SwissCovid in the Perspective of Its Goals

SERGE VAUDENAY, EPFL, Switzerland
MARTIN VUAGNOUX, base23, Switzerland

SwissCovid is the Swiss digital contact tracing app, which was deployed to help fighting against the COVID-19 pandemic. After a year of activity, it is high time to evaluate how effective it has been in its mission. At the highest peak, about 22% of the Swiss population was actively using SwissCovid. The activity of SwissCovid follows the curve on the number of COVID-19 cases. However, performances are rather poor. About 1% of the cases may have been discovered by SwissCovid and much less than 2% of SwissCovid alerts may have been useful, while SwissCovid generates 5% of the quarantines. The measure of the proximity of each encounter and its duration are also imprecise. It further comes with security and privacy issues: adversaries can inject false alerts for SwissCovid users and users can be tracked. On top of that, SwissCovid contributes to strengthen the monopoly of Apple and Google and to make users and their data captive of these giants. It also digs the digital divide. Contrarily to the original plan, the implementation is not open source and the law was twisted to fit the constraints by Apple and Google. Therefore, SwissCovid did not meet its goals.

CCS Concepts: • **Security and privacy** → **Social aspects of security and privacy**; **Domain-specific security and privacy architectures**; *Distributed systems security;*

Additional Key Words and Phrases: COVID-19, digital contact tracing

ACM Reference format:
Serge Vaudenay and Martin Vuagnoux. 2022. SwissCovid in the Perspective of Its Goals. *Digit. Threat.: Res. Pract.* 3, 3, Article 29 (September 2022), 17 pages.
https://doi.org/10.1145/3480465

1 INTRODUCTION

The 2020 pandemic set up the theater for the development and massive deployment of *digital contact tracing*. One dominant technology is maintained and controlled by Apple and Google: **Google and Apple Exposure Notification** (**GAEN**, throughout this article). It relies on people wearing smartphones with Bluetooth and network connections turned on. Close contacts between people create some information (such as tracking information, date, duration, and power of signal) which is stored in the smartphones. If someone is diagnosed with COVID-19, this person can initiate a report by his phone and people who have met this person can be notified that they may be at risk of having been infected.

SwissCovid is the Swiss digital contact tracing app. It is based on GAEN and was deployed on June 25, 2020. It followed the research made by the **DP-3T project (Decentralized Privacy-Preserving Proximity Tracing)**.

Authors' addresses: S. Vaudenay, EPFL, IC/LASEC, Station 14, Lausanne, 1015, Switzerland; email: serge.vaudenay@epfl.ch; M. Vuagnoux, base23, case postale 6482, Geneva 6, 1211, Switzerland; email: martin@vuagnoux.com.

2576-5337/2022/09-ART29 $15.00
https://doi.org/10.1145/3480465

The project officially started on April 1, 2020, within the **PEPP-PT project** (Pan-European Privacy-Preserving Proximity Tracing). However, DP-3T and PEPP-PT split a few days later following a dispute on design strategies [21].

In smartphones, permanent access to Bluetooth is blocked by Apple and Google unless the app runs in foreground and drains the battery. The DP-3T vs. PEPP-PT conflict was arbitrated by these giants who declared they would jointly support the development of a decentralized system. This was enough for governments to leave PEPP-PT and move to the DP-3T solution. However, instead of giving Bluetooth access to the app, Apple and Google decided to implement their protocol into GAEN and to give access to GAEN only.

DP-3T continued to develop the technology based on GAEN in heavily stressful conditions and tight schedules. There has been room for evaluation and debate in the Swiss Parliament. As we cover in Section 5.3, the Parliament adopted the legal frame for SwissCovid within the **Law on Epidemics (LEp)** whose spirit has been immediately twisted in the **Ordinance on the Proximity Tracing System for Coronavirus (OSTP)**. Finally, SwissCovid was launched five days later, on June 25, 2020. It promised to save lives while maintaining high security and privacy standards. Data would never leave the smartphone, alerts would be raised only in risks of infection, and usage would be fully voluntary.

Most of the shortcomings of SwissCovid have been known since the beginning: that the adoption rate would limit the impact, that detecting proximity by Bluetooth was not reliable, that there would be false positives and false negatives, that GAEN would not be open source, that there would be security and privacy issues [20, 22]. These alarming facts did not necessarily imply that SwissCovid would be useless, but they have been absent from, or downplayed in, public communication.[1] We believe that this was made possible by the scientific lobbying of developers. There is a big intersection between people in DP-3T and the Swiss National COVID-19 Science Task Force which releases science briefs and advises the Federal Government. DP-3T also benefits from the support of the EPFL infrastructure and its media service which introduces a strong bias in the scientific communication to the press. We review all these issues in this article.

We believe that the move to the SwissCovid solution was a bad choice. Its effectiveness was at most marginal. It did not avoid strong measures such as mandatory teleworking, closure of restaurants, ban on social activities in person, and so on. Besides, the ethical balance is negative. White and van Basshuysen actually concluded that the choice for a decentralized system was ethically wrong [23]:

> "The privacy advocates' arguments have been influential in debates about digital contact tracing and, backed by Apple and Google's strategy to make it difficult to produce a centralized app which can function effectively on their smartphone systems, they have apparently led policy makers in most countries to implement decentralised systems. It follows from our risk assessment that these policy makers may well have been misled, as centralised systems are in fact the option that could minimise overall ethical risks."

After a year of activity, we believe it is time to look back at the original goals and the current results. In this article, we survey the SwissCovid structure and performances. We list the identified security and privacy issues. Finally, we review the goals of SwissCovid and show that they were not met.

2 THE GAEN/SWISSCOVID STRUCTURE

The infrastructure relies on users, smartphones, and servers. Smartphones use GAEN (installed by default) and a Government-developed app which plays the role of an interface between GAEN, the user, and servers.

[1] https://lasec.epfl.ch/people/vaudenay/swisscovid/ownanalysis.html — Own Analysis of SwissCovid.

2.1 GAEN

GAEN is installed by default on recent smartphones using Android or iOS.[2] GAEN communicates with other GAENs over Bluetooth, and with the app in the same smartphone. We describe how it works. For that, we distinguish several operations: *broadcast*, *reception*, *report*, and *matching*. GAEN maintains two databases for used keys and received keys. Entries are erased after 14 days.

Broadcast. Every rolling period (which is 24h), GAEN selects a new random key called the ***Temporary Exposure Key*** **(**TEK**)**. We denote by τ the time this key is used by the first time. Actually, τ is an integer obtained by dividing the Unix Epoch Time by 600, so that it is incremented every 10 minutes. TEK is stored together with τ in the GAEN database of used keys.

Let t be a current time during the rolling period starting at time τ. Every t (in 10-minute integer format) defines a new key called the **Rolling Proximity Identifier (**RPI**)**. This key is derived from TEK and t following an AES-based cryptographic function

$$\mathsf{RPI} = f(\mathsf{TEK}, t).$$

GAEN broadcasts over Bluetooth 4 times per second the string RPI||AEM where AEM is the **Associated Encrypted Metadata** which is computed by

$$\mathsf{AEM} = g(\mathsf{TEK}, \mathsf{RPI}) \oplus \mathsf{metadata}$$

and metadata encodes the power π by which the Bluetooth signal is sent. Each RPI||AEM is repeated many times during 15 minutes (on average; this duration is randomized) before it is replaced by a new one.

