has highlighted the need for policies that mitigate the spread, without relying on pharmaceutical interventions, that is, without the medical assurance of the recovery process. These policies typically entail lockdowns and mobility restrictions, having thus nonnegligible socio-economic consequences for the population. In this paper, we focus on the problem of finding the optimum policies that “flatten the epidemic curve” while limiting the negative consequences for the society, and formulate it as a nonlinear control problem over a finite prediction horizon. We utilize the model predictive control theory to design a strategy to effectively control the disease, balancing safety and normalcy. An explicit formalization of the control scheme is provided for the susceptible–infected–susceptible epidemic model over a network. Its performance and flexibility are demonstrated by means of numerical simulations.

I. INTRODUCTION

The ongoing COVID-19 pandemic has highlighted the key role played by public health authorities in enacting nonpharmaceutical interventions (NPIs) to “flatten the epidemic curve” when no effective pharmaceutical treatments such as vaccines are available [1], [2]. However, NPIs typically entail the implementation of harsh measures, including lockdowns and restrictions of personal freedom of movement, which may yield severe socio-psychological and economic consequences [3]. Thus, they should be implemented keeping a reasonable balance between safety and normalcy. To this aim, the development of accurate tools to predict the course of an epidemic and evaluate the impact of different interventions has become a task of paramount importance for the scientific community, aiming at assisting public health authorities in their decisions.

The mathematical modeling of epidemics has emerged as a valuable framework to perform such a task [4], [5], [6], [7], [8]. Relevant examples can be found in the useful insights provided into the ongoing COVID-19 pandemic [9], [10], [11], [12], [13], [14], [15], [16], [17], [18]. Within this framework, network models have become popular as they

allow us to capture the relation between human mobility and the spatial spread of a disease. Of particular interest is the problem of understanding how to effectively control the spread of an epidemic disease on a network by acting on the nodal dynamics and on the network structure.

In the literature, such problems have often been addressed assuming limited changes in the network structures, that is, by studying how to re-arrange the network structure or distribute antidote in order to increase the population’s resistance against a possible epidemic. Important results have been found by using geometric programming [19]; distributed algorithms have been recently proposed to address this problem [20], [21]. However, the ongoing health crisis has highlighted the importance of having control schemes, which take into account the dynamic evolution of the epidemic process, and can thus be updated online, as the outbreak evolves. While a considerable body of literature has been proposed to study dynamical control strategies for vaccination campaigns and antidote distribution using optimal control theory [22], [23] and Model Predictive Control (MPC) [24], [25], [26], limited results are available in the absence of effective treatments, that is, when the control action has to focus on the contagion mechanisms rather than on the recovery mechanism. Recently, motivated by the ongoing COVID-19 pandemic, some efforts have been devoted to bridging this gap by proposing feedback control interventions and leveraging MPC [13], [14], [16], [18].

Here, inspired by these works, we propose an optimal control approach in order to address the problem of mitigating an epidemic process spreading by means of regional policies that entail activity reductions and targeted mobility restrictions. This holistic approach to NPIs constitutes a key contribution of this paper and differs from other works in the literature that typically focus on optimizing a specific intervention policy (e.g., social distancing in [27]) only. Specifically, we consider a discrete-time deterministic Susceptible–Infected–Susceptible (SIS) epidemic model on a network in which two types of control actions are included to mitigate the contagion: actions to reduce the social activity in some regions of the network (modeling, for instance, lockdown measures, closures of activities, and curfews), and policies to limit or ban travels between specific nodes (e.g. cities, provinces, regions, countries). Then, an optimal control problem is formulated to find a strategy that mitigates the spread of the disease in the network, limiting the negative consequences of NPIs. The proposed optimal control problem takes into account important features, such as the

C. Cenedese is with the Department of Information Technology and Electrical Engineering, ETH Zürich, Zurich, Switzerland. L. Zino, M. Cao and M. Cucuzzella are with the Faculty of Science and Engineering, University of Groningen, Groningen, the Netherlands. M. Cucuzzella is also with the Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia, Italy. E-mails: ccenedese@ethz.ch, {lorenzo.zino, m.cao}@rug.nl, michele.cucuzzella@unipv.it.

The work by L. Zino, and M. Cao is partially supported by the European Research Council (ERC-CoG-771687) and the Netherlands Organization for Scientific Research (NWO-vidi-14134). The work by C. Cenedese is supported by SNSF under NCCR Automation.

balance between safety and normalcy, the need of keeping the epidemics under control, and the increase of socio-economic costs associated to the implementation of NPIs, caused by accumulating factors. Different from many control schemes proposed in the literature (see for example [4], [7] for an overview), our formalization does not (necessarily) have the objective to eradicate the disease (which may be extremely costly—and practically unfeasible—using only NPIs). In contrast, it allows the controller to set an acceptable prevalence of the disease (which may depend on the hospital capacity and may vary across the nodes), providing thus a framework that might realistically be adopted to assist public health authorities in their decision toward mitigating epidemic outbreaks. Finally, we propose several simulations in which we highlight the performance of the proposed control strategy, the benefits and disadvantages of farsighted policies in comparison to myopic ones, and perform a sensitivity analysis on how the objectives of the policy maker affect the final solution.

