Is Periodic Pseudo-randomization Sufficient for Beacon Privacy?

Liron David, Avinatan Hassidim, Yossi Matias, and Moti Yung

Google Research

Abstract. In this paper, we investigate whether the privacy mechanism of periodically changing the pseudorandom identities of Bluetooth Low Energy (BLE) beacons is sufficient to ensure privacy.
We consider a new natural privacy notion for BLE broadcasting beacons which we call "Timed-sequence-indistinguishability" of beacons. This new privacy definition is stronger than the well-known indistinguishability, since it considers not just the advertisements' content, but also the advertisements' broadcasting times which are observable in the physical world.
We then prove that beacons with periodically changing pseudorandom identities do not achieve timed-sequence- indistinguishability. We do this by presenting a novel privacy attack against BLE beacons, which we call the "Timer Manipulation Attack." This new time-based privacy attack can be executed by merely inserting or reinserting the beacon's battery at the adversary's chosen time. We performed this attack against an actually deployed beacon.
To mitigate the "Timer Manipulation Attack" and other attacks associated with periodic signaling, we propose a new countermeasure involving quasi-periodic randomized scheduling of identity changes. We prove that our countermeasure ensures timed-sequence indistinguishability for beacons, thereby enhancing the beacon's privacy. Additionally, we show how to integrate this countermeasure in the attacked system while essentially preserving its feasibility and utility, which is crucial for practical industrial adoption.

Keywords: Privacy, IoT, broadcasting, BLE, Ephemeral IDs, beacons, beacons advertisement, pseudorandom advertisement, time-based attacks, countermeasures, randomized scheduling.

1 Introduction

A beacon is a small, wireless, battery-operated Bluetooth Low Energy (BLE) device which broadcasts a BLE advertisement along with its MAC address. Within each BLE advertisement, there exists a distinctive field known as the identity, intended to be unique to each device and link the beacon with its owner. All other fields within the advertisement remain constant or appear as randomized (pseudorandom) ciphertexts, with the MAC address also exhibiting a randomized appearance.

To make the entire advertisement look random, beacons employ a cryptographically randomized ephemeral identity (instead of the fixed identity). The beacon changes its identity at fixed intervals, typically every predefined time period T, e.g., 15 minutes (with some only changing identities when the beacon is in proximity to its owner). Namely, each beacon broadcasts Ephemeral Identities (EIDs) that are pseudorandom and refreshed regularly.

Each identity-change forces a change of the random-looking fields associated with an advertisement, namely the MAC and re-randomized ciphertexts. Consequently, following an EID change, the entire new advertisement and MAC look unrelated to the previous advertisement and MAC. The necessity of identity changes in periodic fashion is paramount for optimizing several performance parameters within the system, a topic we delve into with further detail in Section 2.3. The pseudorandom identities can be correlated with the specific beacon through a unique key shared exclusively between the beacon and its owner. Consequently, the owner possesses the capability to recognize the EIDs of its beacon, while to any other observer the beacon's EIDs look random.

When the beacon is far from its owner, the owner retrieves its beacon's whereabouts through BLE-to-IP gateways, typically smartphones, which happen to be in the beacon's vicinity. These BLE-to-IP gateways, called observers, collect the beacon's broadcasts and forward them to a cloud server along with their own

location information, which may sometimes be encrypted. The cloud server then either routes the observer's location directly to the beacon's owner, or otherwise stores it within its database for the owner to request it later on.

BLE beacons, in fact, are becoming highly prevalent in the Internet of Things (IoT). Major examples include Apple's Airtags [1,10], Google's coming beacons whose protocols are associated with the recently announced Find My Device protocols [16,15] (an evolution of Eddystone-EID [7] and [5]), and Samsung's SmartTag [24].

Apple's Airtags do not change their EIDs when far from their owners in order to enable rogue detection. Without changing EIDs, an Apple's Airtag may be exposed to simple tracking attacks based on its stable ID. Namely, Airtags do not achieve the well-known indistinguishability privacy when far from their owners. Regarding Google's beacons, they didn't publish their beacon product yet, so we cannot relate to their coming beacons.

Samsung's SmartTags, as we've seen in our experiments, keep changing their EIDs when they are far from their owners, and therefore achieve indistinguishability privacy. A major outcome of this paper is to show that when the smartTags are far from their owners, they can nevertheless be exposed to tracking based on their EIDs-change times. That is, while cryptographic pseudorandomization achieves indistinguishability, it may be not enough to prevent tracking.

Specifically, we define a stronger privacy notion that better captures real-world tracking scenarios in the physical world. In this tracking scenario the adversary hears the beacons' signals as well as their broadcasting times and can control the beacons' initialization times. This is in contrast to the weaker well-known indistinguishably definition in which the adversary only hears the beacons' signals.

The real-world tracking scenario we cope with in this work is as follows: Adversary $\mathcal{A}$ is provided with m beacons, one of which is special. The adversary can control these beacons' initialization times, that is, the time at which each beacon starts broadcasting its ephemeral identities. After initialization, adversary $\mathcal{A}$ monitors the signals and the broadcasting times of these m beacons. The adversary then loses contact with these beacons for a period of T_{off}. When contact is back, adversary $\mathcal{A}$ again monitors their signals and broadcasting times. The goal of adversary $\mathcal{A}$ is to identify the special beacon among the m beacons.

Here are two real-life examples that motivate the above tracking scenario:

1. **Tracking a special box along the supply chain**: Alice's company sells some equipment to Bob's company. The equipment that arrives to Bob's company is organized in similar-looking boxes, with beacons attached to these boxes to assure traceability of the boxes. One of the items that Bob buys from Alice's company is very special. A passive adversary at Bob's company, "EveB," who is in charge of collecting the received boxes, wants to identify the box with the special item among all the similarly-looking boxes without opening the boxes, therefore based on the beacons' signals only. To this end, EveB cooperates with "EveA," an adversary at Alice's company who is responsible for attaching the beacons to the boxes and therefore can control the beacons' initialization times.
2. **Tracking a special suitcase traveling from Airport-A to Airport-B**: Alice travels from airport-A to airport-B with a special item hidden in a suitcase which looks like most of the suitcases in the airport. EveB, who is waiting at airport-B's baggage claim, aims to identify and snatch the suitcase with the special item among all the similar-looking suitcases. To this end, EveB collaborates with EveA who works at airport-A and who is in charge of collecting suitcases after check-in and loading them into the airplane. EveA therefore can control the beacons' initialization times before loading.

To formally model the above real-world tracking threat, we propose a new privacy definition based on a cryptographic game which considers both beacons' signals and broadcasting times. We call this new privacy definition "timed-sequence- indistinguishability" of beacons. We then show that beacons which periodically change their ephemeral identities do not achieve the stronger privacy of timed-sequence- indistinguishability. We do this by presenting a new time-based privacy attack against smartTags which we call the "Timer Manipulation Attack." We then propose a new effective time-scheduling countermeasure that mitigates the "Timer Manipulation Attack" and similar attacks related to periodic signalling, and we prove that beacons using our countermeasure achieve the stronger privacy of timed-sequence- indistinguishability. This

countermeasure, along with its analysis, is proposed as the main privacy enhancing technique we put forth in this work.

The new "Timer Manipulation Attack" is a concrete threat to the privacy of beacons and to the best of our knowledge, the literature regarding beacons and their privacy, so far, has ignored the possibility of adversarial use of identity broadcasting times. The attack and the work presented in this paper has already influenced Google's coming beacon signalling protocol where randomization has been added to he ID changing times. Further, we hope this work raises the above issue to, both, designers of existing systems and future designers attempting privacy-by-design of broadcasting elements with pseudorandom ephemeral IDs and messages.

It is worth noting that while this paper concretely deals with BLE, the methodologies discussed are based on general methods and are equally applicable to any future broadcasting network technology.

1.1 Contributions

In this paper, we show that periodically changing the pseudorandom identities of BLE beacons is not sufficient to ensure their privacy in real-world scenarios. Instead, we propose a new randomized time-scheduling. The central contribution is therefore the randomized countermeasure and its analysis. In more details, this paper:

1. Presents "timed-sequence- indistinguishability," a new privacy definition for beacons which also considers broadcasts' times, hence better captures tracking scenarios in the physical world.
2. Shows the "Timer Manipulation Attack" which we performed against Samsung's SmartTag. This attack demonstrates that changing the pseudorandom identities periodically does not achieve the timed-sequence- indistinguishability for beacons.
3. Suggests a new countermeasure that mitigates the attack while it essentially retains the performance and other system required properties of the periodic scheduling. Our countermeasure is a new random quasi-periodic time-scheduling mechanism for a beacon to change its ephemeral identity.
4. Proves that beacons using the quasi-periodic time-scheduling countermeasure achieve timed-sequence-indistinguishability.
5. Shows that, in principle, the new countermeasure can be efficiently accommodated in Samsung's SmartTag system. The accommodation maintains the utility and feasibility of the underlying system (originally derived from its periodic nature)– a fact which is extremely important for future deployment and adoption of the mechanism.
6. Finally, analyses the desired concrete parameters of the new countermeasure which allow the system to achieve good performance while maintaining privacy, and presents a numerical simulation studies with these concrete parameters.

