Presence Detection Based on Bluetooth Low Energy Using the Exposure Notification Service

Yannis Greve and Prof. Dr. Frank Oldewurtel
Department of Electrical Engineering and Information Technology,
South Westphalia University of Applied Sciences
Hagen, Germany
E-Mail: oldewurtel.frank@fh-swf.de

***Abstract*—Detecting if a person is present offers a variety of use cases in a smart home environment. These start at energy saving methods, such as reducing the room temperature when no one is at home, and go as far as automated routines to turn off unnecessary lighting. In this paper, we propose a method to detect a user's presence based on the Bluetooth Low Energy Advertising of his mobile device. This method is based on an analysis of the Bluetooth behavior of common smartphones, which identifies the Exposure Notification Service (ENS) to provide a proper source of Bluetooth Advertising messages. The ENS was introduced in April 2020 to solve the too-long-too-close-problem in the fight against the COVID-19 pandemic. Given that the ENS provides periodic advertising by the mobile device, the user does not need to take any action to be identified by a presence detection system. Using a random resolvable Bluetooth address, this method can detect a user by identifying the used identity resolving key. This paper proposes a procedure to implement this method in low power smart home devices, ensuring that the impact on the devices' resources is as small as possible. The proposed method provided an automatic, reliable, scalable and low-cost way of detecting the user's presence. In addition it can count and identify the present users to enable new use cases.**

I. Introduction

Looking at the current energy and climate crisis, smart home environments provide strong economic and ecological use cases (e.g., reducing the room temperature at night or dimming lights after sunset). As an extension of reducing the room temperature at night, the smart home system can shutdown high energy services when they are not needed. Common use cases are the so-called coming home and leaving home functions. These are triggered by the smart home user's presence, and therefore require presence information. Currently in the presence detection context smart homes rarely operate automatically and usually require additional devices or components. Solving such technical issues would result in a significant competitive advantage. In this paper, we propose a method for detecting the presence of users in a smart home environment based on the Bluetooth Low Energy fingerprint of the users smartphones. The contribution of this paper is an automatic, reliable, and scalable method that is usable in new and existing products, which distinguishes it from the current state of the art. Instead of locating a Bluetooth device in an artificial and prepared environment, this paper proposes a way to detect a device's presence without any requiring interaction or application running on it. At last, the presence detection is performed on the tracking device without the use of any servers and no additional beacons are required. To our knowledge there is no approach reported which targets this goal

The remainder of the paper is organized as follows. This section introduces technical solutions for presence detection. Section II discusses related work. Section III presents the results of experiments considering the Bluetooth Low Energy behavior of common smartphones. Section IV proposes our novel method that exploits Bluetooth advertising provided by the exposure notification service for detecting the presence of the user's devices. Section V discusses the developed method, describes its advantages and disadvantages, and outlines the new use cases that it provides. Finally, Section VI gives an overview of the enabled problem, the solution and draws conclusions.

A. Current Solutions

In the first step, a short overview of current solutions of a presence detection in a smart home environment is given to point out their limitations. A first solution, which represents the state of the art, is a wall installed switch that the users can press to indicate that they are coming or leaving their home. Even though this is a highly reliable solution, the manual operation appears as a disadvantage from a user's point of view. Additionally, these solutions require extra components because the wall switch itself needs to be attached and the required electrical wiring installed. This limits the usage of this first solution in existing buildings.

A second approach is the installation of motion sensors (e.g. using passive infrared senors) to detect the presence of a user. These motion sensors can be installed in the ceiling of each room, where it monitors the user's presence through their motions. As an extension of the first solution, additional components are required but an automatic detection of the user is provided. Because the sensor has to maintain a visual connection to the user, at least one sensor per room is needed. Therefore, the component and installation costs rise in correlation to the size of the living space that they monitor. The visual sensing technology also limits the reliability of this solution because presence detection within blind spots of convoluted living space is not possible.

