Lend Me Your Beam: Privacy Implications of Plaintext Beamforming Feedback in WiFi

Rui Xiao[§], Xiankai Chen[§], Yinghui He[¶], Jun Han[‡], and Jinsong Han[§*]
[§]Zhejiang University, [¶]Nanyang Technological University, [‡]KAIST
{ruixiao24, xiankaichen}@zju.edu.cn, yinghui.he@ntu.edu.sg, junhan@cyphy.kaist.ac.kr, hanjinsong@zju.edu.cn

***Abstract*—In recent years, the proliferation of WiFi-connected devices and related research has led to novel techniques of utilizing WiFi as sensors, i.e., capturing human movements through channel state information (CSI) perturbations. While this enables passive occupant sensing, it also introduces privacy risks from *leaked WiFi signals* that attackers can intercept, leading to threats like *occupancy detection*, critical in scenarios such as burglaries or stalking. We propose *LeakyBeam*, a novel and improved *occupancy detection attack* that leverages a new side channel from WiFi CSI, namely beamforming feedback information (BFI). BFI retains victim's movement information, even when transmitted through walls, and is easily captured since BFI packets are unencrypted, making them a rich source of privacy-sensitive information. Furthermore, we also introduce a defense mechanism that obfuscates BFI packets, requiring minimal hardware changes. We demonstrate *LeakyBeam*'s effectiveness through a comprehensive real-world evaluation at a distance of 20 meters, achieving true positive and negative rates of 82.7% and 96.7%, respectively.**

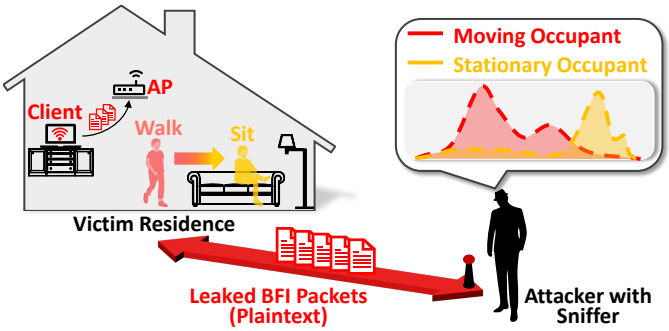

Fig. 1: Figure depicts an attack scenario using *LeakyBeam*, where an attacker uses a distant, passive sniffer to capture BFI packets from the victim's commercial WiFi client devices, leaked outside the residence, to infer the occupancy state by identifying both moving and stationary occupants.

I. Introduction

In recent years, the prevalence of WiFi-connected devices, such as laptops, mobile phones, and smart speakers, has significantly increased [1]. Consequently, we are surrounded by the WiFi signals emitted by these devices. As individuals move within their homes or offices, these WiFi signals, specifically their channel state information (CSI), are perturbed, thereby implicitly capturing information about the occupants [2]. This phenomenon presents opportunities for passive and non-intrusive occupant sensing [3], [4], [5], [6], [7].

This WiFi sensing capability, however, also introduces significant privacy risks from *leaked WiFi signals* that travel beyond their intended boundaries. Attackers within the wireless communication range can intercept these leaked signals, leading to immediate privacy threats through reconnaissance attacks, such as adversarial human motion sensing [8], [9], [10]. For example, an attacker outside the victim's house could determine whether the victim is present or not – i.e., *occupancy detection* – which is particularly concerning in real-world scenarios like targeted burglaries or stalking [11], [12], [13]. Studies show that 86% of burglars try to avoid encounters with occupants [14]. Knowing the presence or absence of occupants can empower criminals to execute their plans when they are confident no one is home, posing a serious threat to personal security and privacy.

Fortunately, a crucial factor mitigates this potential threat. An attacker, typically located outside the victim's residence, can only sniff WiFi signals via a *through-wall channel*. These signals are significantly attenuated as they pass through walls and further weaken with increasing distance from the residence [15], [16], [17]. This attenuation significantly diminishes the signal-to-noise ratio (SNR) and retains only limited privacy-sensitive information, thereby reducing the overall risk posed by this side channel.

In light of this, we pose the following question: *Is it possible to design a more sophisticated technique that allows an attacker to accurately obtain occupancy information of the victim's house even from a further distance outside the house?* To answer this question, we present *LeakyBeam*, a novel *occupancy detection attack*. *LeakyBeam* leverages a new side channel from WiFi CSI, namely the beamforming feedback information (BFI). Unlike direct measurements of analog CSI, which suffer from signal attenuation, the channel information encoded in *digital BFI packets* retains the victim's movement

*Corresponding author.

Network and Distributed System Security (NDSS) Symposium 2025
24-28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-8-3
https://dx.doi.org/10.14722/ndss.2025.240005
www.ndss-symposium.org

information even when transmitted through walls over considerable distances. This is attributed to the digital encoding of the BFI content into bits and its subsequent formatting in accordance with established WiFi protocols, ensuring that the content remains clear and undistorted [18]. Attackers can easily sniff the WiFi network to capture this side channel information, as BFI packets are transmitted in *plaintext* and are not encrypted [19]. Figure 1 depicts *LeakyBeam*'s attack scenario, where the attacker places a WiFi sniffer device outside of the victim's house from a distance away. The sniffer continuously captures the BFI packets exchanged between client devices and the AP within the victim's residence, aiming to deduce occupancy states.

However, designing *LeakyBeam* comes with two unique and significant challenges. The first challenge is environment and device ambiguity. The WiFi channel is influenced not only by human presence but also by the specific environment and device layout under the context of the multipath effect – a well-known issue in WiFi sensing [20]. Prior research has proposed various methods, including environment-specific training, device calibration, and optimizing device layouts to enhance SNR [21], [22], [23]. However, these methods are generally infeasible for attackers, who as external observers have limited capabilities and lack detailed knowledge of the interior environment of the victim's residence. Furthermore, the signal variations induced by human presence can be particularly subtle, especially when the victim is stationary. This subtlety, compounded by environmental and device ambiguity, presents significant detection challenges. To solve this challenge, we delve into the extraction of subtle human-induced dynamics from the WiFi channel, applicable even when the environment and layout are unknown. Our solution is a unified attack framework capable of detecting both moving and stationary victims by jointly leveraging both the amplitude and phase of BFI. This enables accurate detection even with an attacker's limited capabilities.

Another challenge is the complexity of BFI signal processing, particularly in mitigating phase offsets. BFI packets are derived through the application of Singular Value Decomposition (SVD) on CSI, resulting in a distinct structure. Traditional phase offset removal techniques, effective for CSI, are ineffective when directly applied to BFI due to its SVD-processed structure [24], [25], [26]. Furthermore, BFI undergoes compression, leading to a coarse-grained representation. We overcome this challenge by developing a robust feature extraction method that effectively bridges phase-offset-affected BFI to the conjugate multiplication of CSI, leveraging the antenna diversity of AP to remove phase offsets. The dynamic components attributable to human presence are further isolated and fused across Orthogonal Frequency-Division Multiplexing (OFDM) subcarriers, using subcarrier diversity to enhance the granularity and accuracy of occupancy detection.

Capitalizing on digital packets, *LeakyBeam* exhibits several favorable characteristics including its robustness to signal attenuation and external interference (§II-B). Its passive detection nature makes it non-intrusive and thus extremely difficult to identify. Moreover, the widespread accessibility of BFI – with over 73% of commercial off-the-shelf (COTS) devices compliant with 802.11ac and newer standards transmitting BFI packets [19] – positions *LeakyBeam* as a potential widespread threat. The simplicity of capturing BFI packets, such as using Wireshark for wireless packet sniffing, contrasts with the limited number of devices capable of CSI sniffing (only about 6%) [19]. The ease of accessing these plaintext BFI packets significantly lowers the device and technical barriers for attackers, collectively enhancing the practicality of *LeakyBeam* for covert and continuous occupancy monitoring.

In addition to designing a novel occupant detection attack, we design a robust defense mechanism that preserves the plaintext transmission of BFI while effectively countering the aforementioned attack. Specifically, we obfuscate both the amplitude and phase of BFI by applying a temporally varying random transformation, rendering BFI unintelligible to attackers, while still allowing the AP to recover the original information. This countermeasure requires minimal hardware modifications because it piggybacks on the spatial mapping mechanism of AP. Furthermore, it leaves client devices completely unaffected, making it highly practical for adoption.

We evaluate *LeakyBeam* attack on eight diverse APs by capturing around 1.6 million BFI packets over 49 hours across nine deployment settings varying environments and device layouts. We comprehensively evaluate *LeakyBeam*'s performance over different distances, clients, background traffic, and WiFi configurations. *LeakyBeam* proves effective even with its sniffer positioned 20 meters from the victim's residence, while also showing strong resistance to external interference. Overall, *LeakyBeam* demonstrates an average true positive rate of 82.7% and a true negative rate of 96.7%, demonstrating its effectiveness across popular devices and AP vendors. In short, our work makes the following contributions:

- We introduce *LeakyBeam*, a practical adversarial occupancy detection system utilizing the BFI side channel.
- We present the design and implementation of *LeakyBeam* that overcomes the challenges of environment and device ambiguity and complex BFI signal processing by leveraging both antenna and subcarrier diversities to derive robust features.
- We develop and validate a novel defense mechanism that capitalizes on the existing spatial mapping capabilities of WiFi APs, providing an effective countermeasure to potential attacks.

This work also sheds light on potential vulnerabilities in emerging high-frequency WiFi technologies, such as the 60 GHz mmWave 802.11ad/ay standards. These technologies extensively utilize directional beams due to the high attenuation at mmWave frequencies [27], which makes them susceptible to similar security risks. We aim to encourage further research by drawing attention to these vulnerabilities.

II. Backgrounds

We first present the preliminary of the BFI side channel.

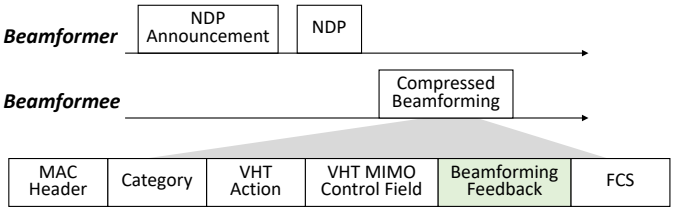

Fig. 2: Figure depicts the explicit channel measurement process: the beamformer transmits an NDP to the beamformee, which then computes and returns the steering matrix in a compressed format to the beamformer.

