SaDiF: Spoofing Attack on BLE Direction Finding Based Localization System

Runting Zhang[1], Yijie Li[2*], Dian Ding[1*], Hao Pan[1], Yongzhao Zhang[3], Xiaoyu Ji[4], Jiadi Yu[1], Guangtao Xue[1], Yi-Chao Chen[1]

[1]Shanghai Jiao Tong University, [2]National University of Singapore, [3]UESTC, [4]Zhejiang University
johnson_zrt@sjtu.edu.cn,yijieli@nus.edu.sg,{dingdian94,panh09}@sjtu.edu.cn,zhangyongzhao@uestc.edu.cn,xji@zju.edu.cn,{jiadiyu,gt_xue,yichao}@sjtu.edu.cn

ABSTRACT

Bluetooth Low Energy (BLE) direction finding, a feature introduced in BLE version 5.1, enables precise localization through Angle of Arrival (AoA) estimation. However, this advancement introduces new risk to BLE direction finding based localization system. Specifically, the AoA estimation based on phase sampling of constant-tone-extension (CTE) is susceptible to the signal injection attack. This paper presents SaDiF, a feasible spoofing attack mechanism to mislead the locators into mistaking the positioning result as a continuous path. By eavesdropping on BLE packets and injecting attack signals containing pre-designed disturbing phase shift, SaDiF subtly alters the AoA estimation without detection, thus interfere the localization results. Moreover, SaDiF address the challenges posed by hardware imperfections by proposing an injection timing optimization to improve attack robustness. Extensive experiments demonstrates the effectiveness of SaDiF in successfully attacking multiple BLE targets in real-time scenarios. In conclusion, our findings reveal critical security risks in BLE direction finding feature and provide insights into strengthening its defenses.

CCS CONCEPTS

• **Networks** → **Location based services**; • **Security and privacy** → **Mobile and wireless security**;

∗ Yijie Li and Dian Ding are corresponding authors.

MobiHoc '25, October 27–30, 2025, Houston, TX, USA

ACM ISBN 979-8-4007-1353-8/2025/10…$15.00
https://doi.org/10.1145/3704413.3764426

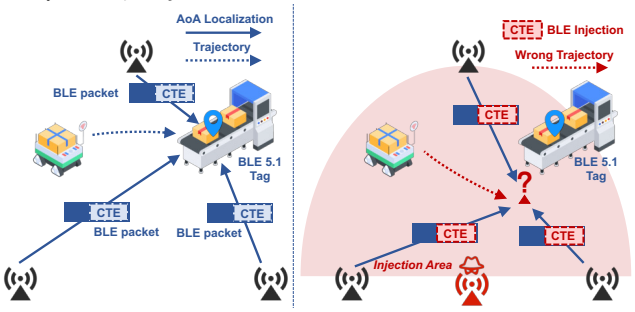

Figure 1: SaDiF: manipulate the BLE direction finding system to mislead locators of a wrong victim location.

KEYWORDS

BLE Direction Finding, Spoofing Attack, Localization

ACM Reference Format:
Runting Zhang[1], Yijie Li[2*], Dian Ding[1*], Hao Pan[1], Yongzhao Zhang[3], Xiaoyu Ji[4], Jiadi Yu[1], Guangtao Xue[1], Yi-Chao Chen[1]. 2025. SaDiF: Spoofing Attack on BLE Direction Finding Based Localization System. In *The Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '25), October 27–30, 2025, Houston, TX, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3704413.3764426

1 INTRODUCTION

Bluetooth-based localization has become a widely adopted solution in tracking applications. Among them, Bluetooth Low Energy (BLE) technology stands out due to its advantages of low cost, low power consumption and scalable solution. Traditional BLE beacon-based (e.g., ibeacon [2]) approaches have already been deployed in industrial asset tracking and automation applications [13, 28], which can achieve meter-level accuracy. As the demand for more precise and reliable positioning systems grows, Bluetooth Special Interest Group (Bluetooth SIG) released version 5.1, which includes "**direction finding**" feature [3] to accurately estimate the Angle of Arrival/Departure (AoA/AoD). This capability enables decimeter-level and has already opened up new opportunities across plenty of applications including logistics, industry, healthcare, retail, and events [7, 23].

In addition to improving accuracy, ensuring the security of localization is critical, especially for the emerging fully automated factory proposed by the leading supply chains (Amazon [5], etc.). These systems rely heavily on autonomous logistics robots that operate without direct supervision. Specifically, the *targets* (i.e., BLE tags) are installed on goods or transport containers. These tags transmit signal regularly, then the *locator* estimates the specific location of the item and report the location to the logistic robots. Based on location information, robots perform tasks such as sorting or handling. Once the locator fails to correctly report the location of the goods, the work of the logistics robots will be in chaos. Therefore, localization systems are mainly becoming the prime targets for malicious actors attempting to spoof, intercept, or manipulate signals to gain unauthorized access.

Unfortunately, the novel direction finding feature brings new risk to BLE positioning systems. The risk comes from the characteristics that the AoA estimation in direction finding feature relies on antenna switching to sample the phase. However, this process lacks a device authentication mechanism, the attackers can maliciously inject unexpected phase shifts (called **signal injection attack**) without changing the packet content to spoof the positioning results. Furthermore, direction finding targets commonly exploit periodic advertising procedures [3], which allow attackers to eavesdrop the BLE packets and prepare for their attacks.

In this paper, we proposed a feasible attack named SADIF, which aims to manipulate the positioning system, falsely reporting the victim device's (typically BLE tags or IoT sensors) location and creating a continuous, misleading movement trajectory. Such movement is along a controlled continuous path, making the attack difficult to detect by the localization system. The main assumptions of SADIF includes two aspects: 1) the location of the locators and victims are known and the locator will report to the automatic logistics robots in real time. Location of the victims can be obtained through visual methods or existing RF positioning mechanisms by the attacker. 2) The victim device is trackable, either through a fixed BLE MAC address or identifiable by manufacturer IDs in the advertising payload.

Specifically, the illustration of SADIF is shown as Fig. 1. The victim device periodically advertises to the locator, which samples the phase information of the constant-tone-extension (CTE) to estimate the device's AoA. SADIF eavesdrops on the BLE packets and then sends a delicately designed attack signal to interfere the CTE phase estimation. As a result, the locators will recognize the CTE with a disturbing phase shift and be misled into an incorrect localization result. To improve the functionality and concealment of our attack, we addressed following challenges:

Firstly, SADIF has to bypass the existing BLE defense mechanism, which forbids the replacement of one or more BLE packets or the interference of the Protocol Data Units (PDUs) during the localization process [11, 26]. This mechanism prevents traditional Man-In-The-Middle (MITM) attacks [4] that send tampered packets or forge information. However, the direction finding process estimates the AoA by sampling the phase shift of the CTE, whose structure does not include any authentication mechanism. Therefore, we carefully crafted a CTE-based signal injection mechanism based on changing the phase sampling of CTE at the right time. Specifically, SADIF first focuses on precisely eavesdropping on the signal segments particularly related to the CTE parts and then SADIF will inject a disturbing phase instead of an entire packet to enable a wrong AoA estimation. Since there is no change in the original packet information except for the phase sampled by the locators, the signal injection attack enables a silent AoA manipulation without being detected.

