

# Privacy-Preserving and Sustainable Contact Tracing Using Batteryless Bluetooth Low-Energy Beacons

**Pietro Tedeschi** | Hamad Bin Khalifa University

**Kang Eun Jeon** | The Hong Kong University of Science and Technology

**James She** | Hamad Bin Khalifa University

**Simon Wong** | The Hong Kong University of Science and Technology

**Spiridon Bakiras and Roberto Di Pietro** | Hamad Bin Khalifa University

**Contact tracing is the technology choice of reference to address the COVID-19 pandemic. Many of the current approaches have severe privacy and security issues and fail to offer a sustainable contact tracing infrastructure. We address these issues introducing an innovative, privacy-preserving, sustainable, and experimentally tested architecture that leverages batteryless Bluetooth low-energy beacons.**

Smartphone-based contact-tracing protocols<sup>1</sup> have been adopted by many countries to help fight the spread of COVID-19. Most practical implementations today follow a common modus operandi: mobile devices continuously broadcast pseudorandom Bluetooth low-energy (BLE) packets 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, to identify potential contagion events. The main challenge of contact-tracing solutions is related to location tracking. Indeed, to signal the presence of a smart device, the devised solutions constantly spread pseudorandom packets via Bluetooth. Furthermore, every device maintains its own contact list by storing the signals broadcast by other devices. 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.

Even though the user's packets are pseudorandom 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 Apple/Google privacy bug found in their COVID-19 exposure notification framework.<sup>2</sup> The highlighted native privacy and energy concerns in the existing solutions undermine the very purpose of contact-tracing applications, hindering their adoption by the general public.<sup>3</sup>

To mitigate the aforementioned privacy and energy issues, we propose the deployment of a lightweight and wide-scale contact-tracing infrastructure consisting of BLE transmitters. The packets transmitted by these devices would replace the smartphone-generated packets but 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 packets, thus gradually building a record of their precise location over time. Then, the exposure notification process would develop as in most standard decentralized BLE-based protocols. The benefits of our architecture are threefold: 1) unconditional privacy for users since their devices are not emitting any information; 2) reduced energy requirements for smartphones, which translate into a longer battery life; and 3) the potential for more accurate proximity

*Digital Object Identifier 10.1109/MSEC.2021.3115497*  
*Date of current version: 19 October 2021*

detection because of the presence of multiple (fixed) BLE transmitters.

To facilitate an easy and wide-scale deployment, BLE beacons are typically battery powered (similar to sensor network deployments). This latter point would trigger the issue of periodic battery replacement, which, in turn, would considerably increase the operational and maintenance costs of the infrastructure.

Such overhead is further amplified in large-scale deployment cases. As an example, the Hong Kong International Airport had to deploy more than 17,000 beacons to provide indoor navigation services. By using BLE beacons as a contact-tracing infrastructure, energy consumption on user smartphones for broadcasting pseudorandom packets is offloaded to the beacon infrastructure. More importantly, smartphones are not transmitting any information, so the users' privacy is unconditionally preserved—something hardly possible with existing device-to-device contact-tracing protocols. Further, an infrastructure-based network can still support distributed protocols/frameworks of various contact-tracing solutions.

## Contributions

We first show that energy-harvesting, batteryless beacons are an affordable, reliable technology with respect to the operating cycle. We conducted an investigation on harvesting different types of energy sources, such as those from light, heat, and radio frequency (RF), and we also considered the corresponding energy-harvesting architecture. Later, we embedded them within a comprehensive, viable architectural proposal to support contact tracing; finally, we showed experimental results supporting our findings. We also shed light on the tradeoff between the broadcast frequency and transmit power that could affect the contact-tracing performance and batteryless beacon sustainability for target tracing accuracy. We also provide a thorough discussion on the performance, efficiency, and security and privacy properties of our solution and conclude by highlighting future research directions.