A. Beamforming Feedback Basics

We first briefly introduce the beamforming feedback mechanism in WiFi communication.

Beamforming in WiFi. Beamforming technology focuses WiFi signals toward the receiver in specific directions, thereby substantially enhancing the received signal strength and augmenting the throughput. Specifically, the transmitter, also called beamformer, would employ **steering matrices** to focus energy precisely towards the receiver, termed beamformee. In an $N_{TX} \times N_{RX}$ MIMO system with K subcarriers and N_{STS} spatial-temporal streams, the steering matrix for the k-th subcarrier is denoted as $\mathbf{V}_k \in \mathbb{C}^{N_{TX} \times N_{STS}}$.

Steering Matrix Computation. The steering matrix $\mathbf{V}_k$ is derived from CSI, denoted as $\mathbf{H}_k \in \mathbb{C}^{N_{RX} \times N_{TX}}$[1]. This process involves the SVD of the CSI[2]:

$$\mathbf{H}_k = \mathbf{U}_k \Sigma_k \mathbf{V}_k^{\dagger}, \tag{II.1}$$

where $\mathbf{U}_k$ and $\mathbf{V}_k$ are unitary matrices. $\mathbf{V}_k$, representing the steering matrix for beamforming, only retains the first N_{STS} columns for use. Σ_k is a diagonal matrix containing the singular values, which are real and positive, representing the strength and quality of the channel at each subcarrier.

Compressed Beamforming Feedback Packets. From the 802.11ac standard, acquiring the steering matrix necessitates the use of *explicit channel measurement frames*. This process, shown in Figure 2, starts with the beamformer transmitting a Null Data Packet (NDP), which includes a known Long Training Field (LTF), denoted as $\mathbf{X}_k \in \mathbb{C}^{N_{STS} \times N_{STS}}$, along with an NDP Announcement to the beamformee. When the beamformee receives the signal $\mathbf{Y}_k \in \mathbb{C}^{N_{RX} \times N_{STS}}$, it calculates the CSI $\mathbf{H}_k$ using the known LTF and derives steering matrix $\mathbf{V}_k$. The steering matrix $\mathbf{V}_k$ will then be **compressed** [28] and sent back to the beamformer in compressed beamforming feedback packets, which we refer to as *BFI packets* for simplicity. Another information contained in the BFI is the subcarrier-averaged stream gain, Λ, which is the arithmetic mean of the SNR values over subcarriers derived from Σ_k.

[1]In CSI $\mathbf{H}_k$, the element at (p, q) describes how the amplitude and phase of a signal in subcarrier k changes when the signal propagates from the p-th antenna of the transmitter to the q-th antenna of the receiver.

[2]We use $(\cdot)^{\dagger}$ to denote the conjugate transpose operation.

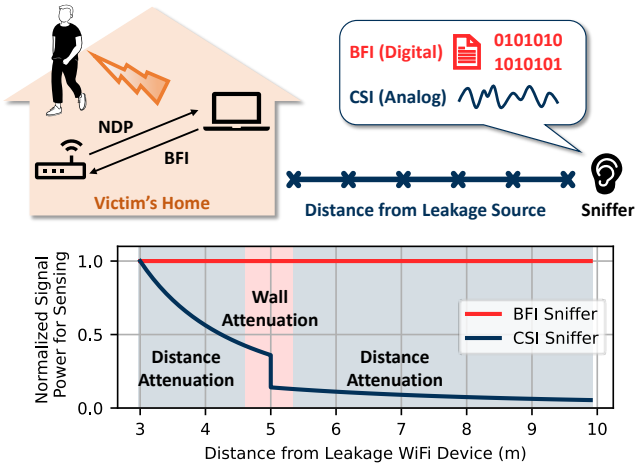

Fig. 3: Figure depicts the advantages of BFI sniffing over CSI sniffing. CSI weakens significantly with wall and distance interference, while BFI, based on digital packet sniffing, maintains its integrity regardless of environmental attenuation.

B. Distinguishing BFI as a Digital Side Channel

LeakyBeam sets itself apart by operating as a ***digital side channel***, leveraging the sniffing of digital BFI packets. This approach provides distinct advantages over traditional physical side channel attacks that intercept *analog leakage signals*, such as those measuring analog CSI signals.

Robustness to Attenuation (Long-range Capability). As conceptually depicted in Figure 3, analog leakage signals, such as CSI, experience rapid signal degradation due to distance and physical barriers, such as the exterior walls of a victim's home [15], [16]. In contrast, BFI's effectiveness is nearly ***unaffected by attenuation*** due to its reliance on digital transmissions. This ensures the integrity of the side channel data regardless of the sniffer's distance from the source or any intervening physical barriers, as long as the sniffer remains within packet sniffing range. Moreover, our experiments in §V reveal that WiFi devices are typically configured by their manufacturers to transmit BFI packets using a low Modulation and Coding Scheme (MCS) with simpler modulation and stronger error correction. While low MCS enhances signal robustness against noise and interference, it inadvertently *benefits attackers* by significantly extending their operational range. Our experiment in §V-C1 demonstrates that this effective range of BFI can extend beyond 20 meters, while CSI suffers significant performance degradation due to attenuation.

Robustness to External Interference. *LeakyBeam*'s digital nature also provides robustness to external interference. External interference, such as pedestrian motion near the sniffer, often adversely affects analog CSI-based side channels by introducing unwanted signal variation. In contrast, *LeakyBeam* captures digital packets where the sensing data, already processed and encapsulated within the victim's residence, remains unaffected by external disturbances as it travels to the sniffer.

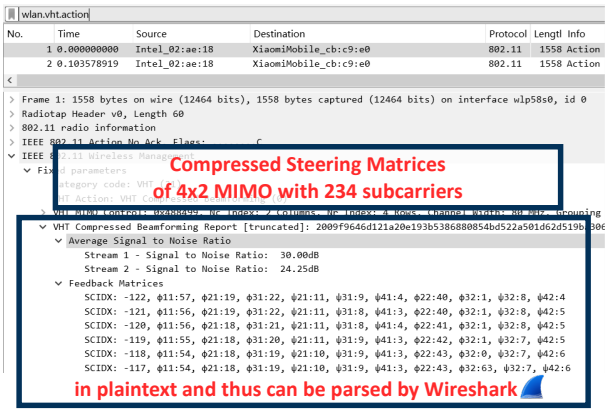

Fig. 4: Figure depicts the steering matrix of 4×2 MIMO in the BFI packet captured and parsed by Wireshark using a single-antenna sniffer.

Our experiments in §V-C2 will further validate *LeakyBeam*'s robustness to interference, enhancing its reliability as a long-range digital side channel.

Accessibility of Plaintext BFI Packets. Unlike CSI sniffing which is only supported by three chipset families [19], [29], [30], [31], BFI packets can be captured using standard wireless packet sniffing tools like Wireshark, as depicted in Figure 4. This allows for the capture of fine-grained MIMO channel information with basic hardware (e.g., a single-antenna device sniffing 4×2 MIMO with 234 subcarriers). The plaintext BFI transmission enables simple parsing and access, significantly lowering device requirements and technical barriers for attackers and increasing the risk of unauthorized data exploitation.

Stealthiness through Passive Sniffing. Unlike attacks relying on active signal or packet injection [32], [10], [33], [34], *LeakyBeam* relies solely on passive sniffing of victim's BFI packets, making it more covert and hard to detect.

C. Privacy Implication

The BFI side channel presents significant privacy threats due to its long-range capabilities, robustness to interference, accessibility, and stealthiness, making *LeakyBeam* a practical and potent tool. We envision multiple scenarios with potential privacy risks: (1. Neighborhood Surveillance) A neighbor could use a laptop to continuously monitor nearby residences. (2. Pre-Burglary Reconnaissance) Thieves might use this technology to discreetly survey homes before attempting break-ins. (3. Espionage Tactics) Spies could deploy drones equipped with sniffers to covertly gather detailed occupancy data over sensitive military or government facilities. Given the ubiquitous deployment of WiFi, *LeakyBeam* represents a low-cost and pervasive reconnaissance attack that could lead to significant and unintended consequences, necessitating increased awareness and protective measures.

III. Threat Model

Attacker's Goal and Capabilities. The *goal* of the attacker is to launch an occupancy detection attack by exploiting BFI packets leaked from the victim's WiFi device. The attacker's *capabilities* include the interception of WiFi packets and the extraction of BFI. This interception process can be carried out by using a device equipped with a network interface card (NIC) in monitor mode and network analysis tools such as Wireshark [35]. The sniffer can have an arbitrary number of antennas (e.g., a single antenna is feasible). Upon the interception of packets, the extraction of BFI involves parsing the relevant information from the packet as specified in the 802.11 protocol. Notably, since BFI packets are sent in *plaintext*, the attacker can extract this information without needing access to secure details such as the cipher keys of the victim or AP.

Assumptions. We assume the victim's WiFi device utilizes the 802.11 ac/ax standards, which incorporate explicit beamforming capabilities. The victim's AP is assumed to have more than one antenna; however, we do not place such a requirement on the number of antennas for client devices. Our analysis remains agnostic to the bandwidth and quantization levels of BFI packets. *LeakyBeam* predominantly considers downlink beamforming with the AP as beamformer, a standard practice as uplink beamforming is less common [19], [36]. We also assume stationary client devices; movement in these devices, which generally causes significant BFI variation, would typically indicate human presence. However, autonomous moving devices like robot vacuums could potentially cause false positives, which may be seen as a limitation of our work.

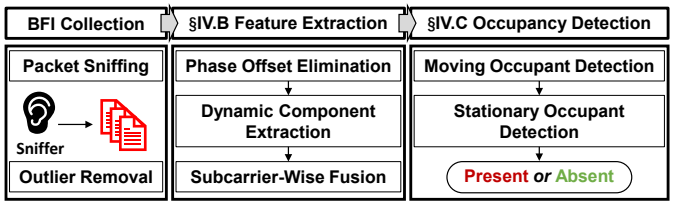

Fig. 5: Figure depicts *LeakyBeam*'s design overview. After sniffing and removing outlier BFI packets, *LeakyBeam* derives phase-offset-free features and fuses across subcarriers to detect both moving and stationary occupants and ultimately determine occupancy state, i.e., occupant present or absent.

IV. Attack Design of *LeakyBeam*

We now present *LeakyBeam*'s attack design.