Secondly, SADIF should maintain a continuous and deterministic attack path rather than a set of random points, making it appear as the victim is still being localized normally. This requires a dedicated phase shift injection design to avoid unacceptable localization deviations under real-world scenarios containing various hardware deployments of locators. To this end, we proposed an attack signal generation with a phase shift filtering algorithm to strictly regulate the signal generation process based on known information of locations and deployment specifications. We simulate the corresponding angular algorithm when attacked under noise environment to re-consider the phase shift of potential attacks and retain a reliable phase shift scheme for attack signal generation to output a stable AoA result for multiple locators.

Finally, to enhance the robustness of the attack, SADIF should perform a stable fake location, avoiding abrupt jumps in location. However, the hardware imperfections (e.g., jitter effect [19]) pose challenges in determining the specific timing of the attack signal injection, otherwise the deviation in timing will cause unstable localization offset. To address this challenge, we proposed an injection timing optimization based on a pre-measured transmission delay caused by hardware clock jitters of the victims. The timing-optimized injection ensures the phase shift closely matches the intended attack as simulated conditions. By dynamically controlling the timing, SADIF minimizes interference with other IQ sampling slots, thereby ensuring the stability of the attack.

Our contributions can be summarized as follows:

- We are the first to reveal the risks to the localization system based on official BLE direction finding feature. We proposed SADIF, a feasible attack without changing the original transmitting packet.
- Our proposed attack maintains a continuous and deterministic spoofing path, thus reducing the risk of the attack being detected.

- We proposed an optimization-based injection time estimation to confront the hardware imperfections to ensure the robustness of the attack.
- Extensive experiments demonstrates the effectiveness of SaDiF in attacking multiple BLE targets in real-time scenarios. We also proposed potential defense mechanisms, which provides valuable reference for strengthening BLE security in the future.

2 RELATED WORK

2.1 BLE Physical-Layer Security

BLE physical-layer attacks. Due to the Over The Air characteristics of BLE physical signal transmission, attacks against the physical layer can be grouped into: signal eavesdropping and signal injection. For signal eavesdropping, [1] brute-forces the parameter to calculate the channel hopping sequence that enables the eavesdropping. Attacks [29] target on the device discovery phase of BLE to sniff the advertising messages from peripheral devices with unchanged MAC address for user tracking. For signal injection, [24] discovers a nearby device that is in the non-discoverable mode when receiving responses. Existing BLE physical-layer attacks provide feasible methods to eavesdrop and inject signal to BLE communications. Most attacks focus on the privacy tracking and cannot be used in manipulating direction finding systems, while others suffer from corresponding defenses.

BLE physical-layer defenses. Defenses at the physical layer tend to detect or prevent the signal injection and eavesdropping by leveraging the physical features of the BLE signal from a specific device, and generally contain a fingerprinting stage and an identification stage [27]. BlueID [11] learns from the unique clock skew feature to detect the unauthorized devices by checking the preamble of each BLE packet. BlueShield [26] uses the signal strength in the physical signal to defend against signal injection attacks when detecting an inconsistent RSSI variation. BLE-guardian [9] jams the communication channel to avoid signal eavesdropping attacks from malicious devices. Existing BLE physical-layer defenses effectively detect and guard the physical-layer signal injection and eavesdropping attacks by judging different physical features in BLE signal transmissions. These defenses pose challenges and restrictions to our SaDiF design, forcing us to target on those unprotected new segments introduced in the direction finding protocol and design an effective method to manipulate the localization result.

2.2 Spoofing Attacks on Wireless Localization System

Spoofing attacks on the localization system have been widely studied among various wireless protocols and technologies, including the well-known GPS [17], Wi-Fi [25] and UWB [14]. Although existing wireless spoofing attacks inspire our SaDiF to some extent, they cannot work on the BLE direction finding localization system due to the internal difference between wireless protocols and system deployments.

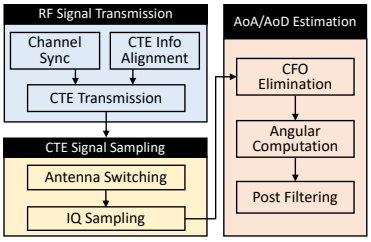

Figure 2: Direction Finding process overview.

Considering the BLE direction finding feature, Cominelli et al.[6] once controlled the phase-delay of the receiver by modifying phase after switching time to compromise correct AoA detection. However, their implementation is limited to manually controlled transceivers (USRPs at TX and RX) that differ from real-deployed locators and targets. With regard to the new packet structure, packet flow regulation, and the angular estimation process for standard BLE direction finding system, our SaDiF has to design a feasible attack plan that differs from existing spoofing attacks.

3 BACKGROUND

BLE Direction Finding Feature. To meet the increasing demand for precise location services, Bluetooth introduced the direction finding feature in version 5.1 [3]. This allows BLE devices, the locators, to determine the direction of signals received from another device, known as the target. Core Specification enables the AoA and AoD to be used in either connectionless or connection-oriented communication (suggesting AoD in connectionless and AoA in connection-oriented mode). However, connection-oriented mode is rarely used in large-scale deployment due to its limited support for concurrent targets (maximum 20-30 concurrent connections).

AoA Estimation Process. As Fig. 2 shows, the direction finding process consists of three main parts: RF signal transmission, CTE signal sampling, and AoA estimation. In the RF signal transmission, the target and locator prepare for the CTE transmission by synchronizing the time and channel, meanwhile aligning the CTE Info indicating the detailed CTE structure. When the locator samples the CTE signal from the target, it operates an antenna switching mechanism based on the aligned CTE structure and generates IQ samples on different antennas. The entire IQ samplings for the CTE is given to the localization server to conduct an AoA estimation.

In brief, the AoA is calculated by measuring phase differences between the signals sampled by different antennas. For a multi-element antenna array, the phase difference $\Delta\phi$ between two antennas is $\Delta\phi = \angle\,(I_1 + jQ_1)/(I_2 + jQ_2)$, where I_1 and Q_1 are the in-phase and quadrature components of the signal received by antenna 1, and I_2 and Q_2 are those received by antenna 2. This phase difference ($\Delta\phi$) is then used

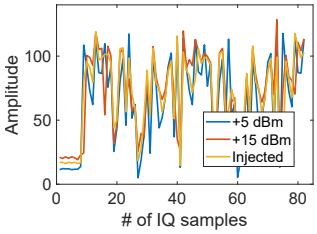

Figure 3: Amplitude fluctuation in sampled CTE.

