Can the COVID-19 epidemic be controlled on the basis of daily test reports?

Francesco Casella, *Member, IEEE*

***Abstract*—Short answer: not much, and only with an overly cautious approach. The paper presents a suitable mathematical model of the process for feedback control analysis and uses well-known results from control theory to prove that suppression strategies based on daily test reports can be effective if enacted very early, while mitigation strategies, including trying to achieve herd immunity, are likely to fail.**

***Index Terms*—COVID-19, control theory, limitations of control.**

I. Introduction

THE first outbreak of the COVID-19 [1] virus epidemic took place in China, starting in December 2019, possibly even earlier. The outbreak has since then spread to Western countries, sparking a debate on what are the most appropriate measures to control it.

The unique features of COVID-19 make this task particularly challenging. The virus is new, so there is no previous immunity in the general population, nor there are effective vaccines or cures, except life support in the most critical cases.

On one hand, the mortality rate, although still uncertain, is much lower than that of previous virus outbreaks such as the MERS and SARS. Many infected people show little or no symptoms, but are contagious and spread the virus to other people that can develop much more serious symptoms.

On the other hand, a significant number of affected subjects eventually develop severe bilateral pneumonia, which requires hospitalization, oxygen ventilation, and intensive or sub-intensive care, and that can lead to respiratory failure and death if not treated properly. Such treatment is effective for a majority of patients, particularly those less than 70 years old, but has the potential to overwhelm and disrupt the public health systems even as a modest fraction of the population is affected at any point in time. Such a disruption could leave most patients with severe pneumonia without crucial life support, thus increasing the fatality ratio of COVID-19 by several times, and must thus be avoided at all costs to avoid dire consequences [2].

Two approaches have been advocated to deal with the outbreak. The first is *suppression*: rigorous, sometimes outright draconian, social distancing measures are taken by national governments, such as closing schools and public places, issuing stay-at-home orders, closing non-essential industrial and commercial activities, banning all kind of non-essential travel, etc. The goal in this case is to reduce the reproduction number R_t, which is the number of persons an infectious person infects on average, below unity, causing the outbreak to subside. This approach was followed very thoroughly by China, where the original outbreak developed [1], effectively suppressing the epidemic in a couple of months time. This approach was subsequently followed, albeit with much less draconian restrictions, by most Western European countries, in particular by Italy, who was hit first and hardest. Some countries, most notably South Korea, adopted a suppression strategy based on extensive tracking and early isolation of infected people, in order to achieve the same effects of general lock-down in a less socially disruptive way.

Although suppression has proven to be very effective, it leaves open the question of what strategy to adopt once the epidemic has been tamed, since it leaves most of the population still vulnerable to the virus and thus prone to a second, and possibly even more, wave(s) of disease outbreak.

The second approach is *mitigation*: the idea is to let the epidemic run its course in a controlled way, eventually leading to herd immunity, while at the same time ensuring that the capacity of the public health system is not overcome. This approach was initially spearheaded by the UK government [3], which later changed the strategy to suppression after the public release of report [2], which predicted 250,000 casualties in the UK and 1.1-1.2 million casualties in the USA, if the mitigation approach was adopted.

The goal of this paper is to provide fundamental insight on this problem by adopting a control-theoretical analysis framework. To this aim, a suitably simplified dynamic model of the process of epidemic evolution is formulated, which is backed by experimental evidence, alongside with a representation of the decision-making process which is amenable to analysis by means of basic control theory.

The main result of this analysis is twofold. On the one hand, *suppression* strategies can be successful, if enacted promptly, much earlier than one would do by looking at the latest daily reports of reported cases and deaths, and with drastic enough measures. On the other hand, *mitigation* strategies are likely to fail, due to the combination of fast unstable dynamics, time delays in measurements, and process uncertainty, and may possibly be used as a last resort option only if special care is taken to reduce those adverse features as much as possible.

The paper is structured as follows: Section II introduces a control-oriented model of the epidemic during the time interval when the above-mentioned policies are taken, based on available daily data regarding the number of new cases. In Section III, the two above-mentioned strategies are analysed in terms of feedback control. Section IV draws conclusions from

F. Casella is with Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133 Milano, Italy e-mail: francesco.casella@polimi.it

A first version of this manuscript was published on arXiv on Mar 16, 2020, https://arxiv.org/abs/2003.06967.

the control-theoretical analysis with some recommendations for decision makers and future research.

II. Modelling

A. Derivation of the model

The mathematical theory of epidemics was started almost 100 years ago by W.O. Kermack and A. G. MacKendrick in their seminal paper [4], that introduced the SIR compartmental model. Since then, extensions were developed in various directions: by increasing the number of compartments, by accounting for more phenomena to the model, by distinguishing among different age cohorts, by adding spatial structure and explicit modelling of contagion paths, and by considering stochastic behaviour, see e.g. [5] for a comprehensive review.