ISBN: 978-3-903176-61-4

© 2024 International Federation
for Information Processing (IFIP).
ISBN: 978-3-903176-61-4

Looking at existing smart home devices, only a few have sensors that can be used to detect the user's presence. Additionally, the majority of smart home devices are not installed in situations where they could supervise a whole room. Meanwhile, most devices do have a radio interface that is used to communicate with other smart home devices or with commissioning devices. Smart home devices using a radio interface following the IEEE standard 802.11 (better known a WLAN) can already use this interface to detect the presence of users, as demonstrated in [1]. To do so, the devices need to access the user's local network to scan its traffic for the user's devices. The smart home device can use the scanned traffic to verify the user's presence through messages to their mobile device exchange within the local network. This presence detection does not rely on any additional components but is a solution based on software. On the downside, WLAN requires a lot of resources (e.g., memory, processing power, and energy) and most available platforms use dedicated WLAN transceivers to match these requirements. Therefore, these solutions have a negative impact on costs per unit for smart home devices and also limit its use in battery-powered devices.

In summary the current solutions to detect the presence of users in a smart home environment are limited by their automatic functionality, their ability to implement it in already installed devices and by their costs per unit and installation. This paper evaluates the possibility to eliminate these major disadvantages by using a Bluetooth Low Energy based Method of presence detection. To achieve this, it focuses on Bluetooth low energy because it is a widespread technology with 4 million devices sold in 2021 alone [2]. Especially for smart home applications, the use of Bluetooth low energy is expected to rise by 440% within the next 5 years. It is also designed to work on resource constrained and low-cost devices, and is implemented in most used smartphones.

II. Related Work

Bluetooth presence and asset tracking is a well analyzed field, and there have recently been a large number of new perceptions. For example, the authors of [3] describe a method to track gallery visitors to provide further information about the exhibitions on the visitors' smartphones. To achieve this, the exhibitions are equipped with Bluetooth beacons that send an identifier via periodic Bluetooth advertising messages. The visitor's smartphone runs a custom application that receives these advertising messages and then extracts the exhibition's identifier. By comparing the signal strength (RSSI) of the received messages, the application can determine which exhibitions the user is the closest to and show some information about the exhibit on the smartphone's screen.

The authors of [4] tie in with this method. Their paper presents a Bluetooth beacon-based detection method to prevent traffic accidents between pedestrians and larger vehicles, such as trucks. In this setup, the vehicle is equipped with Bluetooth beacons and the pedestrians carry a receiver that scan for the beacons' messages. By analyzing the beacon signal's strength, the receiver can calculate its position in relation to the vehicle. While constantly calculating this position, the receiver will trigger an acoustic signal to warn the user if a pedestrian enters the vehicle's blind spot. The authors conclude that their method is very reliable but they have to divide the vehicle's surroundings into large chunks of detectable zones. Therefore, the pedestrians' receivers will warn them when they get as close as 8 meters.

The authors of [5] expanded this method by using multiple devices within a fixed test environment and were able to design a technique to calculate a user's absolute position. In their paper, they demonstrate how to place multiple beacons in a room and make a user's smartphone calculate it's position based on the received message's signal strength. Tracking is based on estimating the distance to each Bluetooth beacon using a matching systems model. In the next step, a trilateration using the distance to each beacon can be performed. However, because this is a complex calculation, the data is forwarded to a server performing the calculation. In addition, the smartphone's application has to constantly collect data and forward it to the server, which affects the ability to run in the background and reduces the battery's lifetime.

The authors of [6] are able to design a location service by switching the roles of sending and receiving. They present a method to locate patients using multiple Bluetooth sensors placed in a medical environment. While the patients wear Bluetooth dongles, the sensors scan for their periodic messages. Following the procedure of the previous work, the authors show that using the signal strength of the patient's beacon collected by each sensor can locate the patient with a success rate of 97%. The author's work shows that a highly accurate location of Bluetooth Beacons is possible, if the system's environment is well designed. However, running the authors' neuronal networks requires too much computational power to do it on the Bluetooth sensor. Therefore, the locating calculations are performed on a central server.

