Privacy-Preserving Presence Tracing for Pandemics Via Machine-to-Machine Exposure Notifications

Christos Laoudias*, Marios Raspopoulos†, Stefanos Christoforou‡, and Andreas Kamilaris‡§
*KIOS Center of Excellence, University of Cyprus, Nicosia, Cyprus
†University of Central Lancashire, Cyprus Campus, Cyprus
‡CYENS Centre of Excellence, Nicosia, Cyprus
§University of Twente, Enschede, The Netherlands

***Abstract*—At the onset of Covid-19 several Mobile Contact Tracing Applications (MCTA) were deployed and in many cases contributed to curbing the pandemic by triggering Exposure Notifications (EN) to users who were in proximity to infected users. Recently, a number of MCTA were enhanced with Digital Presence Tracing (DPT) functionality in an effort of the public health authorities to break infection chains mostly in indoor crowded spaces and manage super-spreading events (e.g., concerts, parties). That is, alerting individuals who visited the same place or attended the same event with infected users. This is typically implemented by scanning a QR code at the venue entrance. In this work, we present a DPT solution that relies on EN-Hubs, i.e., Bluetooth-enabled IoT devices, that propagate EN in a machine-to-machine fashion reaching all visitors/attendants seamlessly through their MCTA. The proposed solution removes the overhead of issuing, managing, and scanning QR codes every time people visit a place. In addition, it can be conveniently retrofitted to existing nation-wide MCTA offering DPT capabilities with limited implementation cost.**

I. Introduction

Soon after the outbreak of the Covid-19 pandemic, Digital Contact Tracing (DCT) emerged as a complementary tool to assist Public Health Authorities (PHA) around the globe with the cumbersome task of identifying the contacts of infected citizens through conventional contact tracing. To address privacy concerns and maintain high accuracy in DCT at the same time, the vast majority of Mobile Contact Tracing Applications (MCTA) deployed around the world follow the privacy-preserving Bluetooth-based approach that relies on the Google-Apple Exposure Notification (GAEN) framework[1]. Essentially, MCTA enable users to receive Exposure Notifications (EN) in case they have been in contact (i.e., distance of less than 2 meters for more than 15 minutes) with an infected user that has shared this information through his/her app. Recent field studies at national scale report that MCTA can effectively identify and notify users [1], thus complementing conventional tracing. By the end of March 2022, 73 territories around the world have deployed GAEN-based MCTA[2].

While MCTA rely on *proximity tracing* to notify possible contacts of an infected user, another requirement emerged recently by policy-makers, which relates to anonymously alerting individuals who had been at the same place (especially indoors) with the infected user, without necessarily meeting the distance condition. This is referred to as *presence tracing*, also known as 'bidirectional', 'retrospective' or 'backward' tracing, as opposed to 'forward' contact tracing delivered by the early MCTA deployments. The rationale is that if PHA can locate the source of the infection of the current positive case, then other undetected cases could possibly be found around this original source. This is critical for breaking possible infection chains in indoor spaces where many people come together (e.g., offices, theaters/cinemas, restaurants, bars/clubs, public transportation, etc.) [2] and importantly in the so called *cluster* or *super-spreading* events (e.g., concerts, sport games, wedding/birthday parties, etc.), which include high people's mobility and interaction in limited space [3]. According to recent studies based on modeling [4], backward tracing more than doubles the effectiveness of forward tracing in terms of reducing the virus reproduction number, while adding backward tracing could make forward contact tracing 2-3 times more effective in the UK context [5].

The digital counterpart, i.e., Digital Presence Tracing (DPT), follows different approaches: i) app users voluntary reporting their location or attendance to a venue/event, e.g., keeping a logbook of places visited in the past few days; ii) correlating location coordinates collected by GPS-based apps with public places or events; iii) "check-in" to a place using the smartphone camera to scan a location-specific QR code; and iv) devices exchanging Bluetooth signals is a similar fashion as MCTA. Dedicated protocols for DPT to securely generate, use, and trace QR codes have been developed to preserve users' location privacy. These include the *CrowdNotifier* [6] and the *Cluster Exposure Verification (Cléa)* [7] protocol. The former is part of the *SwissCovid* MCTA, while the latter has been integrated into the *TousAntiCovid* French MCTA for recording visits to public spaces.

The work of C. L. has been supported by the European Union's Horizon 2020 research and innovation Programme under grant agreement No 739551 (KIOS CoE) and from the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy.

The work of A. K. and S. C. has been has been funded from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 739578 and the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy.

