Lighthouses: A Warning System for Super-Spreader Events

Leonie Reichert*, Samuel Brack*, Björn Scheuermann*†
*Humboldt University of Berlin, Department of Computer Science
{leonie.reichert, samuel.brack, scheuermann}@informatik.hu-berlin.de
†Alexander von Humboldt Institute for Internet and Society, Berlin

***Abstract*—Super-spreader events where one person infects many others have been a driving force of the COVID-19 pandemic. Such events often happen indoors, such as in restaurants, at choir practice or in gyms. Many systems for automated contact tracing (ACT) have been proposed, which will warn a user when they have been in proximity to an infected person. These generally fail to detect potential super-spreader events as only users who have come in close contact with the infected person, but not others who also visited the same location, are warned. Other approaches allow users to check into locations or venues, but these require user interaction.**

We propose two designs how broadcast-based ACT systems can be enhanced to utilize location-specific information without the need for GPS traces or scanning of QR codes. This makes it possible to alert attendees of a potential super-spreader event while still remaining private. Our first design relies on cooperating "lighthouses" which cover a large area and send out pseudonyms. These are recorded by visitors and published by the health authority (HA) in case of an infection. The second design has lighthouses actively communicating with HAs after retrospectively detecting an infected visitor to warn everyone whose stay overlapped.

***Index Terms*—COVID-19, Contact Tracing, Privacy-enhancing technologies, Super-Spreader Detection**

I. INTRODUCTION

Since the outbreak of COVID-19, many infections have been traced back to so-called super-spreader events [1], [2] where a single, sometimes asymptomatic, person infects many others at the same time. Such events are often reported to happen indoors or in locations with tightly packed crowds, such as rehearsals [3], weddings [4], political events [5] or restaurants [6]. With manual contact tracing, incidents of this kind can be detected in retrospect by health authorities (HAs) by conducting interviews with infected patients. When multiple people have been infected at the same place and time, a super-spreader event is likely. In anticipation of a large number of undiscovered infections, the HA can then quarantine and test everyone who was also there.

During the last months, apps have been developed and rolled out with the idea to facilitate and speed up the process of contact tracing by automating it. The most notable approach GAEN [7], supported by smartphones running Android or iOS, relies on a decentralized design where no data about uninfected users is collected centrally. This system was designed to allow users to be warned early in case they have been in contact with an infected person. Here, users publish ephemeral pseudonyms over Bluetooth Low Energy (BLE) and collect those of others. They regularly check the public list of pseudonyms of infected users to determine if any of their past contacts is mentioned. While GAEN provides good privacy for users and can not be repurposed to create surveillance infrastructure, it has been criticized as useless for HAs and not making their task easier [8].

Picture a super spreader event at a restaurant. An infected person I uses a GAEN-based contact tracing app. Users of the app who have been in close proximity to I, will receive a high risk warning by the app as expected. Users who are seated further away than a certain threshold value (usually 2 m) might receive a weak warning. But due to the indoors situation and, e. g., bad air circulation, their risk might be higher than suggested. Users who are out of reach of I's BLE will not receive a warning through the app, even if they might actually be at risk under the given circumstances.

Next, assume the infected person I did not use the contact tracing app. Consequentially, no warnings are created through the app. In some jurisdictions, restaurant operators are required to have customers fill out paper list with contact information [9]. Manually tracing contacts with these paper lists is time consuming and error prone because data might be unreadable, incorrect or incomplete. Due to growing numbers of infections some HA do not have the capability anymore to fully conduct manual contact tracing. Some even ask residents who believe or know that they are infected to warn their contacts themselves and ask them to self-isolate [10].

To solve the problems described in the two scenarios, we propose two designs for a lighthouse super-spreader warning systems that extends the existing GAEN framework. Multiple lighthouses, which are simply BLE-capable smartphones, cooperate to cover a large (indoor or closely-packed) area. They send out lighthouse pseudonyms which are recorded by users. Due to the widespread availability of BLE sensors in smartphones produced during the last years [11], this infrastructure is easy to setup. Our first design relies on a simple broadcasting mechanism. In our second design, lighthouses also collect user pseudonyms and actively check if for any of the past visitors an infection warning is issued. If that is the case, the lighthouses will contact the HA to upload all relevant recorded user pseudonyms. The design focus was placed on user privacy (especially for uninfected users) and usability. The system extends broadcast-based approaches to automated

contact tracing (ACT) such as GAEN [12], but can also be used for message-based proposals such as CAUDHT [13] and Ovid [14]. Our main contributions are

- Two designs to improve ACT applications by handling data regarding past visited locations. Users are warned in case they have visited a location during the same time an infected person was there.
- Only pseudonymous, ephemeral data is passed from an infected user to the HA which does not reveal the user's location history.
- The distribution of warnings does not require any human interaction from the HA. The HA can also manually trigger warnings for infected persons without app.

