Contact Tracing by Giant Data Collectors: Opening Pandora's Box of Threats to Privacy, Sovereignty and National Security

Antoine Boutet, Claude Castelluccia, Mathieu Cunche, Alexandra Dmitrienko, Vincenzo Iovino, Markus Miettinen, Thien Nguyen, Vincent Roca, Ahmad-Reza Sadeghi, Serge Vaudenay, et al.

Contact Tracing by Giant Data Collectors: Opening Pandora's Box of Threats to Privacy, Sovereignty and National Security

Antoine Boutet[1], Claude Castelluccia[1], Mathieu Cunche[1], Alexandra Dmitrienko[2], Vincenzo Iovino[5], Markus Miettinen[3], Thien Duc Nguyen[3], Vincent Roca[1], Ahmad-Reza Sadeghi[3], Serge Vaudenay[4], Ivan Visconti[5], and Martin Vuagnoux[6]

[1]INRIA Privatics, France
[2]JMU Würzburg, Germany
[3]Technical University of Darmstadt, Germany
[4]EPFL, Switzerland
[5]University of Salerno, Italy
[6]base23, Geneva, Switzerland

November 2020

Abstract

Many countries have introduced digital contact tracing apps to fight the COVID-19 pandemic. Such apps help to identify contacts between potentially infectious persons automatically and thus bear the promise of reducing the burden on manual contact tracers and increase tracing accuracy in situations in which people have difficulties identifying with whom they have been in contact.

A number of different proposals for digital contact tracing systems have been made or deployed, ranging from heavily centralized to completely decentralized approaches, each with its own advantages and disadvantages in terms of tracing effectiveness and impact on user privacy. During the phase of highly dynamic evolution of these approaches, surprisingly, Google and Apple established an unprecedented friendship and agreed on a very special scheme for contact tracing, realizing this in the form of an API called GAEN that they quickly integrated into their mobile operating systems. A multitude of nationally rolled out tracing apps are now based on the GAEN approach.

In this paper, we revisit such apps and the GAEN API on which they are built. In particular, we point out a number of very problematic aspects and threats that the GAEN approach creates through its security and privacy weaknesses but also through the threats that it poses on technological sovereignty and the public health system.

1 Introduction

The corona virus pandemic has had the world in its grip for months. The number of infections is rising and the second wave is rolling. Reliable and efficient contact tracing has therefore become more important than ever. In many countries, digital contact tracing apps on smartphones have already been enrolled with the hope to significantly support manual tracing in breaking infection chains and preventing the virus from spreading further[1]. Depending on privacy regulations and the perceived importance of data protection in individual countries, different approaches have been applied. These approaches are mainly divided into two categories, i.e., centralized and decentralized, based on what type of information about the users (and their social graph) is shared with the organization running the backend server of the tracing app (usually a governmental organization like a CDC).

While the first countries (predominantly in Asia) that deployed tracing apps adopted centralized approaches, and extensively collected sensitive user information (e.g., names, addresses, mobile phone numbers, location), a widespread and heated debate on user privacy broke out in Europe and the USA[2]. As a result, it became a matter of academic debate/competition as well as national pride who will deploy the first and/or the best privacy-preserving contact tracing solution. In this turmoil of evolving contact tracing approaches, somewhat surprisingly, Google and Apple established an unprecedented friendship and agreed on their very special decentralized system for contact tracing interface called

[1]Although the usefulness of tracing apps is also questioned (see, e.g., Ross Anderson's view [3])

[2]In the course of this debate about 300 security and privacy researchers from 26 countries signed an open letter criticizing the privacy risks of centralized contact tracing approaches, advocating privacy-preserving solutions whenever better privacy can be obtained without penalizing effectiveness [27].

Exposure Notification API (GAEN) [4] which they rapidly integrated into their mobile operating systems.

Despite their somewhat questionable track record with regard to user privacy and being ordered millions of euros in fines for breaching EU privacy laws [37, 42] as well as ongoing billion-dollar lawsuits in the US [29], Apple and Google positioned themselves to the forefront of 'privacy-preserving' contact tracing by publicly announcing their support for privacy-conscious academic decentralized proposals like DP-3T [35]. By doing so, these companies managed to maneuver themselves into a position in which they currently *de facto* control an oligopoly for contact tracing in many European countries. Conveniently, this co-incides with, e.g., plans of Google's mother company to establish a presence in the health insurance market [11]. Although the documentation of the API is openly available, access to it is heavily controlled by corporate policy so that in each country access is granted only to one single organization that needs to be approved by the corresponding national government. We believe that this dominance is part of the problem we are pointing out to in this paper. A number of governments such as Germany (Corona-Warn-App) [41], Italy (Immuni) [30], and Switzerland (SwissCovid) [39], to name some, have contracted local companies (some with millions of Euros) to develop apps that specifically use the GAEN API. In several cases (e.g., it is clear in UK but pretty obvious also in other countries) the decisions to use GAEN were in contrast with the desires of governments that would have strongly preferred a different design for their automatic contact tracing systems [7]. However, Google and Apple have – referring to supposed privacy issues – systematically refused to offer a flexible Bluetooth API for supporting contact tracing solutions, thereby imposing their own solution on several national governments. This has taken place despite the fact that several major problems with GAEN have been documented in the past 6 months. By the time of writing, 33 countries and US states have already enrolled tracing apps based on GAEN API, or are planning to do so [25].

In this context, France is an exception, officially declaring that resorting to the initiative of Apple and Google would seriously question the sovereignty of the state as the sole authorized entity to control sensitive health information[3]. The decision of Apple and Google to unilaterally impose their contact tracing solution as the sole technical standard in this area goes against the explicit will of several democratically elected European governments [7] and has significant weaknesses going against the best interests of the public, as it

- exposes citizens to severe security and privacy problems (cf. Section 2) and consequently poses crucial threats on our public heath system and even on national security,
- lacks transparency and clear design rationale,
- threatens national technological sovereignty and innovation.