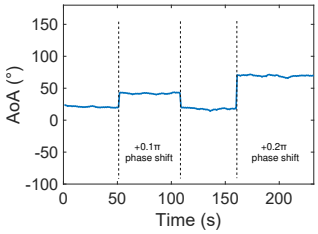

Figure 4: AoA varies with CTE phase shift.

to estimate the AoA based on the geometry of the antenna array. For a simple two-antenna array, the AoA is given by $\theta = \sin^{-1}(\lambda\Delta\phi/d)$, where λ is the signal wavelength, d is the distance between the two antennas. In real deployed direction finding systems, the Carrier Frequency Offset (CFO) elimination is first conducted to the generated raw phase to eliminate the error from hardware imperfections. The CFO is estimated using IQ samples during the reference period with the linear regression method. The AoA is then calculated utilizing angular algorithms (e.g., Multiple Signal Classification (MUSIC)) [15] with the post filtering afterwards.

4 OVERVIEW

4.1 Key Observation

Amplitude fluctuates in CTE-period signal. We observe that the IQ samples received by locators during the CTE period suffer from an unexpected amplitude fluctuation. Unlike the RSSI fluctuation that commonly results from environmental influences (e.g., multipath, RF interference), the amplitude fluctuation of CTE samples mainly comes from the hardware imperfection of the antenna switching process of the locator (RF switch in the antenna array), which leads to a huge and unexpected amplitude fluctuation among samples.

As shown in Fig. 3, the amplitude of the CTE signal ($A = \sqrt{I^2 + Q^2}$) fluctuates violently especially during the antenna switching period (#8-82 IQ samplings) without any obvious pattern of amplitude change. Also, common amplitude judgments (e.g., average amplitude) fail to determine an altered CTE with different Tx power (increased from +5 to +15 dBm) due to the unexpected fluctuation and an inherent normalization of IQ samplings to integers within $[-127, 128]$. We further compare the amplitude when the original CTE is injected. A new signal of a randomized $\{0, 1\}$ sequence overshadows the original signal due to the capture effect [8], and the CTE amplitude is hard to detect whether a signal injection occurs. Besides, the signal strength of CTE is not included in the RSSI calculation of a BLE packet [3]. This key observation inspires us to directly manipulate the CTE signal, which hides from potential defenses including RSSI detections or amplitude judgments.

Phase shift in CTE deviates the AoA output. The phase shift of the CTE signal from the target (TX), even only the reference period signal in CTE, leads to a different AoA result. Cominelli et al.[6] once noticed that the phase shift during the entire CTE period may lead to a different AoA output. Here, we further evaluate that even modifying the CTE reference period signal will result in a different AoA value by impacting on the CFO elimination process.

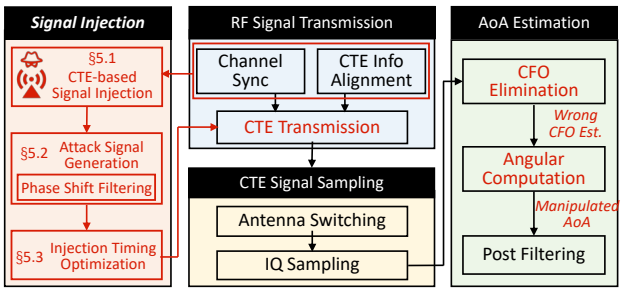

Figure 5: System overview.

As shown in Fig. 4, we collect the AoA outputs for a fixed target, but introducing different additional phase shifts during the CTE reference period. Given an additional phase shift of $+0.1\pi$ and $+0.2\pi$, the AoA values are instantly influenced and change from the ~24.5° to ~48.2° and ~76.3°, when the positions of targets and locators are all fixed (AoA ground truth of 25°). The second observation provides a chance to manipulate the direction finding process with a slight modification to the phase of the CTE reference period signal.

4.2 Threat Model

Attack Objectives. We consider the users or administrators of the BLE direction finding localization system as our victims. In general, a high-accuracy (dm-level) indoor localization service is provided to the supported BLE devices through a triangulation process using AoA. The goal of SADIF is to mislead the localization result of the target BLE devices, specifically, to cause the positioning coordinates to deviate and form a continuous and misleading movement strategy.

To define the objective of our SADIF, we first formulate the process of a normal direction finding system. Consider a BLE victim device V using direction finding localization system with the ground truth movement trajectory of $\{P_{V,GT}(t)\}$ with respect to time t, assume that the nearby n locators (in general $n = 3$ locators with the highest RSSI account for the triangulation in real deployed systems) are fixed at locations: $P_{Loc} = [P^1_{Loc}, P^2_{Loc}, ..., P^n_{Loc}]$. At time t_i, the measured AoA from the n locators are: $\theta_V(t_i) = [\theta^1_V(t_i), \theta^2_V(t_i), ..., \theta^n_V(t_i)]$. The location $P_{V,DF}(t_i)$ of V at time t_i is then derived by the triangulation algorithm: $P_{V,DF}(t_i) = Tri(\theta_V(t_i), P_{Loc})$. The measured trajectory of V is formed as: $\{P_{V,DF}(t)\}$ satisfying: $|P_{V,DF}(t_i) - P_{V,GT}(t_i)| \leq \epsilon_{DF}$, where ϵ_{DF} represents the acceptable localization error (dm level).

The goal of our SADIF is to manipulate the localization result of V: $P'_{V,DF}(t_i) = Tri(\theta'_V(t_i), P_{Loc})$ by influencing the corresponding AoA measurements: $\theta'_V(t_i) = [\theta'^1_V(t_i), \theta'^2_V(t_i), ..., \theta'^n_V(t_i)]$. To create a misleading and continuous trajectory

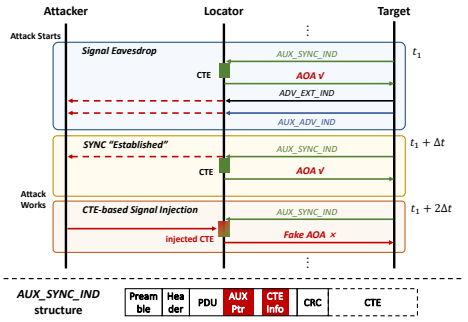

Figure 6: CTE-based signal injection workflow: Signal eavesdropping, Synchronization established and Signal injection.

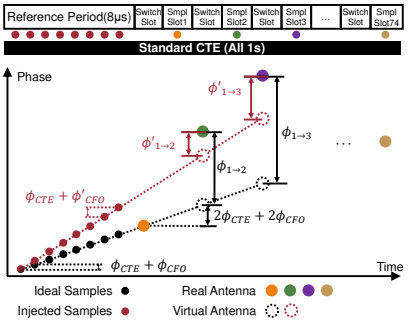

Figure 7: CFO estimation with additional phase shift by injecting the attack signal.

