Location Contact Tracing: Penetration, Privacy, Position, and Performance

LUCA BEDOGNI, University of Modena and Reggio Emilia, Italy
FEDERICO MONTORI, University of Bologna, Italy
FLORA SALIM, University of New South Wales, Australia

The 2020 COVID-19 pandemic radically changed the world and how people interact, move, and behave. Following a lockdown that was imposed worldwide, although with different timing, Mobile Contact Tracing Apps (MCTAs) were proposed to digitally trace contacts between individuals while gradually releasing mobility constraints mandated to contain the spread of disease. General concern for privacy regarding the use of GPS data shifted the efforts toward distributed applications, which use Bluetooth technology to trace proximity and potential infections. Nonetheless, GPS data would help more health operators to understand where hotbeds are and to what extent the spread is progressing and at what pace. In addition to these issues, in this work we take a closer look at the major pillars of MCTA: Penetration, Privacy, Position, and Performance. We focus on (i) how the penetration rate affects the ability of a tracing application to work; (ii) the proposal of a novel method of tracing, which builds on the GPS technology; (iii) how the position of infections is beneficial to rapidly reduce the infection; and (iv) the discussion of the effects of such paradigms in different scenarios.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing**; • **Security and privacy** → *Privacy protections*; • **Computing methodologies** → Simulation evaluation;

Additional Key Words and Phrases: Mobile Contact Tracing, privacy, performance evaluation

ACM Reference format:
Luca Bedogni, Federico Montori, and Flora Salim. 2022. Location Contact Tracing: Penetration, Privacy, Position, and Performance. *Digit. Gov. Res.: Pract.* 3, 3, Article 18 (October 2022), 13 pages.
https://doi.org/10.1145/3532669

1 INTRODUCTION

In 2020, the world was hit by the COVID-19 pandemic, which caused a large portion of the world population to quarantine, a massive amount of deaths, many companies and small businesses to shut down, and, thus, immense damage to the worldwide economy. There is a tremendous amount of work in progress in several countries seeking key enablers for recovery in order to prevent smaller outbreaks of contagion to proliferate in the recovery phase. Human contact tracing has been used extensively in history in order to understand the recent storyline of people with whom an infected person has been in contact and how. Subsequently, authorities would get in contact with such people and warn them or put them into isolation for a period in accordance with guidelines on the illness. This method found its application in epidemic outbreaks of illnesses such as human immunodeficiency

Authors' addresses: L. Bedogni, University of Modena and Reggio Emilia, Via Campi 213/A, 41125 Modena, Italy Italy; email: luca.bedogni@unimore.it; F. Montori, University of Bologna, Mura Anteo Zamboni 7, 40126 Bologna, Italy; email: federico.montori2@unibo.it; F. Salim, University of New South Wales Sydney, Engineering Rd, Kensington NSW 2052, Australia; email: flora.salim@unsw.edu.au.

2639-0175/2022/10-ART18 $15.00
https://doi.org/10.1145/3532669

virus, Ebola, smallpox, and Severe Acute Respiratory Syndrome (SARS), the latter of which is closely related to the COVID-19 pandemic. Due to the unprecedented consequences of the COVID-19 outbreak, human contact tracing is no longer sufficient, as it does not scale up to the needs of entire countries that experienced a pandemic crisis. In these critical times, researchers leverage mobile devices as a key enabler for reopening, with particular focus on digitizing contact tracing practices into Mobile Contact Tracing Applications (MCTAs). MCTAs leverage mobile personal devices to host applications able to record contacts with other individuals in possession of the same technology. In the case of COVID-19, being in "contact" means a vicinity in space and time. The technologies in use are, for the vast majority, Bluetooth and GPS. In the case of Bluetooth, spatial closeness is recorded user-to-user: a temporary token is assigned to each device and, when a contact occurs, the ID of the neighbor is recorded in a phone log. It has the advantage of better preserving location privacy [6]. However, it fails to identify spatially possible hotbeds. Moreover, it can cause potentially detectable infections to be ignored (false negatives) or too many noninfected people to be notified of infection (false positives) [9]. In the case of GPS, users periodically share their location and whenever a user tests positive, all other users who have been in the same place at the same time are notified. This approach can indeed identify possible sources of outbreaks; however, it requires users to share their location, hence, raising potential privacy issues. That said, the use of GPS would help in understanding disease spread, as it would make possible the identification of areas of interest in which the disease is more diffuse. Regardless of the approach used, the penetration of the application is a major concern, as people need to actively install it and maintain the related interface (Bluetooth, GPS) constantly active, even if it is not their previous habit. In fact, all available MCTAs to date are individual mobile apps that need an independent installation. A second, yet not less important, concern is the privacy preservation of the users, as releasing data at this granularity can unveil sensitive information with possibly critical consequences. A third concern involves the need of health and social institutions to gain information on where the contagion spreads the most due to inadequate countermeasures in certain locations. Finally, a fourth concern is that these applications should be used and tuned in a way that false positives and false negatives are avoided as much as possible. We identify these concerns as the four major pillars of MCTA; we will consequently refer to them as the four Ps: *Penetration*, *Privacy*, *Position*, and *Performance*. In this work, we focus our analysis on understanding how each of these pillars affect MCTAs and how different proposals compare with each other.

With respect to these challenges, in this article, we evaluate the possibility of deploying GPS-based and privacy-preserving MCTAs by evaluating both contact tracing and location tracing. We also analyze applications deployed in different countries to see how they compare. We show the possible consequences of selecting different implementations of MCTAs in a simulated environment, outlining guidelines that can help to determine different approaches for different scenarios.

Penetration is certainly a major concern of MCTAs and tracing applications in general, as we will show here. Privacy is of utmost importance given that in case there may be vulnerabilities, the user may not install it, hence, making it useless. We will also show how Position is fundamental, particularly within low-penetration scenarios. Finally, we evaluate different tracing options, showing the benefits and issues for each option.

