Test and Contain: A Resource-Optimal Testing Strategy for COVID-19

Jakob Jonnerby[1] , **Philip Lazos**[2] , **Edwin Lock**[3] , **Francisco Marmolejo-Cossío**[3] , **C. Bronk Ramsey**[4] , **Divya Sridhar**[5]

[1]Department of Physics, University of Oxford
[2]Department of Computer Science, Sapienza University of Rome
[3]Department of Computer Science, University of Oxford
[4]School of Archaeology, University of Oxford
[5]Department of Zoology, University of Oxford

Abstract

We propose a novel testing and containment strategy to limit the spread of SARS-CoV2 while minimising the impact on the social and economic fabric of countries struggling with the pandemic. Our approach recognises the fact that testing capacities in many low and middle-income countries (LMICs) are severely constrained. In this setting, we show that the best way to utilise a limited number of tests during a pandemic can be found by solving an allocation problem. Our problem formulation takes into account the heterogeneity of the population and uses pooled testing to identify and isolate individuals while prioritising key workers and individuals with a higher risk of spreading the disease. In order to demonstrate the efficacy of our strategy, we perform simulations using a network-based SIR model. Our simulations indicate that applying our mechanism to a population of $10,000$ individuals with only 1 test per day reduces the peak number of infected individuals by approximately 27%, when compared to the scenario where no intervention is implemented, and requires at most 2% of the population to self-isolate at any given point.

1 Introduction

Our work sets out a new testing and containment strategy to limit the spread of COVID-19 while minimising the impact on the social and economic fabric of countries struggling with the pandemic. In LMICs, testing resources for viral infections are often extremely scarce.[1] At the same time, testing is an essential tool for combating epidemics, as it provides crucial estimates of virus prevalence and allows for the identification of asymptomatic and symptomatic infected individuals, which forms the basis of many containment policies. In settings where comprehensive individual testing is infeasible, radically new testing policies are needed to maximise the reach of an extremely limited number of tests. Mexico, for instance, only has the ability to conduct up to 72 tests per day per one million inhabitants (as of 6 July 2020). [13]

[1]Testing is constrained by many factors including manpower, laboratory equipment, availability of reagents, and logistics.

Work in the testing literature has largely focused on determining the number of tests required to execute a given testing and containment strategy. Turning the problem on its head, we instead consider the problem of identifying a mechanism that maximises the benefit of a fixed weekly testing budget. Our approach [7] treats this challenge as a resource allocation problem aimed towards minimising infection while mitigating the effects of state-organised containment measures on personal well-being and the economy. The heterogeneous strategy we propose classifies individuals according to their exposure to the virus and their cost of containment. The latter can refer to social cost, e.g. a healthcare worker who by self-isolating is unable to perform essential duties, or financial cost, similar to [10]. Using this classification, our approach performs a series of group tests and subsequently isolates all individuals contained in groups that test positive while allowing the other individuals to resume social and economic activities.

In group testing, multiple samples are pooled into a single test to learn whether any given person in the pool is infected or, more informatively, if none are infected. There exists a substantial literature on group testing in the Computational Learning Theory community [1–4, 15] and the underlying method has also been applied in the fight against HIV and other diseases [11, 14]. Group testing has been verified experimentally [16] and has been proposed as possible way towards testing large parts of the population [6]. One of the most compelling benefits of group testing is its ability to amplify the reach of a limited number of tests to be able to test more people. A key parameter in this respect is choosing the optimal size of the groups to be tested. Our work addresses how a population can be intelligently segmented, how tests can be allocated across segments to maximise socioeconomic benefit, and what group size should be used for each population segment.

Our mechanism can be used to exit lockdown and is especially useful for monitoring the population for asymptomatic carriers who may unknowingly spread the disease. [9] Moreover, we emphasise that our proposed mechanism provides considerable flexibility for policymakers to optimise the balance between virus containment and socio-economic welfare.

