Exposure Notification System May Allow for Large-Scale Voter Suppression

Rosario Gennaro*, Adam Krellenstein†, James Krellenstein‡

September 2, 2020

1 Introduction

Exposure Notification is a system designed by Google and Apple ("GAEN", in the following) for notifying individuals when they have been exposed to SARS-CoV-2 by coming in contact with someone who has tested positive for the virus [13, 3]. It is closely related to the "hybrid decentralized proximity tracing" protocol from the Decentralized Privacy-Preserving Proximity Tracing (DP-3T) project [25]. GAEN is intended to be a minimalistic and privacy-preserving system for automated anonymized exposure notifications, to complement manual contact-tracing efforts in an efficient and highly scalable fashion. In this document, we do not distinguish carefully between GAEN and DP-3T because of their similarity.

The Exposure Notification system is characterized by its strong emphasis on preserving the privacy and anonymity of users, and and the strict limits it places on the potential for abuse comprising unauthorized data collection or mass surveillance. Within GAEN, no user-identifying data is ever uploaded to the central server; users establish their proximity exclusively peer-to-peer and anonymously, with the sole purpose of knowing whether they have been in contact with an individual who may later be deemed to have been infected [25].

The design choices of the protocols in question, which makes them robust against data collection attacks, unfortunately also make them particularly susceptible to data injection by malicious parties. In particular, these protocols allow for a determined attacker to generate false exposure notifications on a mass scale in an undetectable and unpreventable manner. We believe that the potential societal consequences of such an attack have been heretofore underestimated, though we are not the first to have expressed similar concerns [21, 22, 15].

*The City College of New York. rosario@ccny.cuny.edu

†adam@krellenstein.com

‡The PrEP4All Collaboration. james@prep4all.org

2 Analysis

It is well understood that the GAEN protocol allows for the possibility that a malicious attacker can generate **fake exposure events**, specifically *via* e.g. relay, replay and "inverse Sybil attacks"[9]. The DP-3T whitepaper considers this to be a necessary limitation of any contact-tracing system that relies exclusively on the strength of a radio signal for defining proximity [25]. Indeed, a number of features of the GAEN protocol, specifically relating to its strict privacy requirements (and consequent lack of recorded location data), exacerbate the potential for the generation of fake exposure events. These attacks have been discussed by the cryptographic academic community at length [27, 26, 9, 17] and we summarize them here in the context of the GAEN system.

2.1 Proximity

An attacker may use powerful radio transmitters to make targets believe they are proximal to an infected person when they are in fact very far away [25]. Moreover, because the receiving device relies on unauthenticated assertions of the transmitting device's broadcast strength, a malicious broadcaster that under-reports this value may seem even closer than would otherwise be the case [28].

2.2 Device Ownership

There is no authentication mechanism for informing the central server of a new infection beyond the use of a one-time authentication code. With such a code, having reverse-engineered the API, an attacker may be able to upload the past fourteen days of Temporary Exposure Keys (TEKs) and interval numbers ($seed_w$ and w, in DP-3T's terminology) from an *arbitrary device* to the server, which will then proceed to disseminate those keys to every single active device within the geographical area that it serves. That is, there is no way of knowing whether the infected individual themselves was actually in possession of the particular device over the previous fourteen-day period. There are some proposed mechanisms for reducing the attack surface here, namely by committing to the TEKs when an individual takes the COVID-19 test (pp.18–21, [23]) but none of them protects against an attack which is planned before the test is taken.

2.3 Positive Test Result

The developers of GAEN strongly advise public health departments implementing their system to put in place a strict verification protocol that ensures only positive individuals are able to upload their TEKs and that they do so in a timely fashion. However, even if such a protocol has all of the desired security properties, it is not expected to be difficult for an attacker to acquire an otherwise valid authorization code illicitly (p.36, [23]). Indeed, there are many practical methods by which an attacker could do so, including purchasing one on the black market or having a

single infected attacker get tested multiple times (as of this writing, there are over 500,000 tests performed daily in the United States, likely with highly variable security protocols) [16]. In any reasonable implementation, attackers have plenty of time to acquire valid authorization codes from individuals, and a highly motivated attacker could even be willing to infect themselves intentionally. Given that SARS-CoV-2 tests in the U.S. often take a week or more to return results, the attack could likely be effectively launched any time during an extended period leading up to Election Day and the targeted individuals may isolate themselves and eschew voting in person.