The contents of this paper are organized as follows. First, we introduce relevant research and existing systems in Section II. In Section III, some definitions are given and criteria relevant for system design are motivated. Next, the design based on passive lighthouses is introduced in Section IV. Afterwards in Section V, the more complex approach using active communication is proposed. Section VI presents possible attacks against both designs and defense mechanisms. The discussion in Section VII considers improvements to the system, especially regarding usability.

II. Related Work

Research on contact tracing has greatly evolved during the 2020 pandemic. Many approaches regarding *automated contact tracing* (ACT) have been proposed in literature or deployed for real world usage. The main focus was on smartphone apps using *Bluetooth Low Energy* (BLE) where users exchange ephemeral pseudonyms with one another in the background. In an earlier work [15], we classified approaches to ACT using BLE in multiple classes. Broadcast-based ACT approaches have been the most widespread due to the involvement of Google and Apple [7] and integration of the *Google Apple Exposure Notification* (GAEN) framework into the respective mobile operating systems. In GAEN, when a person falls ill, they upload their past pseudonyms to the server of the local HA. The HA publishes the uploaded data so all users can regularly download and compare it with the log of pseudonyms they recorded from other users. Everyone who has been in contact with an infected user during their infectious time period will recognize pseudonyms in the download and thereby learn that they are at risk of being infected. Risk assessment is done by evaluating the estimated distance and exposure time to the infected pseudonym. The decentralized nature of GAEN and it's focus on user privacy stops the HA from finding out which locations have been infection hotspots.

There is a range of ACT apps which use GPS data to compare a user's location traces with those of infected individuals to determine who is at risk [15], [16]. These systems generally lack privacy, as they reveal private data of the user and their habits to the HA. Systems that do rely on GPS data but have privacy protection through cryptographic techniques [17], [18] are not yet performant or scalable enough for real world usage. These ACT apps do not require infrastructure or cooperation of business operators as both GPS Data and BLE pseudonyms can be collected at any location.

Another approach to inform people about infection risks is having them check-in to a place. The most simple approach for finding out if a sick person might have infected others at a specific location is by maintaining physical lists where new arrivals write down their names and contact information. Health authorities can then later contact the operator of a place to retrieve the list for manual contact tracing. For example in Germany, this is mandatory during the pandemic for businesses in some states [9]. Many automated or digital solutions use an analogous concept.

In China, a system called Health Codes [19] was rolled out in 2020, where users are provided with a QR code in their app based on their risk of having contracted COVID-19. People with a green code are allowed to travel, but have to get their code scanned on entry and departure to restaurants, public transit, hotels and apartment blocks. It is not always transparent what causes a green code to switch to yellow or red. The system is effectively mandatory as there is little consideration for people not using it or who do not have a smartphone.

Singapore's SafeEntry [20] system also has users scan some form of identity proof to enter a location. This could be an ID card, a QR code generated by a corresponding user-side app or a token. Users can also select a nearby locations to check-in. Check-in data is stored centrally on a government server for 25 days. This data is used for manual and automatic contact tracing processes. The system is inter-operable with Singapore's ACT app TraceTogether [21].

New Zealand's NZ Tracer app [22] is also a tool to warn people if there has been a COVID-19 case in a location they visited. Here, a location generates QR codes which are presented at the entrance. Users can scan the code and store the corresponding information locally. If during contact tracing the HA finds that an infected person visited a location that displays QR codes, it will publish the corresponding information. The app, which regularly check for updates, will warn the user if they checked into a location the same time a person with COVID-19 did. It will not tell them where, but presumably a malicious, tech-savy user would be able to identify it.

A group of researchers has presented their solution Crowd-Notifier [23] to partially digitize the paper-based process used in many European countries. Here, operator of businesses or organizers of events can generate three QR codes: for entry, exit, and tracing. People coming to the location or event can scan the entry code with their app on arrival, which will locally store location and time, protected through encryption. They can also scan the exit code when leaving, although this step is not necessary. If the HA now detects that an infected person visited a certain location or event, they contact the operator for both the paper lists and the tracing QR code. Some information inside the tracing QR code is only readable when decrypted by the public key of the HA. It distributes this information to all users, which will check their own history of locations with this data. If they visited a location during the relevant

time frame, they will be notified by the app. CrowdNotifier cryptographically hides from everyone but the people that have been there during the same time that an infected user visited a specific location or event. No centralized storage of user data is required, since the process of checking ones risk is done locally.

III. Design Space

A successful super-spreader warning system should fulfill a series of requirements that ensure the system is helpful for the HA in containing the outbreak of an epidemic situation. First, we introduce relevant definitions.