1.1 Preliminaries

Our approach exploits the heterogeneity of the population when deciding how to allocate tests. To model this, we consider a population $[n] := \{1, \dots, n\}$ of size n which is partitioned into k disjoint categories $C_1, \dots C_k$ and let n_i denote the number of individuals in category C_i. Our mechanism tells groups of individuals to self-isolate for a certain duration of time when the individuals may be infected. In some cases, however, a policy maker may wish to isolate an entire category on a permanent basis.[2] In order to keep track of such an *exogenous* action, we let the parameter $S_i \in \{0, 1\}$ indicate whether the entire category C_i has been isolated.

By assumption, all individuals in category C_i are independently infected with probability p_i (and healthy with probability $q_i = 1 - p_i$). Every individual in C_i also has an integral or rational 'exposure parameter' $d_i \geq 0$. One possible interpretation of d_i is the number of other people that the individual is in regular contact with when not in self-isolation, which corresponds to the number of neighbours they have in a social network. Consequently, a higher value of d_i means that an individual has a higher probability of being infected in the first place, and in the case of infection, a higher expected number of individuals they can propagate the disease to if they are not in self-isolation.

All individuals in C_i also have a rational cost of containment denoted $\gamma_i \geq 0$. For example, a category consisting of healthcare workers will have a high γ_i value, whereas a category with individuals in other professions such as software engineering who are able to to work from home will have a low γ_i value. The cost of containment may also include financial cost. For example, a category consisting of daily wage labourers may have a high γ_i value as they do not have the economic means of maintaining self-isolation.

In practice, d_i, p_i and γ_i values vary for each individual in a given population. There exist a variety of means to estimate these values, such as using information from symptom tracking apps, generating exposure heat maps of a geography from location data, or simply using existing government data about a population. Furthermore, the determination of these values can go in hand with a clustering of the population to fit the categorisation mentioned above. We refer the reader to our pre-print [7] for more details.

Finally, we let T denote the number of test kits that are made available per day - this is our *testing budget*. Each test kit can be used to perform a single individual or group test. The maximum feasible group size for pooled tests is determined by biological and practical constraints, and we denote it by G.

Our mechanism allocates ℓ_i tests to each category C_i such that $\sum_{i \in [k]} \ell_i \leq T$. For each category it decides the size, g_i, of the group tests, i.e. the number of people that are pooled together in a single group test. Note that our mechanism may decide to perform individual tests for some category C_i, which is equivalent to setting g_i to 1. For any category C_i, we say that we are performing *uniform group testing* with *granularity* $g_i \geq 1$ and *scope* $\ell \geq 1$ if we perform group tests of ℓ_i *disjoint* groups of size g_i. Note that this implies that $\ell_i g_i \leq n_i$.

[2] This was the case, for instance, in the UK, where all individuals who did not fit into the 'keyworker' category were compelled by law to follow a lockdown order.

1.2 Our Testing and Allocation Mechanism

Our testing and allocation mechanism performs periodic (e.g. weekly) and proactive testing across the population while prioritising population categories with higher exposure or containment cost by allocating more tests and setting smaller group sizes, respectively. We then solve the optimisation problem defined in Section 1.3 to decide how many group tests ℓ_i to allocate to each category C_i, and the group sizes g_i for these tests. Having determined these parameters, we perform uniform group testing with scope ℓ_i and granularity g_i for each category. If a group test returns negative, we allow all individuals in the group to resume social and economic activity, as they cannot infect others. Otherwise, if the group test returns positive, we quarantine all individuals in the group.

1.3 Test Allocation as an Optimisation Problem

Recall that the aim of our mechanism is to contain the spread of the virus while minimising the impact on the social and economic fabric of the country. Formally, we achieve the former objective by minimising the number of infected individuals that are not covered by a group test, while the latter objective can be achieved by minimising the number of healthy individuals that are quarantined unnecessarily as they form part of a group that tests positive. Moreover, our two objectives are weighted by the exposure and the containment cost, respectively. Note that our objectives compete with each other in a certain sense: setting larger group sizes increases the reach of testing and leads to fewer untested but infected individuals, while setting smaller group sizes reduces the number of healthy individuals that are quarantined unnecessarily. Recalling that n_i denotes the number of agents in category C_i and that performing ℓ_i uniform testing of granularity g_i covers $g_i \ell_i$ individuals in category C_i, we define $r_i := n_i - g_i \ell_i$ as the number of individuals in category C_i that are untested. For each category C_i, we define a function $\theta_i(g)$, which denotes the *optimisation cost* of administering a test at granularity g to individuals in C_i. If $S_i = 0$, i.e. if the entire category C_i is not permanently isolated, we let $\theta_i(g) = (\gamma_i q_i - d_i p_i)g - \gamma_i g q_i^g$. Otherwise, if $S_i = 1$, we let $\theta_i(g) = -\gamma_i g q_i^g$.