See the GAEN specifications for more details [1, 2].

Reception. Every 3–5 minutes, GAEN scans Bluetooth for a few seconds and collects all received RPI||AEM. It adds in the reception database the tuple $(\mathsf{RPI}, \mathsf{AEM}, t, p)$ where t is the scanning time and p is the power of the received Bluetooth signal. (See the GAEN specifications for more details [1].)

Report. When triggered by the app (it is intended to happen when the user is diagnosed with COVID-19), GAEN releases the list of stored (TEK, τ) values with τ larger than a date (which is defined by the app based on the possible earliest contagious date, which is determined by the health authority). Following this, the app is supposed to interact with the user who may decide to remove the TEK of some selected dates, then to submit the remaining pairs to the server.

Matching. Based on a new batch of (TEK, τ) pairs provided by the app (obained from the server), GAEN derives all RPI from those TEKs using the function f and checks if one appears in the database of the received RPI. Each match corresponds to a $(\mathsf{RPI}, \mathsf{AEM}, t, p)$ record. The consistency of t with the value used to derive RPI is verified with a tolerance of ± 2 hours. For each record, AEM can be decrypted using the function g to obtain metadata and extract π. Then, the signal attenuation $\pi - p$ is computed. Hence, each match adds an attenuation value in the list. GAEN sorts these values in three buckets: those with attenuation below a threshold t_1 (corresponding to close proximity), those with attenuation between t_1 and a threshold t_2, and those with attenuation larger than t_2 (corresponding to an excessive distance). For each matching, a duration of encounter is assigned. This duration is set to the elapsed time since the last scan, rounded to the minute, and upper bounded to 5 minutes. The sum of all durations is computed in each of the three buckets and returned to the app.

2.2 SwissCovid

The app is in charge of downloading the new TEK reported on the server, deciding whether to alert the user, and interacting with the user to upload a report to the server.

[2]GAEN does not work on excessively old smartphones. It is absent from deGoogled Android smartphones. It is also removed in recent Chinese smartphones due to U.S. regulation.

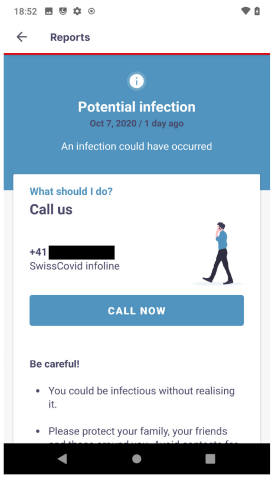

Fig. 1. Screen captures of SwissCovid raising an alert [10].

Routine Connections. The SwissCovid app connects to the server every 6 hours to get some configuration parameters which are used in the matching phase. Normally, these parameters are always the same. Every 4 hours (although this may vary), SwissCovid also retrieves new uploads to the server. As those uploads are posted at a given time on the server, SwissCovid can ask for only the ones which are posterior to the last check. SwissCovid obtains a new batch of (TEK, τ) pairs from the server and provides it to GAEN. For both connections, the answer is actually made by a **CDN (Content Delivery Network)**; for SwissCovid, it is the Amazon one.

Another routine task of SwissCovid is to submit a *fake* report to the server. This is to hide to the network when a genuine report would be issued, which is to hide that the user was diagnosed for privacy reasons. The waiting delay before issuing a fake report is random. The probability distribution is exponential. On average, it is 5 days.

Report. The upload operation is the main technical task of the SwissCovid app. Uploading requires the approval from the health authorities. Approval is given by the medical infrastructure to the user in the form of a 12-digit code (called *covidcode*). This code is a one-time authorization code which remains valid for 24 hours. It is generated randomly by another server: the covidcode server. The user can enter the covidcode in the app. The app will ask for verification of this code to the covidcode server who will return a JWT (*JSON Web Token*) authorization token. The app will then get the report from GAEN and forward it to the server together with this JWT. The server will verify the JWT to accept the report as genuine.

Matching. SwissCovid gets new TEKs from the server and provides them to GAEN for matching. SwissCovid sets the matching parameters, as instructed by the server: to $t_1 = 55$dB and $t_2 = 63$dB for GAEN. GAEN returns a total exposure duration for each of the three buckets. SwissCovid assigns these durations to coefficients $w_1 = 1$ for close proximity, $w_2 = 0.5$ for the intermediate one, and $w_3 = 0$ for the too far distances. These coefficients are also inserted by the server in the configuration parameters. (See the report from the federal authorities for more information [3].) Hence, close proximities fully count while intermediate proximities count for half and far proximities do not count. The total weighted duration of encounter is computed and compared to 15 minutes. If it is larger, the user is notified. Figure 1 shows a screenshot of an alert raised by SwissCovid.

3 SWISSCOVID PERFORMANCES

3.1 The Precision of Bluetooth Distances

Bluetooth is not made to measure distances. GAEN actually tries to guess the distance based on signal attenuation. (There is also an empirical calibration offset to the attenuation which depends on the hardware which is used.)

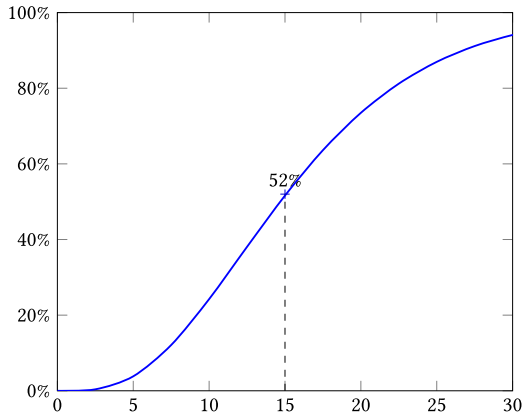

Fig. 2. Probability that an encounter in a moving situation is believed to have occurred longer than 15 minutes based on its real total duration.

The t_1 and t_2 thresholds are based on some lab-condition experiments [3]. There, the following probabilities was reported:

distance of encounter	1.5 m	2 m	3 m
Pr[attenuation $< t_1$]	57.3%	51.6%	45.6%
Pr[attenuation $< t_2$]	89.6%	87.5%	84.2%

We can see that the dependence on the distance is rather weak. Actually, Leith and Farrell [13] have shown in a real-life condition (a tram) that there was no correlation between the distance and the proximity deduction by GAEN.

3.2 The Precision of Exposure Time

The method to measure the exposure time is adapted to the situation where people do not move too much. In a moving situation, we can show that the precision is very loose. Indeed, GAEN scans on average once every 4 minutes so we need four scans to count for an exposure of 15 minutes. When people do not move, this works well. But in a moving condition, the probability that a scan occurs when people are close has an effect. The number of scans with close proximity we collect when the total duration of the proximity lasts $\lambda \times 4$ minutes follows a Poisson distribution with parameter λ: the probability that proximity happens during exactly k scans is $\frac{\lambda^k}{k!}e^{-\lambda}$. Hence, the probability to collect at least four scans is

$$\Pr[\text{spot at least 4 times}|\text{exposure of } \lambda \times 4\text{min}] \approx 1 - \left(1 + \lambda + \frac{\lambda^2}{2} + \frac{\lambda^3}{6}\right) e^{-\lambda}$$

We plot this probability in terms of the actual duration in Figure 2. Hence, in a moving situation, an effective proximity of 15 minutes in total is spotted with a 52% probability.

