

Covid-19 and Flattening the Curve: A Feedback Control Perspective

Francesco Di Lauro^{ID}, István Zoltán Kiss^{ID}, Daniela Rus^{ID}, *Fellow, IEEE*,
and Cosimo Della Santina^{ID}, *Member, IEEE*

Abstract—Many of the policies that were put into place during the Covid-19 pandemic had a common goal: to flatten the curve of the number of infected people so that its peak remains under a critical threshold. This letter considers the challenge of engineering a strategy that enforces such a goal using control theory. We introduce a simple formulation of the optimal flattening problem, and provide a closed form solution. This is augmented through nonlinear closed loop tracking of the nominal solution, with the aim of ensuring close-to-optimal performance under uncertain conditions. A key contribution of this letter is to provide validation of the method with extensive and realistic simulations in a Covid-19 scenario, with particular focus on the case of Codogno - a small city in Northern Italy that has been among the most harshly hit by the pandemic.

Index Terms—Emerging control applications, network analysis and control, large-scale systems.

I. INTRODUCTION

DEFINING and implementing social distancing protocols (SD) is a significant challenge with economical, political, and scientific considerations. The definition of a clear or optimal goal remains unclear. As an example, consider the direct reduction of deaths by Covid-19. Imposing this goal requires the harshest measures possible, for an indefinite period of time. According to the available models [1] a monotonic relationship exists between this cost function and the SD level. Yet, this strategy has many potential drawbacks. First, extreme levels of lockdown are unsustainable in the long run, due to the vast range of pernicious secondary effects (e.g., poverty [2], mental illnesses [3]) which in turn are themselves

Manuscript received August 17, 2020; revised October 22, 2020; accepted November 8, 2020. Date of publication November 19, 2020; date of current version December 10, 2020. The work of Francesco Di Lauro and István Zoltán Kiss was supported by the Leverhulme Trust for the Research Project under Grant RPG2017-370. The work of Cosimo Della Santina was supported by the TU Delft COVID-19 Response Fund. Recommended by Senior Editor L. Menini. (Corresponding author: Cosimo Della Santina.)

Francesco Di Lauro and István Zoltán Kiss are with the Department of Mathematics, University of Sussex, Brighton BN1 9QH, U.K.

Daniela Rus is with the MIT Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Cosimo Della Santina is with the Cognitive Robotics Department, Delft University of Technology, 2628 CD Delft, The Netherlands, and also with the Institute of Robotics and Mechatronics, German Aerospace Center (DLR), 82234 Weßling, Germany (e-mail: cosimodelasantina@gmail.com).

Digital Object Identifier 10.1109/LCSYS.2020.3039322




Fig. 1. The aim of this letter is to devise a control strategy that achieves the curve flattening goal, which should result in a curve similar to the green one. The two alternative extreme cases are shown as comparison: the result of no SD is shown in red, and of full lockdown in yellow.

associated with a rise in mortality. Additionally, relaxing or lifting control after a harsh lockdown may lead to a second wave, possibly more critical than the first one [4]. Another strategy would be to let the epidemic spread freely (red curve in Fig. 1) to get herd immunity as fast as possible. This is also hardly acceptable, as it would lead to higher mortality [5], and to a prolonged stress of the health care system. The “*flattening the curve*” strategy provides a third option, which promises to combine the benefits of the two extremes [6]. The key idea (of which Fig. 1 provides a visual representation) is to allow some level of disease spreading, while ensuring that people seeking medical assistance can access the health care system.

A vast pre-Covid-19 pandemic literature [7] on designing controllers for dealing with epidemics already exists. However, none of these works tackled the curve flattening goal, since no pandemics before threatened to overburden the healthcare system on such a large scale. In the context of Covid-19, open loop optimal control is proposed in [8] for selecting the optimal timing of a time-limited lockdown, and in [9] the authors find a trade-off between number of deaths and damage to the economy. Yet, feed-forward strategies are quite prone to uncertainties naturally associated with epidemics [10]. More robust strategies have been proposed, relying on feedback control. A linear controller is proposed in [11]. A fast switching strategy with duty cycle selected through a slow feedback is discussed in [12]. In [13], the loop is closed by periodically re-planning the optimal action, in a model-predictive-control fashion. An explicit formulation of curve flattening is instead provided in [14], where an open loop strategy is devised so to optimally reduce the infectious peak. An interesting alternative is discussed in [15], where a trade-off between the health care

and the socio-economic cost of the pandemic is proposed, and the limited capacity level of intensive care units is imposed as a constraint. Both these solutions are open loop.

