Optimization of Resource-Constrained Policies for COVID-19 Testing and Quarantining

In memoriam of Benny Chor: A scientist, teacher, and mensch

Konstantin Berestizshevsky, Koffi-Emmanuel Sadzi, Guy Even, and Moni Shahar

Abstract—We address the problem of controlling the COVID-19 contagion with a limited number of PCR-tests. We developed a tool that can assist policy makers in decisions as well as in justifying these decisions. Our tool consists of: A stochastic disease model, a compressed representation of interactions between people via a graph that scales well to large populations, policies for selecting PCR-testees per day, and a simulator that simulates the spread of the COVID-19 while taking into account the testing and quarantine decisions of the chosen policy.

The graph model includes features that help determine the infection risk of individuals. We consider both external infection (inflicted by people outside the studied community) as well as internal infection. The graph model and known infections induce weights to people. These weights are used to select the testees per day in a greedy algorithm and in a linear-programming optimization algorithm.

Our simulations indicate a reduction in total morbidity of $30-50\%$ using the optimization algorithm compared to random sampling. A reduction of up to 40% in peak morbidity is achieved compared to random sampling. We also studied the efficiency of quarantining in various policies.

Index Terms—Disease model, optimization, simulation

I. Introduction

BEFORE he introduction of vaccines, the rapid contagion of the COVID-19 pandemic has been counter-measured by PCR-testing, quarantining, and lock-downs. Each counter-measure bears a cost, incurs discomfort, and disrupts life. Policy decision makers must cope with the dilemma of how to control the contagion effectively and how to justify the decisions so that public trust is not eroded.

This research was initiated when we were approached by the managers of a chain of assisted living facilities in Israel. Each facility houses about two hundred people and employs roughly two hundred people (administration, maintenance, kitchen workers, caregivers, etc.). Residents, most of them in their eighties, are highly vulnerable to spread of

the pandemic and protecting them became a national goal. Epidemic investigations in the facilities revealed that infection was mostly introduced by employees. The initial goal was to detect COVID-positive individuals as soon as possible (upon discovery of an outbreak, the whole house was tested and quarantined for 14 days). However, the number of allocated PCR-tests were insufficient, raising the need to develop a sampling policy for the efficient allocation of PCR-tests. We, therefore, developed a sampling policy that is based not only on a person's probability of getting infected, but also on the potential to infect others. This potential captures the patterns and strength of the interaction between people.

Building on this experience, we propose an approach that takes into account the exposure patterns between individuals. This approach is designed to apply to organizations consisting of several hundreds of people, and scales well to much larger communities (e.g., school systems, towns, as well as states).

We developed a tool that can aid policy makers, help justify decisions, run what-if scenarios, and find desirable trade-offs between controlling the contagion and preserving daily life and the economy. This tool consists of four main parts:

1) Disease model that specifies the disease stages and time per stage.
2) A graph for modeling the interactions between people (i.e., *community graph*).
3) A simulator that simulates the contagion.
4) A policy for testing the people, as well as quarantining and releasing them from the quarantine (in the literature, such policies are often referred to as "interventions" [1]).

One can view our tools as an application of methods from the area of networks. The disease model is a stochastic model that is based on interactions between people. Our graph model is a compressed representation of interactions between people that scales well to large populations. Our simulator provides a platform for simulation the contagion, incorporating the decisions of the policy (i.e., who to test? who to quarantine?), and logging the state of each person in every day. Below, we elaborate on each part.

A. Disease Model

The disease model for COVID-19 (per individual x) is depicted in Fig. 1. The model is based on a state diagram consisting of 7 states: "Susceptible" (i.e., healthy and un-infected to date), "infected" (but not detectable yet in a PCR-test), "detectable" (infected and will return a positive result if tested), "contagious pre-symptomatic", "contagious
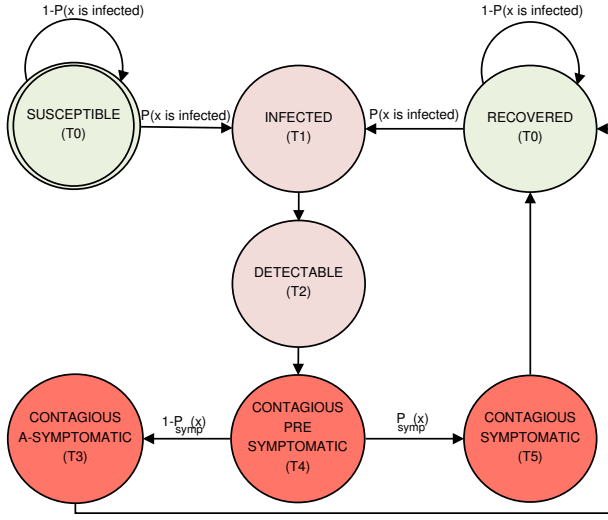
Fig. 1. COVID-19 disease model of an individual x. The model consists of 7 states with stochastic transitions between the states. The label T_i in a state is a random variable that specifies the number of days that x stays in the state. Transition probabilities appear as edge labels. The states colored red are states in which the individual is contagious. We consider an individual to be *positive* if the state is "infected", "detectable", or in one of the contagious states.

Fig. 2. Example of community graph with the set of people containing 14 people (9 students, 3 teachers a manager and a secretary) and 3 groups (Class 1, Class 2, and Management). Each person can belong to 0 or more groups.

asymptomatic", "contagious symptomatic", "recovered". The duration (measured in days) per state is a random variable, denoted by T_i and its distribution is further described in Section II-B. After the duration of a state ends, a transition to the next state takes place. Transition probabilities appear next to the edges. For example, a person x in the "contagious pre-symptomiatic" state stays in this state for T_4 days. After T_4 days, with probability $P_{\text{symp}}(x)$ the state transitions to "contagious symptomatic", and with probability $1 - P_{\text{symp}}(x)$ the state transitions to "contagious asymptomatic". We note that for the states "susceptible" and "recovered", $T_0 = 1$, and every day, a stochastic transition to the "infected" state takes place. One advantage of this state machine is that it models one of the main challenges in coping with the COVID-19 pandemic, namely, the high occurrence of contagious asymptomatic people that spread the disease unconsciously (especially among young people). We note that the transition probabilities depend on the features of individual x.

B. Community Graph

The community graph is a hypergraph $G = (\mathcal{X}, \mathcal{D})$, where $\mathcal{X}$ denotes the set of people, and $\mathcal{D}$ is a family of subsets of people. Each group models a subset of people that meet (e.g., pupils in the same class, teachers sharing a teachers' room). We refer to two people x and y as *colleagues* if there exists a group $d \in \mathcal{D}$ such that both x and y belong to d. We attach attributes, called *risk factors* to people and groups. Attributes of individuals are classified as either external or internal. External attributes include, for example: family size, usage of public transportation, mobility. Internal attributes quantify the individual risk within a group. For example, a kindergarten teacher is more likely to be infected than the children. Group attributes quantify the density/proximity/duration of interac-
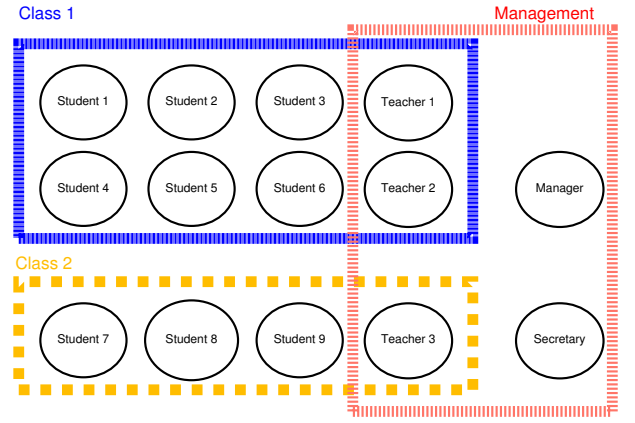
tion caused by factors such as ventilation, joint eating, inability to adhere to social distancing. See Fig. 2 for an example of a community graph.

The main advantage in the community graph is that it offers a compressed representation of the interactions between people. The more common approach is to represent interactions by a graph over people (i.e. *people-graph*) in which edges model the interaction between pairs of people. A people-graph can have a quadratic number of edges, rendering the simulation of large communities impractical. For example, in the extreme case that the people-graph is a complete graph, the community graph is linear and contains one vertex per person and a single group that contains all the people.

C. Simulator

The simulator performs a day-by-day simulation of the disease spread among the people as well as of the progress of the disease in each of the people infected. The simulator deals with two sources of infection: internal and external. An *external infection* is the event that person $x \in \mathcal{X}$ is infected by person $y \notin \mathcal{X}$ (namely, y is outside the set of the people $\mathcal{X}$ that the model focuses on). An *internal infection* is the event that person $x \in \mathcal{X}$ is infected by colleague y (i.e., $y \in \mathcal{X}$ and $\{x, y\} \subset d \in \mathcal{D}$). Each step of the simulator corresponds to a day. Let $X(x, t)$ denote the event that person x is externally infected on day t. The probability of $X(x, t)$ is derived from the external risk factors of the individual. We treat the events $X(x, t)$ as independent for all people x that are COVID19-negative before day t. Internal infection of a person x takes into account all the positive colleagues of x. The probability of infection of x within a group d is influenced by the the group's risk and the internal risk of x. A quarantined individual is excluded from the simulation and cannot infect others.

D. Policies

As interventions, we consider five policies:
1) No policy (i.e., no quarantining at all).
2) Symptom based (quarantine only people in the state "contagious symptomatic".

3) Random testing (RAND) (symptom-based plus random sampling and quarantining of positive people).
4) Risk-Factor-Greedy (RFG) (symptom-based plus greedy sampling and quarantining of positive people).
5) Optimization based (OPTIMIZATION) (symptom-based plus sampling based on optimization and quarantining of positive people).

Testing is applied in the policies RAND, RFG and OPTIMIZATION, and is subject to a budget of B tests per day. Note that RFG and OPTIMIZATION rely on the community-graph and the risk factors of people and groups. A detailed explanation of these policies is given in Section III.