This definition of θ_i formalises the following rationale. If C_i is not under self-isolation, untested individuals may in fact be infected, and they pass the virus to other individuals proportional to their exposure. If a group test is performed, healthy individuals may be within the group test and forced to self-isolate unnecessarily, incurring a cost of γ_i. On the other hand, if C_i is under self-isolation, untested individuals may in fact be healthy, and hence remain under unnecessary self-isolation and incur a containment cost of γ_i. On the other hand, if a group test returns a positive result, there may be healthy individuals in the group who once again are forced to self-isolate unnecessarily. For further details on the construction of θ_i, we refer to our pre-print [7].

With this in hand, we can express the overall cost incurred from testing segment C_i with ℓ_i tests at granularity g_i by $\ell_i \theta_i(g_i)$. Consequently, we can consider how to test according

to the following optimisation problem:

$$\begin{aligned} \min_{g,\ell} \sum_{i=1}^{k} \ell_i \theta_i(g_i) \\ \text{s.t.} \sum_{i=1}^{k} \ell_i \leq T, \qquad g_i \leq G, \\ \ell_i g_i \leq n_i, \qquad g_i, \ell_i \in \mathbb{N}. \end{aligned} \tag{1}$$

The objective function is separable (as the sum of individual loss functions per segment) and linear in ℓ_i for constant g_i. Hence, if we fix group sizes g_i, we can formulate our optimisation problem as an integer linear programme (ILP) in ℓ with k free variables and $k+1$ inequality constraints.

This suggests a straightforward approach to solving (1). For every feasible group size vector $g = (g_1, \ldots, g_k)$, solve the resulting ILP for $\ell = (\ell_1, \ldots, \ell_k)$; return the combination of g and ℓ that minimises the objective. More importantly, the separable nature of (1) implies that the optimal allocation of the T tests can be computed via a simple greedy algorithm, which we state in Algorithm 1. Intuitively, once the g_i's have been set, the resulting ILP becomes a weighted sum of the ℓ_i's, which in turn have to sum up to T. Therefore, we proceed by gradually increasing the ℓ_i with the highest weight $\theta_i(g_i)$ until the $\ell_i g_i \leq n_i$ constraint is met. As there are at most G^k different ways to fix g, this method works in practice if we have few population categories and constrain the set of possible group sizes for each category's tests. Notice that the **for** loop in Algorithm 1 can be run entirely in parallel.

Algorithm 1 Optimal Segmented Uniform Group Testing

Require:
Possible group sizes: $R_G \subseteq [G]^k$
Iterating over Granularities:
$OPT \leftarrow \infty$
for $g \in R_G$ **do**
  $\ell_i \leftarrow 0$ for $i \in [k]$
  Compute σ, an ordering of C_i with respect to increasing $\theta_i(g_i)$ values
  $T_r \leftarrow T$, $i \leftarrow 1$
  while $T_r > 0$ **do**
    Set ℓ_{σ_i} to its maximum possible value.
    $\ell_{\sigma(i)} \leftarrow \min\left\{T_r, \left\lfloor \frac{n_{\sigma(i)}}{g_{\sigma(i)}} \right\rfloor\right\}$
    $T_r \leftarrow T_r - \ell_{\sigma(i)}$, $i \leftarrow i+1$
  Update the optimal solution
  if $\sum_{i=1}^{k} \ell_i \theta_i(g_i) < OPT$ **then**
    $OPT \leftarrow \sum_{i=1}^{k} \ell_i \theta_i(g_i)$
    $g^*, \ell^* \leftarrow g, \ell$