3.3 On Collecting Performance Data

The *Federal Office for Public Health* (FOPH) collects lots of data which are useful to estimate the performance of SwissCovid. Obviously, FOPH publishes the daily number of COVID-19 positive cases.[3] Next, there are

[3] https://www.bag.admin.ch/bag/en/home/krankheiten/ausbrueche-epidemien-pandemien/aktuelle-ausbrueche-epidemien/novel-cov/situation-schweiz-und-international.html — Coronavirus: Situation in Switzerland.

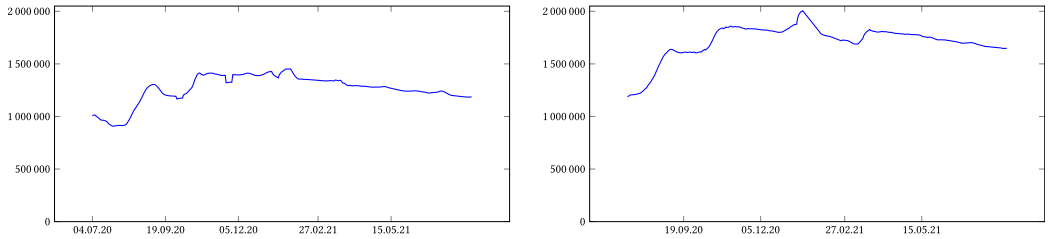

Fig. 3. Evolution of the number of activations of SwissCovid apps (averaged over a 7-day period) using two method: based on parameter queries (left) and based on fake report submissions (right).

SwissCovid-related data.[4] There is the number of *activations*, the number of *entered covidcodes*, the number of *calls to the infoline* by users following a notification by SwissCovid. FOPH also once reported the number of positive cases for which the test was motivated by a SwissCovid notification, as discussed below.

Activations. Measuring the daily number of activations was not easy. At the beginning, FOPH was doing this by dividing the number of queries to the server for parameters by 4, because apps make such queries every 6 hours. This leads to very imprecise results. Next, FOPH switched to multiplying the number of fake reports to the server by 5, because apps submit fake reports every 5 days on average. We plotted the evolution of the number of activations with time in Figure 3 using the two methods. As we can see, it has always been below 2 million. (The population of Switzerland is of 8.5 million.) There are sometimes glitches (e.g., in mid-January, the figures with the new method jumped by 25% in 2 days for no reason, after which FOPH suspended the count for 10 days with no explanation). Otherwise, this is rather stable, although slowly decreasing. In April 2021, the population actively using SwissCovid was estimated to be around 21%. As of August 2021, it was reduced to 19%.

Activity. The number of entered covidcodes measures how many users who were positively tested submitted a report. This gives an indication on the proportion of cases which are in the SwissCovid system. (In April 2021, this ratio was of 9%; it dropped to 3% in mid-August 2021.) The number of calls to the infoline underestimates the number of users who received a notification.

We plotted these two curves together with the number of cases on Figure 4. We plotted with no scale on purpose to show that all curves are nearly the same. For instance, for each date, #codes is approximated by #cases the day before multiplied by a scaling factor of 11%.

The singularity of the number of calls in December 2020 is due to a change in the system: instead of being invited to call an infoline, alerted users are now invited to fill up an online survey. The glitch in December is probably counting people who tested the survey.

Since the number of activations is too stable, we cannot really see the influence of it on the activity. We can deduce from this the following empirical rule:

> When the number of activations is constant, the number of entered covidcodes and the number of calls are proportional (possibly with a delay) to the number of cases.

3.4 Survey-Based Performance

People involved in the development, the deployment, and the promotion of SwissCovid did several surveys with the evident goal to prove that SwissCovid was useful and to influence more people to adopt it.

[4] https://www.experimental.bfs.admin.ch/expstat/en/home/innovative-methods/swisscovid-app-monitoring.html — SwissCovid App Monitoring.

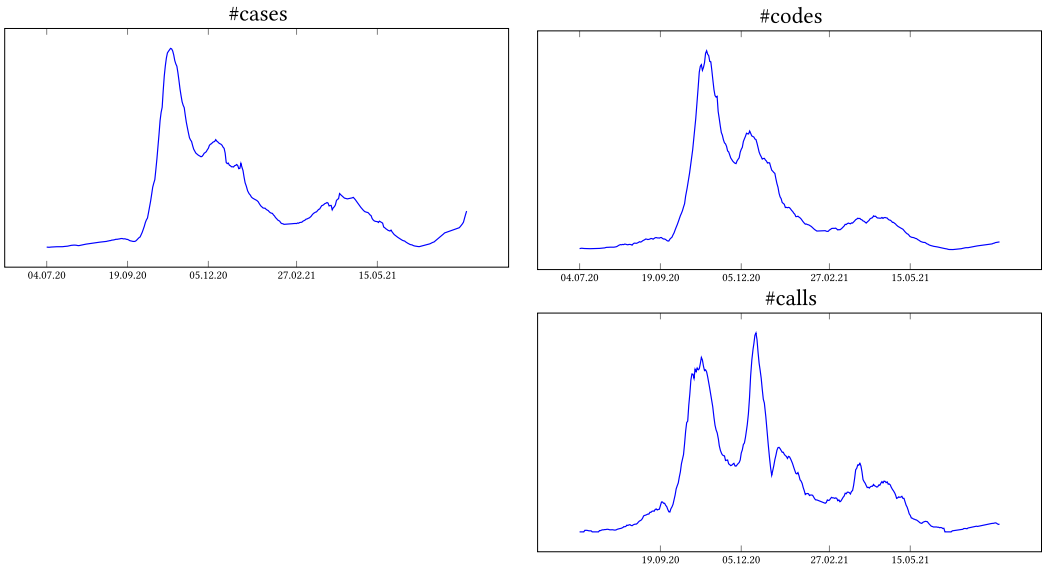

Fig. 4. Activity of SwissCovid (averaged over a 7-day period): number of COVID-19 positive cases in Switzerland, number of entered covidcodes in the SwissCovid system, and number of calls received by the infoline from people who were alerted by SwissCovid.

Adoption. The total number of downloads of SwissCovid is above 3 million for a country of 8.5 million inhabitants. However, the number of "activations" had a peak close to 1.9 million in early November 2020 (if we exclude the abnormal peak in January 2021). That represented 22% of the Swiss population.

In September-October 2020, Von Wyl et al. [25] performed a survey to investigate the reasons behind the lack of use of the app. The survey indicates that 36.8% of the surveyed perceived a lack of usefulness. Then, 22.8% of the interviewees did not have a proper phone to run the app. In the third position, 22.4% of them were concerned about privacy. The remaining 18% came with various other reasons such as not knowing the app, having doubts about the reliability of the technology, worries about battery usage, and the like.