An important issue is when to release a person from quarantine. A person is quarantined if symptomatic or found positive by a test. There are a few options for deciding when a quarantine should end (if symptoms disappear). One option, chosen in Israel, is to quarantine for a fixed amount of time (e.g., 14 days later reduced to 10). We chose the following option: start testing every two days starting from the fourth day of quarantine (provided that symptoms have disappeared) until a negative test is obtained. The rational is to avoid unnecessary quarantines (since a day of quarantine costs more than a test). We note that the testing of quarantined people is not counted in the budget of tests. A discussion regarding the choice of the release from quarantine appears in Section III-B.

We note that one can view epidemiological investigations [2] as a greedy approach. Namely, people exposed to COVID-positive individuals are considered to be in high risk and are either quarantined or tested. One method that assisted epidemiological investigations is based on tracing interactions with the help of a mobile-app based [3].

E. Contribution

We provide a framework for assessing the effectiveness of various testing and quarantining policies in controlling the COVID-19 contagion. Our framework introduces novel quarantine efficiency metrics and enables one to measure peak and total morbidity as well as justifiable application of quarantines as a function of the number of tests per day. Our framework is applicable to communities of various scales (e.g., single school, town, or even states). The source code of our framework is publicly available. [1]

Our disease model differs from previous approaches mainly in the distinction between symptomatic and asymptomatic contagious people. Partial observability, of whether a person is asymptomatic yet contagious without testing, is a major challenge in combating the COVID-19 pandemic (as rates of asymptomatic yet contagious illness are especially high in youngsters [4]). In addition, we model the short time (about a day) between detectability and contagiousness. The duration of disease states and transition probabilities are based on recent data adding to the realistic aspects of the disease model.

We consider two policies for testing with limited budgets: a greedy approach and a linear-programming based approach. These approaches rely on risk factors of individuals as well

as groups. We evaluate and compare these policies with random sampling over realistic disease spread dynamics. In addition, we demonstrate that an optimized approach reduces the sensitivity to PCR test errors and explore the impact of approaches to release people from quarantine.

Our simulator differentiates between external and internal infection. Indeed, external infections are inflicted by people that do not belong to the investigated population (and do not appear in the community graph). Moreover, epidemic investigations have shown that external infection occurs continuously over time and introduces new variants and creates new regions of high morbidity.

F. Related Work

The SEIR model in epidemiology considers four states: "susceptible", "exposed" (i.e., latent yet not contagious), "infectious" (i.e., contagious), and "recovered" [1], [5]. The SEIR is designed to model a period of incubation in which an individual is infected yet does not spread the disease. The SEIR model is considered inadequate for modeling the COVID-19 pandemic [6]. The SEIR model and its variants (such as the SIR and the SIS models [7]) are accompanied by differential equations that model the dynamics but do not address the nonuniform specific interaction patterns between people.

There is a significant line of work that focused on a population-level disease spread modeling, rather than on an individual-level disease spread modelling we use in the current work. Such works include the study of Carcione *et al.* [8] based on an SEIR model where the propagation of the disease is modeled using differential equations. The works of Silva [9] and the work of Stevens [10] developed a "billiard ball" simulation that model interactions between people using collisions between billiard balls with random velocities. A different family of simulations consider a graph of interactions that is usually a two-dimensional grid [11], [12].

The work by Newman [13] described a solution of a disease spread model in a network of individuals. However, the networks discussed in the paper of Newman were random graphs, such as graphs generated from a power-law distribution. In our paper we deal with structured community networks that are modelled by hypergraphs where the vertices are people and the hyper-edges are groups. Infection takes place only if these two individuals share a common group. This structure of a network is common in various real-world communities such as daycare centers, schools, assisted living facilities and workplaces in general. We hope that our model of disease spread will be found useful in real-world scenarios. We also note that hypergraph-based disease spread models have been gaining interest recently [14]–[16].

Optimization-based approaches were applied in the field of epidemic control by Deng *et al.* [17]. In their study, they used a bipartite-graph-based community structure of people and places they visit, analogous to our hypergraph-based community structure. Deng *et al.* employed a linear programming approach to sample a subset of people to be vaccinated. Our work differs from the work of Deng *et al.* mainly by the

Symbol	Domain	Description
$\mathcal{X}$		Set of all the people in the organization
$\mathcal{D}$	Power set $(\mathcal{X})$	Set of all the groups in the organization
$risk_{\text{ext}}(x)$	$[0,1]$	The risk of a person x to get infected from the outside of the organization.
$risk_{\text{int}}(x)$	$[0,1]$	The risk of a person x to get infected from the inside of the organization due to personal susceptibility (not membership in groups)
$risk_{\text{int}}(d)$	$[0,1]$	Risk of infection due to membership to group d
$\lambda(x)$	$[0,1]$	Discount factor (influenced by time since recovery or quarantine)
$\tilde{\lambda}(x)$	$[0,1]$	Discount factor used by the policy (influenced by time since known recovery, time since negative test, or quarantine)
$\delta(x)$	$[0,1]$	Contagiousness of person x
B	$\mathbb{N}$	Daily budget of PCR-tests

fact that we deal with testing and quarantining, rather than with vaccination. Additionally, we examine the quarantine efficiency (namely, the overlap between the quarantine period and illness period).

Machine-Learning based approaches learn the infection probabilities of individuals based on their features and reported interactions [18], [19]. Testees are chosen by combining graph neural networks and reinforcement learning.

II. DISEASE MODEL

This section describes the disease model we developed (Fig. 1 for a depiction of the state diagram of the model). The model employs individual parameters per person. In the subsections that follow, we describe the personal and group features (i.e., attributes) that affect susceptibility as well as infectiousness.

A. Personal and Group Features

The community graph is a hypergraph $G = (\mathcal{X}, \mathcal{D})$, where $\mathcal{X}$ denotes people and $\mathcal{D}$ denotes subsets of people that meet. We attach static features (i.e., attributes) per individual and per group. Individual features are partitioned to external and internal features. *External features* describe risks associated with the individual's life style. In our implementation, we used the following external features: Number of working places, usage of public transportation, mobility (e.g., number of home visits per month in a boarding school), attendance of special events (e.g., weddings). *Internal features* describe risks associated with the individual's role in the organization (not captured by membership in groups). In our implementation, we used the following internal features: Degree of interaction with people in the organization (e.g., caregiver vs. neighbor) and duration of interactions. *Group features* describe the risk associated with belonging to the group. In our implementation, we used the following group features: Quality of ventilation, density, and duration of interaction.

B. Duration per State

Transitions in the state diagram are stochastic. We attach a random variable T_i per state and per individual, called *state*

Symbol	Description
$spread(x)$	A random variable that determines the spreading degree of x (this value is unknown to the policies)
α	Vector of attribute weights related to the internal risk
β	Vector of attribute weights related to the external risk
γ	Vector of attribute weights related to group risks

duration, that describes the number of days that a person remains in a state before a transition to the next state takes place. The state durations are random variables with the following distributions.

1) $T_0 = 1$. Every day, person x may transition from the states "susceptible" or "recovered" to the state "infected".
2) $T_1 \sim N(1, 0.2)$.
3) $T_2 \sim N(2, 0.3)$.
4) $T_3 \sim N(8, 1)$.
5) $T_4 \sim 2 + Exp(0.5)$.
6) $T_5 \sim N(8, 1)$.

For T_1, the average number of days before detection is 1. This is based on the work of Levine-Tiefenbrun *et al.* [20] that connects the viral load and the error rate. The duration of $T_1 + T_2$ is commonly referred to as the *latent period* of the disease, and the duration of $T_1 + T_2 + T_4$ is referred to as the *incubation period*. For COVID-19, the mean latent period was estimated as 3.3 days, and the mean incubation period was estimated to be 6.8 days [21]. Therefore, we tuned the parameters of T_2 and T_4 random variables to fit these reports, by having mean values of 2 and 4, respectively. For the variance of both T_1 and T_2 we chose numbers much smaller than 1, so that, for a day-by-day simulation, the duration of these stages can be viewed as deterministic. For T_3 and T_5 we based the distribution on the quarantine period [22]. For T_4 we based the numbers published on the official World Health Organization website [23].

C. Transition Probabilities

There are three states with out-degree 2 (i.e., forks): "Susceptible", "recovered", and "contagious pre-symptomatic". The choice of the next state (once the duration of the current state expires) is stochastic and is specified by the probabilities $P(x \text{ infected})$ and $P_{\text{symp}}(x)$.

We use the following notation. Let $x \in \mathcal{X}$ denote a person. Let $ctg(x) \in \{0,1\}$ denote a predicate that indicates whether x is in one of the a contagious state (either contagious pre-contagious, contagious asymptomatic, or contagious symptomatic). Let $quarantined(x) \in \{0,1\}$ denote a predicate that indicates whether x is quarantined. Let $spread(x)$ denote a random variable that determines the spreading degree of x. Let $\delta(x)$ denote the *contagiousness* of person x, defined by:

$$\delta(x) \triangleq spread(x) \cdot ctg(x) \cdot (1 - quarantined(x)). \quad (1)$$

Let $t_{\text{rec}}(x)$ denote the number of days that have elapsed since x entered the "recovered" state (if the state of x has never reached "recovered" yet, then $t_{\text{rec}}(x) = \infty$). Let $f_{\text{pos}}(t) : \mathbb{N} \to [0,1]$ denote a function that quantifies the susceptibility of a

person to an infection as a function of the time from recovery (i.e., $f_{\mathrm{pos}}(t)$ is a non-decreasing function that quantifies the "memory" of the immune system). In our implementation, we ruled out reinfection by defining:

$$f_{\mathrm{pos}}(t) \triangleq \begin{cases} 0 & \text{if } t < \infty \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$

One could easily modify the definition to account for the limited memory of the immune system or new variants that might cause reinfection. The f_{pos} function can be also used to quantify the reduced susceptibility due to a vaccination, however we do not address the effect of vaccination in this study. Let $\lambda(x) \in [0,1]$ denote a *personal susceptibility discount factor* on the risk of person x.

$$\lambda(x) = f_{\mathrm{pos}}(t_{\mathrm{rec}}(x)) \cdot (1 - quarantined(x)) \quad (3)$$

In other words, $\lambda(x) = 0$ if x has already been infected or is quarantined, else it is 1. Person x may be infected by a person outside $\mathcal{X}$, in which case we say that x is *infected externally*, or by a person in $\mathcal{X}$, in which case we say that x is *infected internally*.