A. Design Overview

LeakyBeam's design detects both stationary and moving occupants (victims) using the BFI side channel. As shown in Figure 5, the process starts with sniffing BFI packets, filtering outliers by SNR, and then extracting phase-offset-free features (§IV-B). The extracted features allow the attacker to conduct occupancy detection, which first utilizes amplitude[3] variations as a motion indicator to detect moving victims, and then detect stationary victims by analyzing breathing patterns through BFI phase information (§IV-C). *LeakyBeam* generates

[3]The terms "amplitude" and "phase" used in relation to BFI specifically refer to those attributes of the steering matrix within the BFI.

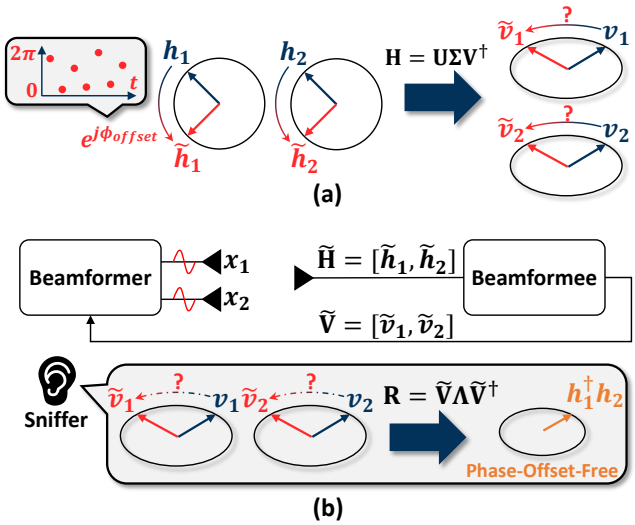

Fig. 6: Figure depicts phase offset when a beamformer has two antennas. (a) shows that temporal phase offsets in $\mathbf{H} = [h_1, h_2]$ result in corresponding shifts in the steering matrix $\mathbf{V} = [v_1, v_2]$. (b) shows the sniffer's phase offset removal using AP's antenna diversity.

an "occupant present" output if either the moving or stationary detection modules identify an occupancy; otherwise, it concludes an "occupant absent" output. Designing *LeakyBeam* attack involves two main challenges:

Challenge I: Modeling Human Presence under Environment and Device Ambiguity. WiFi signals are inevitably susceptible to ambiguities caused by multipath effects that vary with each environment and device layout. Traditional WiFi sensing approaches typically rely on environment-specific calibration, optimized device layout, or extensive training to compensate for these effects [37], [21], [22], [23], [38] – approaches that are impractical in adversarial scenarios where attackers are non-invasive, external observers. This complexity is particularly pronounced when detecting stationary victims, who induce only *subtle temporal variations* in the signal. To solve this challenge, *LeakyBeam* presents a unified detection framework that leverages *environment-independent* features derived from both the BFI *amplitude* and *phase* components, allowing for reliable occupant detection of both moving and stationary individuals, all without requiring environment-specific calibration (§IV-C).

Challenge II: Coarse-Grained BFIs with Random Phase Offsets. BFI, specifically the steering matrix $\mathbf{V}_k$ in it, is coarse-grained as it is compressed and quantized [28]. When we employ the phase information for fine-grained detection, an additional challenge arises from random phase offsets that vary across packets. Specifically, CSI measurements suffer from phase offsets because commodity WiFi transceivers are not tightly synchronized [39]. This time-variant random phase offset distorts the CSI, which in turn affects the BFI phase.

Considering the multipath effect, for the k-th subcarrier and i-th packet of antenna pair $\mathbf{s} = (p, q)$, connecting the p-th transmit and q-th receive antennas, the CSI can be written as:

$$\mathbf{H}_k(i, \mathbf{s}) = \sum_{l=1}^{L} \alpha_k^l(i, \mathbf{s}) e^{-j2\pi f \tau_k^l(i,\mathbf{s})}, \tag{IV.1}$$

where α_l and τ_l represent the attenuation factor and time-of-flight of the l-th propagation path. Commodity WiFi receivers and transmitters often exhibit synchronization mismatches, introducing a temporal random phase offset in CSI [39]. We use $\tilde{(\cdot)}$ to denote variables affected by phase offset:

$$\tilde{\mathbf{H}}_k(i, \mathbf{s}) = e^{-j\phi_{\text{offset}}} \sum_{l=1}^{L} \alpha_k^l(i, \mathbf{s}) e^{-j2\pi f \tau_k^l(i,\mathbf{s})}, \tag{IV.2}$$

where the phase offset term, ϕ_{offset}, is further expanded as:

$$\phi_{\text{offset}} = k(\epsilon_{\text{STO}} + \epsilon_{\text{SFO}}) + \epsilon_{\text{CFO}} + \epsilon_{\text{GR}}, \tag{IV.3}$$

and ϵ_{STO}, ϵ_{SFO}, and ϵ_{CFO} denote the symbol timing offset, sampling frequency offset, and carrier frequency offset between transceivers [40], [41], [42]. ϵ_{GR} denotes the offset introduced by Givens rotations during BFI computation [28].

The steering matrix $\mathbf{V}_k$, derived as the right singular matrix of $\mathbf{H}_k$, is consequently affected by these fluctuating phase offsets, resulting in a modified version denoted as $\tilde{\mathbf{V}}_k$. As shown in Figure 6(a), the inherent complexities introduced by the SVD process complicate the mitigation of phase offsets on $\tilde{\mathbf{V}}_k$. Traditional phase offset removal techniques that work directly on $\mathbf{H}_k$ do not effectively translate to $\tilde{\mathbf{V}}_k$ [24], [25], [39]. To address these challenges, we developed a novel feature extraction method tailored to remove the impact of phase offsets on $\tilde{\mathbf{V}}_k$ using the antenna-diversity of APs. The dynamic components attributable to human motion are further isolated and fused across subcarriers (§IV-B), thereby enhancing the accuracy and reliability of *LeakyBeam*'s detection even under coarse-grained BFI.

B. Phase-Offset-Free Feature Extraction

This module processes coarse-grained BFIs to produce a robust feature for further analysis, addressing Challenge II. We first introduce a novel variable, denoted as $\mathbf{R} = \tilde{\mathbf{V}}\Lambda\tilde{\mathbf{V}}^\dagger$, which is effective in canceling out phase offsets in BFI measurements, as shown in Figure 6(b).

Phase Offset Elimination. We first explain why our constructed $\mathbf{R}$ can cancel phase offset. Recall from Eq. II.1 that $\tilde{\mathbf{H}}_k = \tilde{\mathbf{U}}_k\Sigma_k\tilde{\mathbf{V}}_k^\dagger$ and $\tilde{\mathbf{U}}^\dagger\tilde{\mathbf{U}} = I$. Assuming $\Lambda = \Sigma_k^2$, we have:

$$\mathbf{R}_k = \tilde{\mathbf{V}}_k\Lambda\tilde{\mathbf{V}}_k^\dagger = \tilde{\mathbf{V}}_k\Sigma_k^2\tilde{\mathbf{V}}_k^\dagger = \tilde{\mathbf{V}}_k\Sigma_k\tilde{\mathbf{U}}_k^\dagger\tilde{\mathbf{U}}_k\Sigma_k\tilde{\mathbf{V}}_k^\dagger, \tag{IV.4}$$

Utilizing the property of conjugate transpose, i.e., $\mathbf{V}_k\Sigma\mathbf{U}_k^\dagger = (\mathbf{U}_k\Sigma\mathbf{V}_k^\dagger)^\dagger$, we can rewrite the equation as:

$$\mathbf{R}_k = (\tilde{\mathbf{U}}_k\Sigma_k\tilde{\mathbf{V}}_k^\dagger)^\dagger(\tilde{\mathbf{U}}_k\Sigma_k\tilde{\mathbf{V}}_k^\dagger) = \tilde{\mathbf{H}}_k^\dagger\tilde{\mathbf{H}}_k = \mathbf{H}_k^\dagger\mathbf{H}_k, \tag{IV.5}$$

indicating that $\mathbf{R}_k$ effectively represents the conjugate multiplication of $\mathbf{H}_k$. The temporal random phase offsets are the same across antennas on beamformer as they share the same RF oscillator [24], thus the conjugate multiplication between two antennas can remove the random phase offset. Importantly,

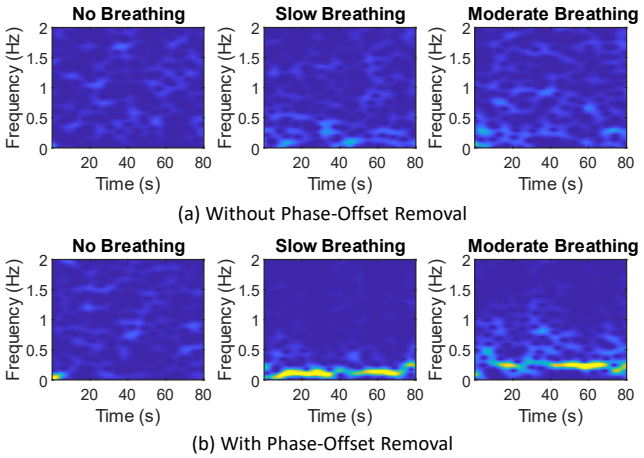

Fig. 7: Figure depicts the effect of phase offset removal on stationary occupancy detection: (a) without phase offset removal, breathing patterns are indistinguishable, and (b) with phase offset removal, breathing patterns become discernible, allowing the detection of stationary occupants.

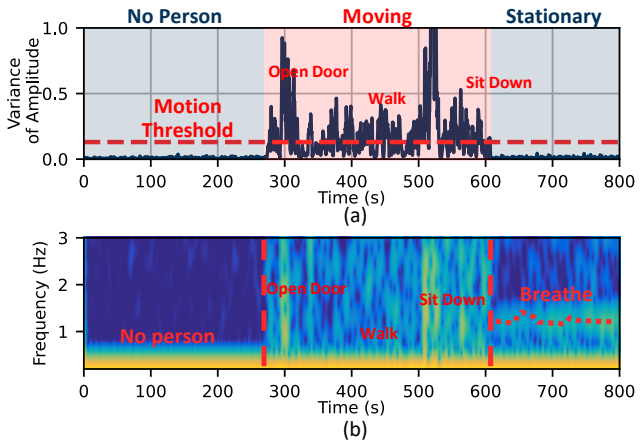

Fig. 8: Figure depicts the overview of *LeakyBeam*'s occupancy detection scheme, which (a) detects moving victims through variance in BFI amplitude and (b) detects subtle BFI phase changes caused by stationary victims.

this phase offset elimination does not require the attacker's access to the victim's devices, making it highly applicable in adversarial settings. It does, however, assume that the victim's AP (beamformer) is equipped with multiple antennas. Since APs with beamforming capability are inherently equipped with multiple antennas, this limitation should have minimal impact. Consequently, $\mathbf{R}_k$, being devoid of phase offset, becomes a robust and practical foundation for *LeakyBeam* attack.