return g^*, ℓ^*, OPT

2 Modelling Testing Allocation During an Epidemic Process

To verify the efficacy of our testing allocation mechanism, we developed a simple network-based Susceptible-Infected-Recovered (SIR) model on a heterogeneous population using the graph-tool Python library [12]. We used this to model the impact of an uneven distribution of the exposure on the epidemiological process. The exposure parameter d_i was identified with the connectivity (also known as the degree) of a node in the network model. The network was initiated as a random geometric graph, meaning that nodes are scattered at random in a 2-dimensional space and connected to their neighbours within a predefined constant radius. We also introduced two other types of nodes, representing key workers, by connecting them to 20 other nodes chosen at random from all nodes. 20 of the key worker nodes represented health care workers, with a high cost of containment, while 480 represented marketplace workers, with a lower cost of containment. In total $10,000$ nodes were initiated for the simulation.

The SIR model is run for a fixed number of (discrete) time steps. In order to ensure high temporal resolution, we let each day consist of 10 time steps and run the simulation for 200 days. At each time step, an infected node recovers with probability γ. Moreover, if it is not self-isolating, it infects each susceptible neighbouring node with probability β. The intervention mechanism consists of performing a fixed number of group tests each day and enforcing social isolation for everyone in a group that tests positive for the virus. The parameters $\beta = 0.02$ and $\gamma = 0.0427$ were chosen such that average number of secondary infections and time until recovery are $R_0 \sim 2.5$ [5] and $t_{\text{recovery}} = 14$ days [8].

We explored the impact of different testing allocation scenarios on the epidemic outcome, paying particular attention to the peak number of infected individuals and to the number of quarantined individuals at any given point in time. The scenarios were constructed with the testing capacity of a LMIC in mind, with 1 test available per $10,000$ inhabitants each day. Self-isolation was conditional on being part of a group that tested positive, and the allocation of tests was determined using two different strategies. Following our testing and containment strategy, the population was categorised into two segments depending on their exposure. Each segment was further subdivided into "key workers" and "non-key workers". We then applied Algorithm 1 in order to allocate tests each day to the different segments according to their probability of infection, exposure, and cost of containment. In the 'Random Sampling' testing strategy, used as a control strategy, testing was conducted in groups of 10 that were sampled uniformly at random from the population, independently of an individual's exposure or cost of containment.

We assumed an initial infection rate of 0.1% in the population and that testing started from day 10. For each scenario, a population of $10,000$ was simulated over 200 days, and was initiated at random 100 times, which allowed us to obtain the mean and the standard deviation of our simulations.

2.1 Results

Figure 1 shows the number of infected and quarantined individuals for the three strategies: no intervention, random sampling and our resource maximising mechanism. When no interventions were performed, the total number of infected individuals was found to be $1,800$ at maximum, or 18% of the population. Our testing and containment strategy shows a reduction by $(27 \pm 5)\%$ of the peak height in the number

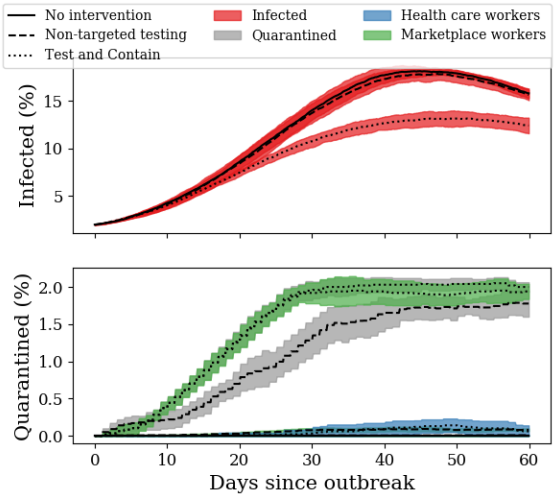

Figure 1: Top: Number of infected for different testing and containment strategies at different stages of the epidemic. Bottom: Total number of quarantined, and number of quarantined health care workers and marketplace workers during the pandemic. The coloured bands represent the standard deviation of each quantity.

of infected compared to the baseline scenario. The random sampling strategy reduced the peak height by $(1.5 \pm 4)\%$ compared to the baseline. The total number of quarantined during the peak were approximately 2% in both cases.