Such models, see e.g. [2], [6], can be extremely detailed and sophisticated, but all share a common feature: they are *not* first-principle models in the sense of mechanical or electrical system models, that are based on very precise and highly repeatable physical laws, but rather crucially depend on empirical coefficients that need to be tuned *a posteriori* on relevant historical data. Their *a priori* predictive power is thus inherently limited when dealing with COVID-19, for which there is no prior experience, there are no effective drugs or vaccines, and only social distancing methods are available to limit the spread of the disease and its consequences.

This is clearly shown by the results of the study [7], which tries to estimate the effect of various types of government interventions onto the relative reduction of the reproduction number R_t, by applying Bayesian methods to data from 11 European countries. The initial reproduction number R_0 is estimated to be well above 2, possibly even about 4 in some countries. Hence, a reduction of the reproduction number by at least 60-70%, possibly more, is necessary to avoid the exponential growth of cases.

The main result of that study (see [7], Fig. 4) is that lockdown leads to an average reduction of R_t by 50%, school closure by 20%, other measures around 10%. However, 95% confidence intervals on the reduction factors are huge, e.g. 10% to 80% reduction for lockdown, 0% to 45% reduction for school closure, fundamentally undermining their usefulness for predictive models. This problem is inherent to the requirement of a large enough data set to be statistically significant, which requires to put countries with very different social habits and very different interpretations of the same measure (e.g. lockdown) in the same data set, inevitably leading to very large uncertainty in the estimation of their effects.

Furthermore, the quality and homogeneity of data used to tune those models are often highly questionable: different countries adopt different standards to determine who gets tested and when, and possibly change them over time; some data get lost or are reported with extra delay because of clerical mismanagement; some countries or regions may report lower numbers than real because of political pressures. Even bona-fide reporting systems may be completely overwhelmed and fail to provide reliable data. For example, the number of additional deaths on municipal records during March 2020 with respect to March 2019 in some areas of Lombardy, Italy, are two to four times larger than officially reported COVID-19-related deaths.

It is then a compelling evidence that any public policy based on such models cannot be applied blindly, but must be adapted and corrected based on the observed outcome, which currently means the daily reports of individuals that are reported positive to swab tests. The crucial question is then: *is feedback control feasible at all in such a system*?

In order to answer this question, a suitably simplified model of the epidemic is used. This model is not accurate enough to perform open-loop predictions, but is good enough to capture the fundamental dynamics that is relevant for the success (or failure) of the feedback policy. This kind of simplification is a well-established common practice in control engineering.

The basic SIR model can be formulated as follows:

$$\frac{dS}{dt} = -\frac{\beta I S}{N} \tag{1}$$

$$\frac{dI}{dt} = \frac{\beta I S}{N} - \gamma I \tag{2}$$

$$\frac{dR}{dt} = \gamma I \tag{3}$$

where N is the total population, S is the number of people Susceptible to infection, I is the number of Infectious people, that can transmit the disease to others, and R is the number of Resistant people, which are immune to the disease because of genetics, vaccination, or immunity acquired after contracting the disease. β and γ are (possibly time-varying) parameters that account for the likelihood of susceptible people to get infected by infectious people, and of infectious people to eventually recover and become Resistant, per unit time. Given the short time spans involved (a few weeks) and the relatively low mortality rate (a few percent at most), deaths and births are neglected. Immigration and emigration are also neglected for simplicity. In this case, adding the three equations leads to

$$\frac{d(S+I+R)}{dt} = 0. \tag{4}$$

Hence, the last equation can be replaced by an algebraic equation, making the model effectively second-order:

$$R(t) = N - S(t) - I(t). \tag{5}$$

The COVID-19 outbreak in Europe has several specific features, that are relevant for the process modelling.

1) COVID-19 is a new virus, so the vast majority of the world population has never been exposed to it yet.
2) No vaccine is available or expected in the near future.
3) No really effective cure has been found yet.
4) The ratio of deceased people over positive tested subjects is strongly country-dependent, ranging from about 2% (Germany), to 4% (China) to about 10% (Italy, France, Spain).
5) The actual mortality ratio with respect to infected people is expected to be much lower, since many subjects show no or mild symptoms, do not report to hospitals, and are thus not tested for the virus, but are still reportedly contagious. The ratio α between positive tested subjects and actually infectious subjects is uncertain, strongly dependent on country-specific testing policies, probably

around one order of magnitude. In the case of the Hubei province outbreak it was estimated that $\alpha = 0.05$ [8].