2.4 Geographic Uniqueness

The devices which broadcast their IDs over Bluetooth do not need to be limited to a specific location. TEKs from a single device may be cloned onto an arbitrary number of other devices, which may then in turn be disseminated widely across the target geographical region. Even if an attacker couldn't share the TEKs—for instance because of OS-level protections, such as DeviceCheck and SafetyNet—the attacker could simply repeat the Bluetooth signals *via* Internet-connected Bluetooth transmitters, where roughly two hours of latency is permitted (p.9, [14]). Coupled with the potential for radio boosting, the lack of geographic uniqueness of a broadcasting device allows for a single set of TEKs to be disseminated across a very wide and non-contiguous area. Indeed, it is a specific non-goal of DP-3T that it be able to detect hotspots of transmission [25]. Of course, were location history shared with the central server, then it would be trivial for the server host to sanity-check the uploaded data and ensure that the devices of infected individuals always maintained a plausible travel pattern over the period of time in question.

2.5 Impossibility of Detection or Mitigation

The above limitations of the Exposure Notification protocol allow for the practically unlimited scaling of fake exposure events, and to targeted areas, using mundane techniques readily available to well-funded attackers. The protocol provides no guarantees regarding the actual proximity of devices, the one-to-one mapping between infected individuals and devices, or that a single "device" can't effectively be in many places at the same time.

This is true of many possible protocols for contact tracing based on peer-to-peer proximity-detection. What is particularly problematic about the Exposure Notification system and DP-3T protocol is: not only are there no guarantees present for any of the above conditions, *there is no way for anyone to detect when one of those conditions has been broken* (except well after the fact, if there's a detectable rise in the number of tests performed within a region and a corresponding drop in positivity rate).[1] All interactions between users, and

[1] As of August 26, the CDC is no longer recommending that exposed individuals get tested if they do not present symptoms [8]. If these guidelines are followed, it would both eliminate the only signal one would have that an attack on GAEN had been perpetrated, as well as

all exposure notifications potentially generated by those interactions, exist only locally, on each user's device. There is no possible mechanism by which the host of the central server—or any other entity—could detect *any* patterns of abuse (impossible travel patterns, number of contacts per device, etc.); there is insufficient global coordination, as the server is not trusted with any private data at all [25]. Within the protocol, health authorities do have the power to revoke previously distributed TEKs, if one or more keys are determined to be associated with an attack. However, there's no possible mechanism allowing these health authorities to identify these keys, even after they have realized that something is wrong.

As a consequence, there is no way for anyone to prevent Exposure Notification from being used to create fake exposure notifications for large numbers of individuals in multiple, targeted locations retroactively. Because of the long infectious period associated with COVID-19, attackers have on the order of fourteen days to fraudulently "expose" as many individuals in the targeted locations as possible, and only have them be notified at the end of that two-week period that they have been exposed to a deadly pathogen.

3 The Swing State Attack

In this section we describe one particular type of attack for **voter suppression**. Consider, for example, an attacker who chooses to target swing voting districts in the United States, where there are often strongly demarcated geographic boundaries between "blue" and "red" neighborhoods. In the 2016 US presidential election, for instance, the margin of victory was very small in a number of states [7] and the ability to suppress even a small fraction of voters in one camp could be successful in determining the outcome of the vote in any of those states.

The attacker could generate a large number of fake exposures in particular neighborhoods during the two-week period leading up to election day, and then during the days before the election upload to that area's central server the TEKs used to generate those fake exposures. One would expect that a non-trivial fraction of individuals who are alerted that they have been exposed to SARS-CoV-2 might decide to isolate themselves and not go to their local polling place and vote as they otherwise would have. By the time any health authorities are able to detect any anomalous behavior (in particular, a large increase in test demand and a drop in positivity rates, if the CDC guidelines are not followed [8]), it would likely be far too late to restore the normal electoral process.

4 The Post Office Attack

Another possible way to exploit GAEN/DP-3T and compromise the security of the election process would be to attack the voting-by-mail process which is expected to play a much larger role in the 2020 US Elections [5]. Alarms are

dramatically reducing the usefulness of the GAEN system, at the very least.

already being raised about the possibility of a terrorist attack on the Postal Service right before the election [24].