"Early Evidence" of Effectiveness. On September 4, 2020, Salathé et al. issued a paper which was subsequently revised and published on December 16th [18]. The title 'Early Evidence of Effectiveness" announced some evidence of effectiveness and seemed to promise that more would come later. This paper essentially surveyed the SwissCovid activities during a given period. 12, 456 positive cases were found: 1,645 of which entered covidcodes and 1,695 calls were received at the infoline. As already mentioned, these numbers can now be followed on a daily basis.

What was new in this survey was the count of the reasons for people to be tested. Indeed, for each positive case, the medical infrastructure is required to fill out a clinical report.[5] This form included a field to indicate the reason. For 7,842 forms, this field was filled. It is a tick-box form with four possible answers: having symptoms (6,380 answers), required by investigation (487 answers), suggested by SwissCovid (41 answers), or other reason to specify. Based on these 41 SwissCovid answers, the authors quickly deduced that those cases would not have been found without SwissCovid and concluded that SwissCovid is useful.

[5]https://www.bag.admin.ch/dam/bag/fr/dokumente/mt/msys/covid-19-meldeformular-ambulant.pdf.download.pdf/OFSP_covid19_formulaire-de-declaration_patients-ambulatoires.pdf — Coronavirus disease COVID-19.

Table 1. The SwissCovid Contribution in the Canton of Zurich in October-November 2020

period	#cases	#codes	#notified	#"discovered"	#dis./#cases	#codes/#dis.	#dis./#notified
1–30.9	1923	324	1374	30	1.56%	10.8	2.18%
1–15.10	3736	422	2457	67	1.79%	6.3	2.73%
16–31.10	13 185	1426	5886	78	0.59%	18.3	1.33%

Table 2. The SwissCovid Contribution in Switzerland in October-November 2020

period	#cases	#codes	#notified	#"discovered"	#dis./#cases	#codes/#dis.	#dis./#notified
1–30.9	11 463	1867	4633	67	0.58%	27.9	1.45%
1–15.10	22 555	2916	7020	121	0.54%	24.1	1.72%
16–31.10	100 057	12 453	18 708	180	0.18%	69.2	0.96%

Actually, there is no evidence of any causality between ticking SwissCovid in this form and being a discovered case. One comment by Dehaye[6] states that people put in quarantine by contact tracers at the time were not required to make a test (which they would have to pay for otherwise) if they had no symptoms but were encouraged to have a free test if they were alerted by SwissCovid. So, it is very likely that those 41 cases were already in quarantine when they were alerted by SwissCovid. This objection, posted on September 13, 2020, was never acknowledged or commented on by the authors. So, it seems that the real number of discovered cases might rather be much lower than 41, even including discovered cases that did not get tested but took voluntary actions. As discussed below, most of the time, people alerted by SwissCovid were contacted by a contact tracer before the alert. In what follows, we write "discovered" under double quotes to remind that the count may be strongly biased (overestimated).

Overall, this analysis is missing an evaluation of the human cost behind these results: how many people were alerted for the number of identified cases. One can always find 41 cases by alerting enough people at random. What would change is the "cost" in terms of number of alerted people, which is not analyzed.

The SwissCovid Contribution in the Canton of Zurich. A paper by Menges et al. [15] studies the activities in the state of Zurich. This survey focuses on the figures in Zurich State and has to make lots of estimates since available figures are not always split in states. They consider three periods: September 2020, first half of October, and second half of October, which was the rise of the second COVID-19 wave in Switzerland, as we can see in Figure 4. Like in the previous study [18], they estimated the number of "discovered cases". As the figures in Table 1 show, the fraction of the cases which are "discovered by SwissCovid" is about 1%. Of course, we do not expect this technology to discover 100% of cases but this 1% figure shows that the impact of SwissCovid is at most marginal. Again, this low performance is claimed to be caused by the low adoption rate. Assuming the dependency is quadratic, doubling the adoption rate (which would go beyond the best hopes of FOPH) would multiply it by 4, which would still be anecdotal.

To estimate the number of SwissCovid alerts, the authors used a survey which concluded that 53.3% of people who were alerted by the app call the infoline. We can deduce the fraction of the alerted users who become "discovered" cases, which we approximate to 2%. This is a ratio which should normally be independent of the adoption rate of SwissCovid, and should consequently not increase.

What is new in the study is the estimate of the number of alerts and the number of quarantines. The number of people put in quarantine by contact tracers in Zurich State was 3,000. The number of people put in quarantine after calling the infoline was estimated to be 166 (756 times a nation-wide ratio of 22% coming from another

[6] https://github.com/digitalepidemiologylab/swisscovid_efficacy/issues/1 — Easy to misinterpret the situation of the 26 reporting SwissCovid as cause for the test.

survey). Hence, 5% of quarantines were made by SwissCovid. As we can see, it contributed more to quarantines (5%) than to "discovering" cases (around 1%). Furthermore, few alerts gave a "discovered" case (2%).

What is a bit alarming is that the two ratios we computed drop as the number of cases increases. This seems to suggest that the relative usefulness of SwissCovid decreases when we would need it more to help overwhelmed contact tracers. Another alarming result is that figures in Zurich seem to overestimate the ones of entire Switzerland, as Table 2 shows. Saying that 1% of the cases in Switzerland are "discovered" by SwissCovid and that less than 2% of alerts are useful is very generous.

The paper was silent about this in both cases. Instead, it focused on the meaningless #discovered/#codes ratio which was generously estimated to 1 : 10.9 (we can see it was 1 : 69.2 in Switzerland by end of October) and concludes:

> "Promoting use of the app, increasing automation of the DCT notification cascade, and connecting rapid antigen testing with DCT—while maintaining the privacy-preserving and voluntary nature of DCT—could further enhance the speed of the notification cascade and increase compliance of exposure-notified app users."

On SwissCovid Sending to Quarantine Faster. Ballouz et al. [7] continued the study of the Zurich cohort to analyze how fast were alerted SwissCovid users going to quarantine compared to others. More precisely, this study looked at contact cases which were alerted by SwissCovid within less than 7 days and which are non-household. It is indeed assumed to be useless to study the effect on any contact tracing tool on household contact cases, and overly delayed notification are assumed to be useless (although they credit the utility of SwissCovid in previous studies [15, 18]). Then, it is observed that those alerted users went to quarantine roughly one day before others. The article concludes:

> "Our analysis showed that non-household contacts notified by the app started quarantine earlier than those not notified by the app. This provides important evidence that DPT apps have an impact on the timely interruption of transmissions chains."

This statement is obviously wrong for multiple reasons.

- The first sentence is overly simplified. It omits "non-household" and "notified within less than 7 days" and it seems to exclude that the contact case could have been aware of their risk status irrespective of the SwissCovid alert.
- The second sentence suggests that SwissCovid-alerted people going faster to quarantine provides evidence that SwissCovid has an impact on the speed to quarantine. Actually, what the study shows is a correlation instead of a causality. It provides no such evidence. Another interpretation is that people who decided to use SwissCovid are simply more concerned about the disease transmission and that they would have gone faster to quarantine in any way.

