

Privacy-Preserving and Sustainable Contact Tracing Using Batteryless BLE Beacons

Pietro Tedeschi*, Kang Eun Jeon[†], James She*[†], Simon Wong[†], Spiridon Bakiras*, Roberto Di Pietro*

*Division of Information and Computing Technology, College of Science and Engineering,

Hamad Bin Khalifa University — Doha, Qatar

[†]HKUST-NIE Social Media Lab., The Hong Kong University of Science and Technology — Hong Kong

Email: *{ptedeschi, sbakiras, pshe, rdipietro}@hbku.edu.qa, [†]{kjeon, eejames, tywongbf}@ust.hk

Abstract—Contact tracing with mobile applications is an attractive approach for many governments and industry initiatives to address the COVID-19 pandemic. However, many approaches today have severe privacy and security issues, and many of them also fail to offer a sustainable contact tracing infrastructure due to the demanding energy consumption. This work makes several contributions towards overcoming these limitations. First, we propose a privacy-preserving architecture for contact tracing that leverages a fixed infrastructure of BLE beacon transmitters. Second, we evaluate the feasibility of adopting batteryless or energy-harvesting BLE beacons to make this architecture more sustainable and green. Finally, we identify practical research challenges and opportunities for academia and industry to advance and realize the proposed privacy-preserving and sustainable contact tracing architecture.

I. INTRODUCTION

Smartphone-based contact tracing protocols [1] have been adopted by many countries in order to help fight the spread of COVID-19. Most practical implementations today follow a common *modus operandi*: mobile devices continuously broadcast pseudo-random Bluetooth Low Energy (BLE) beacons that are received and stored by other devices in the communication range; subsequently, the collected data are reconciled in either a centralized or decentralized fashion, in order to identify potential contagion events. However, this approach not only increases the energy burden on the user’s smartphone—via constant BLE scanning and broadcasting operations—but also inherently imperils user privacy.

Indeed, even though the user’s beacons are pseudo-random and change every few minutes, there is still a vulnerability window that allows an eavesdropping adversary to track the user’s location. Such concerns are further amplified by incorrect software implementations, such as the Google/Apple’s privacy bug found in their COVID-19 exposure notification framework. This particular bug failed to synchronize the change in the pseudo-random beacons with the OS’s periodic change of the Bluetooth MAC address, thus allowing adversaries to correlate previous and new pseudo-random beacons and “continuously trace” the victim’s location. The above highlighted native privacy and energy concerns in existing solutions undermine the very purpose of contact tracing applications, hindering their adoption by the general public [2], [3].

To mitigate the aforementioned privacy and energy issues, we propose the deployment of a lightweight wide-scale contact

tracing infrastructure, consisting of cheap BLE transmitters. The beacons transmitted by these devices would replace the smartphone-generated beacons, but would still allow for accurate proximity tracing for the purpose of exposure notification. In particular, the users’ smartphones would constantly intercept and store the infrastructure-based beacons, thus gradually building a record of their precise location over time. Then, the exposure notification process would develop as in most standard BLE-based protocols. The benefits of this approach are threefold: (i) unconditional privacy for users, since their devices are not emitting any information; (ii) reduced energy requirements for smartphones, which translates into longer battery life; and, (iii) potential for more accurate proximity detection, due to the presence of multiple (fixed) BLE transmitters.

Nevertheless, to achieve the objective of an easy and ubiquitous deployment, the BLE devices must be battery-powered. This latter point would trigger the issue of periodic battery replacement, which in turn would considerably increase the operational and maintenance cost of the infrastructure. Such overhead is further amplified in large-scale deployment cases. As an example, the Hong Kong International Airport had to deploy over 17,000 BLE devices to provide indoor navigation services.

Contributions In this paper, we provide a complete solution to address the above challenges. In particular, we first show that energy-harvesting, batteryless BLE beacons, are a cheap, reliable technology with respect to the operating cycle. In particular, we conducted an investigation on harvesting different types of energy sources, such as light, heat, and RF, and also considered the corresponding energy harvesting architecture. Later, we embedded them within a comprehensive, viable architectural proposal to support contact tracing, and, finally, we showed experimental results supporting our findings. A thorough discussion about the performance, efficiency, and security and privacy properties of our solution is also provided, while the paper concludes by highlighting future research directions.