## Related Work

In the following, we summarize the related work in the field, focusing on contact-tracing approaches and energy-harvesting technologies. We have adopted the following terms throughout the article:

- *BLE packet*: a broadcast packet sent using the BLE protocol
- *BLE beacon*: a piece of specialized hardware (not necessarily a smart device) that simply broadcasts BLE packets
- *luXbeacon*: a BLE beacon with energy-harvesting capabilities to promote a self-sustainable operation.

## Digital Contact-Tracing Solutions

Several governments, research institutes, and companies are now 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 who have 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 claimed positive device uploads its own transmitted BLE packets to the authorities' server. These packets 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 data against their own contact logs). Notable examples of decentralized contact-tracing protocols are Apple/Google's framework<sup>4</sup> and the decentralized privacy-preserving proximity tracing protocol.<sup>5</sup>

**Hybrid protocols.** In a hybrid architecture, data collection follows a decentralized approach; that is, each device maintains its private contact logs and does not disclose anything to the authorities. However, in hybrid protocols, the packets 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; thus, exposure notification is performed by the authorities in a centralized manner. Typical examples of hybrid solutions are BlueTrace<sup>6</sup>—first adopted by Singapore—and the pan-European privacy-preserving proximity tracing protocol.<sup>7</sup>

**Internet of Things-based protocols.** Internet of Things (IoT)-based protocols employ an infrastructure of 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. IoTrace<sup>8</sup> is the only IoT-based solution to date. Under IoTrace, mobile devices are not required to scan the BLE channels for broadcast packets sent by other devices. Instead, they simply broadcast their own packets, 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 packets to/from the centralized server, using 4G/LTE communications. While the architecture introduced in this article falls under the umbrella of IoT-based

solutions, its core functionalities are very different from the ones provided by IoTrace.

### Energy-Harvesting Technologies for IoT Applications

A BLE beacon can be configured with different advertising interval and transmit power values.<sup>9</sup> The advertising interval determines the temporal spacing for the broadcast of beacon packets, 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 because of various security and privacy requirements. For example, a static beacon may easily be spoofed or tracked; therefore, cryptographically secure hashing algorithms are often implemented on the device's firmware to periodically randomize the broadcast of beacon packets.<sup>10</sup> However, such an operation also leads to an increased energy consumption and a reduced lifetime.

To address these issues, we conducted an investigation on harvesting different types of energy sources (such as light, heat, and RF), and we also considered the corresponding energy-harvesting architecture. *luXbeacon* is a BLE beacon that can harvest and store ambient

light energy for energy-neutral operation.<sup>11</sup> It can operate in an indoor lighting environment with a minimum luminosity of 100 lux and is composed of six major components, as shown in Figure 1:

1. The solar panel harvests ambient light energy to power the load. The AM-1815CA solar cell is optimized to harvest the visual light spectrum.
2. A power management integrated circuit (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 (on the order of nanoamperes).
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 future energy deficit.
5. The Bluetooth IC is used to broadcast the BLE beacon to the surrounding devices.
6. A backup battery is used to power the *luXbeacon* when there is not enough ambient light energy to harvest and sustain its operation.

### The Threat Model

In a BLE-based contact-tracing application, the main threat to privacy is an eavesdropping adversary that collects all of the transmitted packets. For instance, the adversary may be equipped with either a software-defined radio 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 packets in the surrounding area. The attacker can also tag the packets with a time stamp and geolocation 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 regard to the virus).

Alternatively, active adversaries may try to replay or relay previously transmitted packets to disrupt the operation of the contact-tracing network. For example, an 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 standard cryptographic primitives adopted in the pseudorandom packet generation functions.




Figure 1. The circuit board and casing design of *luXbeacon*.

## 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.

### The System Architecture

The following entities are involved in the proposed architecture:

**luXbeacon:** This BLE-based IoT device is equipped with specialized hardware for ambient-light energy harvesting. Every luXbeacon device broadcasts pseudorandom packets to the surrounding mobile devices.

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

**Hospital:** The hospital is the authorized medical facility that performs COVID-19 infection testing. If a user tests positive, the health professionals are given permission to access their mobile device and forward the stored packets to the central authority.