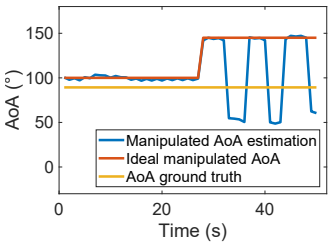

Figure 8: Manipulated AoA suffers from unexpected deviation.

of V, the measured trajectory $\{P'_{V,DF}(t)\}$ must satisfy:

$$\begin{cases} Misleading: |P'_{V,DF}(t_i) - P_{V,GT}(t_i)| \geq d_{atk}, & \forall t_i \geq t_0 \\ Continuity: |P'_{V,DF}(t_{i+1}) - P'_{V,DF}(t_i)| \leq d_{step}, & \forall t_i \geq t_0 \end{cases}$$

where d_{atk} represents the distance deviation sufficient for misleading purposes and d_{step} limits the moving step of V between adjacent measurements. t_0 corresponds to the time when our SADIF begins. Since we focus on the logistics scenarios, we empirically choose $d_{atk} = 3m$ to match the distance between cargo shelves and conveyor belts, which is also out of the error range (dm) of standard direction finding systems. We determine the moving step d_{step} with the restriction of: $d_{step} < v_V \cdot (t_{i+1} - t_i) + \epsilon$, which is set to $0.5m$ (relatively loose) to fit general cases.

Attacker's Assumptions. When launching SADIF, we assume that attackers have the following capabilities: access the locations of both victim devices and all nearby locators; use easily available devices whose transmission power is higher than that of the victim devices; acquire the antenna array size and type on the locators and the switching pattern. Location of the victims can be determined through visual methods or common RF localization technologies (RSSI, TDoA, AoA) by the attacker. Considering the restrictions on attackers, they are prohibited from physical access to any device or hardware in the deployed direction finding localization system; cannot decrypt the encrypted BLE packet or tamper with standard BLE protocol; cannot move the attacker's device arbitrarily.

5 SADIF DESIGN

Our SADIF aims to manipulate the direction finding localization result and create a misleading and continuous trajectory of the victim by manipulating the AoA estimations. As shown in Fig. 5, the attacker first acquires essential synchronization parameters and information by eavesdropping the ongoing direction finding communications, then prepares for the signal injection attack in the CTE transmission at the proper time and channel (§5.1). Meanwhile, the acquired information is used for the attack signal generation, with a phase shift filtering algorithm to better select the appropriate phase modifications (§5.2). After optimizing the injection timing to increase the system robustness (§5.3), the CTE-based signal injection is finally conducted during the CTE transmission. The tampered CTE signal, after being sampled by locators, will lead to wrong CFO estimations and eliminations in localization server and finally output a manipulated AoA value.

5.1 CTE-based Signal Injection

To bypass the existing BLE defenses, we have to avoid being noticed by victims (e.g., a simple system DoS attack), or defending by existing methods [11, 26]. Practical BLE physical-layer defenses that exploit either BLE packet RSSI or preamble to identify the legal and malicious devices, can distinguish whether a BLE packet has been replaced by attackers. Therefore, the replacement of any one or more packets completely (e.g., MitM attack) or partially in the payload (e.g., signal injection attack) within the standard direction finding process will be countered by defenses [4].

To deal with the above challenge, we propose the CTE-based signal injection mechanism aimed at the unique vulnerabilities of the direction finding protocol. Instead of replacing any packet transmitted by the victim target, the attacker manages to locate the essential CTE-period signal from target and launch a signal-injection attack during CTE period (more specifically, the reference period of the whole CTE). Existing defense mechanisms are therefore ineffective since the packet preamble remains legal and unchanged (still transmitted from the original target) and the RSSI behaves normally without obvious fluctuations (CTE period signal not included in RSSI calculation) [3].

The detailed workflow of the CTE-based signal injection mechanism is shown in Fig. 6, containing three parts: 1) **Signal eavesdropping**: the attacker first launches the signal eavesdropping on the standard direction finding process of the victim device. Since the victim device is trackable to the attacker, the attacker is able to capture the advertising packet *ADV_EXT_IND* at the BLE advertising channel. By separating the necessary parameters from the *AUXPtr* segments in

Algorithm 1: Attack Signal Generation

Input: Location of the victim target P_V and nearby n locators $P_{Loc} = [P^1_{Loc}, P^2_{Loc}, ..., P^n_{Loc}]$, antenna array size $S_{ar} = M \times N$, CTE slot interval $T_{CTE} = 1$ *or* $2\mu s$, spectrum peak ratio bound r_b.
Output: Manipulated CFO difference $\Delta\phi_{CFO}$ by the attack signal for signal injection.

for $i = 1 : n, i \in \mathbb{Z}^+$ **do**
  $\Phi_i = \{\phi_{a_1 \to a_{k+1}}, k \in [1, S_{ar}], k \in \mathbb{Z}^+\} = Sim(P_V, P^i_{Loc}, S_{ar}, T_{CTE})$;
for $j = 1 : 72, i \in \mathbb{Z}^+$ **do**
  $\phi_{temp} = j \cdot \pi/36$;
  for $i = 1 : n, i \in \mathbb{Z}^+$ **do**
    $\Phi'_i = \{\phi_{a_1 \to a_{k+1}} - k \cdot \phi_{temp}\}$;
    $\{\theta^i_V, r^i_p\} = MUSIC(\Phi'_i, S_{ar})$;
  if $min(\{r^i_p\}) < r_b$ **then** *continue*;
  $P_{V,M} = Tri(P_{Loc}, \{\theta^i_V\})$;
  if $P_{V,M}$ *is checked appropriate* **then** *break*;
return $\Delta\phi_{CFO} = \phi_{temp}$;

payload, the attacker is now able to accordingly receive a series of packets in the periodic advertising process from the victim device. 2) **SYNC established**: the signal eavesdropping ends and the synchronization is considered "established" when an *AUX_SYNC_IND* packet is successfully captured (at $t_1 + \Delta t$) with a CTE at the end of the packet. The attacker then acquires the time offset of the next *AUX_SYNC_IND* as well as its payload length and contents, and decodes the essential *CTEInfo* segments for CTE structure parameters. 3) **Signal injection**: the attack finally works when the next *AUX_SYNC_IND* is transmitted from the victim (at $t_1 + 2\Delta t$), and the pre-generated attack signal is injected to the original CTE period, which leads to a fake AoA output.

5.2 Attack Signal Generation

To ensure that the misleading trajectory of victim is continuous and logical, the manipulated AoA must be stable and deterministic. In real-scenario tests, however, the manipulated AoA results are likely to show an large deviation from their expected outputs in simulated cases.

