

Poster: Can the Exposure Notification Framework Expose Personal Information?

Kazuki Nomoto
Waseda University
nomoto@seclab.jp

Mitsuaki Akiyama
NTT
akiyama@ieee.org

Masashi Eto
NICT
eto@nict.go.jp

Atsuo Inomata
Osaka University
inomata@mail.osaka-u.ac.jp

Tatsuya Mori
Waseda University/NICT
mori@seclab.jp

Abstract—In this study, we assess the privacy risk of de-anonymizing attacks against the Google/Apple Exposure Notification (GAEN) Framework, which is one of the most popular/widespread COVID-19 contact tracing system. The key idea of the attack is to associate the anonymized rolling proximity identifier (RPI) emitted from each device with the face image of the device owner by using a monitoring device equipped with the BLE receiver and camera. An attacker can obtain a large number of pairs of RPIs and facial images. Then, the attacker can obtain RPIs derived from the temporary exposure keys (TEKs) of COVID-19 positive individuals and match them with the observed pairs to obtain facial images of COVID-19 positive individuals. We evaluated the feasibility of the de-anonymizing attack with the extensive field experiments and revealed that it was possible to uniquely link face images of a smartphone owner captured by a camera with the RPI information contained in received BLE frames, in the range of 2m to 10m, implying the success of the de-anonymizing attack.

I. INTRODUCTION

The contact tracing systems are becoming popular as a promising countermeasure against COVID-19. Among the several contact tracing frameworks, decentralized contact tracing system using Bluetooth low energy (BLE) has attracted attentions because it can realize contact tracing with a built-in privacy protection mechanism. The most widespread implementation of the BLE-based contact tracing system is the Exposure Notification Framework (GAEN framework) developed by Google and Apple as it is installed on the mobile OS platforms with the highest market share, namely iOS and Android; it has been adopted by health authorities in 38 countries around the world. We have surveyed the number of installations of GAEN-based apps and found that as of January 2021, more than 20 million installations have been reported in Japan, Germany, and the UK, respectively [1], implying the significant impact of the framework across the world.

There have been many reported cases of discrimination and stigmatization due to the fact that people have been infected with COVID-19 [2]. Therefore, privacy protection should be a top priority in the digital contact tracing technologies. With these backgrounds, this study poses the following research question: *Can the Exposure Notification Framework Expose Personal Information?* More specifically, we assess the privacy risk of de-anonymizing attacks against the GAEN Framework. The key idea of the de-anonymizing attack is to associate the rolling proximity identifier (RPI) contained in the BLE

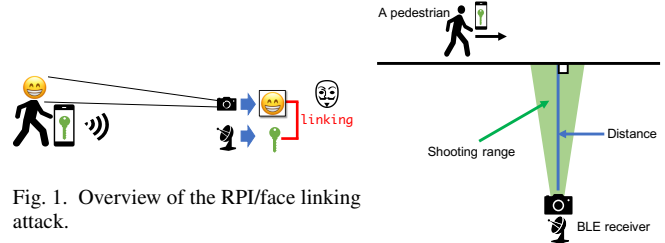


Fig. 1. Overview of the RPI/face linking attack.

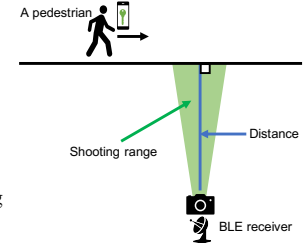


Fig. 2. Setup of the linking attack.

frame emitted from a mobile phone with the face image of the device owner captured by the camera. To this end, an attacker sets up a device equipped with a camera and a BLE receiver and keeps capturing the data on the street. In the GAEN framework, the temporary exposure keys (TEKs) of the positive person, which are collected and managed by the health authority in each country, are periodically distributed to the contact tracing applications implementing the framework. Upon receiving a TEK, the attacker derives the corresponding RPIs and checks against the previously stored RPIs collected from nearby devices. If the attacker finds an RPI associated with the TEK of COVID-19 positive person, they now get a face image of the positive person.

In this work, we focus on the first stage of the attack, i.e., linking an RPI emitted from a device and the facial image of the device owner; i.e., linking attack shown in Fig. 1. We note that the establishing the second stage, associating TEKs of positive individuals with the locally collected RPIs should be a trivial task.