The probability that a person x is infected externally is defined as follows. Let $\boldsymbol{z}(x)$ denote the vector of external features of x. Let β denote a vector of weights related to external risks. Define

$$risk_{\mathrm{ext}}(x) \triangleq \langle \boldsymbol{\beta}, \boldsymbol{z}(x) \rangle. \quad (4)$$

The probability of external infection is given by:

$$P(x \text{ infected externally}) \triangleq \lambda(x) \cdot risk_{\mathrm{ext}}(x). \quad (5)$$

The probability that a person x is infected internally is defined as follows. Let $\boldsymbol{\alpha}$ denote a vector of weights related to internal risks. Let $\boldsymbol{v}(x)$ denote the internal features of person x. Let $\boldsymbol{v}(d)$ denote the group features of group d. Let $\boldsymbol{\gamma}$ denote a vector of weights related to group risks. We define the internal risk factors per person and group as follows.

$$risk_{\mathrm{int}}(y) \triangleq \langle \boldsymbol{\alpha}, \boldsymbol{v}(y) \rangle \quad (6)$$

$$risk_{\mathrm{int}}(d) \triangleq \langle \boldsymbol{\gamma}, \boldsymbol{v}(d) \rangle \quad (7)$$

and the probability of person x infecting person y in group d is:

$$P_d(x,y) \triangleq \begin{cases} \delta(x)\lambda(y)risk_{\mathrm{int}}(y)risk_{\mathrm{int}}(d) & \text{if } \{x,y\} \subseteq d \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Note that $\delta(x) > 0$ implies that x is contagious and not quarantined. We assume that infection events by the same individual are independent. Namely, the event that y infects x is independent of the event that y infects x' if $x \neq x'$. In addition, we assume that the events of non-infection of x by y in different groups are independent. Namely, if x and y jointly belong to more than one group, then the events of non-infection in each group are independent. The probability that x infects y is therefore given by

$$P(x \text{ infects } y) = 1 - \prod_{d \in \mathcal{D}} (1 - P_d(x,y)). \quad (9)$$

The probability that x is infected internally is given by

$$P(x \text{ infected internally})$$
$$= 1 - \prod_{y \in \mathcal{X}} (1 - P(y \text{ infects } x)). \quad (10)$$

Since x may be infected internally or externally, we have:

$$P(x \text{ infected}) = 1 - (1 - P(x \text{ infected internally}))$$
$$\times (1 - P(x \text{ infected externally})). \quad (11)$$

The transition of person x from the state "contagious pre-symptomatic" to "contagious-symptomatic" occurs with probability $P_{\mathrm{symp}}(x)$. Based on studies in the literature of the rate of symptomatic people (Oran *et al.* [24] report $40\%-45\%$, and Al-Qahtani *et al.* [25] report 30%–50%), we chose to fix

$$P_{\mathrm{symp}}(x) \triangleq 40\% . \quad (12)$$

for every person $x \in \mathcal{X}$.

Intuitively, higher values of the probability P_{symp} makes the problem easier because symptomatic people are quarantined and do not further spread the contagion. Conversely, lower values of the probability P_{symp} (while all other parameters remain fixed) decreases severe cases and mortality. We, therefore, focus on the reported value of P_{symp}.

D. Tests

PCR-testing is employed to detect infected people [26]. PCR-tests are subject to errors (both false-positive and false-negative). According to Wu *et al.* [27], false-positive rates are 12.16% for recovered people (i.e., state "recovered"). Cohen *et al.* [28] report a false-positive rate of 3.2% for susceptible people (i.e., state "susceptible"). Our simulator adopts these false-positive rates.

False-negative rates depend on the time since infection. Levine-Tiefenbrun *et al.* [20] report that false-negative rates are highest during the first 5 days after exposure (up to 22.8%), after which the false-negative rate reduces to 10.7%. Following this report, we set the false-negative rate in the simulator to:

1) 22.8% in the states "infected", "detectable", and "contagious-pre-symptomatic".
2) 10.7% in the states "contagious-symptomatic" and the "contagious-asymptomatic".

III. TESTING AND QUARANTINE POLICIES

The policies are characterized by combinations of the following rules:

1) Quarantining of symptomatic people (SQ). Rule SQ dictates that every symptomatic person is immediately quarantined. A quarantined person cannot infect others.
2) Daily test samples ($Test = (B,X)$). Rule $Test = (B,X)$ means that, every day, a sample of B people is tested among the un-quarantined people. The sampling algorithm X is one of three sampling algorithms: *Rand*, *RFG*, and *Optimization*. The *Rand* algorithm simply selects a random sample. The *RFG* algorithm

selects the people with the top B weights (described in Section III-D). The *Optimization* algorithm selects B people computed by the optimization procedure described in Section III-E.

3) Testing of quarantined people ($Retest(a, b)$). A quarantined person is tested after $(a - 1)$ days of quarantine. Further tests, if positive, are conducted every b days. After a negative test, a quarantined person is released.

We denote a policy by a three-tuple specifying the chosen rules.

A. No-policy (no $SQ, Test(0, Rand), Retest(\infty, \infty)$)

The no-policy is a trivial policy in which no counter-measures are conducted (no tests and no quarantining). The no-policy is used as a reference of the contagion when no counter-measures are employed. One key feature that can be inferred from the no-policy approach is the herd-immunity. Namely, what is the morbidity ratio after which the epidemic decays in the absence of counter-measures? The herd-immunity threshold obtained from simulations serves as a "sanity-check" of the disease model as well as the simulator.

B. Symptoms-based ($SQ, Test(0, Rand), Retest(4, 2)$)

No tests are applied, but every symptomatic person is immediately quarantined. A quarantined person is tested starting from the 4th day of quarantine every 2 days until found negative. After a negative test, a quarantined person is released. The choice of the low values 4 and 2 allowed the research to focus on the selection of people for testing and isolation, by reducing the impact of recovered people being held in the quarantine. Choosing higher values will inevitably cause longer quarantining of the already recovered people and a reduction in quarantine efficiency (GQE, mPQE). We leave the exploration of a repeat testing for further work, and keep it fixed in our work. We do, however, provide an experimental section where we show the impact of reducing the repeat testing frequency from $Retest(4, 2)$ to $Retest(9, 7)$.

C. Random Sampling ($SQ, Test(B, Rand), Retest(4, 2)$)

Every symptomatic person is quarantined. Every day, a random sample of B people among the unquarantined people is tested. Every positive person is quarantined. A quarantined person is tested starting from the 4th day of quarantine every 2 days until found negative. After a negative test, a quarantined person is released. We refer to this policy, in short, as $Rand(B)$.

D. Risk Factor Greedy ($SQ, Test(B, RFG), Retest(4, 2)$)

The $RFG(B)$ policy is similar to $Rand(B)$ except that the daily sample of B testees is computed greedily (rather than a random sample). The greedy selection of the testees assigns, every day, a weight $w(x)$ to every non-quarantined person $x \in \mathcal{X}$. The B people with the highest weights constitute the sample for that day. We refer to this policy, in short, as $RFG(B)$. We now elaborate on how weights are computed.

The formula for the weight of person x is

$$w(x) \triangleq \tilde{\lambda}(x) \left(risk_{\text{ext}}(x) \right.$$
$$\left. + risk_{\text{int}}(x) \Big(\sum_{d \in \{d' | x \in d'\}} \frac{POS(d, T) + 1}{|d|} \cdot risk_{\text{int}}(d) \Big) \right) . \tag{13}$$

The parameters $risk_{\text{ext}}(x), risk_{\text{int}}(x)$ and $risk_{\text{int}}(d)$ are defined in Section II. The parameter $POS(d, T)$ equals the number of people in group d that were found positive in a test during the past T days. We note that the addition of 1 in the numerator of 13 is to avoid zeroing of the internal risk if no COVID-positive people were detected recently. We also note that the notation $|d|$ stands for the number of people in group d. The parameter $\tilde{\lambda}(x)$ is a discount factor defined by

$$\tilde{\lambda}(x) \triangleq f_{\text{neg}}(t_{\text{neg}}(x)) \cdot f_{\text{pos}}(t_{\text{pos}}(x))$$
$$\times (1 - quarantined(x)) . \tag{14}$$

Note that $\lambda(x)$ is a discount factor based on the ground truth, whereas $\tilde{\lambda}$ is a discount factor computed for the policy based on the information it has access to. The parameters $t_{\text{pos}}(x)$ (resp., $t_{\text{neg}}(x)$) equals the number of days that have elapsed since the last positive test (resp., negative test) of person x. (If no such test result exists, then the parameter is set to ∞). The function $f_{\text{pos}}(t)$ is defined in 2, from which it is implied that a person that was found positive in the past will receive $f_{\text{pos}} = 0$, and a person that has not yet been found positive will have $f_{\text{pos}} = 1$. The function $f_{\text{neg}}(t)$ is defined in 15

$$f_{\text{neg}}(t) \triangleq \begin{cases} 0.1 & \text{if } t = 0 \\ 0.3 & \text{if } t = 1 \\ 0.6 & \text{if } t = 2 \\ 1.0 & \text{if } t \geq 3 \end{cases} . \tag{15}$$

The definition of $f_{\text{neg}}(t)$ implies that as the time since the last negative test grows, the likelihood of of infection increases. The monotone increase of $f_{\text{neg}}(t)$ provides negative incentive to test the same individual day after day. The precise numbers that define $f_{\text{neg}}(t)$ were chosen manually to alleviate the repetitive testing of the riskiest people in the community, however these values can be subject to a rigorous optimization, which we kept out of the scope of this work.