This letter investigates the use of feedback control theory as a tool for engineering an effective curve flattening strategy. We wish to design a simple rule that can be implemented on a local level, without the need of accessing specialized facilities to run complex optimization routines. We perform extensive simulations of epidemics on networks [1], [16], with conditions inspired by real Covid-19 scenarios. This is as far as we know the first time that such analysis is carried out for Covid-19 control related research. We remark that the acceptable level of ‘‘curve flattening’’ is to be decided by policy makers, based upon cost-benefit analysis. However, once an optimal curve has been identified, this letter offers a novel, theoretically-backed strategy that guarantees that the goal of controlling the epidemic curve is achieved.

II. BACKGROUND: MODEL OF THE EPIDEMICS WITH DYNAMIC INTERVENTIONS

Consider a fixed population of N individuals, and a disease spreading among them, through direct contacts. Each individual can be in either of three states: (i) susceptible, meaning that they can be infected by the pathogen; (ii) infected, meaning that they contracted the pathogen and they can now infect other susceptible people; (iii) recovered -and therefore immune, or removed. We denote with $S(t)$, $I(t)$, $R(t)$ the number of people at time t who are susceptible, infected or recovered, respectively. We have that $S(t) + I(t) + R(t) = N$. We can therefore neglect the study of R , as its value can always be recovered from S , I and N . If the population is well mixed, the evolution of the disease can be described by the SIR model

$$\dot{s}(t) = -\beta(t)\iota(t)s(t), \quad \dot{\iota}(t) = +\beta(t)\iota(t)s(t) - \gamma\iota(t), \quad (1)$$

where $s(t)$ and $\iota(t)$ are the system state, indicating respectively the number of susceptible $S(t)$ and infectious $I(t)$, divided by the total population N . Note that, despite its simplicity, the SIR model has proven able to match real data when applied to Covid-19 [6], [14], and it is therefore widely used in the literature. Without loss of generality, we consider that, at $t = 0$, $s + \iota = 1$. The constant $\gamma \geq 0$ defines the transition rate from the pool of infected, to the compartment of recovered/removed. β is the rate at which an infected individual makes disease-transmitting contacts with other people. When SD policies are imposed, the value of β varies, $0 < \beta_m \leq \beta \leq \beta_M$, with β_m corresponding to total lockdown. Therefore β can vary in time, and it is the control input of (1).

III. CONTROL STRATEGY

We propose here a control strategy acting on system (1). As shown by Fig. 2, this architecture is made of two components: (i) an optimal open loop action, and (ii) a feedback controller implementing trajectory tracking.

A. Optimal Curve Flattening Under Nominal Conditions

Our aim here is to introduce a nominal strategy (‘‘Optimal Solution’’ in Fig. 2) for optimally flattening the epidemic curve $\iota(t)$, so to keep the number of infected people ι within the maximum capacity of the health care system, $\iota_{th} > 0$. This can, for instance, be evaluated by considering the percentage of people that will need Intensive Care Units (ICUs), which are probably the most critically limited resources. As discussed in the introduction, enforcing this constraint is of paramount

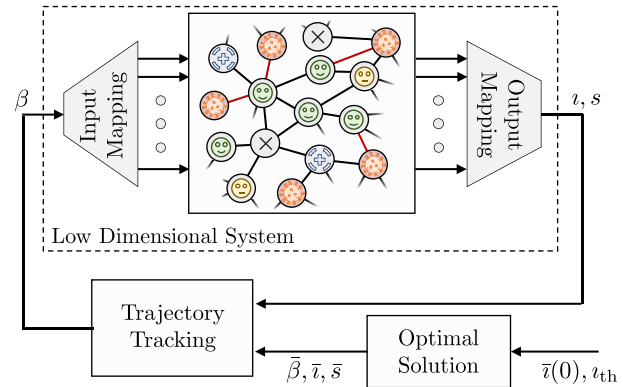


Fig. 2. Block diagram of the strategy proposed in this letter. The input and output maps reduce the high-dimensional dynamics of the outbreak to a simpler evolution of few salient characteristics, namely the prevalence of infected and susceptible ι , s , which are sensible to changes in the level of SD, modelled here as different values of the transmission rate of infection β . A nonlinear feedback controller acts within this representation implementing trajectory tracking of an optimal control policy.

importance, since exceeding it may provoke a critical failure of the healthcare system, leading to a substantial increase in the number of deaths not only from the disease, but also from uncorrelated health issues. On the other hand, we want to keep the level of restriction on the population as low as possible, to minimise secondary negative effects. Note that the curve flattening goal is the result of a careful balance between competing interests, and as such we decide to explicitly impose it as a goal. We consider the case of a constant β . This simplification is instrumental in making the optimal control problem more manageable.

We summarize the above considerations through the optimization problem

$$\max_{\beta \in \mathbb{R}} \beta, \quad \text{s.t. } 0 < \iota(t) \leq \iota_{th} \quad \forall t \text{ and (1).} \quad (2)$$

We now propose a Lemma introducing a general solution to this optimal control problem.