2 The BLE Beacons Ecosystem

To better understand our context, we next describe in detail how the beacons operate and explain why the ephemeral identities, essential for privacy, need to be changed periodically to achieve a good balance between privacy and performance.

2.1 The Players and Configuration

There are four players in the BLE beacon ecosystem, see Figure 1:

– **The beacon:** a BLE device, which broadcasts its advertisement along with its MAC address over the air every time-unit (e.g. a second). Each advertisement includes an ephemeral identity (EID) and additional fields that are either constant values independent of the beacon or pseudorandom ciphertexts. The MAC address is also pseudorandom. The EID is a pseudorandomized identity that changes according to a time-scheduling algorithm. This algorithm typically changes EIDs periodically, which is a straightforward and manageable process (see Section 2.3). When EID changes, both the MAC and the pseudorandom ciphertexts in the beacon's advertisement change too.

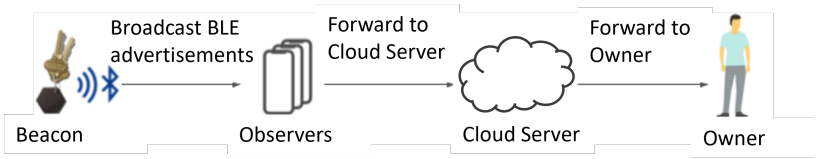

Fig. 1. The BLE Beacon Ecosystem.

- **The beacon's owner:** typically a smartphone or any device with BLE and IP communication support, which has been paired with its beacon and exchanged private keys. When the owner is in the vicinity of its beacon, it receives its beacon's broadcasts directly and can therefore recognize its beacon location. This is particularly useful if the beacon is in its owner's vicinity but it is hidden from its owner's eye. However, if the owner is far from its beacon, the following two participants become necessary:
- **The observer:** an arbitrary BLE-to-IP gateway, typically a smartphone, which happens to be in the beacon's vicinity. The observer forwards the beacon's advertisement, along with its own (encrypted) location, to the cloud server over IP. If the observer hears the same advertisement several times (since for example it stays at the same place near the beacon), the observer forwards the advertisement only once, and ignores any subsequent repetitions.
- **The cloud server:** responsible for forwarding the observer's (and thus the beacon's) location to the owner. The cloud server can operate in one of two ways: (1) as a router, which sends each new beacon's location directly to the beacon's owner. To facilitate this, the cloud server maintains a mapping table that associates beacons' ephemeral identities with owners' IDs; or (2) as a database, which stores each incoming beacon's location and allows the respective owner to query for it.

2.2 Performance Vs. Privacy Requirements

The goal of the BLE beacon ecosystem is to enable communication within the described configuration while achieving a good balance between performance and privacy. That is,

- Performance: Low power consumption. The power consumption of the entire system should be minimized. Specifically, both the computational workload and storage requirements should be kept to a minimum while still ensuring that the system provides accurate, reliable, and secure service.
- Privacy: Short-lived stable EID. The duration for which the EID remains stable should be minimized. A shorter stable EID period reduces the window during which an adversary can track the beacon based on its stable ephemeral identity.

2.3 Why Periodic EID-Change Algorithm?

To achieve an optimal balance between privacy and performance, the beacon should not change its ephemeral identity with every time-unit. While frequent changes enhances privacy, they also lead to excessive and inefficient power usage across all system components:

- For the beacon, which would need to compute an enormous number of identities, leading to frequent battery replacements.
- For the arbitrary smartphone observer, which would need to encrypt and forward its location to the cloud server with every beacon broadcast it receives each time-unit.
- For the router-flavored cloud server, which would need to maintain a huge mapping table from ephemeral identities to their owners, requiring an entry for each beacon and each time-unit in the upcoming period. Additionally, it would need to handle a large number of lookup operations in this extensive mapping table.
- For the database-flavored cloud server, which similarly would need to store a vast number of locations and handle a high volume of queries

– For the owner, which would need to manage an enormous number of (encrypted) locations received from the cloud server.

We can therefore conclude that the beacon should not generate a new ephemeral identity every second, but only once in a while. The question is which time-scheduling for the identity-change will lead to the optimal balance. One candidate for such time-scheduling is the Geometric distribution, where at any second t, there is a probability of 0.5 (say) that the beacon will change its identity at time t. However, this approach will negatively effect both performance and privacy:

1. Performance: It also enables a scenario in which at a certain unpredictable time the beacon changes its ephemeral identity every time-unit. To support this scenario, the system must be designed to handle a potential EID change every second. As previously discussed, this requirement results in significant overall power consumption for the system.
2. Privacy: It enables a scenario in which at a certain unpredictable time the beacon changes its ephemeral identity more rarely. In this case the beacon is certainly exposed to a longer (and theoretically unbounded) period during which the beacon is exposed to tracking based on its stable identity.

As illustrated by the example above, achieving an optimal balance requires that the difference between any two consecutive EID-change times be neither unpredictable nor unbounded and should be easily computable. Unpredictable EID-change times necessitate that the cloud server be prepared for the worst-case scenario of an EID change every second, resulting in significant storage demands and high power consumption for the whole system. Unbounded EID-change times, on the other hand, can lead to extended periods of tracking based on the beacon's stable identity. To mitigate unpredictability and unboundedness, the most straightforward approach is to periodically change the EID every T time-units.

Therefore, using distributions such as Geometric or Poisson for time-scheduling ephemeral identity changes is not suitable for BLE beacon scenarios. Instead, a periodic deterministic approach is needed to ensure smooth, regular, and predictable performance. For example, Apple's AirTag, Samsung's SmartTag, and Google's Eddystone-EID all change their ephemeral identities at fixed intervals of T time-units, where T is predefined in their protocols.

Definition 1. *(T-Periodic Time-scheduling) We say that a time-schedule of EIDs-change is "T-Periodic" if the ephemeral identity is changed every T time-units, see Figure 5 (bottom).*

Note that although the beacon changes its ephemeral identity every T time-units, it broadcasts the same identity at each time-unit, repeating it until the next scheduled change. This approach not only improves system efficiency but also boosts the probability of successfully reaching its owner, thereby enhancing overall system throughput.

3 The Threat Model

In this section, we first provide an informal description of the privacy threat addressed in this work. That is, we give more details on the tracking scenario presented in the introduction. Then, to formally model this threat, we introduce a new cryptographic privacy definition, which we call "timed-sequence-indistinguishability." Achieving this new privacy definition indicates that the threat has been effectively mitigated.

3.1 Informal Description of the Threat

Adversary $\mathcal{A}$ is provided with $m \geq 2$ beacons for a short period of a few hours, and can control their initialization times, that is, the times in which each beacon broadcasts its first EID. After initializing these beacons, adversary $\mathcal{A}$ monitors their EIDs and broadcasting times up to time t. That is, the adversary is provided with $m \geq 2$ sequences of EIDs and their broadcasting with each sequence corresponding to one beacon, while the adversary knows which sequence is associated with which beacon.

At time t, the adversary loses contact with these beacons for a period of T_{off}, which could range from a few hours to a day or two. During this time, the beacons are transported from one location to another, and therefore the adversary does not hear their EIDs. This interval is referred to as "the adversary's off-time period."

Once the contact is back, from time $t + T_{\text{off}}$ adversary $\mathcal{A}$ again monitors their EIDs and broadcasting times. That is, the adversary is provided with m new sequences of EIDs and their broadcasting times, starting from time $t + T_{\text{off}}$. These new sequences correspond to the same m beacons but are presented in a random order, with the adversary not being informed which sequence is related to which beacon. The goal of adversary $\mathcal{A}$ is to identify the special beacon among the m beacons.

Notice that scenarios where the adversary does not lose contact with the beacons at all or loses contact for only a short period so that the special beacon has not yet changed its EID, are not considered interesting. In these cases, the adversary would be able to identify the special beacon with 100% certainty based on the EIDs alone. Thus, the interesting question is whether the adversary can identify the special beacon after a sufficiently long off-time period during which the ephemeral identities of all beacons have changed. Additionally, from a practical perspective, allowing some time for the adversary to lose contact with the beacons is beneficial. This helps ensure that any potential side-channel information associated with the beacons (beyond our model), such as their exact location in a warehouse, becomes unavailable. Therefore, the off-time makes the tracking problem more intriguing, as the adversary has less information on which to base its decision.

Also notice that since the transportation of the beacons takes a few hours or a day or two, the off-time is assumed to be similarly long. During this short off-time period, the clock's drift of each beacon adds only a few seconds. This drift is negligible compared to the 10-minute delay in the identity-change times of the special beacon. For example, a 20 ppm drift would result in about 2 seconds per day, which is negligible compared to the 10-minute delay and can therefore be ignored.