Unsurprisingly, the conclusion of the paper was immediately twisted by the colleagues of the authors from the Swiss Science Task Force and in the media into the great news that SwissCovid was alerting contact cases faster than other means.[7]

We may wonder about the actual speed of SwissCovid in alerting contact cases. A partial answer can actually be found, drowned in the middle of the same paper:

[7] https://iacr.org/submit/files/slides/2021/rwc/rwc2021/1003/slides.pdf — Privacy by Design from Theory to Practice in the Context of COVID-19 Contact Tracing — Slide 26.
https://www.rts.ch/info/suisse/11908431-lappli-swisscovid-reste-plus-utile-que-jamais-estime-son-concepteur.html — L'appli Swisscovid reste plus utile que jamais, estime son concepteur.
https://www.rts.ch/play/radio/on-en-parle/audio/le-devenir-de-lapplication-swisscovid?id=11872353 — Le devenir de l'application Swisscovid — 3:30.

> "Among the 192 close contacts using the app, 38% ($n = 73$) received an app notification within 7 days of the last relevant exposure. Out of these, 12% ($n = 9$) received the notification before being contacted by MCT."

As we can verify elsewhere in the paper [7, Table 3 p.14], this is not a typo: among the isolated contact cases, SwissCovid was faster than manual contact tracing in only 12% of the cases. We believe it definitely proves that the speed at which quarantine started has nothing to do with SwissCovid and that SwissCovid is alerting contact cases slower than other means.

The Epidemiological Impact of the NHS COVID-19 App. Wymant et al. published in Nature another article in *Nature* [24] claiming a big success of the British sibling of SwissCovid: the NHS COVID-19 App. What the media (and the general public) took from this paper is the punchline: "between 300,000 and 600,000 COVID-19 cases have been prevented in the UK alone thanks to the NHS COVID-19 app".[8] This means that the authors estimated how many more people would have caught COVID-19 during the last quarter of 2020 if the app did not exist based on several assumptions. First of all, the computation assumes that a notified person who also becomes positive does not contaminate anyone thanks to the notification but would have contaminated people without the notification from the app. Then, each of these cases would contaminate others using the same ratio as existing cases. Each of this new contaminated cases would also contaminate using the same ratio, and so forth until the end of the studied period. Based on the reproduction rate per location, on several modeling assumptions, and thanks to the help of the exponential growth, the authors reach a number which is either 300,000 or 600,000 (which represents 15–30% of the actual number of cases, based on the official 1.9 million cases during the period[9]) depending on the computation method. Hence, this number is highly speculative.

This article also estimates numbers which are more reliable: there were 16.5 million app users in England and Wales (representing 28% of the population); 560,000 app users have been positively tested; 72% of app user cases consented to reporting (hence, 400,000 covidcodes entering 21% of all cases in the system); each entered report generated 4.4 notifications on average; about 0.27 of these notified users become positive on average; 1.7 million app users received a notification. This last number deviates from the official NHS figures (1.2 million[10]) but the authors maintain that their numbers are correct. Based on these figures, this article essentially assumes that $560{,}000 \times 72\% \times 0.27 = 109{,}000$ cases were discovered by the app during the period of studies, which would be an honorable ratio of 5.5% of all cases. Those allegedly discovered cases may include household contacts (there are 1.2 household contacts per case on average).

We can compare these figures with the Swiss ones. The adoption rate of the app is a bit higher in the U.K. (28% vs. 22%). However, the ratio of positive cases per app user is the same (29% in both cases, with the 16–31.10 period in Table 2). The consent-to-report rate is much higher with NHS (72% vs. 43%). This could be due either to deficiencies in the delivery system of covidcodes in Switzerland or to the fact that it takes less effort to report on the NHS app: instead of entering a 12-digit code, users make test appointments on their app, receive their result on their app, and only have to click to report. The British app generates more notifications per entered covidcode (4.4 vs. 1.5). The higher adoption rate does not fully explain this phenomenon. This could be due to a difference in the calibration of the sensitivity to spot contacts and also to a difference on the population mobility in the two countries during the two periods. Finally, the probability that a notified user becomes positive is dramatically higher in the UK (6% vs 1%), which is surprising because we could have thought that the NHS app was notifying too much. This could potentially reinforce the hypothesis that the alleged app success is not a causality of usage. Instead, it could be linked to the possible tendency of those who volunteered to use the app to be more cautious and aware of their infection risks.

[8] https://actu.epfl.ch/news/contact-tracing-apps-prove-that-they-save-lives/.
[9] https://coronavirus.data.gov.uk/details/cases.
[10] https://stats.app.covid19.nhs.uk/data/covid19_app_country_specific_dataset.csv.

4 SECURITY AND PRIVACY

Security and privacy were promised to reach a very high standard. We discuss security and privacy issues related to SwissCovid [22]. We stress that most of those threats would apply to other GAEN-based automated contact tracing systems.

4.1 False Exposure Notifications

One important risk is that people will receive unjustified exposure notifications which will push them to quarantine. Such false positives can occur randomly, by bad luck, but can also be maliciously created by an adversary in order to get an advantage in putting some person in quarantine or at least under stress. For instance, a lazy student could try to make his teacher receive an alert to skip a test. An adversary could try to get an advantage by stressing competitors. Activists could massively spread false alerts in a wide area. People could also self-inject false alerts to have a good excuse to stay home. There are several ways to inject false alerts. To describe them, we follow the characterization of Iovino et al. [10].

Attack with Real Encounter. The adversary encounters the victim and uses the app normally. Then, it reports the encounter to raise an alert to the victim. Reporting can be done by corrupting the health authority system (to get a covidcode) or by corrupting a user who was just diagnosed and received a covidcode. Avitabile et al. [6] describe how a smart contract could make just-diagnosed people and false alert injectors securely agree without even knowing each other.

Attack with a Simulated App. The adversary is using an ad-hoc Bluetooth device which mimics SwissCovid. This way, the device can be set to cheat with the sending power: it can send at full power but encrypt a lower value in the metadata so that the receiver will think it is in close proximity to the malicious sender. The other advantage is that many adversaries can use devices which are synchronized to send the same RPI. They will look like the same user meeting many people and create false alerts massively.

Reporting can be done as in the previous case. If done by many synchronized adversaries, one of them could sacrifice to get COVID-19 and genuinely report. A terrorist/activist attack would try to make a wide population receive a false notification. Some competitor would try to make many employees of the same company receive notifications.

Replay Attack. The adversary can replay to the victim the Bluetooth messages which were sent by people who are likely to be diagnosed soon, for instance by capturing messages in emergency places such as hospitals. In most of contact tracing applications, there are countermeasures against replay attacks. The countermeasure which is used in GAEN still allows replays for about 2 hours, following the GAEN specifications. Dehaye and Reardon [9] describe how a benign app could make the replay attack from a smartphone without the holder being aware.

Sometimes, a "KISS attack" (Keep It Simple, Stupid) surprisingly works: replay a TEK which was just posted on the server, as long as it is still valid [10]. The SwissCovid infrastructure seems to filter well those keys so that this attack does not work. However, it works well in some neighbor countries like Italy and Austria. Furthermore, interoperability between countries impact other countries like Germany: we can replay a reported and still valid TEK from Austria in Germany and it may trigger false alerts.