Roadmap The remainder of this paper is structured as follows. Section II summarizes the related work on emerging contact tracing protocols and energy-harvesting technologies. Section III presents a threat model to evaluate different tracing protocols and suggest an ideal architecture for preserving privacy. Section IV introduces our proposed contact tracing

architecture, and Section V provides a viability study, performed with qualitative and quantitative evaluations, empirical field tests, simulations, and model analysis. Lastly, Section VI envisions several exciting research challenges and opportunities for the future, and Section VII concludes our work.

II. RELATED WORK

A. Digital Contact Tracing Solutions

Nowadays, several governments, research institutes, and companies are working on exposure notification protocols to limit the spread of infectious diseases, such as COVID-19. Contact tracing is defined as an identification process that aims to track the recent physical contacts of individuals that have been tested positive for the virus. Broadly speaking, existing BLE-based contact tracing protocols can be categorized as follows.

Decentralized Protocols. In a decentralized architecture, users do not share any data with the authorities unless they have a confirmed positive test. In that case, the patient's device uploads its own transmitted beacons to the authorities' server. These beacons are then propagated to the entire contact tracing network, where the individual smartphones perform the exposure notification function in a fully decentralized manner. (By matching the published beacons against their own contact logs.) Notable examples of decentralized contact tracing protocols are Apple/Google's framework [4] and the Decentralized Privacy-Preserving Proximity Tracing (DP-3T) protocol [5].

Hybrid Protocols. In a hybrid architecture, data collection follows the decentralized approach, i.e., each device maintains its private contact logs and does not disclose anything to the authorities. However, in hybrid protocols, the beacons transmitted by the mobile devices are generated by the health authorities. Then, in the event of a positive test, the user's device discloses its contact logs to the authorities, and, therefore, exposure notification is performed by the authorities in a centralized manner. Typical examples of hybrid solutions are BlueTrace [6]—first adopted by Singapore—and the Pan-European Privacy-Preserving Proximity Tracing (PEPP-PT) protocol [7].

IoT-based Protocols. IoT-based protocols employ an infrastructure of cheap IoT devices to facilitate contact tracing. In other words, smartphones no longer interact with each other but rather depend on IoT devices to detect proximity. To this end, IoTrace [8] is the only IoT-based solution to date. Under IoTrace, mobile devices are not required to scan the BLE channels for beacons broadcast by other devices. Instead, they simply broadcast their own beacons, which are received and logged by the IoT infrastructure. The reconciliation mechanism is fully tunable and could range from a decentralized to a centralized one. However, it is worth noting that reconciliation necessitates the transfer of a large number of beacons to/from the centralized server, using 4G/LTE communications. While the solution discussed in this paper falls under the IoT-based protocol umbrella, its core functionalities are very different from the ones provided by IoTrace.

Table I summarizes the characteristics of the most representative solutions—under different contact tracing

architectures—and shows how they compare against the proposed protocol. First, our solution is the only one that replaces part of the smartphones' energetic cost (stemming from beacon transmissions) with renewable energy. This is not possible with IoTrace, because the energy demand for the IoT devices' operations is very high and cannot be supported by energy-harvesting technologies. For the same reason, IoTrace has a significant maintenance/operation cost, due to the involvement of cellular communications and the need for frequent battery replacements.

In terms of privacy, hybrid protocols are the most vulnerable because the users' beacons are generated by the central authorities. As such, a malicious adversary that compromises the centralized server is able to track the movements of all users. On the other hand, decentralized solutions (and IoTrace) are more privacy-preserving because users construct their own beacons that are never revealed unless the user becomes infected with the virus. Nevertheless, the broadcasting of beacons from the mobile devices is, by itself, a privacy risk, as explained previously.