3.2 Formally Modeling the Threat

In this section we formally define our threat model via a cryptographic game. We refer to this new privacy notion as the "timed-sequence-indistinguishability for beacons" (Definition 4). For comparison, we also formally define the well-known indistinguishability game which is a weaker threat model considering the EIDs only (Definition 2).

To prove that the countermeasure achieves timed-sequence- indistinguishability, we define another intermediate definition (Definition 3) similar to timed-sequence-indistinguishability for beacons but this time for truly random beacons. That is, a beacon in Definition 3 is assumed to follow the quasi-periodic time-scheduling algorithm while it broadcasts truly random identities instead of pseudorandom identities.

Recall that m indicates the number of beacons in the game. We define the following notations: For any $\ell \geq 1$ and $i \in [1, m]$, let $\mathsf{EID}_{k_i}(\ell)$ denote the ℓ'th EID of the i'th beacon whose key is k_i and let $s_{k_i}(\ell)$ be the time at which $\mathsf{EID}_{k_i}(\ell)$ was broadcast for the first time. That is, $s_{k_i}(\ell)$ is the time when $\mathsf{EID}_{k_i}(\ell - 1)$ changes to $\mathsf{EID}_{k_i}(\ell)$, with $s_{k_i}(1)$ representing the first time $\mathsf{EID}_{k_i}(1)$ was broadcast. Thus, $s_{k_i}(1)$ indicates the initialization time of the i'th beacon. The values $s_{k_i}(\ell)$ are determined by the time-scheduling of the beacons which is the same for all m beacons. In the notations, above we assume that $\mathsf{EID}_{k_i}(\ell)$ are pseudorandom. We now define equivalent notations for truly random ephemeral identities: For any $\ell \geq 1$ and $i \in [1, m]$, let $\mathsf{R}_i(\ell)$ denote the ℓ'th truly random EID of the i'th truly random beacon, and let $s_i(\ell)$ be the time at which $\mathsf{R}_i(\ell)$ was broadcast for the first time. That is, $s_i(\ell)$ is the time when $\mathsf{R}_i(\ell - 1)$ changes to $\mathsf{R}_i(\ell)$.

We assume that the size of beacon's life-time is polynomial in the security parameter of the cryptography used, denoted by n. Therefore any sequence of ephemeral identities and their broadcasting times is polynomial too. In the rest of the paper, by "polynomial" we mean polynomial in the security parameter n unless said otherwise. In the following definitions we provide the adversary with even more power and let him choose the special beacon and the time in which it loses contact with the beacons.

Definition 2. *(Indistinguishability of Beacons) Let $\mathcal{A}$ be a Probabilistic Polynomial Time (PPT) adversary in the following game:*

– *Adversary $\mathcal{A}$ chooses a special beacon $j \in [1,m]$.*
– *The challenger chooses a random permutation π. For any $i \in [1,m]$, it chooses a random key k_i of size n, and gives to adversary $\mathcal{A}$ the following two sequences, each of a polynomial-size $f(n)$:*
  - $S_i = \{\ \mathsf{EID}_{k_i}(\ell)|$ *for* $1 \le \ell < f(n)\ \}$, *and*
  - $E_i = \{\ \mathsf{EID}_{k_{\pi(i)}}(\ell)|$ *for* $f(n) \le \ell < 2f(n)\ \}$.
– *Adversary $\mathcal{A}$ guesses j' and wins if $j = \pi(j')$, namely if it chooses $E_{j'}$, the real continuation of S_j.*

We say that the beacons are indistinguishable if for any PPT adversary $\mathcal{A}$ (as above), for any j, π and for any inverse polynomial function γ, there exists n_0 such that for any $n > n_0$ it holds

$$\Pr[\mathcal{A} \textit{ wins}] \le \frac{1}{m}\big(1 + \gamma(n_0)\big).$$

Definition 3. *(Timed-Sequence-Indistinguishability of Truly Random Beacons) Let $\mathcal{A}$ be a PPT adversary in the following game:*

– *Adversary $\mathcal{A}$ chooses a special beacon $j \in [1,m]$, sets $s_i(1)$ for every $i \in [1,m]$, and chooses a point $t > \max_{i\in[1,m]} s_i(1)$ in which it starts losing contact with the m beacons. It then sends $j, t, s_i(1)$ for any $i \in [1,m]$ to the challenger.*
– *The challenger chooses a random permeation π. For any $i \in [1,m]$, it chooses a random key k_i of size n and gives to adversary $\mathcal{A}$ the following two polynomial-size sequences:*
  - $S_i = \{\ (\mathsf{R}_i(\ell), s_i(\ell))|$ *for* ℓ *s.t.* $s_i(\ell) < t\ \}$,
    where $s_i(1)$ is determined by the adversary, and
  - $E_i = \{\ (\mathsf{R}_{\pi(i)}(\ell), s_{\pi(i)}(\ell))|$ *for* ℓ *s.t.* $s_{\pi(i)}(\ell) > t + T_{off}\ \}$.
– *Adversary $\mathcal{A}$ guesses j' and wins if $j = \pi(j')$, namely if it chooses $E_{j'}$, the real continuation of S_j.*

We say that the truly random beacons are timed-sequence- indistinguishable if for any PPT adversary $\mathcal{A}$ (as above), for any $j, \pi, t, s_i(1)$, there exists a decreasing function $\delta(r) \xrightarrow{r\to\infty} 0$, and there exists T_0 which is long enough to ensure an identity change such that for any $T_{off} > T_0$ it holds that

$$\Pr[\mathcal{A} \textit{ wins}] \le \frac{1}{m}\big(1 + \delta(T_0)\big).$$

Definition 4. *(Timed-Sequence-Indistinguishability of Beacons) Let $\mathcal{A}$ be a PPT adversary in the following game:*

– *Adversary $\mathcal{A}$ chooses a special beacon $j \in [1,m]$, sets $s_{k_i}(1)$ for every $i \in [1,m]$ and chooses a time $t > \max_{i\in[1,m]} s_{k_i}(1)$ in which it starts losing contact with the m beacons. It then sends $j, t, s_{k_i}(1)$ for any $i \in [1,m]$ to the challenger.*
– *The challenger chooses a random permutation π. For any $i \in [1,m]$, it chooses a random key k_i of size n and gives to adversary $\mathcal{A}$ the following two polynomial-size sequences:*
  - $S_i = \{\ (\mathsf{EID}_{k_i}(\ell), s_{k_i}(\ell))|$ *for* ℓ *s.t.* $s_{k_i}(\ell) < t\ \}$
    where $s_{k_i}(1)$ is determined by the adversary.
  - $E_i = \{\ (\mathsf{EID}_{k_{\pi(i)}}(\ell), s_{k_{\pi(i)}}(\ell))|$ *for* ℓ *s.t.* $s_{k_{\pi(i)}}(\ell) > t + T_{off}\}$.
– *Adversary $\mathcal{A}$ guesses j' and wins if $j = \pi(j')$, namely if it chooses $E_{j'}$, the real continuation of S_j.*

We say that the beacons are timed-sequence-indistinguishable if for any PPT adversary $\mathcal{A}$ (as above), for any $j, \pi, t, s_{k_i}(1)$, and for any γ inverse polynomial function in n, there exists a decreasing function $\delta(r) \xrightarrow{r\to\infty} 0$, and there exists T_0 which is long enough to ensure an identity change such that for any $T_{off} > T_0$ and for any security parameter $n > n_0$ it holds that

$$\Pr[\mathcal{A} \textit{ wins}] \le \frac{1}{m}\big(1 + \gamma(n_0) + \delta(T_0)\big).$$

Important notes regarding the definitions Notice that in Definitions 3 and 4 we say that timed-sequence- indistinguishability is achieved if for *any initialization times* ($s_i(1)$ or $s_{k_i}(1)$ correspondingly), the winning probability is *independent* of the initialization times, the special beacon, the random permutation, and the exact time the adversary loses contact with the beacons. That is, the winning probability is *only dependent* of the off-time period. Therefore:

- By proving that beacons achieve timed-sequence- indistinguishability, we actually prove that their time-scheduling mitigates any attack that involves controlling the initialization times. That is, for any value of $s_i(1)$ where $i \in [1, m]$, the winning probability is $1/m(1+\delta(T_0))$ (and similarly, $1/m(1+\gamma(n_0)+\delta(T_0))$ for any value of $s_{k_i}(1)$).
- To prove that beacons do not achieve timed-sequence- indistinguishability, we need to establish two things: (1) the existence of specific adversarial initialization times, $s_{k_j}(1)$ or $s_j(1)$, for which the beacons' time-scheduling fails to obscure information, thereby resulting in a significantly higher winning probability than $1/m$, and (2) the ability of the adversary to execute a physical attack on deployed beacons to achieve these adversarial initialization times.

4 The Timer Manipulation Attack

In this section we demonstrate that beacons with T-periodic time-scheduling fail to achieve time-sequence-indistinguishability. First, in Lemma 1 we prove the existence of initialization times that allow the adversary to win with a probability of 1. We then show a specific physical attack on deployed beacons in the real-world that enables the adversary to achieve these adversarial initialization times.

