

POSTER: Privacy Guarantees of BLE Contact Tracing for COVID-19 and Beyond: A Case Study on COVIDWISE

Salman Ahmed
IBM Research
Yorktown Heights, NY, USA
sahmed@ibm.com

Ya Xiao
Virginia Tech
Blacksburg, VA, USA
yax99@vt.edu

Taejoong (Tijay) Chung
Virginia Tech
Blacksburg, VA, USA
tijay@vt.edu

Carol Fung
Virginia Commonwealth University
Richmond, VA, USA
cfung@vcu.edu

Moti Yung
Google and Columbia University
Mountain View, CA, USA
motiyung@gmail.com

Danfeng (Daphne) Yao
Virginia Tech
Blacksburg, VA, USA
danfeng@vt.edu

ABSTRACT

Google and Apple jointly introduced a digital contact tracing technology and an API called “exposure notification,” to help health organizations and governments with contact tracing. The technology and its interplay with security and privacy constraints require investigation. In this study, we examine and analyze the security, privacy, and reliability of the technology with actual and typical scenarios (and expected typical adversary in mind), and quite realistic use cases. We do it in the context of Virginia’s COVIDWISE app. This experimental analysis validates the properties of the system under the above conditions, a result that seems crucial for the peace of mind of the exposure notification technology adopting authorities, and may also help with the system’s transparency and overall user trust.

CCS CONCEPTS

• **Security and privacy** → **Pseudonymity, anonymity and untraceability; Privacy-preserving protocols.**

KEYWORDS

Digital Contact Tracing, COVIDWISE, Exposure Notification, GAEN, COVID-19, Privacy, Security.

ACM Reference Format:

Salman Ahmed, Ya Xiao, Taejoong (Tijay) Chung, Carol Fung, Moti Yung, and Danfeng (Daphne) Yao. 2022. POSTER: Privacy Guarantees of BLE Contact Tracing for COVID-19 and Beyond: A Case Study on COVIDWISE. In *Proceedings of the 2022 ACM Asia Conference on Computer and Communications Security (ASIA CCS '22)*, May 30–June 3, 2022, Nagasaki, Japan. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3488932.3527279>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ASIA CCS '22, May 30–June 3, 2022, Nagasaki, Japan

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9140-5/22/05.

<https://doi.org/10.1145/3488932.3527279>

1 INTRODUCTION

COVID-19 has become the most deadly viral outbreak across the globe since the “The Spanish influenza”. Today, containment and mitigation have been the best strategies, at the start, in the absence of vaccination, and then after initial vaccines have been found, as the strategy when new waves of variants and mutations of the virus appear. Contact tracing can greatly help early containment by tracing from people exposed to newly infected patients and isolating them early [7]. The latest advancement in computer technology aids the contact tracing process by tracking individuals’ mobile devices and their proximity using Global Positioning Systems [15], or Bluetooth Low Energy (BLE) beacons [5, 19, 20].

To combat COVID-19 and aid governments and health organizations with contact tracing, technology companies (Google and Apple, in particular) jointly introduced a Bluetooth Low Energy (BLE) technology called Google/Apple Exposure Notification (GAEN) system in April 2020 [9]. The GAEN system uses interoperable BLE signals to broadcast Bluetooth beacons from one device to another when Android/iOS users come nearby. The Bluetooth beacons help track the distance between the users and the duration of users being in close proximity. When one person is diagnosed as COVID-19 positive at the time of the contact or within a valid time frame of the contact (and only then) the system can notify the other users about potential exposure to a COVID-19 positive person.

Researchers have scrutinized the contact tracing technology and warned that adoption of the technology can have privacy and security issues [3, 4, 8, 11, 12, 16, 22], thus perhaps advocating against its wide adoption. However, researchers scrutinized using attacks based on abstract protocol design and abstract adversaries with extreme settings and economically unjustified (i.e., expensive) scenarios, rather than a typical adversary. Most did not verify the systems based on actual device investigation, and none of them try to find out in which scenarios the system is robust enough against a typical attack.

While scrutiny is always important, none of the earlier works assessed the feasibility of the attacks in realistic situations in terms of operation or cost feasibility vs. gain. This work, like other works [13, 14, 17, 18] that evaluate trust, security, privacy, usefulness, traceability, transparency, and reliability, means to fill the gap and investigate

[§]A preliminary version of the work appeared in IEEE Computer [1].