Finally, Table I also shows a quantitative comparison of the energy consumption for the entire contact tracing architecture. Let α and β be the daily RF transmission and receiving costs (including channel scanning), respectively. Also, let γ be the daily cost to communicate with the centralized server over an LTE network. Then, the table shows the total daily energy consumption for a network with n mobile devices and m IoT devices. We expect that $\alpha \ll \beta \ll \gamma$, and $n > m$.

B. Energy-Harvesting Technologies for IoT Applications

A beacon device can be configured with different advertising interval and transmit power values [9]. The advertising interval determines the temporal spacing of the beacons, while the transmit power controls its coverage area. A short advertising interval increases the beacon signal's reliability and enables more accurate distance estimation/localization. However, advertising intervals significantly influence the beacon's overall energy consumption and its lifetime.

In contact tracing applications, the energy demand for the devices is amplified due to various security and privacy requirements. For example, a static beacon may easily be spoofed or tracked so, cryptographically secure hashing algorithms are often implemented on the device's firmware to periodically randomize the broadcasted beacon [10]. However, such an operation leads to increased energy consumption and reduced lifetime.

To address these issues, we conducted an investigation on harvesting different types of energy sources, such as light, heat, and RF, and also considered the corresponding energy harvesting architecture. To this end, the *luxbeacon* is a BLE device that can harvest and store ambient light energy for energy-neutral operation [11]. It can operate in an indoor lighting environment with a minimum luminosity of 100 lux, and is composed of 5 major components, as shown in Fig. 1:

- 1) The solar panel harvests ambient light energy to power the load. The AM-1815 CA solar cell is optimized to harvest the visual light spectrum.

Table I
COMPARISON OF STATE-OF-THE-ART REPRESENTATIVE SOLUTIONS.

LOW: *, MEDIUM: **, HIGH: ***. A ✓ SYMBOL INDICATES THE FULFILLMENT OF A PARTICULAR FEATURE, A ✗ SYMBOL DENOTES THAT THE FEATURE IS EITHER NOT PROVIDED OR NOT APPLICABLE.

Features	Decentralized protocols [4], [5]	Hybrid protocols [6], [7]	IoTrace [8]	This work
Green Energy	✗	✗	✗	✓
Privacy	**	*	**	***
Total Energy Consumption	$n \cdot (\alpha + \beta)$	$n \cdot (\alpha + \beta)$	$n \cdot \alpha + m \cdot (\beta + \gamma)$	$n \cdot \beta + m \cdot \alpha$
Maintenance/Operation Cost	✗	✗	***	*

α : RF transmission cost, β : RF receiving cost, γ : LTE communication cost (with server), n : number of smartphones, m : number of IoT devices.




Figure 1. Circuit board and casing design of luXbeacon.

- 2) The power management IC routes the harvested energy from the solar panel to different parts of the circuit. The S6AE103A board leverages a linear harvesting architecture to achieve a low level of quiescent current (order of nA).
- 3) The primary buffer is a small energy storage unit that is charged first with the harvested energy. The energy in the primary buffer is used to boot-up the Bluetooth IC.
- 4) The supercapacitor is a large energy storage unit, where the harvested energy is stored during an energy surplus. The stored energy is used to offset any energy deficit in the future.
- 5) The Bluetooth IC is used to broadcast the BLE beacon to the surrounding devices.

III. THREAT MODEL

In a BLE-based contact tracing application, the main threat to privacy is an eavesdropping adversary that collects all the transmitted beacons. For instance, the adversary is equipped with either a Software Defined Radio (SDR) with a powerful antenna, or a Bluetooth-compliant transceiver connected to a laptop/smartphone. Thus, the adversary only needs to set the frequency adopted by the Bluetooth communication technology to intercept all BLE beacons in the surrounding area [12]. The attacker can also tag the beacons with timestamp and geo-location information computed by standard GPS or indoor localization methods. An eavesdropping attack aims mostly at compromising the users' privacy by either tracking their

movements or exposing their health status (with regards to the virus).

Alternatively, active adversaries may try to replay or relay previously transmitted beacons in order to disrupt the operation of the contact tracing network. For example, the adversary may try to cause a large number of false-positive exposure notifications. Finally, we assume that the adversary can only perform polynomial-time computations and is unable to break the cryptographic primitives adopted in the beacon generation functions.