An attacker could generate a large number of fake exposures in postal facilities around the country, possibly focusing on facilities that process mail coming from areas with a known political leaning. By creating the impression of a COVID-19 outbreak running through the US Postal Service, an adversary could potentially shut down or greatly impair the functioning of the service, compromising the ability to process and deliver mail-in ballots.

5 Methods

There are two primary ways that we can imagine the above attack being perpetrated:

5.1 Malicious Mobile App

An attacker is able to implement a Swing State or Post Office Attack by inserting malicious functionality into any widely used mobile app that has permission to use the device's Bluetooth and location interfaces. This app, at the very least while it is running in the foreground, would be able to broadcast messages over Bluetooth Low Energy (BLE) that conform to the extremely simple Exposure Notification peer-to-peer proximity detection protocol, which implements no handshake or authentication mechanism. Such an app, not being an approved proximity tracing application, would be unable able to use the OS-level GAEN API—however in practice this simply means that the app would use more power than a whitelisted one.

The conceived malicious app would need to be modified to implement one of the following changes:

1. Include a hard-coded globally shared set of pre-generated TEKs, and use them to generate the Rolling Proximity Identifiers ($EphID_{w,n}$ in DP-3T), which TEKs would later have to be uploaded to the server upon the receipt of a positive test result *not* using an authorized GAEN application.

2. Fetch every twenty-four hours a shared TEK that has been extracted from a device running an authorized GAEN application, and use the collected keys to generate the RPIs.

3. Frequently fetch RPIs themselves as broadcast over BLE by a device running the authorized GAEN application and operating normally. These RPIs are valid for roughly two hours (p.9, [14]). (This method may be necessary if the mobile app and server implement DeviceCheck or SafetyNet vendor-specific device attestation and verification [11, 19]).

The malicious app would broadcast malicious RPIs to nearby devices precisely when the host device is in the desired neighborhoods (and in the message

metadata possibly intentionally understating the power of the broadcast to augment the exposure risk perceived by target devices). At the end of the broadcast window, the attacker acquires the authorization code for a single infected individual (per jurisdiction), and uploads the shared TEKs to the health department's central server.

As of iOS 13.5, only apps that have been whitelisted may broadcast messages that conform to the Exposure Notification Bluetooth protocol, because of security measures at the level of the OS that reserve the relevant service UUID for this purpose [2]. However, there are no known countermeasures of similar design either in older versions of iOS or in Android phones. (As of June 2020, 19% of all iOS devices are not using iOS 13, the latest released version [32]). Before iOS 13, apps did not need any explicit permission at all to use Bluetooth on iPhones and iPads [29]. Nevertheless, this leaves *all* mobile devices running a GAEN app susceptible to message relays performed by the subset of mobile devices that are not running iOS 13.5 or higher.

Popular mobile apps that do request permission to use Bluetooth on iOS 13, and that could therefore be turned malicious in this way include ESPN, Dunkin', Macy's, Sling TV, and FitBit. Were a capable attacker able to compromise even one app capable of broadcasting Bluetooth signals on a significant number of mobile devices, potentially very many mobile devices could readily become a vector for triggering false exposure alerts in a way that would allow for precise geographic targeting. That the attack we describe could be so devastating is a testament to the great risk associated with relying so heavily on unauthenticated peer-to-peer communication about important health information while not allowing for any out-of-band detection or mitigation.

5.2 Physical Attack

In addition to the above software-based attack, it is worth stating that a well-funded attacker would of course also be able to implement a Swing State or Post Office Attack using more traditional methods involving even commercial off-the-shelf Bluetooth beacons and relays. It is outside the scope of this document to estimate the scale at which this attack could occur. However, it is worth noting that commercial Bluetooth radios have ranges on the order of 1–2 km, making heavy use of forward error correction that nevertheless provides sufficient bandwidth for the Exposure Notification broadcast protocol to function (at the maximum range, say, providing transfer rates of around 125 Kbit/s with LE Coded S=8) [4].

As discussed above, there is no practical limit to the number of these Bluetooth beacons that could be deployed across a geographic area served by the same central server, all transmitting the exact same RPIs. The primary countermeasures would have to be directed towards actively seeking out unusually strong Bluetooth signals in political swing districts. There exist proposals for significant protocol modifications that would mitigate these attacks; however they would likely require the use of user location data in order to provide an effective defense [1, 18].