The amount of relevant work to this topic shows that person- and asset-tracking using Bluetooth is a well researched field. Nevertheless, these papers do not present any suitable solution for presence detection in a smart home environment. Therefore this paper will show a method to detect a user's presence using common devices without the need of a laboratory environment and also show that this can be implemented in a resource constrained embedded device. To the best of our knowledge, there is no related work that can provide a solution matching these and the requirement mentioned in the previous chapter.

III. Bluetooth Advertising of Common Smartphones

To implement a presence detection system as described in the previous sections, the Bluetooth Low Energy footprint of the user's mobile end devices can be used. Given that these devices can be from a variety of manufacturers, they can behave differently regarding their Bluetooth socket. Therefore this paper had a deeper look into the Bluetooth Advertising of common mobile end devices to provide a method that works in an in-market environment.

ISBN: 978-3-903176-61-4

The objective of this step is to state out how, when, and what data different smartphones send using Bluetooth Low Energy. To achieve this, the devices under test (DUTs) are placed in a low traffic environment and their sent Bluetooth messages are tacked using a BLE-Sniffer and Wireshark. The influence of different operating systems and versions, (de-)activated Bluetooth sockets and (un-)locked devices were part of these measurements. While performing these measurements, all of the DUTs screens are locked to ensure that the observed behavior is available in the background (i.e. without any interaction of the user).

Looking at the results, that are not displayed in this paper, it gets clear, that not all smartphone send periodic BLE advertising messages. Especially when being locked, some devices do not transmit any messages. To detect a device based on it Bluetooth fingerprint, it needs to send some kind of message. This can be resolved if the device uses an application that implements the Exposure Notification Service (ENS). The ENS service profile was released by the Bluetooth Special Interest group in 2020 to solve the too-long-too-close-problem in the fight against the COVID-19 epidemic [7]. It was designed to detect the physical distance between two devices without sharing the identity of these devices with each other. Because the ENS is handled by the operating system itself and only started by the application that uses it, this process runs completely in the background of the device. The corresponding messages use a random resolvable address and will be sent, even if the device is locked and the application that implements the ENS is closed.

IV. Presence Detecting Algorithm

After the ENS was identified as a source of periodic advertising, a method to detect the device's presence can be derived. The mobile end-device will serve as the transmitter, the detecting smart home device as the receiver of the ENS' messages. Because the ENS is designed to save battery power by sending messages rarely and with reduced transmitting power, the impact on the mobile device's battery is minimal.

A simple solution is based on evaluating the advertising address of the received advertising messages generated by the device's ENS. In this situation, a smart home device (e.g., a switch or a dimming actuator) scans for Bluetooth messages. The user's devices (as described in the previous section) send periodic advertising messages that are used and generated by the ENS. The smart home devices check upon receiving an advertising message. If the used advertising address can be associated with a known device, then the device's user can be assumed to be present. Because the advertising of the ENS is sent with limited transmitting power, receiving an advertising message from a device relates to a physical proximity to the transmitting device.

Due to privacy concerns, the address of the ENS advertising is changing every 15 minutes. The kind of advertising address used by the ENS and is called "resolvable private random address." This address is generated randomly to avoid an eavesdropper tracking a device based on the advertising address by using a device-specific "identity resolving key". In comparison to static or non-resolvable random addresses, resolvable addresses allow an association between the address and the generated device if the corresponding key is known. This procedure is described below.