IV. LUXBEACON CONTACT TRACING

The novelty of the proposed architecture lies in the deployment of a batteryless IoT infrastructure to facilitate privacy-preserving and energy-efficient proximity detection. In the following sections, we describe in detail the operations of the underlying contact tracing protocol.

A. System Architecture

The entities involved in the proposed architecture are the following:

luXbeacon. This is a BLE-based IoT device, equipped with specialized hardware for ambient-light energy harvesting. Every luXbeacon device broadcasts pseudo-random beacons to the surrounding mobile devices.

User. This is a smart device that runs the suggested contact tracing application. The app periodically scans the BLE spectrum for beacons transmitted by the deployed luXbeacon devices. Unlike existing approaches, the app operates in scan-only mode, i.e., it does not transmit any BLE beacons. During exposure notification, the smart devices approximate their relative proximity based on the received beacons from the IoT infrastructure.

Hospital. This is an authorized medical facility that performs COVID-19 infection tests. If a user tests positive, the health professionals are given permission to access his/her mobile device and forward the stored beacons to the central authority.

Authority. This is a trusted party whose role is to store the beacons that were recently collected from the infected users. In a real scenario, this role can be played by the *Ministry of Health*.

B. Protocol Message Flow

The protocol consists of two main tasks, namely, beacon collection and exposure notification. We assume that each

stored beacon at the user’s device contains a timestamp, the luXbeacon’s MAC address, the pseudo-random beacon (ephemeral ID), and the Received Signal Strength Indicator (RSSI). The high-level protocol message flow is as follows:

- 1) Every *luXbeacon* device periodically generates and transmits a pseudo-random BLE beacon, according to some cryptographic primitives, such as a keyed hash function.
- 2) Every *User* collects the beacon(s) transmitted in its surrounding area. Should the *User* test positive, the *User* will send all its stored beacons to the *Authority*.
- 3) Every *User* periodically downloads the up-to-date beacon list from the *Authority*, and checks (locally) if there are common elements between its stored beacons and the received list.
- 4) Finally, for all identified common beacons, the *User* will estimate its relative proximity to that patient, based on the signals’ RSSI.

The protocol message flow is also summarized in Fig. 2.

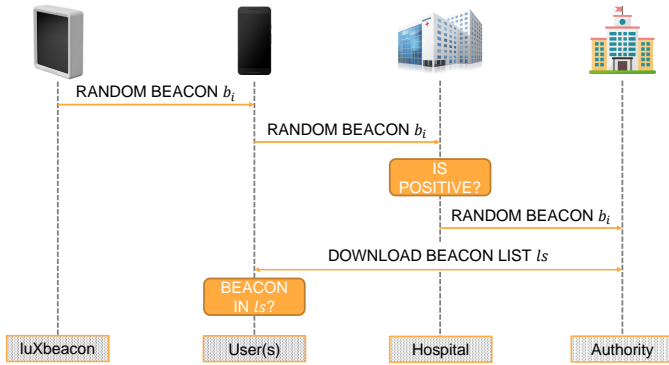


Figure 2. Message flow overview.

C. Contact Detection and Result Notification

In order to accurately detect a close contact between two users, it is critically important to estimate the following two parameters: (i) the distance between the two users; and (ii) the duration of the contact. The distance is essential because, if the two users were practising social distancing and separated by at least 2–3 *m*, the probability of contagion would be extremely low, and therefore, the contact would not be considered significant. Similarly, even if the two users were close enough for a contagion, but only for a period of less than a few minutes, the probability would also be very low. Therefore, the exposure notification function would consider these two variables when determining the threat level of a particular contact event.