3 Discussion

Testing is a limited resource that can have many different and valuable purposes during a pandemic. In this study, we have considered methods of maximising the utility of a limited number of tests in a heterogeneous population. By combining segmented testing with group testing, our proposed testing mechanism can dramatically enhance the utility of a small number of test kits. Our approach is conceptually simple and general. It can be applied in scenarios where countries want to ease out of lock-down, as well as to avoid a possible second peak in the number of infections. Furthermore, it can also be applied for smaller populations such as care homes, refugee camps, jails, etc. where it is more feasible to map social dynamics and estimate relevant parameters to our model in order to optimise testing allocation within the population.

In order to validate our method, we conducted simulations on a network-based SIR model. The results show that our strategy can substantially reduce the peak height in the number of infections even when the number of tests available is severely constrained. This is especially valuable in large outbreaks, where testing capacities are significantly smaller than the affected population. Our simulations do not include implementing any other interventions, such as isolating symptomatic cases, which would further increase the reach of a limited number of tests by mainly testing asymptomatic individuals.

Importantly, our method can be used together with a contact tracing app. Assuming that information could be obtained about each individual's connectivity from the app, testing resources could be focused on those with high connectivity.

3.1 Future Directions

The optimisation problem in Section 1.3 has already proved useful in providing non-trivial testing strategies over a segmented population. However, multiple natural refinements to the model can be made in order to tackle the practical and logistical challenges present in many LMICs.

Incorporating test accuracy Our current optimisation problem assumes that tests are perfectly accurate. In practice, tests have a false positive and false negative rate which can be taken into account with a slight modification to our resource-constrained optimisation framework. To achieve this, we propose the following. Each group test on g individuals is considered to have an "intensity" s, and as a function of the intensity, the test has a false positive rate of $p^+(s)$ and a false negative rate of $p^-(s)$. Intensity can be understood as the result of making s independent tests on the same group, and taking the majority vote of the results as the final group test result. We can thus construct an identical containment policy to that in Section 1.3, but with realistic false positive and false negative rates. We obtain a similar optimisation programme where each category C_i is tested at granularity g_i, scope ℓ_i and intensity s_i, resulting in a testing cost of $\ell_i s_i$ for that segment.

Pseudo Group Testing In practice, performing group testing may not always be possible for various reasons. For instance, while group testing reduces demand for test kits, it still requires the collection of large numbers of test samples, posing significant logistical challenges. A possible alternative to group testing is to *simulate* group testing by selecting a 'representative' subset of each group, selected uniformly at random. The benefits of pseudo group testing are limited if the groups themselves are randomly sampled from a population category or the entire population. However, when data is available to identify and test highly correlated groups, sampling a random subset of the group yields a good estimate of their health status. We note that one natural application of this is to pool households and testing a single individual in each household.

Non-Disjoint Group Tests In the mechanism described in Section 1.2, we perform uniform group testing by applying a group test to ℓ_i *disjoint* groups of size g_i. In practice, ensuring the disjointedness of groups that are tested can be challenging. Instead, we can consider the simpler approach of randomly sampling groups of size g_i with replacement.

It is important to note that this approach may result in individuals being tested more than once in a single round, leading to 'testing fatigue'. Moreover, given the same number of tests, this approach yields weakly less information than uniform group testing. Indeed, the expected utility derived from tests scales sub-linearly with the number of tests. However, non-disjoint group testing may be a conceptually simpler approach when dealing with the increased number of tests required in the setting with imperfect testing, as it allows us to increase the number of tests per segment beyond n_i/g_i.

References

[1] M. Cheraghchi, A. Karbasi, S. Mohajer, and V. Saligrama. Graph-constrained group testing. *IEEE Transactions on Information Theory*, 58(1):248–262, 2012.