Dynamic Component Extraction. By dividing signal paths into static and dynamic ones, $\mathbf{R}$ can be further expressed as:

$$\begin{aligned}\mathbf{R} &= \left(\mathbf{H}^{static\dagger} + \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l\dagger}\right)\left(\mathbf{H}^{static} + \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l}\right) \\ &= \underbrace{\mathbf{H}^{static\dagger}\mathbf{H}^{static}}_{\text{semi-static term}} + \underbrace{\sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l\dagger} \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l}}_{\text{higher-order minima}} \\ &+ \underbrace{\mathbf{H}^{static\dagger} \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l} + \mathbf{H}^{static} \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l\dagger}}_{\text{dynamic term}},\end{aligned} \tag{IV.6}$$

where L_{dynamic} represents the set of dynamic paths and $\mathbf{H}^{static}$ represents the static component. In this decomposition, the "semi-static term" arises from the product of static path components and remains relatively constant over short durations, lacking the transient information sought for occupancy detection. The "higher-order minima", attributed to the interaction of dynamic path components, are minimal and can typically be neglected. The "dynamic term", crucial for capturing human movement, reflects the interaction between static and dynamic paths, embodying the essence of temporal variations. To enhance detection accuracy, we isolate this dynamic information by minimizing the influence of the semi-static component. This is achieved by calculating the mean value of $\mathbf{R}$ over a defined period and subsequently subtracting this mean from $\mathbf{R}$ itself, effectively reducing the influence of the static component, thus:

$$\mathbf{R} = \mathbf{H}^{static\dagger} \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l} + \mathbf{H}^{static} \sum_{l\in L_{\text{dynamic}}} \mathbf{H}^{l\dagger}. \tag{IV.7}$$

The dynamic term can be further presented as follows:

$$\alpha^{static} \sum_{l\in L_{\text{dynamic}}} (\alpha_k^{l\dagger} e^{j\Delta\phi_k^l}) + \alpha^{static\dagger} \sum_{l\in L_{\text{dynamic}}} (\alpha_k^{l} e^{-j\Delta\phi_k^l}), \tag{IV.8}$$

where $\Delta\phi_k^l = 2\pi f(\tau_k^l - \tau_k^{static})$. Therefore, the dynamic feature that reflects the variation caused by the occupant is effectively extracted. The efficacy of phase-offset removal is demonstrated in Figure 7. With the phase offset removed, the spectrogram effectively highlights the presence of stationary victims, as detailed in §IV-C.

Outlier Packet Removal. Prior to the above analysis, we filter out packets with anomalous SNR drops (10 dB or more) compared to adjacent packets. This criterion for outlier detection is based on the observation that such drastic SNR changes are typically not characteristic of normal environmental variations or human movement, and therefore represent noise or errors that could impair subsequent analysis. Note that the required SNR data is readily available in the BFI packets, without the need for physical measurements by the attacker.

Subcarrier-wise Feature Fusion. To counteract the granularity loss due to compression, we utilize the diversity of subcarriers. Specifically, we employ Principal Component Analysis (PCA) to fuse information across all available subcarriers [43]. For instance, in the 20 MHz bandwidth setup of 802.11ac, PCA is applied to all 53 subcarriers. We then select the first principal component, which encapsulates the predominant and consistent signal variations indicative of occupancy states, enhancing the detection accuracy of our system.

C. Occupancy Detection

This module uses the phase-offset-free feature extracted in §IV-B as input to conduct the occupancy detection, identifying both moving and stationary victims, as shown in Figure 8.

Detecting Moving Victims. This submodule aims to detect moving victims by analyzing temporal variations in the BFI amplitude. Specifically, *LeakyBeam* computes the *variance* with a sliding window of the extracted feature **R**. As illustrated in Figure 8(a), movement by a victim results in a significantly larger variance compared to scenarios where no victim is present or the victim is stationary. For each sample window containing ten packets starting at time t, we compute:

$$\sigma_m^w(t) = \mathrm{Var}(\mathbf{R}(t), w), \tag{IV.9}$$

where $\mathrm{Var}(\cdot)$ represents the variance function, and w specifies the sample window of ten packets. Under conditions with a packet rate of 10 Hz, this corresponds to one second. If the packet rate drops (e.g., to 5 Hz), we adjust the duration of the window proportionally to still include ten packets to maintain detection accuracy. We also evaluate the impact of BFI rate with this adaptive sliding window approach in §V-C1.

To decide whether the variance indicates motion, we utilize threshold-based detection, i.e., motion is detected if $\sigma_m^w(t) > u$. The threshold u is calculated based on a set of long-term reference measurement windows, W_{ref}, as follows:

$$u = \mathrm{median}_{w \in W_{ref}}(\sigma_m^w(t)) + C \cdot \mathrm{MAD}_{w \in W_{ref}}(\sigma_m^w(t)), \tag{IV.10}$$

where $\mathrm{median}(\cdot)$ and $\mathrm{MAD}(\cdot)$ denote the median and median absolute deviation, respectively. The parameter C, as a conservativeness factor, is empirically set to 10 for accurate moving victim detection.

Detecting Stationary Victims. This submodule aims to detect stationary victims by analyzing phase changes, as shown in Figure 8(b). While stationary, victims continue to induce subtle phase changes, or Doppler shifts, in WiFi signals due to their breathing movements [44], [45], [20]. To capture these phase changes, we apply the short-term Fourier transform (STFT) to the first principal component. We use a Gaussian window of 10 seconds to enhance the resolution of the resulting spectrogram. In the absence of any occupants, the spectrogram exhibits noise with power scattered broadly across the frequency spectrum, as depicted in Figure 7. In contrast, the presence of a stationary victim, such as someone breathing, creates a distinct frequency response that focuses the power around specific frequencies. We measure the variance of power distribution, denoted as $\sigma_p(t)$, within each analysis window. Occupancy is confirmed when $\sigma_p(t)$ falls below a fixed threshold of 0.05, an empirical parameter that remains unchanged, for at least 50% of the detection window, indicating consistent respiratory patterns.

V. Evaluation

We present the evaluation of *LeakyBeam* through comprehensive real-world experiments, demonstrating its feasibility.

TABLE I: Table enumerates the eight tested AP models, along with their SoC manufacturer, MIMO configuration, the highest supported WiFi standard (ax/ac), and BFI packet rate.

No.	AP Model	SoC	MIMO	BFI Rate
1	Xiaomi AX6000	Qualcomm	[ax] 4×4	9.1 Hz
2	Redmi AX6000	MediaTek	[ax] 4×4	16.9 Hz
3	TP-LINK XDR5430	Qualcomm	[ax] 4×4	9.5 Hz
4	TP-LINK XDR6050	MediaTek	[ax] 4×4	17.0 Hz
5	NETGEAR AX5400	Broadcom	[ax] 4×4	10.0 Hz
6	NETGEAR AX6600	Broadcom	[ax] 4×4	9.8 Hz
7	ASUS AX86U	Broadcom	[ax] 4×4	10.3 Hz
8	D-Link DIR-823X	MediaTek	[ax] 3×3	14.2 Hz

TABLE II: Table enumerates the 10 tested client devices, along with their MIMO configuration and the highest supported WiFi standard (ax/ac).

No.	Client Type	Model	MIMO	Quantization
1	Laptop	Thinkpad X201	[ax] 2×2	[6,4] bits
2	Laptop	Alienware X17R2	[ax] 2×2	[6,4] bits
3	Laptop	ROG Strix G16	[ax] 2×2	[6,4] bits
4	Phone	iPhone 14 Pro	[ax] 2×2	[4,2] bits
5	Phone	Huawei P40 Pro	[ax] 2×2	[6,4] bits
6	Phone	Xiaomi 12S Pro	[ax] 2×2	[6,4] bits
7	Tablet	iPad Air 4	[ax] 2×2	[4,2] bits
8	Camera	Xiaomi CW500	[ax] 1×1	[6,4] bits
9	Smart Speaker	MI Speaker Pro 8	[ac] 1×1	[6,4] bits
10	Smart TV	Dangbei H3S	[ax] 2×2	[6,4] bits

A. Experiment Setup

Sniffer Implementation. Our sniffer setup comprises a Dell XPS 13 laptop running Ubuntu 21.04, equipped with an Intel AX210 NIC [46], [47]. Note that we do not modify the laptop's antenna setup, i.e., relying solely on its default configuration **with no external antennas**. The laptop is configured to monitor mode using the *iwconfig* command. We use Wireshark to capture BFI packets [35], while the decompression of the steering matrix was performed using pyshark based on the 802.11 protocol [48].

Data Collection. We evaluate *LeakyBeam* on a total of eight AP models from six popular AP vendors. These AP models vary in their WiFi chipsets and MIMO specifications [49], [50], [51], as depicted in Table I. Note that APs denoted as [ax] are also compatible with ac protocol. We evaluate all eight APs in three distinct environments, where the placement of clients was altered across three locations within each environment, resulting in a total of nine deployment layouts, as shown in Figure 9. To simulate the real-world attack scenario, the sniffer was positioned outside of the residence. We recruited eight volunteers to act as victims. They were informed that their occupancy state would be deduced but were not briefed on the specifics of the methodology employed. Overall, we collect around 1.6 million BFI packets of a total duration exceeding 49 hours. We conduct this study upon the approval of our institution's Institutional Review Board.

