Decentralized Privacy-Preserving Proximity Tracing: Simplified Overview

8th April 2020. For more information and authors, see the full white paper.

There is growing interest from politicians and epidemiologists around the world in technological approaches to help individuals and countries navigate the COVID-19 pandemic. One suggested approach has been to make use of bluetooth signals on personal smartphones to provide a system that informs users about encounters with individuals who have since tested positive for COVID-19, and provides data to epidemiologists that supports their efforts to model the spread of the pandemic. However, the proposed infrastructures that underlie such *proximity tracing systems* vastly differ in their privacy and security properties. Some proposals may fail to protect highly sensitive data, or can be misused or extended far beyond their initial purpose and lifetime of the crisis. This is all the more important given the global nature of this challenge and the fact that the pandemic reaches across borders and jurisdictions with different levels of fundamental rights guarantees and in times where many governments are functioning under rules of exception.

We currently see **different approaches** emerge across countries and groups:

- **Data grab model**: Approaches that suggest that due to exceptional conditions it is legitimate to obtain location, telecoms, and sensor data that is collected by existing commercial and public infrastructures, centralise this data, and analyse it. They purely rely on legal norms to protect these efforts. This model advocates disproportionate collection of personal data, and assumes legal protections will be sufficient to protect populations which often is not the case.
- **'Anonymised' data approach**: Solutions that propose to anonymise existing location, telecoms sensor data for further use in the pandemic. Anonymisation of personal data is a difficult, if not impossible bar to reach. For location data, for example, this will generally be 'privacy washing', as such rich data is impossible to effectively anonymise. Such solutions also lack in purpose specification and proportionality.
- **Designs to minimise data collection**: Solutions that propose to set up an infrastructure specific to the task that only collects data needed for fulfilling proximity tracing needs of health authorities or epidemiologists. Proposals avoid relying on data collected by existing commercial or public infrastructures that were not set up for the goal of proximity tracing. These solutions can broadly be categorised into two classes: centralized and decentralized models.
    - **Centralized** models attempt to minimise data by generating and keeping track of ephemeral identifiers distributed to users which can be used to construct the contact graph of a user only in the case they are infected. The generation of identifiers and generation of contact graphs are done on a server which is often assumed to be controlled by a government or another trusted entity. This model assumes that the entity running the server shall not misuse the data and capabilities of the server other than when people are infected, for example, at the request of law enforcement, border control or

intelligence agencies. Such protection relies on the protection of the central server which can potentially be repurposed into a ‘data grab’ model.

- **Decentralized** models are designed to keep as much sensitive data on users devices as possible. Methods are introduced to strictly control data flows in order to avoid accumulating any contact data on a centralized server. This means that a server exists but only to enable people to use their own devices to trace contacts. The server is not trusted with sensitive data at all and therefore is not vulnerable to function creep like all the other solutions.

Given the concerns about the effectiveness of legal safeguards, the impossibility of anonymization, and the intrinsic vulnerabilities of centralized data minimization models, we focus on decentralized designs for privacy preserving proximity tracing. As discussed above, centralized designs raise concerns: if they are attacked, compromised or repurposed, they can generate great harm and broadly so. **In order to mitigate these issues, we describe and implement two alternative designs for a proximity tracing system that follow a decentralized approach. Both designs do not require the centralized collection and processing of information on users.** The proposed designs provide build-in, strong, mathematically provable support for privacy and data protection goals and minimise the data required to what is necessary for the tasks envisaged. Furthermore, the designs strictly limit how the system can be repurposed through the application of cryptographic methods and prevent data misuse, for example, by law enforcement.

Both designs operate in four phases:

1. **Installation:** The app is installed on a user’s phone. Each app installation generates a secret piece of data from which it derives a chain of ephemeral broadcast identifiers.
2. **Normal operation:** Each app broadcasts ephemeral identifiers via bluetooth, and records ephemeral identifiers that are broadcast by other apps in the vicinity. Each app rotates the broadcasted identifiers frequently. It is not possible to predict the chain of future identifiers from a given broadcast identifier at any point in time. This prevents a third-party listening out from tracking individuals (e.g. to spot repeat visits to the same place).
3. **Handling infected patients:** If a health authority confirms that a user has been infected, it issues a token that authorises the user to upload data to the backend server. The user, at this point, has to give the explicit permission to submit data to the backend that allows others to reconstruct the identifiers broadcast by her app during the contagious window, i.e. five days before the onset of symptoms. The time period for which this data is made available is determined by the patient together with a health official. This data does not allow backend server or other app installations to identify the infected patient. It is in effect anonymous.
4. **Decentralized contact tracing**: Each app can use the data they download from the backend server to compute privately and on-device whether the app’s user was in physical proximity to an infected person and is potentially at risk of being infected. If they were, the app informs the user to take action.

Additionally, app users can *voluntarily* provide (anonymous) data to epidemiology research centers.

Both designs:

- **Ensure data minimization**. The backend server only observes anonymous identifiers of infected people with no proximity information; health authorities learn no information (beyond when a user manually reaches out to them after notification); and epidemiologists obtain an anonymized proximity graph with minimal information.
- **Prevent abuse of data**. As the different entities in the system receive the minimum amount of information tailored to their requirements, none of them can abuse the data for other purposes, nor can they be coerced or subpoenaed to make other data available.
- **Prevent tracking of non-infected users.** No entity, including the backend server, can track *non-infected users* based on broadcasted ephemeral identifiers.
- **Ensure graceful dismantling.** The system will organically dismantle itself after the end of the epidemic. Infected patients will stop uploading their data to the server, and people will stop using the app. Data on the server is removed after 14 days.

Avoiding the accumulation of sensitive data on a centralised database comes at the 'cost' of localized vulnerabilities elsewhere in the infrastructure. Specifically, we see two types of high-effort attacks on the system which are theoretically possible.

- A tech-savvy adversary could reidentify identifiers of infected people *that they have been physically close to in the past* by i) actively modifying the app to record more specific identifier data *and* ii) collecting extra information about identities through additional means, such as a surveillance camera to record and identify the individuals. This would generally be illegal, would be spatially limited, and high effort.
- A *tech-savvy adversary* deploying an antenna to eavesdrop on Bluetooth connections can learn which connections correspond to infected people, and then can estimate the percentage of infected people in a small radius of 50m.

The two alternative decentralized designs provide a trade-off between bandwidth consumption and privacy. Design 1, described in detail in the white paper, offers a low-cost solution to privacy-preserving proximity tracing but allows a tech-savvy adversary to link the broadcast identifiers reported by an infected individual for the duration of the infectious period. Design 2 addresses this privacy vulnerability, at the cost of increased computation time. Design 2 provides overall better properties, but we will not make a decision about the preferred design before both designs have received sufficient public scrutiny.

Our proposal demonstrates that privacy-preserving approaches to proximity tracing are feasible, and that countries or organisations do not need to accept methods that pose high risks and can be misused. Where the law requires strict necessity and proportionality, and given broad societal support for proximity tracing, the decentralized approach provides an abuse-resistant way to carry it out.

3 April 2020

Contact author: Prof. Carmela Troncoso, EPFL