The rest of the paper is organized as follows. In Section II, we present the SIS epidemic model. In Section III, we discuss the control policies. In Section IV, we propose our optimal control policy. In Section V, we illustrate the results of our numerical simulations. Section VI concludes the paper and outlines the avenues of future research.

A. Notation

We gather here the notation used throughout the paper. The set of nonnegative integer numbers is denoted by \mathbb{N} , and the set of nonnegative real numbers is denoted by \mathbb{R}_+ . Given a vector \mathbf{x} , we denote by \mathbf{x}^\top its transpose and, for a positive definite matrix $S \succ 0$, we denote by $\|\mathbf{x}\|_S := \sqrt{\mathbf{x}^\top S \mathbf{x}}$ the norm of \mathbf{x} weighted by S . With $\mathbf{0}$ and $\mathbf{1}$, we denote the all-0 and all-1 vectors, respectively. Given a matrix S , we denote by $\sigma(S)$ its maximum singular value. S_{ij} and S_i denote the element in the i -th row and j -th column and the i -th row of S , respectively. We use the following notation $\text{col}((x_i)_{i \in [1, \dots, N]}) := [x_1, \dots, x_N]^\top$.

II. SIS EPIDEMICS MODEL OVER NETWORKS

In this section, we introduce the SIS epidemic model, illustrate its main properties, and discuss the assumptions needed to formulate the control problem.

We consider a discrete-time deterministic SIS epidemic model on a network [6]. The network structure arises from the local interactions among several nodes (communities), which represent geographical entities, such as countries, regions, or even cities, thanks to the high flexibility of the model. We assume there are N communities interconnected by P links, which can represent for instance roads, air routes or even simple geographic adjacencies. In general, a link between two communities means that there is a flow of people between them. The network is hereafter formalized via an undirected and connected graph \mathcal{G} , where the communities correspond to the set $\mathcal{V} := \{1, \dots, N\}$ of the nodes. A connection between two communities $i, j \in \mathcal{V}$ is denoted via an edge e_ℓ connecting the two corresponding

nodes, defined as the unordered couple $e_\ell := \{i, j\}$. We assume that all the self-loops $\{i, i\}$, $i \in \mathcal{V}$ are present. The edge set is the collection of all the edges of the graph, i.e., $\mathcal{E} := \{e_1, \dots, e_P\}$. Therefore, the graph is defined as $\mathcal{G} := (\mathcal{V}, \mathcal{E})$. The fraction of infected individuals in community $i \in \mathcal{V}$ at time $t \in \mathbb{N}$ is denoted by $x_i(t) \in [0, 1]$. This quantity describes the temporal evolution of the health state of community i . We assume that the health state evolves according to the dynamics of a discrete-time SIS model [6], i.e., for all time-steps $t \in \mathbb{N}$, we have

$$x_i^+ = (1 - \mu)x_i + (1 - x_i)\bar{\beta}_i \sum_{j \in \mathcal{V}} \bar{A}_{ij}x_j, \quad (1)$$

where $x_i^+ := x_i(t+1)$ and $x_i := x_i(t)$.

Next, we discuss the role and physical interpretation of each parameter appearing in (1).

1) *Recovery rate* $\mu \in [0, 1]$: It is the rate at which the individuals manage to recover from the disease. Here, we assume that the recovery rate is constant across the entire population and cannot be increased by the controller, capturing those epidemics for which a cure is not yet developed, e.g., the early spread of COVID-19.

Assumption 1 (Positive recovery rate): For all $i \in \mathcal{V}$, the recovery rate $\mu > 0$ is constant and strictly positive. \square

2) *Infection rate* $\bar{\beta}_i > 0$: This represents the rate at which the individuals in community $i \in \mathcal{V}$ become infected when they get in touch with others. The higher this value is, the easier people become infected. Such a value can differ among the communities, e.g., due to the implementation of different NPIs. Infection rates are gathered in the n -dimensional vector $\bar{\beta}$. The trajectory of $\mathbf{x}(t) := \text{col}((x_i(t))_{i \in \mathcal{V}})$ can naturally have two possible behaviors, depending on the model parameters. Either it converges to the disease-free equilibrium, i.e., $\mathbf{x}(\infty) = \mathbf{0}$, or the disease-free equilibrium becomes unstable, and the trajectory converges to a (unique) endemic equilibrium with $x_i(\infty) > 0$, for all $i \in \mathcal{V}$ [28], [29]. The threshold between these two regimes depends on whether $\sigma(\text{diag}(\bar{\beta})\bar{A})/\mu$ is smaller or greater than 1. A simpler—network-independent—sufficient (but nonnecessary) condition for the trajectory to converge to an endemic equilibrium can be established by requiring that $\bar{\beta}_i/\mu > 1$ for all $i \in \mathcal{V}$. In this work, we focus on the case in which the disease does not die out naturally, and thus NPIs have to be put in place toward mitigating the epidemic outbreak. Hence, we make the following assumption.