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

### The Protocol Message Flow

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

each stored BLE packet at the user's device contains a time stamp, the luXbeacon's media access control address, a pseudorandom value (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 pseudorandom BLE packet according to a secure keyed hash function.
2. Every *User* collects the packet(s) transmitted in its surrounding area. Should the *User* test positive, the *User* will send all its stored packets to the *Authority*.
3. Every *User* periodically downloads the up-to-date packet list from the *Authority* and checks (locally) if there are common elements between its stored packets and the received list.
4. Finally, for all identified common packets, the *User* will estimate its relative proximity to that claimed positive, based on the signal's RSSI.

The protocol message flow is also summarized in Figure 2. Note that our architecture follows the decentralized exposure notification approach (steps 3 and 4), where each device locally determines whether the user was in close contact with a claimed positive. Assuming that the authorities will always learn the infected person's BLE packets, a further privacy goal is to not disclose these packets to everyone else. Our discussion in the section "Security and Privacy" proposes such an approach that leverages well-known cryptographic protocols.

### Contact Detection and Result Notification

To accurately detect a close contact between two users, it is critically important to estimate the following two parameters: 1) the distance between the two users and

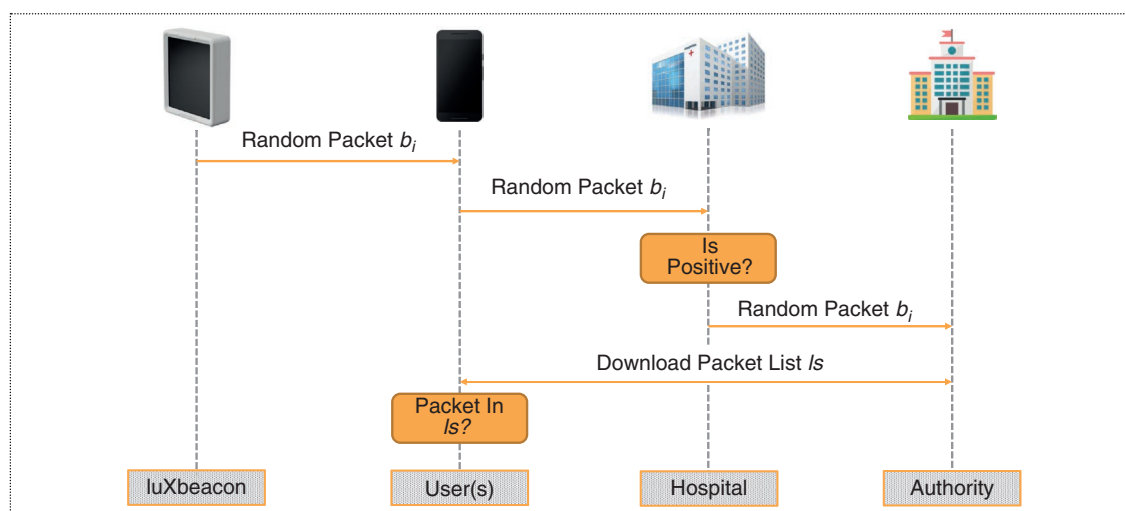


Figure 2. A message flow overview.

2) the duration of the contact. The distance is essential because, if the two users were practicing social distancing and separated by at least 2–3 m, the probability of contagion would be extremely low; 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.

Note 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 packets. However, the RSSI metric is subject to frequent fluctuations caused by various environmental conditions, such as channel state, and fading and shadowing effects from the surrounding physical environment. Therefore, it is pivotal to deploy BLE beacons at a high density to improve the distance estimation accuracy through better triangulation. On the other hand, the duration of contact can be acquired by simply computing (from the available time stamps) the time interval that encloses a certain subset of common packets.