Conceptually, $\tilde{\lambda}$ is 0 for people who are either quarantined or recovered, and it degenerates to $f_{\text{neg}}(t_{\text{neg}}(x))$ for other people (non-quarantined people who have not been known to be infected by COVID-19). The term $f_{\text{neg}}(t_{\text{neg}}(x))$ causes people x who have recently been tested and found negative, to have their $\tilde{\lambda}(x)$ (and consequently $w(x)$) to drop while other candidates receive a higher weight.

The goal of the weight function $w(x)$, is to approximate the probability of x to become infected. The intuition behind the structure of $w(x)$ is that the probability of becoming infected is affected by the fully observable static features of the people and the groups ($risk_{\text{int/ext}}$) as well as by the partially observable information regarding the actual illness state gathered by the policy ($\tilde{\lambda}, POS(d, T)$).

E. OPTIMIZATION $(SQ, Test(B, Optimization), Retest(4, 2))$

The optimization-based approach is identical to $RFG(B)$, except in the way that the daily sample of testees is computed. The daily set of B testees among the non-quarantined people is computed by an optimization algorithm. The optimization algorithm solves a maximization problem in a form of a Linear Program (LP) and rounds the fractions to a $\{0, 1\}$-solution. The LP balances the "coverage" of groups rather than simply choose "riskiest" people.

The unknowns (i.e., variables) of the LP are $s(x) \in [0, 1]$, for every non-quarantined person $x \in \mathcal{X}$. A value $s(x) = 1$ means that x is surely in the sample, whereas, a value $s(x) = 0$ means that x is surely not in the sample.

The LP uses the following coefficients (i.e., knowns):

(i) B - number of testees per day.
(ii) V - the set of unquarantined people.
(iii) E - set of sufficiently heavy groups (see Equation 16).
(iv) $w(x)$ - the weight of x as defined by Equation 13.
(v) $w(d) \triangleq \sum_{x \in d} w(x)$ - weight of group d.

The variable $c(d)$, for group d, measures "coverage" of d and is determined by the values of $s(x)$, for $x \in d$. Namely, $c(d) \in [0, 1]$ describes which portion of the group weight will be removed ("covered") if the currently selected people will be tested. The variable z simply equals $\min_{d \in E} c(d)$, hence maximizing z requires obtaining a fair coverage of all the sufficiently large groups. The set E of sufficiently heavy groups is defined by

$$E \triangleq \left\{ d \in \mathcal{D} \mid w(d) \geq \frac{1}{2|V|} \cdot \sum_{x \in V} w(x) \right\}. \quad (16)$$

We define the set E of sufficiently heavy groups in order to avoid the optimization to be constrained on covering groups whose total weight is below half the average personal weight in the community. Without the restriction of the groups to the set E, the sample selected by the solution of the LP contained people unlikely to become infected, due to their association with small groups.

The formulation of the LP is as follows.

$$
\begin{aligned}
\text{maximize} \quad & z + \frac{1}{100 \cdot |E|} \sum_{d \in E} c(d) \\
\text{subject to} \quad & \frac{1}{w(d)} \cdot \sum_{x \in d} s(x) \cdot w(x) = c(d) \quad \forall d \in E \\
& c(d) \geq z \quad \forall d \in E \\
& \sum_{x \in V} s(x) \leq B \\
& s(x) \in [0, 1] \quad \forall x \in V
\end{aligned}
$$

Note that the second addend in the objective function $1/100 \cdot |E| \sum_{d \in E} c(d)$ is used to provide incentive to utilize the budget B (even if a smaller fractional sample achieves the same value z).

The solution of the LP outputs a fractional-sample, namely, $0 \leq s \leq 1$, for every $x \in \mathcal{X}$. A sample of B people is obtained by randomized rounding as follows:

1) Let $b(x) \in \{0, 1\}$ denote a Bernoulli random variable with probability $s(x)$. Let $B' \triangleq \sum_x b(x)$.
2) If $B' > B$: among the people with $b(x) = 1$, set $b(x) = 0$ for $B' - B$ people with the lowest fractional values $s(x)$.
3) If $B' < B$: among the people with $b(x) = 0$, set $b(x) = 1$ for $B - B'$ people with the highest fractional values $s(x)$.

The sample is taken to be the set $\{x \in \mathcal{X} \mid b(x) = 1\}$.

IV. METRICS

This section describes the quantitative metrics used to asses the performance of the different policies.

Notation: Let $D_x^{(iso)}$ denote the set of days during which person x is isolated. Let $D_x^{(ctg)}$ denote the set of days during which person x is in one of the contagious states. For $p \in \{iso, ctg\}$, let $D^{(p)} \subset \mathbb{N} \times \mathcal{X}$ denote the set

$$D^{(p)} \triangleq \left\{ (x, t) \mid t \in D_x^{(p)} \right\}.$$

Namely, $D^{(iso)}$ denotes the set of person-day pairs (x, t) in which person x is isolated on day t.

1) **Total Morbidity** - total number of infected people is defined by

$$\text{Total Morbidity} \triangleq \left| \{ x \mid D_x^{(ctg)} \neq \emptyset \} \right|. \quad (17)$$

The total morbidity is an important metric for estimating mortality, severe cases of illness, etc.

2) **Peak Morbidity** - maximum number of simultaneously contagious people is defined by

$$\text{Peak Morbidity} \triangleq max_t \left\{ \sum_{x \in \mathcal{X}} \mathbb{I}(t \in D_x^{(ctg)}) \right\}. \quad (18)$$

Keeping the peak morbidity low (i.e., "flattening of the curve") is extremely important to avoid the collapse of the healthcare system. When peak morbidity is bigger than the capacity of hospitals, proper medical treatment is infeasible, thus leading to rejection and prioritization of patients.

3) PQE_x - Personal quarantine efficiency of person x.

$$PQE_x \triangleq \frac{|D_x^{(iso)} \cap D_x^{(ctg)}| + \epsilon}{|D_x^{(iso)} \cup D_x^{(ctg)}| + \epsilon} \quad (19)$$

where $\epsilon \in \mathbb{R}^>$ is arbitrarily small constant (to avoid division by zero). We use $\epsilon = 2.22 \times 10^{-16}$ in our implementation, as it was the smallest representable floating point number on the system we used.

The PQE_x metric is an intersection-over-union metric that captures both the precision of quarantine days of x and the recall of the contagious period of x. This metric is similar to the $F1$ metric in which the denominator assigned half the weight to the symmetric difference.[2] A value of $PQE_x = 1$ implies a perfect decision. Conversely, PQE_x close to 0 indicates a poor choice of quarantine days for person x. There are two causes for mistaken decisions: not quarantining when contagious and quarantining when not contagious. For simplicity, we treat both errors equally (of course, not quarantining when contagious has an adverse effect on controlling the epidemic).

---

[2] $F1 = \frac{|D_x^{(iso)} \cap D_x^{(ctg)}|}{|D_x^{(iso)} \cap D_x^{(ctg)}| + 0.5 \cdot (|D_x^{(iso)} \triangle D_x^{(ctg)}|)}$

4) $\boldsymbol{mPQE}$ - Mean personal quarantine efficiency across all people in $\mathcal{X}$

$$mPQE \triangleq \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} PQE_x \qquad (20)$$

5) $\boldsymbol{GQE}$ - Global quarantine efficiency

$$GQE \triangleq \frac{|D^{(iso)} \cap D^{(ctg)}| + \epsilon}{|D^{(iso)} \cup D^{(ctg)}| + \epsilon} . \qquad (21)$$

As opposed to the $mPQE$ measure, the GQE equally treats the value of all human-days, rather than treating the people equally. Ashcroft *et al.* [29] define a similar metric, referred to as "utility of quarantine", as the ratio between the number of correctly chosen quarantine days and the total number of quarantine days. This metric captures the precision of the quarantine days of chosen by the policy, but does not penalize for a bad recall of the contagious period of the people[3]. GQE metric on the other hand, addresses both the precision and the recall capabilities of the policy at hand.

6) **Policy Efficiency Counters**: Number of human-days during which the person is (i) healthy and free, (ii) contagious and quarantined, (iii) healthy and quarantined, or (iv) contagious and free.

V. EXPERIMENTAL SETUP

A. Simulation Environment

It is hard to theoretically evaluate the ability of an intervention policy to cope with a disease spread due to the complexity of the interactions within the community of people. Therefore, we employ the simulation-driven approach (i.e. empirically-driven, rather than theory driven), which is widely used in the research of epidemic control [30]–[32]. Apart from epidemic control, the simulation an approach is widely used in the research of controlling other complex systems such as power grids [33] or communication networks [34].

The simulator executes a day-by-day simulation of the disease model for all the people in the community-graph. The simulator takes into account the decisions of the policy that may select testees and quarantine people found positive. The simulator and the policy are separate entities, and hence, a policy has only partial information about the ground truth. The input of the policy consists of the community graph, features of the people and the group, budget B of daily tests, results of tests, and the subset of people currently in the state "contagious symptomatic". See Fig. 3 for a high level depiction of the simulation-policy framework. We note that, in addition to the budget of tests B, quarantined people are tested according to the retesting rule (e.g., $Retest(4, 2)$).

---

[3]For example, in a 1-person community, a policy that assigns a total of 2 quarantine days that overlap 2 out of 10 contagious days will result in a 100% utility of quarantine, whereas the GQE and the $mPQE$ metrics will be at 20%
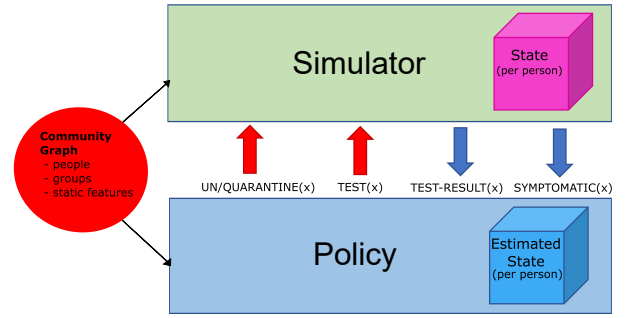

Fig. 3. Simulator-Policy interaction used in our experimentation.