Belated Replay Attack. An outdated message can also be replayed by corrupting the time of the victim's phone, as detailed by Iovino et al. [10]. Figure 1 shows a screenshot of a SwissCovid alert which was obtained this way. There are many options to set up the date of a phone without any physical access or authorization. The easiest way is to corrupt NTP messages from the network to the phone [16]. This literally implements a time machine which can send a phone any time in the past or in the future. This way, the adversary can take a recently uploaded TEK from the server, send the victim to the time when it was used, replay the RPI, wait until the victim comes

back to the present time, download the TEK, match it, and raise an alert. By making several time jumps, the duration of the attack can be reduced to 1 second. (See Iovino et al. [10] for more details.)

4.2 Privacy of SwissCovid Users

GAEN Captive Users. Forcing users to adopt a GAEN-based architecture implies forcing them to adopt Google or Apple systems. On Android, GAEN is part of the *Google Play Services*, a built-in app which is used by other Google apps. It could be removed from Android to obtain a "deGoogled phone" which is still running Android. Using GAEN implies using Google Play Services which makes people more captive of this infrastructure. It helps dominant companies maintain their monopoly [12]. However, this adoption has a privacy cost since Google Play Services is sharing lots of private information with Google [14].

We recall that the pre-standard version of SwissCovid was not using GAEN. It was running well on deGoogled phones, as well as on banned Chinese phones. The standard SwissCovid was thus either forcing users to acquire a recent phone with either Google or Apple systems or excluding them from the system, thus digging the digital divide.

Inherent Bluetooth Issues. The Bluetooth technology is inherently creating privacy threats.

First of all, vulnerabilities allowing hackers to enter into a phone by Bluetooth are discovered frequently. (See, e.g., Antonioli et al. [5].) This is why people often turn off Bluetooth.

Second, observing Bluetooth messages can threaten privacy. It is trivial to detect SwissCovid users and to notice them and hence uncover who decided to use SwissCovid or not. We can further build networks of sensors which can collect Bluetooth signals with position and time. Every smartphone could also take part in this network. This way, we can link consecutive Bluetooth messages and create a database showing the movement of people, who can be traced this manner. One critical element of this attack is the ability to link Bluetooth messages. When they are consecutive, linking them is trivial. (It is sometimes made clear by desynchronized rotations of the message and the MAC address.[11])

Finally, the My Number attack [10] allows the attacker to map each SwissCovid device to a unique RPI number by making them travel to an arbitrary reference date in the far future such as April 1, 2048. The TEK of this day will stay stored forever and be reused to broadcast the same RPI in every time jump to this date. GAEN is thus a good opportunity for surveillance.

Leakage to the Network. Although SwissCovid worked hard in trying to obfuscate communication (true reports are drowned among fake reports), there are always leakages. For instance, SwissCovid reveals its activation to the network. The CDN of Amazon can collect IP addresses and other connection information. Probably, the most severe leakage was about the online survey, which is filled by alerted users through a server that is hosted by Microsoft in London. It reveals when the alleged infection occurred, all data filled by the user, and of course the IP address and other metadata. Officially, SwissCovid does not leak, for the simple reason that the user leaves the SwissCovid app to fill the survey through a browser. The problem has reference INR-14788 and was filled on February 24, 2021. The reaction was:

> "A new information icon will be integrated next to the SwissCovid Guide button. This information icon informs that, he/she is leaving the app and that the contact date will be submitted for better assistance."[12]

Meanwhile, medical data will continue to leave Switzerland, the European Union, and escape GDPR protection.

[11] https://vimeo.com/453948863 — Little Thumb Attack on SwissCovid.
[12] https://www.ncsc.admin.ch/dam/ncsc/de/dokumente/dokumentation/covid-security--test/SwissCovid_Public_Security_Test_Current_Findings.pdf — SwissCovid Proximity Tracing System - Public Security Test.

4.3 Privacy of Reporting People

The decentralized architecture of GAEN creates privacy threats for reporting people.

People who report essentially publish the keys which they used to announce themselves to close-by people. If a malicious person captures their Bluetooth notification (by being close or not), remembers from whom it comes, and sees a matching key on the server, then this person can deduce that the sender was diagnosed. A paparazzi can capture from far away the signal of a celebrity, then wait to see if the celebrity reports. This requires neither proximity nor a long-lasting contact [20].

People can also use "enriched apps" which can store more information than the Bluetooth signal they receive. For instance, they can store the location and precise time, as well as notes from the malicious user. Collected data can even be shared or sold. If the Bluetooth signal of a person who is later diagnosed is captured this way, this can reveal where this person has been and when the person was there. A proof of concept is available as a "Corona Detective".[13]

The "enriched app" can also run on a phone without its user being aware of it. For instance, there are applications which need to scan Bluetooth and share collected data together with localization information with other devices which use the same app. Those apps could easily be changed into a malicious distributed surveillance system.

Any organization which identifies people (hotels at registration desks, companies at entrance doors, shops at checkout) may, at the same time as they identify the person, capture a Bluetooth signal from this person and remember it. This way, a hotel/company/shop can sell health information about a visitor to advertisers.

A video-surveillance system may be enriched with Bluetooth captures. Hence, we may be able to recover visual information of people who have been diagnosed. Actually, a Bluetooth-only surveillance system already collects lots of private information. Baumgärtner et al. [8] describe such a system.

5 GOALS

5.1 DP-3T Goals

The DP-3T project describes the architecture and the "system requirements" in a white paper [19]. They are listed in four categories. We review them in light of what we discussed.

(1) Enabling requirements.
- The system should capture all exposure events (*completeness*).
  👎 We have seen that due to the precision of measuring distance and time, there is a high chance that 15-minute proximities are not detected.
- Reported events must reflect actual proximities (*precision*).
  👎 We have seen that, due to the precision of measuring distance, there is a chance that a non-proximity encounter is reported.
- Events cannot be faked (*authenticity*).
  👎 We have seen that an adversary can easily inject a false alert.
- Adversaries cannot access to the contact history of a user (*confidentiality*).
  👎 As long as this history remains on the phone, access is protected. However, GAEN has access to the contact history and it sometimes leaks out [17].
- Users must be informed about prolonged exposure (*notification*).
  👍 If the prolonged exposure is detected, users are informed.

(2) Privacy.
- Data collection and use should be limited to the purpose of the system (*data use*).
  👎 We have seen that using Bluetooth enables surveillance.

[13] https://www.coronadetective.eu/ — Find out exactly who gave you COVID19.

- Information about individuals should be controlled to avoid unintended information leakage (*controlled inference*).
  👎 Inference can be made by collecting available data. Answers to surveys are collected outside of Switzerland and of the GDPR-protected zone.
- The system should use anonymization of pseudonyms (*protect identities*).
  👍 The RPI and TEK are pseudonyms.
- The system should erase data when no longer useful (*erasure*).
  👎 Data in the far future is never erased; this causes the My Number attack [10].

(3) Scalability.
- The system should scale to billions of users (*scalability*).
  👍 The system is currently used by many millions.
- The system should work across borders (*interoperability*).
  👎 SwissCovid does not work across border due to regulation with the European Union. (The Swiss law on data protection is considered as being insufficient by the EU.) Since March 2021, it became interoperable, but with Germany only.