2 CURRENT WORK AND LIMITATIONS

Since early 2020, there has been an enormous body of work that tackled the problem of contact tracing in society. The main goal is to forecast the individuals who could have been in touch with diseased people, to test those potentially exposed people in advance and, eventually, slow down disease spread. Historically, contact tracing has been done manually, with emergency responders keeping track of contacts among people and testing in advance people who may have contracted the disease [7]. During the COVID-19 pandemic, however, this has become unpractical as the number of infections per day was too high to be monitored manually [4]. This is why digital contact tracing gained interest, as it was possible to rapidly deploy custom applications that would run on mobile devices already carried by most individuals. However, national and international bodies, along with citizens, started to question the privacy issues that may arise by using MCTAs [5]. This led to almost completely abandoning the use of GPS in favor of the more privacy-preserving use of Bluetooth Low Energy (BLE). Following this path, even Apple and Google proposed an innovative framework [22], which enabled low power consumption for these applications and left enough freedom to governments and stakeholders to build custom-tailored applications for their needs.

The current MCTA ecosystem comprises several applications, all of them accounting for different limitations. First, stakeholders, such as health care operators, would largely benefit from the GPS information. In fact, the sole usage of BLE may potentially estimate the total number of potential new cases but does not give any clue as to where and how these cases are distributed. For instance, it would be greatly helpful to understand whether new cases are spread evenly over a whole region, indicating a slower pace for the infection, or are instead more densely concentrated, which indicates possible new hotbeds. A second limitation comes from BLE itself, as there are many studies showing how inaccurate its distance estimation can be [13], leading to a potentially high number of both false positives and false negatives [9]. According to many reports from all over the world, MCTAs are struggling to find true-positive cases. For instance, in Australia, MCTAs have found very few true-positive cases to date despite the application having been installed by millions of users, and the disease is still causing many people to be hospitalized [12]. This example is representative, as we find the worldwide trend to be similar or even worse. A third aspect is that, even though low MCTA penetration rates would still be helpful [15], the installation of such apps and the actions to be taken after being notified of an alleged contagion are mostly voluntary. This means that, without a tangible reward for the users, the penetration rate would remain low, and it is not guaranteed that notified users would actually self-quarantine or call the hospital in order to be tested [14]. There are even proposals for running the contact tracing activities directly into indoor public environments instead of mobile devices. This, in theory, helps solve the problem of low penetration, sometimes also affected by many people not having access to a smartphone [17]. However, it requires a massive amount of money to be deployed. The fourth and final limitation is the fact that current MCTAs are able to notify only those users who have been in direct contact with a positive case and cannot notify users who had contact with a user who had previously been in contact with a positive case (i.e., it is not transitive). A more in-depth discussion on COVID-19 MCTAs can be found in [1].

3 LOCATION CONTACT TRACING

Penetration is certainly a major concern for every application; specifically for MCTAs, it is of paramount importance to correctly and rapidly track the contagion and possible hotbeds. For this reason, incentivization techniques of all sorts have been proposed for similar applications [24]. However, for MCTAs, none of them has been adopted. The number of installed applications is far from reaching the 60% bar [11], which is worldwide recognized as the percentage at which pure MCTAs would work best [9]. In fact, MCTAs do not offer any immediate service or reward besides warning about a possible contagion, which, again, depends on the application penetration and it is most likely to take effect in the long run.

To this end, contact and location tracing applications should focus specifically on how to deal with low penetration, which happens at the beginning of the tracing campaign and may also occur in specific scenarios. Either through rewarding mechanisms that would push individuals to install the tracing applications or by specifically studying solutions for low-penetration scenarios, it is without doubt one of the major challenges that MCTAs have faced. Whenever a massive incentivization policy cannot be adopted, alternative solutions shall be inspected, often at the price of coarser data and overestimation of the danger. Based on this, we inspect an additional tracing mechanism – not necessarily an alternative to MCTA but, rather, complementary – named Mobile Location Tracing (MLT).

It is clear that the sole usage of BLE does not allow MLT, since only contacts happening between two BLE-equipped devices can be tracked, without any information about their current position. By leveraging instead the GPS, it is possible to perform both location and contact tracing at the same time. MCT with GPS implies a slightly coarser monitoring, as the distance at which two persons can be considered in contact cannot – in theory – be as precise as with the BLE. However, at the same time, it also allows consideration of a larger amount of users as potential contacts. In other words, it does not keep track of *close* contacts; rather, it keeps track of *potential* contacts in a given area. On the other hand, MLT does not focus on individuals; rather, it keeps track of locations where infections occur in order to monitor hotbeds and areas that generally need further control. This alternative has been considered in the literature, although less investigated [3]. Upon understanding that there is a possible hotbed in a monitored area, emergency responders may have different options to limit the disease spread: they could, for instance, lock down the area by forcing people to travel alternative routes in order to get to their destination. This solution has the benefit of distributing people instead of aggregating them in a specific location. However, this could potentially shift the problem to another area. Another possibility is to enforce more meticulous monitoring in the identified hotbed area to limit the contagion with additional control managed by emergency responders. Either way, MLT may be seen as an alternative to MCT if used exclusively or it can be used together with MCT, either performed through BLE or GPS. Nevertheless, it is worth noting that in case BLE is used for MCT, the joint usage with MLT would require the tracing application to use both technologies, reducing convenience and audience.

Clearly, the usage of MCT and MLT applications involves different risks and brings several advantages since they are different in terms of technologies, architectures, power consumption, location precision, and privacy implications. In Table 1, we summarize the different possibilities with MCT and MLT applications, taking both BLE and GPS into account, highlighting potential advantages and issues that may arise due to their use.

4 PRIVACY MODEL

We have mentioned privacy extensively over the earlier sections, outlining how it is more easily integrated in Bluetooth-based solutions, whereas it comes with additional effort in design-time engineering of GPS-based apps. In this section, we propose our architectural model that aims to provide possible infection information to users with GPS-based location awareness with a negligible overhead in data transfer and computation. This comes at the price of disclosing tiny bits of personal location information, as location-aware monitoring applications cannot avoid it [2]. However, the question is whether we can unveil as little information as needed to achieve the goals of the model and keep the correlation of the information revealed imperceptible. In this section, we address that aspect as well. The model is designed to fit MCTA specifically. However, we will also discuss how it can be tailored to MLTA at the same time.