Lemma 1: The following is the closed form solution of (2)

$$\beta = -\frac{\gamma}{1 - \iota_{th}} W_{-1} \left(-\frac{1}{e} \frac{1 - \iota_{th}}{1 - \iota(0)} \right), \quad (3)$$

where W_{-1} is the Lambert W function [17], branch -1 .

Proof: Since the cost function is linear in the optimization parameter, the optimal value is to be found on the boundary of the feasible set, i.e., β has to be such that $\max_t \iota(t) = \iota_{th}$.

The maximum value of ι is given by the non-trivial solution of $\dot{\iota}(t) = 0$. Combining this condition with the second equation in (1) yields $s^+ = \gamma/\beta$. Further, we can combine the two equations in (1) into $d\iota/ds = \gamma/(\beta s) - 1$. This nonlinear ordinary differential equation can be solved together with the initial condition $s(0) = 1 - \iota(0)$, $\iota(0)$, to get

$$\iota(s) = \frac{\gamma}{\beta} \ln \left(\frac{s}{1 - \iota(0)} \right) - s + 1. \quad (4)$$

By inverting $\iota(s^+)$ for β , we get the desired optimal value such that $\max_t \iota(t) = \iota_{th}$. The following is a solution for all integer values of j ,

$$\beta = -\frac{\gamma}{1 - \iota_{th}} W_j \left(-\frac{1}{e} \frac{1 - \iota_{th}}{1 - \iota(0)} \right), \quad (5)$$

where $W_j(\bar{a})$ is the j -th branch of the Lambert W function [17]. Each of the branches is built as the solution of

$\bar{a} = W_j e^{W_j}$. Among all of them, only W_{-1}, W_0 have domain within the real line. Moreover, it is always the case that $W_0 > W_{-1}$, which in turn assures that the larger value of β is always reached for $j = -1$, concluding the proof. ■

It is worth noting that the argument of W_{-1} is always between $-1/e$ and 0 since $\iota(0) \leq \iota_{\text{th}}$. This is exactly the range of arguments for which the W_{-1} is well defined [17].

B. Trajectory Tracking Controller

The following Lemma introduces the tracking controller (“Trajectory Tracking” in Fig. 2) implementing the reactive change of the SD level β . Note that in principle this controller is agnostic to the choice of the reference to be tracked, and it is introduced as such. However, in our architecture we consider the output of Lemma 1 as the reference $\bar{\beta}, \bar{\iota}, \bar{s}$.

Lemma 2: Consider $\psi_i, \psi_s \in \mathbb{R}, \psi_s > 0, \psi_i \geq 0$. Consider also $\bar{s}, \bar{\iota}, \bar{\beta} \in \mathbb{R}$ to be a solution of (1) with $\bar{s}(0), \bar{\iota}(0) > 0$. Then, the feedback loop composed by the action

$$\beta(s, \iota, t) = +\psi_i(\bar{\iota} - \iota) - \psi_s(\bar{s} - s) + \frac{\bar{s}\bar{\iota}}{s\iota} \bar{\beta}, \quad (6)$$

and the SIR model (1), is such that (s, ι) converges uniformly w.r.t. the initial time to $(\bar{s}, \bar{\iota})$ if $s(0), \iota(0) > 0$.

Proof: Consider now the linear change of coordinates $x = -(t+s)/\gamma$. Adding up the two equations in (1), yields $\dot{i} + \dot{s} = -\gamma\iota$. We can therefore establish the change of coordinates

$$\iota = \dot{x}, \quad s = -\gamma x - \dot{x}. \quad (7)$$

Combining the latter, with the second equations in (1) allows writing the following equivalent formulation of the SIR dynamics

$$\ddot{x} = -(\gamma x + \dot{x})\dot{x}\beta - \gamma\dot{x}. \quad (8)$$

We take the following control action

$$\beta(x, \dot{x}, t) = -\frac{\gamma\dot{x} + \ddot{x}}{(\gamma x + \dot{x})\dot{x}} + \alpha_p(\bar{x} - x) + \alpha_d(\dot{\bar{x}} - \dot{x}), \quad (9)$$