In addition, while the effectiveness of GAEN-based systems is still unclear, the high costs for citizens[4] in some countries are remarkable. In this paper we aim to summarize various critical aspects of GAEN and also discuss some of the enrolled tracing apps based on GAEN.

Our intention is to nudge the community out of its apparent coma around this topic and to encourage a critical debate on tracing apps based on GAEN and look in other directions. Since Apple and Google have refused so far to make their APIs more flexible, we strongly believe that governments should shut down current systems leaving it to Apple and Google to take full responsibility for the experienced failure caused by their unwillingness to collaborate with democratic governments on this topic.

In particular, we recall the original genuine effort of our colleagues who signed in April 2020 a letter asking for privacy-preserving solutions whenever available without affecting efficacy [27] possibly without exposing the social graph. Those valid demands were unfortunately misinterpreted by governments leading them to capitulate as soon as giant data collectors such as Google and Apple decided to intervene. They are now positioning themselves as a dominant force in the background into the European public health systems by imposing upon us an API that has potentially catastrophic security and privacy issues.

2 Security and Privacy Threats of GAEN

A number of attacks against the security and privacy of GAEN have been propose. In the following, we will review some of the recent attacks.

2.1 Security Threats

Researchers have pointed out several security threats against which the GAEN tracing approach is vulnerable. These include online relay attacks [8, 44, 12, 36, 18], 'time machine' attacks and tracing forgery attacks [5, 45, 15], to name some (see also Appendix

[3] A translated excerpt: "... In addition, resorting to the initiative of Apple and Google would raise serious questions of sovereignty. The Government considers that protecting the health of the French is a task that falls exclusively to the State and not to private international actors. The definition of the contact-tracing algorithm and the capacity of the health authority to have all the statistical data to improve the efficiency of its action, cannot therefore be left in the hands of another entity: it is a question of health and technological sovereignty...".

[4] Several costs should be taken into account, e.g.: developing and maintaining the backedend and the app, updating procedures in the health system, buying a new smartphone to be able to use the system.

5 for more details). All of these attacks aim to sabotage the reliability of the tracing system in order to generate massive false exposure notifications on a large scale that can cause public unrest and strain the health systems as well as cause unnecessary quarantining. Further, [15] discusses the potential threat of using such attacks against democratic processes, e.g., the US elections[5].

2.2 Privacy Threats

Similarly, also a number of privacy threats related to GAEN exist. These include massive data collection [24], or creating movement profiles of infected individuals [46] and [8]. Notice that a privacy threat can arise from anyone (not only the bad government that likes to spy on its citizens) who can "listen" to Bluetooth beacon identifiers announced by smartphones using GAEN. As such, contact tracing systems using GAEN expose personal and medical data to third parties in an uncontrollable manner.

Leith and Farrell [24] provide evidence that Google Play Services, into which the GAEN functionality is embedded, collects and shares extensive sensitive information from GAEN-based app users. In particular, Google Play Services shares detailed information about the phone, e.g., IMEI, WiFi MAC address, hardware serial number as well as information about the apps running on the phone. Further, even under privacy-conscious settings, Google Play Services still connects to Google servers approximately every 20 minutes to potentially locate the phone via its IP address.

Obviously, the above requirements (e.g., a Google account to use Google Play Services) naturally expand the coverage of such data collecting giants[6] (e.g., to users that were more conservative about the use of smart phones might now have a modern smartphone with Bluetooth and GPS always-on with in the hope of protecting their health) and this is happening while Google and Apple seem to be downplaying dangers of attacks on privacy, claiming that they are not a major concern.

These privacy threats are exacerbated by the fact that on the Android platform, due to operating system policy settings, smartphones using tracing apps need to enable the geolocalization service in order to use Bluetooth LE. Even though it is likely true that tracing apps themselves do not use positioning, enabling geolocalization will enable any other apps with appropriate permissions as well as the Google Play Services to use it.

3 Sovereignty and National Security

The widely discussed divergence in data protection policies also gives rise to the question of confidence:

Are Apple and Google sufficiently trustworthy and transparent entities to be eligible to enter as players into public health care systems?

Several countries have publicly released the source code of their apps and backend servers, but this gives a somewhat misleading impression of transparency of the related tracing apps. It is true that availability of source code in general is useful as it allows one to perform in-depth testing of the app, in particular when trying to understand potential bugs/vulnerabilities and possible countermeasures. However, such testing is not possible with contact tracing systems based on GAEN since it is only accessible on phones when used by unmodified and approved national apps. While the source code of the official German, Italian and Swiss apps are public[7], the source code of the operating systems of Apple and Google are more difficult to verify - even by experts. For instance, the Google Play service containing the GAEN functionality on Android devices is closed source and can thus not undergo independent scrutiny. It is not only the lack of transparency, but also the enormous data power of these two companies that continues to raise critical questions on technological sovereignty.

Who has the task of protecting the health of the population? American data monopolies or the government?

For people to have trust in the tracing function, it is essential to protect user privacy and data integrity. The decisive factor is who has control over the data.

What happens if Apple and Google stop supporting their API, or make it purely their own business?

Until now the European healthcare system has been an area that has suffered only marginally from the dominance of giant data collectors. Do we want to give up this independence to Apple and Google who already have access to vast amount of data about their users?

Who guarantees that the data already collected will not be linked and correlated with the contact tracing data?

As already mentioned, the contact tracing interface of Apple and Google still suffers from data protection limitations such as the ability to create move-

[5]The adversary generates massive fake exposure notifications at a post office used for voting-by-mail or in swing voting districts so that the adversary can impair the functionality of vote delivery or prevent people from going to vote.

[6]Clearly Google and Apple already have a significant infrastructure to map social relationships and governments should try to reduce this rather than expanding it.

[7]This transparency has allowed to discover a few bugs, even catastrophic ones (e.g., `https://mrsuicideparrot.github.io/security/2020/07/30/CVE-2020-15957.html`). Nevertheless, more serious bugs have been found recently (e.g., now it is even suggested to open the app once every day since there is a problem in displaying an alert in case of exposure when the app runs in the background only), and thus it is clear that only a few experts gave a look at the source codes.

ment profiles of infected users (see the Appendix for more details on this).