### A Comparison With Related Protocols

Table 1 summarizes the characteristics of the most representative solutions reviewed in the section “Related Work”—under different contact-tracing architectures—and shows how they compare against the proposed protocol. First, our proposed architecture is the only one that replaces part of the smartphones’ energetic cost (stemming from beacon transmissions) with renewable energy, hence having a low impact on the maintenance cost. This is not possible with IoTrace because its energy demand for the IoT devices’ operations is very high and cannot be supported by

energy-harvesting technologies.<sup>12</sup> For the same reason, IoTrace has a high maintenance/operation cost because of 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’ ephemeral IDs are generated by the central authorities, which results in low privacy guarantees. For example, 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 ephemeral IDs that are never revealed unless the user becomes infected with the virus. As such, they guarantee a medium level of privacy. Nevertheless, the broadcasting of packets from the mobile devices is, by itself, a privacy risk, as explained previously. By contrast, our proposed architecture can guarantee a high level of privacy, as further discussed in the section “Security and Privacy.”

Finally, Table 1 also shows a quantitative comparison of the energy consumption for the entire contact-tracing architecture. Let  $\alpha$  and  $\beta$  be the daily RF transmission and receiving costs (including channel scanning), respectively. Also, let  $\gamma$  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  $n > m$ .

### Viability Study

In this section, we revise the different requirements that assure the viability of a contact-tracing protocol, showing that our approach satisfies them all.

### Sustainability

The following section investigates and evaluates the energy efficiency and sustainability of luXbeacon, loaded with the contact-tracing firmware and also

**Table 1. A comparison of state-of-the-art representative solutions. A ✓ symbol indicates the fulfillment of a particular feature; a X symbol denotes that the feature is either not provided or not applicable.**

Features	Decentralized protocols <sup>4,5</sup>	Hybrid protocols <sup>6,7</sup>	IoTrace <sup>8</sup>	This work
Green energy	X	X	X	✓
Privacy guarantee	Medium	Low	Medium	High
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	X	X	High	Low

$\alpha$ : RF transmission cost;  $\beta$ : RF receiving cost;  $\gamma$ : LTE communication cost (with server);  $n$ : number of smartphones;  $m$ : number of IoT devices.

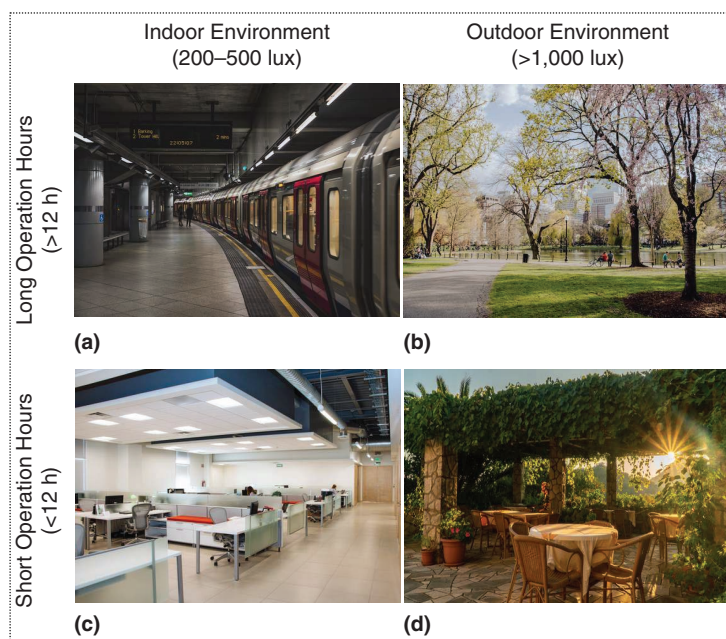
performing the needed cryptographic operations. We first measured the power consumption of the contact-tracing firmware, which proved to consume 12.2  $\mu\text{A}$ , with a 100-ms advertising interval and  $-8$  dBm transmit power. To prove its sustainability and practicality, we deployed a luXbeacon in a real-life environment and monitored the changes in its supercapacitor voltage. The luXbeacon was deployed near a window to harvest both solar and indoor light sources, which can provide sufficient ambient power to support the luXbeacon. The result is shown in Figure 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. This observation further supports the self-sustainability of the luXbeacon in a 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. Figure 4 shows four different possible locations for deployment, with varying lighting conditions and operation hours. It can be seen that, in all social locations, luXbeacon has an extended battery lifetime of at least 70% compared to that of traditional battery-powered BLE beacons. Moreover, luXbeacon proved to be the most beneficial in outdoor deployment scenarios, which are the most difficult locations in which to conduct battery replacement or maintenance operations.