6) A certain fraction of officially reported infectious cases (about 30% in Northern Italy) ends up with severe bilateral pneumonia and respiratory difficulties, that require hospitalization and may require non-invasive CPAP respirators. A certain fraction σ, about 4% in Northern Italy, eventually requires mechanical ventilation and intensive care for a period from a few days to a couple of weeks to sustain the patient's life functions. Absent the possibility of adequate life support, this fraction of patients is more likely to die because of respiratory failure, increasing the mortality rate dramatically.
7) The number of beds, equipment, and personnel which is required to perform artificial ventilation on patients in public health systems is based on normal needs of post-surgery care, trauma care, and care of people with rare but serious diseases like amyotrophic lateral sclerosis, with relatively small extra capacity. The number N_{ic} of COVID-19 patients that can be admitted to intensive and sub-intensive care is thus severely limited. While the actual numbers vary widely by country, the typical order of magnitude is $10^{-4}N$. During the very short time span of the initial phases of the outbreak, this number can be significantly expanded if timely action is taken, but certainly not by orders of magnitude.
8) The initial dynamics of the outbreak is very fast, with doubling times of reported cases of the order of 4 days.
9) The moment an infected patient becomes infectious is still not fully clear. There are some indications that patients, who later report severe symptoms and get tested because of that, started becoming infectious two days before the onset of fever. Furthermore, in most countries people only get tested when they develop serious symptoms, which can delay the moment they are tested by many days. There is thus a time delay τ_t of several days on average between the moment people become infectious, and the moment they get tested. This delay is critical for decision and control.
10) The testing process also introduces a delay in the process. Although in principle it is possible to provide the results of the test in a few hours, the average time to report the results is normally much longer, because of logistic constraints and because of limited availability of analytical equipment. For example, the average time required to obtain the test results in Lombardy during the second week of March 2020 was about one week. This delay τ_r is also critical for decision and control.

Assuming a worst-case scenario, which is required by the precautionary principle given the number of lives at stake, Item 1 suggests to consider $S(0) = N$. The absence of a vaccine (Item 2) implies there is no means to reduce the value of S and increase the value of R by means of vaccination campaigns. Item 3 allows us to consider γ as a constant, at least as a first approximation.

Items 5-6, coupled to Items 1-3, are crucial from the modelling point of view. When the health-care system capacity limit is reached, standard recommended triage practice is to give priority access to intensive care to younger and healthier people, which are likely to recover more quickly (freeing up the scarce resources for other patients), and who are likely to live longer once recovered, while denying access to elderly or otherwise frail individuals, see e.g. the recommendations in [9]. It is the opinion of the author that enacting a public policy with a significant risk of causing this outcome *a priori* is morally unacceptable. Hence, any acceptable control policy should ensure *a priori* that $\alpha\sigma I < N_{ic}$ at all times. With the previously mentioned values, this implies

$$I < \frac{N_{ic}}{\alpha\sigma} \approx 0.05N \tag{6}$$

It is then possible to assume *a priori* that $I \ll S$. Furthermore, since the average time to heal from COVID-19 is between two and six weeks, it is possible to assume that during the first one-two months of the outbreak, the number R of people who recover will be even smaller, so $I + R \ll S$. Since $I + R + S = N$, one can assume in Eq. (2) that S is nearly constant, and approximately equal to N. This assumption decouples Eq. (2) from Eq. (1), leading to

$$\frac{dI}{dt} = (\beta - \gamma)I. \tag{7}$$

Assuming that also β is constant, which is reasonable as long as no significant social distancing measures are taken by the authorities, Eq. (7) has an exponential solution

$$I(t) = I(0)e^{(\beta-\gamma)t} = I(0)e^{rt}, \tag{8}$$

where $r = \beta - \gamma$. The doubling time T_d of the infectious cases is given by

$$T_d = \frac{log(2)}{r}. \tag{9}$$

Note that the initial reproduction number is $R_0 = \beta/\gamma$, estimated between 2 and 4 for COVID-19 [7].

In fact, Eq. (8) refers to the number of *actual* infectious cases, which is largely unknown, see Item 4. However the empirical ratio σ of patients requiring artificial ventilation mentioned in Item 5 is referred to the much lower number $I_t(t)$ of cases that will eventually get *tested positive* to the virus. Given that $I_t(t) = \alpha I(t)$, assuming the ratio α to be constant, it is trivial to prove that also the dynamics of *tested positive* cases I_t obeys the same differential equation

$$\frac{dI_t}{dt} = (\beta - \gamma)I_t, \tag{10}$$

whose solution is

$$I_t(t) = I_t(0)e^{(\beta-\gamma)t} = I_t(0)e^{rt}. \tag{11}$$

The government interventions mentioned earlier (lockdown, school closures, etc.) have the effect of changing the rate of infection β, hence the actual reproduction number $R_t = \beta/\gamma$. These measures are varied and can be applied progressively.