We define our model as a crowd of mobile users, also defined as participants, that run our mobile application on their personal devices. We assume this application to be stand-alone for the purpose of illustrating its functionalities. However, it does not present significant differences in case it is integrated as an add-on with other applications (which can help increase penetration if the host application is already widespread enough). The app has three main functions, as shown in Figure 1. Each of the illustrated functions corresponds to different phases of the participant's life cycle, also altering the participant's role.

4.1 Crowdsensing Phase

In the crowdsensing phase, the mobile application collects data about the locations visited by the user. The application, by construction, does not require a login. Therefore, no personal information about the participant is stored anywhere but locally. The participant periodically sends a single datapoint to a central server, communicating position and timestamp of the device. Note that, again, no data about the participant are disclosed here, except the IP address, which can be forced to change over time. The server, by receiving that datapoint, associates it with a randomly generated hash and stores the couple in an internal database. Subsequently, it sends the generated hash back to the participant, who saves it locally in a personal hash pool. After sending a certain number of datapoints, the participant detains the same number of unique hashes.

Table 1. Overview of Possible Implementations of MCT and MLT Applications

Kind and Tech	Data collected	Data storage	Pros	Cons	Suitable for
MCT with BLE	Anonymous nearby codes	Local or on remote server	Low-energy technology, no location data required	Identification attacks, unreliability of BLE range estimation leads to risks of false positives and false negatives	Indoors
MCT with GPS	User location	Local or on remote server	More resilient to BLE range inaccuracies, reduced risk of false negatives	Coarse identification of contacts, high risk of false positives	Outdoors and low-population areas
MLT with GPS	User location	Remote server	Identification of hotbeds	Possible lockdowns of large areas, no contact tracing	Outdoors and public areas
MLT, MCT with GPS	User location	Remote server	Identification of hotbeds, reduced risk of False Negatives	MCT not effective due to coarse identification of contacts (same as MCT), High risk of False Positives	Outdoors and public areas
MLT, MCT with BLE and GPS	User location and anonymous nearby codes	Local and on remote server	Possibility to identify hotbeds (MLT) and finer identification of contacts (MCT)	Requires two different technologies, medium risk of false positives	Indoors and outdoors

4.2 Infection Phase

Let us imagine that, due to a medical check, one of the participants is found to be infected. Under the approval of the participant, that individual's entire hash pool is sent to the central server. Hashes are then converted to a set of "contagion events" – corresponding to locations and times – and sent into the database for matching. As a result, all hashes that are reasonably close in time and space to a contagion event are marked as "endangered." The conversion from hashes to contagion events is necessary, as multiple hashes that are very close in time and space can be clustered in a single contagion event depending on the granularity of the event itself. In fact, contagion events, as opposed to datapoints, are not punctual. Instead, they are expected to cover a spatial area and time range (with a degree of uncertainty). This also ensures a fastest computation in the matching phase.

4.3 Request and Reward Phase

This phase occurs whenever a participant wants to request information about having been in potential contact with an infected person. Intuitively, this would mean sending all of the hashes to the server for it to check the database regarding whether any of them has been marked as endangered. This, of course, discloses the full trace of the participant, which is undesirable. An immediate solution to this could be sending over one hash at a time. However, this is impractical as, for a complete privacy safeguard, it would mean changing the IP address after every message as well as hindering any rewarding strategy built on this model. In a previous work of ours [18],

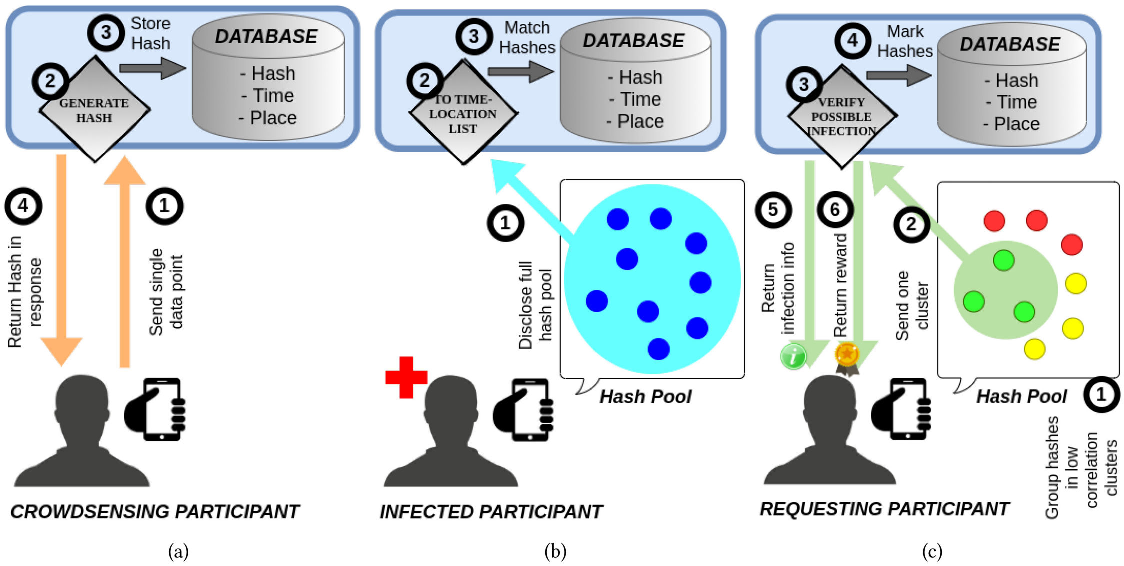

Fig. 1. The three-step process of our proposal. On the left is the Crowdsensing phase, in which users send their locations to the central entity. In the center is the Infection phase, through which the central server performs the matching among users. On the right is the Request and Reward phase, with which users can have the feedback from the central platform regarding possible infections and, if implemented, obtain a reward token.