### Contact-Tracing Accuracy

To measure and reference the radio characteristics of luXbeacon, an experiment was conducted to investigate


its RSS over varying values of transmission power and receiving distance. The luXbeacon’s RSS was measured for 5 min for each distance ranging from 0 to 6 m in 1-m intervals. In Figure 5, it can be observed that each transmit power curve is vertically separated from its neighbor by 5–7 dBm; however, the overall trend of the plot shows evident similarities. It is also noteworthy that the change in RSS is dramatic between 0 and 1 m but less noticeable after a 1-m distance.



**Figure 4.** The expected lifetime of luXbeacon and lifetime extension compared to the battery-powered devices (which last 2.3 years), under varying lighting conditions of social locations. (a) subway: 4 years (x170%), (b) park: 6.5 years (x280%), (c) office: 4 years (x170%), and (d) outdoor restaurant: 6.5 years (x280%). (Source: <https://unsplash.com/>.)



**Figure 3.** The supercapacitor voltage level of luXbeacon deployed in a real environment.



**Figure 5.** The RSS of luXbeacon measured over distance.

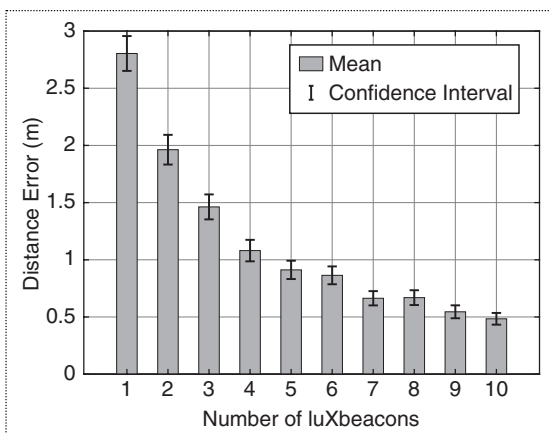
To validate our architecture, we also conducted an extensive simulation campaign using MATLAB 2020b, where we investigated how the random deployment of a varying number (from 1 to 10) of BLE beacons could be leveraged for optimal coverage area and positioning accuracy. We present, for the first time in the literature, an end-to-end system that detects the contact between users based on BLE packet scanning information, namely the RSSI and ephemeral ID.

Our architecture 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 distance error reported in Figure 6. The higher the number of deployed luXbeacons, the lower the distance estimation error between two generic devices. Figure 6 also reports the 95% confidence interval, computed over 10,000 tests, with a luXbeacon transmit power of  $-8$  dBm and a random deployment of two smartphones in an area of  $100 \text{ m}^2$ . Let  $R_1$  and  $R_2$  be two generic receivers; the distance between them,  $\hat{d}$ , is estimated as the maximum absolute difference between arrays  $\mathbf{d}_{R_1}$  and  $\mathbf{d}_{R_2}$ , where each array consists of the estimated distances to each one of the surrounding luXbeacons, as shown in (1). The distances inside the two arrays are estimated by leveraging the relationship between the RSSI and distance (Figure 5) collected from our experimental radio propagation model. Essentially, this approach is an attempt to estimate the distance between two users without knowing the precise locations of the surrounding luXbeacons.

$$\hat{d} = \max(|\mathbf{d}_{R_1} - \mathbf{d}_{R_2}|). \quad (1)$$

### Ephemeral ID Generation

Each ephemeral ID is generated with the SHA-256 hashing function and the XOR ( $\oplus$ ) operation. The generated



**Figure 6.** The error in distance estimation between two receivers.

packet starts with the first 19 B of device-specific information, such as the device ID (18 B) and battery status information (1 B). Then, the packet contains a time stamp of 8 B. Further, we adopted the hashing function  $\mathcal{H}$  on the concatenated 27 B by providing an output of 32 B. Finally, to reduce the size of the hashed data, we split the 32 B of hashed data into two equal parts and then applied the  $\oplus$  operation iteratively to reduce the hashed data to just 4 B.