[2] R. Dorfman. The detection of defective members of large populations. *The Annals of Mathematical Statistics*, 14(4):436–440, 1943.

[3] D. Du and F. Hwang. *Pooling designs and nonadaptive group testing: important tools for DNA sequencing*. World Scientific, 2006.

[4] D. Du, F. K. Hwang, and F. Hwang. *Combinatorial group testing and its applications*, volume 12. World Scientific, 2000.

[5] S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, H. Coupland, T. A. Mellan, H. Zhu, T. Berah, J. W. Eaton, P. N. P. Guzman, N. Schmit, L. Callizo, K. E. C. Ainslie, M. Baguelin, I. Blake, A. Boonyasiri, O. Boyd, L. Cattarino, C. Ciavarella, L. Cooper, Z. Cucunubá, G. Cuomo-Dannenburg, A. Dighe, B. Djaafara, I. Dorigatti, S. van Elsland, R. FitzJohn, H. Fu, K. Gaythorpe, L. Geidelberg, N. Grassly, W. Green, T. Hallett, A. Hamlet, W. Hinsley, B. Jeffrey, D. Jorgensen, E. Knock, D. Laydon, G. Nedjati-Gilani, P. Nouvellet, K. Parag, I. Siveroni, H. Thompson, R. Verity, E. Volz, P. G. Walker, C. Walters, H. Wang, Y. Wang, O. Watson, C. Whittaker, P. Winskill, X. Xi, A. Ghani, C. A. Donnelly, S. Riley, L. C. Okell, M. A. C. Vollmer, N. M. Ferguson, and S. Bhatt. Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries. *Imperial College London*, March:1–35, 2020.

[6] C. Gollier and O. Gossner. Group Testing Against Covid-19. *EconPol Policy Brief*, 2020.

[7] J. Jonnerby, P. Lazos, E. Lock, F. Marmolejo-Cossío, C. B. Ramsey, M. Shukla, and D. Sridhar. Maximising the Benefits of an Acutely Limited Number of COVID-19 Tests. *ArXiv (preprint)*, pages 1–12, 2020.

[8] S. A. Lauer, K. H. Grantz, Q. Bi, F. K. Jones, Q. Zheng, H. R. Meredith, A. S. Azman, N. G. Reich, and J. Lessler. The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application. *Annals of Internal Medicine*, 03 2020.

[9] S. Lee, P. Meyler, M. Mozel, T. Tauh, and R. Merchant. Asymptomatic carriage and transmission of SARS-CoV-2: What do we know? *Canadian Journal of Anesthesia*, 2020.

[10] Matthew Cleevely, D. Susskind, D. Vines, L. Vines, and S. Wills. A workable strategy for Covid-19 testing: Stratified periodic testing rather than universal random testing. *CEPR Press*, 8, 2020.

[11] C. S. Mcmahan, J. M. Tebbs, and C. R. Bilder. Informative Dorfman Screening. *Biometrics*, 68(1):287–296, 2012.

[12] Peixot, Tiago. Graph tool, 2020. [Online; accessed 15-April-2020].

[13] M. Roser, H. Ritchie, E. Ortiz-Ospina, and J. Hasell. Coronavirus pandemic (covid-19). *Our World in Data*, 2020. https://ourworldindata.org/coronavirus.

[14] L. M. Wein and S. A. Zenios. Pooled testing for HIV screening: Capturing the dilution effect. *Operations Research*, 44(4):543–569, 1996.

[15] J. Wolf. Born again group testing: Multiaccess communications. *IEEE Transactions on Information Theory*, 31(2):185–191, 1985.

[16] I. Yelin, N. Aharony, E. Shaer Tamar, A. Argoetti, E. Messer, D. Berenbaum, E. Shafran, A. Kuzli, N. Gandali, O. Shkedi, T. Hashimshony, Y. Mandel-Gutfreund, M. Halberthal, Y. Geffen, M. Szwarcwort-Cohen, and R. Kishony. Evaluation of COVID-19 RT-qPCR test in multi-sample pools. *Clinical Infectious Diseases*, 05 2020.