And who guarantees that other countries do not interfere and that the contact tracing results are not manipulated by attacks?

Recent research has shown that tracing forgery attacks (cf. Sect. 2.1) can even threaten national security [5, 15]. Creating false contacts and other denial of service attacks can be conducted for commercial gains but also to generate public panic and disinformation as pointed out in [5, 8, 45], or can be used to harm, or disturb democratic processes like, the US election by generating massive fake exposures [15]. A government might take wrong actions ending up in affecting negatively the health of its citizens just because data collected through apps based on GAEN are vulnerable to attacks that completely violate the integrity of contact data.

4 Nationally Enrolled GAEN-based Tracing Apps

In this section, we take a look at the selected subset of nationally enrolled GAEN-based tracing apps.

4.1 Corona-Warn-App (Germany)

The first 100 days of the official Corona-Warn-App (CWA) in Germany have already passed. The German government started a nationwide advertising campaign for the CWA and its use was personally recommended by top politicians such as the German Chancellor. According to the Robert Koch Institute (RKI), 21.1 million downloads of the app were recorded at the end of October 2020. The CWA is advertised as the European model app compared to other decentralized as well as centralized tracing apps. The declared aim of the CWA was to efficiently and quickly detect and interrupt infection chains. All users should be reliably and promptly informed about encounters with infected persons. Does the Corona-Warn-App help to slow down the spread of COVID-19? And how effective is the app actually?

Unclear effectiveness. Researchers have investigated the effectiveness of tracing apps [10]. In their study they evaluated 15 automated and semi-automated contact tracing solutions but found no empirical evidence of the effectiveness of automated contact tracing (in terms of the contacts traced or the reduction in transmission). They conclude that contact tracing can help to confine COVID-19 if enough people *use* such an app. Hence, the question arises, do really many people actively use CWA?

According to an online tracker of the Diagnosis Keys of the German Corona-Warn-App[8], roughly an estimated 12-13% of infected users use currently the Corona-Warn-App to share their diagnosis keys with others. Currently, we do not have sufficient information about how many people actually use the Corona-Warn-App and whether it is effective. According to a study by the University of Oxford, tracing apps begin to have an effect as soon as 15 percent of the population participate. According to further scientific estimates, at least 60 percent of the population would have to participate in digital contact tracing for the Corona-Warn-App to achieve a digital equivalent to herd immunity [32].

Let's take a look at the available data: The Corona-Warn-App (CWA) has been downloaded 21.1 million times until now. However, this does not mean that the app is actively used by 21.1 million people. After all, one download does not correspond to one user. Also, this number does not speak much about effectiveness of the app. The number of downloads can vary from country to country. From our point of view, the acceptance and the corresponding downloads of the tracing apps also depend on social factors. These include, for example, the extent to which people trust the government and top politicians when it comes to public health. If, as in Germany, politicians promote an app with a large advertising campaign, it is to be expected that many people will download it - either out of curiosity or out of the belief that the Corona-Warn-App will help them against the pandemic.

The CWA doesn't use a central database and therefore the exact number of app users can only be estimated based on the overall daily number of downloaded TEKs. Furthermore, it is not known how many people are warned by the CWA. The decision whether users share their diagnosis keys after a positive test result is voluntary. According to recent publications on github.io, more than 90.000 users of CWA have done this to date [9].

In this context, what is completely missing is a digital infrastructure that would allow a privacy-preserving processing of information from involved parties like users, health care professionals and authorities. The German CDC, the Robert-Koch-Institut (RKI) would also like to extend the functionality of the Corona-Warn-App and merge various existing applications into a single "Universal-App" [22]. Only with the help of exact information about the user data the effectiveness of the app could be evaluated extensively.

A similar picture can be seen in Italy. There, the warning app Immuni was downloaded by about 9.3 million Italians [17]. That is a bit more than 15 percent of the population. This means that the download numbers are well below the government's target. Spain has so far been the country most affected by the second wave of the pandemic. One of the reasons for the high incidence of infection would seem to be the overloading of the healthcare system. The Spanish authorities were not able to track contacts to the necessary extent. Despite this fact, the

[8] `https://micb25.github.io/dka/`

Radar Covid app published in August has only about four million downloads [28]. In addition, the warning app was not expected to be operational throughout Spain until mid-September [31].

However, in order to evaluate and prove the effectiveness of this method, further prospective studies are needed. Furthermore, such studies usually assume that the integrity of the system is always completely given. However, as shown, relay attacks can completely compromise the integrity of the system and thus its effectiveness.

High costs. Besides the evaluation of its effectiveness, we believe it is equally important to compare the costs of the Corona-Warn-App with its usage and effectiveness. According to the contract between Deutsche Telekom and the German government, the costs will amount to 68 million Euros by the end of 2021 [40]. The development of the Corona-Warn-App by Deutsche Telekom and SAP alone has already costed EUR 20 million. The monthly costs for running the app are estimated at more than three million euros [16]. In comparison to many other countries, a major fraction of the costs goes towards implementing necessary infrastructure that is required for the CWA to work. Implementation of the App itself is only a minor part. Therefore, development costs of a more secure and independent app solution would have only a minor impact on the overall cost of the overall system.

4.2 Immuni (Italy)

The Italian government announced an open call and received at the end of March 2020 more than three-hundred proposals for a national automatic contact tracing system to be provided for free. The selection was performed by a task force that opted for a proposal named "Immuni" of the Italian company "Bending Spoons" that was part of the PEPP-PT [33] consortium[9]. While the initial proposal did not have a decentralized design, about a week after the selection of Immuni, the government announced that the system would be based on GAEN. During that intense week a strong campaign against the system of PEPP-PT started in social media and newspapers in many cases using conspiracy arguments and stressing risks of mass surveillance through automatic contact tracing. Instead, the use of the solution from Apple and Google was suggested (ironically, from the frying pan to the fire).