B. Scenario Data

We examine different policies using 4 scenarios. Each scenario is defined by: (i) The community graph, (ii) the features of each individual (i.e., external and internal features) and (iii) the features of each group.

The community graphs were designed to represent a realistic interaction between students and teachers in a school. For example, the sizes of the classrooms were chosen to be the typical capacity of 30–40 students. The interaction of the people was chosen to be relatively local (students interact within a class much more than across classes. We also consider larger community graphs with 7 schools in which interaction between individuals from different schools is either through families or through groups of friends. Larger communities can be viewed as having another level of hierarchy in their structure.

Individual features included 4 external features (having multiple jobs, usage of public transportation, frequency of appearance in the community and extremely infection-prone activity outside the community) with the respective attribute weights $\boldsymbol{\beta} = (0.1, 0.1, 0.1, 0.7)$. In addition, 3 internal features (Highly interacting person, degree of physical presence of a person, and other infectious-prone activity of a person) were equally weighted by the weight vector $\boldsymbol{\alpha} = (1/3, 1/3, 1/3)$. The spreading degree ($spread(x)$) of people were assigned randomly with 5% of the people having a high value of 0.5 (being super-spreaders) and the rest of the people to values below 0.25. Group features consisted of 5 features (social distancing negligence, poor room ventilation, crowdedness, duration of interaction in the group, and other highly infectious-prone activity), and were weighted by $\boldsymbol{\gamma} = (0.06, 0.06, 0.06, 0.412, 0.412)$. The risk factors and the weights are generic and were tuned to achieve a significant contagion phenomenon in the simulated communities, as will be presented in Figs. 4(a) and 5(a). Determining more realistic weights for these risk factors, as well as discovering novel risk factors, can be of a great contribution to the simulation fidelity. We encourage a research in this direction by providing an open-sourced, easily configurable simulator infrastructure, in which the risk factor weights can be adjusted without reprogramming, but rather using a configuration file.

We now elaborate on the 4 scenarios in more details.

1) Single-School: A single school consists of 150 people, with 4 classes of students approximately of the same size

(i.e., 37 students), where each class is isolated from the rest. Besides students, there are 7 teachers who visit all the classes. In addition to the teachers and the students, there are 3 managers that only interact between themselves, and there is also a cleaner that interacts with everyone in the organization. Average number of people a single person is in a direct contact with: 54.64.

2) Single-School-Mobility: This scenario is the same as the Single-School scenario with the modification that between every two adjacent classes there are 5 pairs of interactions. This modification represents a less strict social distancing rule. Average number of people a single person is in a direct contact with: 62.9.

3) Multischool-Families: A community of 1000 people, consisting of 7 schools of the following form: 4 "Single-School"-type schools and 3 "Single-School-Mobility"-type schools. Schools are indexed 1 through 7, and two schools are *adjacent* if the absolute difference of the indexes is 1. In addition to the schools, there are 300 small groups referred to as *families* with an average of 2.3 children per family. The children of a family attend either a single school or two adjacent schools. Average number of people a single person is in a direct contact with: 57.

4) Multischool-Friends: Same scenario as Multischool-Families scenario, however, instead of the 300 small families, there are 13 groups of *friends*. Groups of friends consist of 10 groups of students, each consisting of approximately 50 students drawn at random from all schools; and 3 groups of teachers with $5, 7$ and 12 teachers each. The teachers are chosen at random from all the teachers in all of the schools. Average number of people a single person is in a direct contact with: 76.

C. Experiments

In every experiment, we evaluate the 5 policies that were presented in section III. The policies are referred to as: No-policy, symptom-based, $Rand(B)$, $RFG(B)$ and $Optimization(B)$. The policies that apply tests are examined with different test budgets.

Initially, on day-1, all people are in the state "susceptible". From this initial configuration, a day-by-day simulation runs for a total duration of either 150 days for the small scenarios (Single-School scenarios) and 300 days for larger scenarios (Multischool scenarios).

Each scenario-policy-budget combination runs in a separate simulation, for which we record the metrics described in section IV. Due to the stochastic nature of the simulation, we repeat the simulation to achieve a statistical significance of the metrics. Specifically, for the small scenarios (Single-School scenarios) each simulation is repeated 50 times, and for the larger scenarios (Multischool scenarios) the simulation is repeated 10 times. We note that the Multischool scenarios have a regular structure (i.e., 7 schools consisting of 4 classes each). Thus, a simulation of a Multischool scenario is actually a simulation of 7 schools with some interactions between the schools. Indeed, simulations of Multischool scenarios exhibit smaller variance in the measured metrics across the 10 runs

of the large scenarios compared to 50 runs of the smaller scenarios (i.e., Single-School).. All the recorded metrics are averaged across the repeating runs, and we report their mean values as well as standard deviations.

In order to make our results more realistic, we added two experiments: (1) Measure the effect of errors in PCR-tests, and (2) measure the effect of retesting of quarantined people. Both these experiments were chosen in order to test the robustness of our optimization based policy. To show the impact of test errors, we introduced an error probability for tests into our simulation (based on numbers from the literature, as discussed in Section II-D), and compared the measures for the Multischool-Families scenario. Effects of retesting of quarantined people was simulated using the Multischool-Families scenario. The default testing policy we chose for quarantined person was $Retest(4, 2)$ (i.e., starting from day 4 of quarantine, the person is tested once every 2 days until found negative). The less strict retesting rule we chose is $Retest(9, 7)$. We dedicate section VI-D for the comparison between these two rules, and the impact they have on controlling the epidemic.

VI. RESULTS

A. Morbidity-over-Time Experiment

To gain an intuition regarding the effect of the policies on the morbidity, we present plots that show the daily number of infected people. To this end, we simulate all 5 policies in a Multischool-Families scenario with a budget of $B = 40$ daily tests, no test errors, and with a re-testing rule of $Retest(4, 2)$. The simulations indicate that the more sophisticated policies exhibit a reduction in the total number of infected people (Fig. 4) as well as in the peak of morbidity (Fig. 5). Compared to the symptom-based policy, the optimized policy $Optimization(40)$ reduced the total morbidity roughly by half, and the peak morbidity by roughly 40%.

B. Policy Comparison

In this section we compare the four metrics (total morbidity, peak morbidity, GQE, mPQE) with respect to the 5 policies. Our goal is to learn which policy is better in terms of morbidity reduction and in terms of quarantine efficiency.

Fig. 6 shows that, to achieve the same bound on total morbidity, the OPTIMIZATION-based policy requires roughly $B/4$ of the daily budget of tests compared to both $Rand(B)$ and the $RFG(B)$ policies (for small population refer to a budget $B \approx 10$, for large population refer to a budget $B \approx 15$).

The reduction of the peak morbidity was also consistently better for the OPTIMIZATION strategy in all the scenarios. As can be seen in Fig. 7, the peak number of infected people is reduced by $5\%-50\%$ by the OPTIMIZATION policy (compared to the other policies).

The reduction in the morbidity is less profound in the Single-School-Mobility scenario rather than in the Single-School scenario and similarly for Multischool-Friends compared to Multischool-Families. Indeed, the connectivity of the community-graph for Single-School-Mobility is higher than

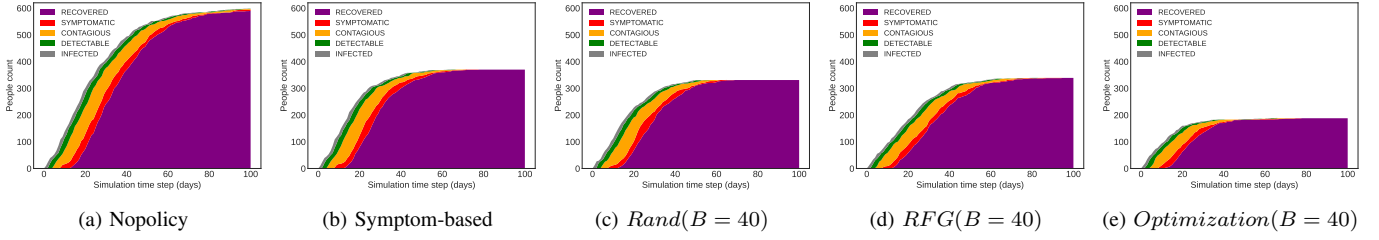(a) Nopolicy (b) Symptom-based (c) $Rand(B = 40)$ (d) $RFG(B = 40)$ (e) $Optimization(B = 40)$

Fig. 4. Cumulative Morbidity over time. Each plot (a)-(e) describes a simulation of a different policy applied in a Multischool-Families scenario (1000 people) using a daily test budget of 40. Note that the purple portion of the plot height at each time step describes the number of recovered people. **The OPTIMIZATION policy decreases the total morbidity roughly by half compared to the Symptom-based policy.**

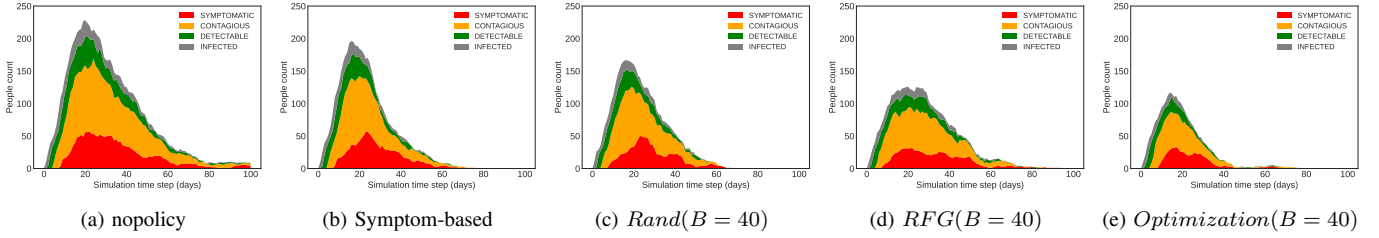(a) nopolicy (b) Symptom-based (c) $Rand(B = 40)$ (d) $RFG(B = 40)$ (e) $Optimization(B = 40)$