Assumption 2 (Disease spreading): For all $i \in \mathcal{V}$, the infection rate satisfies $\bar{\beta}_i > \mu$. \square

3) *Communities interaction* $\bar{A}_{i,j} \in [0, 1]$: People in $j \in \mathcal{V}$ can move to the neighboring communities, denoted by $\mathcal{N}_j := \{\ell \in \mathcal{V} : \{j, \ell\} \in \mathcal{E}\}$, and interact with people there. The population that flows from j to $i \in \mathcal{N}_j$ influences the health state of community i , i.e., x_i . This is modeled in (1) via the weighted adjacency matrix $\bar{A} := [\bar{A}_{ij}]_{i,j \in \mathcal{V}} \in \mathbb{R}_+^{N \times N}$, where \bar{A}_{ij} is the weight associated to the edge (i, j) of \mathcal{G} . The diagonals of \bar{A} represent the part of population that remains in the same community. We now introduce the following assumption on this matrix.




Fig. 1: Schematic of the controlled network SIS model. The time-varying weighted adjacency matrix $A(t)$ determines the interaction that community i has with other communities. The time-varying infection rate $\beta_i(t)$ determines the new infections, while the constant recovery rate μ governs the recovery process.

Assumption 3 (Stochastic and positive diagonals): The weighted adjacency matrix \bar{A} associated with the undirected graph \mathcal{G} is stochastic, i.e., $\bar{A}\mathbf{1} = \mathbf{1}$, and with strictly positive diagonal entries, i.e., $\bar{A}_{ii} > 0$, for all $i \in \mathcal{V}$. \square

In the following section, we discuss the control policies.

III. CONTROL POLICIES

To mitigate the spread of a disease, a policy maker can apply *endogenous* or *exogenous* NPIs to each local community. The former category includes all those measures put in place inside the local community i , e.g., lock-downs, usage of face masks or encouraging social distancing. The latter instead concerns those measures that limit the inflow of people from neighboring communities, i.e., implementation of travel bans or requesting quarantine upon entrance. It is clear that these policies directly affect the dynamics in (1). Specifically, for each community i , the endogenous measures reduce the infection rate $\bar{\beta}_i$, while the exogenous ones act on \bar{A} . The static values of $\bar{\beta}_i$ and \bar{A} refer to the *uncontrolled* evolution of the epidemic, while the intervention of the policy makers transforms these parameters into time-varying functions $\beta_i(t)$ and $A(t)$, which can be seen as inputs to be designed in a (possibly) optimal way to control the epidemic spreading. Thus, the dynamics of the *controlled* evolution of the SIS epidemic become

$$x_i^{\dagger} = (1 - \mu)x_i + (1 - x_i)\beta_i(t) \sum_{j=1}^N A_{ij}(t)x_j, \quad (2)$$

as illustrated in the schematic in Fig. 1.

If the policies put in place are effective, the value of the infection rate should decrease, i.e., $\beta_i(t) \in [0, \bar{\beta}_i]$ for all $t \in \mathbb{N}$. We consider policies that affect the uncontrolled coefficients in (1) linearly. So, we define time-varying functions $\beta_i(t)$ and $A(t)$ in (2) as

$$\beta_i(t) := \bar{\beta}_i - v_i(t), \quad \forall i \in \mathcal{V}, \quad (3a)$$

$$A_{ij}(t) := \bar{A}_{ij} - W_{ij}(t), \quad \forall i, j \in \mathcal{V}, \quad (3b)$$

where $v_i(t) \in [0, \bar{\beta}_i]$. Furthermore, it is reasonable to assume that the policy cannot implement new connections among the community, this translates into the following set of constraints on $W(t) := [W_{ij}(t)] \in \mathbb{R}^{N \times N}$

$$W(t)\mathbf{1} = \mathbf{0}, \quad (4a)$$

$$\bar{A}_{ij} - W_{ij}(t) \geq 0, \quad (4b)$$

$$\bar{A}_{ij} = 0 \implies W_{ij}(t) = 0. \quad (4c)$$

From (3b) and (4), it follows that the time-varying weighted adjacency matrix $A(t)$ defines a directed subgraph $\mathcal{G}(t) \subseteq \mathcal{G}$, which is composed of the same nodes and a subset of the edges of \mathcal{G} , i.e., $\mathcal{G}(t) := (\mathcal{V}, \mathcal{E}(t))$ with $\mathcal{E}(t) \subseteq \mathcal{E}$, so it may not be connected at every $t \in \mathbb{N}$.

IV. OPTIMAL CONTROL POLICY

In this section, we present the main result of this paper on optimal control policies.

A. Problem formulation

The adoption of stringent policies carries costs including both monetary (e.g. recession) and social (e.g. personal restrictions) aspects [3]. In this section, we formalize the problem of choosing the optimal NPIs that minimize the cost while keeping under control the epidemic. Therefore, at each time instant $t \in \mathbb{N}$, we have to design the values of the control actions $W(t)$ and $v_i(t)$ for all time instants k in the prediction horizon $\mathcal{T} := \{t, t+1, \dots, t+T_h\}$, the length of which ($T_h \in \mathbb{N}$) may vary due to the specific epidemic. In the remainder, we use the index k when we refer to an instant belonging to \mathcal{T} and t otherwise.