Immuni has been active since the first week of June, the source code is public (except of course the part implemented by GAEN) and the financial costs have been limited compared to several other countries. However, so far the system has been mainly ignored by the Italian population. The daily number of new diagnosed citizens that have notified other citizens through Immuni about the possible exposures has always been an extremely small fraction of the daily number of infected citizens.

Certainly, the switch in a week from PEPP-PT to GAEN did not give a clear message to citizens about the actual plan of the government about leveraging such apps to contain the spread of the virus. Moreover, the strong accusations about alleged mass surveillance programs might have impacted on the trust of citizens towards such new systems. While the app is open source, the management of the github repository has been problematic and in some cases embarrassing (e.g., some issues have remained open without any answer for weeks[10], the discussions were not moderated and in some cases monopolized by trolls). Contrary to the Corona-Warn-App in Germany, the Immuni app reports statistics about operational and epidemiological data like the number of warnings shown to users back to the central system so it can derive statistics about app usage and epidemiological information [34]. The official documentation of the system presents an ideal world where citizens can use the app preserving privacy and protecting themselves and others. However, the real world is different and unfortunately the increasing threats and attacks that have been documented in the last months have been substantially ignored. Clearly, citizens are asked to use the app without first informing them about the actual risks. Last but not least, Immuni seems to be poorly connected to the regional health system and in many cases local health authorities have preferred to completely ignore Immuni when citizens asked to upload their data through the app.

With the recent acceleration of the infection rate the hope that Immuni could somehow help seems to have vanished and even the recommendations from the Italian government often focus on traditional practices only (e.g., use of masks and social distancing), seemingly avoiding to disorient citizens with something considered ineffective. Nevertheless, precious resources are still spent to keep the system active (e.g., the management of the app is now in charge of Sogei, a company controlled by the Ministry of the Economy and Finance, and 4 million euros have recently been allocated to manage a national call center).

4.3 SwissCovid (Switzerland)

Deployment and usage. SwissCovid, the Swiss app, was developed based on the DP-3T project. In Switzerland, DP-3T is often presented as the project which served as basis for the GAEN protocol of Apple and Google. Most of the development was financed by Swiss academic institutes. Additionally, the Federal Office for Public Health (FOPH) had a

[9]PEPP-PT is based on a centralized tracing approach and was heavily criticized, due to privacy concerns related to possible misuse of data held by the central entity.

[10]https://github.com/immuni-app/immuni-documentation/issues/114

budget of nearly 5 million Swiss francs which was roughly equally split in development costs, exploitation costs, and advertisement costs [1, p. 16].

SwissCovid was officially deployed on June 25, 2020. As of end of October, the application was downloaded 2.7 million times. However, this does not reflect well the true usage of the app. Estimating the number of active users is not an easy task. For this, FOPH uses the fact that the app sends fake reports at random to the server, on average once every 5 days, in order to *hide* to the network when a true report is submitted. (This is one of the rare features which is implemented by the app, the rest being totally outsourced to GAEN.) The daily number of activations is thus the number of received fake reports multiplied by 5. At the end of October, it was 1.8 million. Hence, the adoption rate is pretty high: 21% of activations per inhabitant. (The population of Switzerland is of 8.55 million.) Note that this number is only an estimate which could be inaccurate if, e.g., affected by fake reports that could be generated by malicious entities with the goal to corrupt the tally.

Von Wyl et al. [48] investigated the reasons that prevent more people from using SwissCovid. For 37%, this is due to a perceived lack of usefulness. Indeed, for people who never stay close to an unknown person for several minutes (e.g., because they never use public transport, work from home and make all their meetings by video conference), the app is useless. For 23%, this is simply due to not having a suitable smartphone or operating system (GAEN works neither on old smartphones, nor on recent Chinese smartphones due to US regulations, nor on deGoogled Android phones, and, of course, only on Android and iOS). For 22%, there are concerns about privacy. It is likely that this includes people who do not want to keep Bluetooth activated because of the regular intrusion vulnerabilities which are discovered. There are various other reasons such as concerns about battery usage, doubts about the severity of the pandemic, mistrust in science or the government, etc. Another factor is that SwissCovid does not promise to provide any pleasant information to the user, only bad news in case of encountered infectious contact.

Effectiveness. SwissCovid sets two GAEN parameters which define the sensibility in the proximity detection. These parameters were increased twice since SwissCovid was deployed, so that it would eventually detect device proximity. The first parameter is a signal attenuation threshold set to 55 dB. In a lab experiment it was shown that two devices at a distance of 1.5m have a probability of 57.3% to observe an attenuation less than this threshold [14]. At a distance of 3m, this probability is somewhat lower: 45.6%. This shows that there are good chances that a device in proximity is not detected, or, that a device farther away is detected as being in close proximity. On top of that, a second threshold set to 63 dB accepts contacts at a larger distance but counts the duration of the contact with a coefficient $\frac{1}{2}$. (The probability at 3m is of 84.2% in lab conditions.) This would suggest that SwissCovid captures nearly everything. This is, however, not the case because it scans Bluetooth for a few seconds only every five minutes. A contact appearing between two scans is not detected. A brief contact during one scan (like someone passing by on a corridor) counts as a five-minute contact. Unsurprisingly, there are many false positives and false negatives. An experiment done by Leith and Farell in real conditions in a tram showed that there is little correlation between the attenuation and the distance and that SwissCovid actually detects no exposure [26].

Usefulness. To report infection, a user needs a 12-digit one-time access code which is delivered by the health authorities. FOPH also monitors the number of entered access codes (hence the number of reporting users) and the number of users who called the hotline after having received an alert. At the end of October, there were about 1038 reports and 626 calls daily. At the same time, the number of daily cases was 8028[11].

As we can see, the fraction of cases which are reported in the SwissCovid system is ca. 13%. By using an investigation suggesting that 53% of users who receive an alert call the hotline, von Wyl estimates the number of alerts based on the number of calls [47]. This way, the number of alerts must be around 1200 every day. Hence, each reporting user generated 1.2 alerts on average in this period. How many of them revealed to be positive and not suspected to be at risk otherwise is unknown.