We analyze the unexpected deviations in AoA estimation based on phase shift offsets induced by the attack signal injection on the angle calculation. Fig. 7 shows the manipulated CFO estimation process when injecting an attack signal with additional phase shift. Estimating CFO by linear regression on the reference period sample, the injected samples (red dots) lead to a wrong CFO estimation $\phi'_{CFO} > \phi_{CFO}$ than the ideal original samples (black dots). After CFO elimination and the reference virtual antenna alignment, the phase difference between antennas will be wrongly derived to $\phi'_{1\to2}$ instead of the correct $\phi_{1\to2}$ (antenna #1 and 2, for example)

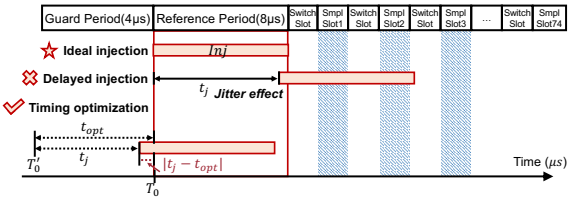

Figure 9: Injection timing optimization mechanism.

The phase difference between the set of derived $\{\phi'_{1\to i+1}\}$ and original $\{\phi_{1\to i+1}\}$ is linear. However, switching between rows/columns in a URA destroys the linear character of the phase difference sequence. Therefore, it is difficult to accurately match the phase difference sequence based on the attack signal injection, and the phase perturbation by ambient noise further increases the deviation, especially in cases of locators equipped with large-size URA antenna arrays using complex angular algorithms. For example, when launching our attack to locators with 4×4 antenna arrays using 2-D MUSIC algorithm for AoA calculation, as Fig. 8 shows, the manipulated AoA estimations show an $> 90°$ deviation to the expected manipulated AoA value when managing to manipulate the AoA from ground truth 89.3° to 145°. The unexpected deviation of manipulated AoA severely challenges our attack to create a continuous misleading trajectory, which makes it easy for the victim to notice the jumping and uncontrollable moving trajectory and realize that he/she is under attack.

As a result, to solve the above challenge, we design a phase shift filtering algorithm to select the appropriate additional phase shift when generating the attack signal. The consistency performance of the output manipulated AoA is judged by its corresponding MUSIC spectrum for angular computation. For potential phase shift that meets the misleading and continuous trajectory requirement in simulation, the peak ratio (highest peak/second local peak) of its MUSIC spectrum is calculated and filtered by our ratio presets to ensure the stability of the output AoA. The complete attack signal generation algorithm (including the phase shift filtering) is shown in Algo. 1. The IQ values of the attack signal to inject in CTE reference period is then derived as: $\{I'+iQ'\} = \{I+iQ\} \cdot angle(\Delta\phi_{CFO} - \phi_{CFO,atk})$, where the ideal IQ value $\{I + iQ\}$ can be modulated according to its known PDU content, and the CFO of the attacker device $\phi_{CFO,atk}$ should be measured in advance and eliminated here.

5.3 Injection Timing Optimization

Although the attack signal generation process is dedicated designed to ensure the consistency of manipulated AoA, the hardware imperfection of the attacker's device may still harm the system robustness. For instance, the transceiver of the attacker performs an unexpected signal injection timing deviation (commonly a μs-level delay) when giving the same signal transmission operations and timing parameters. As shown in Fig. 9, a delayed signal injection may cause an

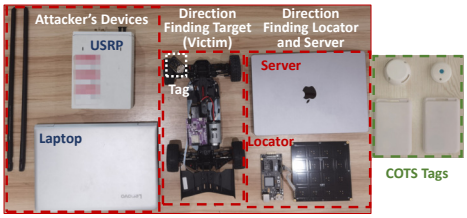

Figure 10: Implementation of SADIF.

Figure 11: Experimental environments.

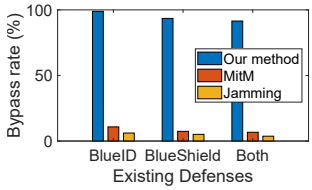

Figure 12: Bypass rate against existing defenses.

incomplete covering of the reference period signal (red zone) and more seriously, unexpectedly influence the IQ samples during the sampling slots (blue zones).

The main cause of the signal injection timing deviation can be attributed to the jitter effect (jitter phenomenon) from the attacker's hardware [19]. Here, we mainly focus on the long term (accumulated) jitter that causes an obvious time delay or deviation when transmitting the attack signal. The long term jitter is defined as the deviation of clock edge-to-edge time interval from the ideal one over more than one clock cycle. The timing jitter of a phase-modulated sinusoid is expressed as: $x(t) = A_c \cos[\omega_c(t + \theta(t)/\omega_c)]$, $\Delta_{jitter} = \theta(t)/\omega_c$.

We design a signal injection timing optimization mechanism for our SADIF. Given the applied CTE structure and the slot interval, we aim to maximize the coverage of the attack signal within the reference period while minimizing the interference with the following sampling slots. As shown in Fig. 9, we optimize our signal injection operation by triggering the transmitting operation of the attack signal t_{opt} ahead (from ideal case T_0 to $T_0' = T_0 - t_{opt}$) and retaining the same phase shift. The optimization interval t_{opt} is derived by solving the following optimization problem:

$$\underset{t_{opt}}{\text{argmax}}: E[L_{ref}(t_{opt}, t_j)] \cdot W_{ref} - \sum_{i=1}^{n_{smpl}} E[L^i_{smpl}(t_{opt}, t_j)] \cdot W_{smpl}$$

subject to: $0 \le t_{opt} \le t_{opt\_max}$, where $L_{ref}(t_{opt}, t_j) = \max(0, 8 - |t_j - t_{opt}|)$ and $L^i_{smpl}(t_{opt}, t_j)$ represent the time interval (μs) of the injection signal that covers the reference period and the i^{th} sampling slot. W_{ref} and W_{smpl} weight the influence on the output AoA error when properly covering the reference period and avoiding the sampling slots. A maximum t_{opt_max} is set to prevent interfering with the former packet payload transmission.

6 EVALUATION

6.1 Experiment Setup and Metrics

Implementation. As shown in Fig. 10, for attacker's devices, we use an Ettus USRP N210 [16] as the attacker's transceiver. The Tx and Rx gains are set to $30dB$ and $20dB$ respectively. Two identical bidirectional antennas of 12dBi are equipped on the USRP RF ports for a higher Tx signal strength. The N210 is connected and controlled by a Lenovo laptop with a 4-core CPU @$2.3GHz$ via Ethernet. The direction finding locator is developed using the Silicon Labs Pro Kit [22] due to its highly programmable console for debugging, and is connected to a Macbook Pro as the localization server. The radio board [20] coupled with the Pro Kit features a dual-polarized antenna array with 4×4 antenna slots. For the victim device, the Silicon Labs EFR32BG22 thunderboard kit [21] with a BLE 5.2 version is used as the tracking tag during the in-lab scenario experiments with multiple parameters available to modify (e.g., direction finding interval, BLE Tx power, etc.). The tag is mounted on a remote car for mobility simulations. For real deployed direction finding systems, some COTS tags with SoCs from various manufacturers (e.g., TI CC2640 [12], Nordic nRF52810 [18]) are also used to evaluate the attack performance in real logistics scenarios.