We can then assume that the time-varying parameter β is in fact a function of a representative manipulated variable $u(t)$, where the value of u indicates the intensity of adopted public health measures on a scale from 0 (no intervention) to 1 (full lockdown and isolation of all individuals). The $\beta(\cdot)$

function is thus monotonously decreasing from the value β_0, observed during the initial outbreak when no social restrictions are enforced, to zero, corresponding to a complete lockdown. Note that the latter situation is a bit unrealistic, since it would require people to also isolate themselves within their households.

If the effect of each public measure in terms of reduction of β (or, equivalently, of R_t) was known exactly, the function $\beta(\cdot)$ would be also known. However, considerable uncertainty is involved in the estimation of the effects of different interventions [7], which then affects the knowledge of $\beta(\cdot)$.

The measured variable $I_r(t)$ of the process is the number of *reported* infected cases. As mentioned in items 8-9, the overall measurement process of reported cases is subject to an average delay of τ_t days between the onset of infectiousness and the moment the swab test is taken, and by another delay of τ_r days before the result of testing becomes available to public authorities.

The control-oriented model of the virus outbreak dynamics is thus the following:

$$\frac{dI_t(t)}{dt} = \left[\beta(u(t)) - \gamma\right] I_t(t) \tag{12}$$

$$I_r(t) = I_t(t - \tau_m), \tag{13}$$

where $\beta(u)$ is an uncertain function, γ is an uncertain constant parameter, τ_t, τ_r are uncertain, possibly time-varying parameters, and $\tau_m = \tau_t + \tau_r$ is the overall measurement delay.

B. Validation and Tuning

The first validation of the model is based on data of the Hubei province outbreak in China [1]. The outbreak initially ran unchecked, until a very strict lockdown (strict stay-at-home order, one person per building allowed to shop for food) was imposed on Jan 23, 2020 in the city of Wuhan, followed by other 15 cities of the province on the next day.

Assuming $R_0 = 3$, the values $\beta(0) = 0.26$ days^{-1} and $\gamma = 0.0867$ days^{-1} give the best fit of the model (red curve) to the data, corresponding to $r = 0.173$ days^{-1} and $T_d = 4.0$ days, see Fig. 1. The lockdown was applied on days 5 and 6, marked by the arrows in the figure, and caused very clearly a delayed effect. Assuming the reduction of β was equally split on those two days, the best fit to the second part of the transient is obtained by assuming an overall 89% reduction of the initial value β_0, and an overall delay $\tau_m = 11$ days.

The second validation case is based on data from the Lazio region in Italy, including Rome, reported by the Italian Civil Protection Department [10], for the period Feb 24 through Mar 29, 2020. This dataset was picked because no restrictions of any sort were applied until March 5 in that region, while Northern Italy, were the outbreak started and made most damage, many subsequent restrictions were applied in sequence, making it difficult to identify their individual effect; this also makes the analysis of the overall numbers from Italy conceptually questionable, because it puts together completely different situations. Starting from Mar 5, 2020, schools, theatres and museums were closed and sports events cancelled in the whole country; from Mar 12, restaurants, bars,