Note that the GAEN system in iOS 13.5 and 13.6 have countermeasures in place to attempt to detect and discount unrealistic Bluetooth signals, such as those generated in relay and relay attacks at the network layer. These countermeasures are based primarily on straightforward heuristics and remain completely undocumented publicly. They may be able to reduce the efficacy of a naïve physical attack by a significant margin, but this is difficult to verify [2]. The short lifetime of an RPI is designed to make broad replay attacks more difficult, yet RPIs can be transmitted almost instantaneously through the Internet to many other devices who can replay them, leaving RPI replay quite feasible.

6 Possible Mitigations

It is worth stating explicitly that numerous modifications to the existing GAEN protocol have been proposed in the cryptographic literature which could possibly effectively mitigate the attacks we describe here. All that we have seen would unfortunately require a substantial effort in design, implementation and review, and moreover they would generally require making some novel trade-offs that, as far as we know, would significantly change one or more properties of the current system.

For instance, GAEN apps could be allowed to make use of mobile location data that would then be cryptographically blinded before transmission in such a way that user privacy would be largely preserved, and this could be used by the central server to de-duplicate encounters by geographic region [9]. There are also ways of preventing a single cryptographic seed from being used to "expose" an unrealistic number of individuals. However, such methods involve expensive public-key cryptographic computations to be performed on user devices, and this additional overhead could possibly hinder adoption of the GAEN system because of its effect on device performance and battery life.

It is possible for GAEN app developers and state authorities rolling out the apps, to implement a centralized logging or monitoring mechanism that alerts public health authorities as each new exposure event is recorded locally [12]. Such a feature could allow a *rate limiting* mitigation where health authorities "monitor exposure notifications for an increase in rates inconsistent with traditional virus models and put in place mechanisms for responding to such anomalies" [6]. Such a feature of course has the potential to compromise the privacy guarantees of the GAEN project, and would be a major departure from the decentralized paradigm on which GAEN is based. Moreover, without an analysis of the methods used to detect inconsistencies and respond to the anomalies, such method would have unknown efficacy. As far as we know, no health authority has implemented this type of mitigation, though it appears that the state of Virginia is collecting such data as part of its CovidWise app, without using it for any mitigation [12].

We do not consider this the correct place to discuss these various proposals in depth; we would rather focus first on treating the threat posed by the adoption of GAEN in its present form. Furthermore, we feel confident that it would be challenging to roll out any of these protocol modifications to wide adoption before

the 2020 U.S. presidential election in November. A more practical mitigation would be to simply turn off exposure notifications, nationally, in the two-week period preceding the election.

7 Criticism

Some individuals doubt the likelihood that the attacks we describe here will occur. Their primary criticism of our analysis revolves around the idea that the attacks are too expensive and complicated to mount, and that an adversary would choose other simpler methods to achieve voter suppression, such as launching COVID-19 fake news campaigns on social media that would discourage people from voting directly (see for example Ron Rivest's comments in [6], though others have made similar points as well). In the absence of sociological studies on how users will react to a COVID-19 exposure notification on their app (a serious problem of its own), especially in comparison to reading fake news in social media, much of these comments is based on pure speculation, and therefore difficult to address.

While a fake news campaign is potentially cheaper and easier to mount, in our opinion the expected return on a Swing State Attack is much higher. A state-sanctioned message of possible exposure, delivered to a particular individual, should be more likely to lead to self-isolation than a fake news campaign on social media that is clearly not directed at anyone in particular. Indeed, if there is not a high likelihood of an individual self-isolating after receiving an exposure notification, then the value of the GAEN system itself should be called into question.[2] Additionally some subsets of the population, if nothing else, may respond differently to different forms of voter suppression attacks.

Finally it is not clear to us how fake news campaigns would be a cheaper alternative to our Post Office attack. Postal workers cannot refuse to go to work because of something they read on social media, but they probably be required to isolate if their state-sanctioned contact tracing app alerts them that they may have been exposed.