Environment and deployment. We use two different experimental environments: an in-lab scenario in the school laboratory to conduct most of our evaluations and a real-logistics scenario in a logistics transfer center to evaluate on the real deployed systems with more occlusion and interference. For the in-lab scenario, experiments are conducted in the lobby and corridor (with around $22.5m \times 14.5m$ area). Locators are deployed on the brackets (~$2.5m$ high) and direction finding tags are fixed on the remote car to imitate the mobile attack scenarios of logistics packet transportation. For the real logistics scenario, the tags are deployed on the packages on the shelves. Locators and the attacker's device are deployed in the aisle between the shelves with a relatively strong occlusion and multipath effect.

Evaluation metrics. We evaluate the performance of SADIF by comparing the deviation between the simulated (expected) manipulated AoA value or localization position (trajectory) and the real calculated AoA output or location (trajectory). A lower deviation shows a greater attack performance with more reliable and consistent spoofing effects, while a huge deviation leads to an uncontrollable manipulated localization performance as well as a large attack failure rate.

6.2 Microbenchmark

CTE-based signal injection. We evaluate our CTE-based signal injection mechanism by comparing the bypass rate against existing BLE physical-layer defenses (BlueID [11] and BlueShield [26]). Two other common attacks (the MitM attack and jamming attack) are also taken into account as baselines. The threshold for BlueID is set to 128 for the best true positive rate, and no other features except the RSSI is considered in

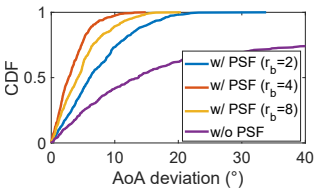

Figure 13: AoA deviation with different generated attack signal.

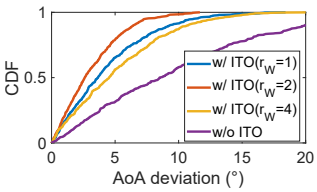

Figure 14: AoA deviation with injection timing optimization.

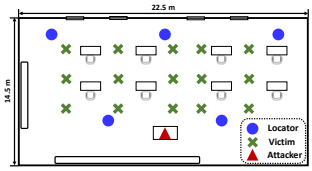

Figure 15: Floor plan for still point attack

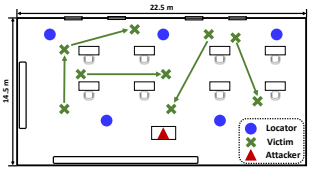

Figure 16: Floor plan for trajectory attack

BlueShield to minimize the false negative rate. Each attack is repeated for ~500 times (around ~20*mins*) to generate the bypass rate. As Fig. 12 shows, our CTE-based signal injection achieves a very high bypass rate (> 91.5%) when encountering both defense mechanisms, while the MitM attack and the jamming attack (bypass rate of ~6.7%, ~3.7%) are ineffective against these defenses. For cases that our method is detected and defended by BlueShield (~6%), one possible reason is the injected signal unexpectedly interferes with the original payload, leading to an RSSI deviation.

Attack signal generation. We compare four different mechanisms for the generation of attack signal: without the phase shift filtering (PSF) algorithm and applying PSF with the AoA MUSIC spectrum peak ratio bound of $r_b = \{2, 4, 8\}$, respectively. AoA deviations are collected for ~ 500 times for each of the evaluated mechanisms. As shown in Fig. 13, four attack signal generation mechanisms get the average AoA deviation of 7.28°, 3.31°, 4.81°, 29.77°, where the generation mechanism with PSF using $r_b = 4$ achieves the best performance with a $90th$ percentile of AoA deviation 6.55°, which fits in line with the acceptable error range in standard direction finding systems (5°-7° AoA error). A high and strict r_b may cause a limited manipulated AoA range, while a low r_b fails to filter out the unstable manipulated AoA.

Injection timing optimization. We compare four different injection timing optimization mechanisms: without any timing optimization and applying our injection timing optimization (ITO) with different weight parameters of $r_W = W_{smpl}/W_{ref} = \{1, 2, 4\}$, respectively. AoA deviations are collected for ~ 500 times in each case. As shown in Fig. 14, four optimization mechanisms achieve the average AoA deviation of 4.41°, 3.16°, 6.19°, 9.86°, and ITO with $r_W = 2$ achieves the best performance with a $90th$ percentile of AoA deviation 6.41°. The selection of r_W indicates the impact on the manipulated AoA when the injected signal covers the reference period or interferes with the sampling slots, thus a balanced

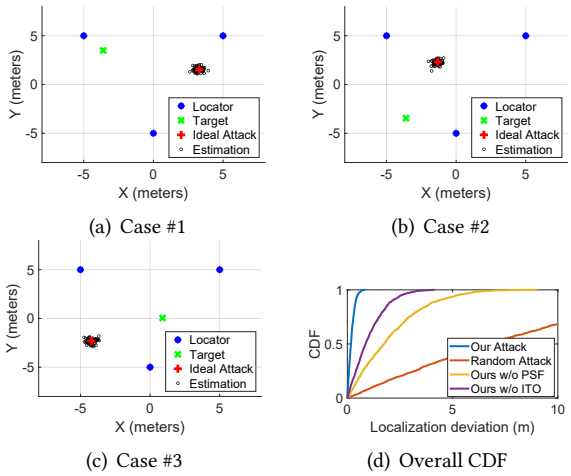

(a) Case #1 (b) Case #2

(c) Case #3 (d) Overall CDF

Figure 17: Attack on still points

r_W appropriately optimizes the timing of signal injection and reduces the deviation of the manipulated AoA.

6.3 Overall Performance

All schemes in our SADIF are enabled to achieve the best overall attack performance. We first evaluate the performance of the still point attack (assume a stationary victim), then conduct an experiment on the real-scenario trajectory attack. Fig. 15, 16 show the floor plan of the overall performance experiments. We also compare the overall performance with three other attack methods: random-position attack to inject a randomly generated attack signal to CTE, and our SADIF without the PSF or ITO components.

Still Point Attack. For the attack on still points, we fix the location of the victim target, and launch our SADIF to manipulate the localization result to an ideal wrong position. The ideal position under attack must be far from the real position to ensure the misleading effect, and all localization measurements must be stable and close to the ideal attack position to serve for the following requirement of trajectory consistency. Here, we fix the parameters for attack signal generation in our SADIF to manipulate the same AoA output. Fig. 17(a), 17(b) and 17(c) show three example cases of the still point attack experiments, where the position estimations (black *o*) gather around the expected ideal attack position (red +) closely, with a large distance (> 5*m*) deviated from the real location of the target (green ×). The overall performance CDF (Fig. 17(d)) for all still point attack experiments with the four attack methods show an average localization deviation of 0.199*m*, 7.943*m*, 2.053*m*, 1.037*m*, respectively. The random-position attack suffers from a high deviation due to its uncontrollable phase shift and output localization result when fitting to a still point. Our SADIF is proved effective and stable when attacking on stationary victims.