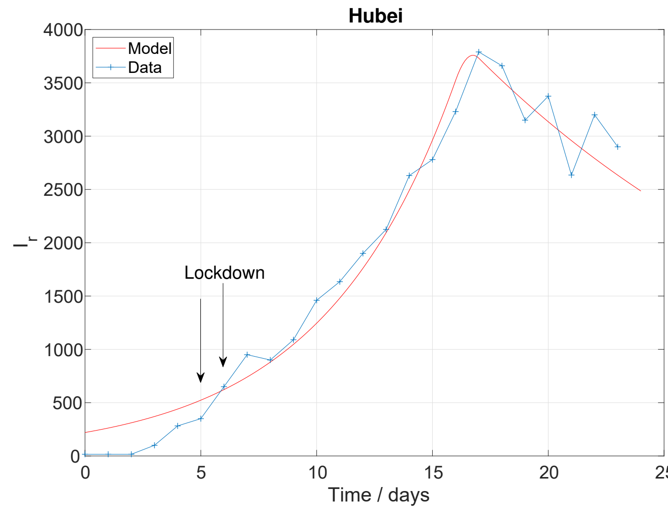

Fig. 1. Model validation: Hubei outbreak

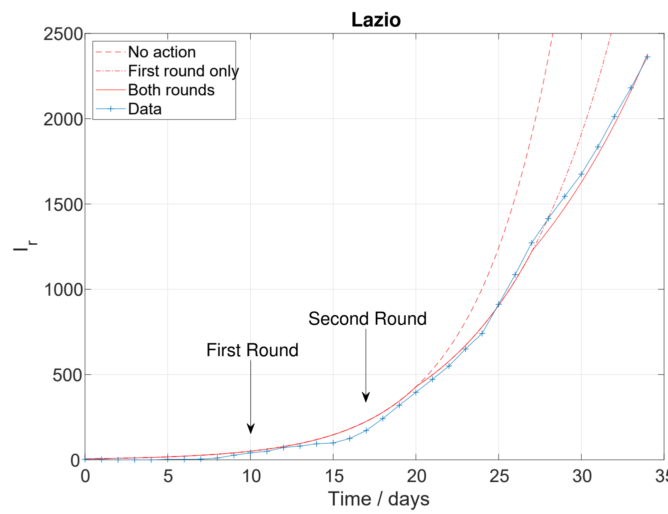

Fig. 2. Model validation: Lazio outbreak

and all commercial outlets except for food and medicines were also closed.

The best fit of the model, red curve of Fig. 2, is obtained by taking $\beta(0) = 0.315$ and $\gamma = 0.105$, corresponding to $r = 0.21$ days^{-1} and $T_d = 4.8$ days, by assuming a 20% reduction of β on day 10 (first round of measures), and a further 17% reduction of β on day 17 (second round of measures), with an overall delay $\tau_m = 10$ days. The delayed effect of the two rounds of public measures is clearly visible when comparing the data with models that do not consider the second one, or both (dashed curves).

III. Control

The effects of the application of the two control policies briefly outlined in the Introduction will now be analysed, based on the model derived in the previous Section.

The common theme behind both can be summarized by the title of the famous 2003 Bode Lecture paper by Gunter Stein: "Respect the Unstable" [11]. Feedback control strategies should not be applied light-heartedly to unstable systems, particularly when a large number of human lives is at stake, as the sobering memory of the Chernobyl disaster discussed in that paper suggests. In the case discussed in this paper, particularly deadly consequences can stem from the combination of unstable dynamics, time delay, and uncertainty.

A. Suppression

The suppression policy can be brutally summarized in the following terms: as soon as $I_r(t)$ reaches a value I_s which is scary enough to decision makers to overcome their reluctance to disrupt the social and economic life of their country, drastic containment measures are taken, which means the following discontinuous control law is applied:

$$u(t) = \begin{array}{ll} 0, & I_r(t) < I_s \\ \bar{u}, & I_r(t) \geq I_s \end{array} \tag{14}$$

As long as $\bar{u}$ is large enough that $r = \beta(\bar{u}) - \gamma < 0$, once the threshold I_s is exceeded the number of *tested positive* cases $I_t(t)$ will start decaying exponentially, until it will asymptotically reach zero in a time of about $-5/r$ days, corresponding to a steady cumulated value of infected people that will eventually heal or die.

The number of *reported* cases $I_r(t) = I_t(t - \tau_m)$ is instead delayed by τ_m days, so it will peak at a value MI_s after τ days, where the multiplier

$$M = exp(r\tau_m) \tag{15}$$

can have quite large values, e.g. $M = 6.7$ for the Hubei province and $M = 10$ for the Lazio region, assuming stronger measures were taken than those reported in the previous Section. A wise choice of I_s should ensure that the number of cases needing intensive care is always such that $\sigma I_t(t) < N_{ic}$, i.e., within the reach of the public health system, possibly with a good safety margin. This requires to pick $I_s < N_{ic}/\sigma M$. Political decision-makers without a training in mathematical modelling may have difficulties in understanding the role and magnitude of factor M and may be caught by surprise once it is too late to act.

The other open problem of this policy is that it is not clear if the equilibrium which is finally reached is sustainable once the containment measures are lifted, or if there is rather a risk that a second outbreak ensues. However, even in the case this unfortunate event takes place, this policy allows to buy precious time to improve the maximum capacity of intensive and sub-intensive health care units, by building or restoring new hospitals, by procuring ventilators and intensive care beds, and by hiring and training personnel to operate them. Given enough time, these measures can strongly mitigate the final death toll and avoid the need of taking wartime-like triage decisions in severely strained hospitals.