with $\alpha_p > \gamma, \alpha_d > 0$ being the gains of a PD-like action. Note that the denominator of the first element of the right hand side of (9) is equal to $s\iota$. Indeed, the hypothesis $s(0), \iota(0) > 0$ implies that $s(t), \iota(t) > 0, \forall t$, since (1) is a strictly positive dynamical system [1]. This produces the closed loop dynamics $\ddot{e} = -(\gamma + \alpha_d(-\gamma x - \dot{x}))\dot{e} - \alpha_p(-\gamma x - \dot{x})e$, where $e = \bar{x} - x$. As mentioned before $s, \iota > 0$, therefore $\gamma + \alpha_d(-\gamma x - \dot{x})\dot{x} = \gamma + \alpha_d s\iota > 0$ and $\alpha_p(-\gamma x - \dot{x})\dot{x} = \alpha_p s\iota > 0$, for all $t < \infty$. Thus, we propose the following energy-like Lyapunov candidate $V(e, \dot{e}, t) = \dot{e}^2/2 + \int_0^e (\alpha_p(-\gamma(\bar{x} - \epsilon) - \dot{\bar{x}})\dot{\bar{x}})\epsilon d\epsilon$, which is positive definite. We perform time differentiation, obtaining through standard manipulations $\dot{V} = -(\gamma + \alpha_d s\iota)\dot{e}^2 \leq 0$, since $s, \iota > 0$. This proves that (e, \dot{e}) are bounded. Yet, it is not enough to prove stability, since the error dynamics is time-variant. Indeed, s, ι are a combination of $x, \dot{x}, \bar{x}, \dot{\bar{x}}$, the latter two being explicit functions of time. To prove asymptotic convergence we invoke Barbalat’s Lemma, which requires to perform a further differentiation $\dot{V} = -(\gamma + \alpha_d s\iota)\dot{e}^2 - \alpha_d(\dot{s}\iota + s\dot{\iota})\dot{e}^2$ which is bounded since it is sum and products of bounded functions. Therefore, (e, \dot{e}) converges asymptotically to $(0, 0)$, and as a consequence (s, ι) converges to $(\bar{s}, \bar{\iota})$.

To conclude the proof, we need to show now that (6) and (9) are equivalent. First, we use (8) to obtain $\ddot{x} = (\gamma\bar{x} + \dot{\bar{x}})\dot{\bar{x}} - \gamma\dot{\bar{x}}$. We then take $\psi_s = \alpha_p/\gamma$ and $\psi_i = \alpha_d - \alpha_p/\gamma$. Finally, we combine these three equations with (7) and (9). This leads to (6), therefore concluding the proof. ■




Fig. 3. Two executions of the proposed control architecture when applied to system (1). Two different choices of control gains ψ_i and ψ_s are considered. The other parameters are $\gamma = 0.1, \beta_M = 0.22, \bar{\iota}(0) = 0.1, \iota_{\text{th}} = 0.12, \iota(0) = 0.14$. Susceptibles are not shown for the sake of space.

We want our control action to remain limited when acting on a neighborhood of $s\iota = 0$. Also, it is not meaningful to act on the system by changing β to values smaller than the one associated with total lockdown $\beta_m > 0$, or greater than the one representing no social distancing $\beta_M > \beta_m$. We therefore introduce the following modification on the ideal controller

$$\beta(s, \iota, t) = \left[\psi_i(\bar{\iota} - \iota) - \psi_s(\bar{s} - s) + \frac{\bar{s}\bar{\iota}}{[s\iota]_\epsilon^\infty} \bar{\beta} \right]_{\beta_m}^{\beta_M}, \quad (10)$$

where $\epsilon > 0$ is a small constant, and $[a]_l^u$ is equal to l or u if $a < l$ or $a > u$ respectively, and equal to a otherwise. Fig. 3 reports two examples of application of the algorithm to the SIR model (1).

IV. NETWORK CONTROL

A. Network Model

We implement two important features in a refined model: (i) people interact through heterogeneous contact structures, i.e., the population is not well-mixed, and (ii) real epidemics have an intrinsic degree of stochasticity, so they cannot be exactly described by (1). We therefore consider stochastic epidemics on networks [1], [16]. A network is a pair (V, \mathcal{E}) , where V is a set of N nodes (or vertices), and \mathcal{E} is a set of edges (or links) connecting nodes, i.e., tuples $\{u, v\}$, where $u, v \in V$. A population contact structure is modelled by a network in which nodes represent individuals, and links are associated with routes of disease transmission between individuals. We consider undirected networks, such that $\{u, v\} \in \mathcal{E} \iff \{v, u\} \in \mathcal{E}$. Figs. 2, 4, 5, show pictorial representations of networks. Here, we focus on a particular well-known class of random networks, i.e., Erdős-Rényi [18], generated as follows: start with N isolated nodes, consider each unique pair of two distinct nodes and connect them with probability $0 \leq p \leq 1$. Hence, the probability of a node having k neighbors follows a binomial distribution $\mathcal{B}(N-1, p)$, $\langle k \rangle = p(N-1)$ being the average degree. Such networks may be considered a very first order approximation of realistic contact structures, as they display sufficient heterogeneity and are easy to implement [1].