Trajectory Attack. For the trajectory attack, we preset the moving trajectory and the velocity of the victim target to

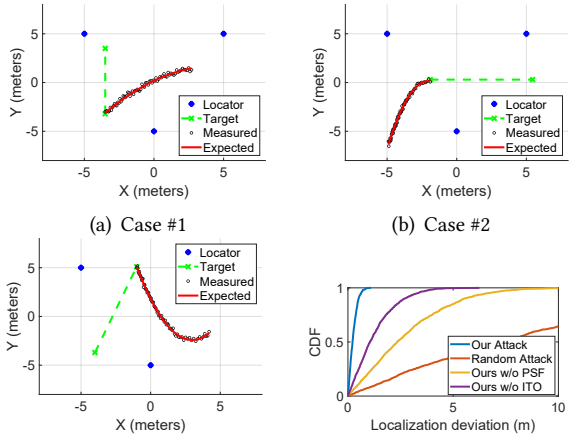

(a) Case #1 (b) Case #2

(c) Case #3 (d) Overall CDF

Figure 18: Attack on trajectory.

perform a uniform linear motion. We aim to create a misleading and continuous spoofing trajectory to gradually deviate from the preset trajectory. Fig. 18(a), 18(b) and 18(c) show three example cases of the trajectory attack experiments. The measured locations show a continuous and different trajectory from the original preset trajectory (green), and close to the expected manipulated trajectory (red, smoothed). The overall performance CDF (Fig. 17(d)) for all trajectory attack experiments with the four attack methods show an average localization deviation of $0.250m$, $8.06m$, $2.57m$, $1.14m$, respectively. The random-position attack fails to fit to a continuous path with random output positions. Our SADIF performs the best and succeeds in manipulating trajectory for the moving target with a low localization deviation (within the deci-meter error range of standard direction finding).

6.4 Impact of Factors

Impact of antenna array size and type. We modify the locator firmware to control the available antenna elements and its corresponding antenna switching pattern. We test the antenna arrays of 1×4 ULA using 1-D MUSIC angular algorithm for AoA and 2×2, 4×4 URA using 2-D MUSIC (azimuth as AoA). Fig. 19(a) shows that for locators using 1×4 ULA, 2×2, and 4×4 URA, our SADIF achieves the average manipulated localization deviation of $0.2157m$, $0.2799m$, $0.3521m$, respectively. Given the localization error of $\sim 0.3m$ in standard direction finding systems, our SADIF is proved effective to attack on various antenna sizes and types, even on those complex 4×4 URA deployed in real systems.

Impact of RF interference. We consider the following interference cases: light RF interference with no other BLE or Wi-Fi users nearby, medium RF interference with several Wi-Fi mobile phone users to explore news, and heavy RF interference with Wi-Fi downloading and Bluetooth wireless headphones using. The standard direction finding system is not disturbed under these interference scenarios. Fig. 19(b) shows that our system achieves the average localization deviation of $0.193m$, $0.204m$, $0.223m$, for the light, medium and

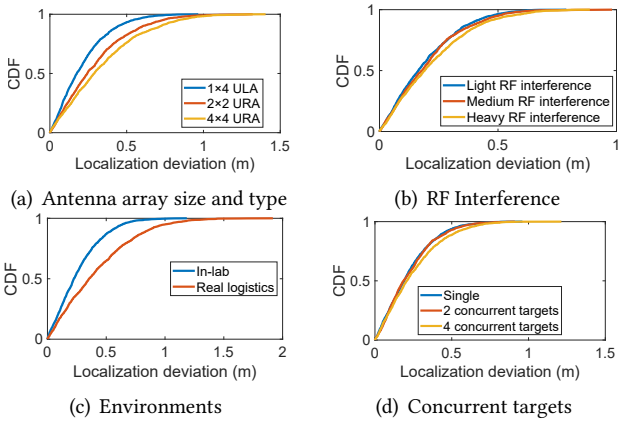

(a) Antenna array size and type (b) RF Interference

(c) Environments (d) Concurrent targets

Figure 19: Impact of different factors.

heavy interference cases respectively. The RF interference causes limited impact on the manipulated localization performance, as long as the direction finding is still running.

Impact of environments. Two distinct environments are considered: the in-lab scenario and the real logistics scenario (Fig. 11). Fig. 11 shows that our system achieves the localization deviation of $0.258m$ and $0.407m$ in average under these environments. The real logistics scenario brings a higher manipulated localization deviation that still stays within the decimeter-level direction finding error range, which showcase the effectiveness of our SADIF in complex environments.

Impact of concurrent targets. To test our attack performance on multiple victims (concurrent localization targets), our SADIF is launched under three different scenarios with 1, 2 and 4 concurrent targets running direction finding. All localization tags are set to the connectionless mode with periodic advertising time interval of $\sim 200ms$, which fits the update rate requirement for standard direction finding systems [3]. Fig. 19(d) shows that the average localization deviations for different concurrent targets are $0.214m$, $0.225m$, $0.247m$, respectively. Our SADIF is effective to perform the attack on multiple victims (at least 4) at the same time without exposing an obvious localization deviation increase.

7 COUNTERMEASURE AND CONCLUSION

Countermeasure. The key idea is to additionally apply a physical-layer device identification process on the CTE period signal to determine whether it is maliciously tampered or not. Considering the CTE signal characteristics of amplitude deviations due to antenna array RF switch imperfections and phase pattern differs by changing locations, it is invalid to simply judge from signal amplitude or phase deviations. An effective method relies on the CTE-period CFO estimation results, a stable physical-layer hardware imperfection fingerprint to uniquely identify the specific target device.

To be specific, for a set of IQ samples generated from locators, the CFO estimation of the complete CTE-period signal (instead of simply the reference period signal) should be calculated in advance to the following angular computation.

Considering a common setup of maximum 160μs-long CTE signal of 1μs slot period (aiming for maximum 82 IQ samples per packet), the CFO estimations are derived as:

$$\phi_{CFO}(n) = \begin{cases} N/A, & n = 1, \\ \phi_n - \phi_{n-1} - \pi/2, & 8 \geq n > 1, \\ \phi_n - \phi_{n-N} - N \cdot \pi, & 82 \geq n > 8. \end{cases}$$

where ϕ_i represents the calculated phase for the $i-th$ IQ sampling and N is the number of antennas within a complete antenna switching round. An empirical bound for the standard deviation of CFO estimations is set to $500Hz$ (based on [10]) to compare the intra- and inter-packet CFO estimations and determine whether a packet is fully spoofed or partially injected, e.g., our SADIF.