The CovTracer-E12N project has received funding from the European Union's Emergency Support Instrument programme under grant agreement No. CYPRUS-LC-015948.

[1]Exposure Notifications API, https://bit.ly/3Dm2cdm

[2]GAEN-based mobile apps worldwide, https://bit.ly/34AQFGw

Several national MCTA offer DPT capabilities for event registration, including UK's *NHS COVID-19* and Germany's *Corona-Warn-App (CWA)* among others. Even though significant effort is being put to ensure cybersecurity and *privacy-by-design* in QR-based DPT solutions, including encryption [6] and dynamically changing QR codes [7], there is still significant overhead to generate and manage the QR codes, while user interaction is needed to scan a code every time a new location is visited (sometimes leading to long queues). Thus, this process inevitably creates misconceptions and concerns regarding location privacy that may reduce their acceptance.

These challenges are addressed by DPT solutions that follow the fundamentally different fourth approach. For instance, the *Lighthouses* system deploys dedicated custom devices that cover an area of interest and broadcast proprietary messages over Bluetooth to notify app users who shared the same area with an infected user without the need for QR codes [8]. However, it introduces unnecessary overhead due to the additional Bluetooth messages broadcasted by the underpinning proprietary DPT protocol. This creates data management overheads that PHA need to handle.

We envision a QR-free DPT system to go hand-in-hand with DCT without the need for a new DPT-specific Bluetooth protocol that introduces undesirable overheads, complexity, and increased message traffic flow. In this work, we propose a solution that leverages the existing GAEN protocol and underlying Bluetooth signalling to deliver efficient, reliable and cost-effective DPT capabilities on top of DCT. Notably, it can be seamlessly integrated into existing GAEN-based MCTA; thus, it inherits the cybersecurity and data privacy guarantees of the GAEN framework.

Assuming that a GAEN-based MCTA is released across a country, the main idea is to deploy *static* Bluetooth-equipped Internet of Things (IoT) devices to adequately cover indoor spaces and venues. These devices act as *exposure notification hubs*, i.e., upon receiving an EN due to a user that reported infection in his/her MCTA, the IoT device A will automatically report "infection" that will trigger notifications in neighboring IoT devices B, C, etc., as well as other users that where close to device A, but possibly not meeting the proximity requirements of DCT with the infected user. After a few iterations, the "wave" of EN will cover the entire indoor space, thus serving as a DPT system. To this end, our contributions in this work are twofold:

- We introduce our privacy-preserving DPT approach and present the system architecture considering full integration with an existing nation-wide MCTA, i.e., the *CovTracer-Exposure Notification (CovTracer-EN)* system, released in Cyprus[3].
- We assess the presence tracing capabilities in a simulated variable-size indoor environment and provide insights about the performance observed.

The rest of the paper is structured as follows: Section II overviews the necessary background for DCT and the exposure risk calculation using the GAEN framework. The joint contact and presence tracing system is presented in Section III including the proposed DPT solution and integration with existing CovTracer-EN. Next, the indoor simulation environment for larger-scale experimentation and evaluation is discussed in Section IV. Finally, Section V provides concluding remarks and directions for future work.

[3]CovTracer-EN official website, https://bit.ly/3x0Yipo

II. Background on DCT

The GAEN framework provides a secure Bluetooth scanning and message exchange with nearby devices (e.g., AES128-based encryption is used). The underlying API enables cross-device interoperability between Google Android and Apple iPhone devices. In practice, GAEN-based MCTA regularly search for proximity information from other users. Essentially, an MCTA-equipped mobile device transmits custom beacon messages through Bluetooth Low Energy (BLE) to preserve battery life. These beacons are picked up by nearby devices running a MCTA and serve as proxy for distance, i.e., they suggest device owner proximity, while the message attenuation level indicates the likelihood that nearby devices have been within a certain distance to the device sending the beacon. Moreover, counting the number of these messages, the receiver estimates the duration of a contact.

In particular, GAEN-based MCTA generate a random Temporary Exposure Key (TEK) once a day. Based on the TEK, Rolling Proximity Identifiers (RPI) are generated and updated approximately every $10\,\mathrm{min}$ to $30\,\mathrm{min}$ to preserve privacy. These RPIs are appended to the messages together with other encrypted metadata such as the wireless transmission power level. In case a positively tested user opts to share this information with his/her consent through the MCTA, the personal TEKs, i.e., infected keys, are uploaded to the MCTA backend server and stored securely together with the TEKs of other infected users. All MCTA users frequently connect to the backend (e.g., a few times per day) and request to download the latest TEKs and use them to identify and match RPIs that were received via beacons and stored locally on the device before. If there is a match, the values reported by GAEN are used to calculate the Exposure Score (ES) that quantifies the infection risk in Meaningful Exposure Minutes (MEM), based on the attenuation levels of the messages and the duration of the contact[4]. In case the cumulative daily ES in MEM is sufficiently high, exceeding a predefined threshold set by epidemiologists (e.g., 15 MEM in the CovTracer-EN app), then EN is triggered on the exposed user's phone device.