B. Epidemic Model on Network

We consider a SEIRD model for disease spreading, in which nodes are divided into compartments representing their status




Fig. 4. Pictorial representation of SEIRD dynamics on a network. The process is a continuous-time Markov chain. Each infected (and infectious) node spreads the disease to its susceptible neighbors at rate β_n until no longer infectious. A node that has been successfully infected, becomes first exposed, then infectious itself. Its ultimate destiny is either dying (with probability p_D), or fully recovering (with probability $1 - p_D$). The rate of each event is given on the continuous arrows.

with respect to the disease: S (susceptible), E (exposed), I (infected/infectious), R (recovered), or D (deceased). **Fig. 4** illustrates the possible transitions of a susceptible node that is in contact with two infectious neighbors. Compared to a SIR model (see Section II), we add an exposed class to account for individuals who have been infected but are not yet infectious (biologically known as incubation phase). We also allow for infected individuals to either survive or die. Outbreaks are modeled as Markovian processes on the generated network, in which an infected node spreads the disease, via links, to its susceptible neighbors at a constant rate β_n , turning them into exposed. At a constant rate γ_E , an exposed node becomes infected, and stops infecting at a constant rate γ_I , after which two outcomes are possible: either full recovery (transition into R), with probability $1 - p_D$, or death with probability p_D (transition into D). Nodes in compartments R and D play no further role in the dynamics. Further, p_D depends on the prevalence of the disease, to model increased mortality in case of saturation of the health care system. Control interventions in this model are implemented as changes in the value of β_n . At time $t = 0$, $I(0) = N_I(0) \ll N$ randomly chosen nodes are infected. The remaining ones are initialized as susceptible. We use a Gillespie algorithm [19] adapted to networks [1] to simulate this process. In **Fig. 5** we show a realization of an outbreak on a network of modest size, to highlight how the topology impacts the dynamics.

C. Input and Output Maps

To connect the controller to the network model, we introduce two maps, as shown in **Fig. 2**. Such mappings are general, and they could be used in conjunction with different control techniques relying on similar input-output description of the pandemic. The output map extracts s and i from the full state of the network by counting as s the fraction of nodes either S or E , and as i the fraction of I . The input map provides expressions for the control input on the network level β_n given the output of the controller $\beta(s, i, t)$. With the aim of evaluating the input map, we turn to the adaptation of β to networks. From (1) we get

$$N\dot{i} = \beta N i s - \gamma N i \Rightarrow \dot{i} = \beta I \frac{S}{N} - \gamma I. \quad (11)$$

The term $\beta I S/N$ represents the total infectious pressure in the ODE model. This quantity drives the whole infectious process,

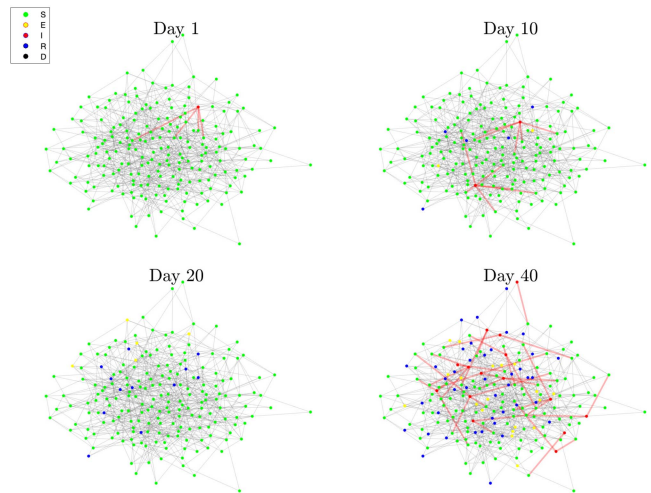


Fig. 5. Simulation of a SEIRD outbreak on a Erdős-Rényi network of size 200, with average degree $E[k] = 7$. A single node at day 1 spreads infection to its neighbors (red edges), which in turn become first exposed, then infected, and eventually recover or die. The network is drawn in such a way that nodes with fewer links are on the periphery. The effect of the topology on the disease is particularly evident on such nodes, as only a few of them gets infected compared to central ones.

and it is crucial that the map preserves it. On the network, the infectious pressure is given by the infectious pressure β_n times the number of links between infected and susceptible nodes, which is a random variable that depends on which nodes are infected/recovered and on the topology of the network. Therefore, implementing an exact mapping would require to impose a different SD level on each individual, depending on the degree of its social interactions. Although well defined in theory, this is clearly not implementable in practice. To overcome this issue, we introduce the so-called mean-field approximation [1], [16]. On average, an infected node is connected to $\langle k \rangle$ neighbors, of which we assume that a proportion S/N is susceptible - where $\langle \cdot \rangle$ is the expected value operator. Hence, we set the number of $S-I$ links as $\langle k \rangle I S/N$. We derive β_n as a simple linear function of β

$$\beta_n I(t) \langle k \rangle \frac{S(t)}{N} \simeq \frac{\beta}{N} I(t) S(t) \Rightarrow \beta_n \simeq \frac{\beta}{\langle k \rangle}. \quad (12)$$