Fig. 5. Daily Morbidity evolution over time (flattening the curve). Each plot (a)-(e) describes a simulation of a different policy applied in a Multischool-Families scenario (1000 people) using a daily test budget of 40. Note the different Y-axis scale of the different plots. **The OPTIMIZATION policy reduces the peak by roughly 40% compared to the Symptom-based policy.**

(a) Single-School (b) Single-School-Mobility

(a) Single-School (b) Single-School-Mobility

(c) Multischool-Families (d) Multischool-Friends

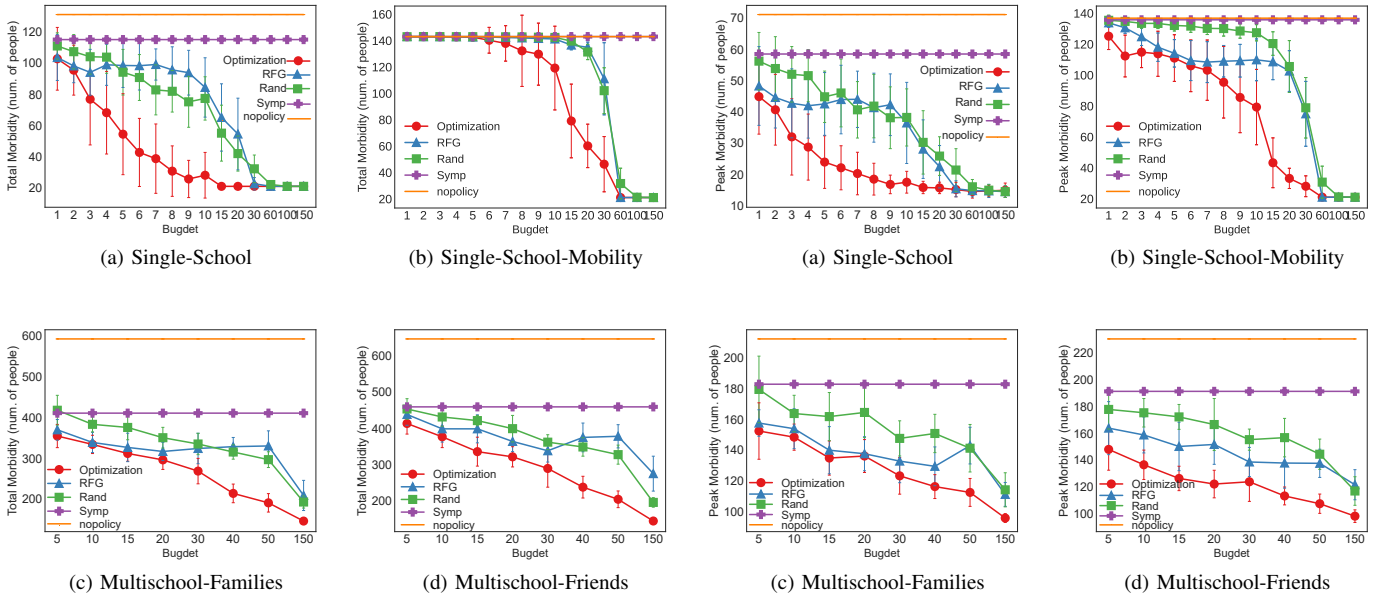(c) Multischool-Families (d) Multischool-Friends

Fig. 6. **Total Morbidity** as a function of the daily test budget, in 4 different scenarios (a)-(d) and 5 policies (nopolicy, Symptom-based, RAND, RFG, OPTIMIZATION). No test errors were simulated. We applied a re-testing rule of $Retest(4, 2)$ for all the policies (except for no-policy). The no-policy graph indicates that herd immunity is reached with roughly $80 - 65\%$ morbidity depending on the population size and community-graph.

Fig. 7. **Peak Morbidity** as a function of the daily test budget, in 4 different scenarios (a)-(d) and 5 policies (nopolicy, Symptom-based, RAND, RFG, OPTIMIZATION). No test errors were simulated. We applied a retesting rule of $Retest(4, 2)$ for all the policies (except no-policy).

the connectivity-graph of Single-School. The same observation holds for the Multischool-Friends community-graph (which introduces large groups of friends that span across all the schools in the community) compared to the community-graph of Multischool-Families. A stronger connectivity of the community graph results in (i) requirement of a larger test-

budget in order to achieve a control over the disease (ii) lower personal quarantine efficiency.

It is easy to achieve low morbidity by increasing the number of quarantined people. We quantify the efficiency of quarantining using the intersection-over-union metrics. The GQE metric presented on Fig. 8 represents the global efficiency of the quarantine decisions, and these results show

(a) Single-School

(b) Single-School-Mobility

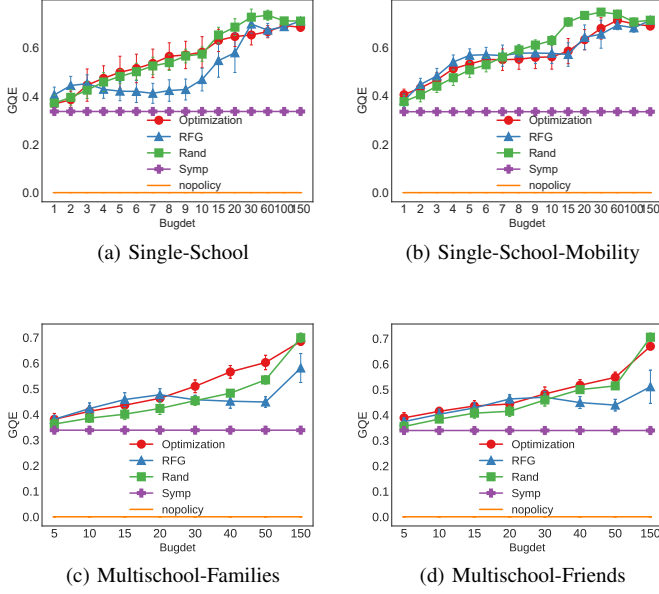(c) Multischool-Families

(d) Multischool-Friends

Fig. 8. **GQE** as a function of the daily test budget, in 4 different scenarios (a)-(d) and 5 policies (nopolicy, Symptom-based, RAND, RFG, OPTIMIZATION). No test errors were simulated, and the re-testing rule was $Retest(4, 2)$ for all the policies.

(a) Single-School

(b) Single-School-Mobility

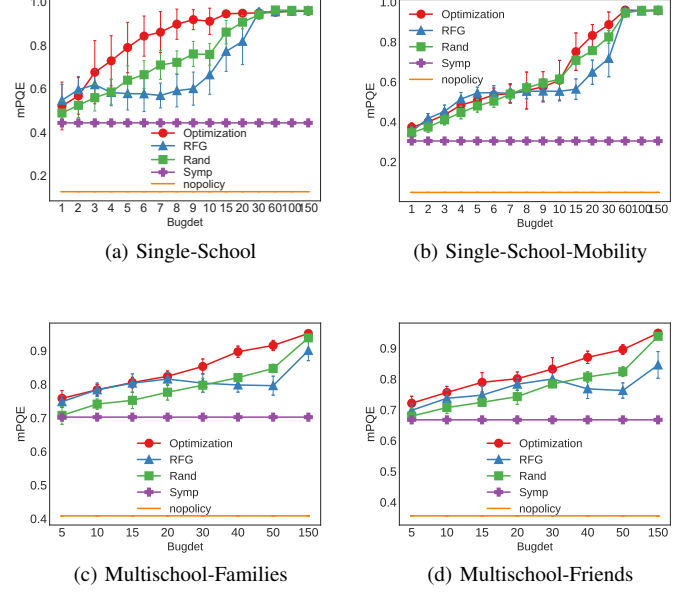(c) Multischool-Families

(d) Multischool-Friends

Fig. 9. **mPQE** as a function of the daily test budget, in 4 different scenarios (a)-(d) and 5 policies (nopolicy, Symptom-based, RAND, RFG, OPTIMIZATION). No test errors were simulated, and the rule of re-testing was $Retest(4, 2)$ for all the policies.
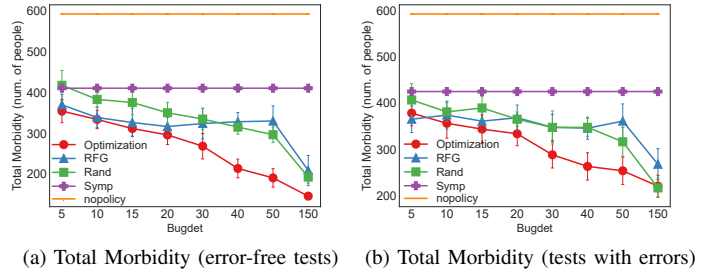
that RAND and OPTIMIZATION policies are similarly efficient. The OPTIMIZATION and the RFG policies are especially prevalent in lower test budgets, whereas the RAND policy shows a competitive efficiency when using larger budgets ($B > 10$) in smaller communities (150 people). This is explained by the bias the OPTIMIZATION has towards the more risky people, whereas the RAND policy uniformly samples the people regardless of their risk factors. When the budget is large enough the uniform sampling pays off as it helps identifying the infected people that were less likely to become infected.

The mPQE metric, as shown on Fig. 9, is the highest for the OPTIMIZATION policy. The mPQE metric represents a fair efficiency metric with respect to the individuals. The mPQE metric values are generally higher than the the values observed for GQE due to the people that never became contagious, and hence they were easily not quarantined and contributed a perfect score of 1.0 to the mean value. The fact that the OPTIMIZATION policy leads both in morbidity reduction and global and personal efficiency indicates the superiority of $Optimization$ over the other policies. Namely, the OPTIMIZATION policy manages to reduce the morbidity while maintaining individual fairness and community-level efficiency.