Retail vs Wholesale Attacks. In an attempt to quantify cost and return of an attack, discussions in the election security research community classify attacks in two categories that we will call for simplicity *retail* and *wholesale*[3]. In a retail attack, the cost of the attack is proportional to the number of votes affected, e.g. buying votes. In a wholesale attack the cost is much less than the impact, e.g. hacking the tabulation computer or the software running on voting machines. Voting solutions which offer substantial advantages may be accepted even if susceptible to retails attacks: for example some countries may accept and encourage voting by mail since it increases voting turnout (especially during a

[2] Theoretically, individuals might be expected either to get tested *without* self-isolating, or to self-isolate except to go to a polling station, in which case the app could still have significant utility; however, we consider these two eventualities to be somewhat far-fetched.

[3] We thank Ron Rivest for pointing us to this terminology and encouraging us to discuss our attacks within this framework [20].

pandemic) even if it allows for vote buying and coercion (as opposed to voting in the privacy of a booth). On the other hand wholesale attacks should never be tolerated.

Another criticism of our attack is that it is a retail attack and therefore the benefits of deploying contact tracing apps susceptible to such an attack greatly offset the risk. Glossing over the fact that alternative contact tracing technologies with similar benefits, but without the threats, could be deployed (as discussed above), we strongly disagree that our attack is a retail one. As discussed above, a single malicious update to an app could turn millions of devices into relay beacons, each broadcasting to several users. The app can be chosen to affect specific populations of known political leanings (e.g. hack apps used by a given party for political activism, volunteering etc.) The cost of such an attack is "constant" in the number of targeted individuals, so the impact could be very large. Similarly the Post Office attack (carried out by either a malicious app or physical means) is also wholesale, since shutting down a few strategically chosen postal facilities has the potential of suppressing the votes of a very large fraction of the population.

It is difficult to classify the Swing State attack carried by physical means as either retail or wholesale; it seems to be somewhere in between. The cost seems to be wholesale (a single broadcasting antenna), the physical circumstances seem to be retail (the adversary must be in the physical proximity of many people). But the opportunities to achieve the retail cost without a real penalty abound (e.g. the adversary could perform the attack at crowded Black Lives Matter demonstrations, political rallies etc.). This seems qualitatively much different than the retail cost of say vote-buying which requires the adversary to physically interact with each manipulated voter.

The Adversary An implicit assumption in the criticism of the attack being too costly or complicated to implement is that the adversary has a limited "budget" to spend in attacking the U.S. elections and therefore will choose other, cheaper methods. We strongly reject this assumption: as shown by recent history, foreign nations are more than willing to interfere in U.S. elections [31], using many different means. A nation state should be modeled as an adversary with infinite (or very large) budget, not subject to significant limitations on cost. The question therefore is not if the attack is too costly or complicated (as those are not concerns for an adversary with very large budgets) but if it is a rational choice for an adversary to choose this attack.

We believe that to attack the upcoming U.S. elections, adversaries will most likely choose a mixed strategy that will consist of several elements among all available to them, strategically deployed to obtain maximum effect. Naturally, we expect disinformation campaigns to play an important role in the adversary strategy as they did in past election cycles. We believe, however, that if the GAEN-based contact tracing system is widely adopted by the U.S. before Election Day, that the adversary strategy will update accordingly to incorporate voter suppression through the attacks described in this paper. As listed above, the

reasons are many including:

- the strong legitimacy and value of a state-sanctioned notification to exposure to a dangerous virus
- the ability to manipulate populations not easily reachable *via* disinformation campaigns
- the impossibility to trace the perpetrator of the attack (or even detect the attack itself) due to the decentralized nature of the GAEN system

The "value" of the attack for the adversary is also increased by the byzantine nature of the U.S. election system which features extreme geographic polarization compounded with the winner-takes-all Electoral College. This makes the final outcome extremely sensitive to otherwise small perturbations in voter turnout. As Gomez et al. [10] showed, the ability of *rain* to depress voter turnout in select districts may have had a critical impact on the outcome of at least two U.S. presidential elections. In our opinion, it is reasonable to assume that the impact of a state-sanctioned COVID-19 exposure notification is on the same order of rain on individuals' propensity to vote, if not greater.

Additionally, we must keep in mind that the goal of the adversary may not just be to change the outcome of the election, but rather simply cast doubt on the results. In a democratic system, public trust in the integrity of the election is as important as whether or not the election's results were changed by some attack. Tampering with state-sanctioned public health notification systems has a much larger destabilizing value than launching a fake news campaign on social media, in terms of the potential for compromising citizens' faith in the democratic process.