Each population $i \in \mathcal{V}$ is assumed to establish a desired trajectory $\hat{x}_i(k)$, for all $k \in \mathcal{T} \setminus \{t\}$, that is the fraction of infected population $x_i(k)$ considered to be acceptable at time k . An interesting case is the one in which the value of $\hat{x}_i(k)$ is greater than 0 at the beginning of \mathcal{T} , and it decreases over time, i.e., $\hat{x}(k) \geq \hat{x}(k+1)$ for all $k, k+1 \in \mathcal{T} \setminus \{t\}$. Its slope suggests how aggressive the desired community's policies should be. Notice that the terminal value $\hat{x}_i(t+T_h+1)$ is not necessarily constrained to be 0. In fact, it is reasonable that some communities accept a small fraction of infected in exchange for relaxed NPIs.

Next, let us define the vector of all policies put in place at time instant t by agent i as the $(N+1)$ -dimensional vector $\mathbf{u}_i(t) := \text{col}(v_i(t), W_{i1}(t), \dots, W_{iN}(t))$. The cost for each community i consists of two antagonizing components. The first is the *health-care cost* $c_i^{\text{HC}}(x_i(t); \hat{x}_i(t))$ due to the presence of more infected than the desired ones $\hat{x}_i(t)$, which is individually decided by each community and communicated to the health authority. If the optimal policies leads to $x_i(t) \leq \hat{x}_i(t)$, then the cost is 0; otherwise it is assumed quadratic in the difference with respect to the desired value, and thus it reads as

$$c_i^{\text{HC}}(x_i(t); \hat{x}_i(t)) := q_i(t) (\max\{0, x_i(t) - \hat{x}_i(t)\})^2,$$

with $q_i(t) > 0$ being a (possibly time-varying) weight. On the other hand, we consider a (quadratic) *control cost* associated with the implementation of control policies. So, the global cost that the N communities face over the prediction horizon reads as

$$J(\mathbf{x}, \mathbf{u}) := \sum_{k \in \mathcal{T}} \sum_{i \in \mathcal{V}} c_i^{\text{HC}}(x_i(k+1); \hat{x}_i(k+1)) + \|\mathbf{u}_i(k)\|_{S_i(t)}^2,$$

where $S_i(k) \succ 0$ is the diagonal matrix of the $N+1$ weights associated to the NPIs $\mathbf{u}_i(k)$. The fraction of infected and the adopted policies of all the populations at t are $\mathbf{x}(t)$ and $\mathbf{u}(t) := \text{col}((\mathbf{u}_i(t))_{i \in \mathcal{V}})$, respectively. The pair $\mathbf{x} = \text{col}((\mathbf{x}(t+1))_{t \in \mathcal{T}})$ and $\mathbf{u} = \text{col}((\mathbf{u}(t))_{t \in \mathcal{T}})$ denotes in compact form all the variables involved. In our general formulation, the weights associated to the health-care cost (i.e., $q_i(t)$) and to the control cost (i.e., $S_i(t)$) are time-varying, since the same fraction of infected individuals or the same level of NPIs may yield different costs depending on the timing, for example due to the increasing hospital preparedness or the accumulation of socio-economic costs. It is worth noticing that only $2P - 2N$ values of $W(k)$ can be freely designed at each time instant $k \in \mathcal{T}$. In fact, for each one of the $P - N$ off-diagonal edges, we can set $W_{ij}(k)$ and $W_{ji}(k)$; then the weight associated with the self-loop is constrained by (4a). Thus, the number of elements of $\mathbf{u}(t)$ to be optimized is $2P - N$.

Finally, we cast the problem of designing the optimal policies \mathbf{u}^* for the control dynamics (2)

$$\begin{aligned} (\mathbf{x}^*, \mathbf{u}^*) &= \underset{\mathbf{x}, \mathbf{u}}{\text{argmin}} J(\mathbf{x}, \mathbf{u}), \\ &\text{s.t. } (\mathbf{x}, \mathbf{u}) \in \Omega, \end{aligned} \quad (\mathcal{P})$$

where

$$\Omega := \{(\mathbf{x}, \mathbf{u}) \mid v_i(t) \in [0, \bar{\beta}_i], (2), (3), (4) \text{ hold } \forall k \in \mathcal{T}\}.$$

This problem belongs to the class of Nonlinear Model Predictive Control (NL-MPC), due to the highly nonlinear and nonconvex controlled SIS dynamics in (2). It is well known that these problems are hard to solve in their original form. Nevertheless, we establish that a feasible solution to (\mathcal{P}) always exists, and thus that the problem is worth to be studied.

Proposition 1 (Solution existence): For every initial condition $\mathbf{x}(t) \in [0, 1]^N$, there exists at least one pair $(\mathbf{x}^*, \mathbf{u}^*)$ that is a solution to (\mathcal{P}) . \square

Proof: From the constraints' definition (2) and (4), it follows that the pair $\mathbf{u} = \mathbf{0}$ and \mathbf{x} , obtained via (1), always lies in Ω , hence $\Omega \neq \emptyset$. The box constraints on the control inputs and the fact that the dynamics (2) are positively invariant with respect to $[0, 1]^N$, see [28, Lem. 1], implying that Ω is bounded. For every constraint, the set of points satisfying it is closed. Therefore, Ω is the intersection of closed sets, and hence it is closed itself. Since J is continuous, we invoke the Weierstrass theorem to conclude the existence of $(\mathbf{x}^*, \mathbf{u}^*) \in \Omega$ that minimizes J , and, in turns, solves (\mathcal{P}) . \blacksquare

B. NL-MPC solution algorithm

In the literature, there are several approaches to solve NL-MPC. Specifically, in order to solve nonlinear constrained optimization problems with differentiable cost functions, the most popular approaches are based on the Sequential Quadratic Programming (SQP). As shown in [30, Ch. 18], [31] and reference there in, these methods provide excellent convergence properties and ensure a fast convergence to a (local) optimum of the original nonlinear problem.