[§]The opinions and statements in this work (performed as a project within an academic setting) are personal, and do not necessarily represent the employer of this author.

the systems in a balanced way, by inspecting the actual system, and by assessing also strengths and not only weaknesses. Specifically, we perform an analysis of GAEN with two focus points: i) ensuring that the GAEN library code and contact tracing app code protect user privacy, and ii) investigating the privacy shortcomings/flaws in the design and implementation of GAEN, if any. We investigate the above in the context of Virginia’s COVIDWISE [21] app due to its high adaptability rate during this study [2]. This investigation can be useful to understanding the mitigation capability of the system against typical attacks as an explanation toward the system adoption—currently, in future waves of the COVID-19 pandemic, or future pandemic outbreaks.

2 DESIGN OVERVIEW OF GAEN

GAEN uses BLE technology due to its wide availability on a smartphone. Figure 1 illustrates the interactions between a user, the exposure notification system, and the app. The heart of GAEN is a 16-byte random key called Temporary Exposure Key (TEK) generated every 24 hours to make it hard for attackers to track infected users. The Bluetooth beacon’s payload includes an identifier called Rolling Proximity Identifier (RPI) derived from a TEK as an AES encryption key. When a user is infected, its TEKs for the past 14 days are uploaded to the server, and users pull TEKs of infected people from the server, produce a day’s RPIs, and match against their device stored RPI of that day and time, done locally on a device to detect exposure. TEK being a daily key, makes it impossible to link RPIs between days.




Figure 1: Interactions of user and contact tracing app with the exposure notification system.

2.1 Overview of GAEN’s privacy design

Out of the 17 GAEN API calls that allow the interactions between the GAEN system and the app, two APIs – *getTemporaryExposureKeyHistory()* and *provideDiagnosisKeys()* – deal with potentially sensitive information. The first API call fetches the TEKs of the last 14 days from the on-device data store and provides them to the app for uploading to the key server. The second API call inserts one or more batches of TEKs into the on-device data store.

2.2 Threat models and privacy guarantees

We consider four threat levels to discuss GAEN’s privacy guarantees: i) walking trail model, ii) your neighbor model, iii) stalker model, and iv) organized crime model. We define and categorize the threat levels based on attackers’ privilege levels on accessing RPI beacons in real-world scenarios, which are compatible with the assumptions in the existing literature [3, 4, 8, 11, 12, 16, 22]. In the walking trail and your neighbor models, an adversary sniffs a limited amount of beacons for obtaining RPIs, whereas an adversary in a stalking model sniffs a small number of BLE beacons to obtain RPIs. Finally, in an organized crime model, we assume that an adversary can compromise a smartphone, set up a large-scale infrastructure to sniff BLE beacons, and hack health care systems to obtain PINs to share positive information. We discuss the threat levels of these attack scenarios in Table 1.

The privacy guarantees of GAEN and the contact tracing app lie in five key aspects: i) preventing tracking, ii) generating TEKs without using any personally identifiable information (PII), iii) sharing COVID-19 positive diagnostic information without revealing any user information, iv) preventing attackers from obtaining any PII, even if attackers get access to the TEKs, and v) users’ ability to turn ON/OFF GAEN based on their discretion. Furthermore, based on the principle of least privilege, TEKs never leave a user’s device unless the user tests positive.

3 GAEN’S PRIVACY W.R.T. THREAT MODELS

As with all security solutions, the privacy guarantees of GAEN are relative. There certainly exist extreme scenarios (e.g., [3, 6, 10, 16, 20, 22]) where attackers may learn additional information. If an adversary has access to RPIs, TEKs, and RPI date-time information read by (say) thousands of users, then the adversary can profile a user’s movement [3, 16, 20].

Table 1 summarizes the attack difficulty and the leak severity in GAEN under multiple (increasing) threat categories. The first three models capture the most typical threat scenarios (representing small-scale individuals or group effort), in which, it turns out, GAEN leaks no sensitive information. The *Stalker I* model only reveals the approximate number of nearby GAEN users, still not posing any privacy threat. A reported attack [22] relied on the asynchronous change of Bluetooth addresses and RPIs, which is represented in the *Stalker II* model (ID 4) in Table 1. However, this attack no longer works, since GAEN requires the Bluetooth address and RPI to change synchronously, which we experimentally confirmed by extracting around 11k random Bluetooth addresses and RPI pairs from the advertising packets over three days.