(4) Feasibility.
- The system should solely rely on existing and well-established technology.
  👍 It works.
- The system should rely on widely available hardware.
  👍 It only uses the existing smartphones of the owners.

Before the latest version of the white paper (May 25, 2020), it was required to provide data to epidemiologists with the consent of users. This requirement was dropped.

5.2 DP-3T-Initiated Academic Joint statement

On April 19, 2020, members of DP-3T wrote a joint statement which was signed by more than 300 people from the research community[14] This joint statement calls on policymakers to consider several principles. The statement was written in the context of DP-3T lobbying and fighting against the centralized solutions of PEPP-PT. It was written in a biased manner, suggesting that centralized systems would necessarily allow governments to perform mass surveillance while their systems are privacy-preserving by authoritative claims [21]. Besides this obviously biased presentation, the letter concludes with four principles:

> "The following principles should be at least adopted going forward:
> - Contact-tracing apps must only be used to support public health measures for the containment of COVID-19. The system must not be capable of collecting, processing, or transmitting any more data than what is necessary to achieve this purpose.
> - Any considered solution must be fully transparent. The protocols and their implementations, including any sub-components provided by companies, must be available for public analysis. The processed data and if, how, where, and for how long they are stored must be documented unambiguously. Such data collected should be minimal for the given purpose.
> - When multiple possible options to implement a certain component or functionality of the app exist, then the most privacy-preserving option must be chosen. Deviations from this principle are only permissible if this is necessary to achieve the purpose of the app more effectively, and must be clearly justified with sunset provisions.

[14] https://www.esat.kuleuven.be/cosic/sites/contact-tracing-joint-statement/ — Contact Tracing Joint Statement.

- The use of contact-tracing apps and the systems that support them must be voluntary, used with the explicit consent of the user, and the systems must be designed to be able to be switched off, and all data deleted, when the current crisis is over."

We can see that GAEN-based systems fail to respect the second principle as GAEN is not open.

5.3 Legal Basis

The Swiss Parliament updated the Law on Epidemics (LEp) on June 20, 2020 in order to set the legal frame for SwissCovid. This resulted in two new articles (60a and 62a) and the modification of two other articles (80 and 83). Article 60a introduces the system, defines and limits its purpose, and states some requirements. One main requirement is that the source code of all components of the system should be publicly available and verifiable (Art.60a al.5e).[15] It is quite surprising as GAEN is not verifiable and has no available source code.[16] However, the Government is in charge of the details (Art.60a al.7).

Hence, the Government released the *Ordinance on the Proximity Tracing System for Coronavirus* (OSTP) on June 24.[17] OSTP Article 2 defines the *components* of SwissCovid through an exhaustive list in which GAEN is not present. GAEN appears in Article 5 where it is said that the SwissCovid app completes some tasks "*with the help of an interface of the operating system*" (Art.5 al.2). Hence, GAEN appears as a helper and its existence is acknowledged as if it was a necessary component of the phone. The article lists the tasks which the app must complete: the generation of the daily TEK, the exchange over Bluetooth, the download of the keys, the matching, and the notification. Except for the download and notification tasks, however, it is fairer to say that these tasks of the app are *outsourced* to GAEN and that GAEN is not merely *helping* the app. Finally, Article 5 explicitly says that GAEN does not need to have a public source code (Art.5 al.3).

It looks clear to us that the spirit of the law was twisted to accommodate GAEN in SwissCovid although the letter of the law looks correct. Based on the fact that both the LEp and OSTP have been written at the same time, we can claim that the Parliament has been deceived.

OSTP also defines which proximities should be identified: proximities at a distance up to 1.5 meters for a total period of at least 15 minutes. Clearly, SwissCovid is loosely respecting that. Since SwissCovid falls into the category of medical devices, it is also subject to the regulation by Swissmedic.[18] The goal of such regulation is to enforce transparency and to make sure people can consent with awareness of the pros and cons. We wonder to what extent the lack of information concerning false positives and false negatives to spot contacts at a 1.5-meter distance lasting 15 minutes is critical for compliance.

Another important parameter is defined by OSTP: the starting date for reporting the TEK should be set to two days before symptoms appear, as it is believed to correspond to the contagious period. This requirement prevents SwissCovid from being useful for *backward contact tracing*. As a matter of fact, it is known that an important fraction of cases do not contaminate anyone [4] so trying to break the transmission chain this way is not so effective. It was already known that since May 5, 2020, contact tracing should focus on backward contact tracing to find the super-spreaders [11]. Instead, QR-code-based contact tracing was added in "SwissCovid 2", in June 2021 only. Besides, its voluntary usage does not fulfill the requirement for mandatory contact tracing in some places. Consequently, restaurants use another contact tracing system.

[15] https://www.admin.ch/opc/fr/classified-compilation/20071012/index.html#a60a — Loi Fédérale sur la lutte contre les maladies transmissibles à l'homme.

[16] Although an *excerpt* of the source code was later published.

[17] https://www.admin.ch/opc/en/classified-compilation/20201730/index.html — Ordinance on the Proximity Tracing System for the Sars-CoV-2 coronavirus.

[18] Added note: On July 1, 2021, Article 23bis was added in the COVID-19 Ordinance 3. it states that SwissCovid is no longer subject to the regulation about medical devices. This clearly eases the compliance effort but weakens transparency. https://www.fedlex.admin.ch/eli/cc/2020/438/fr#art_23.

LEp specifies that the Government should shut down SwissCovid if it appears that it is *not efficient enough to fight against the pandemic* (Art.60a al.8). However, no way to assess this efficiency is defined and no calendar to assess it is proposed.

6 CONCLUSION

It is remarkable to see how DP-3T cried wolf and defied governments who were suspected of putting in place measures for mass surveillance. While DP-3T intended to strengthen privacy, it instead offered a boulevard for giant data collectors to impose GAEN and secure their monopolies. It succeeded in making governments deploy a system which provides little useful information about the pandemic. It also ended up making the Swiss Government having to twist the law. To believe SwissCovid would work was already an act of faith as it was clear that Bluetooth would be imprecise and that adoption would not be total. Persisting in the mistake eventually leads to forging proofs, blaming non-users for the low performances of SwissCovid, and teaching inadequate privacy lessons to people. We hope this case will help prevent similar mistakes in the future (such as suggesting to people who are concerned about the privacy of SwissCovid to stop using social media).

REFERENCES

[1] Exposure Notification. Bluetooth Specification. v1.2 April 2020. https://blog.google/documents/70/Exposure_Notification_-_Bluetooth_Specification_v1.2.2.pdf.

[2] Exposure Notification. Cryptography Specification. v1.2 April 2020. Apple & Google. https://blog.google/documents/69/Exposure_Notification_-_Cryptography_Specification_v1.2.1.pdf.

[3] SwissCovid Exposure Score Calculation. Version of 11 September 2020. Hosted by: Federal Office of Information Technology, Systems and Telecommunication FOITT. https://github.com/admin-ch/PT-System-Documents/raw/master/SwissCovid-ExposureScore.pdf.