A. GAEN-based Risk Calculation

Based on the exposure configuration documentation, we formulate and compute the ES as[5]:

$$ES = w_i \cdot t_i + w_n \cdot t_n + w_m \cdot t_m + w_o \cdot t_o, \qquad (1)$$

[4]Exposure Risk Value Calculation, https://apple.co/35tMNL6

[5]The ES function contains additional multiplicative terms including the *Infectiousness* weight that is higher for those TEKs closer to the symptoms onset date of the infected user and the *Report type* weight that is higher if the user has a confirmed test vs a self-diagnosis. However, for simplicity in this work we set both weights equal to 1 (100%) and ES comes down to (1).

		Immediate	Near	Medium	Other
Narrower Net 1.0	*Threshold*	<55 dB	<63 dB	<70 dB	--
	Weight	150%	100%	40%	0%
Wider Net 1.0	*Threshold*	<55 dB	<70 dB	<80 dB	--
	Weight	200%	100%	25%	0%

Fig. 1. Common configurations of BLE attenuation thresholds and weights.

where the parameters t_i, t_n, t_m, and t_o are the exposure durations in the attenuation ranges (buckets) $B_i : [0, a_i)$, $B_n : [a_i, a_n)$, $B_m : [a_n, a_m)$, and $B_o : [a_m, \infty)$, while w_i, w_n, w_m, and w_o are their corresponding user-defined weights. Essentially, the range of attenuation values is divided into four buckets based on user-selected thresholds a_i, a_n, and a_m. The buckets coarsely approximate the proximity between two devices as *Immediate*, *Near*, *Medium*, and *Other*, i.e., very far. This design decision is taken by the GAEN team to handle the inherent uncertainties in the BLE signals that do not allow for accurate estimation of distance from attenuation values. If $ES \geq T_e$, where T_e is the exposure duration threshold, then the user is notified through the MCTA. For instance, $T_e = 15\,\text{min}$ in CovTracer-EN.

All these parameters can be selected by the PHA through extensive experimentation, e.g., to minimize *false negatives* (i.e., undetected critical contacts) and maximize *true positives* (i.e., detected critical contacts), as in the case of the CWA [9]. Luckily, there are some thoroughly-studied and well-documented recommendations for the parameter values that are fine-tuned to provide different performance in terms of the exposures captured and the volume of notifications triggered.

Figure 1 depicts the thresholds and weights for two configurations that were recommended by the Risk Score Consortium in November 2020 and released as open source by the Linux Foundation Public Health[6]. Essentially, the *Narrower Net* configuration captures some fraction of close contacts and limits the number of further-distance exposures captured; thus, triggering fewer notifications. In contrast, the *Wider Net* configuration captures most close-contact exposures and a non-negligible amount of further-distance exposures; thus, triggering more notifications.

For instance, assuming the Wider Net configuration with $T_e = 15\,\text{min}$, if a user spends $8\,\text{min}$ in immediate distance or $2\,\text{h}$ in medium distance from a positive case, then a notification will be triggered because $SE = 16$ and $SE = 30$, respectively.

III. Joint Contact & Presence Tracing System

A. Outline of the proposed DPT solution

An indoor office environment featuring a room with plasterboard walls and an open-plan area with two rows of cubicles is illustrated in Fig. 2. Two EN-Hubs, i.e., Bluetooth-equipped IoT devices (e.g., cheap Android smartphones, Android boxes, Raspberry Pi running Android, etc.) denoted H_1 and H_2 are running the GAEN protocol with the Wider Net configuration, as shown in Fig. 1, to increase coverage. The three co-centric circular discs represent the *Immediate*, *Near*, and *Medium* attenuation ranges, respectively[7]. The dashed circle denotes larger attenuation, while the red-orange-green colors indicate the weights associated with the attenuation ranges.