Fig. 2: Control scheme for the design and implementation via receding horizon of optimal NPIs to flatten the pandemic curve.

The SQP is an iterative algorithm in which, during each iteration p , a candidate optimal trajectory $(\mathbf{x}^p, \mathbf{u}^p)$ is computed as the solution of a Quadratic Programming (QP), obtained by linearizing the constraints and approximating the cost via a quadratic function, see [30, Alg. 18.3]. The linearization is performed with respect to the trajectory $(\tilde{\mathbf{x}}^{p-1}, \tilde{\mathbf{u}}^{p-1})$ defined as

$$(\tilde{\mathbf{x}}^{p-1}, \tilde{\mathbf{u}}^{p-1}) = (\tilde{\mathbf{x}}^{p-2}, \tilde{\mathbf{u}}^{p-2}) + \alpha^{p-1}(\mathbf{x}^{p-1}, \mathbf{u}^{p-1}),$$

where $(\mathbf{x}^{p-1}, \mathbf{u}^{p-1})$ is the solution of the QP solved at the previous iteration, and α^{p-1} can be computed for example via line search as in [30, Eq. 18.28]. If the candidate solution satisfies some convergence conditions then $(\mathbf{x}^*, \mathbf{u}^*) = (\mathbf{x}^p, \mathbf{u}^p)$; otherwise a new iteration is performed.

The sole nonlinearity in the constraints of (\mathcal{P}) is associated with the dynamics in (2). The effectiveness of the NPIs depends on the predictive accuracy of the linearized dynamics. This difference is studied via numerical simulations in Section V, where it is shown that the linearized dynamics allow to obtain a very high prediction accuracy for a sufficiently long prediction horizon. Thus, the used NL-MPC algorithm generates an effective solution for the control problem.

Finally, in Figure 2 we depict the complete control scheme to solve the problem of designing optimal NPIs to control an SIS type dynamics. Specifically, at each time instant t an instance of (\mathcal{P}) is cast. The optimal trajectory $(\mathbf{x}^*, \mathbf{u}^*)$ is computed over \mathcal{T} via SQP. Finally, a receding horizon approach is implemented by applying only NPIs associated to the first instant, i.e., $\mathbf{u}^*(t)$, and then the loop starts again. This implementation allows us to minimise the prediction error inherently present in the prediction of the model.

V. SIMULATIONS

In this section, we present several numerical simulations that validate the procedure proposed to design optimal NPIs and provide insightful information on how different parameter choices (e.g., the weights on the control and the length of the prediction horizon) influence the final optimal policies.

A. Reference tracking and reproduction number

The simulations are performed over a randomly generated connected network with $N = 4$ that is described by a weighted adjacency matrix with off-diagonal nonzero entries chosen from a normal distribution with mean 0.1 and variance of 0.2, and diagonal entries such that $\bar{\mathbf{A}}\mathbf{1} = \mathbf{1}$, obtaining

$$\bar{\mathbf{A}} = \begin{bmatrix} 0.7 & 0.17 & 0 & 0.13 \\ 0.42 & 0.31 & 0.16 & 0.11 \\ 0 & 0.12 & 0.88 & 0 \\ 0.28 & 0.1 & 0 & 0.62 \end{bmatrix}. \quad (5)$$




Fig. 3: In (a), we show the trajectories of the uncontrolled dynamics $x_i^{uc}(t)$ (red dashed curve), of the reference $\hat{x}_i(t)$ (blue solid curve), and of the controlled dynamics $x_i(t)$ obtained via the control scheme in Figure 2 (red solid curve) for all $i \in \mathcal{V}$. In (b), we plot the amount of endogenous (blue) and exogenous (green) control applied to every community $i \in \mathcal{V}$ to obtain $x_i(t)$, respectively $\beta_i(t)$ and $\|W_i(t)\|_2^2$.