This is a valid first-order approximation, that is known to give an upper estimate of the true $S-I$ link count (see [1], [16]), which in our case can only translate in a more conservative control strategy. This expression connects a SIR model (1) to a stochastic SIR on networks, rather than a stochastic SEIRD, as we want. Hence, we need to add an additional layer that conciliates γ_E and γ_I with γ in the SIR model. To do so, we first consider the time to full recovery (or death) of an individual who has been infected in a SEIRD model. This is a random variable exponentially distributed with rate $\gamma_E \gamma_I / (\gamma_I + \gamma_E)$. We set the controller γ to this value. To find the infection rate, we use the definition of R_0 [1] for both models, i.e., $R_0 = \beta/\gamma$ for the SIR, and $R_0 = \tilde{\beta}/\gamma_I$ for SEIRD (we momentarily use $\tilde{\beta}$ to distinguish it from the β in the SIR), and we impose that they are equal. This yields $\tilde{\beta} = \beta(\gamma_I + \gamma_E)/\gamma_E$. Finally, combining this expression with (12), gives β_n as

$$\beta_n = \frac{\tilde{\beta}}{\langle k \rangle} = \frac{\beta}{\gamma_E} \frac{\gamma_I + \gamma_E}{\langle k \rangle}. \quad (13)$$

TABLE I
PARAMETERS USED FOR SIMULATIONS IN SECTION V

β_{\max}	0.0227	β_{\min}	0.0057
γ_E	0.25	delay (days)	$\mathcal{N}(\{3, 7, 20\}, 1)$
γ_I	0.1428	noise (signal)	$\mathcal{N}(0, 0.1)$
N	16000	Hospitalization rate	0.02
T_f (days)	240	ν_{th} (%)	0.025
I_0	800	p_D if $\nu \leq \nu_{th}$	0.005
S_0	15200	p_D if $\nu > \nu_{th}$	0.02
$\langle k \rangle$	19	policy update (days)	$\{1, 7, 15\}$

V. SIMULATIONS


On top of the complexity introduced by the network dynamics, we consider several non-ideal behaviors to better approximate a real-world scenario. Note that none of these effects are considered in the controller design, and therefore are to be seen as uncertainties.

- Unknown random delay affects measurements, which changes every time the controller is executed. This models the difficulties in getting on-line estimates of prevalence through daily swab tests.
- Policy update is allowed at a fixed rate, to mimic real life scenarios in which policy makers are reluctant to apply different degrees of restrictions too frequently.
- Quantization of the possible levels of the network control level β_n , based on Italian mobility data [20]. Policy makers can realistically implement only limited control actions. We use 5 distinct, equally spaced, levels, from $\beta_{\min} > 0$ to β_{\max} . We set $\beta_{\min} = 0.25\beta_{\max}$.
- We introduce measurement noise of the signal, proportional to its value, to model uncertainty in the estimation of the prevalence when the epidemic is out of control.


For the tuning of the model parameters we consider the case of Codogno, which has been the first city in Lombardy with a diagnosed case of Covid-19. We have used Google data for the number of people in Codogno and the hospital capacity. We considered realistic parameters for incubation period [21], infectious period [21], hospitalization rate [22], infection fatality rate [23], and social network connectivity [24]. All the parameters are reported in Table I. The initial condition is set to $I_0 = 800$, to model a delayed recognition of the presence of the disease, and simulations are run for $T_f = 240$ days.

Figs. 6 and 7 show the evolution of infected ι , deaths D/N , and prescribed SD $\beta(s, t, \iota)$, for the case where policy can change once a week and delay between testing and results is on average 4 days. We report the results when using the proposed feedback action $\beta(t, s, \iota)$, and, as comparison, the evolution of the uncontrolled epidemics ($\beta_n = \beta_{\max}$) and of a one on-off intervention lasting for 60 days, during which $\beta_n = \beta_{\min}$. Susceptible percentages s are not shown for the sake of space. We aggregate results from 100 simulations, each one run on a different network realization.

We evaluate the performances of the controller in various settings, in which we act on two main parameters, namely, the delay in knowledge of the status and the frequency at which the control policy can be changed. The former one can take values of $\{3, 7, 20\}$ days, while the latter moves between $\{1, 7, 15\}$ days. We consider all the possible combinations of these parameters. We cannot report here the complete results of our simulations, for the sake of space. We report instead some relevant performance indexes in Fig. 8 - namely the reduction in social distancing compared to 60 days full lockdown, and reduction in deaths with respect to not applying any strategy. We observe that the controller performs well



(a) Infected ι



(b) Deaths D/N

Fig. 6. Prevalence of infected and dead nodes for the considered simulation scenario. It is shown here the case in which the policy changes only once every week, and the average delay in measurements is set to 3 days. All the other values are as in Tab. I.




Fig. 7. Level of SD $\beta(t, s, \iota)$ as a function of time. The average output of the controller across 100 simulations, when policy changes every 7 days and delays in data are 3 days, is shown together with its lower and upper quartiles (Q1-Q3). We also report for comparison a 60 days full-lockdown strategy, and the feedforward action $\bar{\beta}$.

on average even in the most extreme cases. Yet, we observe increased dispersion as we increase delays and reaction times. The use of the controller consistently induces a reduction of over 32% of deaths in the worst case, and, in the best tested case, of 63%.