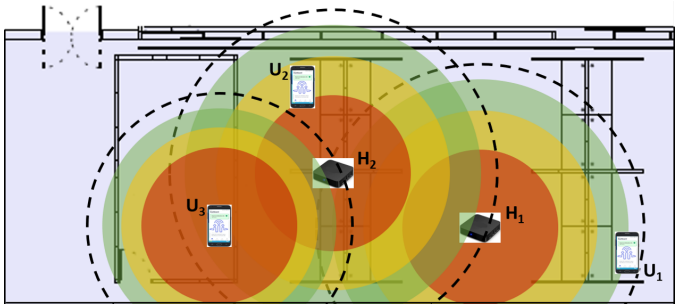

Fig. 2. Example deployment of EN-Hubs in the proposed DPT solution.

There are also three employees who carry their smartphones running a GAEN-based MCTA with the Narrow Net parameter configuration according to the recommendation of the local PHA for DCT. Note the smaller sizes of the discs for U_3 that correspond to the lower thresholds and the smaller attenuation ranges compared to the EN-Hubs. In case user U_1, who did not move from his desk the whole day, reports infection through his/her MCTA, then H_1 receives EN because U_1 is in the *green* range of H_1 for sufficient time. Subsequently, H_1 acting as an EN-Hub automatically reports "infection" that triggers an EN on H_2 because H_1 is in the *green* range of H_2. In a similar fashion, H_2 reports "infection" that triggers an EN on U_3 because H_2 is in the *green* range of U_3. In addition, U_2 receives EN because H_2 is in the *orange* range of U_2 (the range discs of U_2 are omitted for clarity).

Note that if another user was sitting on his/her desk at the top right for the whole day, then he/she would not receive EN neither from U_1 nor from H_1 because they are both in the *Other* range of that user with zero weight. This could be easily addressed with additional EN-Hubs; however, the optimal number and placement of the EN-Hubs are beyond the scope of this work.

B. Integrated DCT and DPT architecture

The architecture for joint DCT and DPT is illustrated in Fig. 3. The right part (brown color) is the simplified block diagram of the CovTracer-EN [10] national MCTA of Cyprus that consists of the following four main components.

SNOW Platform: is used by the Ministry of Health for managing the positive cases in Cyprus, conducting conventional contact tracing, and issuing One-Time-Passwords (OTP) to CovTracer-EN users that opt to share their infected TEKs.

[6]Configuring EN Risk Scores for COVID-19, https://bit.ly/3iVCgvT

[7]This is only for illustrating the attenuation ranges. In open-space signal propagation, doubling the transmitter-receiver distance leads to 6 dB additional attenuation.

CovTracer-EN App: is a light-weight and user-friendly GAEN-based client running on the users' smartphones for triggering notifications upon potential contact with another infected user. The app requests and downloads the infected TEKs as well as the user configuration parameters from the Backend automatically twice per day, whenever a data connection is available. It also allows a user to report infection with his/her consent and upload the infected TEKs to the Backend for other users to check for possible exposures.

Verification Server: is responsible for i) creating the OTP to enable the infected user to share his/her TEKs and ii) verifying that the OTP submitted together with the TEKs is valid.

CovTracer-EN Backend: manages i) the collection of the TEKs submitted by infected users, ii) storing them in a database in case the OTP is valid, and iii) the distribution of the infected TEKs to all app users together with the configuration parameters for computing the ES value daily on each smartphone to trigger a notification, if and when needed.

The left part in Fig. 3 shows the EN-Hub IoT device that implements the proposed DPT solution together with its associated data flow (gray arrows). The EN-Hub runs a slightly modified version of the original CovTracer-EN, having two key differences: i) it requests separate EN-Hub configuration parameters twice per day, together with the same infected TEKs downloaded by all CovTracer-EN users and ii) upon receiving a notification, it automatically reports infection, i.e., uploads its "infected" TEKs to the CovTracer-EN Backend bypassing the OTP creation and verification process.

Let's assume that Bob has CovTracer-EN enabled on his smartphone and goes to the theater with some friends. The theater is equipped with EN-Hub devices that adequately cover the entire hall for the purpose of DPT. Alice has also installed CovTracer-EN on her smartphone and attends the same show sitting several rows away from Bob. At the next day, Bob is tested positive for Covid-19. The following steps take place that enable Alice and other attendants to receive a notification (without any indication that this was due to Bob) and follow PHA recommendations, e.g., self-isolation, testing, etc.; thus, stopping a super-spreading event at the earliest stage.