The uncontrolled infection rate $\bar{\beta}_i$ of each node is selected uniformly and randomly in the interval $[0.3, 0.6]$, obtaining $\bar{\beta} = [0.3, 0.59, 0.3, 0.45]$. The recovery rate is chosen to be $\mu = 0.15$. So, the uncontrolled dynamics converge over time to an endemic equilibrium [28]. Note that, since μ is the inverse of the time for recovering, if one time-step corresponds to a day, then this setup may represent a disease like seasonal flu, for which (approximately) 6 days are enough to recover. On the other hand, if one time-step is interpreted as one week, the setup may capture diseases with longer recovery time, e.g., gonorrhea. The initial fraction of infected for each population is set as $x(0) = [0.65, 0.55, 0.75, 0.40]$, that is, a scenario of an endemic disease. We denote the trajectory of the fraction of infected individuals in the uncontrolled dynamics (1) by $x^{uc}(t)$. The desired reference trajectory $\hat{x}_i(t)$ of each community is assumed to start close to $x_i(0)$ and linearly decrease until it reaches its terminal value at $t = 20$ that is $\hat{x}(20) = [0.1168, 0.0548, 0.0856, 0.1175]$. The length of the prediction horizon is set equal to $T_h = 10$ and we will discuss its optimal value in Section V-B. The weights $q(t)$ and $S(t)$ are chosen constant in time, uniform across the different nodes, and equal to $q_i(t) = 1$, for all $i \in \mathcal{V}$, while S is so that all off-diagonal terms are equal to 0, the diagonal terms that correspond to entries of v are equal to 0.2 and those that correspond to entries of W are equal to 0.05. Note that the weights are such that applying the same

amount of exogenous and endogenous control respectively have similar costs.

In Figure 3a, we present the trajectory of $x_i(t)$, obtained from (2), by following the control scheme in Figure 2 (red solid curves), compared with the value of the uncontrolled SIS dynamics (red dashed curves), and the reference (blue solid curve). As expected from having a higher weight on the healthcare cost, the value of $x_i(t)$ remains relatively close to the desired one $\hat{x}_i(t)$. From Figure 3b, it is clear that, in this scenario, acting on local restrictions as lock-down and social distancing is more effective than implementing travel bans. This is consistent with empirical observations during the ongoing COVID-19 pandemic, suggesting that travel bans are more effective in the early stages of an epidemic outbreak, i.e., when $x(0)$ is close to $\mathbf{0}$ [15]. In the first instances, the control action is more intense due to the decreasing reference and, after $t = 20$, it is relaxed. This can also be assessed via the ratio β_i/μ depicted in Figure 4a. In fact, during the initial time instants, all this ratios decrease reaching values lower than 1. This suggests that, by adopting the NPIs implemented in the first phase for the entire duration of the simulation, the trajectory would converge to the disease-free equilibrium, at the cost of dramatically increasing the social and economical impact of the control.

In Figure 4a, we also report the evolution of the singular value $\sigma(B(t)A(t))/\mu$ where $B(t) := \text{diag}(\beta(t))$, which is known to determine the behavior of the epidemic process [5],




Fig. 4: In (a), we show the evolution of the ratio $\beta_i(t)/\mu$ for all the communities and the value of $\sigma(B(t)A(t))/\mu$. In (b), we compare the trajectory predicted by the linearized SIS model, i.e., $x^l(t)$ (red) and the one obtained from the open loop controlled nonlinear dynamics $x^{nl}(t)$ (blue).

[28]. It is also strictly related to the effective reproduction number R_t , which can be obtained by multiplying such a quantity by the fraction of susceptible individuals in the population. As expected, the value at which $\sigma(B(t)A(t))/\mu$ settles is above 1, since the terminal value of the reference $\hat{x}(t)$ is greater than the disease-free equilibrium. It is interesting to notice that the optimal control $W(t)$ tends to isolate the communities by increasing the weight of each self-loop and decreasing the off-diagonal entries. Analyzing more complex graphs may lead to the emergence of different patterns in the exogenous control actions, and this is left for future research.

In Figure 4b, the accuracy of the linearized model used to solve (\mathcal{P}) at time $t = 4$ is shown. By solving the optimization problem, we obtain the trajectories $\mathbf{x}^*(t)$ and $\mathbf{u}^*(t)$ for $t \in \mathcal{T} = \{4, \dots, 14\}$, which we denote by $\mathbf{x}^l(t)$. We compute $\mathbf{x}^{nl}(t)$ by applying the whole \mathbf{u}^* to (2) and letting it evolve in an open loop. The estimation of the nonlinear behavior of the dynamics is almost perfect during the first six time instants of \mathcal{T} , while it becomes less accurate during the second half. An even longer prediction horizon would exacerbate this difference affecting the performance, as we show in the next section.

Overall, the policies computed by the proposed control scheme generate noteworthy effects in controlling the epidemics by achieving the two main objectives of the proposed approach, i.e., an acceptable level of infected and at a low social and economical price. Nevertheless, it is evident that the optimal solution to the problem (\mathcal{P}) highly depends on the choice of the weights $q_i(t)$ and $S_i(t)$ in the cost function $J(\mathbf{x}, \mathbf{u})$. To better understand the effect of the endogenous and exogenous measures on the (steady-state) health state of the overall population, we have performed a sensitivity analysis of the weights applied to the control action. For all i and t , we set $q_i(t) = 1$ and $S_i(t)$ equal to the block diagonal matrix that has on the diagonal s_v and $\frac{s_w}{N}I$, which are the weights on $v_i(t)$ and $W_i(t)$, respectively. In order to compare the performance obtained with different weights,




Fig. 5: Results of the sensitivity analysis with respect to the endogenous (s_v) and exogenous (s_w) control costs. The heat-map represents the value of the index π , defined in (6), for the different values of the control cost parameters.