B. Mitigation

1) Policy statement: Proponents of the mitigation strategy for the COVID-19 outbreak argue that, in the absence of a vaccine, the majority of people should get infected and become immune one way or the other, until herd immunity is achieved. This requires $R > NR_0/(R_0 + 1)$, which is about 60-75% of the population given the current estimates for R_0. This outcome should be achieved by initially letting the epidemic run free, and then introducing the right measures at the right time to control the outbreak and ride through it as fast as possible, while avoiding an excessive number of critical cases, that could overwhelm the public health system.

2) Mathematical formalization: In mathematical terms, the first step to enact this strategy is to compute a reference trajectory $I^0_r(t)$ for the reported cases $I_r(t)$, and a corresponding optimal control policy $u^0(t)$, obtained by the application of public measures whose effect on β is accurately calibrated. This reference trajectory should guarantee to reach the herd immunity condition as fast as possible, while ensuring that $\sigma I_t(t) < N_{ic}$ at all times, meaning that the health system is never exceeding its capacity limit to provide intensive care to *all patients* who need it. Note that this requires to take the delays τ_r and τ_t into proper consideration.

Such a reference trajectory can be obtained by means of constrained nonlinear optimization, using sophisticated models of the COVID-19 epidemic evolution such as [2], that can be much more accurate than the simple SIR model presented in the previous Section.

It is quite obvious that the unstable nature of the state trajectories while $r > 0$, i.e. before the peak of the epidemic is reached, makes a completely open-loop implementation of this policy infeasible. The reference trajectory should then be followed by adapting the public policy measures $u(t)$ in real time, based on the observed values of the reported cases $I_r(t)$. In fact, every government pays extreme attention to the new daily reports of $I_r(t)$, and rightfully so. This corresponds in principle to closing a *feedback loop* to stabilize the unstable reference trajectory.

In the following sub-sections it will be shown that even a fairly sophisticated feedback control law cannot manage to stabilize the trajectory close to the reference, because of the adverse nature of the process dynamics identified in the previous Section. This casts very serious doubts on the feasibility of all mitigation approaches, which will in fact use a much less sophisticated feedback control strategies to achieve the same goal.

3) Trajectory controller design: The process model, linearized around the reference trajectory $u^0(t)$, $I^0_r(t)$, reads:

$$\frac{d\Delta I_t(t)}{dt} = \beta'(u^0(t))I^0_t(t)\Delta u(t) + \left[\beta(u^0(t)) - \gamma\right]\Delta I_t(t) \tag{16}$$

$$\Delta I_r(t) = \Delta I_t(t - \tau_m) \tag{17}$$

where β' indicates the derivative of β with respect to u.

By making the overly optimistic assumptions that the parameters γ, σ, and τ_m are constant and perfectly known, and that the function $\beta(u)$ that expresses the effects of public policy decision on the infection rate β is time-invariant, monotonously decreasing, smooth, and perfectly known, one could compensate the very strong nonlinearity of the process behaviour, by designing a linear controller with transfer function $C(s)$ with suitable gain scheduling, that will result in a linear and (almost) time-invariant loop dynamics, and add its output to the reference trajectory $u^0(t)$ to stabilize it.

One should also account for a further delay τ_c of one-two days within the controller, corresponding to the time which is necessary to collect the data, make decisions which are potentially disruptive for the social fabric and the economy, communicate them effectively to the public, and give the public time to actually implement them. The feedback control law can then be described as:

$$u(t) = u^0(t) + \frac{1}{\beta'(u^0(t-\tau_c))I_r^0(t-\tau_c)} u_f(t-\tau_c) \quad (18)$$

$$u_f(s) = C(s)I_r(s), \quad (19)$$

The loop transfer function of this system reads:

$$L(s) = C(s)\frac{\mu}{1-(\beta(u^0(t))-\gamma)s}e^{-s(\tau_m+\tau_c)}. \quad (20)$$

where μ is the ratio between the *actual* (unknown) value of $\beta'(u^0(t))I_t^0(t)$ and its *reference* value. In ideal conditions, $\mu = 1$, though results from [7] suggest this parameter is affected by very significant uncertainty.

This notation is slightly improper, since $\beta(u^0(t))$ is in general time-varying for non-trivial reference trajectories. For the sake of the subsequent analysis, we assume that this parameter changes over a time scale which is much longer than the time scale of the closed-loop system feedback response, a common assumption when dealing with gain-scheduling control, and thus consider it as a constant.

The loop transfer function reveals the very dangerous nature of this process, which has an unstable pole with time constant T, a time delay τ, and a highly uncertain gain μ:

$$T = \frac{1}{\beta(u^0)-\gamma} = \frac{1}{r} = \frac{T_d}{log(2)} \quad (21)$$

$$\tau = \tau_t + \tau_r + \tau_c. \quad (22)$$

The observed order of magnitude of those parameters at the beginning of the outbreak is $T = 4 \div 7$ days and $\tau = 12 \div 14$ days. Anyone familiar with basic control theory will immediately recognize this situation as a guaranteed recipe for disaster [11], [12].