Lemma 1. *Beacons with T-periodic time scheduling do not achieve timed-sequence- indistinguishability.*

Proof. We show that there exist initialization times allowing the adversary to win with a probability of 1. For example, adversary $\mathcal{A}$ in Definition 4 can set $s_{k_i}(1) = v$ for any beacon $i \neq j$ and $s_{k_j}(1) = v+10$ for the special beacon j. Since the time-scheduling is T-periodic, for any $i \in [1, m]$ and $\ell > 0$ such that $s_{k_i}(\ell) > t + T_{\text{off}}$, it follows that $s_{k_i}(\ell) \mod 15 = v$ for $i \neq j$ and $s_{k_j}(\ell) \mod 15 = v + 10$. Since adversary $\mathcal{A}$ receives only a permutation of the sequences, and assuming each sequence i is represented by $S_i = \{(x_r^i, y_r^i)\}_{r>0}$, the adversary returns j' such that for any $r > 0$ it holds $y_r^{j'} \mod 15 = v + 10$, therefore wins with a probability of 1.

We show that the above adversarial initialization times can be achieved with real-world deployed beacons. To illustrate this, we conduct a physical attack, which we call the "Timer Manipulation Attack," on Samsung's SmartTags [24], which is currently the only beacon known to change its ephemeral identities when far from its owner. We show how an adversary can easily set the initialization times of SmartTags by inserting/reinserting their batteries.

We begin by describing our experiment with a single SmartTag beacon. In this experiment, we used the Thingsup BLE Beacon Scanner [21] to read the MAC addresses and broadcasting times of Samsung's SmartTags. Since each change in the EID forces a change in the MAC address, we interpret the times at which the MAC addresses change as the times when the EID changes.

We first initialized this beacon at time 12:21, and we saw that the MACs (hence the EIDs) changed every 15 minutes, namely at times 12:36 and 12:51. See the MAC-change times depicted in Figure 2 while the first MAC depicted by number 1, the second by number 2, etc. At 19:43 we reinserted the battery of the beacon. We recorded its broadcasts and saw that the MAC of the beacon, after the reinsertion, changes at times 19:58 and 20:13. See the MAC-change times depicted in Figure 3. We repeated this experiment several times, each time the MAC-change time was adjusted to the battery reinsertion time.

This experiment leads to two key conclusions:

1. The times at which a beacon's EIDs change are determined by its initialization time. Specifically, if the beacon is initialized at time v, then EIDs will change every 15 minutes starting from v, i.e., $v, v+15, v+30$ and so on.

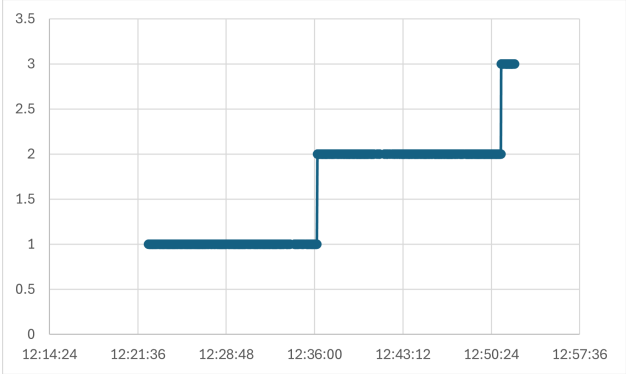

Fig. 2. The MAC (EID) change-times with the beacon initialized at 12:21. The x-axis is the time in format hour:min:sec and the y-axis shows the MAC addresses, labeled sequentially (1 for the first MAC, 2 for the second, etc.)

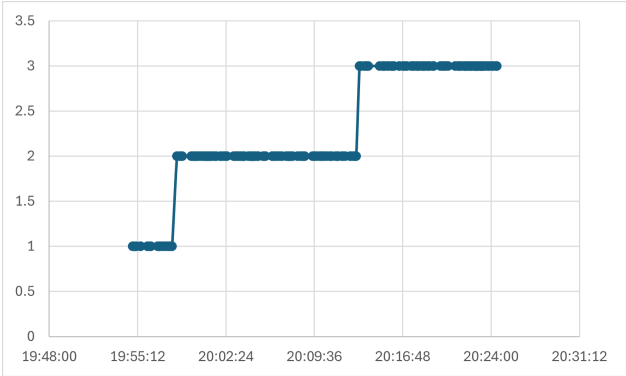

Fig. 3. The MAC (EID) change-times given the battery is reinserted at time 19:43. The x-axis represents time (hour:min:sec), and the y-axis shows MAC addresses, numbered sequentially.

2. Inserting or reinserting the battery is a very simple and straightforward way to initialize a beacon.

Building on these conclusions, we now introduce the "Timer Manipulation Attack," which allows an adversary in our threat model to successfully identify the special beacon based on the times when the beacons' batteries are inserted and based on their broadcasting times. We illustrate this attack with the real-world example outlined in the introduction, where a special box is tracked along the supply chain: EveA knows which box holds the special item. Therefore when packaging the boxes and attaching their beacons, EveA inserts the battery of each beacon attached to a non-special box at time v, and inserts the battery of the beacon attached to the special box with a delay of 10 minutes, namely at time $v + 10$.

Note that the number of boxes arriving at Bob's company is usually not enormous. However, even if there are many boxes such that it becomes difficult for EveA to insert the batteries of all beacons at approximately the same time, then EveA can divide the non-special boxes into small groups. Then EveA can insert the batteries of each group at "effectively" the same time as follows: in the first group, all batteries are inserted at time v, in the second group all batteries are inserted at time $v + 15$, in the third group at time $v + 30$ and so one. Since the beacons change their EIDs every 15 minutes, all beacons in each group will change their EIDs at approximately the same time $v + 15 \cdot x$ for $x \geq 0$. Only the special box will change its EIDs at time $v + 10 + 15 \cdot x$. The 10-minute delay is significantly longer than the minor differences in the battery insertion times of the different groups, which are measured in seconds.

When the boxes arrive at Bob's company, EveB hears the beacons' EIDs and their broadcasting times. Given an EID-change at time z, EveB can determine whether the beacon is attached to the special box

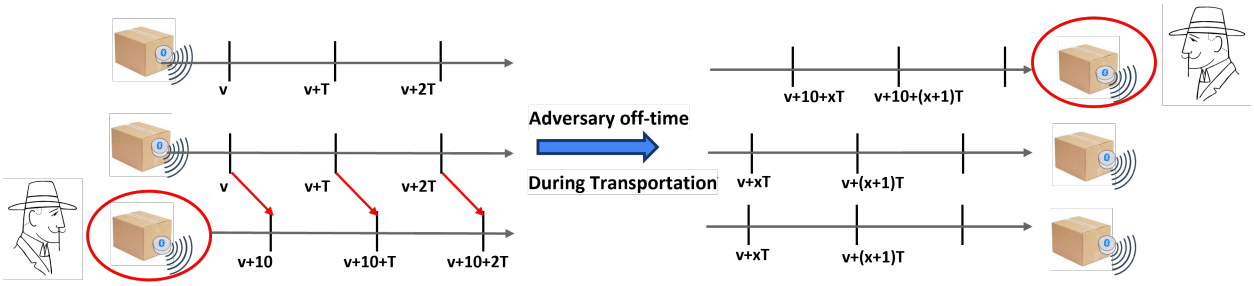

Fig. 4. (Left) The boxes at Alice's company, (Right) the boxes at Bob's company: The attack violates privacy in our scenario.

or not (despite the EIDs being pseudorandom) as follows: if ($z \mod 15$) == ($v \mod 15$), then the EID is not associated with the special box, otherwise if ($z \mod 15$) == ($v + 10 \mod 15$) it is associated with the special box. Hence, due to this delay in the identity-change times, identifying the beacon attached to the special box becomes straightforward.

As shown in [24], the entire pool of privacy IDs of Samsung's SmartTag can be collected within a few days. However, in our scenario, EveA and EveB are not required to hear the beacons' broadcasts for a long time but only for a couple of hours or even less. Consequently, the probability of encountering a repeated EID within this short period is low, meaning the chance of identifying the special beacon in our threat model based on a repeated EID is small. In contrast, as we demonstrated above, the probability of identifying the special beacon based on the timing of EID-changes is 100%.

5 Countermeasure: Quasi-Periodic Time-Scheduling

Deterministically changing identities at regular intervals of T, as is done in T-periodic time-scheduling, preserves the initialization times and therefore maintains the 10-minute delay.

To mitigate the "Timer Manipulation Attack" and any other attacks that exploit adversarial initialization times to gain an advantage, we need a countermeasure: a new time-scheduling method that obscures the adversary's initialization times without complicating the system or degrading its performance. Specifically, this new time-scheduling should ensure that for any given initialization times, the winning probability remains $1/m$ with only a negligible error that depends solely on the off-time period and not on the initialization times themselves. This approach will ensure that beacons using our proposed time-scheduling method achieve timed-sequence- indistinguishability.