we introduce the following index:

$$\pi := \frac{\|\mathbf{x}(\infty) - \hat{\mathbf{x}}(\infty)\|}{\|\mathbf{x}(\infty) - \mathbf{x}^{uc}(\infty)\|}. \quad (6)$$

As shown in Figure 5, we obtain larger values of π (i.e., the system converges close to the endemic equilibrium \mathbf{x}^{uc}) when the weights of the control actions are large, while smaller values of π (i.e., the system converges close to the desired reference $\hat{\mathbf{x}}$) arise if s_v and s_w are small. Furthermore, from Figure 5 one can observe that changing the weight s_v makes the index π change more significantly than changing the weight of the exogenous actions. Then, we can conclude that the endogenous actions are more effective than the exogenous ones. However, we can also observe that when the cost of implementing for instance lock-downs is very high (i.e., larger values of s_v), putting in place travel bans is very beneficial.

B. Myopic vs predictive policies

Next, we consider the problem of choosing the optimal length for the prediction horizon. This is a critical choice that a policy maker has to perform. In fact, if T_h is too small the policies will be myopic failing to prepare in time for the future evolution of the epidemics. On the other hand, longer prediction horizons will require an increased computational effort. Moreover, due to the difference between the real dynamics and the ones used in (\mathcal{P}) , a long prediction horizon may lead to inaccurate estimation and consequently to incorrect precautionary policies, see Figure 4b. Therefore, we believe that there is an optimal length for the prediction horizon.

To validate this claim via simulations, we generate 50 different scenarios in which \bar{A} , $\bar{\beta}$ and \hat{x} are randomly chosen as described in the previous section, while the rest of the coefficients of the problem do not change. As index of the performance we consider the sum over the whole simulation of the cost actually experienced by the communities, we denote it simply by J . Figure 6a clearly depicts a Pareto front in which, after an initial reduction of J due to the growing T_h , there is a diminishing return. Moreover, for even greater values of T_h the performance starts to slowly get worse. The computational time of each one of the simulation grows




Fig. 6: In (a), we illustrate the value of J , i.e., the sum of the cost over the whole simulations, as the prediction horizon T_h grows. In (b), we show the computational time for one simulation for growing values of T_h . In both figures, the vertical bars represent 95% confidence intervals computed over 50 independent realizations of the random variables \bar{A} , $\bar{\beta}$, and \hat{x} .

linearly with the value of T_h , as shown in Figure 6b, this can be a serious issue in the case of large scale networks. Therefore, the optimal value of T_h for the simulated cases is in the range 5 – 7. In fact, it is large enough to benefit from the initial steep increment in terms of performance but not that large to make the estimations less accurate.

VI. CONCLUSION AND OUTLOOK

In this paper, the problem of designing optimal NPIs to ensure containing an epidemic on a network while minimizing the costs due to restrictive policies has been cast as a nonlinear constrained optimization problem over a prediction horizon. We have proposed a solution via a linearization technique and a receding horizon strategy, in which the real evolution of the disease is used as a feedback to cast a new instance of the optimization problem and compute the upcoming policies that have to be implemented.

The performance of the proposed control scheme has been demonstrated via preliminary numerical simulations. Specifically, the good accuracy of the linearized dynamics for sufficiently long time horizons guarantees accurate farsighted forecasts, avoiding thus the risks of relying on myopic policies. Our numerical findings have shown nontrivial solutions for the optimal control strategies, depending on the costs associated with the implementation of the different types of NPIs and on the requirements on the desired evolution of the epidemic outbreak. We finally would like to stress that, while the current formulation has been developed using an SIS epidemic model, the same control scheme can be extended to more complex epidemic models, such as those used to capture the features of the ongoing COVID-19 pandemic [9]. Such an extension, and the analysis of the proposed control scheme in real-world scenarios is an important direction for the future research.