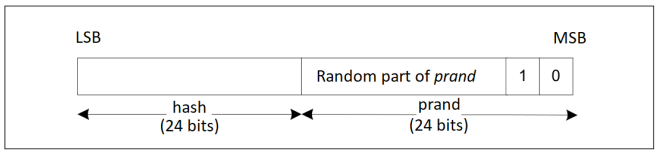

Fig. 1. Structure of a random resolvable address. Adapted from [8]

The random resolving address is separated into two parts, as seen in 1. The lower 24 Bits of the 48 Bit address contain a random number (prand), while the lower Bits contain a hash value. The transmitted hash value can be calculated from the random number using the so-called "Identity Resolving Key" (IRK). Every device using a random resolvable address owns at least one of these 128 Bit symmetric keys, which is described by the Bluetooth Low Energy specification. If a device owns an IRK, then it will exchange it when pairing in a Bluetooth Low Energy connection with another device. Because a simple comparison of the random resolvable address of a received message with a known address does not work due to its randomness, the resolving mechanism can be used.

To achieve this, the receiving device has to calculate the hash from the received random number and the IRK itself. If it matches with the received hash, then the advertising address can be matched to the device using the IRK used in the calculation. The calculation of the hash is specified as a decryption of the to zero-padded random number with the IRK using a AES-128 [9] blockcipher. Only the 24 least significant bits of the decrypted value are used as the calculated hash. Equation (1) sums up this resolving calculation:

$$H = AES_{128}(IRK,\ prand_{ZeroPadded})\ mod\ 2^{24} \quad (1)$$

The probability of collisions rises due to the reduction of the hash size to 24 Bits. But there are still 2^{24} = 16777216 different possible hash values and usually just one IRK per device, resulting in a sufficiently low collision probability for low energy devices. Because the IRK is only shared with trusted devices in the pairing process, the ability to imitate another device by creating a random resolvable address using its IRK is only given to those trusted devices [8].

The standardized resolving mechanism can be used as the base of the presence detection. If the scanning smart home device knows the IRK of the user's mobile device, then it can resolve its advertising address; and therefore match the address with the user's device and verify its presence. Upon reception, the scanning smart home device tries to resolve the received advertising address. Using every known IRK from every device, the hash value is calculated based on the received random number. If the calculated hash value

ISBN: 978-3-903176-61-4

```
Bluetooth Message received
extract advertising address
if advertising address is random private or static random then
    return // Do not handle the message
end if
if Address is found in address cache then
    return // Do not handle the message
else
    if Address is random resolvable then
        for all Identity Resolving Keys do
            extract random number prand from address
            H_calculated ← AES_128(IRK, prand) mod 2^24
            if H_calculated == H_received then
                if Device using this IRK is already present then
                    Remove Link from old address and link new address to the user's device
                else {User's device not present}
                    Mark device using the IRK as present
                    Link the address to the user's device
                end if
            end if
        end for
        Add address to address cache
    end if
end if
```

Fig. 2. Detect device's presence

matches the received value, then the device using this IRK is marked as present and the resolving stops. Consequently, a received random resolvable address can be matched with known devices.

Because a presence detection system should be able to handle multiple users, every received random resolvable address has to be resolved with every known IRK. Therefore, the time taken by resolving received addresses rises with the number of users whose presence shall be detected. Using the ENS, which sends its advertising messages every 100 milliseconds, a smart home device has to take care of multiple resolving attempts per second. If there is more than one transmitting device nearby, which is a realistic scenario in urban areas, then the rate of resolving tries per seconds rises even further. In such cases, the presence detecting devices are faced with the need to resolve hundreds of addresses every second, which results in the device taking a long time to calculate the hash values. As mentioned earlier, the resolving mechanism is based on a decryption using the standardized AES blockcipher, meaning that the time that this step takes is highly optimized using libraries such as Mbed TLS [10]. For this reason, the only way to reduce the time that a presence detecting device spends resolving addresses is by reducing the total amount of resolving attempts. Therefore, an optimized method, which reduces the number of resolutions performed for every transmitting device to one, is provided in 2.