we defined a privacy-preserving mechanism for Mobile Crowdsensing (MCS) scenarios that need to reward users for fixed-sized sets of observations. In such work, participants owning N hashes, respective to their observations, could claim a single reward by "spending" a fixed amount k of hashes. The proposed solution groups the N hashes in k-sized clusters by minimizing the intra-cluster correlation. The correlation between two datapoints can be defined in various ways depending on the field of application. This ensures that, by sending over to the server one cluster at a time, the significant information disclosed is minimized and negligible. In fact, the only information that the server (or a hypothetical man-in-the-middle) can infer is that such k uncorrelated datapoints belong to the same person. In this article, we propose an analogous mechanism: the participant clusters one's hash pool by minimizing the intra-cluster correlation and sends the resulting sets of hashes at distant points in time. The server, upon receiving the set of hashes, verifies whether they are marked as endangered and sends back the information. This operation is idempotent, that is, the participant can check the status of these k hashes multiple times, as the situation can change over time. Optionally, the first time a certain cluster is received, the participant may also receive a reward in a form that depends on the type of campaign. In case such a practice is foreseen, after this first time, the hashes are marked as "rewarded" in the database in order to ensure the uniqueness of this operation.

Note that the mechanism presented in this section can be customized based on the requirements of the scenario or the hosting application. If, for example, the actions taken in the crowdsensing phase, that is, sending one datapoint at a time, are impractical or incompatible, then datapoints could be clustered in advance with the same method presented in the request and reward phase, leaving the mechanism unaltered at the price of a slight delay in monitoring the phenomenon. Similarly, according to the privacy preferences of the infected individuals, we could apply the same method to the hash pool disclosure occurring in the infection phase. This way, hashes are not disclosed all at once. Instead, the same clustering strategy is applied. Furthermore, note that the inclusion of MLT in this scenario needs running through the first two phases only. However, if the second phase is implemented by clustering the hashes of the infected participant, then the estimate of how many different infected individuals have been inside a certain location of interest would be less precise as we cannot discriminate regarding whether hashes belonging to two clusters are actually belonging to the same person. This would necessarily affect the metrics for deciding what actions should be taken against the interested location in MLT (e.g., how many infections are needed to shut it down).

5 EVALUATION

5.1 Simulated Environment

The framework presented in this article can be used in different scenarios: to monitor large deployments such as cities or to monitor smaller areas in which hotbeds may show up due to heavy population density. It is then important to understand how to behave in these different setups. For this purpose, we evaluate different tracing proposals by combining existing mechanisms as well as those proposed in this article.

- **Mobile Location Tracing (MLT):** Here, GPS is leveraged to monitor possible hotbeds, which are consequently shut down upon identifying a possible disease spread (i.e., a certain number of infections). In this case, participants will not receive any notification about their contact with positive individuals but will have to share their GPS data in case they test positive. Clearly, this solution cannot detect specific contacts between individuals. Rather, it monitors places for higher density of individuals, which may lead to an increased infection.
- **BLE Mobile Contact Tracing (BLE-MCT):** Here, people monitor contacts with other people with the application installed. In case anyone tests positive, all previous contacts will be notified and eventually quarantined. This follows the standard and widespread approach of MCTAs, although this behavior fails to understand whether there are places in the city where infections are more severe than in other places.
- **GPS Mobile Contact Tracing (GPS-MCT):** In this case, GPS is leveraged to cover a larger area of contact compared with BLE. The notification system is similar to **BLE-MCT**, although the actual tracing uses the privacy-preserving mechanism that we presented in Section 4.
- **Mobile Location Tracing and BLE Contact Tracing (MLT-BLE-MCT):** This is a combination of two proposals, as BLE is used to monitor close contacts, while the GPS is instead leveraged to understand possible hotbeds. Although not technologically convenient for battery powered devices, it is presented as a term of comparison with the other proposals.
- **Mobile Location and GPS Contact Tracing (MLT-GPS-MCT):** In this case, BLE is not used at all, as GPS serves both as contact tracer and to identify possible hotbeds. This proposal certainly catches all of the relevant hotbeds in a city, although not specifically monitoring local contacts.

In addition to these proposals, we also evaluate the case in which no tracing is performed to provide a baseline for our experiments to show the improvements that tracing applications may offer.

To perform our analysis, we created a custom simulator programmed with the infection parameters reported in [10]. We then created a custom dataset, with users moving within a city with a density of 1200 inhabitants/km^2, similar to what can be found in many cities worldwide. This choice was driven by the fact that a real dataset that fits our requirements for duration and granularity is not available in the literature. However, we believe that, although a custom dataset is not an accurate representation of reality, it can still yield significant insight on the difference in performance among the configurations presented rather than giving absolute values. We introduce 10 infected users at the beginning and run our system considering different application penetration rates. Participants move around the area according to the Random Waypoint Model (RWM), with a relative speed of 1 m/s. The area is a 500 × 500 m square and subdivided in 1 m-sided squares. We then instantiated 3 square areas, each with a side of 25 m, which represent the Points-of-Interest (PoIs). Whenever users select a new location according to the RWM, they select a location within a PoI with a significantly high probability. The simulation spans over 2 months. Some of the infected individuals are asymptomatic, while all others show

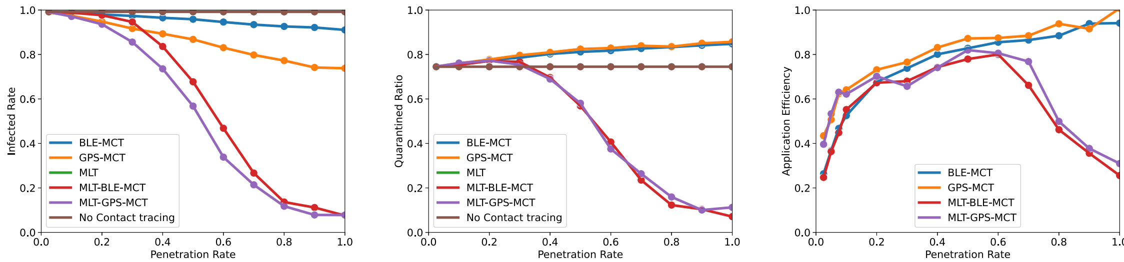

Fig. 2. Infected Rate is shown in (a), the Quarantined Rate is shown in (b), and the overall Application Efficiency is shown in (c).

their symptoms after a certain period of time and unavoidably quarantine themselves. The GPS detects contacts within a radius of 50 m, whereas the BLE detection radius is set to 4 m, but with a 25% success ratio, according to [13, 20]. Furthermore, when the central system detects a too high number of infections within a defined area, the whole area is locked and, assuming an intensive exploitation of responders, no additional infections are allowed within the area. Although we performed a basic simulation scenario compared with a real environment, we are mostly interested in qualitative and comparative results. In fact, when seeking that kind of results, multi-agent modeling environments such as this one have been widely accepted as representative [8].