Some attack scenarios in Table 1 have rather strong assumptions regarding the complexity of the attack setup and demand huge resources. For example, attackers in the *Organized Crime I* model (ID 5) require TEKs, aggregated metadata (such as date, time, interaction graph, social graph, address, location, etc.), jailbroken/rooted device, and fake contact tracing app to de-anonymize infected users [3, 16, 20]. While obtaining TEKs through a jailbroken/rooted device might be feasible, imitating a contact tracing app is rather difficult. To imitate a contact tracing app, an attacker needs to somehow fool or bypass the authorization system. In addition, a malicious entity cannot fool the contact tracing app to accept forged TEK export files.

Table 1: Privacy leak and severity of leak in GAEN against realistic and complex threat models and their assumptions

ID	Threat Level	Attack Difficulty	Attack Requirement	Attack Goal	Info Leaked	Severity if Leak	Refs
1	Walking Trail	Low	Access to one RPI (common scenario)	Any information about a user	None	None	—
2	Your Neighbor	Low	Access to 0-5 RPIs from 3-5 victims considering neighbors come nearby 0-5 times a day (common scenario)	Any information about a user	None	None	—
3	Stalker I	Low	Access to at least on RPIs in a 10 to 20-minute time window from 5-10 victims	To estimate the number of GAEN users around an attacker	Approximate number of nearby GAEN users	None	[10]
4	Stalker II	Medium	1. Access to RPIs from at least one victim. Tracking a victim for an hour requires all RPIs in that hour 2. Continuity of RPI reception from a victim	To continuously track a user	None (Not trackable based on our observation)	None	[22]
5	Organized Crime I	High	1. Access to unlimited RPIs with location data from 10+ victims 2. Access to published TEKs through jailbreaking or rooting attacker's phone or imitating a contact tracing app 3. Aggregated data for each 10-20-minute time window: date, time, interaction graph, social graph, addresses, location type (residential, workplace, library, etc.), surveillance cameras	To profile movements of infected users and de-anonymize them	Imprecise de-anonymization (precision decreases with increasing number of profiles)	Medium	[20], [3], [16]
6	Organized Crime II	High	1. Access to a victim's smartphone through hacking 2. Storage protection bypass	To obtain the victim's infection status	Information whether the victim is infected or not	Medium	[6]

For maintaining the back-end key server, an authorized contact tracing entity (e.g., Virginia Department of Health) must create a signing key to sign the TEK export files and share the corresponding public key with Google/Apple – ensuring information authenticity. The attack represented by the *Organized Crime II* model in Table 1 (ID 6) is difficult to launch in practice, as it requires the hacker to gain access to the victim's smartphone [6].

4 CONCLUSIONS

Our findings confirmed that GAEN preserves privacy in a comprehensive collection of typical threat scenarios. Compromising user privacy by exploiting GAEN requires an unlikely, complex, or costly attack setup, e.g., compromising a victim's smartphone, mounting many Bluetooth devices, correlating with additional victim information, or rogue access to the healthcare systems. Besides, the built-in authorization, permission, and policy-enforcement mechanisms in GAEN add an extra layer of difficulty against the proposed attacks in the literature. To summarize, and in light of our findings, our article aims at helping people understand and appreciate GAEN's privacy protection and encourage them to adopt GAEN-based contact tracing. This knowledge can be extremely powerful as it will enable us to effectively manage the rest of the current (2020-2022) and future pandemics and, in turn, help reduce unnecessary casualties due to enhanced automatic contact tracing and its advantages, especially given the initial estimates of effectiveness [7, 23].

ACKNOWLEDGMENTS

This work has been supported by the Virginia Commonwealth Cyber Initiative (CCI).