8 Conclusion

We believe that the GAEN/DP-3T system was designed with a strong emphasis on the concerns for individual privacy and security, and the risk it creates for individuals is well managed, as the consequences of a false exposure notification to a specific individual are minor (namely, the need to get tested and self-isolate). Unfortunately, the protocol design creates an unusual risk for society at large.

The literature on GAEN and DP-3T, and the original proposals themselves, often explicitly contemplate an attacker that might wish to create mass disruption from fake exposure notifications [25], but they haven't sufficiently analyzed: 1) the potential scale of such an attack, 2) the impossibility of mitigating such an attack using traditional means, and 3) the eventuality that such an attack might be used to alter the outcomes of democratic elections *via* voter suppression.

With GAEN as it is currently being implemented and rolled out around the world [30], a single valid authorization code could be used to "expose" millions of users in hundreds of locations simultaneously, and it would be impossible to detect or prevent the attack while it is in progress. This, coupled with the recent history of numerous attempts by foreign governments to interfere in US

elections [31], should worry those who are considering adopting the Exposure Notification system for wide use.[4]

References

[1] DP-3T. *Preventing Inverse-Sybil Attacks #295*. GitHub Issues. URL: https://github.com/DP-3T/documents/issues/295 (visited on 08/19/2020).

[2] Apple, Inc. Personal Communication. July 28, 2020.

[3] Apple, Inc. *Privacy-Preserving Contact Tracing*. URL: https://www.apple.com/covid19/contacttracing (visited on 08/19/2020).

[4] Bluetooth SIG, Inc. *Bluetooth Core Specification Version 5.0 Feature Overview*. URL: https://www.bluetooth.com/bluetooth-resources/bluetooth-5-go-faster-go-further/ (visited on 08/19/2020).

[5] Dan Boneh. Personal Communication. July 24, 2020.

[6] Michael del Castillo. "Google And Apple Downplay Possible Election Threat Identified In Their Covid-19 Tracing Software." In: *Forbes* (Aug. 27, 2020). URL: https://www.forbes.com/sites/michaeldelcastillo/2020/08/27/google-and-apple-downplay-possible-election-threat-identified-in-their-covid-19-tracing-software.

[7] David Catanese. "The 10 Closest States in the 2016 Election." In: *U.S. News & World Report* (Nov. 14, 2016). URL: https://www.usnews.com/news/the-run-2016/articles/2016-11-14/the-10-closest-states-in-the-2016-election (visited on 08/19/2020).

[8] Centers for Disease Control and Prevention. *Test for Current Infection*. https://www.cdc.gov/coronavirus/2019-ncov/testing/diagnostic-testing.html\#who-should-get-tested. Aug. 24, 2020. (Visited on 08/26/2020).

[9] Crypto Group at IST Austria. *Inverse-Sybil Attacks in Automated Contact Tracing*. Cryptology ePrint Archive, Report 2020/670. https://eprint.iacr.org/2020/670. 2020.

[10] Brad T. Gomez, Thomas G. Hansford, and George A. Krause. "The Republicans Should Pray for Rain: Weather, Turnout, and Voting in U.S. Presidential Elections." In: *Journal of Politics* 69 (3 2007), pp. 649–663.

[11] Google Exposure Notifications Server. *Remove DeviceCheck and SafetyNet #507*. GitHub Pull Request. URL: https://github.com/google/exposure-notifications-server/pull/507 (visited on 08/19/2020).

[4] It is worth noting that in this document the authors contemplate only security threats associated with the explicit design of the GAEN system: if any of its various implementations additionally has particular security *bugs*, then an attacker may be able to exploit the GAEN system to the end of voter suppression, independent of the limitations of the protocol which are knowable ahead of time and in a way that subverts mitigations for these limitations which are later developed.

[12] Google, Inc. Personal Communication. Aug. 3, 2020.

[13] Google, Inc. *Exposure Notifications: Using technology to help public health authorities fight COVID-19*. URL: https://www.google.com/covid19/exposurenotifications/ (visited on 08/19/2020).

[14] Google, Inc. & Apple, Inc. *Exposure Notification: Cryptography Specification*. https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ExposureNotification-CryptographySpecificationv1.2.pdf. Apr. 29, 2020. (Visited on 08/19/2020).