In Figure 2, we show the Infected Rate (Figure 2(a)) and the Quarantined Rate (Figure 2(b)) for the different proposals. First, the GPS has a beneficial effect on the ratio of those infected in MCT, as it allows for a broader monitoring compared with BLE. All Tracing Applications increase their performance with the growth of the penetration rate, although those with the MLT expose a steep slope at a penetration rate of around 0.5, whereas GPS-MCT and BLE-MCT are slower. It is also worth mentioning that for lower penetration rates, in the [0,0.3] range, the GPS-MCT offers better performance compared with MLT solutions except for MLT-GPS-MCT, which uses the notification from GPS contacts as well. Given penetrations of different MCTAs, it is even more important to show good results in that range instead of higher penetrations.

The quarantined ratio, shown in Figure 2(b), shows an interesting behavior. As the penetration rate increases, the quarantined ratio of both BLE-MCT and GPS-MCT increases as well. This means that the inumber of infected is increasing, and although BLE-MCT and GPS-MCT help in containing the spread of disease, they are not as fast as other MLT applications, which can act faster with a higher Penetration Rate. This means that there are less people to quarantine simply because there are less people infected, as Figure 2(a) has also shown. Overall, Tracing Applications that leverage GPS are indeed more effective in stopping the infection compared with BLE-only Tracing Applications.

We also evaluate Tracing Applications in terms of True Positives (TPs), individuals selected to be quarantined and are indeed positive; False Positives (FPs), quarantined individuals who test negative; True Negatives (TNs), individuals who are free to move and not infected; and False Negatives (FNs), positive individuals not detected by the MCTA and, therefore, free to roam. The last are of utmost importance, because they are the main cause of infection spread. In order to better illustrate this, in Figure 2(c) we show the application efficiency, which we define as $\frac{TP-SQ}{SQ} \cdot \frac{1}{PEN}$, where SQ is the number of Self-Quarantined individuals, that is, those who have not gone into quarantine because the MCTA told them to do so, but because they have shown symptoms, and PEN is the Penetration Rate. In other words, it shows how well the Tracing Applications work with respect to the increase in penetration. We can observe that after 0.6 of penetration, applications relying on MLT do not gain any efficiency from an increased Penetration Rate, as the app is so popular that several locations may close; instead, MCT-only applications always benefit from the increasing penetration. We also note that the MLT-GPS-MCT

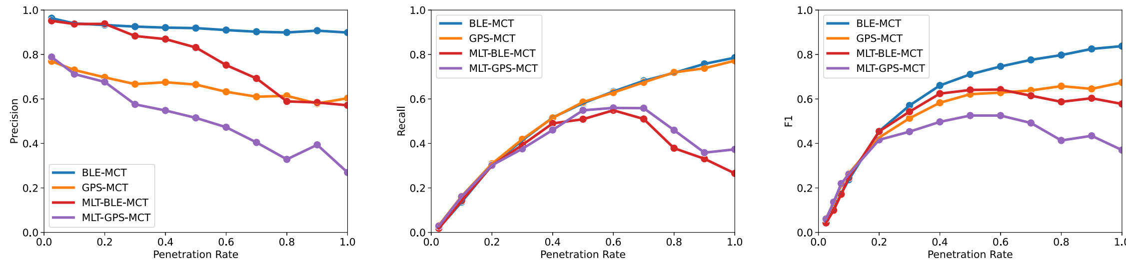

(a) (b) (c)

Fig. 3. The application efficiency is shown in Figure 2(a), the recall is shown in (b), and the F1-measure is shown in (c).

follows a similar trend as the GPS-MCT at an early stage but drops as penetration increases due to an increased number of locked down spots.

In Figure 3, we show three different performance metrics which evaluate how precisely Tracing Applications selected individuals to notify, quarantine, or what places to shut down. All of these metrics assess how well the Tracing Applications foresee individuals who may have been in contact with a positive individual without putting into quarantine an excessive amount of individuals or leaving positives untested.

The Precision, defined as $\frac{TP}{TP+FP}$, shows how precise Tracing Applications are in identifying only those individuals which are indeed positive. In other words, a high precision score means that few negative individuals would be quarantined, however possibly leaving positive individuals untested. The metric is shown in Figure 3(a), which shows that for MCT Precision is almost nondependent on the Penetration Rate. Applications relying on MLT instead show a lower precision as the penetration increases, thus, raising the number of FP (negative individuals put into quarantine). This happens because, while the Penetration Rate increases, it also increases the total area monitored and possibly the number of places recognized as hotbeds.

The recall is defined as $\frac{TP}{TP+FN}$, showing how well Tracing Applications are able to detect all positive individuals at the cost of possibly quarantining negatives. The metric is shown in Figure 3(b), showing exactly what we described for the precision. At low penetration rates, all Tracing Applications work equally well. However, after 0.5 of Penetration Rate, applications relying on MLT drastically drop in performance. This happens since as the number of locked down places increases, individuals need to find alternative routes to their destination, distributing them sparsely over the area of interest. Hence, contacts that occur in sparse regions are unlikely to be real, while MCT applications tend to overestimate.

Finally, the F1 score combines the recall and the precision into a single metric, defined as $\frac{2 \cdot PRECISION \cdot RECALL}{PRECISION+RECALL}$, showing the balance between the two. Figure 3(c) summarizes well our previous discussions, as Tracing Applications at medium penetration rates start to expose different F1 scores, confirming our previous comments.