4) Control feasibility: Well-known results from the theory of limitations of control can now be applied. In order to guarantee some robustness of the system performance against the large gain uncertainty of the process, the Bode plot of $\left|L(j\omega)\right|$ should maintain a roughly constant slope over a sufficiently wide interval around the crossover frequency ω_c, thus approximating Bode's ideal loop transfer function. The analysis reported in [12], Sect. 4.6, leads to conclude that in the most favourable case the product of the unstable pole p and of the time delay T should be $pT < 0.326$ for the process to be controllable. If one wants to limit the maximum norm of the sensitivity function $M_s < 2$, for increased robustness, the limit is even more stringent, namely $pT < 0.156$. Using the notation of this paper, these conditions become:

$$\frac{\tau}{T_d} < 0.47 \quad (23)$$

$$\frac{\tau}{T_d} < 0.225 \quad (24)$$

In other words, even under very optimistic assumptions on the knowledge of the process parameters, the process is controllable only if the total feedback loop delay is less than half of the doubling time of the epidemic, preferably less then one quarter. In the two validation cases reported in the previous section, the delay is more than double the doubling time, making the feedback control strategy utterly infeasible even in this idealized case. This refers to the initial phases of the outbreak, when $\beta \approx \beta_0$.

The exact details of mitigation policies have not been made public. What is understood is that some trajectory has been planned, and it will then be followed by changing the public health measures $u(t)$ when certain thresholds of the number of reported cases $I_r(t)$ are crossed, possibly also considering the number of new daily reported cases, which is related to $dI_r(t)/dt$. This would be an extremely crude step-wise approximation of a proportional-derivative controller $C(s)$, which is hardly going to perform much better than a carefully designed gain-scheduled linear controller.

Of course there is no theorem that can be directly invoked to prove that *any* feeback control policy would not suffer from the same limitations of a carefully scheduled linear controller. However, limitations of control in the case of unstable processes with large delays and uncertainty are inherent to the unfavourable nature of the process dynamics and not to the specific type of controller employed. In principle, a well-designed nonlinear, possibly time-varying controller could achieve somewhat better performance, but the large amount of uncertainty in the knowledge of the process behaviour makes this proposition completely impractical.

The analysis also clearly indicates under which conditions feedback control of the epidemic based on daily swab test reports may be feasible, which may give precious indications for the handling of the re-opening phase, once the suppression strategy has been successful in stopping the outbreak.

On one hand, it is essential to reduce the delay τ as much as possible, which could be in principle achieved if one had widespread instant-testing technology that could be applied to statistically relevant swaths of the population. This could probably halve the typical values of τ to about one week. On the other hand, the application of substantial, though not utterly draconian, measures, such as in the case of Lazio, could reduce $\beta(u^0)$ by another factor two/four, bringing the system to the brink of controllability, albeit with a very thin robustness margin.

IV. Conclusions

Governments all the world over are faced with very challenging life-or-death decisions regarding the management of the COVID-19 epidemic, involving the balance between public health and economic issues. In order to take such decisions, they rely on expert advice based on the results of epidemiological mathematical models and on daily reports of numbers of infectious people, based on swab test results.

This paper analyses the problem from a control systems perspective, casting it as a feedback control problem, as it actually is, and using a simple model that captures the control-relevant dynamic behaviour of sufficiently homogeneous territories, in which certain public health measures are applied. The model was validated in two cases, one corresponding to a draconian lockdown in the Hubei province of China, the other one to the application of severe, but milder social distancing measures in the region of Lazio, Central Italy.

The main results of the paper are the following:

- The suppression strategy can be effective, but requires a full understanding of the multiplicative effect of measurement delays, the factor M in Eq. (15), to correctly decide when it is time to enforce strict enough social distancing and lockdown measures.
- Mitigation strategies leading to herd immunity are not viable, because they require to let the outbreak run loose at the beginning of the transient to pick up high enough numbers of infected subjects, and the process simply cannot be controlled in such conditions, with high risk of catastrophic runaway scenarios.
- Mitigation strategies could in principle be applied to manage the re-opening phase after the outbreak has been effectively suppressed, but they would require extensive availability of instant testing equipment, as well as significant social distancing measures to ensure that the unstable dynamics of the epidemic process is at least two-three times slower than it is without any measure enforced. Even in this case, the control problem would be very difficult, and would probably have a significant likelihood of runaway scenarios, with eventual collapse of public health systems and unacceptable loss of life.