- *Location*: A location is a place where people gather in groups at certain times, e. g., restaurants, bars or sports venues. Events like demonstrations and outdoor markets can be seen as locations. Locations are prone to super-spreading because of the high number of people in the same area during the same time.
- *Operator*: An operator is responsible for operating one or multiple locations.
- *Visitor*: A visitor is a person visiting a location.
- *ACT App*: An existing application for BLE-based ACT using a broadcast approach (e.g. apps using GAEN like Corona-Warn App [24]).

There are several important properties that a system should fulfill to be useful as a super-spreader warning system. To ensure that users do trust the system and do not avoid or circumvent it (in case of mandatory usage requirements), user privacy should be one of the main design goals. Tools that can be turned into surveillance infrastructure will lower the adoption rate of such a system [25]. With the principle of data minimization only epidemiologically necessary data should be collected.

Furthermore, an ACT system for detecting super-spreader events should speed up the HA's contact tracing substantially and help to streamline contact tracing processes by automating exposure notifications for large amounts of visitors. Depending on the remaining resources of an HA, notified visitors can be contacted by the HA and ordered to quarantine or even asked to warn their other contacts themselves.

Apart from being trusted by users and helpful to the HA, the system should also fulfill certain functional requirements.

- Visitors of a location with a positive COVID-19 case should be notified, even when they did not record a contact directly. To minimize false positives, time slots should be recorded so that only those visitors are notified that were at the location during the presence of an infected visitor.
- The deployment of ACT apps during the first half of 2020 have shown that such an app will not be used by everyone. Therefore, it is important that a compatible fallback system exists that can be used to notify visitors that are not using the app. This can be achieved by further emphasizing manual contact tracing lists for these visitors.
- The system should work in indoor and outdoor settings. While restaurants, bars, and other similar venues are suspected to increase the spreading of the virus, outdoor events can result in spreading events as well [5]. User trust through guaranteed privacy is especially important in such settings.
- To enhance privacy against attackers that have access to a visitor's phone after a notification was issued, it should not be possible to derive the exact location or time without any additional information. This requirement can protect visitors against attackers in their close personal circle to learn that they visited certain events that might result in social outcasting. Such fears were observed with visitors of a gay club in South Korea where a significant amount of visitors wrote down fake contact information [26].

The idea for super-spreader detection presented in this paper is an enhancement for existing systems for ACT that rely on BLE and a broadcasting approach [15].

IV. Passive Lighthouses

Let's assume a restaurant operator wants to keep their business open during the pandemic. Currently, some jurisdictions require operators to log arrival time, departure time and contact information of all guests [9]. This is often done by handing out pieces of paper to guests and storing the filled out ones for when the HA requires to see them. This process is time consuming and bothersome for both customers and the HA. Customers have to write down their information quite frequently and might be worried that their data will be not protected correctly by the restaurant or repurposed later by law enforcement agencies [27]. For HAs, the paper trail is difficult to work with and might not contain useful information if the operator did not enforce the policy or if visitors provided fake information. For this reason, in some places this paper trail is rarely used [28]. Digital solutions have been rolled out so that customers can insert their information using online forms, but data is still stored in a single place in clear text. The current system, both paper-based and digital, relies on the HA to request customer logs and then manually notify all people mentioned in the logs that are at risk. This is a time consuming process and barely feasible in an situation where in some locations over 90% of all infections can not be traced due to the lack of contact tracing staff [29].

We propose a lighthouse system that can be easily integrated into existing ACT infrastructure. Once it is set up, there is no additional interaction required from either HA or operator. In the following, we will explain how an existing ACT application like GAEN can be extended to support our super-spreader warning mechanism.

A. Setup and Operation

Operators can setup *lighthouses*, i. e., smartphones where a dedicated app is installed, in their locations. Lighthouses continuously emit ephemeral pseudonyms (lighthouse pseudonyms LPs) over BLE. They are organized in groups to

cover areas larger than the reach of BLE. LPs are generated randomly and are distinguishable from BLE pseudonyms broadcast for regular ACT by an additional transmitted prefix. After a certain time T_{duty}, e. g., 30 min, a new LP is generated and broadcast. Since lighthouses themselves do not have any privacy requirements, this time span can be significantly longer than the rotation periods used in ACT. This reduces the load on the system as users would upload less LPs for the same visit to a location.

When a visitor arrives at the location who uses a compatible ACT app, they will broadcast their ephemeral pseudonyms (Ps) and collect those of other users. To take part in super-spreader detection, visitors will additionally collect LPs transmitted by the lighthouse. To mitigate false positives, it is important to consider how long a signal is detected. LPs are only stored if they are received for a duration T_{thres}, e. g., 10 min. This way people simply passing by a location will not accidentally store LPs of places they have not actually visited.