It is worth noting that both variables can be trivially estimated by the proposed architecture. First, the distance between two users can be approximated by observing and comparing the RSSIs of their common beacons. However, the RSSI metric is subject to frequent fluctuations due to various environmental conditions, such as channel state, and fading and shadowing effects from the surrounding physical environment. Therefore, it is pivotal to deploy beacons at a




Figure 3. Supercapacitor voltage level of luXbeacon deployed in a real environment.

high density, in order to maximize the distance estimation accuracy. For example, if BLE devices are deployed with 4 *m* separation, the system would be able to detect contacts within 4 *m* with high precision. On the other hand, the duration of contact can be acquired by simply computing (from the available timestamps) the time interval that encloses a certain subset of common beacons.

V. VIABILITY STUDY

A. Sustainability

The following section investigates and evaluates the energy efficiency and sustainability of luXbeacon, loaded with the contact tracing firmware—also performing the needed cryptographic operations. To this end, we first measured the power consumption of the contact tracing firmware, which proved to consume 12.2 μA , with 1000 *ms* advertising interval and -8 *dBm* transmit power. In order to prove its sustainability and practicality, we deployed a luXbeacon in a real-life environment and monitored the changes in its supercapacitor voltage. The luXbeacon was deployed near a window, such as to harvest both solar and indoor light sources. The result is shown in Fig. 3, where the luXbeacon continuously charges and discharges its supercapacitor. It can also be observed that the supercapacitor voltage will never be lower than 2.7 *V*—the luXbeacon’s operating voltage being 1.8 *V*. Such observation further supports the self-sustainability of the luXbeacon for the contact tracing application.

To generalize our results, the lifetime of luXbeacon for various social locations was predicted using the lighting conditions of the locations. The predictions were made based on the measured energy consumption of the contact tracing firmware and also the power output of the solar panel. Fig. 4 shows 4 different possible locations for deployment, with varying lighting conditions and operation hours. It can be seen that in all social locations, luXbeacon is capable of



Figure 4. Expected lifetime of luXbeacon and lifetime extension compared to the battery powered devices under varying lighting conditions of social locations.

extending its lifetime by at least 170% with respect to battery-powered devices. Moreover, luXbeacon proved to be the most beneficial in outdoor deployment scenarios, which are the most difficult locations to conduct battery replacement or maintenance operations.

B. Contact Tracing Accuracy

BLE beacon infrastructure has been widely used for various indoor localization applications. With recent advancements in machine learning techniques, RSSI-based indoor localization and distance estimation have been shown to be reliable and accurate [13]. To validate our solution, we conducted an extensive simulation campaign using MATLAB©2020b, where we investigated how the random deployment of a varying number (from 1 to 10) of BLE beacon devices could be leveraged for optimal coverage area and positioning accuracy. Indeed, we were the first, to the best of our knowledge, to develop an end-to-end system that detects the contact between users based on beacon scanning information, namely RSSI and ephemeral ID. Our solution allows us to first estimate the distance of the users from the deployed BLE devices. From this information, our method then triangulates each user's position and estimates the distance between any two users with the accuracy reported in Fig. 5. The cited figure also reports the 95% confidence interval, computed over 10,000 tests, with a luxBeacon TX power of -8 dBm and a random deployment of two smartphones in an area of 100 m^2 . Let R_1 and R_2 be two generic receivers; the accuracy is estimated as the Maximum Absolute Difference (MAD) between the distance vectors of each receiver, i.e., d_1 and d_2 . The distances are computed by leveraging the relationship between RSSI values and distances collected from our experimental radio propagation model.

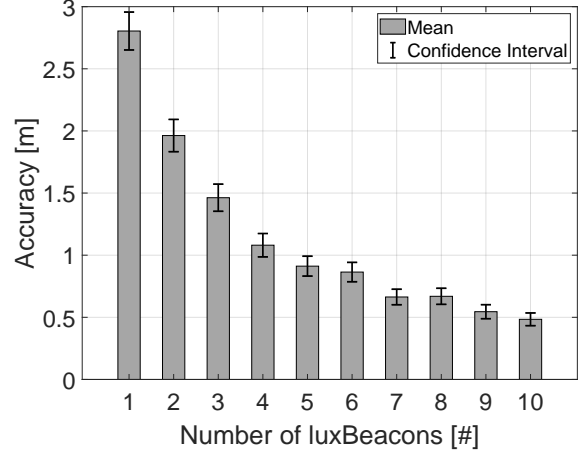


Figure 5. Estimated distance between two receivers — accuracy.