REFERENCES

- [1] N. Haug *et al.*, “Ranking the effectiveness of worldwide COVID-19 government interventions,” *Nature Human Behaviour*, vol. 4, no. 12, pp. 1303–1312, 2020.
- [2] K. Prem *et al.*, “The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study,” *The Lancet Public Health*, vol. 5, no. 5, pp. e261–e270, 2020.
- [3] G. Bonaccorsi *et al.*, “Economic and social consequences of human mobility restrictions under COVID-19,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 27, pp. 15 530–15 535, 2020.
- [4] C. Nowzari, V. M. Preciado, and G. J. Pappas, “Analysis and Control of Epidemics: A Survey of Spreading Processes on Complex Networks,” *IEEE Control Systems Magazine*, vol. 36, no. 1, pp. 26–46, 2016.
- [5] W. Mei, S. Mohagheghi, S. Zampieri, and F. Bullo, “On the dynamics of deterministic epidemic propagation over networks,” *Annual Reviews in Control*, vol. 44, pp. 116–128, 2017.
- [6] P. E. Paré, C. L. Beck, and T. Başar, “Modeling, estimation, and analysis of epidemics over networks: An overview,” *Annual Reviews in Control*, vol. 50, pp. 345–360, 2020.
- [7] L. Zino and M. Cao, “Analysis, Prediction, and Control of Epidemics: A Survey from Scalar to Dynamic Network Models,” *IEEE Circuits and Systems Magazine*, 2021, to appear (arXiv:2103.00181).
- [8] M. Ye, J. Liu, C. Cenedese, Z. Sun, and M. Cao, “A network sis meta-population model with transportation flow,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 2562–2567, 2020, 21st IFAC World Congress.
- [9] G. Giordano *et al.*, “Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy,” *Nature Medicine*, vol. 26, no. 6, pp. 855–860, 2020.
- [10] F. Casella, “Can the COVID-19 Epidemic Be Controlled on the Basis of Daily Test Reports?” *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1079–1084, 2021.
- [11] G. C. Calafiore, C. Novara, and C. Possieri, “A time-varying SIRD model for the COVID-19 contagion in Italy,” *Annual Reviews in Control*, vol. 50, pp. 361–372, 2020.
- [12] M. Gatto *et al.*, “Spread and dynamics of the COVID-19 epidemic in Italy: Effects of emergency containment measures,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 19, pp. 10 484–10 491, 2020.
- [13] F. Della Rossa *et al.*, “A network model of Italy shows that intermittent regional strategies can alleviate the COVID-19 epidemic,” *Nature Communications*, vol. 11, no. 1, 2020.
- [14] R. Carli, G. Cavone, N. Epicoco, P. Scarabaggio, and M. Dotoli, “Model predictive control to mitigate the COVID-19 outbreak in a multi-region scenario,” *Annual Reviews in Control*, vol. 50, pp. 373 – 393, 2020.
- [15] F. Parino, L. Zino, M. Porfiri, and A. Rizzo, “Modelling and predicting the effect of social distancing and travel restrictions on COVID-19 spreading,” *Journal of the Royal Society Interface*, p. 20200875, 2021.
- [16] J. Köhler, L. Schwenkel, A. Koch, J. Berberich, P. Pauli, and F. Allgöwer, “Robust and optimal predictive control of the COVID-19 outbreak,” *Annual Reviews in Control*, vol. 51, pp. 525–539, 2021.
- [17] M. Bin *et al.*, “Post-lockdown abatement of covid-19 by fast periodic switching,” *PLOS Computational Biology*, vol. 17, no. 1, pp. 1–34, 01 2021.
- [18] F. Parino, L. Zino, G. C. Calafiore, and A. Rizzo, “A model predictive control approach to optimally devise a two-dose vaccination rollout: A case study on COVID-19 in Italy,” *International Journal of Robust and Nonlinear Control*, 2021.
- [19] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, “Optimal Resource Allocation for Network Protection Against Spreading Processes,” *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 99–108, 2014.
- [20] E. Ramírez-Llanos and S. Martínez, “Distributed discrete-time optimization algorithms with applications to resource allocation in epidemics control,” *Optimal Control Applications & Methods*, vol. 39, no. 1, pp. 160–180, 2018.
- [21] V. S. Mai, A. Battou, and K. Mills, “Distributed Algorithm for Suppressing Epidemic Spread in Networks,” *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 555–560, 2018.
- [22] S. Eshghi, M. H. R. Khouzani, S. Sarkar, and S. S. Venkatesh, “Optimal Patching in Clustered Malware Epidemics,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 283–298, 2016.
- [23] P. Grandits, R. M. Kovacevic, and V. M. Veliov, “Optimal control and the value of information for a stochastic epidemiological SIS-model,” *Journal of Mathematical Analysis and Applications*, vol. 476, no. 2, pp. 665 – 695, 2019.
- [24] F. Sélley, Á. Besenyei, I. Z. Kiss, and P. L. Simon, “Dynamic Control of Modern, Network-Based Epidemic Models,” *SIAM Journal on Applied Dynamical Systems*, vol. 14, no. 1, pp. 168–187, 2015.

- [25] J. Köhler, C. Enyioha, and F. Allgöwer, "Dynamic Resource Allocation to Control Epidemic Outbreaks A Model Predictive Control Approach," in *2018 Annual American Control Conference*, 2018, pp. 1546–1551.
- [26] N. J. Watkins, C. Nowzari, and G. J. Pappas, "Robust Economic Model Predictive Control of Continuous-Time Epidemic Processes," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 1116–1131, 2020.
- [27] M. M. Morato, S. B. Bastos, D. O. Cajueiro, and J. E. Normey-Rico, "An optimal predictive control strategy for COVID-19 (SARS-CoV-2) social distancing policies in Brazil," *Annual Reviews in Control*, vol. 50, pp. 417–431, 2020.
- [28] P. E. Paré, J. Liu, C. L. Beck, B. E. Kirwan, and T. Başar, "Analysis, Estimation, and Validation of Discrete-Time Epidemic Processes," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 1, pp. 79–93, 2020.
- [29] F. Liu, S. CUI, X. Li, and M. Buss, "On the stability of the endemic equilibrium of a discrete-time networked epidemic model," vol. 53, no. 2, 2020, pp. 2576–2581, 21st IFAC World Congress.
- [30] J. Nocedal and S. J. Wright, "Sequential quadratic programming," *Numerical optimization*, pp. 529–562, 2006.
- [31] K. Schittkowski, "NLPQL: A fortran subroutine solving constrained nonlinear programming problems," *Annals of Operations Research*, vol. 5, pp. 485–500, 1986.