5.2 Real-World Analysis

In this section, we present a comparison of real-world applications categorized into voluntary and mandatory, showing their effectiveness regarding the number of cases in different countries. We conclude our analysis by discussing current trends in location and contact tracing.

5.2.1 Location and Contact Tracing Applications Effectiveness. To analyze the effectiveness of MLT and MCTAs, we have considered the different applications proposed and utilized in several countries around the world. We have also distinguished between voluntary applications, meaning that there were no strict installation restrictions on the population, and mandatory applications, meaning countries that imposed installation restrictions or more strict mobility regulations for those who did not install the application.

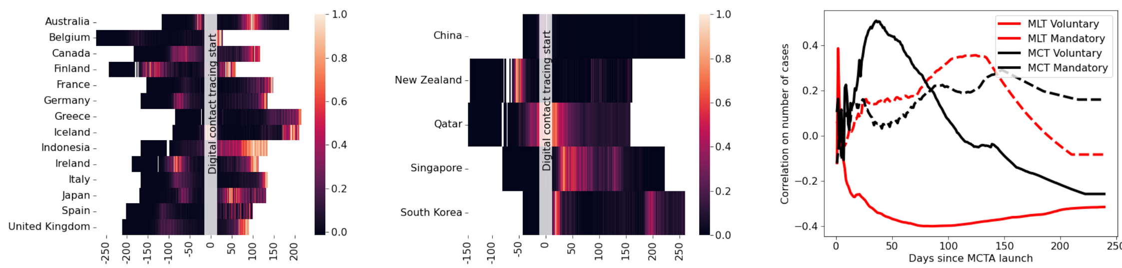

(a) (b) (c)

Fig. 4. Experiments on real-world Location Contact Tracing applications. (a) presents countries in which voluntary applications were proposed; (b) shows the same for countries with mandatory applications. The correlation of the number of cases for the different MCT and MLT applications is provided in (c).

A first analysis that we performed, shown in Figures 4(a) and 4(b), takes into account the change in daily cases since the adoption of the application. Different countries had a different application deployment date. We considered a relative time before and after the application deployment, which allows us to align the deployment date across all the countries. We then took the maximum number of cases of the considered period and marked it as 1, and computed the correlation with all the other days, which then present values between 0 (0 cases) and 1 (same number of cases as the maximum). Regarding voluntary applications, it is evident that, globally, the peak in the number of cases is well after the deployment of the application, possibly meaning that its introduction did not contribute significantly to the decrease in the number of cases. Looking instead at Figure 4(b), we see a different trend, in which the highest number of cases was before the deployment of the application or shortly after it. We note that this behavior may not be totally related to the application itself since in different countries there were also other limitations that went beyond the sole use of the application. This behavior is also confirmed by Figure 4(c), which presents the correlation on the number of cases for 4 different application types: MLT Voluntary and Mandatory, and MCTA Voluntary and Mandatory. In this case, for each country, we correlate the daily number of cases against the cases on day 0, that is, the application deployment date. What it is possible to see is that after a few days, 4 trends emerge: countries with the MLT Mandatory application immediately show a decrease in the number of daily cases, whereas MCT Mandatory applications take more time to experience a similar trend. For voluntary applications, after almost 5 months, daily cases continued to rise regardless of the technology used. However, Voluntary MLT applications showed a slightly better result after 5 months.

We want to reiterate that these results may be hindered by all of the other restrictions in place in these countries. In fact, countries that have imposed mandatory applications also generally adopted more strict measures to contain the infections. Nevertheless, our results present real evidence of the trends in daily cases in countries with different policies and using different location and contact tracing technologies. It is also evident from Figure 4(c) that countries with MLT applications deployed experienced a more rapid decrease of daily cases compared with countries with MCT applications.

5.2.2 Current Trends and Limitations. Concerning the COVID-19 situation at the time of writing, we are witnessing a parallel trend that somehow resembles the MLT paradigm: check-in–based contact tracing. In such a descending phase of the pandemic emergency, the need for monitoring restricted communities feels more crucial and more efficient than applying tracing at large [19]. There is an emerging trend for tracing applications in which users are requested to advertise their presence as they check in to a definite place (usually indoors), for instance, by scanning a QR code. A nationwide example is the COVID Tracer application in New Zealand,[1] but also

[1]https://tracing.covid19.govt.nz/.

the NHS COVID-19 app in the United Kingdom[2] or the app in South Korea. This is more practical when infected individuals are rarer and contacts are less likely to take place outside these locations. Clearly, this paradigm requires voluntary participation by users, which makes it inherently different from MLT and MCTAs, which are mostly opportunistic [16]. Nevertheless, check-in–based contact tracing shares the same base concepts with MLT, as the monitoring of communities is based on location. Therefore, MLT concepts could still have an applicability in such a recovery phase. In addition, many countries with high vaccination rates have now entered the consolidation, or even the post-vaccination, phase (e.g., Australia)[3] with the requirements for contact tracing and check-ins already phased out as of this writing. Mobile contact tracing has become a legacy of the coronavirus pandemic; the strategies, limitations, and challenges presented in this article are lessons learned that we hope are useful for planning for future surveillance strategies beyond applications in the public health context.

6 DISCUSSION AND CONCLUSION

This article has discussed GPS-based tracing approaches, both in terms of MCT and MLT, for tackling the challenge of the worldwide COVID-19 outbreak by addressing fundamental pillars. We have witnessed that many MCTAs have been proposed and, in many cases, also adopted by governments as an official counterthrust, such as Italy[4] and France[5]. Nevertheless, these efforts have mostly been in vain, none of the MCTAs had a sufficient penetration at the time of writing, with perhaps the exception of Iceland [21], and, even then, no relevant result emerged. In this article, we proposed a comprehensive solution that looks at the current issues in MCTAs from different sides. It is worth specifying that our proposal is not necessarily concurrent with BLT-based solutions. Rather, it can complement them by potentially providing additional information about where and when contagion clusters occur. We try to consider all factors that influence a good outcome of the application. Factors are briefly summarized and discussed in this section. We identified Penetration, Privacy, Position, and Performance, referred to as the four pillars (the four Ps) of contact tracing.