1) During Bob's interview with the contact tracers an OTP is requested for CovTracer-EN.
2) The Verification Server creates the 12-digit OTP.
3) The OTP is sent to Bob via SMS.
4) Bob's infected TEKs and OTP are submitted.
5) The Verification Server checks the OTP.
6) The OTP is verified and Bob's infected TEKs are stored.
7) Bob's TEKs are distributed by the Backend and downloaded by EN-Hubs and regular users. Bob's friends, who were sitting next to him, will receive notification due to proximity with him as part of the usual DCT.
8) The EN-Hub that was closest to Bob receives a notification and automatically uploads its "infected" TEKs. When these TEKs are distributed, the notification is propagated to other EN-Hubs, as well as regular users who were not in proximity with Bob.

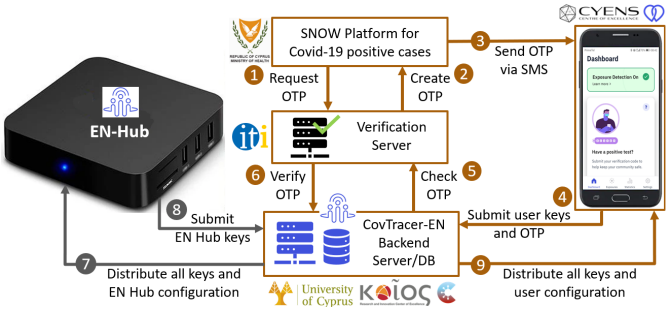

Fig. 3. Architecture of the integrated DCT and DPT solution.

9) Eventually, a notification is triggered on Alice's smartphone due to the "infected" TEKs of the EN-Hub that was closer to her (possibly several hops away from Bob).

We investigated the feasibility of the proposed solution using the *Test* environment of CovTracer-EN. This is an exact copy of the *Production* system that allows to test changes and upgrades in a fully controlled manner before pushing them to the real system. We used four commercial Android smartphones: two devices U_1 and U_2 acting as users had the original CovTracer-EN installed, while the other two devices H_1 and H_2 serving as EN-Hubs were running the modified CovTracer-EN that behaves as described previously. The devices were placed along a straight line on a bench with approximately $6\,\mathrm{m}$ distance between the EN-Hubs. Each user device was placed approximately $2\,\mathrm{m}$ away from the corresponding EN-Hub, i.e., the distance separating U_1 from U_2 was approximately $8\,\mathrm{m}$.

After $15\,\mathrm{min}$ a "fake" infection was reported from U_1 and its TEKs were uploaded to the Backend in the *Test* environment. Subsequently, the nearby H_1 received EN and its TEKs were automatically shared. Due to the limited exposure time and the distance between the user devices, U_2 was not notified. Subsequently, H_2 received EN (which was due to the "infected" TEKs of H_1) and uploaded its TEKs. Finally, U_2 was notified about a possible exposure due to H_2. This would not be possible without the EN-Hubs.

IV. Indoor Simulation Environment

To carry out a more extensive evaluation of our DPT solution, a 3D presence tracing simulator was implemented in MATLAB. It comprises a 3D random walk mobility model and a deterministic channel model. The mobility model builds upon the open source 2D random walk model [11] that is purely open space, not considering any walls or obstructions in the environment. Our model extends the mobility to 3 dimensions, while it introduces the functionality for mobile users to bounce back once they hit a wall of the environment. Additionally, the model introduces *static* nodes, i.e., EN-Hubs, which the original model did not include. The simulator estimates the instantaneous power between all nodes (users and EN-Hubs) by using a custom-built Ray Tracing Model [12].

Our setup is a typical $15\,\mathrm{m} \times 50\,\mathrm{m}$ office environment comprising sixteen $20\,\mathrm{m}^2$ offices, three $40\,\mathrm{m}^2$ conference rooms, a large horizontal corridor, and a smaller vertical corridor, as

shown in Fig. 4. All external walls including the floor and ceiling are assumed to be made of reinforced concrete, while internal walls are made of plasterboard and doors are wooden. All nodes use the $2.4\,\mathrm{GHz}$ BLE frequency band including 35 equally spaced EN-Hubs, i.e., 1 inside each room (not shown) and 1 every $10\,\mathrm{m}$ along the corridor, as well as 100 pseudo-randomly placed regular users, i.e., at least 2 in each room, 10 in the corridor and some outside the building.

The users move for a period of $20\,\mathrm{min}$ with a variable speed ranging randomly between $1.2\,\mathrm{m\,s^{-1}}$ to $1.9\,\mathrm{m\,s^{-1}}$ in all 3 dimensions, while the altitude of any user's mobile device can vary between $1.2\,\mathrm{m}$ to $1.9\,\mathrm{m}$ above ground. They walk in intervals of $2\,\mathrm{s}$ to $6\,\mathrm{s}$ before deciding to either change direction or take a pause for a period of up to $3\,\mathrm{s}$. If during their mobility they encounter a wall, then they bounce back according to Snell's law of reflection. The attenuation between all nodes is estimated every $2\,\mathrm{s}$ using the channel model, which was configured to account for up to 2^{nd} order reflections and all the possible refractions through the building walls. Typical electrical properties for the building materials were assumed.