To achieve this, we introduce randomization and relax the periodicity constraint, resulting in a probabilistic quasi-periodic time-scheduling algorithm. This algorithm approximates the periodicity requirement by allowing the time between any two consecutive EID-changes to vary between a and $b > a$, rather than being fixed at T, where a, b are close to T and $a \leq T \leq b$. We refer to this algorithm as the "$[a, b]$-periodic time-scheduling," see Figure 5 (top). In Section 7 we will discuss how to choose a and b to ensure that our new quasi-periodic time-scheduling algorithm provides sufficient randomization to obscure the delay, thus achieving privacy, while maintaining performance close to that of the original periodic scheduling.

Figure 6 illustrates the intuition behind our countermeasure. Our new quasi-periodic time-scheduling for beacons ensures that, regardless of the initialization times used, the adversary in our threat model cannot identify the special beacon. A formal proof of this is provided in Section 6.

Formally, $[a, b]$-periodic time-scheduling is defined as follows:

Definition 5. *($[a, b]$-Periodic Time-scheduling) We say that a time-scheduling algorithm is "$[a, b]$-Periodic" for $a < b$ if the time-interval between any two consecutive EID-changes of a beacon is an i.i.d (independent and identically distributed) random variable uniformly distributed in the range $U[a, b]$.*

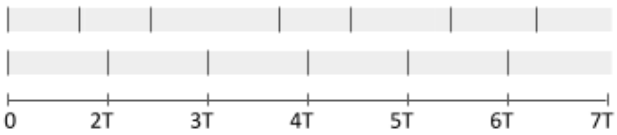

Fig. 5. (Top) $[a, b]$-periodic; (Bottom) T-periodic

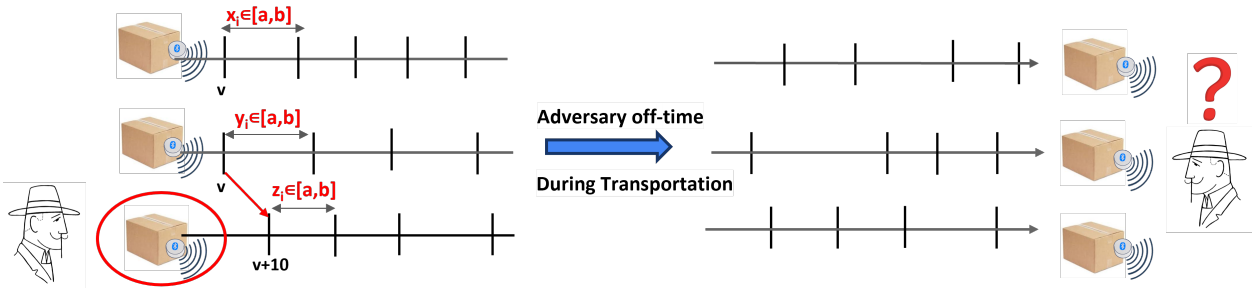

Fig. 6. (Left) The boxes at Alice's company, (Right) the boxes at Bob's company: Our countermeasure mitigates the attack.

Our countermeasure is straightforward to implement. Since our $[a, b]$-periodic time-interval is no longer fixed, it only requires generating a new random value to determine the next time-interval. For $[a, b]$ values close to T, this additional step in the algorithm compared to the original T-periodic time-scheduling is negligible, as discussed in Section 7.

6 The Privacy Proof

In this section we prove our main theorem which shows that beacons with $[a, b]$-periodic time-schedule achieve the timed-sequence -indistinguishability. Recall that the ephemeral identities are the results of a pseudorandom function. We will start by proving timed-sequence- indistinguishability for truly random beacons which broadcast random identities according to an $[a, b]$-periodic time-schedule. The privacy for real beacons with pseudorandom identities will then follows from the pseudorandomness of the EID values by by-now standard indistinguisability arguments.

Theorem 1. *Truly random beacons (i.e., beacons which use a truly random source for assigning random fields) with $[a, b]$-periodic time-scheduling achieve the timed-sequence indistinguishability according to Definition 3.*

Proof. Let j be the special beacon, t be the time at which the adversary loses contact with the m beacons, π be a random permutation, and $s_i(1)$ be the initialization time of beacon $i \in [1, m]$. According to Definition 3, j' is the adversary's guess, and the winning probability is the probability that the adversary correctly guesses j' which is $\pi^{-1}(j)$ out of all possible options for $j' \in [1, m]$. Since the identities are truly random, each sequence of identities is equally likely to be generated by any beacon. Therefore we can ignore the random values of the broadcasts and focus instead on the times of the broadcasts.

In a $[a, b]$-periodic time-scheduling scheme, each EID-change time depends solely on the previous EID-change time and on a random time-interval, which is uniformly chosen from $\in [a, b]$. Therefore, to determine which sequence E_i among $i \in [1, m]$ is the continuation of S_j, we only need to examine the first EID-change in each sequence E_i (where $i \in [1, m]$), and calculate the probability that this EID-change continues from the last EID-change in sequence S_j where the first EID-change at S_j starts at time $s_j(1)$. See Figure 7.

To this end, we define the following notations: for any $i \in [1, m]$, let t_i denote the time of the first EID-change in sequence E_i, and let t' denote the time of the last EID-change in sequence S_j. With these

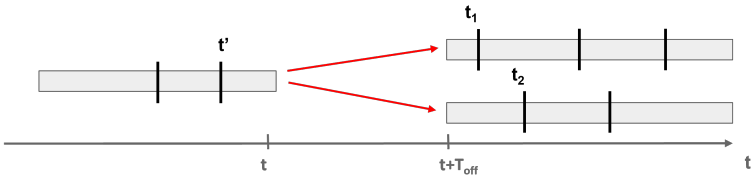

Fig. 7. Given the last EID-change is at time t', the adversary's goal is to determine which scenario has a higher probability: to have an EID-change at t_1 or t_2.

notations, the probability that sequence E_i is a continuation of sequence S_j is the probability of having an EID-change at time t_i given that there was an EID-change at time t'. Therefore, for any $i \in [1, m]$, let $\Pr[t_i|t']$ be the probability of having an EID-change at time t_i given that there was an EID-change at time t', where each time-interval between any two consecutive EID-changes within the range $[t', t_i]$ is chosen randomly and uniformly. Using this notation, the winning probability is

$$\Pr[\mathcal{A} \text{ wins}] = \frac{\Pr[t_{\pi^{-1}(j)}|t']}{\sum_{i=1}^{m} \Pr[t_i|t']}.$$

The probability $\Pr[t_i|t']$ does not depend on the absolute values of t' and t_i, but only on the difference between them. That is, $\Pr[t_i|t']$ represents the probability that a sum of random variables, each uniformly distributed in the range $[a, b]$, equals $t_i - t'$. For simplicity, we introduce a new notation, $p[r]$, which depends only of the difference r. That is, $p[r]$ is the probability of having an EID-change at time $x + r$ given that there was an EID-change at time x, for any $x \geq 0$. Therefore

$$p[t_i - t'] := \Pr[t_i|t'].$$

More specifically, let $p[r, k]$ denote the probability that a sum of exactly k random variables, each uniformly distributed in the range $[a, b]$, equals $r > 0$. Thus we get

$$p[r] = \sum_{k=-\infty}^{\infty} p[r, k]. \tag{1}$$

Using the $p[r]$ notation, the winning probability is

$$\Pr[\mathcal{A} \text{ wins}] = \frac{p[t_{\pi^{-1}(j)} - t']}{\sum_{i=1}^{m} p[t_i - t']}.$$

To prove Theorem 1, we need to prove that the winning probability converges to $1/m$ as the off-time increases. To this end, in Lemma 2 we prove that $p[r]$ converges to $1/\mu$ as r increases, where μ is the expected value of the uniform distribution $U[a, b]$. That is,

$$p[r] \xrightarrow{r \to \infty} 1/\mu.$$

Since the off-time period starts at t where $t > t'$ and ends at $t + T_{\text{off}}$ where $t + T_{\text{off}} < t_i$ it holds that $t_i - t' > T_{\text{off}}$ for any $i \in [1, m]$ (see Figure 7). Therefore the winning probability converges to $1/m$ as the off-time increases, that is:

$$\Pr[\mathcal{A} \text{ wins}] = \frac{p[t_{\pi^{-1}(j)} - t']}{\sum_{i=1}^{m} p[t_i - t']} \xrightarrow{T_{\text{off}} \to \infty} \frac{1/\mu}{\sum_{i=1}^{m} 1/\mu} = \frac{1}{m}. \tag{2}$$

That is, the winning probability converges to $1/m$ regardless of t, j, π and the initialization times $s_i(1)$. Hence, truly random beacons using our $[a, b]$-periodic time-schedule achieve timed-sequence- indistinguishability.