Furthermore, as depicted in Figure 10, we compare *LeakyBeam*'s performance against CSI-based baseline, focusing on their long-range attack performance (§V-C1) and their robustness against external interferences (§V-C2). We also evalu-

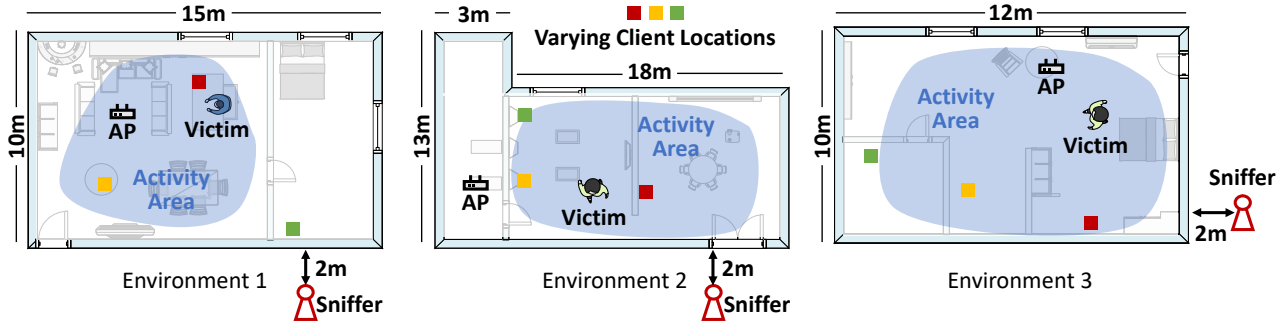

Fig. 9: Figure depicts the nine deployment scenarios across three environments, with three client locations per environment.

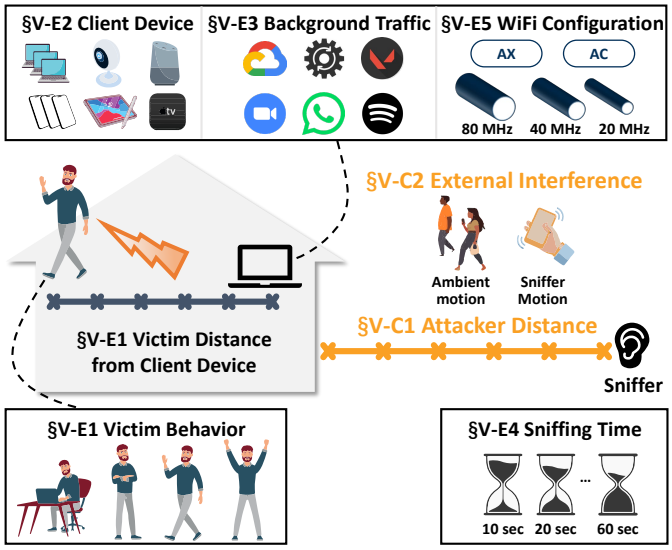

Fig. 10: Figure depicts the setup of *LeakyBeam*'s experiment conditions, specifically when comparing with CSI baseline (§V-C) and varying different conditions (§V-E).

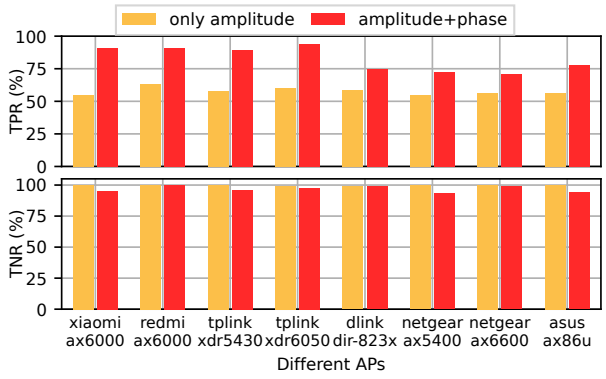

Fig. 11: Figure depicts the individual TPRs and FPRs for the eight AP models averaged over nine deployment settings.

ate *LeakyBeam*'s performance over various factors including different victim behaviors with varying distances from client device (§V-E1), diverse client devices (§V-E2), background traffic types(§V-E3), sniffing time (§V-E4) and WiFi configurations (§V-E5).

Performance Metrics. We define **True Positive Rate (TPR)** and **True Negative Rate (TNR)** to evaluate *LeakyBeam*'s performance on occupancy detection.[4] We consider a BFI trace to be a positive sample if *LeakyBeam* recognizes it as indicative of the presence of one or more victims (and a negative sample otherwise). Hence, we define TPR as the proportion of positive samples correctly identified in scenarios where a victim is actually present, and TNR as the proportion of negative samples correctly identified in scenarios where no victim is present. Besides TPR and TNR, we also utilize **Balanced Accuracy (annotated as Accuracy for abbreviation)**, where Accuracy $= \frac{1}{2}$ (TPR + TNR).

[4]TPR, TNR, False Positive Rate (FPR), and False Negative Rate (FNR) are commonly used in binary classification. FPR and FNR can be derived from TPR and TNR, with FPR = 1 - TNR and FNR = 1 - TPR.

B. Overall Attack Performance

To ensure diversity and minimize bias in our evaluation, *LeakyBeam*'s overall performance is assessed over eight AP models in three distinct environments. The client device placements varied across three distinct locations within each environment, creating nine different deployment layouts as depicted in Figure 9. This setup was used to evaluate *LeakyBeam*'s performance across varying conditions of environment and device ambiguity. The APs' specifications are shown in Table I. We collect 54 samples per layout per AP: 18 negative samples with no victim presence, 18 positive samples with a moving victim, and another 18 positive samples with a stationary victim, each lasting 30 seconds. This results in 32 hours of data comprising over 1.1 million packets.

The results are depicted in Figure 11, presenting TPRs and TNRs for each AP, averaged across nine layouts. We present two sets of results: one set for occupancy detection using only the moving occupant detection module based on amplitude, and another using both the moving and stationary occupant detection modules, which incorporates the use of offset-free phase. With only amplitude-based moving occupant detection, the TPR and TNR are 57.9% and 99.7%, respectively. This outcome occurs because the amplitude-only approach fails to detect stationary victims, thus yielding a relatively low TPR. However, with the integration of phase information, the TPR improves to 82.7% with a significant increase of 24.8% with a

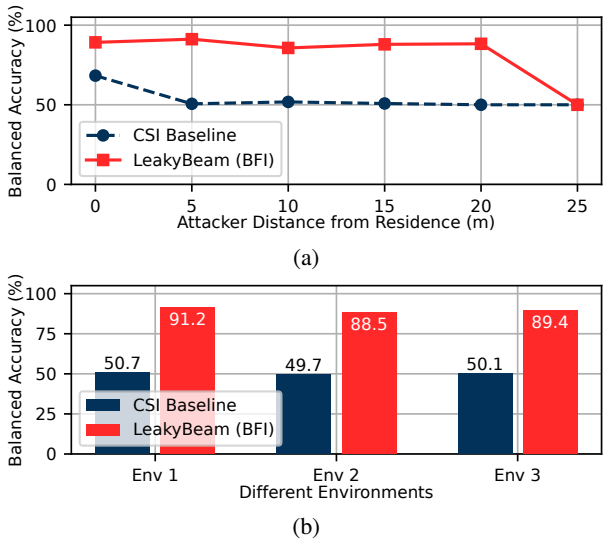

Fig. 12: Figure depicts *LeakyBeam*'s performance compared with CSI-based adversarial sensing attack under (a) varying distances and (b) different environments.

minor 3% decrease in TNR as a trade-off. This improvement demonstrates the effectiveness of *LeakyBeam*'s utilizing phase information to detect subtle channel variation caused by stationary victims. Particularly notable is the performance of the TP-LINK XDR6050, which achieved the best results with a TPR of 93.8% and a TNR of 97.6%. This superior performance may be attributed to more accurate BFI measurements and higher BFI rates (will be further explored in §V-E3). Overall, these results demonstrate *LeakyBeam*'s practicality and underscore the considerable privacy risks posed by the BFI side channel, affirming the need for effective countermeasures.

C. Comparison with CSI Baseline

Recall from §II-B that *LeakyBeam* represents a significant advancement over analog CSI-based attacks, particularly in its robustness to signal attenuation and external interferences. To further validate this claim, we compare *LeakyBeam* against a state-of-the-art CSI-based adversarial sensing attack [8]. For consistency in our experiments, we employ the same laptop equipped with an AX210 NIC for both BFI and CSI sniffing, thereby minimizing variations due to different radio configurations. CSI data acquisition is conducted using PicoScenes [52], a specialized tool for CSI research, as Wireshark does not support CSI sniffing. The experiments are conducted with Xiaomi AX6000 AP.

1) ***Impact of Attack Distance****:* We conduct experiments by varying the sniffer's distance from 0 to 25 meters from the victim's residence. At each distance, both the BFI and CSI sniffers are positioned identically to ensure a fair comparison. As the sniffers are placed outside of the residence, all measurements suffer from wall attenuation. Results shown in Figure 12(a) reveal that *LeakyBeam* maintains a high detection accuracy—around 88.3% even at 20 meters—owing to the digital nature of BFI packets whose content remains unaffected

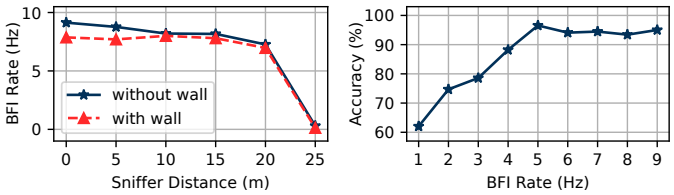

Fig. 13: Figure depicts (a) the sniffed BFI rates under different distances and (b) the detection accuracy under different BFI rates.

by the distance of the attacker. We observe a sharp decrease in accuracy at 25 meters where the sniffer fails to capture the leaked BFI packets. We will further detail the analysis of the packet loss. Our current setup does not use external antennas. Therefore, advanced attackers could further increase the attack range by employing external, or even directional, antennas.

Conversely, the CSI-based approach achieves only 68.3% accuracy, even at 0 meters, due to two primary factors: Firstly, the CSI baseline relies solely on amplitude variance to detect motion, which inherently leads to its *disability to detect stationary victims*. Secondly, wall attenuation significantly reduces the variance disparity in CSI signals when victims are present versus absent, insufficient to establish a reliable decision boundary for occupancy detection. This attenuation effect becomes more pronounced with increased distance, causing the performance to degrade to around 50.6%—equivalent to a random guess—at just 5 meters. Extended evaluations at 5 meters across different environments, depicted in Figure 12(b), consistently showed CSI baseline performance around 50%, substantially lower than *LeakyBeam* by 39.5%. This performance discrepancy highlights the superiority of *LeakyBeam* in real-world attack scenarios involving distance and physical barriers.