At the beginning, user U_{52} located in the lower left room reports infection (green diamond in Fig. 4) and remains inside the room for the whole simulation, as indicated by the red dotted trajectory. At the end, we check which users and EN-Hubs receive notifications taking into consideration that EN-Hubs are propagating notifications as described in Section III-A. The ES values are computed according to (1) assuming the Narrow Net configuration for the user devices, while the EN-Hubs use the Wider Net configuration.

By activating and de-activating various combinations of EN-Hubs, we found that for the above GAEN configurations, all EN-Hubs placed along the corridor receive EN and report "infection" one after the other (purple triangles), thus causing users in their surrounding rooms to be notified (yellow stars). This is due to the low attenuation caused by the plasterboard internal walls to the EN-hub transmissions allowing the users inside the rooms to receive strong signals from them and accumulate risk beyond the exposure threshold. On the other hand, users that are outside the building (blue circles) remain unnotified as the attenuation through the external concrete walls is much higher.

We demonstrate the importance of EN-Hub deployment in Fig. 5, where we selectively disabled EN-Hubs H_{25} and H_{27} (gray triangles). Assuming the same user mobility profiles as before, we make 2 important observations: i) the users in the rooms next to those 2 EN-Hubs do not get notified and ii) EN-Hub H_5 in the vertical corridor remains unnotified (red triangle) as the presence of 2-3 internal walls between itself and other EN-Hubs reduces the signal not allowing to the risk exposure to build up and exceed the threshold. In this context, the presence of H_{25} becomes vital to propagate the EN to H_5 and subsequently to other users. However, users U_{47}, U_{112} and U_{115} (black dashed line) in the close vicinity of H_5 do receive EN because they have moved in the horizontal corridor and close to other "infected" EN-Hubs for enough time as indicated by blue dotted trajectory of U_{115}.

V. Conclusions

We presented a Bluetooth-based privacy-preserving DPT approach that removes the burden of issuing, managing, and scanning QR codes to notify users who visited a place or attended an event where an infected person was also present at the same time. The proposed solution can be conveniently combined with an existing GAEN-based MCTA deployed at national level, thus offering DPT capabilities on top of DCT with minimal integration effort.

As part of our future work we plan to validate the proposed solution in a large-scale indoor setting with several real devices. We will also investigate and answer the following important questions: Does the use of EN-Hubs raise additional cybersecurity and/or data privacy issues? What is the optimal number and placement of EN-Hubs? What is the optimal parameter configuration for EN-Hubs with respect to the associated MCTA configuration given by the PHA? In the last two questions optimal refers to achieving high DPT performance at the same level as QR-based solutions, while keeping the number of EN-Hubs as low as possible.

References

[1] C. Wymant, L. Ferretti, D. Tsallis, M. Charalambides, L. Abeler-Dörner, D. Bonsall, R. Hinch, M. Kendall, L. Milsom, M. Ayres *et al.*, "The epidemiological impact of the NHS COVID-19 app," *Nature*, vol. 594, no. 7863, pp. 408–412, 2021.

[2] Y. Li, H. Qian, J. Hang, X. Chen, P. Cheng, H. Ling, S. Wang, P. Liang, J. Li, S. Xiao, J. Wei, L. Liu, B. J. Cowling, and M. Kang, "Probable airborne transmission of sars-cov-2 in a poorly ventilated restaurant," *Building and Environment*, vol. 196, p. 107788, 2021.

[3] S. L. Miller, W. W. Nazaroff, J. L. Jimenez, A. Boerstra, G. Buonanno, S. J. Dancer, J. Kurnitski, L. C. Marr, L. Morawska, and C. Noakes, "Transmission of sars-cov-2 by inhalation of respiratory aerosol in the skagit valley chorale superspreading event," *Indoor Air*, vol. 31, no. 2, pp. 314–323, 2021.

[4] W. J. Bradshaw, E. C. Alley, J. H. Huggins, A. L. Lloyd, and K. M. Esvelt, "Bidirectional contact tracing could dramatically improve COVID-19 control," *Nature communications*, vol. 12, no. 1, pp. 1–9, 2021.

[5] A. Endo *et al.*, "Implication of backward contact tracing in the presence of overdispersed transmission in covid-19 outbreaks," *Wellcome open research*, vol. 5, 2020.