[15] Ian Levy. *High level privacy and security design for NHS COVID-19 Contact Tracing App*. National Cyber Security Centre. May 3, 2020. URL: https://www.ncsc.gov.uk/files/NHS-app-security-paper%20V0.1.pdf.

[16] Johns Hopkins University & Medicine. *Daily State-by-State Testing Trends*. Aug. 19, 2020. URL: https://coronavirus.jhu.edu/testing/individual-states (visited on 08/19/2020).

[17] Krzysztof Pietrzak. *Delayed Authentication: Preventing Replay and Relay Attacks in Private Contact Tracing*. Cryptology ePrint Archive, Report 2020/418. https://eprint.iacr.org/2020/418. 2020.

[18] Benny Pinkas and Eyal Ronen. *Hashomer – A Proposal for a Privacy-Preserving Bluetooth Based Contact Tracing Scheme for Hamagen*. GitHub. Apr. 27, 2020. URL: https://github.com/eyalr0/HashomerCryptoRef/blob/master/documents/hashomer.pdf (visited on 08/19/2020).

[19] ProteGO Safe iOS App. *DeviceCheckService.swift*. GitHub. https://github.com/ProteGO-Safe/ios/blob/8f625fbdada98f522c220062d11171d0bdff694e/safesafe/Services/DeviceCheckService.swift. (Visited on 08/19/2020).

[20] Ron Rivest. Personal Communication. Aug. 29, 2020.

[21] Ashkan Soltani. Twitter. May 5, 2020. URL: https://twitter.com/ashk4n/status/1257688292377587712?s=21 (visited on 08/19/2020).

[22] Ashkan Soltani, Ryan Calo, and Carl Bergstrom. *Contact-tracing apps are not a solution to the COVID-19 crisis*. Brookings Institute. Apr. 27, 2020. URL: https://www.brookings.edu/techstream/inaccurate-and-insecure-why-contact-tracing-apps-could-be-a-disaster/ (visited on 08/19/2020).

[23] The DP-3T Project. *Secure Upload Authorisation for Digital Proximity Tracing*. GitHub. Apr. 30, 2020. URL: https://github.com/DP-3T/documents/blob/master/DP3T%20-%20Upload%20Authorisation%20Analysis%20and%20Guidelines.pdf.

[24] The Times Editorial Board. "Editorial: Attacking the U.S. Postal Service before an election is something a terrorist would do." In: *The Los Angeles Times* (Aug. 4, 2020). URL: https://www.latimes.com/opinion/story/2020-08-04/undermining-postal-service-voting-trump (visited on 08/19/2020).

[25] Carmela Troncoso et al. *Decentralized Privacy-Preserving Proximity Tracing*. 2020. arXiv: 2005.12273 [cs.CR].

[26] Serge Vaudenay. *Analysis of DP3T*. Cryptology ePrint Archive, Report 2020/399. https://eprint.iacr.org/2020/399. 2020.

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

[28] Serge Vaudenay and Martin Vuagnoux. "Analysis of SwissCovid." In: (2020). URL: https://chaosticino.ch/docs/20200605--vaudenay%2Bvuagnoux--analysis-of-swisscovid.pdf.

[29] Chris Welch. "Here's why so many apps are asking to use Bluetooth on iOS 13." In: *The Verge* (Sept. 19, 2019). URL: https://www.theverge.com/2019/9/19/20867286/ios-13-bluetooth-permission-privacy-feature-apps (visited on 08/19/2020).

[30] Wikipedia. *Exposure Notification — Wikipedia, The Free Encyclopedia*. http://en.wikipedia.org/w/index.php?title=Exposure\%20Notification&oldid=973638465. [Online; accessed 19-August-2020]. 2020.

[31] Wikipedia. *Russian interference in the 2016 United States elections — Wikipedia, The Free Encyclopedia*. http://en.wikipedia.org/w/index.php?title=Russian\%20interference\%20in\%20the\%202016\%20United\%20States\%20elections&oldid=973831536. [Online; accessed 19-August-2020]. 2020.

[32] Joe Wituschek. "iOS 13 has been installed on 92% of iPhones released in the last 4 years." In: *iMore* (June 19, 2020). URL: https://www.imore.com/ios-13-has-been-installed-92-iphones-released-last-4-years (visited on 08/19/2020).