Conclusion. In this paper, we highlight significant vulnerabilities in the official BLE direction finding system introduced in BLE version 5.1. We proposed SADIF, an effective spoofing attack mechanism that effectively exploits the inherent vulnerabilities in the phase sampling process of constant-tone-extension (CTE) signals. By carefully injecting signals with disturbing phase shifts, SADIF can successfully manipulates the AOA estimation, misleading existing BLE localization systems into reporting continuous and incorrect trajectories. Extensive experiments demonstrated the effectiveness of SADIF. Finally, we also proposed practical countermeasures to strengthen the resilience of such BLE direction finding systems against such spoofing attacks.

ACKNOWLEDGMENTS

This work is supported by Natural Science Foundation of Shanghai (No. 25ZR1401171, 24ZR1430600), and in part by National Natural Science Foundation of China (No. 61936015, No. 6240071246).

REFERENCES

[1] Wahhab Albazrqaoe, Jun Huang, and Guoliang Xing. Practical bluetooth traffic sniffing: Systems and privacy implications. MobiSys '16, page 333–345, New York, NY, USA, 2016. ACM.

[2] Apple Developer. iBeacon. https://developer.apple.com/ibeacon/, 2024.

[3] Bluetooth SIG. Core specification v5.1. https://www.bluetooth.com/specifications/specs/core-specification-5-1/, 2019.

[4] Romain Cayre, Florent Galtier, Guillaume Auriol, Vincent Nicomette, Mohamed Kaâniche, and Géraldine Marconato. Injectable: Injecting malicious traffic into established bluetooth low energy connections. In *DSN 2021*, pages 388–399, 2021.

[5] Supply Chain. Amazon's bid to revolutionise warehouse automation, 2024.

[6] Marco Cominelli, Paul Patras, and Francesco Gringoli. Dead on arrival: An empirical study of the bluetooth 5.1 positioning system. In *WiNTECH '19*, pages 13–20, 2019.

[7] Mitsubishi Materials Corporation. Mitsubishi materials corporation tracks 80,000 tonnes of copper products in its challenging metal-heavy environment, 2023.

[8] Simon Erni, Martin Kotuliak, Patrick Leu, Marc Roeschlin, and Srdjan Capkun. Adaptover: adaptive overshadowing attacks in cellular networks. MobiCom'22, page 743–755, NY, USA, 2022. ACM.

[9] Kassem Fawaz, Kyu-Han Kim, and Kang G. Shin. Protecting privacy of BLE device users. In *USENIX Security 16*, pages 1205–1221, Austin, TX, August 2016. USENIX Association.

[10] Hadi Givehchian, Nishant Bhaskar, Alexander Redding, Han Zhao, Aaron Schulman, and Dinesh Bharadia. Practical obfuscation of ble physical-layer fingerprints on mobile devices. In *IEEE S&P 2024*, pages 2867–2885, 2024.

[11] Jun Huang, Wahhab Albazrqaoe, and Guoliang Xing. Blueid: A practical system for bluetooth device identification. In *IEEE INFOCOM 2014*, pages 2849–2857, 2014.

[12] Texas Instruments. Cc2640: Simplelink™ 32-bit arm cortex-m3 bluetooth® low energy wireless mcu with 128kb flash. https://www.ti.com/product/CC2640/, 2016.

[13] Liheng Jiang, Yongzhao Zhang, Ting Chen, Yi-Chao Chen, Dian Ding, Yijie Li, Fenghua Xu, Liwei Guo, Jingwei Li, Xiong Li, Jiguo Yu, and Xiaosong Zhang. A survey on acoustic sensing in the metasurface era: Challenges, advances, and applications. *Fundamental Research*, 2025.

[14] Patrick Leu, Giovanni Camurati, Alexander Heinrich, Marc Roeschlin, Claudio Anliker, Matthias Hollick, Srdjan Capkun, and Jiska Classen. Ghost peak: Practical distance reduction attacks against HRP UWB ranging. In *31th USENIX Security Symposium (USENIX Security 22)*, pages 1343–1359, Boston, MA, August 2022. USENIX Association.

[15] Xiang Li, Shengjie Li, Daqing Zhang, Jie Xiong, Yasha Wang, and Hong Mei. Dynamic-music: Accurate device-free indoor localization. In *Proceedings of the 2016 ACM Ubicomp*, pages 196–207, 2016.

[16] USRP N210. https://www.ettus.com/all-products/un210-kit/, 2024.

[17] Harshad Sathaye, Martin Strohmeier, Vincent Lenders, and Aanjhan Ranganathan. An experimental study of GPS spoofing and takeover attacks on UAVs. In *31th USENIX Security Symposium (USENIX Security 22)*, pages 3503–3520, Boston, MA, August 2022. USENIX Association.

[18] Nordic Semiconductor. nrf52810: Bluetooth 5.4 soc supporting bluetooth low energy, 2019.

[19] M. Shimanouchi. An approach to consistent jitter modeling for various jitter aspects and measurement methods. In *Proceedings International Test Conference 2001 (Cat. No.01CH37260)*, pages 848–857, 2001.

[20] Silicon Labs. Bg22 dual polarized antenna array radio board, 2024.

[21] Silicon Labs. Efr32bg22 thunderboard kit. https://www.silabs.com/development-tools/thunderboard/thunderboard-bg22-kit, 2024.

[22] Silicon Labs. Silicon labs bg22 bluetooth dual polarized antenna array pro kit. https://www.silabs.com/development-tools/wireless/bluetooth/bgm22-pro-kit, 2024.

[23] DGS Transports. Optimising logistics handling – cargovis indoor positioning at dgs transports, 2023.

[24] Tyler Tucker, Hunter Searle, Kevin Butler, and Patrick Traynor. Blue's clues: Practical discovery of non-discoverable bluetooth devices. In *IEEE S&P 2023*, pages 3098–3112, 2023.

[25] Kang Wang, Shuhua Chen, and Aimin Pan. Time and position spoofing with open source projects. *black hat Europe*, 148:1–8, 2015.

[26] Jianliang Wu, Yuhong Nan, Vireshwar Kumar, Mathias Payer, and Dongyan Xu. BlueShield: Detecting spoofing attacks in bluetooth low energy networks. In *RAID 2020*, pages 397–411, San Sebastian, October 2020. USENIX Association.

[27] Jianliang Wu, Ruoyu Wu, Dongyan Xu, Dave Jing Tian, and Antonio Bianchi. Sok: The long journey of exploiting and defending the legacy of king harald bluetooth. In *IEEE S&P 2024*, pages 2847–2866, 2024.

[28] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. md-track: Leveraging multi-dimensionality for passive indoor wi-fi tracking. In *MobiCom 2019*, pages 1–16, 2019.

[29] Heng Zhang, Amiya K. Maji, and Saurabh Bagchi. Privacy in the mobile world: An analysis of bluetooth scan traces. CPSIOTSEC'20, page 61–65, New York, NY, USA, 2020. ACM.