B. Infected Visitors

In broadcast-based ACT approaches, an infected user will upload their used pseudonyms P after being diagnosed. For our super-spreader warning system, the infected user will also upload all LPs they have seen during the relevant time period. Prefixes of LPs are removed before upload. Users can opt out of uploading specific or all LPs. The HA will publish both Ps and LPs. Users will regularly download the all published pseudonyms and locally check if any of the LPs or Ps they have recorded in the past matches. To improve performance an idea mentioned by the authors of DP3T [30] can be used. The HA stores all uploaded pseudonyms in a Cuckoo filter which can be downloaded. After downloading the table, users can check if their recorded Ps and LPs cause a hash collision. To keep failure probability of the Cuckoo filter low, the HA has to create a new filter after some time.

Unlike in normal ACT, proximity information can be ignored for LPs. If a signal from a lighthouse was recorded, it will be used during this step. As soon as a recorded LP appears in the downloaded table, the user is automatically notified that their visit to an event or place overlapped with that of an infected individual. This means they are warned even if they have not been in proximity of the infectious person. This is useful as especially in indoor locations with bad air circulation, where proximity is not the only indicator of an high infection risk [31].

After receiving a warning the user will know that they might be at risk of having contracted the disease and should quarantine until tested negative. More specific, the user will learn only the time at which they could have gotten infected, but not the name of the location nor which infected individual caused the alarm. Users who have not visited this location or visited it at a different time will not detect a hash collision and therefore no alarm is raised. Since risk assessment is done locally, the HA does not learn the location history or identity of users at risk.

C. Infected Person without App

In case the HA discovers during manual contact tracing that an infected person visited a location, this information can also be integrated into the warning system. The HA can ask the location operator to upload the LPs for the corresponding time period to their servers using a one-time token generated by the HA. This prevents misuse of operators and ensures that only locations with confirmed infected persons can upload LPs.

D. Combining Multiple Lighthouses

If a lighthouse only consist of as single device, not much is gained compared to normal BLE-based ACT. Users that see the lighthouse are also likely to see each other. But especially for indoor locations the area to be covered can be bigger than the the reach of BLE. We therefore suggest to set up multiple lighthouses which synchronize their LPs to cover a larger area. For this purpose, LPs have to be created by a single master lighthouse and communicated to the helper lighthouses. This can be done by letting the master create a communication key. This key can subsequently be installed on the helpers and used to create a secure communication channel, e. g., by using a chat protocol. This requires all lighthouses to be connected to the internet. An offline solution using Bluetooth pairings of the synchronized lighthouses can also be implemented. In such a case, the master displays a QR code that is scanned by the helpers to establish a connection over Bluetooth or other local channels, like Wifi Direct.

If a restaurant operator for example wants to cover multiple floors of their restaurant, they can setup one group of lighthouses per floor each with their own master.

E. Warning Users that entered at a later time

It might be helpful to also be able to warn people who did not have an overlap with the infected person, but arrived shortly after the infected person left. Depending on the durability of the virus and the ventilation of the location, new arrivals might still be at risk of getting infected [31]. For this reason, it is convenient if the duration T_{duty} is rather long. So if an infected person left during the beginning of the duty cycle of an LP_t, but stayed long enough to for T_{thres} to be surpassed, users that arrived toward the end of T_{duty} of LP_t will receive a warning. In case the infected person left towards the end of T_{duty} of LP_t, users that arrived during the cycle of the following pseudonym LP_{t+1} will not receive a warning. We therefore suggest to make these cycles overlap and advertise both LPs during that overlap period. The overlap has to at least as long as T_{thres}.

F. False positives

One problem with this passive design is that people who have left before the infected person arrived but recorded the same LP, will also receive a warning even though their risk is very limited. The longer T_{duty}, the more people will receive a false warning in case an infected person arrives towards the end of the duty cycle of an LP. Duration T_{duty} should therefore be not too long. As we see, this optimization criteria

is contrary to what was written in section IV-E. One option to combat this issue would be to allow the operator to adjust T_{duty} dependent on the average duration a visitor stays at their location. The following section will introduce a different approach for which the false positive rate is expected to be lower.

V. Active Lighthouses

To only warn people who have not left when the infected user arrived and thereby minimize the false positive rate, lighthouses can become actively involved in the process. Since visitors also send out pseudonyms P, due to the functionality of the ACT app, these can be recorded by the lighthouses. Setup and operation is similar as described in section IV-A, with only minor differences. When a lighthouse and a visitor can receive the other's pseudonym, they generate a secret S by using P and LP as input for a Diffie-Hellmann key exchange [32]. The lighthouse will store S, P and timestamp T, the visitor only needs to store S.