C. Tests with Errors

We present the results of Multischool-Families scenario simulation with 5 different policies in presence of realistic test errors. The simulation followed the error probabilities as defined in section II-D. Fig. 10 shows the total morbidity comparison between the simulation of error-free testing and testing with errors. The plots demonstrate that the superiority

(a) Total Morbidity (error-free tests) (b) Total Morbidity (tests with errors)

Fig. 10. **Impact of tests with errors** - on the total morbidity, for scenario **Multischool-Families**.

of $Optimization$ policy is sustained in the presence of testing errors.

D. Repeated Testing Rules

The releasing from the quarantine in our previous experiments was based on the $Retest(4, 2)$ re-testing rule. In this section we compare the total morbidity and the global quarantine efficiency (GQE) achieved by the $Retest(4, 2)$ rule against a $Retest(9, 2)$ rule.

Results presented in Fig. 12 show that, as expected, reducing re-testing of quarantined does not affect the morbidity. However, it does degrade the efficiency, as can be seen from the difference between Figs. 12(c) and 12(d). This can also be explicitly viewed in a form of increased isolation days incurred for the healthy people (Fig. 11).

E. Running Time Complexity

The complexity of the running time can be decomposed into two components: (i) The simulation overhead which accounts

<div style="columns:2">

	OPTIMIZATION	RFG	RAND	Symptom-based	nopolicy
Single-School (150 people)	0.08	0.05	0.04	0.03	0.02
Multischool (1000 people)	1.99	1.70	1.64	1.39	1.18

for infection probability computation for each person and the maintenance of the illness states; (ii) the policy overhead, which accounts for the selection of candidates for testing from the non-quarantined people and the performing the repeat testing of the quarantined people.

The Simulation overhead is dominated by the infection probability computation which is $\mathcal{O}(|\mathcal{X}| \cdot \Delta)$, where Δ denotes the average number of people that each individual interacts with.

The policy overhead for the RAND policy is $\mathcal{O}(|\mathcal{X}|)$ since only a random choice of candidates for testing is required. The policy overhead for RFG policy increases this to $\mathcal{O}(|\mathcal{X}| \cdot (|\mathcal{D}| + \log |\mathcal{X}|))$ due to weights computation (13) and sorting. The complexity of the OPTIMIZATION policy, which relies on the same weight computation as RFG, further increases due to the need to solve a linear program. The running time of the simulation is reported in Table III.

F. Community Structure Challenges

Two key parameters of a community which have a high impact on the spread of the infection are average distance between pairs of people over the community graph (average number of hops) and the small-set expansion-factor (i.e., for a small random set U of nodes, $|U|$ plus the number of neighbors of U divided by $|U|$). The lower the average distance and the higher the expansion factor compared to the size of the community - the faster the contagion.

Table IV summarizes these parameters for the 4 community-graphs that we considered. The expansion factors of each community with respect to group sizes $|U| \in \{5, 10, 20, 50\}$ were computed by taking the average of expansion factors of 100 randomly selected subgroups U. The 2 smaller communities (i.e. the 150-people Single-School and Single-School-Mobility) exhibit a much lower average distance between two random people, and their expansion factor suggests that a random set of 5 people has an immediate connection to approximately 125 people (out of the total 150). In this case, the contagion happens extremely fast, rendering various policies less effective. On the other hand, 2 larger communities (i.e. the 1000-people Multischool-*) have a higher average distance and they do not expand small subsets to almost the entire population size. Therefore, the contagion becomes more manageable in the Multischool-* communities and this explains the overall more profound contribution of the proposed testing and isolation policies, as was observed in this section.

We remark that graphs such as preferential attachment [35] graphs and stochastic block graphs have high expansion and small average distance. Indeed, the main contribution of social distancing or limiting meetings to small numbers of people is in reducing the average distance and expansion factor parameters.

	Diameter	Average min distance	Expansion factor (group size)			
			(5)	(10)	(20)	(50)
Single-School	2	1.6467	25.8	14.4	7.4	2.99
Single-School-Mobility	2	1.5911	26.3	14.41	7.4	2.99
Multischool-Families	5	3.8352	49.47	41.99	31.96	18.11
Multischool-Friends	4	2.3109	61.67	51.7	37.01	18.66

(a) Multischool-Families under $Retest(4, 2)$ rule

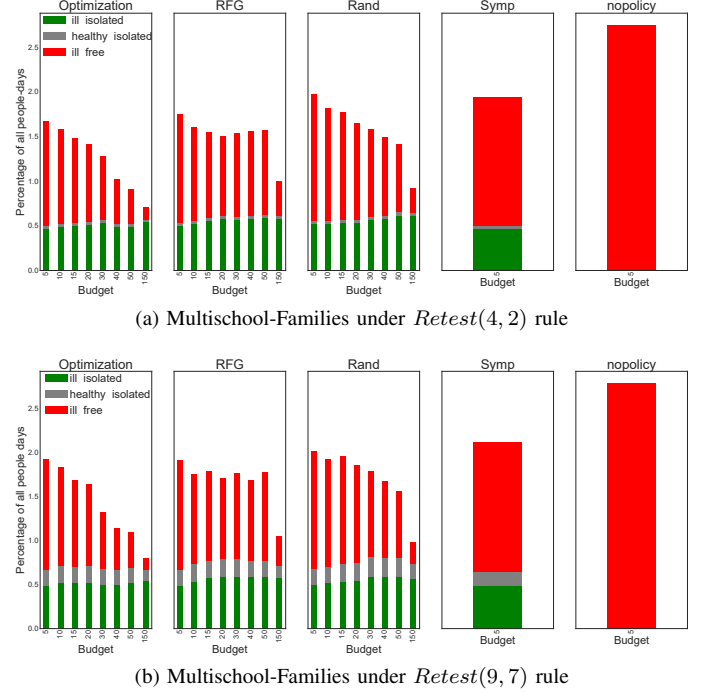(b) Multischool-Families under $Retest(9, 7)$ rule

Fig. 11. **Policy Efficiency Counters** metrics under a $Retest(4, 2)$ rule (11a) and under a $Retest(9, 7)$ rule (11b), both in a error free simulation of **"Multischool-Families"** scenario.

VII. CONCLUSION

We consider a realistic simulation of COVID-19 contagion over a community-graph. With a limited daily budget of PCR-tests, it is beneficial to rely on the structure of the community-graph to select testees. This has been consistently observed throughout our quantitative comparison between policies which demonstrated that the $Optimization$ policy is the most effective in reducing peak/total morbidity as well as maximizing quarantine efficiency. Namely, the policies that disregard the community structure (nopolicy and Symptom-based, which do not perform tests; and $Rand$ which applies tests randomly regardless of the grouped structure of the community) performed worse both in terms of quarantine efficiency and in terms of morbidity reduction. The success of the RFG and $Optimization$ stems from the community structure-aware weight functions, according to which the policies rank the people for testing. Further superiority of the $Optimization$ policy over the $Rand$ is achieved thanks to the constraint that forces a fair coverage (via testing) of all the groups in the community. The fair coverage mitigates outbreaks across the community.

In addition, we show that our conclusions are robust to realistic PCR testing errors, and we describe the relation

</div>

(a) Total Morbidity, $Retest(4, 2)$

(b) Total Morbidity, $Retest(9, 7)$
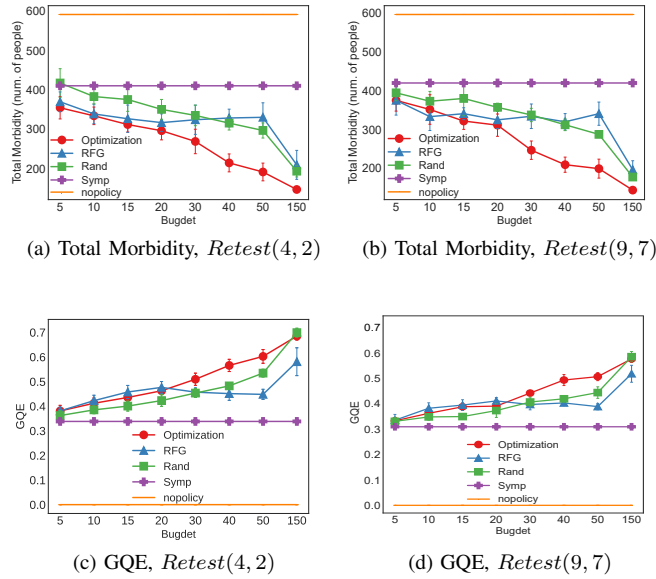

(c) GQE, $Retest(4, 2)$

(d) GQE, $Retest(9, 7)$

Fig. 12. **Re-test strategies of the quarantined people** testing every 2 days after a 4 day quarantine ($Retest(4, 2)$) against weekly tests after a 9-day quarantine ($Retest(9, 7)$). Scenario **"Multischool-Families"**, without test errors.

between the effectiveness of the testing policies and the connectivity metrics (average distance and expansion factor) of the communities at hand.

VIII. DISCUSSION

During peak morbidity in Israel, the daily budget of tests reached $0.5 - 2\%$ of the population. It seems that exceeding this budget is not feasible for large populations. The community-graph models the interactions between people, and therefore, can help perform investigations of "chains-of-contagion". A common policy is to test or quarantine people that have interacted with positive people. We note that such a policy may require too many tests (i.e., testing much more than 1% of the population per day). The proposed policies (RFG, OPTIMIZATION) carry out such epidemiological investigations indirectly. Indeed, the presence of a positive individual in a group increases the weight of the group members, and hence, the likelihood of these individuals to be selected for testing.