VI. DISCUSSION

Our approach resulted in a strategy able to keep the curve below the health care capacity when uncertainty is low, with increased variability when delays and other inaccuracies in measuring become important. From this analysis, it appears clear that is crucial to have a reliable estimate of the current prevalence of the disease. This is of course the downside of closed loop strategies, i.e., that the controller becomes less reliable as the quality of measurements deteriorates. Instead, it is

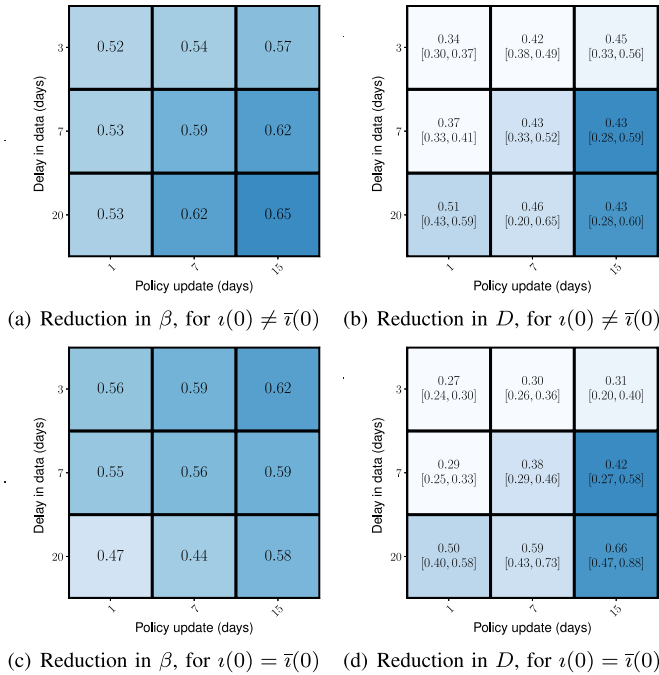


Fig. 8. Heat maps reporting (a) the average reduction in β_n , normalised by $\int_0^{T_f} \beta_n(t) dt$ in the reference scenario, and (b) the average reduction in deaths, normalised by the average number of deaths in the reference scenario, across different conditions. Colors in (b) follow the width of the first and third percentile (reported in the cells under the average). Both the indices are defined so that the smaller the better.

worth noting that, given low delays in data, updating policies every 15 days has a limited impact on the performance of the controller. Interestingly, increasing delays (or control updates frequency), does not have a major impact on the average performance of simulations, in terms of reduction of mortality. However, this result might be misleading, because the variance between different realisations gets higher as the delay increases, meaning that the controller becomes unreliable if applied to an individual realization. This suggests that the crucial quantity for control is on-line prevalence estimation. Therefore, this analysis confirms that, when implementing control policies based on daily testing data, policy makers should ideally have access to the exact state of the system. Clearly, this is far from being a realistic assumption. Still, our results demonstrate that periodic loop closure might be a viable solution also in a more realistic scenario.

At the same time, we observe a relevant outcome in all our simulations, namely that when control acts on an outbreak that has already reached a significant proportion of the population, the advisable strategy is to go into full lockdown until the epidemic curve is brought down to acceptable levels, and then to gradually relax and adjust control measures, according to the estimated prevalence.

VII. CONCLUSION AND FUTURE WORK

This preliminary work showed that a simple feedback action can improve the robustness and the effectiveness of an optimal policy for epidemic control, even in presence of quite non ideal behaviors in the system and in measurements. The effectiveness of strategies based on control for dealing with epidemics is still an open topic, with respected academics having opposite positions [7], [25]. Although our results are far from being readily applicable, they provide a new piece to this intricate puzzle. Future work will be devoted to use more reliable

input maps (and possibly theoretical models for the controller), improve control design with robust and adaptive techniques, include other sources of lags and uncertainties, use more realistic network models - possibly dynamic networks, the ultimate goal being engineering a sound model that could be useful when it comes to decision making for governments.