However, a study from the promoters of SwissCovid [38] claims that SwissCovid was useful to discover 65 cases (out of 12 456 during September) in Switzerland, while there were 1695 calls (hence about 3200 alerts by the same approximation technique as above). This is based on a survey of clinical reports for people who were tested positive: one tick-box field indicates the reason to be tested. Most of the time it is because of symptoms, but it can be due to investigations, an alert by SwissCovid (which allows to get a free test by law), or other reasons to specify. The number of forms with "SwissCovid" indicated as a reason was 41. As many forms were not filled, the true number of SwissCovid-motivated cases was estimated to 65 (hence 0.52% of all cases). This does not prove that tested people were not aware about the risk to be contaminated otherwise. Actually, they could have been in quarantine already, but used the

[11] Average over the 7-day period October 28–November 3
https://www.experimental.bfs.admin.ch/expstat/fr/home/methodes-innovation/swisscovid-app-monitoring.html
https://www.bag.admin.ch/bag/en/home/krankheiten/ausbrueche-epidemien-pandemien/aktuelle-ausbrueche-epidemien/novel-cov/situation-schweiz-und-international.html

alert as a reason to get a test. Nevertheless, even by assuming that those 65 cases are genuinely discovered by SwissCovid, we can see that 65 useful cases over 3200 alerts means 2% of useful alerts. In the investigation by von Wyl [47], this ratio was determined to be 1.75%, i.e., slightly lower. The same survey also shows that SwissCovid generated 5% of the quarantines. Hence, the 'cost' in terms of quarantines is certainly much bigger than the benefit in terms of new discovered cases.

An interesting point on SwissCovid (and maybe other apps) is that they utilize a Content Delivery Network (CDN) provided by Amazon. The need for using a CDN for a local service to a population of 8.55 million inhabitants is questionable. Our investigation has shown that depending on where users are located (or which VPN they use), a local Amazon server will respond (for example in France, the server is based in Netherlands). The content is obtained from the servers of Swiss Federal offices and signed by them so that Amazon cannot tamper with the content, but Amazon can still exploit traffic metadata.

A lost race. Since May 2020, people have been warning that trying to chase for people who may have been infected by someone who was diagnosed may be a lost race [21]. Indeed, it was shown that only 19% of the cases are responsible for 80% of the transmissions [2]. Hence, *forward contact tracing* is likely to fail. Lambert [23] estimates that if the adoption rate is lower than 70%, there is no chance automated contact tracing can lower a reproduction rate R_0 from 1.3 down to 1. Something which could be more efficient to defeat the pandemic is to make *backward contact tracing*: to try to identify the origin of the contamination of people who were diagnosed. This could indeed discover people who contaminate many people without even knowing. As the SwissCovid infrastructure is made in a way that people report keys which were used when they started to be contagious (hence, after contamination), this would require some important changes.

Broken promises. The DP-3T project was aiming at objectives [43] which have not been met:

- to give data to epidemiologists — this goal was actually dropped after it became apparent that the GAEN implementation would not allow it;
- to be open source — the implementation of the protocol is in GAEN which is not open source;
- to be decentralized — it is rather GAEN-distributed in local storage;
- to have no false positives — it has;
- to have no false negatives — it has;
- to make false encounters impossible — relay attacks work very well;
- to be privacy-preserving — users can be tracked over Bluetooth and diagnosed users can be identified;
- and to be interoperable with other countries — the Swiss law on data protection is currently insufficient for the European countries to accept interoperability.

The law says that SwissCovid should be terminated if it reveals to be insufficient to fight against the pandemic. However, no objective criteria were defined to assess this and no calendar was set to make the assessment. As a matter of fact, the developers of SwissCovid are members of the Swiss COVID-19 Science Task Force which makes recommendations for the Federal government. Hence, developers are also auditors and the voice of science. The promoters often claim utility with shallow arguments. They acknowledge that performance is low and blame the low acceptance for this, or the length taken by the local health authorities for testing and delivering the access codes.[12] They also criticize an inappropriate debate about privacy threats for SwissCovid and ask for more faith from people. However, the privacy debate is the one they created by deciding to prevent authorities from being able to collect data which could have been useful, and making promises for privacy protection which were not met.

5 Conclusion

Digital contact tracing is one helpful tool to support manual tracing efforts in the current pandemic situation in order to contain the infection process and thereby save human lives and reduce the pandemic's adverse impact on member states' societies and economy. However, tracing approaches adopted in a number of member states are based on the Exposure Notification API (GAEN) from Apple and Google. This approach involves significant technological and political problems. First, as shown by a number of researchers, GAEN from Apple and Google has *fundamental security and privacy problems*. Second, current GAEN-based apps hand over an unpredictable amount of power in the form of user data to giant data collectors and threatens not only the technological sovereignty of member states but also opens up their public health systems to the influence of these technology giants. Third, decentralized tracing apps based on GAEN do not seem to help, since due to its decentralized nature it is not possible to understand whether and how they are used by users and whether they are effective. Hence, these tracing approaches are *missing a comprehensive digital infrastructure* enabling privacy-preserving feedback on the

[12] On October 31, the average number of days between the beginning of symptoms and the delivery of an access code was 5.85 days, the median being between 4 and 5.

operation and effectiveness of these apps. Hence, we need substantial changes to the current solutions in the form of a *comprehensive technology update* that includes the replacement of GAEN, in order to establish trust in the digital contact tracing systems and enable them to become effective against the COVID-19 pandemic.

One of the main issues with the widespread involuntary adoption of the GAEN solution is that it stopped the continuous process of improvement of digital contact tracing and its security in general. Some suggested improvements such a PRONTO-C2 [6] have not been deployed (because they are not in the interests of Apple and Google?). One should also not forget that as an alternative, Apple and Google could have instead of or in addition to GAEN given an optimized access to their Bluetooth APIs, but they chose not to do so. Thus, fully state-controlled apps such as StopCovid (FR) did have technical issues (excessive battery drain, having to run the app in the foreground, etc.) that are impossible to resolve. Apple and Google should have given access to their Bluetooth APIs and countries should have demanded this more adamantly from these big vendors. However, they have not done so.