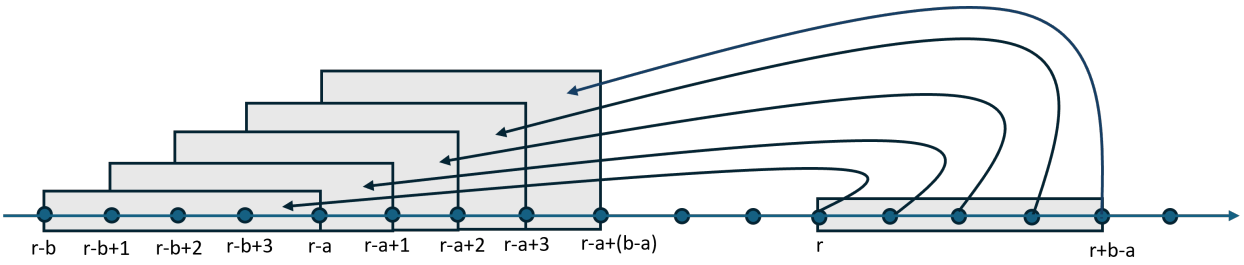

Fig. 8. The value $p[r]$ depends on range $[p[r-b], ..., p[r-a]]$, the value of $p[r+1]$ depends on range $[p[r-b+1], ..., p[r-a+1]]$, and so on.

Lemma 2. *$p[r]$ converges to $1/\mu$ as r increases, where μ is the expected value of the uniform distribution $U[a,b]$. That is,*

$$p[r] \xrightarrow{r\to\infty} 1/\mu.$$

Proof. Recall from Equation 1 that $p[r]$ is sum of $p[r,k]$ for any k. Therefore to calculate $p[r]$ we need first to calculate $p[r,k]$. We define $p[r,k]$ recursively: $p[0,0]=1$, $p[0,k]=0$ for any $k\neq 0$, and $p[r,k]=0$ for any $r<0$. For any $r>0$, we get

$$p[r,k] = \frac{1}{b-a+1}\sum_{i=a}^{b} p[r-i,k-1]. \tag{3}$$

We substitute Equation 3 in Equation 1 and get for any $r>0$

$$p[r] = \frac{1}{b-a+1}\sum_{i=a}^{b} p[r-i] \tag{4}$$

where $p[0]=1$ and $p[r]=0$ for any $r<0$.

Since $p[r]$ is the average of all $p[r-i]$ for any $i\in[a,b]$, as long as not all the elements in the sum are identical, $p[r]$ is smaller than the maximum value of $p[r-i]$ and greater than the minimum value. Specifically,

$$\min_{i\in[a,b]} p[r-i] < p[r] < \max_{i\in[a,b]} p[r-i].$$

This proves convergence.

Now we prove that the convergence value is $1/\mu$. To do this, we sum both sides of Equation 4 for any $r\in[1,R]$ with a chosen R.

$$\sum_{r=1}^{R}(b-a+1)p[r] = \sum_{r=1}^{R}\sum_{i=a}^{b} p[r-i].$$

We express the right-hand side of the equation as a single summation rather than a double summation, while ignoring all $p[i]$ for $i<0$ since their value is zero. We therefore get

$$\sum_{i=1}^{R}(b-a+1)p[i] = \sum_{i=0}^{R-b}(b-a+1)p[i] + \sum_{i=1}^{b-a}(b-a+1-i)p[R-b+i].$$

After eliminating terms that appear on both sides, we get

$$\sum_{i=1}^{b-a} i\cdot p[R-b+i] + \sum_{i=1}^{a}(b-a+1)p[R-a+i] = b-a+1.$$

Let L denote the convergence value. As R approaches infinity, all $p[R-i]$ for $i \in [a,b]$ converge to L. We therefore get

$$\sum_{i=1}^{b-a} i \cdot L + \sum_{i=1}^{a} (b-a+1)L = b-a+1.$$

The convergence value L is therefore

$$L = \frac{2}{a+b} = \frac{1}{\mu}.$$

7 Time-Sequence-Indistinguishability Vs. Performance Parameters Trade-off

An interesting question is how to choose (a,b) to achieve a good balance between timed-sequence- indistinguishability, low power consumption, and short-lived stable EID.

To achieve timed-sequence- indistinguishability after a relatively short off-time, we need to choose (a,b) such that $p[r]$ converges relatively quickly to $1/\mu$. The rate of convergence of $p[r]$ determines the duration for which the adversary needs to be off to avoid distinguishing the beacon. As defined in Equation 4 and as illustrated in Figure 8, the value $p[r+b]$ depends on the elements at indices $[r, r+b-a]$. These elements, in turn, depend on the elements at indices $[r-b, r+b-2a]$. To reduce the convergence rate, we need: (1) a to be relatively small so that the index $r+b$ will be close to indices $[r, r+b-a]$, and (2) $b = 2a$ so that the indices $[r, r+b-a]$ will be consecutive to the indices $[r-b, r+b-2a]$, namely $r = r+b-2a$. Therefore, we ask $b = 2a$.

To achieve good balance between performance and privacy, we would like a and b to be close to the original T, thereby achieving performance and privacy similar to that of a T-periodic schedule. Specifically,

- Performance: Low power consumption. In the worst case, each EID is changed every a time-units instead of T time-units. Consequently, the beacon should choose a pseudorandom value in range $[a,b]$ every a time-units. Keeping a close to T will minimize overhead for both the beacon and the system as a whole. Note that, in addition to generating a new EID occasionally and broadcasting an EID every second as in the T-periodic time-scheduling, the beacon in the $[a,b]$-periodic time-scheduling should also generate a new random choice from $[a,b]$ every time an EID is generated. However, this additional operation is negligible compared to the beacon's computation.
- Privacy: Short-lived stable EID. In the worst case, the EID remains stable for b time-units instead of T time-units. Therefore keeping b close to T will minimize this period.

To conclude: to achieve performance and privacy as close as possible to the original T-periodic, while achieving indistinguishably after a short off-time, b should be set to $2a$, and both a and b should be relatively small and close to T. The optimal trade-off is achieved with $a \leq T \leq b$ where $b = 2a$.

8 Simulation Studies

We proved above that beacons with quasi-periodic time-schedule achieve time-sequence- indistinguishability. That is, we showed that there exists T_0 such that for any off-time longer than T_0, indistinguishability is mitigated.

In this section, we show that, actually, for concrete real-world parameters this T_0 is relatively short. To this end, we first implemented $p[r]$ calculations and analyzed it for various concrete values of a and b where $b = 2a$. The common choice for T in existing beacons that use T-periodic time-scheduling is between 15 and 20 minute, therefore we selected values for $[a,b]$ within this range. The results are summarised in Figure 9. In this simulation study we assume that the time-unit is one second. Namely, if for example $a = 10$ minutes and $b = 20$ minutes, then $x \in [a,b]$ represents a random second chosen uniformly within the interval $[600, 1200]$ seconds.

The blue line in Figure 9 represents $p[r]$ with the x-axis corresponding to r. As illustrated in the figure, $p[r]$ converges to $1/\mu = 2/(a+b)$ as r increases, and it converges rapidly. To demonstrate it, we added

two reference lines: (1) the red line which indicates $(1+0.05)/\mu$ and (2) the purple line which indicates $(1-0.05)/\mu$. For $a=10, b=20$ minutes, it is observed that after approximately one hour, $p[r]$ is close to $1/\mu$ up to $0.05/\mu$, namely the blue line of $p[r]$ is in between the reference lines.

Another conclusion from the simulations, which aligns with the discussion above, is that the convergence rate of $p[r]$ is faster for smaller a. For instance, when comparing different values of $[a, 2a]$, the convergence rate of $a=10$ minutes is faster than that for $a=12$ minutes.

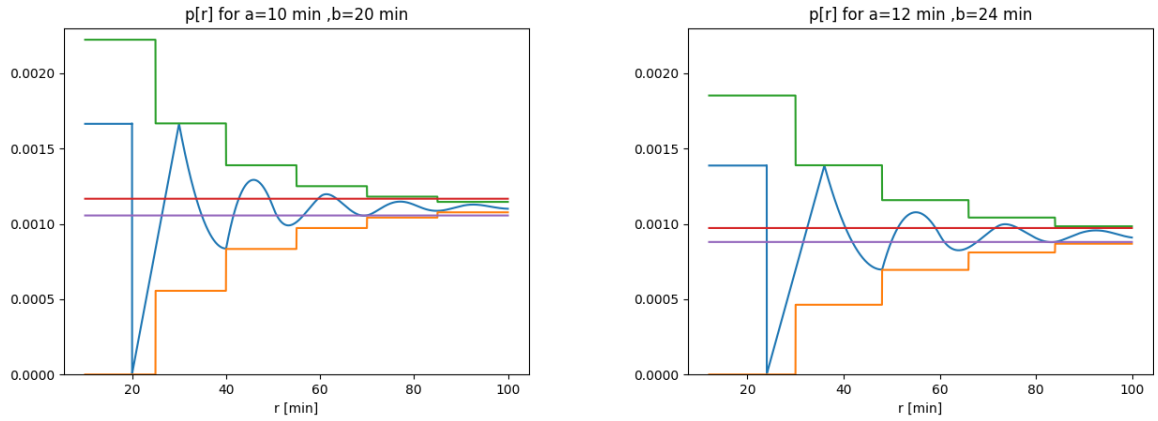

Fig. 9. For different values of $[a, b]$ where $b=2a$: (blue) $p[r]$; (green) upper bound; (orange) lower bound; (red) $(1+0.05)/\mu$; (purple) $(1-0.05)\mu$.