II. ATTACK SCENARIO

Figure 2 presents the setup for the RPI/face linking attack. An attacker uses a BLE receiver to monitor the BLE frames generated by the GAEN-activated smartphones of people in the vicinity and extracts the RPI information from the frames. At the same time, the attacker takes continuous pictures of the pedestrians. As shown in the figure, the goal of an attacker is to link the target's facial image with the RPI generated by the GAEN application implemented on the target's smartphone, which walks towards the location where a camera and a BLE receiver are installed in front of it. The attack is achieved by the following process. First, an attacker collects

TABLE I
LIST OF EQUIPMENTS

Target's Device	Model
Smartphone	Apple iPhone XR / iOS 14.3
GAEN app	COCOA 1.2.1
Attacker's Device	Model
Computer	Panasonic CF-AX2 (laptop) / Ubuntu 20.04.1 LTS
BLE Receiver	EDUP 600M
Antenna	LP0965 Log Periodic PCB Antenna
Camera	BUFFALO BSW505MBK





Fig. 3. Time – Signal strength Graph (distance = 2 m)
Fig. 4. Time – Signal strength Graph (distance = 10m)

the BLE frames observed during the time window when the camera captures the face images of a pedestrian. We note that face recognition can be automated with the computer vision technologies. Next, the attacker extracts the BLE frames that contain RPI information from the collected BLE frames. For each RPI, multiple BLE frames should be observed. If only one RPI is observed in a time window, the attacker can uniquely link that RPI to the face images of a pedestrian. If more than one RPI is observed in a time window (as shown in Figure 3), the BLE frames corresponding to each RPI are examined, and the RPI corresponding to the group of frames with increased signal strength in the time window is adopted.

III. REAL-WORLD EXPERIMENTS

In order to evaluate the feasibility of the attack scenario presented in the previous section, we show the experimental results performed in the real world. List of equipments used in the experiments are shown in Table I. Figures 3 and 4 present the results. In this experiment, the distance defined in Fig. 2 was set to 2 m (Fig. 3) and 10 m (Fig. 4), respectively. In both cases, an attacker was able to uniquely identify the RPI associated with the person. We notice that there is an inherent tradeoff between distance and the length of the time window in which a linking attack is possible. That is, shorter distance results in clearer face images and the reception of BLE frames with larger intensities, but a shorter time window for linking attack. Conversely, if the distance is longer, the face image obtained becomes coarser and the strength of the received signal will be weaker, but the time window for the linking attack will be longer. Evaluating the conditions under which the linking attack is optimized is left for a future study.

IV. DISCUSSION

Countermeasure

The RPI/face linking attack can be mitigated by adjusting the signal strength, RPI transmission frequency/duration, and the protocol extension. By properly adjusting the signal strength and frequency, it becomes difficult for an attacker to establish the linking attack.

Future Work

As implied in Figure 3, if the time window available for capturing the pedestrian's face image is small, an attacker may not be able to identify the increase in signal strength of the BLE frames corresponding to the pedestrian's RPI; i.e., the linking attack may fail. In such a case, the attacker may be able to increase the attack success probability by increasing the number of BLE receivers. Evaluation of such an approach is left for future study.

We demonstrated the feasibility of the attack using iPhone, which had relatively strong signal strength. On the other hand, it is known that several Android devices adopt low signal strengths. We will study possible attack strategies when the signal transmission strength of the target device is weak.

Ethical Considerations

This attack has been shared with Google/Apple, the developer of GAEN, as a potential risk. We obtained the consent of the participants prior to the experiment.

V. CONCLUSION

This study examined the feasibility of a de-anonymizing attack on the GAEN framework. Our experimental results demonstrated that it was possible to uniquely link face images of a smartphone owner captured by a camera with the RPI information contained in received BLE frames, in the range of 2m to 10m. That is, by linking the temporary key of a COVID-19 positive person with the RPI, it is possible to obtain the face image of a person whose RPI has been observed by the attacker. Studying the optimal attack conditions, evaluating the feasibility of the attack under more complex conditions with a large number of pedestrians, and implementation and evaluation of effective countermeasures are left for future research.

Acknowledgements A part of this work is supported by the Security Innovator Training Program, SecHack365 [5], which is operated by the National Institute of Information and Communications Technology (NICT).

REFERENCES

- [1] Kazuki Nomoto, Github `exposurennotification_survey`, https://github.com/nomokazu/exposurennotification_survey/blob/main/downloadNumber.md, (Accessed on 01/25/2021).
- [2] UNICEF, WHO and IFRC, Social stigma associated with the coronavirus disease (COVID-19), <https://www.unicef.org/documents/social-stigma-associated-coronavirus-disease-covid-19>, (Accessed on 01/25/2021).
- [3] Google, Github - `exposure-notifications-server`, <https://github.com/google/exposure-notifications-server/blob/main/tools/export-analyzer/main.go>, (Accessed on 01/23/2021)
- [4] Apple and Google, Exposure Notification Cryptography Specification, https://blog.google/documents/69/Exposure_Notification_-_Cryptography_Specification_v1.2.1.pdf, (Accessed on 01/23/2021).
- [5] NICT, Young Security Innovator Training Program SecHack365, <https://sechack365.nict.go.jp/>, (Accessed on 01/25/2021).