From our perspective, further academic discussion and profound studies are urgently needed. Governments should not avoid the debate for the fear of admitting negative results, but rather promote honest discussion and critical evaluation in the interest of all citizens concerned.

References

[1] *20.040 message from the Federal Council. Message concernant la modification urgente de la loi sur les épidémies en lien avec le coronavirus (Systéme de proximité)*. https://www.admin.ch/opc/fr/federal-gazette/2020/4361.pdf. Schweizerische Eidgenossenschaft, May 20, 2020.

[2] Dillon Adam et al. "Clustering and superspreading potential of severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) infections in Hong Kong". In: (2020).

[3] Ross Anderson. *Contact Tracing in the Real World*. https://www.lightbluetouchpaper.org/2020/04/12/contact-tracing-in-the-real-world/. Apr. 12, 2020.

[4] Apple and Google. *Exposure Notification: Cryptography Specification, v1.2*. https://www.apple.com/covid19/contacttracing. Apr. 2020.

[5] Gennaro Avitabile, Daniele Friolo, and Ivan Visconti. *TEnK-U: Terrorist Attacks for Fake Exposure Notifications in Contact Tracing Systems*. Cryptology ePrint Archive, Report 2020/1150. https://eprint.iacr.org/2020/1150. 2020.

[6] Gennaro Avitabile et al. *Towards Defeating Mass Surveillance and SARS-CoV-2: The Pronto-C2 Fully Decentralized Automatic Contact Tracing System*. Cryptology ePrint Archive, Report 2020/493. https://eprint.iacr.org/2020/493. 2020.

[7] Dorothee Bär et al. *Corona Apps: Die globalen Konzerne haben eine Chance verpasst*. https://www.faz.net/aktuell/politik/inland/corona-apps-die-globalen-konzerne-haben-eine-chance-verpasst-16785681.html. June 8, 2020.

[8] Lars Baumgärtner et al. "Mind the GAP: Security & Privacy Risks of Contact Tracing Apps". In: *arXiv preprint arXiv:2006.05914* (2020).

[9] Michael Böhme. *Statistical Analysis of the daily diagnosis keys of the German COVID-19 Tracing-App (Corona-Warn-App)*. https://micb25.github.io/dka/. Oct. 28, 2020.

[10] Isobel Braithwaite et al. "Automated and partly automated contact tracing: a systematic review to inform the control of COVID-19". In: *The Lancet Digital Health* 2.11 (Nov. 2020). https://doi.org/10.1016/S2589-7500(20)30184-9, e607–e621. ISSN: 2589-7500. DOI: 10.1016/S2589-7500(20)30184-9.

[11] Kristen V. Brown. *Alphabet's Verily Plans to Use Big Data for Health Insurance*. https://www.bloomberg.com/news/articles/2020-08-25/alphabet-s-verily-plans-to-use-big-data-as-health-insurance-tool. Aug. 25, 2020.

[12] Andrew Crocker, Kurt Opsahl, and Bennett Cyphers. *The Challenge of Proximity Apps For COVID-19 Contact Tracing*. 2020. URL: https://www.eff.org/de/deeplinks/2020/04/challenge-proximity-apps-covid-19-contact-tracing.

[13] Paul-Olivier Dehaye and Joel Reardon. *SwissCovid: a critical analysis of risk assessment by Swiss authorities*. https://arxiv.org/abs/2006.10719. 2020. arXiv: 2006.10719 [cs.CR].

[14] Federal Office of Information Technology, Systems and Telecommunication FOITT. "SwissCovid Exposure Score Calculation". In: (Sept. 2020). eprint: https://github.com/admin-ch/PT-System-Documents/blob/master/SwissCovid-ExposureScore.pdf.

[15] Rosario Gennaro, Adam Krellenstein, and James Krellenstein. *Exposure Notification System May Allow for Large-Scale Voter Suppression*. https://static1.squarespace.com/static/5e937afbfd7a75746167b39c/t/5f47a87e58d3de0db3da91b2/1598531714869/Exposure_Notification.pdf. 2020.

[16] Golem. *Entwicklung von Corona-App kostet 20 Millionen Euro*. https://www.golem.de/news/bundesregierung-entwicklung-von-corona-app-kostet-20-millionen-euro-2006-149033.html. June 11, 2020.

[17] Italian Government. *Immuni*. https://www.immuni.italia.it/dashboard.html. Oct. 28, 2020.

[18] Yaron Gvili. *Security Analysis of the COVID-19 Contact Tracing Specifications by Apple Inc. and Google Inc.* Cryptology ePrint Archive, Report 2020/428. https://eprint.iacr.org/2020/428. Apr. 2020.

[19] Giles Hogben. *Presentation at Privacy Aspects of Contact Tracing Workshop Session 1*. https://satcfrontier.eng.uci.edu/events/privacy-contact-tracing/. Oct. 2, 2020.

[20] Vincenzo Iovino, Serge Vaudenay, and Martin Vuagnoux. *On the Effectiveness of Time Travel to Inject COVID-19 Alerts*. Cryptology ePrint Archive, Report 2020/1393. https://eprint.iacr.org/2020/1393. 2020.

[21] Sadamori Kojaku et al. *The effectiveness of backward contact tracing in networks*. 2020. arXiv: 2005.02362 [q-bio.PE].

[22] Stefan Krempl. *Corona-Pandemie: Robert Koch-Institut will eine App für alles*. Heise.de. https://www.heise.de/news/Corona-Pandemie-Robert-Koch-Institut-will-eine-App-fuer-alles-4915824.html. Sept. 30, 2020.

[23] Amaury Lambert. "A mathematical assessment of the efficiency of quarantining and contact tracing in curbing the COVID-19 epidemic". In: *medRxiv* (2020). DOI: 10.1101/2020.05.04.20091009. eprint: https://www.medrxiv.org/content/early/2020/05/08/2020.05.04.20091009.full.pdf.

[24] Doug Leith and Stephen Farrell. *Contact Tracing App Privacy: What Data Is SharedBy Europe's GAEN Contact Tracing Apps*. https://www.scss.tcd.ie/Doug.Leith/pubs/contact_tracing_app_traffic.pdf. Oct. 30, 2020.