### Challenges and the 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.

### Infrastructure and Maintenance Costs

The proposed architecture requires a large deployment of IoT devices that are quite affordable when produced at mass scale. As such, the main cost of the infrastructure will be determined 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 because of non-uniform lighting conditions, and so will the battery life. Such a phenomenon would lead to an asynchronous expiry of the battery lifetime, which may cause additional complications and difficulties in managing the infrastructure. To address such issues, it would be ideal to investigate and design energy-aware firmware that is capable of load balancing to match its battery usage rate with that of nearby beacons.

A cost-benefit analysis for the proposed architecture should evaluate 1) the efficiency of this new architecture in terms of resources, 2) its effect on social well-being, and 3) how social costs and benefits can be monetized. When the luXbeacon is mass produced, its unit cost will be  $\approx$  US\$15, including the casing and hardware, which is comparable to off-the-shelf battery-powered beacon devices that cost  $\approx$  US\$30.<sup>13</sup> It is also relevant to analyze the best deployment plan to cover the most crowded areas. Further, comparing our architecture to other BLE-based approaches from the maintenance and application reliability perspective, the one-time cost to build the entire infrastructure can be considered relatively low.

### 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 beacon. Since the luXbeacon's broadcasting frequency (the advertising interval) is limited by the availability of harvestable ambient energy, the contact-tracing accuracy may be affected by scarce energy resources and the deployment environment. It would be imperative to study the relationship between luXbeacon's operational configurations—namely the 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 (such as user mobility). As future work, an evaluation of the luXbeacon's advertising interval and transmit power (that is, 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 of the advertising interval, taking into account the area density as well as the beacon key update frequency.

### Security and Privacy

From a privacy perspective, the architecture follows the privacy-by-design approach. Indeed, offloading the packet broadcast operation to the fixed hardware infrastructure avoids the “data leakage” issue for users because their mobile devices are not transmitting any information. Therefore, our architecture 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 their stored BLE packets 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 packets, the server could engage in a two-party private-set intersection protocol<sup>14</sup> with individual users. The protocol's output would reveal (to the user) the common ephemeral ID set but nothing else. It is also imperative to perform an experimental study to assess the effectiveness and computational cost of exposure notification in this privacy-preserving setting.

Further, compared to IoTrace, our architecture provides better security for data at rest because no user information is stored on the luXbeacon(s). 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 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 packets submitted by a new claimed positive. 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(s) that are known to the authorities.

### Discussion

Besides the privacy-preserving benefit of the proposed architecture, the energy consumption of user smartphones is also reduced compared to existing methods. The reason is that they only need to carry out Bluetooth scanning operations instead of both scanning and broadcasting. While BLE is a low-energy system compared to traditional Bluetooth, it involves a reasonably power-intensive operation. Continuous scanning would negatively affect the smartphone's battery life and, therefore, degrade the users' experience or even force them to uninstall the contact-tracing app. Therefore, it is extremely important to consider the energy consumption at the end devices when developing the contact-tracing system.