[4] Dillon C. Adam, Peng Wu, Jessica Y. Wong, Eric H. Y. Lau, Tim K. Tsang, Simon Cauchemez, Gabriel M. Leung, and Benjamin J. Cowling. 2020. Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nat. Med.* 26, (2020), 1714–1719. https://doi.org/10.1038/s41591-020-1092-0

[5] Daniele Antonioli, Nils Ole Tippenhauer, Kasper Rasmussen, and Mathias Payer. 2020. BLURtooth: Exploiting cross-transport key derivation in Bluetooth classic and Bluetooth low energy. 24 September 2020. Preprint arXiv:2009.11776 [cs.CR], 2020. https://arxiv.org/abs/2009.11776.

[6] Gennaro Avitabile, Daniele Friolo, and Ivan Visconti. 2020. TEnK-U: Terrorist attacks for fake exposure notifications in contact tracing systems. Cryptology ePrint Archive: Report 2020/1150. IACR. http://eprint.iacr.org/2020/1150.

[7] Tala Ballouz, Dominik Menges, Hélène E. Aschmann, Anja Domenghino, Jan S. Fehr, Milo A. Puhan, and Viktor von Wyl. 2020. Digital proximity tracing app notifications lead to faster quarantine in non-household contacts: Results from the Zurich SARS-CoV-2 Cohort Study. medRxiv December 23, 2020. medRxiv 2020.12.21.20248619. Cold Spring Harbor Laboratory Press. https://doi.org/10.1101/2020.12.21.20248619

[8] Lars Baumgärtner, Alexandra Dmitrienko, Bernd Freisleben, Alexander Gruler, Jonas Höchst, Joshua Kühlberg, Mira Mezini, Markus Mi-ettinen, 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. Preprint arXiv:2006.05914 [cs.CR], 2020. https://arxiv.org/abs/2006.05914.

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

[10] Vincenzo Iovino, Serge Vaudenay, and Martin Vuagnoux. 2021. On the effectiveness of time travel to inject COVID-19 alerts. In *Topics in Cryptology (CT-RSA'21): The Cryptographers' Track at the RSA Conference 2021*, Virtual Event, Lecture Notes in Computer Science 12704, pp. 422–443, Springer-Verlag, 2021. (IACR eprint 2020/1393 https://eprint.iacr.org/2020/1393) https://doi.org/10.1007/978-3-030-75539-3_18

[11] Sadamori Kojaku, Laurent Hébert-Dufresne, Enys Mones, Sune Lehmann, and Yong-Yeol Ahn. 2021. The effectiveness of backward contact tracing in networks. *Nature Physics* (2021). https://doi.org/10.1038/s41567-021-01187-2

[12] Marjolein Lanzing. 2020. Contact tracing apps: An ethical roadmap. 29th September 2020. *Ethics Inf. Technol.* (2020). https://doi.org/10.1007/s10676-020-09548-w

[13] Douglas J. Leith and Stephen Farrell. 2020. Measurement-based evaluation of Google/Apple exposure notification API for proximity detection in a light-rail tram. *PLoS ONE* 15(9) 2020. https://doi.org/10.1371/journal.pone.0239943

[14] Douglas J. Leith and Stephen Farrell. 2020. Contact tracing app privacy: What data is shared by Europe's GAEN contact tracing apps. 18 July 2020. https://www.scss.tcd.ie/Doug.Leith/pubs/contact_tracing_app_traffic.pdf.

[15] Dominik Menges, Hélène E. Aschmann, André Moser, Christian L. Althaus, and Viktor von Wyl. 2021. A data-driven simulation of the exposure notification cascade for digital contact tracing of SARS-CoV-2 in Zurich, Switzerland. April 30, 2021. *JAMA Network Open*. 2021;4(4):e218184. https://doi.org/10.1001/jamanetworkopen.2021.8184

[16] Shinjo Park, Altaf Shaik, Ravishankar Borgaonkar, and JeanPierre Seifert. 2016. White rabbit in mobile: Effect of unsecured clock source in smartphones. In *Proceedings of the 6th Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM@CCS 2016)* (Vienna, Austria, October 24, 2016). ACM, New York, 2016, pp. 13–21. https://doi.org/10.1145/2994459.2994465

[17] Joel Reardon. 2021. Why Google should stop logging contact-tracing data. https://blog.appcensus.io/2021/04/27/why-google-should-stop-logging-contact-tracing-data/.

[18] Marcel Salathé, Christian L. Althaus, Nanina Anderegg, Daniele Antonioli, Tala Ballouz, Edouard Bugnion, Srdjan Čapkun, Dennis Jackson, Sang-Il Kim, James R. Larus, Nicola Low, Wouter Lueks, Dominik Menges, Cédric Moullet, Mathias Payer, Julien Riou, Theresa Stadler, Carmela Troncoso, Effy Vayena, and Viktor von Wyl. 2020. Early evidence of effectiveness of digital contact tracing for SARS-CoV-2 in Switzerland. *Swiss Medical Weekly.* 2020;150:w20457. https://doi.org/10.4414/smw.2020.20457

[19] 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 Capkun, David Basin, Jan Beutel, Dennis Jackson, Marc Roeschlin, Patrick Leu, Bart Preneel, Nigel Smart, Aysajan Abidin, Seda Guerses, 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. Version: 25th May 2020. Preprint arXiv:2005.12273 [cs.CR], 2020. https://arxiv.org/abs/2005.12273.

[20] Serge Vaudenay. 2020. Analysis of DP3T. *Cryptology ePrint Archive:* Report 2020/399. *IACR*. http://eprint.iacr.org/2020/399.

[21] Serge Vaudenay. 2020. Centralized or decentralized? The contact tracing dilemma. *Cryptology ePrint Archive*: Report 2020/531. IACR. http://eprint.iacr.org/2020/531.

[22] Serge Vaudenay and Martin Vuagnoux. 2020. *Analysis of SwissCovid.* 2020, June 5. https://lasec.epfl.ch/people/vaudenay/swisscovid/swisscovid-ana.pdf.

[23] Lucie White and Philippe van Basshuysen. 2021. Privacy versus public health? A reassessment of centralised and decentralised digital contact tracing. *Sci. Engin. Ethics* 27, article 23, 2021. https://doi.org/10.1007/s11948-021-00301-0

[24] Chris Wymant, Luca Ferretti, Daphne Tsallis, Marcos Charalambides, Lucie Abeler-Dörner, David Bonsall, Robert Hinch, Michelle Kendall, Luke Milsom, Matthew Ayres, Chris Holmes, Mark Briers, and Christophe Fraser. The epidemiological impact of the NHS COVID-19 app. *Nature*, 2021. https://doi.org/10.1038/s41586-021-03606-z

[25] Viktor von Wyl, Marc Höglinger, Chloé Sieber, Marco Kaufmann, André Moser, Miquel Serra-Burriel, Tala Ballouz, Dominik Menges, Anja Frei, and Milo Alan Puhan. 2021. Drivers of acceptance of COVID-19 proximity tracing apps in Switzerland: Panel survey analysis. *MIR Public Health Surveill.* 2021;7(1):e25701. Published 2021 Jan. 6. https://doi.org/10.2196/25701

Received November 2020; revised June 2021; accepted August 2021