[25] Doug Leith and Stephen Farrell. *Testing Apps for COVID-19 Tracing (TACT) - TEK Survey*. https://down.dsg.cs.tcd.ie/tact/tek-counts/. Oct. 30, 2020.

[26] Douglas J. Leith and Stephen Farrell. "Measurement-based evaluation of Google/Apple Exposure Notification API for proximity detection in a light-rail tram". In: *PLOS ONE* 15.9 (Sept. 2020), pp. 1–16. DOI: 10.1371/journal.pone.0239943. URL: https://doi.org/10.1371/journal.pone.0239943.

[27] Open letter. *Joint Statement on Contact Tracing: Date 19th April 2020*. https://drive.google.com/file/d/1OQg2dxPu-x-RZzETlpV3lFa259Nrpk1J/view. Oct. 30, 2020.

[28] Mallorca Magazin. *Immer mehr Menschen laden sich spanische Corona-Warn-App herunter*. https://www.mallorcamagazin.com/nachrichten/lokales/2020/09/15/83757/immer-mehr-menschen-laden-spanische-corona-warn-app-herunter.html. Sept. 15, 2020.

[29] Carrie Mihalcik. *Google faces $5 billion lawsuit for tracking people in incognito mode*. https://www.cnet.com/news/google-faces-5-billion-lawsuit-for-tracking-people-in-incognito-mode/. CNET.com, June 3, 2020.

[30] Italy Ministero della Salute. *Immuni - Exposure Notifications Italy*. https://apps.apple.com/it/app/immuni/id1513940977?l=en. June 7, 2020.

[31] OVB Online. *Spanien entwickelt Corona-Warn-App "Radar Covid"*. https://www.ovb-online.de/weltspiegel/spanien-corona-warn-app-bluetooth-nutzer-covid-19-sars-cov-2-zr-90022661.html. Aug. 13, 2020.

[32] University of Oxford. *Digital contact tracing can slow or even stop coronavirus transmission and ease us out of lockdown*. https://www.research.ox.ac.uk/Article/2020-04-16-digital-contact-tracing-can-slow-or-even-stop-coronavirus-transmission-and-ease-us-out-of-lockdown. Apr. 16, 2020.

[33] PEPP-PT. *pepp-pt*. 2020. URL: https://www.pepp-pt.org/content.

[34] *Privacy-Preserving Analytics*. https://github.com/immuni-app/immuni-documentation/blob/master/Privacy-Preserving%20Analytics.md. Nov. 23, 2020.

[35] DP-3T project. *Decentralized Privacy-Preserving Proximity Tracing*. https://github.com/DP-3T/documents. 2020.

[36] *Replay attack "in the past"*. https://github.com/immuni-app/immuni-app-android/issues/278. 2020.

[37] Mathieu Rosemain. *France fines Google $57 million for European privacy rule breach*. https://www.reuters.com/article/us-google-privacy-france-idUSKCN1PF208. Jan. 21, 2019.

[38] Marcel Salathé et al. "Early Evidence of Effectiveness of Digital Contact Tracing for SARS-CoV-2 in Switzerland". In: *medRxiv* (2020).

[39] Government of Switzerland. *SwissCovid App*. https://bag-coronavirus.ch/swisscovid-app/. 2020.

[40] Deutsche Telekom. *Vertragsdokumente zur Corona-App mit der Telekom und SAP*. https://fragdenstaat.de/anfrage/vertragsdokumente-zur-corona-app-mit-der-telekom-und-sap/513354/anhang/Vertrag%20Telekom.pdf. June 3, 2020.

[41] Deutsche Telekom and SAP. *Corona-Warn-App - The Official COVID-19 Exposure Notification App for Germany*. https://github.com/corona-warn-app. June 7, 2020.

[42] *The Swedish Data Protection Authority imposes administrative fine on Google*. https://edpb.europa.eu/news/national-news/2020/swedish-data-protection-authority-imposes-administrative-fine-google_en. Mar. 11, 2020.

[43] Carmela Troncoso et al. "Decentralized privacy-preserving proximity tracing". In: *arXiv preprint arXiv:2005.12273* (2020).

[44] Serge Vaudenay. *Analysis of DP-3T*. Cryptology ePrint Archive, Report 2020/399. Apr. 2020. URL: https://eprint.iacr.org/2020/399.

[45] Serge Vaudenay. *Centralized or Decentralized? The Contact Tracing Dilemma*. Cryptology ePrint Archive, Report 2020/531. https://eprint.iacr.org/2020/531. May 2020.

[46] Serge Vaudenay and Martin Vuagnoux. *Analysis of SwissCovid*. 2020. URL: https://lasec.epfl.ch/people/vaudenay/swisscovid/swisscovid-ana.pdf.

[47] Viktor von Wyl. "The contribution of the SwissCovid digital proximity tracing app to pandemic mitigation in the Canton of Zurich". In: ().

[48] Viktor von Wyl et al. "Drivers of acceptance of COVID-19 proximity tracing apps in Switzerland". In: *medRxiv* (2020). DOI: 10.1101/2020.08.29.20184382. eprint: https://www.medrxiv.org/content/early/2020/11/12/2020.08.29.20184382.full.pdf. URL: https://www.medrxiv.org/content/early/2020/11/12/2020.08.29.20184382.

Appendix

In this section we provide more details about the security and privacy attacks on GAEN.

Relay Attacks. In a relay attack (also known as *wormhole* attack) an adversary aims to generate fake encounters (contacts) resulting in a massive number of false exposure notifications. The attack is performed by capturing the temporary pseudonymous identifiers, called *Rolling Proximity Identifiers (RPIs)*, emitted by tracing apps at a particular location and sending them via the Internet to other (distant) places where the RPIs are *replayed* so that other devices will capture these relayed RPIs. If any of the relayed RPIs originates from a person that will report herself as infected, the recipients of the RPIs will receive false exposure notifications, even though a real contact has not taken place. To maximize the effect of this attack the adversary will gather RPIs from locations with expected high infection rates (e.g., locations known to suffer from a COVID-19 outbreak) and relay them to other crowded places in big cities, e.g, shopping centers or railway stations, i.e., the adversary creates a *wormhole* to capture and broadcast RPIs. For instance, Baumgärtner et al. [8] have conducted attacks in three cities (Darmstadt, Frankfurt, and Marburg) in Germany on the Corona-Warn-App (Germany) as well as SwissCovid apps. They show that an adversary can effectively use regular smartphones to capture and relay RPIs among those cities.