When a user gets infected, they upload all their own past pseudonyms P as done in regular broadcast-based ACT. They additionally upload all secrets S from the time they were contagious. These will not be made public by the HA. Lighthouses regularly check with the HA which pseudonyms P have been uploaded by infected users and published. If a lighthouse recognizes one P_i from it's history, this means that an infected person has visited the location. The lighthouse then directly contacts the HA. To prove to the HA that it can provide meaningful data, the lighthouse will authenticate itself using the secret S_i which corresponds to P_i. The HA compares the lighthouse's S_i with the one uploaded by the infected user and if they are the same allows the lighthouse to upload LPs from that time to warn other users.

The lighthouse will upload all pseudonyms P of visitors that had an overlap with the infected person's stay. More specifically, it will check the time when P_i was recorded first and last, and uploads all P that fall into this time period. It can be useful to also upload some P that have been recorded shortly after. To ensure that no information is leaked about the location of the lighthouse and thereby about the location history of the infected person, it is necessary that all communication with the HA is conducted though an anonymisation service such as Tor [33]. Users that have been at a location during the same time or shortly after an infected user visited will be informed about their risk. They will not learn the pseudonym of the infected individual that caused the alarm unless they came in close contact and recorded the corresponding pseudonym P. Users that have not visited the location or left before the infected person arrived will not learn that there has been a possible outbreak.

A. Breaking the Link

In order to break the link in the upload from a lighthouse between pseudonyms of visitors that visited the location at the same time, a blind signature scheme similar to the one in our work CAUDHT [13] can be used. Instead of directly uploading relevant visitor pseudonyms P, the lighthouse fetches a blind signature [34] for each relevant P. Afterwards, the lighthouse holds for each P a signature $sig_{HA}(P)$ from the HA. The HA did not learn the value of P nor of $sig_{HA}(P)$. Now, the lighthouse uploads all tuples $(P, sig_{HA}(P))$ using different connections through the anonymisation network. To mitigate timing attacks, it can spread out it's upload as described in [15]. By checking the signature, the HA knows that the uploader has authenticated themselves earlier. Only if the signature is valid, the HA accepts uploaded visitor pseudonyms and publishes them. Users download this data and check locally if they discover one of their own past pseudonyms. If they do, they will know that during the time this pseudonym was active, they visited a location that was potentially an infection hotspot.

B. Infected Visitor without App

It can happen that an infected person visits the location who does not use any ACT app and thereby does not send out any pseudonyms P for the lighthouses to collect. If the HA discovers such a case during manual contact tracing, it demands from the operators of locations identified as relevant during the interview with the infected person that they upload all at-risk pseudonyms. For this purpose, it passes the corresponding time period and an additional secret token to the location operator. The operator manually inserts both in the master lighthouse, which will use the token to authenticate itself with the HA and upload all recorded user pseudonyms from the time period.

C. Combining Multiple Lighthouses

In this design where lighthouses actively communicate with the HA, cooperation between lighthouses works differently. A private communication channel has to be established between lighthouses, e. g., by sharing a authentication code between lighthouses at setup time and relying on an encrypted chat as described in section IV-D. This channel is used for lighthouses in the role of a helper to report recorded tuples of (P_i, S_i, T_i) back to the master lighthouse. The master lighthouse stores all recorded data and takes the responsibility of communicating with the HA.

VI. Security Considerations

To understand the security threats to the designs for the presented super-spreader warning systems, we now present several attacks. We will not discuss threats to broadcast-based ACT apps such as GAEN, but only new attack surfaces introduced by our lighthouse systems.

A. Mapping of Locations to Lighthouse Pseudonyms

It would be harmful, if the HA could find out for arbitrary users which places they have visited. Since in the passive design only pseudonymous LPs, stripped from all static prefixes, are communicated to HA by the infected user this attack is mostly mitigated. Important is hereby, that LPs are derived locally by lighthouses, are changed frequently and do not contain hidden information about their creator.

Healthy users never upload any data. In the active design, infected users upload their secrets S. If the HA does not know the LP from which S was derived, it can not use this information. As long as the lighthouse, which also knows S, communicates anonymously with the HA, no information about the nature of the location, and thereby about the infected users location history, is leaked. If the lighthouse does not use an anonymisation service, no information is leaked that would not be also recorded during manual contact tracing.

If the HA now wants to map LPs to locations, it has to actively collect data. In case of the passive design it could continuously issue requests to locations to upload their LPs. Uploading LPs requires manual interaction from the location operator. Such an attack could therefore be easily detected and would result in a lack of trust and abandonment of the system by location operators. Another way to get access to LPs of locations is eavesdropping on the BLE band. Placing the necessary infrastructure in all possible locations would be rather expensive. But an HA can single out certain locations of interest and record LPs there. This would allow it in both designs to identify if an infected individual visited a certain place during a specific time based on the uploads. But since this information is collected about infected individuals during the manual process as well, not much is gained by the HA.