C. Ephemeral Beacons Generation

Each ephemeral ID (beacon) is generated with the SHA256 hashing function and the \oplus operation. The generated beacon starts with the first 19 bytes of device-specific information, such as device ID (18 bytes) and battery status information (1 byte). Then, the beacon contains a timestamp of 8 bytes. Further, we adopted the hashing function \mathcal{H} on the concatenated 27 bytes by providing an output of 32 bytes. Finally, in order to reduce the size of the hashed data, we split the 32 bytes of hashed data into two equal parts, and then we applied the \oplus operation iteratively in order to reduce the hashed data to just 4 bytes.

VI. CHALLENGES AND ROAD AHEAD

In the following sections, we describe the research challenges from the perspectives of security and privacy, infrastructure maintenance, and localization accuracy. We also outline the limitations of our proposed solution.

A. Infrastructure and Maintenance Costs

The proposed architecture requires a large deployment of cheap IoT devices. Being the devices cheap, especially when produced at scale, the main cost of the infrastructure will be given by its maintenance. To this end, the adoption of energy-harvesting technologies, such as luXbeacon, reduces the maintenance cost significantly if deployed in an environment with sufficient light. However, in those environments that may not have enough light to enable energy-neutral operation, the energy consumption rate may vary, and so will the battery life. Such a phenomenon would lead to asynchronous expiry of battery lifetime, which may cause additional complications and difficulties in managing the infrastructure. A cost-benefit analysis for the proposed architecture should evaluate: (i) the efficiency of this new architecture in terms of resources; (ii) its effect on social well-being; and (iii) how social costs and benefits can be monetized. The luXbeacon IoT device

has a cost of ~ 30.00 USD per unit, including the casing and hardware—its cost will be a fraction of it when mass-produced—, it is relevant to analyze the best deployment plan to cover the most crowded areas. Further, it is worth noticing that, comparing our solution to other BLE-based approaches from the maintenance and application reliability perspective, the one/time cost to build the entire infrastructure can be considered extremely low.

B. Tracing Performance

BLE beacon infrastructures have been widely used for various indoor localization applications. Many investigations have been performed on techniques that could enhance the positioning accuracy of a user in an environment with densely deployed IoT devices. However, very few studies exist concerning the energy consumption of a BLE device. Since the luXbeacon’s broadcasting frequency is limited by the availability of harvestable ambient energy, the contact tracing accuracy may be affected by the scarce energy resources and the deployment environment. It would be imperative to study the relationship between luXbeacon’s operational configurations—namely advertising interval and transmit power—with accuracy. Furthermore, the deployment method of the luXbeacon infrastructure may further be explored for optimal coverage area and positioning accuracy. Additionally, a method to accurately detect significant contacts between users must be investigated (e.g., user mobility). As future work, an evaluation of the luXbeacon’s transmission frequency and transmit power (i.e., the coverage area) correlated to the density of a particular zone is needed to achieve better performance in terms of energy consumption, communication efficiency, and hardware sustainability. This analysis allows for an implementation of a self-adaptive solution that permits tuning the transmission frequency (i.e., the delay between two consecutive data transmissions), taking into account the area density as well as the beacon key update frequency.

C. Security & Privacy

From a privacy perspective, the architecture follows the privacy-by-design approach. Indeed, off-loading the beacon broadcast operation to the fixed hardware infrastructure avoids the “pebble dropping” issue for users since their mobile devices are not transmitting any information. Therefore, this approach makes it infeasible for an eavesdropping adversary to track users. However, if a user has a positive COVID-19 test, the authorities have to publish his/her stored beacons to a public database for the purpose of exposure notification. As such, the user’s recent location history is disclosed to the entire network. To this end, it is important to consider cryptographic techniques in the exposure notification function. In particular, instead of publishing the user’s beacons, the server could engage in a two-party private-set intersection protocol [14] with individual users. The protocol’s output would reveal (to the user) the common beacon set, but nothing else. It is also imperative to perform an experimental study in order to assess the effectiveness and computational cost of exposure notification in this privacy-preserving setting.