Analysis on BFI Packet Loss. Despite its resilience to attenuation, *LeakyBeam* can still suffer from *packet loss* at excessive distances. We assess how attack distance influences sniffed BFI rates by varying distances from 0 to 25 meters, as shown in Figure 13(a), considering scenarios with and without concrete walls. The findings reveal that when packets traverse a concrete wall, the average sniffed BFI rate decreases by about 0.7 Hz. Notably, even at a distance of 20 meters through a wall, the sniffed BFI rate remains as high as 7 Hz, with 76% of packets successfully retained, with a sharp decrease when the distance is further extended. This performance is likely because BFI packets are transmitted at a low MCS by default, which enables their capture from substantial distances.

To further assess the **impact of BFI rate** and packet loss, we evaluated detection accuracy by downsampling BFI traces from 9 Hz to rates between 1 Hz and 8 Hz. As shown in Figure 13(b), when the BFI rate exceeds 5 Hz, detection accuracy consistently exceeds 92.2%, demonstrating that *LeakyBeam*'s adaptive sliding window size effectively accommodates varying BFI rates. Additionally, *LeakyBeam*'s

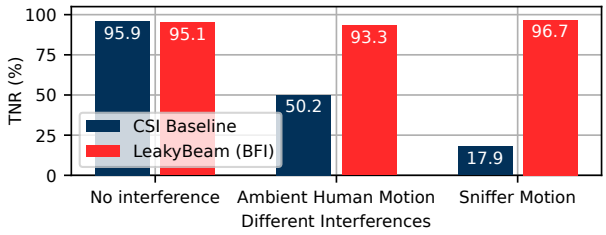

Fig. 14: Figure depicts the impact of external interferences on TNRs of CSI baseline and *LeakyBeam*.

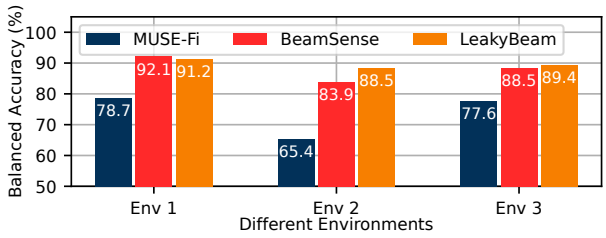

Fig. 15: Figure depicts *LeakyBeam*'s performance with previous BFI sensing methods.

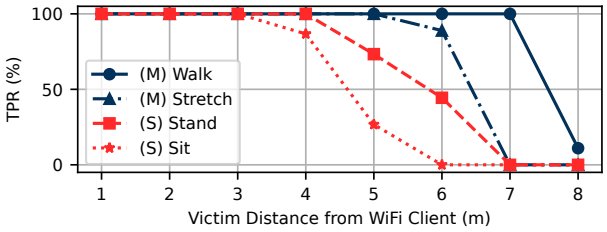

Fig. 16: Figure depicts *LeakyBeam*'s TPRs for different victim behaviors across different distances. (M) stands for motion behaviors and (S) stands for static behaviors.

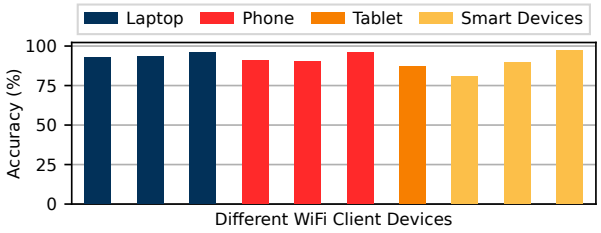

Fig. 17: Figure depicts *LeakyBeam*'s performance on 10 different client devices of four categories.

ability to maintain high accuracy at a 20-meter distance is now reasonable, given that the effective BFI rate of 7 Hz sustains accuracy. When the BFI rate drops below 5 Hz, we observe a positive correlation between the BFI rate and detection accuracy, with lower rates leading to reduced accuracy. Given this correlation, attackers could potentially estimate a confidence score based on the captured BFI rate when using the adaptive sliding window approach. They may also set a threshold to avoid detection when the BFI rate falls below a certain level, ensuring that lower BFI rates at extended distances, such as more than 25 meters, do not compromise detection quality.

2) ***Impact of External Interference****:* Besides the long-range capability, another notable advantage of *LeakyBeam* is its robustness to external interference. Specifically, while external interference typically induces unwanted variations in CSI signals—potentially misclassifying non-occupant activities as occupant detections—BFI remains unaffected by such disturbances. To systematically evaluate this characteristic, we consider two prevalent types of interference: ambient human motion near the sniffer, and motion of the sniffer itself (e.g., when hand-held by an attacker). We show the TNRs in Figure 14, showing that the CSI baseline is susceptible to external interferences as expected while confirming *LeakyBeam*'s robustness to these interference types.

D. Comparing LeakyBeam with Previous BFI Sensing Works

We now compare *LeakyBeam* with state-of-the-art BFI works, particularly BeamSense [19] and MUSE-Fi [53], which enable ubiquitous human sensing using BFI. Unlike MUSE-Fi, which relies on the *raw* steering matrix $\tilde{\mathbf{V}}$ for sensing, *LeakyBeam* utilizes the extracted feature $\mathbf{R}$ for occupant detection. As demonstrated in Figure 15, *LeakyBeam* achieves an average accuracy of 89.7% across environments, surpassing MUSE-Fi's accuracy of 73.9%. This significant improvement of 15.8% in accuracy highlights the benefits of our approach, which effectively mitigates phase offset and isolates dynamic signal components, leading to more reliable detection of stationary victims and overall enhanced system performance.

Meanwhile, BeamSense innovatively uses *bi-directional BFI* to reconstruct CSI for sensing, thus can achieve a comparative accuracy of 88.2%. However, BeamSense is *severely constrained* for attack purposes due to its reliance on bi-directional BFI. While downlink BFI can be easily extracted, uplink BFI is only accessible from 1% of WiFi clients, a limitation acknowledged by BeamSense. In our efforts to reproduce BeamSense for comparison, following the instructions from the BeamSense authors, the client's NIC had to be hacked and modified to enable beamformer capability to produce bi-directional BFI. While this approach is feasible in legitimate sensing scenarios, it is highly challenging and often impractical in adversarial settings where client devices are not physically accessible. In contrast, *LeakyBeam* avoids these constraints and assumptions, while maintaining effectiveness and practicality for attackers.

E. Differing Experimental Conditions

We evaluate *LeakyBeam*'s performance across several factors. By default, we perform our experiments on a representative AP - Xiaomi AX6000, and ThinkPad X201 as victim's device, in the first environment. We use *Accuracy* as the primary metric, calculated on a dataset *balanced* between negative and positive cases.

1) ***Impact of Victim Behavior and Distance****:* The victim's behavior and distance from the WiFi distance significantly influence signal variations. Behaviors with motion, such as walking, generally result in more pronounced signal variations.

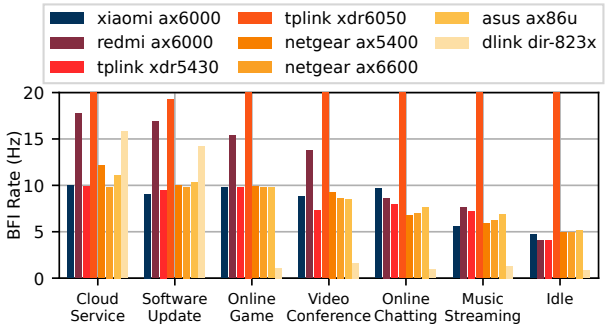

Fig. 18: Figure depicts the BFI rates of different background applications. For enhanced visualization, data values exceeding 20 Hz have been clipped.

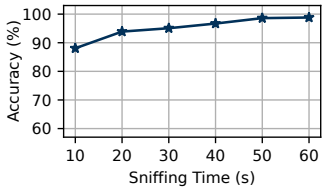

Fig. 19: Figure depicts detection accuracy under different sniffing durations.

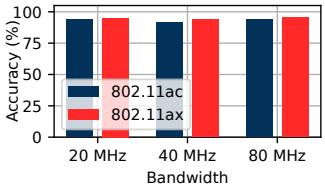

Fig. 20: Figure depicts detection accuracy under different WiFi configurations.

Similarly, closer proximity to the WiFi source strengthens signals reflected off human body, enhancing these variations. We evaluated the impact of these two critical factors by analyzing four common behaviors—walking, stretching (moving behaviors), and standing, sitting (stationary behaviors)—across distances from 1 to 8 meters. The results are shown in Figure 16. The results reveal that walking allows for detection up to 7 meters, whereas sitting, the least detectable behavior, maintains an accuracy of 86.7% at a distance of 4 meters, demonstrating the effectiveness of *LeakyBeam* to detect occupancy state.

2) ***Impact of Client Device****:* The performance of *LeakyBeam* varies depending on the client device (beamformee), as each device exhibits unique characteristics in its MIMO configurations and quantization levels. We conduct tests across a representative set of 10 common client devices, including laptops, phones, tablets, and smart devices such as cameras, speakers, and TVs. The specifications of these devices are detailed in Table II. The results in Figure 17 reveal an average accuracy of 91.4%. Notably, the smart speaker demonstrates the lowest accuracy, achieving only 80.8% due to its limited WiFi capabilities in various aspects, including a single antenna, relatively low traffic, and a bandwidth of merely 20 MHz. Overall, *LeakyBeam* proves to be highly effective, with seven out of the 10 tested devices achieving an accuracy higher than 89%. These results underscore *LeakyBeam*'s potential threat and its capability to effectively discern occupancy across a diverse array of commonly used wireless devices.

3) ***Impact of Background Traffic****:* The BFI rate can vary depending on the real-life background traffic, which ultimately influences *LeakyBeam*'s performance. To explore this impact, we conducted tests using six representative network applications: video conferencing, music streaming, online gaming, online chatting, cloud services, and software updating. These applications were run on all eight APs, with results displayed in Figure 18. Our analysis reveals that cloud services generate the highest average BFI rate at 13.4 Hz, while music streaming produces the lowest at 8.1 Hz. Among individual APs, only the Dlink AP exhibits a notably low BFI rate of 1.3 Hz in scenarios involving moderate to low network traffic. In contrast, all other seven APs consistently report a BFI exceeding 5 Hz, meaning an accuracy higher than 92.2% from Figure 13(b), demonstrating their susceptibility to *LeakyBeam* attacks with varying background traffic types. We also investigate conditions where the client is idle, and 7 out of 8 APs exhibit a BFI rate exceeding 4.1 Hz, corresponding to an detection accuracy of 88.2% from Figure 13(b). Additionally, we observed that when the channel is dynamic, i.e., with a moving occupant, the BFI rate could further increase by an average of 2.25 Hz across various background applications and APs, making it more susceptible to *LeakyBeam*, as detailed in Appendix A.