Dehaye and Reardon [13] propose an effective way of realizing such attacks in the context of their critical analysis of the *SwissCovid* app also based on the GAEN API. They demonstrate how a malicious Software Development Kit (SDK) could be used to inject malicious functionality for relaying RPIs in a presumed secure application incorporating such a library, thereby transforming devices of benign users into malicious relay stations for RPIs without the knowledge of the device owners.

Time-machine attacks As demonstrated by Iovino et al. [20], it is possible for an adversary in proximity of a victim to remotely manipulate the clock of the victim's phone. This way, the adversary can 'send' the victim's phone to the past, replay an RPI derived from a (outdated) diagnosis key advertised on the contact tracing server, and then wait until the victim's clock is restored to the correct time. The tracing app will then see that the replayed RPI belongs to a diagnosis key and raise an alert. This attack takes only a few seconds and requires $10-equipment [20].

There are several ways to take control of the clock of a remote phone. The simplest configuration assumes that the adversary and the victim are connected to the same Wi-Fi network and that information from the data network does not overrule the clock setup (it could be jammed otherwise). In this setting, the phone adjusts its clock by making NTP queries which can be intercepted. Some phone configurations do not make NTP queries so often but they can be triggered by denial of service attacks. The naïve attack uses one time jump and beams the key for 15 minutes over Bluetooth. It can be boosted by making several time jumps at very short intervals when the app scans Bluetooth. Indeed, a successful attack can be staged using scans done during a few seconds only every five minutes. This way, the whole attack takes only 20 seconds (but the alert may come a few hours later).

To boost the attack, an amplifier for Bluetooth

can be used. Since broadcasts are unidirectional, they can thus be sent to the victim devices from farther away[13]. Since the attack needs only 20 seconds of exposure, we can imagine a car driving in the target area, broadcasting outdated RPIs and remotely changing date and time of mobile phones with the rogue base station. This is a very feasible attack. All smartphones we tested are vulnerable to it and we successfully tested the attack on Corona-Warn-App (DE), Immuni (IT), SwissCovid (CH) as well as NHS-covid-19 (UK).

Tracing Forgery: The Terrorist Attack. In a tracing forgery attack, the adversary attacks the integrity of a GAEN-based system by colluding with infected users to upload adversary-chosen Temporary Exposure Keys (TEKs) to the tracing server. As a result, every user who has captured RPIs derived from those TEKs will receive a false exposure alarm. This class of attacks exploit infected users who want to monetize their infection status, i.e., by uploading TEKs chosen by the adversary or selling their TANs (Transaction Authentication Numbers) required for authenticating their infection status to the tracing server for money. For example, Avitabile et al. [5] propose a smart-contract based on-line market where the adversary and infected users can trade TANs in an anonymous way but with integrity guarantees.

Movement profiling attacks. Although apps built on the GAEN system typically do not explicitly capture or record the true identity of individual users, they allow, however, to collect movement profiles of infected users that use the system to warn others and potentially even identify these individuals. As suggested, e.g., by Vaudenay and Vuagnoux [46], recent research by Baumgärtner et al. [8] has shown that such attacks are possible, since infected users need to upload their so-called Temporary Exposure Keys (TEKs) to the tracing server. The temporary identifiers broadcast by the tracing apps into their vicinity, so-called Rolling Proximity Identifiers (RPIs), are derived from these TEKs. With the help of the TEKs, it is therefore possible to link all RPIs of a particular person for the duration of 24 hours. This makes it possible to use recorded RPI observations in a particular area to construct movement profiles of infected persons. Since movement profiles of persons are typically unique and differ from person to person, further identifying information about users can be extracted. For example based on where a user located during the night can be used to draw conclusions about their place of residence. If one looks at the predominant location during working hours, it is likely possible to identify a person's workplace. The more pieces of such information one can combine, the more likely it is that one can uniquely identify the person.

A threat to privacy from such an attack could come from anyone who can install observation nodes in an area to take advantage of RPI information. The GAEN technology thus exposes personal and medical data to arbitrary third parties. Google is aware of this threat and has said it is considering reducing the validity period of TEKs to 6 hours in an attempt to mitigate it [19]. However, so far, nothing has been changed.

The attack implementation of Baumgärtner et al. used commercially available and inexpensive tools such as Bluetooth sniffers (applicable as an app on smartphones or Raspberry Pis) to ensure the necessary physical proximity. We have also demonstrated that for some smartphones (half of the tested smartphones), COVID-19 app users can be traced continuously due to a bug in GAEN that has been confirmed by Google[14]. It is likely related to the used Bluetooth drivers. This bug causes that when the RPI is updated, the BD-ADDR of the Smartphone is not (for the duration of one or more packets). Thus, an adversary can use this mismatch in the updating schedule to bind the new RPI with the previous one and consequently trace users for more than 10 minutes. For further information please refer to a demonstration video showing the exploitation of this vulnerability entitled 'Little Thumb' that is available online[15].

Some have voiced the opinion that this bug is not related to GAEN as such. However, this is disputed by the fact that the previous version of the SwissCovid app (called prestandard) that did not NOT use GAEN is NOT vulnerable to this issue. The same applies to StopCovid of France. Google has communicated that they are still trying to fix this issue.

[13]Moreover, an attack variant utilizing a rogue GSM base station has a potential a radius of 35km (maximum range in GSM)
[14]https://github.com/google/exposure-notifications-internals/commit/8f751a666697c3cae0a56ae3464c2c6cbe31b69e
[15]https://vimeo.com/453948863