Penetration. The first point to discuss is, as mentioned in the introduction, the penetration of the tracing application. Whether tracing contacts or locations, or both, we have shown how in order to rapidly control disease spread, it is mandatory to obtain high penetration rates. However, we have also shown that, by monitoring places and in general leveraging data from the GPS, it is also possible to better assess the contagion, hence, reducing the infected ratio. In the current situation, there is no way that simply proposing a stand-alone MCTA would work; we also must not forget that the MCTA needs to be operative quickly, without necessarily waiting for the conscious response of the crowd. This obviously affects the way in which data privacy is preserved, which leads us to the second point of discussion.

Privacy. The second point to discuss relates to the privacy concerns that arise when sharing location data that is fine-grained both in space and time. We presented a framework for sharing such data, preserving its spatial precision and guaranteeing that only a definite subset of positions can be attributed to the same person. Also, connections are anonymous and pull based, meaning that the client initiates the "infection check" on that subset of information, and the server does not have a way to contact them directly. This ensures that each subset is in no way related to other subsets, maintaining anonymity for users. The extent to which this paradigm works and does not cause unpleasant side-effects is entirely due to parameter tuning for different scenarios, leading us to the next points of discussion.

Position. As we have shown, the Position enables better monitoring for any Tracing Application, as it allows understanding of where hotbeds originate in order to react faster and more precisely. Either as a possibility

[2] https://covid19.nhs.uk/.
[3] https://www.australia.gov.au/national-plan.
[4] https://www.immuni.italia.it/.
[5] https://www.economie.gouv.fr/stopcovid.

to enable MLT or to use as an addition to MCT, it boosts the ability to better monitor and handle the disease spread. Although GPS possibly carries more privacy concerns compared with BLE, we have shown how it would be possible to leverage on it with Privacy-Preserving techniques. Overall, both BLE and GPS present strengths and weaknesses. Their use should be evaluated depending on the scenario, on the penetration rate, and on the government's policies.

Performance. The fourth point to discuss is performance. The way in which certain parameters are set can influence any unwanted behavior. First, concerning GPS tracking, we have shown that the "notification area" size of the GPS boosts performance, as contacts have to be considered as such in areas larger than 1 m in order to avoid false negatives, which are way more dangerous than false positives. This creates unavoidably false positives; therefore, the size of this area should be appropriately tuned according to the scenario and the environment. This also has an interesting side-effect on privacy, since once this size has been established, participants can report their position using cloaked areas (i.e., with lesser precision) of that same size without significantly changing the overall performance and considerably improving their privacy as their precise position is slightly obscured. We also have shown that GPS in general contains the infection way better than BLE does by introducing, as stated earlier, an acceptable degradation of precision without affecting the recall. Finally, we have shown the effectiveness of MLT, which can be, when tuned properly, an efficient tool for preventing localized outbreaks without incurring a massive shutdown. A total lockdown would reduce the FN number but at the cost of shutting down shops and constraining people who may have not been in contact with positive individuals. Hence, more selective lockdowns are preferable, as in MLT. However, appropriately tuning the parameters is of utmost importance, as MLT can easily (and rapidly) lead to a corner case similar to a total lockdown, especially at high penetration.

Discussion and future challenges. Future research in this area should be devoted mostly to understanding how to correctly balance the trade-off between the different characteristics we have identified. Maintaining privacy is crucial and should be discussed at design time when deploying such applications. It is also challenging, as different regulations exist in different parts of the world, making the privacy aspect difficult to standardize. Penetration is directly related to how such applications are advertised and how citizens may react to them. Governments should clearly state the benefits of such applications to citizens, and the trust that citizens have in their own government has a strong impact on their adoption [23]. Regarding position, future efforts should go toward understanding how to achieve accurate localization of devices without exposing private information of citizens. Finally, performance is a consequence of the discussed choices: if MCTAs are widely adopted (high penetration) and provide accurate localization of hotbeds (accurate position) while keeping private data of citizens safe (high privacy) then achieving high performance will be possible eventually.

REFERENCES

[1] Nadeem Ahmed, Regio A. Michelin, Wanli Xue, Sushmita Ruj, Robert Malaney, Salil S. Kanhere, Aruna Seneviratne, Wen Hu, Helge Janicke, and Sanjay Jha. 2020. A Survey of COVID-19 Contact Tracing Apps. arXiv:cs.CR/2006.10306.

[2] Mohammad Abu Alsheikh, Yutao Jiao, Dusit Niyato, Ping Wang, Derek Leong, and Zhu Han. 2017. The accuracy-privacy trade-off of mobile crowdsensing. *IEEE Communications Magazine* 55, 6 (2017), 132–139.

[3] Nikos Askitas, Konstantinos Tatsiramos, and Bertrand Verheyden. 2020. *Lockdown Strategies, Mobility Patterns and COVID-19*. IZA Discussion Papers 13293. Institute of Labor Economics (IZA). Retrieved July 16, 2022 from https://ideas.repec.org/p/iza/izadps/dp13293.html.

[4] A. Barrat, C. Cattuto, M. Kivelä, S. Lehmann, and J. Saramäki. 2020. Effect of manual and digital contact tracing on COVID-19 outbreaks: A study on empirical contact data. *medRxiv* (2020). https://doi.org/10.1101/2020.07.24.20159947

[5] Lars Baumgärtner, Alexandra Dmitrienko, Bernd Freisleben, Alexander Gruler, Jonas Höchst, Joshua Kühlberg, Mira Mezini, Richard Mitev, Markus Miettinen, Anel Muhamedagic, Thien Duc Nguyen, Alvar Penning, Dermot Frederik Pustelnik, Filipp Roos, Ahmad-Reza Sadeghi, Michael Schwarz, and Christian Uhl. 2020. Mind the GAP: Security & privacy risks of contact tracing apps. (June 2020), *arXiv e-prints arXiv:cs.CR/2006.05914*.