[6] W. Lueks, S. Gürses, M. Veale, E. Bugnion, M. Salathé, K. G. Paterson, and C. Troncoso, "CrowdNotifier: Decentralized privacy-preserving presence tracing," *Proceedings on Privacy Enhancing Technologies*, vol. 2021, no. 4, pp. 350–368, 2021.

[7] V. Roca, A. Boutet, and C. Castelluccia, "The Cluster Exposure Verification (Cléa) Protocol: Specifications of the Lightweight Version," 2021.

[8] L. Reichert, S. Brack, and B. Scheuermann, "Lighthouses: A warning system for super-spreader events," in *2021 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2021, pp. 1–6.

[9] S. Meyer, T. Windisch, A. Perl, D. Dzibela, R. Marzilger, N. Witt, J. Benzler, G. Kirchner, T. Feigl, and C. Mutschler, "Contact tracing with the exposure notification framework in the German Corona-Warn-App," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2021, pp. 1–8.

[10] P. Isaia, C. Laoudias, A. Kamilaris, and C. G. Panayiotou, "CovTracer-EN: The Journey of Covid-19 Digital Contact Tracing in Cyprus," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2021, pp. 1–8.

[11] D. Rawat. Two dimensional random walk - matlab central file exchange. [Online]. Available: https://www.mathworks.com/matlabcentral/fileexchange/14864-two-dimensional-random-walk

[12] M. Raspopoulos and S. Stavrou, "Frequency selective buildings through frequency selective surfaces," *IEEE Transactions on Antennas and Propagation*, vol. 59, no. 8, pp. 2998–3005, 2011.

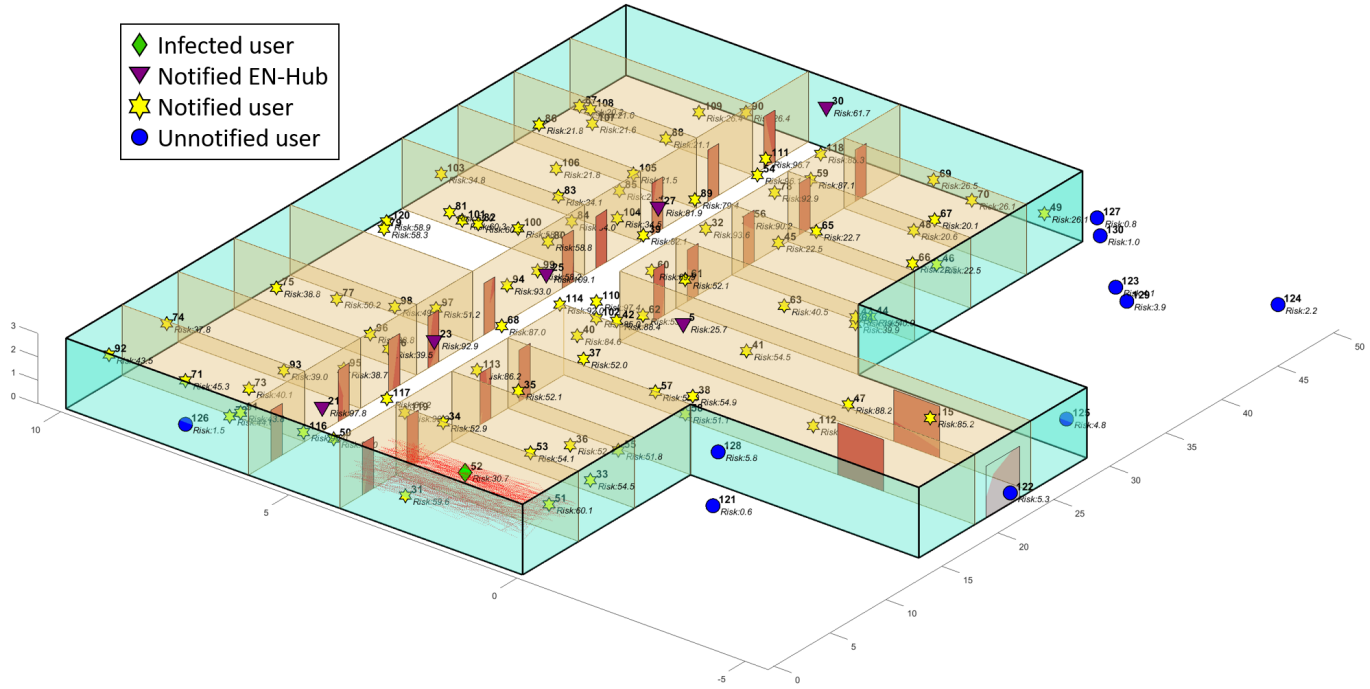

Fig. 4. 3D view of the indoor simulation environment: infected user (green diamond), his travelled trajectory within the bottom left room (red dotted line), notified EN-Hubs along the horizontal and vertical corridors that become "infected" and propagate the notifications (purple triangles), users that are notified either due to the infected user or the "infected" EN-Hubs (yellow stars), and users outside the walls who do not receive notification because of the high signal attenuation (blue circles).