Further, compared to IoTrace, our solution provides better security for data at rest because no user information is stored on the luXbeacon devices. However, a critical security challenge is to find and analyze the right countermeasures to mitigate replay attacks. Specifically, a malicious adversary may deploy rogue luXbeacon devices in order to manipulate the protocol’s proximity detection module. To this end, we should investigate the feasibility of detecting counterfeit beacons at the centralized server by analyzing the beacons submitted by a new patient. The analysis would consider the timing information, the beacons’ ephemeral IDs (which are generated based on secret luXbeacon IDs), and the locations of the luXbeacon devices that are known to the authorities.

D. Discussion

While BLE is a low-energy system compared to traditional Bluetooth, scanning is still a reasonably power-intensive operation. Continuous scanning would negatively affect the battery’s life, and therefore, degrade the user’s experience or even force them to uninstall the contact tracing app. The energy consumption of the Bluetooth scanning operation depends on many factors, such as the Bluetooth SoC, the hardware design, the scanning parameters, and the number of scannable Bluetooth devices in the vicinity. Based on the nRF51822 SoC, an active and continuous scanning operation consumes 40 mW, whereas the broadcasting operation consumes at most 600 μ W. As reported in [15], the power consumption of Bluetooth scanning is similar to that of Wi-Fi during web browsing. We should note that the scanning operation is duty-cycled at the OS level in order to reduce excessive power consumption. Therefore, it is of paramount importance to design the mobile app by taking into consideration the OS-related operations. Additionally, while a longer beacon broadcast cycle favours sustainability, it negatively impacts the proximity detection accuracy. There is a need to balance this trade-off, while also maintaining a low luXbeacon TX power.

VII. CONCLUSION

Digital contact tracing can play a vital role in limiting the spread of deadly viruses. However, its effectiveness is dependent upon its adoption by a large majority of the general public. To this end, privacy and energy-efficiency are two important metrics that can motivate users to participate in the contact tracing network. Our work makes a significant contribution towards this goal, by proposing an energy-efficient and privacy-preserving architecture for contact tracing. The proposed solution leverages a dense deployment of batteryless IoT devices that constantly broadcast BLE beacons for the purpose of proximity detection. We have shown that batteryless IoT has a reliable operating cycle and proved that their deployment can help improve detection accuracy. The proposed architectural design enjoys low maintenance cost, reduces energy consumption on the user side, greatly improves distance accuracy estimation, and provides privacy by design. Finally, we have summarized the most important research challenges and directions that need to be addressed by the academia and industry, towards the development of IoT based privacy-preserving and efficient contact tracing.

ACKNOWLEDGMENTS

This publication was partially supported by awards NPRP 11S-0109-180242 from the QNRF-Qatar National Research Fund, a member of The Qatar Foundation. The information and views set out in this publication are those of the authors and do not necessarily reflect the official opinion of the QNRF.