[6] Luca Bedogni, Shakila Khan Rumi, and Flora Salim. 2020. Modelling memory for individual re-identification in decentralised mobile contact tracing applications. arXiv:cs.CY/2010.05514.
[7] T. Berge, A. J. Ouemba Tassé, H. M. Tenkam, and J. Lubuma. 2018. Mathematical modeling of contact tracing as a control strategy of Ebola virus disease. *International Journal of Biomathematics* 11, 07 (2018), 1850093. https://doi.org/10.1142/S1793524518500936
[8] Erik Cuevas. 2020. An agent-based model to evaluate the COVID-19 transmission risks in facilities. *Computers in Biology and Medicine* 121 (2020), 103827.
[9] Danielle Currie, Cindy Peng, David Lyle, Brydie Jameson, and Michael Frommer. 2020. Stemming the flow: How much can the Australian smartphone app help to control COVID-19? *Public Health Research & Practice* (2020). Retrieved July 16, 2022 from https://www.phrp.com.au/issues/june-2020-volume-30-issue-2/stemming-the-flow-how-much-can-the-australian-smartphone-app-help-to-control-covid-19/.
[10] Shuo Feng, Zebang Feng, Chen Ling, Chen Chang, and Zhongke Feng. 2021. Prediction of the COVID-19 epidemic trends based on SEIR and AI models. *PLOS ONE* 16, 1 (2021), 1–15. https://doi.org/10.1371/journal.pone.0245101
[11] France24. [n.d.]. France's COVID-19 tracing app fails to engage, chalking up roughly 1.5 million userss. https://www.france24.com/en/20200623-france-s-covid-19-tracing-app-fails-to-engage-chalking-up-roughly-1-5-million-users. Accessed: 2021-11-05.
[12] The Guardian. [n.d.]. Australia's Covidsafe coronavirus tracing app works as few as one in four times for some devices. Retrieved May 11, 2021 from https://www.theguardian.com/australia-news/2020/jun/17/covid-safe-app-australia-covidsafe-contact-tracing-australian-government-covid19-tracking-problems-working. Accessed: 2021-11-05.
[13] Hsu-Chun Hsiao, Chun-Ying Huang, Bing-Kai Hong, Shin-Ming Cheng, Hsin-Yuan Hu, Chia-Chien Wu, Jian-Sin Lee, Shih-Hong Wang, and Wei Jeng. 2020. An empirical evaluation of Bluetooth-based decentralized contact tracing in crowds. arXiv:cs.CY/2011.04322.
[14] Web24 Journal. [n.d.]. "Prisoner" of Immuni: It is the complaint of a 63 year old from Bari. Retrieved May 11, 2021 from https://www.web24.news/u/2020/06/prisoner-of-immuni-it-is-the-complaint-of-a-63-year-old-from-bari.html.
[15] Michelle Kendall, Luke Milsom, Lucie Abeler-Dörner, Chris Wymant, Luca Ferretti, Mark Briers, Chris Holmes, David Bonsall, Johannes Abeler, and Christophe Fraser. 2020. Epidemiological changes on the Isle of Wight after the launch of the NHS Test and Trace programme: A preliminary analysis. *The Lancet Digital Health* 2, 12 (2020), e658–e666.
[16] Uichin Lee and Auk Kim. 2021. Benefits of mobile contact tracing on COVID-19: Tracing capacity perspectives. *Frontiers in Public Health* 9 (2021), 586615.
[17] Mohamed Mokbel, Sofiane Abbar, and Rade Stanojevic. 2020. Contact tracing: Beyond the apps. *SIGSPATIAL Special* 12, 2 (2020), 15–24.
[18] Federico Montori and Luca Bedogni. 2020. A privacy preserving framework for rewarding users in opportunistic mobile crowdsensing. In *7th International Workshop on Crowd Assisted Sensing, Pervasive Systems and Communications (CASPer'20)*.
[19] Ichiro Nakamoto, Sheng Wang, Yan Guo, Weiqing Zhuang, et al. 2020. A QR code–based contact tracing framework for sustainable containment of COVID-19: Evaluation of an approach to assist the return to normal activity. *JMIR mHealth and uHealth* 8, 9 (2020), e22321.
[20] Khuong An Nguyen, Zhiyuan Luo, and Chris Watkins. 2020. Epidemic contact tracing with smartphone sensors. *Journal of Location Based Services* 14, 2 (2020), 92–128. https://doi.org/10.1080/17489725.2020.1805521
[21] Technology Review. [n.d.]. Nearly 40% of Icelanders are using a COVID app–and it hasn't helped much. Retrieved May 11, 2021 from https://www.technologyreview.com/2020/05/11/1001541/iceland-rakning-c19-covid-contact-tracing/.
[22] Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, Marcel Salathé, James Larus, Edouard Bugnion, Wouter Lueks, Theresa Stadler, Apostolos Pyrgelis, Daniele Antonioli, Ludovic Barman, Sylvain Chatel, Kenneth Paterson, Srdjan Čapkun, David Basin, Jan Beutel, Dennis Jackson, Marc Roeschlin, Patrick Leu, Bart Preneel, Nigel Smart, Aysajan Abidin, Seda Gürses, Michael Veale, Cas Cremers, Michael Backes, Nils Ole Tippenhauer, Reuben Binns, Ciro Cattuto, Alain Barrat, Dario Fiore, Manuel Barbosa, Rui Oliveira, and José Pereira. 2020. Decentralized privacy-preserving proximity tracing. arXiv:cs.CR/2005.12273.
[23] YouGov. [n.d.]. International Covid-19 tracker update. Retrieved May 11, 2021 from https://yougov.co.uk/topics/international/articles-reports/2020/05/18/international-covid-19-tracker-update-18-may.
[24] Xinglin Zhang, Zheng Yang, Wei Sun, Yunhao Liu, Shaohua Tang, Kai Xing, and Xufei Mao. 2015. Incentives for mobile crowd sensing: A survey. *IEEE Communications Surveys & Tutorials* 18, 1 (2015), 54–67.

Received February 2021; revised November 2021; accepted April 2022