B. Detecting Closeness on a Social Graph

In the passive design the HA can identify that visits of two users to the same location overlap if they become infected and upload the same LP. Such an overlap might indicate that they know each other or are in the same social circles. This information about infected individuals is also recorded by the HA during manual contact tracing. But since it can leak private information, users can decide to not upload LPs from certain locations or times. To hide their identity, an infected user can also use an anonymity network for their upload. This works as long as upload tokens, usually required for proving to the HA that the uploader is actually infected, are not directly linkable to the user. Some token schemes are discussed in [15]. In the active design, knowing two secret keys S_A and S_B, the HA can not derive if they were recorded at the same time and location. It can only verify a guess for an LP it possesses.

C. Crowd Control

In the passive design, if the HA does learn LPs of a fraction of locations, these could be used for crowd control. For example, it could make sure that on election day a certain demographic is confined at home as they believe to the infectious by publishing LPs of places frequented by these people. In the second design, this attack vector does not apply as users will check for their own past pseudonyms. To publish these pseudonyms, the HA would need the operators of a location to cooperate.

D. Fake Hotspots

A competitor might want to use the lighthouse system to harm a location operator by making the location appear as an infection hotspot. For the design with passive lighthouses, this could be done by collecting LPs and sneaking them into the upload of an infected individual. With the active design, the lighthouse would record most keys of visitors published by the HA. This allows it to perform a sanity check before contacting the HA.

E. Extortion

Another cause of fake infection hotspots can be extortion, as reported in Korea [35]. Infected users demand money for not visiting a location or for not uploading the corresponding LPs (in the passive design) or S (in the active design). All systems where users check into locations can make operators target of such an attack.

F. Network Observer

A network observer is capable of seeing all data that is communicated over the internet between lighthouses and the HA as well as between the HA and users. In the passive design, a network observer will be able to tell who is infected as it can see who uploaded data to the HA. This problem is inherent to broadcast-based ACT approaches. The authors of DP3T [30] proposed to introduce probabilistic cover traffic where any user might communicate with the HA in a way that is not distinguishable from real traffic by an observer. In the active designs, lighthouses communicate with the HA to upload recorded data. To hide the lighthouse's location, Tor is used which makes cover traffic obsolete. If Tor is not used, lighthouses would also have to contact the HA at random so that an observer is not able to tell which locations had a recent outbreak. The communication path from HA to the user in both designs is safe from this attacker as all users receive the same data.

VII. Discussion

A. Neighbors of Locations with Lighthouses

Proximity is only partially considered when visitors try to detect lighthouses, as they will always choose the one that is closest to them. Locations often have neighbors, who live next door but might not come in. These neighbors will detect the installed lighthouses and will be warned in case of an outbreak at the location even though they are not at risk.

This can be partially mitigated by setting a threshold for the distance to the lighthouse, so that visitors will only consider lighthouses which are less than, e. g., 5 m away. To ensure that all visitors can still interact with lighthouses even when for example seated in a corner, the operator has to ensure good coverage. Another option for how to mitigate alarming neighbors would be to have lighthouses transmit a static identifier (e. g., a prefix used for forming groups as discussed earlier), that will not be uploaded to the HA. This allows neighbors to blacklist certain lighthouses for which LPs will not be recorded. To make it more easily usable, this could be done with one simple button press that places all currently received prefixes of LPs on a ignore list.

B. Recording the Correct Lighthouse Pseudonyms

Assume a location has setup for each of their floors separate groups of lighthouses each with their own master. A visitor might detect multiple LPs at the same time, even those from lighthouses that are on a different floor. So that in case of an infection people on a different floor do not receive a warning it is necessary that users only record the LP that was the closest for at least the duration T_{thres}. If several lighthouses are equally close or the error of the proximity measurement is too large to make a meaningful decision, pseudonyms of multiple lighthouses can be stored. By that it is possible to distinguish between different locations that are separated by a wall to minimize false positives. Location changes are still recorded because the signal of the new location will be perceived stronger than the one of the previous location.

C. Integration with Paper Lists

Since pseudonyms of infected users are published by the HA (as by design of broadcast-basted ACT apps), lighthouses can automatically check if an infected person has visited their location. If a visit of an infected person is detected, the master lighthouse can prompt the operator and inform them that they have to provide their paper trail to the HA for manual contact tracing. In the first design this means that lighthouses need to scan LPs published by the HA for their own past LPs. In the setting where lighthouses are active, this detection is already done. This process speeds up detection times for the HA as the operator will approach them instead of the other way around, informing them about an outbreak which would otherwise maybe be detected days later. It also helps the HA in the way that they do not actively have to request lists from operators and then wait until they have replied.