We suggest two directions for further research based on learning. The first direction deals with the problem that the accumulation of data regarding the features of individuals and groups is costly, time-consuming, and error-prone. Further work should consider methods for extracting such features from available sources such as location from phones, registration upon entry to groups, and filling of online forms. Determining the weights of such features can be a great contribution to the overall fidelity of the simulator environment, and worth investigating as well. The second direction suggests to employ our simulator as an environment simulation for training a reinforcement-learning agent. Our hope is that such an agent can outperform the OPTIMIZATION-policy.

REFERENCES

[1] B. Tang *et al.*, "The Effectiveness of Quarantine and Isolation Determine the Trend of the COVID-19 Epidemics in the Final Phase of the Current Outbreak in China," *International J. Infectious Diseases*, vol. 95, pp. 288–293, 2020.

[2] S. E. F. Yong *et al.*, "Connecting Clusters of COVID-19: An Epidemiological and Serological Investigation," *The Lancet Infectious Diseases*, vol. 20, no. 7, pp. 809–815, 2020.

[3] N. Ahmed *et al.*, "A Survey of COVID-19 Contact Tracing Apps," *IEEE Access*, vol. 8, pp. 134577–134601, 2020.

[4] P. Poletti *et al.*, "Association of Age With Likelihood of Developing Symptoms and Critical Disease Among Close Contacts Exposed to Patients With Confirmed SARS-CoV-2 Infection in Italy," *JAMA Netw. Open*, vol. 4, no. 3, pp. e211085–e211085, Mar. 2021.

[5] H. W. Hethcote, "The Mathematics of Infectious Diseases," *SIAM review*, vol. 42, no. 4, pp. 599–653, 2000.

[6] L. Peng, W. Yang, D. Zhang, C. Zhuge, and L. Hong, "Epidemic Analysis of COVID-19 in China by Dynamical Modeling," *arXiv preprint arXiv:2002.06563*, 2020.

[7] I. Z. Kiss *et al.*, "Mathematics of Epidemics on Networks," *Cham: Springer*, vol. 598, 2017.

[8] J. M. Carcione, J. E. Santos, C. Bagaini, and J. Ba, "A simulation of a covid-19 epidemic based on a deterministic seir model," *Frontiers Public Health*, vol. 8, p. 230, 2020.

[9] P. Silva, "Agent-based Simulation of COVID-19 Health and Economical effects," 2020. [Online]. Available: https://towardsdatascience.com/agent-based-simulation-of-covid-19-health-and-economical-effects-6aa4ae0ff397

[10] H. Stevens, "Why Outbreaks Like Coronavirus Spread Exponentially, and How to "Flatten the Curve"," 2020. [Online]. Available: https://www.washingtonpost.com/graphics/2020/world/corona-simulator/

[11] K. Simler, "Outbreak," 2020. [Online]. Available: https://meltingasphalt.com/interactive/outbreak/

[12] Y. S. By Joe Fox and A. E. Fe, "How Epidemics Like COVID-19 End," 2020. [Online]. Available: https://www.washingtonpost.com/graphics/2020/health/coronavirus-how-epidemics-spread-and-end/

[13] M. E. Newman, "Spread of Epidemic Disease on Networks," *Physical review E*, vol. 66, no. 1, p. 016128, 2002.

[14] Á. Bodó, G. Y. Katona, and P. L. Simon, "SIS Epidemic Propagation on Hypergraphs," *Bulletin Mathematical Biology*, vol. 78, no. 4, pp. 713–735, 2016.

[15] I. Iacopini, G. Petri, A. Barrat, and V. Latora, "Simplicial Models of Social Contagion," *Nature Commun.*, vol. 10, no. 1, pp. 1–9, 2019.

[16] G. F. de Arruda, G. Petri, and Y. Moreno, "Social Contagion Models on Hypergraphs," *Physical Review Research*, vol. 2, no. 2, p. 023032, 2020.

[17] Y. Deng, S. Shen, and Y. Vorobeychik, "Optimization Methods for Decision Making in Disease Prevention and Epidemic Control," *Mathematical Biosciences*, vol. 246, no. 1, pp. 213–227, 2013.

[18] E. A. Meirom, H. Maron, S. Mannor, and G. Chechik, "How to Stop Epidemics: Controlling Graph Dynamics with Reinforcement Learning and Graph Neural Networks," *arXiv preprint arXiv:2010.05313*, 2020.

[19] V. Kompella *et al.*, "Reinforcement Learning for Optimization of COVID-19 Mitigation Policies," *arXiv preprint arXiv:2010.10560*, 2020.

[20] M. Levine-Tiefenbrun *et al.*, "Association of COVID-19 RT-qPCR Test False-Negative Rate with Patient Age, Sex and Time Since Diagnosis," *medRxiv*, 2020. [Online]. Available: https://www.medrxiv.org/content/early/2020/11/17/2020.10.30.20222935

[21] S. Zhao *et al.*, "Estimating the Generation Interval and Inferring the Latent Period of Covid-19 from the Contact Tracing Data," *Epidemics*, vol. 36, p. 100482, 2021.

[22] "World health organization." [Online]. Available: https://www.cdc.gov/coronavirus/2019-ncov/if-you-are-sick/quarantine.html

[23] "World health organization." [Online]. Available: https://www.who.int/emergencies/diseases /novel-coronavirus-2019/question-and-answers-hub/q-a-detail/coronavirus-disease-covid-19

[24] D. P. Oran and E. J. Topol, "Prevalence of Asymptomatic SARS-CoV-2 Infection: A Narrative Review," *Annals Internal Medicine*, vol. 173, no. 5, pp. 362–367, 2020.

[25] M. Al-Qahtani *et al.*, "The Prevalence of Asymptomatic and Symptomatic COVID-19 in A Cohort of Quarantined Subjects," *International J. Infectious Diseases*, vol. 102, pp. 285–288, 2021.

[26] J. M. Bartlett and D. Stirling, "A Short History of the Polymerase Chain Reaction," *PCR protocols*. Springer, 2003, pp. 3–6.

[27] X. Wu *et al.*, "A Follow-up Study Shows that Recovered Patients with Re-positive PCR Test in Wuhan May Mot be Infectious," *BMC medicine*, vol. 19, no. 1, pp. 1–7, 2021.

[28] A. N. Cohen and B. Kessel, "False Positives in Reverse Transcription PCR Testing for SARS-CoV-2," *medRxiv*, 2020. [Online]. Available: https://www.medrxiv.org/content/early/2020/05/20/2020.04.26.20080911

[29] P. Ashcroft, S. Lehtinen, D. C. Angst, N. Low, and S. Bonhoeffer, "Quantifying the Impact of Quarantine Duration on COVID-19 Transmission," *Elife*, vol. 10, p. e63704, 2021.

[30] P. Amar, "Pandæsim: An Epidemic Spreading Stochastic Simulator," *Biology*, vol. 9, no. 9, p. 299, 2020.

[31] F. D. Sahneh, A. Vajdi, H. Shakeri, F. Fan, and C. Scoglio, "GEMFsim: A Stochastic Simulator for the Generalized Epidemic Modeling Framework," *J. Comput. Sci.*, vol. 22, pp. 36–44, 2017.

[32] A. Kuzdeuov *et al.*, "A Network-based Stochastic Epidemic Simulator: Controlling COVID-19 with Region-specific Policies," *IEEE J. Biomed. Health Inform.*, vol. 24, no. 10, pp. 2743–2754, 2020.

[33] P. H. Nardelli *et al.*, "Models for the Modern Power Grid," *European Physical J. Special Topics*, vol. 223, no. 12, pp. 2423–2437, 2014.

[34] A. Varga and R. Hornig, "An Overview of the OMNeT++ Simulation Environment," in *Proc. ACM Simutools*, 2008, pp. 1–10.

[35] R. Albert and A.-L. Barabási, "Statistical Mechanics of Complex Networks," *Reviews Modern Physics*, vol. 74, no. 1, p. 47, 2002.

Moni Shahar received his B.Sc. degree in 1993 in Mathematics and Physics. He received his M.Sc. in Computer Science in 1997 and Ph.D. in Electrical Engineering in 2005, all from Tel-Aviv University. In the last 25 years he has held a variety of algorithmic leadership positions in the high-tech industry. His work spans a variety of areas including machine vision, NLP, algorithmic finance, machine learning and applied physics. In parallel to his industrial job, he teaches in the Department of Statistics and has 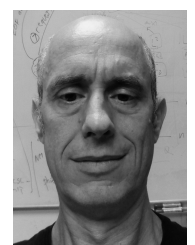published over 20 scientific papers. Currently he is the scientific manager of the AI and Data Science Center of Tel-Aviv University (TAD).

Konstantin Berestizshevsky received his B.Sc. in Electrical Engineering and Computer Science in 2014 and the M.Sc. in Electrical Engineering in 2016, both from the Tel-Aviv University. He is currently working on his Ph.D. in Tel Aviv University. His main interests are efficient Deep Learning models and applications of Deep Learning in various fields including recommender systems, computer vision, smart power grid and analytical chemistry.

Koffi-Emmanuel Sadzi is a B.Sc. student in Electrical and Electronics Engineering at Tel-Aviv University. His fields of interest are statistics, computer simulations, communications and signal processing.

Guy Even received his B.Sc. degree in 1988 in Mathematics and Computer Science from the Hebrew University. Received his M.Sc. and D.Sc. in Computer Science from the Technion in 1991 and 1994, respectively. Dr. Even spent his Post-Doctorate in the University of the Saarland during 1995-1997. Since 1997, he is a faculty member in the School of Electrical Engineering in Tel-Aviv University. Prof. Even is interested in algorithms and their applications in various fields. He has published papers on the theory of VLSI, approximation algorithms, computer arithmetic, online algorithms, frequency assignment, scheduling, packet-routing, linear-programming decoding of LDPC codes, and data structures. He is on the editorial board of *Theory of Computing Systems*. Together with Moti Medina, Dr. Even wrote a digital hardware textbook titled *Digital Logic Design: A Rigorous Approach* published in 2012 by Cambridge University Press.