REFERENCES

- [1] Salman Ahmed, Ya Xiao, Taejoong Tijay Chung, Carol Fung, Moti Yung, and Danfeng Daphne Yao. 2022. Privacy Guarantees of Bluetooth Low Energy Contact Tracing: A Case Study on COVIDWISE. *Computer* 55, 2 (2022), 54–62.
- [2] Alejandro De La Garza. 2021. COVID-19 Contact Tracing App Adoption Rates. <https://time.com/5905772/covid-19-contact-tracing-apps/>. Accessed on April 14, 2021.
- [3] Lars Baumgärtner, Alexandra Dmitrienko, et al. 2020. Mind the GAP: Security & Privacy Risks of Contact Tracing Apps. *arXiv preprint arXiv:2006.05914* (2020).
- [4] Antoine Boutet, Claude Castelluccia, et al. 2020. *Contact Tracing by Giant Data Collectors: Opening Pandora's Box of Threats to Privacy, Sovereignty and National Security*. Ph. D. Dissertation. EPFL, Switzerland; Inria, France; JMU Würzburg, Germany; University of
- [5] Ran Canetti, Yael Tauman Kalai, Anna Lysyanskaya, Ronald L. Rivest, Adi Shamir, Emily Shen, Ari Trachtenberg, Mayank Varia, and Daniel J. Weitzner. 2020. Privacy-Preserving Automated Exposure Notification. *IACR Cryptol. ePrint Arch.* (2020), 863. <https://eprint.iacr.org/2020/863>
- [6] Justin Chan, Shyam Gollakota, Eric Horvitz, Joseph Jaeger, Sham Kakade, Tadayoshi Kohno, John Langford, Jonathan Larson, Sudheesh Singanamalla, Jacob Sunshine, et al. 2020. Pact: Privacy Sensitive Protocols and Mechanisms for Mobile Contact Tracing. *arXiv preprint arXiv:2004.03544* (2020).
- [7] L. Ferretti. 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science*, 8 May 2020 368: <https://www.science.org/doi/10.1126/science.abb6936> (2020).
- [8] Adam Krollenstein Rosario Gennaro, Adam Krollenstein, and James Krollenstein. 2020. Exposure Notification System May Allow for Large-Scale Voter Suppression. *arXiv preprint arXiv:2005.12273* (2020).
- [9] Google and Apple Inc. 2021. Exposure Notifications: Using Technology to Help Public Health Authorities Fight COVID-19. <https://www.google.com/covid19/exposurenotifications/>. Accessed on April 14, 2021.
- [10] Rémy Grünblatt. 2021. Stop Covid Detector 3000. https://github.com/rgrunbla/Stop_Covid_Detector_3000. Accessed on April 14, 2021.
- [11] Yaron Gvili. 2020. Security Analysis of the COVID-19 Contact Tracing Specifications by Apple Inc. and Google Inc. *IACR Cryptol. ePrint Arch.* 2020 (2020), 428.
- [12] Vincenzo Iovino, Serge Vaudenay, and Martin Vuagnoux. 2020. *On the Effectiveness of Time Travel to Inject COVID-19 Alerts*. Technical Report. Cryptology ePrint Archive, Report 2020/1393. <https://eprint.iacr.org/2020/1393>.
- [13] Yong Jin Park, Jae Eun Chung, and Dong Hee Shin. 2018. The structuration of digital ecosystem, privacy, and big data intelligence. *American Behavioral Scientist* 62, 10 (2018), 1319–1337.
- [14] Yong Jin Park and Donghee Don Shin. 2020. Contextualizing privacy on health-related use of information technology. *Computers in Human Behavior* 105 (2020), 106204.
- [15] Leonie Reichert, Samuel Brack, and Björn Scheuermann. 2020. Privacy-Preserving Contact Tracing of COVID-19 Patients. *IACR Cryptol. ePrint Arch.* 2020 (2020), 375.
- [16] Otto Seiskari. 2021. Paparazzi Attack PoC. <https://github.com/oseiskar/coronasniffer>. Accessed on April 14, 2021.
- [17] Dong-Hee Shin. 2010. The effects of trust, security and privacy in social networking: A security-based approach to understand the pattern of adoption. *Interacting with computers* 22, 5 (2010), 428–438.
- [18] Dong-Hee Shin. 2010. Ubiquitous computing acceptance model: end user concern about security, privacy and risk. *International Journal of Mobile Communications* 8, 2 (2010), 169–186.
- [19] Hallam Stevens and Monamie Bhadra Haines. 2020. TraceTogether: Pandemic Response, Democracy, and Technology. *East Asian Science, Technology and Society: An International Journal* 14, 3 (2020), 523–532.
- [20] Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, et al. 2020. Decentralized Privacy-preserving Proximity Tracing. *arXiv preprint arXiv:2005.12273* (2020).
- [21] Virginia Department of Health. 2021. COVIDWISE, Official Contact Tracing App in Virginia. <https://www.vdh.virginia.gov/covidwise/>. Accessed on April 14, 2021.
- [22] Martin Vuagnoux. 2020. Little Thumb Attack on SwissCovid. <https://vimeo.com/453948863>. Accessed on April 14, 2021.
- [23] C. Wymant. 2021. The epidemiological impact of the NHS COVID-19 app. *Nature* 594, 408–412, <https://doi.org/10.1038/s41586-021-03606-z> (2021).