Next, we investigate the convergence rate of the winning probability as a function of the off-time. To achieve this, we first investigate the convergence rate of $p[r]$ and establish both upper and lower bounds for $p[r]$. Based on our simulation, we show that $p[r]$ converges fast. Specifically, we show that the difference between $p[r]$ and the convergence value $1/\mu$ decreases by a factor of 2 every $1.5a$ starting at $r=a$, that is

$$|p[r]-\frac{1}{\mu}| \leq \frac{1}{\mu}\cdot\left(\frac{1}{2}\right)^{\lfloor\frac{r-a}{1.5a}\rfloor}. \tag{5}$$

After an off-time of T_{off}, since it holds that $t_i - t' > T_{\text{off}}$ for any $i \in [1,m]$, we get

$$1/\mu(1-\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor}) \leq p[t_i-t'] \leq 1/\mu(1+\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor}).$$

Hence

$$\begin{aligned}\Pr[\mathcal{A}\text{ wins}] = \frac{p[t_{\pi^{-1}(j)}-t']}{\sum_{i=1}^{m} p[t_i-t']} &\leq \frac{1/\mu(1+\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor})}{\sum_{i=1}^{m} 1/\mu(1-\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor})} \\ &= \frac{1}{m}\left(1+\frac{2\cdot\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor}}{1-\frac{1}{2}^{\lfloor\frac{T_{\text{off}}-a}{1.5a}\rfloor}}\right).\end{aligned} \tag{6}$$

In Figure 10, the green and the orange lines represent the upper and the bound (respectively) as defined in Equation 5. As can be seen, these bounds are very tight. For example, for $a=10, b=20$ minutes, $p[r]$ approaches $1/\mu$ within $0.05/\mu$ after one hour, while the upper and lower bounds approach $1/\mu$ within $0.05/\mu$ after approximately 1 hour and 10 minutes.

Figure 10 shows the winning probability as a function of the off-time period T_{off} as illustrated in Equation 6. We show this for $m = 4$ and $m = 100$, using the same values of a as before: $a = 10$ minutes and $a = 12$ minutes. The x-axis of Figure 10 is the off-time starting from $T_{\text{off}} = 3a$, since the lower bound of $p[r]$ is zero in the range $[a, 2.5a]$. As shown, for $m = 4$ the winning probability is approximately $1/m = 0.25$ after 2 hours for $a = 10, b = 20$ minutes and 2.5 hours for $a = 12, b = 24$ minutes. This indicates that an adversary would need only a relatively short off-time to avoid winning the game with a probability significantly higher than $1/m$. In real-life scenarios, transportation of physical goods typically takes a few hours to a day or two. Therefore, guaranteeing that the adversary cannot win after 2.5 hours is quite practical.

In addition, as expected, since the multiplicative error depends only on T_{off} and not on m, both the winning probability for $m = 4$ and the winning probability for $m = 100$ get to $1/m(1 + \lambda)$ for some λ, after the same off-time.

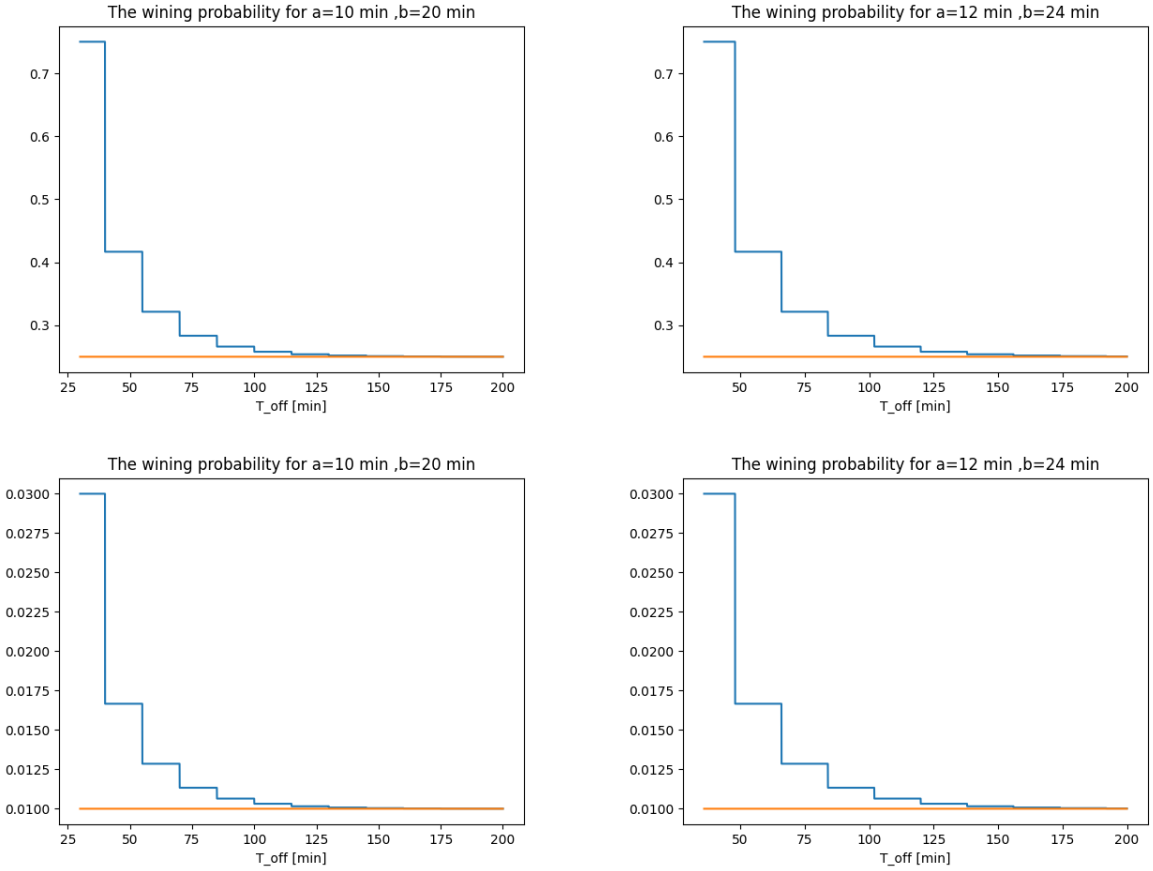

Fig. 10. (Blue) An upper bound on the winning probability for $b = 2a$ as a function of the off-time T_{off}; (orange) $1/\mu$. The top two figures are for $m = 4$ and the bottom two figures are for $m = 100$.

9 Integrating Our Quasi-Periodic in Real BLE Location-Tracking Systems

Next, we concentrate on the integration of the countermeasure in Samsung system in light of the transition from T-periodic schedule to $[a, b]$-periodic schedule.

Any SmartTag's identity is chosen from a pool of identities which are unique for this SmartTag. The integration of the new time-scheduling is therefore trivial with almost no cost at all: the SmartTag's pool of identities remains the same, but instead of choosing a new identity from the pool after T units of time, it first chooses $x \in [a, b]$ and then chooses a new identity from the pool after x units of time.

10 Related Work

To the best we know, there is no work in the literature which explicitly focuses on the indistinguishability of periodic broadcasting beacon devices (or other periodic broadcasting methods, dealing with messages that are pseudorandom and ephemeral) based on the time-scheduling of the broadcasts and its implications to privacy. In particular, no one has pointed out such a timing attack in T-periodic broadcasting beacons.

Next we survey some seemingly related works that deal with time side channels which, in fact, are quite different from our investigation, both in terms of problems and solutions.

There is a substantial body of work in the literature addressing time-side channels from various perspectives, such as exploiting time-information for extracting secret keys [17,4] or for revealing the underlying distribution from which elements are sampled [3]. There is also literature focusing on the time-side channel of unlinkability or anonymity, however it only addresses non-periodic signals. For example, Tramer et. al. [22] deal with privacy-focused crypto-currencies, such as Zcash or Monero which aim to provide strong cryptographic guarantees for transaction confidentiality and unlinkability. They show that a remote party can link all transactions that send funds to a user, by measuring the response time of that user's P2P node to certain requests. Rohrer et. al. [20] focus on the Lightning Network which is a scaling solution for Bitcoin that promises to enable rapid and private payment processing. In Lightning, multi-hop payments are secured by utilizing Hashed Time-Locked Contracts (HTLCs) and encrypted on the network layer by an onion routing scheme to avoid information leakage to intermediate nodes. The work in [20] shows, however, that the privacy guarantees of the Lightning Network may be subverted by an on-path adversary conducting timing attacks on the HTLC state negotiation messages.

The above works are different from our case: they do not deal with periodic signaling, and what can be seen as the time-side-channel we deal with is completely different and not related to the above cases. In addition, the above works do not suggest any countermeasure.