The application of the precautionary principle, a fundamental staple of European Union legislation, suggests to take the issue of controllability of the process very seriously. Given the limitations exposed in this paper, it is the opinion of the author that the safest way out of the COVID-19 epidemic is a combination of suppression and very aggressive research towards a vaccine and an effective cure of the severe pneumonia caused by the virus, which is the main cause of death and of the potential collapse of public health systems.

That said, in the absence of effective pharmaceutical solutions in the medium-long term, any exit strategy should be carefully studied with the tools of control theory, which may possibly suggest viable solutions that are not obvious to epidemiologists and physicians.

One such example, which is already discussed in this paper, is the awareness of the absolutely crucial role of reducing the measurement delay for the success of control policies.

Another example concerns the availability of reliable antibody tests that could be applied to statistically significant samples of the general population, allowing to provide a reliable estimate of the real state of the system, in particular regarding the number of subjects $R(t)$ that have recovered from the disease without showing significant symptoms and being tested for the virus. Such knowledge may help understanding by how much the contagion rate will be reduced because the term S/N in Eq. (1) is significantly less than one. The number of recovered subjects is believed to be much higher than the number of reported cases, but is currently largely unknown.

A control strategy based on such a measurement could be classified as state-feedback control, which control practitioners will immediately recognize as more effective than output feedback subject to large gain uncertainty and time delay.

Acknowledgment

The author is grateful to colleagues Gianni Ferretti, Alberto Leva, Marco Lovera, Luigi Piroddi, Paolo Rocco, and Andrea Zanchettin of Politecnico di Milano, for their reviews and constructive criticism, provided during the preparation of this paper, and to Paolo Mantegazza of Politecnico di Milano for help with data collection. He is also grateful to colleagues Renato Casagrandi and Marino Gatto of Politecnico di Milano, for earlier discussions about this topic.

References

[1] Z. Wu and J. M. McGoogan, "Characteristics of and Important Lessons From the Coronavirus Disease 2019 (COVID-19) Outbreak in China: Summary of a Report of 72 314 Cases From the Chinese Center for Disease Control and Prevention," *JAMA*, 02 2020.

[2] N. F. et al., "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," Imperial College London, Tech. Rep., 2020.

[3] D. Hunter, "Covid-19 and the stiff upper lip – The pandemic response in the United Kingdom," *The New England Journal of Medicine*, March 2020.

[4] W. Kermack and A. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A*, vol. 115, pp. 700–721, 1927.

[5] H. Hethcote, "The mathematics of infectious diseases," *SIAM Review*, vol. 42, pp. 599–653, 2000.

[6] G. Giordano, F. Blanchini, R. Bruno, P. Colaneri, A. D. Filippo, A. Di Matteo, and M. Colaneri, "A sidarthe model of covid-19 epidemic in italy," *Submitted to Nature Communications*, vol. n/a, no. n/a, 2020. [Online]. Available: https://arxiv.org/abs/2003.09861

[7] S. F. et al., "Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries," Imperial College London, Tech. Rep., 2020.

[8] J. Read, J. R. Bridgen, D. A. T. Cummings, A. Ho, and C. P. Jewell, "Novel coronavirus 2019-nCoV: early estimation of epidemiological parameters and epidemic predictions," *medRxiv*, pp. 1–11, 2020.

[9] Italian Society of Anesthesiology, Analgesia, Resuscitation and Intensive Care, "Raccomandazioni di etica clinica per l'ammissione a trattamenti intensivi e per la loro sospensione in condizioni eccezionali di squilibrio tra necessità e risorse disponibili (in Italian)," SIAARTI, Tech. Rep., 2020. [Online]. Available: http://www.siaarti.it/SiteAssets/News/COVID19%20-%20documenti%20SIAARTI/SIAARTI%20-%20Covid19%20-%20Raccomandazioni%20di%20etica%20clinica.pdf

[10] D. della Protezione Civile, "Repository of COVID-19 outbreak data for Italy," https://github.com/pcm-dpc/COVID-19, 2020.

[11] G. Stein, "Respect the unstable," *IEEE Control Systems Magazine*, vol. 23, no. 4, pp. 12–25, 2003.

[12] K. Åström, "Limitations on control system performance," *European Journal of Control*, vol. 6, no. 1, pp. 2–20, 2000.

Francesco Casella got his master's degree *cum laude* in Electronic Engineering and his PhD in Information and Control Engineering from Politecnico di Milano, where he is senior assistant professor in Control Systems and Advanced Process Control. His main research interests are modelling and control of energy systems and equation-based, object-oriented modelling of engineering systems. He is the Director of the Open Source Modelica Consortium and a member of the Board of the Modelica Association.