Lighthouses could also directly contact the HA when the past presence of an infected user is detected and communicate the name of the location and the time. But since it is assumed that the lighthouse system is voluntary, the cooperation of operators is required. Any system which automatically forwards data to the HA and reports locations might not enjoy good trust and widespread utilization.

D. Usability

An important feature of the proposed lighthouse system compared to the ones discussed in section II is usability. Users do not have to do any scanning when entering a location, record or reveal their GPS traces but will still receive location specific warnings. This makes the system accessible for example for people who have difficulties using their phone or who are blind.

For usability reasons it is also important that the visitor's application can run in the background. The passive design without prefixes would not require any changes to the current GAEN framework. All other proposals discussed in this paper would require changes.

VIII. Conclusion and Future Work

In this paper we presented a system for sending location-specific super-spreader warnings to users by extending existing BLE broadcast-based system for ACT such as GAEN. It extends the use case of ACT and serves as a tool to deliver exposure notifications quicker than with manual notifications after a location was determined as a potential infection hotspot. No GPS data has to be collected, only BLE is used to exchange pseudonyms between users and lighthouses set up by operators. Lighthouses can cooperate to cover larger areas and thereby warn people about infected users even if they did not see this users pseudonyms.

We presented two designs with different false positive rate and different privacy guarantees. The first one relies on pseudonyms of lighthouses to be recorded by users and in case of infection to be distributed using the existing broadcast infrastructure. In the second design, lighthouses communicate with the HA when they recognize a past visitor as infected and will upload the recorded pseudonyms of everyone whose visit overlapped. Both designs are compatible with existing contact tracing apps and only require minor changes in the existing infrastructure.

References

[1] L. Wang, X. Didelot, J. Yang, G. Wong, Y. Shi, W. Liu, G. F. Gao, and Y. Bi, "Inference of person-to-person transmission of covid-19 reveals hidden super-spreading events during the early outbreak phase," *Nature communications*, vol. 11, no. 1, pp. 1–6, 2020.

[2] S. Chang *et al.*, "Mobility network models of covid-19 explain inequities and inform reopening," *Nature*, pp. 1–8, 2020.

[3] L. Hamner *et al.*, "High sars-cov-2 attack rate following exposure at a choir practice," 2020, accessed: 16. November 2020. [Online]. Available: www.cdc.gov/mmwr/volumes/69/wr/mm6919e6.htm

[4] P. Mahale *et al.*, "Multiple covid-19 outbreaks linked to a wedding reception in rural maine," 2020, accessed: 16. November 2020. [Online]. Available: www.cdc.gov/mmwr/volumes/69/wr/mm6945a5.htm

[5] New York Times, "White house is not tracing contacts for 'super-spreader' rose garden event," 2020, accessed: 10. November 2020. [Online]. Available: www.nytimes.com/2020/10/05/health/contact-tracing-white-house.html

[6] K. A. Fisher *et al.*, "Community and close contact exposures associated with covid-19 among symptomatic adults $\geq$ 18 years in 11 outpatient health care facilities," 2020, accessed: 16. November 2020. [Online]. Available: www.cdc.gov/mmwr/volumes/69/wr/mm6936a5.htm

[7] Google and Apple, "Privacy-preserving contact tracing," 2020, accessed: 11. September 2020. [Online]. Available: www.apple.com/covid19/contacttracing

[8] W. Post, "Apple and google are building a virus-tracking system. health officials say it will be practically useless." 2020, accessed: 16. November 2020. [Online]. Available: www.washingtonpost.com/technology/2020/05/15/app-apple-google-virus/

[9] P. Roos, "No personal data, no food? the new german covid-19 regulations and their data-protection relevance for the food and drink industry," 2020, accessed: 16. November 2020. [Online]. Available: https://digital.freshfields.com/post/102g7db/no-personal-data-no-food-the-new-german-covid-19-regulations-and-their-data-pro

[10] City of Berlin, "Special quarantine regulations in mitte and neukölln," 2020, accessed: 16. November 2020. [Online]. Available: www.berlin.de/en/news/coronavirus/6323014-6098215-special-quarantine-regulations-in-mitte-.en.html

[11] Bluetooth SIG, Inc., "2020 bluetooth market update," 2020, accessed: 28. April 2020. [Online]. Available: www.bluetooth.com/bluetooth-resources/2020-bmu/

[12] Google and Apple, "Exposure notification - bluetooth specification," 2020, accessed: 18. May 2020. [Online]. Available: www.covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ExposureNotification-BluetoothSpecificationv1.2.pdf

[13] S. Brack, L. Reichert, and B. Scheuermann, "Decentralized contact tracing using a dht and blind signatures," Cryptology ePrint Archive, Report 2020/398, 2020.

[14] L. Reichert, S. Brack, and B. Scheuermann, "Ovid: Message-based automatic contact tracing," Cryptology ePrint Archive, Report 2020/1462, 2020, https://eprint.iacr.org/2020/1462.