REFERENCES

- [1] N. Ahmed *et al.*, “A Survey of COVID-19 Contact Tracing Apps,” *IEEE Access*, vol. 8, pp. 134 577–134 601, 2020.
- [2] H. Cho *et al.*, “Contact Tracing Mobile Apps for COVID-19: Privacy Considerations and Related Trade-offs,” 2020.
- [3] R. Sun *et al.*, “Vetting Security and Privacy of Global COVID-19 Contact Tracing Applications,” 2020.
- [4] Apple Google. (2020) Privacy-Preserving Contact Tracing. (Accessed: 2021-03-07). [Online]. Available: <https://www.apple.com/covid19/contacttracing>
- [5] “Decentralized Privacy-Preserving Proximity Tracing: Overview of Data Protection and Security,” <https://github.com/DP-3T/documents/blob/master/DP3TWhitePaper.pdf>, 2020, (Accessed: 2021-03-07).
- [6] J. Bay *et al.*, “BlueTrace: A privacy-preserving protocol for community-driven contact tracing across borders,” *Government Technology Agency-Singapore, Tech. Rep.*, 2020.
- [7] PEP-PT Team. (2020) Pan-European Privacy-Preserving Proximity Tracing. (Accessed: 2021-03-07). [Online]. Available: <https://www.pepp-pt.org/>
- [8] P. Tedeschi *et al.*, “IoTrace: A Flexible, Efficient, and Privacy-Preserving IoT-enabled Architecture for Contact Tracing,” *IEEE Communications Magazine*, 2021.
- [9] K. E. Jeon *et al.*, “BLE Beacons for Internet of Things Applications: Survey, Challenges, and Opportunities,” *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 811–828, 2018.
- [10] A. Zidek *et al.*, “Bellrock: Anonymous Proximity Beacons From Personal Devices,” in *2018 IEEE International Conf. on Pervasive Computing and Communications (PerCom)*, 2018, pp. 1–10.
- [11] K. E. Jeon *et al.*, “luXbeacon—A Batteryless Beacon for Green IoT: Design, Modeling, and Field Tests,” *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5001–5012, 2019.
- [12] A. B. Dar *et al.*, “Applicability of mobile contact tracing in fighting pandemic (COVID-19): Issues, challenges and solutions,” *Computer Science Review*, vol. 38, p. 100307, 2020.
- [13] C. H. Lam *et al.*, “Improved Distance Estimation with BLE Beacon Using Kalman Filter and SVM,” in *2018 IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.
- [14] E. De Cristofaro *et al.*, “Practical Private Set Intersection Protocols with Linear Complexity,” in *Financial Cryptography and Data Security*, R. Sion, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 143–159.
- [15] A. Carroll *et al.*, “An Analysis of Power Consumption in a Smartphone,” in *Proceedings of the 2010 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIXATC’10. USA: USENIX Association, 2010, p. 21.

BIOGRAPHIES

Pietro Tedeschi is PhD Student at HBKU-CSE-ICT, Doha-Qatar. He received his Master’s degree with honors in Computer Engineering at Politecnico di Bari, Italy. He worked as Security Researcher at CNIT, Italy, for the EU H2020 SymbIoTe. His research interests cover security issues in UAVs, Wireless, IoT, and Cyber-Physical Systems.

Kang Eun Jeon is currently working toward his Ph.D. degree in the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology (HKUST) since 2015. Before joining the HKUST-NIE Social Media Lab, he received his B.Eng. degree in Electronic Engineering also in HKUST. His research interests include self-sustaining, secure, and social BLE beacons for the Internet of Things applications.

James She is the founding director of HKUST Social Media Lab., and affiliated with the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology. His current research areas include Social and Multimedia Computing, Data Science and AI for Visual Creativity, IoT for Sustainable, Smart and Interactive Systems. James is also the associate editor for *IEEE TRANSACTIONS ON MULTIMEDIA*, and *ACM Transactions on Multimedia Computing, Communications and Applications*.

Simon Wong is currently working toward his M.Phil. degree in the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology (HKUST) since 2019. His research interests include technologies on machine learning and signal processing for IoT devices.

Spiridon Bakiras is associate professor of cybersecurity at HBKU-CSE, Doha-Qatar. His research interests include Security and Privacy, Applied Cryptography, and Spatiotemporal Databases. He held teaching and research positions at Michigan Technological University, the City University of New York, the University of Hong Kong, and the Hong Kong University of Science and Technology. He is a recipient of the U.S. National Science Foundation CAREER award.

Roberto Di Pietro, ACM Distinguished Scientist, is Full Professor in Cybersecurity at HBKU-CSE. His main research interests include security and privacy for distributed systems (e.g., Blockchain, Cloud, IoT, OSNs), virtualization security, and applied cryptography. Other than being involved in M&A of start-up, he has been producing 230+ scientific papers and 20+ patent applications over the cited topics, has co-authored three books, and contributed to a few others. In 2011-2012 he was awarded a Chair of Excellence from University Carlos III, Madrid. In 2020 he received the Jean-Claude Laprie Award for having significantly influenced the theory and practice of Dependable Computing.