The extended method is based on resolving a received address only once because the same random number will always lead to the same calculated hash using a given list of IRKs. To achieve this, the device uses caches to store the received and already resolved random resolvable addresses. As seen in 2, upon reception of a message, the smart home devices will check if the address can be found in their address cache. If the address can be found, then the device will not try to resolve it and ignore the received message, ensuring that no unnecessary resolving is performed. Otherwise, the device will try to resolve the address using all IRKs known. Once resolved, the address is placed in the address cache, whether or not the resolving was successful. If the resolving was successful, then the address now stored in the cache is linked to the known user, who's IRK was used. Consequently, the smart home devices can verify the presence of a user's device without resolving its address, simply by finding it and its resolved identity in the cache.

If the messages of an address that is linked to a user's device stop, then two cases are possible: either the user's device is out of range of the scanning smart home device or the device has changed its random resolvable address by choosing another random number. In the first case, the smart home did not receive any messages from the user's device for a certain time because the device has left the smart home's proximity. The smart home device can now mark the user's device, and therefore the user, as absent. In the second case, the transmitting device stopped using the already resolved address and generated a new address. Because the user's device will use the new address immediately, the smart home device will successfully resolve the new address before the receiving timeout of the old address runs out. The smart home device will not only link the user's identity to the new address and consider them to still be present but will also delete the old address and its link from the address cache. Under the ENS, this change of the used random resolvable address happens every 15 minutes.

Similar to the random resolvable addresses, this technique can also be used to detect a device using static addresses. This allows the presence of devices such as wearables, headphones, or Bluetooth dongles to be detected. Because these addresses do not need to be resolved, a comparison between the received address and the addresses stored in the address cache provides sufficient information to confirm the presence of the device.

The presented method was implemented in a prototypical system to verify its function. For this purpose, a small smart home system of three Bluetooth mesh devices based on a Silicons Labs EFR32BG21 platform was set up as senn in 3. The system was able to detect the presence of multiple devices without the need to unlock any of them. To synchronize the presence of a user between multiple smart home devices, a consensus protocol that is based on Raft [11] was implemented. It also exchanged the IRKs between all smart home devices if a user performs a pairing with one of them.

ISBN: 978-3-903176-61-4

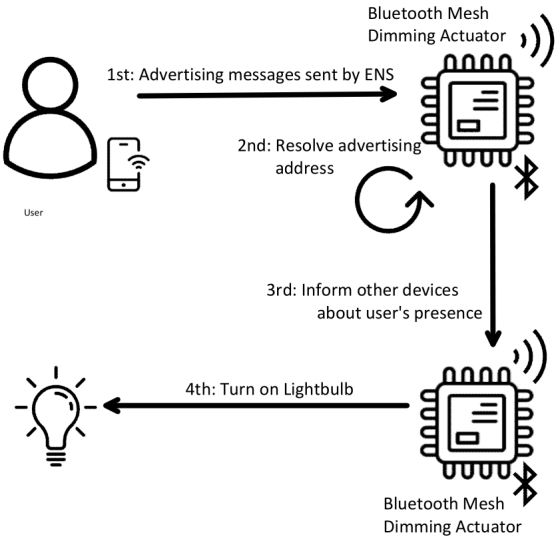

Fig. 3. Structure of implemented prototype

V. Discussion

The unique attribute of the designed method is that it is able to detect the user without requiring any interaction from the user. Not only does the interaction-less detection provide comfort for the user but it also comes without the need of any additional components, such as sensors. The software-based method can be implemented in every smart home device capable of Bluetooth Low Energy communication, even through updates in existing devices. This way a presence detection functionality can be implemented in both new and existing devices for the same production costs. The automatic detection also provides advantages in the reliability of the method. Since there are no black spots, there is no need to have multiple devices in a single room to cover it entirely. However, the method relies on the user to carry their mobile device with them at all times. Additionally, that mobile device needs to have its Bluetooth adapter turned on and it requires an application that uses the ENS. If these requirements are not met, then it is not possible to detect the user. Even though the impact on the battery power of the user's device was considered, sending periodic Bluetooth messages still needs additional power from the battery.