[15] ——, "A survey of automatic contact tracing approaches," Cryptology ePrint Archive, Report 2020/672, 2020.

[16] National Informatics Centre, Ministry of Electronics & Information Technology, Government of India, "Aarogya Setu Mobile App," 2020, accessed: 16. September 2020. [Online]. Available: www.mygov.in/aarogya-setu-app/

[17] L. Reichert, S. Brack, and B. Scheuermann, "Privacy-preserving contact tracing of covid-19 patients," Poster Session at the 41st IEEE Symposium on Security and Privacy, 2020.

[18] A. Berke, M. A. Bakker, P. Vepakomma, R. Raskar, K. Larson, and A. S. Pentland, "Assessing disease exposure risk with location histories and protecting privacy: A cryptographic approach in response to A global pandemic," *CoRR*, vol. abs/2003.14412, 2020.

[19] Wired, "China's health code system shows the cost of controlling coronavirus," 2020, accessed: 12. November 2020. [Online]. Available: www.wired.co.uk/article/china-coronavirus-health-code-qr

[20] Government of Singapore, "Safeentry," 2020, accessed: 12. November 2020. [Online]. Available: www.safeentry.gov.sg/

[21] ——, "TraceTogether," 2020, accessed: 06. April 2020. [Online]. Available: www.tracetogether.gov.sg

[22] Ministry of Health New Zealand, "Nz covid tracer app," 2020, accessed: 12. November 2020. [Online]. Available: www.health.govt.nz/our-work/diseases-and-conditions/covid-19-novel-coronavirus/covid-19-resources-and-tools/nz-covid-tracer-app

[23] W. Lueks *et al.*, "Crowdnotifier - decentralized privacy-preserving presence tracing," 2020, accessed: 12. November 2020. [Online]. Available: github.com/CrowdNotifier/documents

[24] Deutsche Telekom AG and SAP SE, "Corona-warn-app," www.github.com/corona-warn-app/cwa-documentation, 2020, accessed: 13. May 2020.

[25] F. Buder *et al.*, "Adoption rates for contact tracing app configurations in germany," 2020, accessed: 8. September 2020. [Online]. Available: www.nim.org/en/research/research-reports/adoption-rates-contact-tracing-app

[26] H. Shin and J. Smith, "South korea scrambles to contain nightclub coronavirus outbreak," www.reuters.com/article/us-health-coronavirus-southkorea/south-korea-scrambles-to-contain-nightclub-coronavirus-outbreak-idUSKBN22N0DA, 2020, accessed: 11. May 2020.

[27] Reuters, "German restaurants object after police use covid data for crime-fighting," 2020, accessed: 16. November 2020. [Online]. Available: www.reuters.com/article/us-health-coronavirus-germany-privacy-idUSKCN24W2K6

[28] S. Nachrichten, "Nur eine nachverfolgung in drei monaten," 2020, accessed: 17. November 2020. [Online]. Available: www.stuttgarter-nachrichten.de/inhalt.gaestelisten-in-der-stuttgarter-gastronomie-nur-eine-nachverfolgung-in-drei-monaten.d032fd7b-1083-4ef8-b6f2-199f678347d9.html

[29] "Gesundheitsämter scheitern an der kontaktnachverfolgung," 2020, accessed: 16. November 2020. [Online]. Available: www.rbb24.de/panorama/thema/2020/coronavirus/beitraege_neu/2020/11/gesundheitsaemter-kontakte-nachverfolgung-unklar.html

[30] C. Troncoso *et al.*, "Decentralized Privacy-Preserving Proximity Tracing - Version: 25 May 2020," 2020, accessed: 28. May 2020. [Online]. Available: www.github.com/DP-3T/documents/blob/master/DP3T%20White%20Paper.pdf

[31] G. Correia, L. Rodrigues, M. Silva, and T. Gonçalves, "Airborne route and bad use of ventilation systems as non-negligible factors in sars-cov-2 transmission," *Medical Hypotheses*, p. 109781, 2020.

[32] W. Diffie and M. Hellman, "New directions in cryptography," *IEEE transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, 1976.

[33] R. Dingledine, N. Mathewson, and P. F. Syverson, "Tor: The second-generation onion router," in *USENIX*. San Diego, CA, USA: USENIX, 2004, pp. 303–320.

[34] D. Chaum, "Blind signatures for untraceable payments," in *CRYPTO*. Plenum Press, New York, 1982, pp. 199–203.

[35] T. Guardian, "'more scary than coronavirus': South korea's health alerts expose private lives," 2020, accessed: 17. November 2020. [Online]. Available: www.theguardian.com/world/2020/mar/06/more-scary-than-coronavirus-south-koreas-health-alerts-expose-private-lives