4) ***Impact of Sniffing Time****:* Our standard sniffing duration is set at 30 seconds. To assess the influence of different sniffing durations on system performance, we vary the sniffing time from 10 to 60 seconds. The results, shown in Figure 19, reveal that *LeakyBeam* maintains an accuracy of over 92% when the sniffing duration extends to 20 seconds or more. This demonstrates the system's efficiency in quickly detecting a victim's occupancy state.

5) ***Impact of WiFi Configuration****:* We evaluate its impact by testing both ax and ac protocols across varying bandwidths (20, 40, 80 MHz). The number of subcarriers associated with each bandwidth differs significantly—for example, 52 subcarriers for 20 MHz and 234 for 80 MHz under ac protocol. More subcarriers might lead to more spatial information being transmitted, potentially increasing the risk of privacy leakage. The results, detailed in Figure 20, demonstrate that the average accuracies for the ac and ax protocols are 92.87% and 95.06%, respectively. The accuracy variation across different bandwidths is a minor 2.6%, indicating that both ax and ac protocols maintain high performance even at the lowest bandwidth of 20 MHz, despite having fewer subcarriers. This demonstrates the robustness of *LeakyBeam* and highlights its potential as a significant threat in various WiFi environments.

VI. Defense Design

Given the potential threat for *LeakyBeam*-like attacks that exploit ubiquitous WiFi devices for unauthorized sensing and surveillance, we propose a defense mechanism featuring AP-defined spatial-temporal obfuscation. Designing this defense, however, is non-trivial, as it must effectively protect privacy while also ensuring practicality for seamless integration into existing WiFi APs and client devices. Specifically, our defense design includes the following features:

■ **Effective Privacy Preserving.** Our defense can obfuscate both the amplitude and phase of the steering matrix, rendering it *completely unintelligible* to potential attackers.

■ **Minimal Impact on Communication.** Leveraging the mathematical properties of the beamforming sounding process, our defense ensures that original steering matrices are *recoverable* by AP from their obfuscated versions through a *computationally efficient* method—multiplication by the matrix used for obfuscation. This preserves the integrity of beamforming while maintaining optimal network performance.

■ **Minimal Hardware Modification:** Our defense strategy *piggybacks on the existing spatial mapping mechanisms* embedded within the WiFi protocol, thus avoiding significant hardware modifications at the AP level. Additionally, from the perspective of client devices, they remain *unaffected* and *agnostic* to these defense modifications. This allows them to operate as usual without any alterations or performance impacts, ensuring seamless integration with the vast number of existing WiFi devices.

A. AP-Defined Spatial-Temporal Obfuscation

We now present the design details of our defense.

AP-Defined Signal Obfuscation Mechanism. Recall from §II-A, the receiver device calculates the CSI $\mathbf{H}_k$ using the known LTF $\mathbf{X}_k$ and derives steering matrix $\mathbf{V}_k$. In our defense design, instead of directly transmitting the LTF, our approach involves transmitting transformed $\mathbf{X}'_k = \mathbf{Q}_k^{obf}\mathbf{X}_k$, where $\mathbf{Q}_k^{obf} \in \mathbb{C}^{N_{TX}\times N_{TX}}$ is a *unitary* matrix for signal obfuscation. This transformation on LTF piggybacks on the spatial mapping mechanism[5] of the AP, thereby minimizing the hardware modification. This approach is also permissible under the 802.11 protocol, which specifies that spatial mapping is "not restricted" [28]. Our obfuscation matrix, $\mathbf{Q}_k^{obf}$, is designed to vary randomly across each BFI packet, enabling effective spatial-temporal obfuscation.

As the beamformee is unaware of the obfuscation matrix $\mathbf{Q}^{obf}$ decided by AP, the received signal at each beamformee antenna appears to be superimposed as if originating from a "virtual antenna". Consequently, the measured CSI is modified by this obfuscation and can be expressed as:

$$\mathbf{H}_k^{obf} = \mathbf{H}_k\mathbf{Q}_k^{obf}, \tag{VI.2}$$

where $\mathbf{H}_k^{obf}$ describes the channel between the virtual antennas and receiving antennas. Thus, an obfuscated steering matrix, denoted as $\mathbf{V}_k^{obf}$, will be sent back to the beamformer.

AP's BFI Recovery Mechanism from Obfuscation. Our subsequent analysis demonstrates that after receiving the BFI impacted by obfuscation, i.e., $\mathbf{V}_k^{obf}$, we can recover the original steering matrix $\mathbf{V}_k$ using

$$\mathbf{V}_k = \mathbf{Q}_k^{obf}\mathbf{V}_k^{obf}. \tag{VI.3}$$

Proof: Given that $\mathbf{H}_k = \mathbf{U}_k\Sigma_k\mathbf{V}_k^{\dagger}$, $\mathbf{H}^{obf}$ can be expressed as:

$$\mathbf{H}_k^{obf} = \mathbf{U}_k\Sigma_k\mathbf{V}_k^{\dagger}\mathbf{Q}_k^{obf} = \mathbf{U}_k\Sigma_k\left(\mathbf{Q}_k^{obf\dagger}\mathbf{V}_k\right)^{\dagger}. \tag{VI.4}$$

Since $\mathbf{Q}_k$ is a unitary matrix, $\mathbf{Q}_k^{obf\dagger}\mathbf{V}_k$ is also unitary. Therefore, the SVD of $\mathbf{H}_k^{obf}$ retains $\mathbf{U}_k$ as the left singular matrix. The right singular matrix of $\mathbf{H}_k^{obf}$ differs from that of $\mathbf{H}_k$ by a pre-multiplication by $\mathbf{Q}_k^{obf\dagger}$, i.e., $\mathbf{V}_k^{obf} = \mathbf{Q}_k^{obf\dagger}\mathbf{V}_k$. Thus, as outlined in Eq. VI.3, the AP can accurately recover the original $\mathbf{V}_k$ using this transformation. Meanwhile, attackers, with no knowledge of $\mathbf{Q}_k^{obf}$, cannot recover $\mathbf{V}$ from their sniffed $\tilde{\mathbf{V}}$. We present the detailed security analysis in Appendix B. Importantly, this method does not necessitate the client's awareness of $\mathbf{Q}_k^{obf}$, thereby eliminating the need for secure preliminary exchanges or synchronized algorithms between the AP and the client. This ensures that the approach remains compatible with unmodified beamformee devices, maintaining system functionality while enhancing security through spatial-temporal obfuscation.

[5]Given that the number of data streams might not align with the number of transmitting antennas, a mapping matrix can be applied for allocating the N_{STS} data streams across N_{TX} transmitting chains for onward transmission. The signal transmitted by AP can thus be expressed by:

$$\mathbf{X}'_k = \mathbf{Q}_k\mathbf{X}_k, \tag{VI.1}$$

where $\mathbf{Q}_k \in \mathbb{C}^{N_{TX}\times N_{STS}}$ represents the spatial mapping matrix.

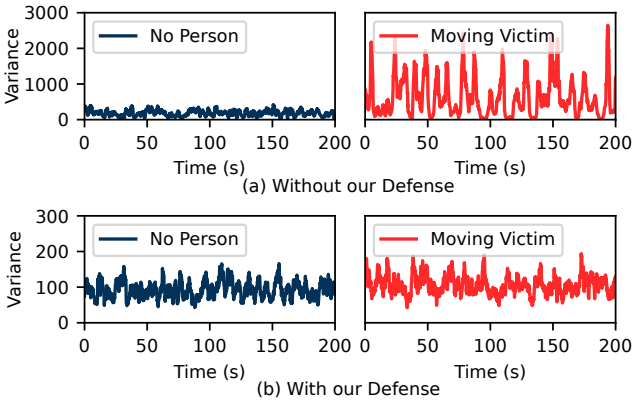

Fig. 21: Figure depicts the amplitude variance of BFI (a) without and (b) with our defense.

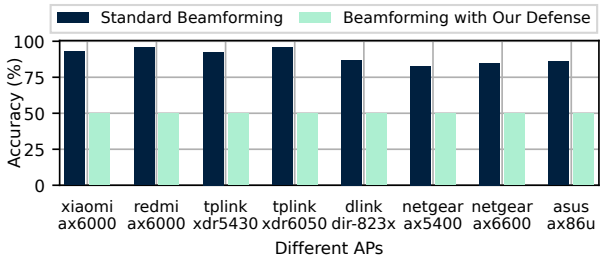

Fig. 22: Figure depicts the performance of our defense mechanism on occupancy detection, reducing accuracy to around 51%—the level of a random guess.

TABLE III: Comparison of Packet Error Rate (PER) and Throughput (TP) under different beamforming settings.

	No Beamforming		Standard Beamforming		Beamforming with our Defense	
SNR(dB)	PER	TP (Mbps)	PER	TP (Mbps)	PER	TP (Mbps)
20.7	9.54e-1	14.36	1.90e-3	309.49	2.26e-3	309.38
22.7	1.00e-1	278.95	1.20e-4	310.04	9.00e-5	310.05
24.7	3.00e-5	310.07	<1e-5	310.08	<1e-5	310.08

B. Defense Performance

Figure 21 illustrates the signal variance without and with our defense applied. With our defense, the variance between scenarios with no person and a moving victim becomes visually indistinguishable. To further evaluate the effectiveness of our defense, we assess its impact on the *LeakyBeam* attack, as depicted in Figure 22. The results demonstrate a significant degradation in the accuracy of *LeakyBeam*, dropping from 89.7% to approximately 51%, which is nearly equivalent to random guessing. This notable reduction in attack efficacy is consistent across all eight APs tested, highlighting the robustness and general effectiveness of our defense in mitigating the privacy risks posed by *LeakyBeam*.

C. Impact on Communication

Following Eq. VI.3, the AP should recover the original steering matrix $\mathbf{V}$ without discrepancies. However, in practical scenarios, the BFI undergoes quantization before being transmitted back to the beamformer. This quantization can cause the $\mathbf{V}$ recovered by the AP to *deviate slightly* from the originally quantified $\mathbf{V}$. To evaluate its impact on communication, we conducted simulations using MATLAB WLAN Toolbox across three scenarios: no beamforming, standard beamforming, and beamforming with our defense, incorporating quantization effects in a multipath fading channel. Results depicted in Figure 23 indicate that the variance of the constellation plot in both beamforming scenarios is nearly identical and significantly lower compared to the non-beamforming scenario, suggesting minimal disruption from our defense. Further analysis focuses on two key metrics: packet error rate and throughput. The simulations were configured with an MCS of index 4, bandwidth of 80 MHz, 4 transmit antennas, 2 receive antennas, with LDPC channel coding and the TGax Model-B for fading under three different SNR conditions. The channel was randomized across simulation instants. The results, shown in Table III, demonstrate that our defense maintains comparable performance to standard beamforming, emphasizing its practicality in preserving communication integrity while protecting user privacy. We provide additional simulation results in Appendix C.