Given that this method is based on resolving the received addresses to check if the messages originate from the user's device, the smart home device can identify the present user. Therefore, this method provides two new pieces of information besides the presence itself: which and how many users are present. This generates a variety of new use cases that these devices can take part in. For example, the user's identity can be used to apply user specific configurations, such as room temperature or lighting. These applications can be relevant for shared offices, where the furniture can be set to a preferred height. Also, Smart homes can use the information to set the user's preferred warm water temperature when using a warm water boiler. Another example is preemptively adapt heating or cooling devices to prevent the room's temperature from drifting depending on how many users are present. If the designed method can be combined with a more precise location detection, as shown in [6], even more use cases from the smart home context can be created.

VI. Conclusion

This paper introduced a method for detecting the presence of the user's mobile device in a smart home environment. Based on Bluetooth advertising messages, smart home devices can detect the presence of the user's device through the presence of its advertising messages . Using the ENS in background processing, the user's device still transmits advertising messages even when it's locked.

Furthermore this work presents an implementation of this method that has been applied to resource constrained and low-cost embedded devices. Because the presence detection procedure is a software-based solution, no additional components are needed. Therefore, it provides a scalable solution, that can be used not only in a laboratory environment but is also interoperable with commercially available devices. In contrast to existing methods, the advertising address included in the ENS messages is used to identify the transmitting device, making it possible not only to identify the present users, but also to count them. This enables many more use cases, that cannot be realised by conventional presence detection methods.

References

[1] A. Freyer, "Monitor," Jan. 2022, (accessed on 2022-10-03). [Online]. Available: https://github.com/andrewjfreyer/monitor

[2] Bluetooth SIG, "2021 Market Update," 2021, (accessed on 2022-09-25). [Online]. Available: https://www.bluetooth.com/wp-content/uploads/2021/01/2021-Bluetooth_Market_Update.pdf

[3] M. Pušnik, M. Galun, and B. Šumak, "Improved Bluetooth Low Energy Sensor Detection for Indoor Localisation Services," *Sensors*, vol. 2336, 2020.

[4] N. De Raeve, M. de Schepper, J. Verhaevert, P. Van Torre, and H. Rogier, "A Bluetooth-Low-Energy-Based Detection and Warning System for Vulnerable Road Users in the Blind Spot of Vehicles," *Sensors*, May 2020.

[5] J.-H. Huh and K. Seo, "An Indoor Location-Based Control System Using Bluetooth Beacons for IoT Systems," *Sensors*, vol. 2917, Dec. 2017.

[6] G. Tang, Y. Yan, C. Shen, X. Jia, M. Zinn, Z. Trivedi, A. Yingling, K. Westover, and S. Jiang, "Development of a real-time indoor location system using bluetooth low energy technology and deep learning to facilitate clinical applications," *Medical physics*, vol. 47, pp. 3277–3285, 08 2020.

[7] G. Inc., "Exposure Notifications," 2021, (accessed on 2022-10-03). [Online]. Available: https://www.google.com/covid19/exposurenotifications/

[8] Apple Inc., Ericsson AB Intel Corporation, Lenovo (Singapore) Pte. Ltd., Microsoft Corporation, Nokia Corporation, and Toshiba Corporation, "Bluetooth Core Specification," 2019, (accessed on 2022-10-01). [Online]. Available: https://www.bluetooth.com/specifications/bluetooth-core-specification/

[9] National Institute of Standards and Technology (NIST), "Announcing the Advanced Encryption Standard (AES)," 2001, (accessed on 2022-08-17). [Online]. Available: https://csrc.nist.gov/csrc/media/publications/fips/197/final/documents/fips-197.pdf

[10] L. Limited, "Mbed tls," 2022, (accessed on 2022-10-03). [Online]. Available: https://www.trustedfirmware.org/projects/mbed-tls/

[11] D. Ongaro and J. Ousterhout, "In Search of an Understandable Consensus Algorithm," in *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, ser. USENIX ATC'14. USA: USENIX Association, 2014, p. 305–320.

ISBN: 978-3-903176-61-4