REFERENCES

- [1] I. Z. Kiss, J. C. Miller, and P. L. Simon, *Mathematics of Epidemics on Networks: From Exact to Approximate Models*. Heidelberg, Germany: Springer, 2017.
- [2] A. Goolsbee and C. Syverson, *Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020*, Nat. Bureau Econ. Res., Cambridge, MA, USA, 2020.
- [3] A. K. M. I. Bhuiyan, N. Sakib, A. H. Pakpour, M. D. Griffiths, and M. A. Mamun, "COVID-19-related suicides in bangladesh due to lockdown and economic factors: Case study evidence from media reports," *Int. J. Mental Health Addiction*, pp. 1–6, May 2020, doi: [10.1007/s11469-020-00307-y](https://doi.org/10.1007/s11469-020-00307-y).
- [4] S. Xu and Y. Li, "Beware of the second wave of COVID-19," *Lancet*, vol. 395, no. 10233, pp. 1321–1322, 2020.
- [5] R. A. Armstrong, A. D. Kane, and T. M. Cook "Outcomes from intensive care in patients with COVID-19: A systematic review and meta-analysis of observational studies," *Anaesthesia*, vol. 75, no. 10, pp. 1340–1349, 2020.
- [6] L. Thunstrom *et al.*, "The benefits and costs of flattening the curve for COVID-19," in *Proc. SSRN*, 2020, Art. no. 3561934.
- [7] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Analysis and control of epidemics: A survey of spreading processes on complex networks," *IEEE Control Syst. Mag.*, vol. 36, no. 1, pp. 26–46, Feb. 2016.
- [8] F. Di Lauro, I. Z. Kiss, and J. C. Miller, *Optimal Timing of One-Shot Interventions for Epidemic Control*. Cold Spring Harbor, NY, USA: Cold Spring Harbor Lab. Press, 2020, doi: [10.1101/2020.03.02.20030007](https://doi.org/10.1101/2020.03.02.20030007).
- [9] D. Demasse *et al.*, *Optimal COVID-19 Epidemic Control Until Vaccine Deployment*. Cold Spring Harbor, NY, USA: Cold Spring Harbor Lab. Press, 2020, doi: [10.1101/2020.04.02.20049189](https://doi.org/10.1101/2020.04.02.20049189).
- [10] F. Di Lauro *et al.*, "The impact of network properties and mixing on control measures and disease-induced herd immunity in epidemic models: A mean-field model perspective," 2020. [Online]. Available: [arXiv:2007.06975](https://arxiv.org/abs/2007.06975).
- [11] G. Giordano *et al.*, "Modelling the COVID-19 epidemic and implementation of population-wide interventions in italy," *Nat. Med.*, vol. 26, pp. 855–860, Apr. 2020.
- [12] M. Bin *et al.*, "On fast multi-shot epidemic interventions for post lockdown mitigation: Implications for simple COVID-19 models," 2020. [Online]. Available: [arXiv:2003.09930](https://arxiv.org/abs/2003.09930).
- [13] J. Köhler *et al.*, "Robust and optimal predictive control of the COVID-19 outbreak," 2020. [Online]. Available: [arXiv:2005.03580](https://arxiv.org/abs/2005.03580).
- [14] D. H. Morris *et al.*, "Optimal, near-optimal, and robust epidemic control," 2020. [Online]. Available: [arXiv:2004.02209](https://arxiv.org/abs/2004.02209).
- [15] A. Charpentier *et al.*, "COVID-19 pandemic control: balancing detection policy and lockdown intervention under ICU sustainability," 2020. [Online]. Available: [arXiv:2005.06526](https://arxiv.org/abs/2005.06526).
- [16] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, "Epidemic processes in complex networks," *Rev. Mod. Phys.*, vol. 87, no. 3, p. 925, 2015.
- [17] R. M. Corless *et al.*, "On the lambert W function," *Adv. Comput. Math.*, vol. 5, no. 1, pp. 329–359, 1996.
- [18] B. Bollobas, *Random Graphs* (Cambridge Studies in Advanced Mathematics), 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2001.
- [19] D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *J. Phys. Chem.*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [20] E. Pepe *et al.*, "COVID-19 outbreak response, a dataset to assess mobility changes in italy following national lockdown," *Sci. Data*, vol. 7, no. 1, pp. 1–7, 2020.
- [21] Y. Ling *et al.*, "Persistence and clearance of viral RNA in 2019 novel coronavirus disease rehabilitation patients," *Chin. Med. J.*, vol. 133, pp. 1039–1043, May 2020.
- [22] World Health Organisation. (2020). *Media Statement: Knowing the Risks for COVID-19*. [Online]. Available: <https://www.who.int/indonesia/news/detail/08-03-2020-knowing-the-risk-for-covid-19>
- [23] G. Meyerowitz-Katz and L. Merone, "A systematic review and meta-analysis of published research data on COVID-19 infection-fatality rates," *Int. J. Infectious Diseases*, vol. 101, pp. 138–148, Dec. 2020.
- [24] A. Melegaro, M. Jit, N. Gay, E. Zaghenni, and W. J. Edmunds, "What types of contacts are important for the spread of infections? Using contact survey data to explore European mixing patterns," *Epidemics*, vol. 3, no. 3, pp. 143–151, 2011.
- [25] F. Casella, "Can the COVID-19 epidemic be controlled on the basis of daily test reports?," *IEEE Control Syst. Lett.*, vol. 5, no. 3, pp. 1079–1084, Jul. 2021.