The energy consumption of the Bluetooth scanning operation depends on many factors, such as the Bluetooth system on chip (SoC), the hardware design, scanning parameters, and 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  $\mu$ W. As reported in Carroll et al.,<sup>15</sup> the power consumption of Bluetooth scanning is similar to that of Wi-Fi during web browsing. To address such issues, the latest smartphone hardware and operating systems have implemented several mechanisms to minimize energy consumption. We should note that the scanning operation is duty cycled at the operating system level to reduce excessive power consumption. Therefore, it would be important to design a mobile app considering the energy consumption through an intelligent framework that requires minimum scanning operation. Additionally, while a longer BLE packet broadcast cycle favors sustainability, it negatively impacts the proximity detection accuracy. There is a need to balance this tradeoff, while also maintaining a low luXbeacon transmit power. Today, BLE is the de facto standard for the most prominent contact-tracing solutions in the literature. Further research on various communication technologies is needed; for example, ultrawideband carriers may be used to increase the proximity tracing accuracy, privacy, and reliability.



The human and economic impacts of the COVID-19 pandemic have shown the need for novel technological solutions to tackle similar events in the future. Digital contact tracing can play a vital role in limiting the spread of deadly viruses. However, its effectiveness depends on 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 toward this goal by proposing an energy-efficient and privacy-preserving architecture for contact tracing. The proposed architecture leverages a dense deployment of batteryless IoT devices that constantly broadcast BLE packets for the purpose of proximity detection. We have shown that batteryless IoT devices have a reliable operating cycle and proved that their deployment can help improve detection accuracy. The proposed architectural design has a 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 academia and industry toward the development of IoT-based privacy-preserving and efficient contact tracing. ■

### Acknowledgments

The authors would like to thank the anonymous reviewers, who helped improve the quality of the article. This publication was partially supported by awards NPRP-S-11-0109-180242 from the Qatar National Research Fund, a member of the Qatar Foundation, and the North Atlantic Treaty Organization Science for Peace and Security Program MYP G5828 project “SeaSec: DronNets for Maritime Border and Port Security.”

### References

1. N. Ahmed et al., “A survey of COVID-19 contact tracing apps,” *IEEE Access*, vol. 8, pp. 134,577–134,601, July 2020. doi: 10.1109/ACCESS.2020.3010226.
2. A. Hayward, “Privacy bug found in Apple, Google COVID-tracing framework,” Decrypt, 2020. [Online]. Available: <https://decrypt.co/40765/privacy-bug-found-apple-google-covid-tracing-framework> (accessed Aug. 30, 2021).
3. H. Cho, D. Ippolito, and Y. W. Yu, “Contact tracing mobile apps for COVID-19: Privacy considerations and related trade-offs,” 2020, arXiv:2003.11511.
4. “Privacy-preserving contact tracing,” Apple Google, 2020. Accessed: Mar. 7, 2021. [Online]. Available: <https://www.apple.com/covid19/contacttracing>
5. “Decentralized privacy-preserving proximity tracing: Overview of data protection and security,” GitHub, San Francisco, 2020. [Online]. Available: <https://github.com/DP-3T/documents/blob/master/DP3TWhitePaper.pdf> (accessed Mar. 7, 2021).
6. J. Bay et al., “BlueTrace: A privacy-preserving protocol for community-driven contact tracing across borders,” Government Technol. Agency, Singapore, Tech. Rep., 2020. [Online]. Available: <http://tiny.cc/qo1kuz>
7. “Pan-European Privacy-Preserving Proximity Tracing,” Wix.com, 2020. Accessed: Mar. 7, 2021. [Online]. Available: <http://github.com/pepp-pt>
8. P. Tedeschi, S. Bakiras, and R. Di Pietro, “IoTrace: A flexible, efficient, and privacy-preserving IoT-enabled architecture for contact tracing,” *IEEE Commun. Mag.*, vol. 59, no. 6, 2021, pp. 82–88. doi: 10.1109/MCOM.001.2000729.
9. K. E. Jeon, J. She, P. Soonsawad, and P. C. Ng, “BLE beacons for internet of things applications: Survey, challenges, and opportunities,” *IEEE Internet Things J.*, vol. 5, no. 2, pp. 811–828, 2018. doi: 10.1109/JIOT.2017.2788449.
10. A. Zidek, S. Tailor, and R. Harle, “Bellrock: Anonymous proximity beacons from personal devices,” in *Proc. 2018 IEEE Int. Conf. Pervasive Comput. Commun. (PerCom)*, pp. 1–10. doi: 10.1109/PERCOM.2018.8444603.
11. K. E. Jeon, J. She, J. Xue, S. Kim, and S. Park, “luXbeacon—A batteryless beacon for green IoT: Design, modeling, and field tests,” *IEEE Internet Things J.*, vol. 6, no. 3, pp. 5001–5012, 2019. doi: 10.1109/JIOT.2019.2894798.
12. P. Tedeschi, S. Sciancalepore, and R. Di Pietro, “Security in energy harvesting networks: A survey of current solutions and research challenges,” *IEEE Commun. Surveys Tuts.*, vol. 22, no. 4, pp. 2658–2693, 2020. doi: 10.1109/COMST.2020.3017665.
13. “Smart Beacon,” Kontakt.io, 2021. Accessed: Aug. 24, 2021. [Online]. Available: <https://store.kontakt.io/our-products/30-smart-beacon-sb16-2.html>
14. E. De Cristofaro and G. Tsudik, “Practical private set intersection protocols with linear complexity,” in *Financial Cryptography and Data Security*, R. Sion, Ed. Berlin: Springer-Verlag, 2010, pp. 143–159.
15. A. Carroll and G. Heiser, “An analysis of power consumption in a smartphone,” in *Proc. 2010 USENIX Conf. USENIX Annu. Tech. Conf. (USENIXATC’10)*, p. 21.