There is an area of works in which time-side-channel is used for achieving *linkability*, however linkability (like authentication) is the exact opposite of our goal which is achieving *un-linkability* (i.e., privacy) despite the existing time-side-channel. In this line of work, Kohno et al. [18] recognized that clock skews are useful for remotely fingerprinting networked devices. They show that different devices have different skews and these skews are stable enough over time. As a result, numerous works were published using the clock-skew as a fingerprint, for example, Jana et al. [13] used clock skew to fingerprint wireless devices (the opposite of privacy) and to detect fake wireless access points. Similarly, Huang et al. [11] explore the use of clock skew of a wireless local area network access point (AP) as its fingerprint to detect unauthorized APs quickly and accurately. Other methods for fingerprinting broadcasting services are given in [9,12].

To ensure anonymity for Internet users, several low-latency anonymous networks (e.g., Onion Routing [19], Tor [8]) have been developed to obscure the identity and communication between parties. However, the need for low latency introduces vulnerabilities to timing attacks. These attacks exploit the timing correlation between the original and anonymized flows to establish a link between them. Currently, all practical low-latency anonymous networks are susceptible to timing attacks, as no such network can entirely eliminate the timing correlation between the original and anonymized traffic. A body of research in the literature aims to understand and quantify the negative impact of timing attacks on low-latency anonymous networks, as exemplified by studies such as [14,23].

During the COVID-19 pandemic, several BLE-based contact tracing systems were developed, with the Google-Apple Exposure Notification (GAEN) system [2,6] being a primary example. This system involves broadcasting periodic BLE signals to identify contacts with individuals who have tested positive for COVID-19. Since these BLE-based contact tracing systems are phone-based, the periodic signals are synchronized according to a global clock available on the phones.

11 Conclusion

This paper shows a new privacy attack against BLE beacons, specifically against Samsung's SmartTag. The attack enables an adversary to control the times in which the beacon's EIDs changes. Therefore, in many feasible scenarios, the attack enables the adversary to mark its beacon-of-interest by choosing a unique identity-change times to its beacon-of-interest.

To claim privacy against such attack and other timed-base privacy attacks, we propose a new privacy definition. We call this definition timed-sequence- indistinguishability. The new definition considers the broadcasts' times (in addition to the traditionally considered cryptographic pseudorandom content of advertisements) and therefore it is stronger than the well-known indistinguishability definition (a privacy property which directly translates to the cryptographic pseudorandomness of the functions deriving EIDs and encrypted values). We then propose a countermeasure and prove that using this countermeasure, the beacons are timed-sequence-indistinguishable. The goal of the definition is to be able to claim that any time manipulation in order to violate privacy is mitigated.

We further show how to integrate our countermeasure in Samsung's system efficiently (it is important when proposing a change to an established practice like a periodic EID changes, to assure its feasibility and efficiency and we follow this practice in our work). Finally, we present an extensive simulation study regarding actual potential parameters and evaluate how fast they achieve privacy; beyond the useful simulation, analytically analyzing convergence rate is an interesting open question.

We conclude by saying that, in fact, there are two ways to view our work: On the one hand, it treats an existing beacon systems (Samsung's SmartTags) and can be viewed as a privacy attack or at least a privacy comment on an existing Ephemeral ID system, and a way to protect against the identified attack and similar attacks. On the other hand, the work is, perhaps, the first to raise the feasibility of time based attacks in the BLE beacons broadcasting domain and therefore as a warning signal to future designers of such and similar systems to consider time of broadcasts in their privacy evaluation (hence our general definition which includes times of signal is a way to force designers to pay attention to the issue). Investigating further scenarios where time information violates privacy, and where mitigation like ours applies, are left as an open question for future investigations.

Acknowledgment

We thank David Lazarov and Omer Berkman for discussions on beacons.

References

1. Apple. FindMy, 2020. `https://www.apple.com/newsroom/2021/04/apples-find-my-network-now-offers-new-third-party-finding-experiences`.
2. Apple, and Google. Privacy-preserving contact tracing. `https://covid19.apple.com/contacttracing`.
3. Ben-Dov, Y., David, L., Naor, M., and Tzalik, E. Resistance to Timing Attacks for Sampling and Privacy Preserving Schemes. In *4th Symposium on Foundations of Responsible Computing (FORC 2023)* (Dagstuhl, Germany, 2023), K. Talwar, Ed., vol. 256 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 11:1–11:23.
4. Ben-Dov, Y., David, L., Naor, M., and Tzalik, E. Are your keys protected? time will tell. In *5th Conference on Information-Theoretic Cryptography, ITC 2024, August 14-16, 2024, Stanford, CA, USA* (2024), D. Aggarwal, Ed., vol. 304 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 3:1–3:28.
5. David, L., Berkman, O., Hassidim, A., Lazarov, D., Matias, Y., and Yung, M. Cryptiny: Compacting cryptography for space-restricted channels and its use-case for IoT-E2EE. Cryptology ePrint Archive, Paper 2024/1128, 2024. `https://eprint.iacr.org/2024/1128`.
6. David, L., Hassidim, A., Matias, Y., and Yung, M. Scaling up gaen pseudorandom processes: Preparing for a more extensive pandemic. In *European Symposium on Research in Computer Security* (2022), Springer, pp. 237–255.

7. David, L., Hassidim, A., Matias, Y., Yung, M., and Ziv, A. Eddystone-EID: Secure and private infrastructural protocol for ble beacons. *IEEE Transactions on Information Forensics and Security* (2022).
8. Dingledine, R., Mathewson, N., Syverson, P. F., et al. Tor: The second-generation onion router. In *USENIX security symposium* (2004), vol. 4, pp. 303–320.
9. Givehchian, H., Bhaskar, N., Herrera, E. R., Soto, H. R. L., Dameff, C., Bharadia, D., and Schulman, A. Evaluating physical-layer ble location tracking attacks on mobile devices. In *2022 IEEE symposium on security and privacy (SP)* (2022), IEEE, pp. 1690–1704.
10. Heinrich, A., Stute, M., Kornhuber, T., and Hollick, M. Who can find my devices? security and privacy of apple's crowd-sourced bluetooth location tracking system. *arXiv preprint arXiv:2103.02282* (2021).
11. Huang, D.-J., Teng, W.-C., Wang, C.-Y., Huang, H.-Y., and Hellerstein, J. M. Clock skew based node identification in wireless sensor networks. In *IEEE GLOBECOM 2008-2008 IEEE Global Telecommunications Conference* (2008), IEEE, pp. 1–5.
12. Huang, J., Albazrqaoe, W., and Xing, G. Blueid: A practical system for bluetooth device identification. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications* (2014), IEEE, pp. 2849–2857.
13. Jana, S., and Kasera, S. K. On fast and accurate detection of unauthorized wireless access points using clock skews. *IEEE transactions on Mobile Computing 9*, 3 (2009), 449–462.
14. Jin, J., and Wang, X. On the effectiveness of low latency anonymous network in the presence of timing attack. In *2009 IEEE/IFIP International Conference on Dependable Systems & Networks* (2009), IEEE, pp. 429–438.
15. Kay, E. 5 ways to use the new find my device on android. `https://blog.google/products/android/android-find-my-device/`.
16. Kleidermacher, D. How we built the new find my device network with user security and privacy in mind. `https://security.googleblog.com/2024/04/find-my-device-network-security-privacy-protections.html`.
17. Kocher, P. C. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *Advances in CryptologyCRYPTO96: 16th Annual International Cryptology Conference Santa Barbara, California, USA August 18–22, 1996 Proceedings 16* (1996), Springer, pp. 104–113.
18. Kohno, T., Broido, A., and Claffy, K. C. Remote physical device fingerprinting. *IEEE Transactions on Dependable and Secure Computing 2*, 2 (2005), 93–108.
19. Reed, M. G., Syverson, P. F., and Goldschlag, D. M. Anonymous connections and onion routing. *IEEE Journal on Selected areas in Communications 16*, 4 (1998), 482–494.
20. Rohrer, E., and Tschorsch, F. Counting down thunder: Timing attacks on privacy in payment channel networks. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies* (2020), pp. 214–227.
21. Thingsup. Thingsup ble beacon scanner and logger, 2021. `https://play.google.com/store/apps/details?id=io.thingsup.blescanner&hl=en_SG&gl=US`.
22. Tramèr, F., Boneh, D., and Paterson, K. Remote side-channel attacks on anonymous transactions. In *29th USENIX Security Symposium (USENIX Security 20)* (2020), pp. 2739–2756.
23. Wang, X., Chen, S., and Jajodia, S. Network flow watermarking attack on low-latency anonymous communication systems. In *2007 IEEE Symposium on Security and Privacy (SP'07)* (2007), IEEE, pp. 116–130.
24. Yu, T., Henderson, J., Tiu, A., and Haines, T. Privacy analysis of samsung's crowd-sourced bluetooth location tracking system. *arXiv preprint arXiv:2210.14702* (2022).