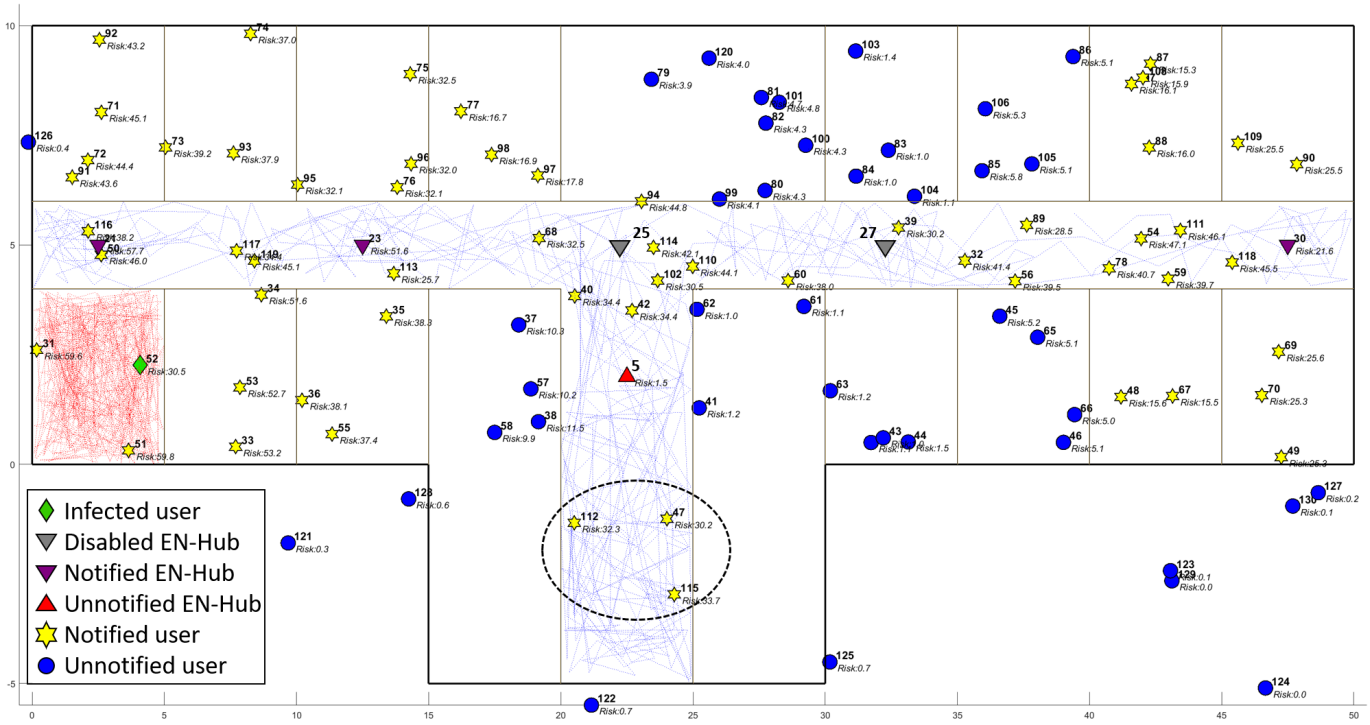

Fig. 5. 2D view of the indoor simulation environment: infected user (green diamond), his travelled trajectory within the bottom left room (red dotted line), disabled EN-Hubs H_{25} and H_{27} (gray triangles), notified EN-Hubs along the horizontal corridor that become "infected" and propagate the notifications (purple triangles), unnotified EN-Hub H_5 in the vertical corridor (red triangle), users notified either due to the infected user or the "infected" EN-Hubs (yellow stars), and unnotified users either outside the walls because of the high signal attenuation or inside the walls in the vicinity of the disabled EN-Hubs (blue circles). Some users in the vertical corridor (back dashed line) are notified because of their partial movement along the horizontal corridor during the simulation, e.g., the trajectory of user U_{115} (blue dotted line).