**Pietro Tedeschi** is a Ph.D. student in the Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Doha, 5825, Qatar. His research interests cover security issues in unmanned aerial vehicles and wireless, Internet of Things, and cyberphysical systems. Tedeschi received his master’s degree with honors in computer engineering at Politecnico di Bari, Italy. He is an Associate Member of IEEE. Contact him at [ptedeschi@hbku.edu.qa](mailto:ptedeschi@hbku.edu.qa).

**Kang Eun Jeon** is a Ph.D. student in the Social Media Lab, the Hong Kong University of Science and

Technology (HKUST), Kowloon, Hong Kong. His research interests include self-sustaining, secure, and social Bluetooth low-energy beacons for Internet of Things applications. Jeon received his B.Eng. degree in electronic engineering at HKUST. He is a Student Member of IEEE. Contact him at kejeon@ust.hk.

**James She** is the founding director of the Social Media Lab, the Hong Kong University of Science and Technology and an associate professor in the Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Doha, 5825, Qatar. His current research interests include social and multimedia computing, data science, artificial intelligence for visual creativity, and the Internet of Things for sustainable, smart, and interactive systems. She received a Ph.D. in electrical and computer engineering from the University of Waterloo, Canada. He is an IEEE Member. Contact him at eejames@ust.hk.

**Simon Wong** is an M.Phil. student at the Hong Kong University of Science and Technology, Kowloon, Hong Kong. His research interests include technologies on machine learning and signal processing for Internet of Things devices. Wong received an M.A. in philosophy from Hong Kong University of Science and Technology. He is a Member of IEEE. Contact him at simonwong201293@gmail.com.

**Spiridon Bakiras** is an associate professor of cybersecurity in the Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Doha, 5825, Qatar. His research interests include security and privacy, applied cryptography, and spatiotemporal databases. He has 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. Bakiras received a Ph.D. in electrical engineering from the University of Southern California, Los Angeles. He is a Member of IEEE. Contact him at sbakiras@hbku.edu.qa.

**Roberto Di Pietro** is a full professor of cybersecurity in the Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Doha, 5825, Qatar. His main research interests include security and privacy, distributed systems, virtualization, and applied cryptography. He is an Association for Computing Machinery Distinguished Scientist and a 2020 recipient of the Jean-Claude Laprie Award in Dependable Computing. Di Pietro received a Ph.D. in computer science from the Università La Sapienza, Rome, Italy. He is a Senior Member of IEEE. Contact him at rdipietro@hbku.edu.qa.