D. Other Countermeasures.

We also propose two other potential defense strategies against *LeakyBeam* attack. First, we propose to *encrypt BFI Packets*. Encrypting these packets, for example, using WPA3, could effectively eliminate the leakage of BFI. However, this approach would necessitate firmware updates across all WiFi devices to support WPA3 – a formidable task given the billions of existing and legacy devices. Additionally, encryption introduces computational overhead that could burden resource-constrained devices. Similarly, while 802.11w protocol offers management frame protection through encryption and integrity checks, it significantly increases computational demands and is not universally supported, especially by common household devices [54].

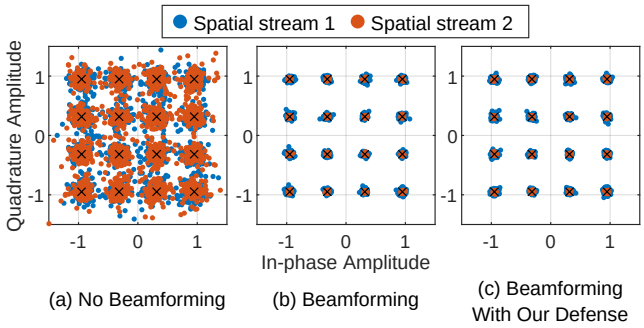

Fig. 23: Figure depicts the constellation plot comparing signal distribution across different beamforming conditions: (a) without beamforming, (b) with standard beamforming, and (c) with beamforming using our defense mechanism. Both beamforming conditions include quantization effects.

Second, we propose *injecting fake BFI packets for obfuscation* by either the client or the AP, each approach presenting potential issues. If the client injects fake BFIs, the AP might not be able to distinguish the fake packets from genuine ones, potentially using an incorrect steering matrix for beamforming and severely degrading communication performance. Conversely, if the AP injects fake BFI packets, an advanced attacker might still distinguish the fake packets from the ones from client devices by analyzing the channel characteristics [9].

VII. Extension to Occupant Tracking

We now discuss how *LeakyBeam* can be extended from *occupancy detection* to *tracking*. *LeakyBeam* can perform effective occupancy detection using existing WiFi clients in the residence as proximity sensors. The activation of these devices correlates with human presence within their detection range, typically at room level. Given the prevalence of WiFi devices, *LeakyBeam* can potentially evolve into an *adversarial tracking* system, if precise location data of WiFi devices within a residence are available. Previous research has demonstrated the feasibility of acquiring such locality information [33], [55], [8]. Hence, the trajectory of victims can be detected and inferred based on their proximity to active WiFi devices, posing a greater privacy threat.

We present a feasibility study to demonstrate the potential extension. Specifically, we conduct a preliminary evaluation in a two-story residential building spanning approximately 600 square meters. The arrangement, as detailed in Figure 24, consisted of five rooms, each equipped with one WiFi client device. The WiFi AP is placed in Room 1, while a passive

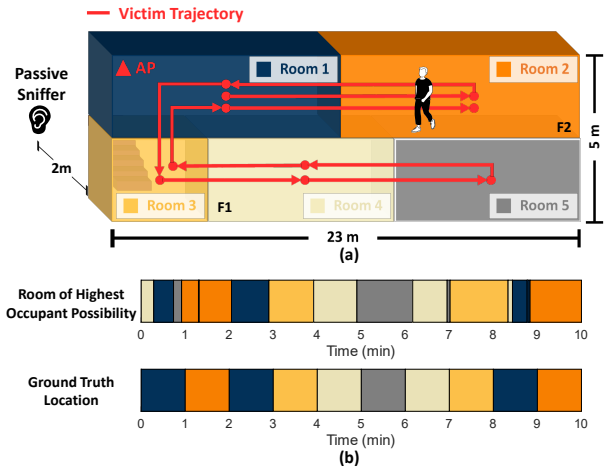

Fig. 24: Figure depicts (a) experimental setup in a two-floor building with five rooms, each containing a client device, and a passive BFI sniffer outside the building; (b) the comparison between the room of highest occupant possibility against the ground truth victim location.

WiFi sniffer is positioned two meters from the exterior of the house to collect BFI. We trace the trajectory of the victim over ten minutes as depicted by the red solid arrows in Figure 24(a). We also present the precision in correlating the victim's position with the nearest WiFi device in Figure 24(b). Through this study, we demonstrate the potential of extending *LeakyBeam* to an adversarial tracking system when combined with additional location knowledge of clients devices [33], [55], [8], thereby raising substantial privacy concerns.

VIII. Related Work

WiFi Sensing. Extensive research has focused on leveraging WiFi CSI for ubiquitous sensing applications, including gesture recognition, fall detection, and intrusion detection [56], [20], [57], [58], [59]. Recent research has explored BFI as a novel sensing medium [19], [36], [60], [61]. Early efforts predominantly focused on exploiting BFI amplitude [62], [63]. BFMSense [36] employed a BFM-ratio to counteract scaling factors in the feedback. While BeamSense [19] realizes CSI recovery from BFI, it requires bi-directional BFI, which is seldom available due to the scarcity of uplink BFI from most devices. And these approaches require direct control and calibration on the WiFi devices, which does not apply to *LeakyBeam*'s adversarial sensing scenario. With these constraints in mind, *LeakyBeam* innovatively leverages antenna and subcarrier diversities and derive robust features, enabling robust adversarial occupancy detection.

Adversarial WiFi Sensing. The WiFi signals can also be utilized by attackers to infer sensitive information [64], [65], [66], e.g., device location inference [67], [33] and keystroke and password inference [68], [69], [70], [71], [72], [73]. Recently, WiKI-Eve [62] extends keystroke attacks on phone by utilizing BFI packets. *LeakyBeam* fundamentally differs from WiKI-Eve as WiKI-Eve uses a *device-based* approach, requiring the victim's direct, physical interaction with the device transmitting BFI packets [74]. In contrast, *LeakyBeam* employs a *device-free* sensing method that utilizes signal reflections from the victim, bypassing the need for direct interaction. However, *LeakyBeam* faces the challenge of more subtle signals and being more prone to environmental ambiguities and multipath effects – challenges that *LeakyBeam* innovatively addresses through its design.

Defending Physical-based Adversarial Sensing. To counter physical-based adversarial sensing, Zhu et al.[8] employ a fake AP to inject cover packets, which might be compromised by adversaries who can distinguish the channel characteristics unique to the fake AP [9]. MIMOCrypt [75] secures privacy by encrypting CSI, necessitating key distribution and decryption at the client's end. PhyCloak [76] and Aegis [77] use a full-duplex radio to introduce channel variation. RF-Protect [10] and IRShield [9] adopt customized intelligent reflecting surfaces to obfuscate and prevent passive and active adversarial motion sensing. Unlike these approaches, *LeakyBeam* reuses the existing spatial mapping mechanism in WiFi protocols for spatial-temporal obfuscation, avoiding the need for additional costly hardware or complex cryptographic processes on client devices, thus ensuring minimal disruption while effectively preserving privacy.

IX. Conclusion

Through *LeakyBeam*, we demonstrate that plaintext BFI packets can pose significant privacy risks when exploited for adversarial human sensing. Our implementation and comprehensive real-world evaluation across eight APs in various deployment conditions confirm the practicality and threat of such attacks. Finally, we propose an effective defense strategy involving AP-based spatial-temporal obfuscation that safeguards privacy with minimal overhead. This work aims to spur further research into the privacy aspects of beamforming technologies, including for high-frequency WiFi like 60 GHz mmWave, which may rely heavily on beamforming.

Acknowledgment

We sincerely thank our anonymous reviewers for their valuable feedback. This paper is supported by the National Natural Science Foundation of China under grants U21A20462 and 62372400, "Pioneer" and "Leading Goose" R&D Program of Zhejiang under grant No. 2024C03287, and the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT (MSIT) under grant RS-2024-00464269.

References

[1] WiFi Alliance, "Wi-Fi by the numbers: Technology momentum," 2023, https://www.wi-fi.org/beacon/the-beacon/wi-fi-by-the-numbers-technology-momentum-in-2023.

[2] H. Cai, B. Korany, C. R. Karanam, and Y. Mostofi, "Teaching RF to sense without RF training measurements," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT'20*, vol. 4, no. 4, pp. 120:1–120:22, 2020.

[3] Z. Yang, Y. Zhang, K. Qian, and C. Wu, "Slnet: A spectrogram learning neural network for deep wireless sensing," in *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI'23*. USENIX Association, 2023, pp. 1221–1236.

[4] C. Wu, F. Zhang, Y. Fan, and K. J. R. Liu, "Rf-based inertial measurement," in *Proceedings of the ACM Special Interest Group on Data Communication, SIGCOMM'19*. ACM, 2019, pp. 117–129.

[5] F. Adib and D. Katabi, "See through walls with wifi!" in *ACM SIGCOMM 2013 Conference, SIGCOMM'13*. ACM, 2013, pp. 75–86.

[6] Q. Pu, S. Gupta, S. Gollakota, and S. N. Patel, "Whole-home gesture recognition using wireless signals," in *The 19th Annual International Conference on Mobile Computing and Networking, MobiCom'13*. ACM, 2013, pp. 27–38.

[7] J. Zhang, Z. Tang, M. Li, D. Fang, P. Nurmi, and Z. Wang, "Crosssense: Towards cross-site and large-scale wifi sensing," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking, MobiCom'18*. ACM, 2018, pp. 305–320.

[8] Y. Zhu, Z. Xiao, Y. Chen, Z. Li, M. Liu, B. Y. Zhao, and H. Zheng, "Et tu alexa? when commodity wifi devices turn into adversarial motion sensors," in *27th Annual Network and Distributed System Security Symposium, NDSS'20*. The Internet Society, 2020.

[9] P. Staat, S. Mulzer, S. Roth, V. Moonsamy, M. Heinrichs, R. Kronberger, A. Sezgin, and C. Paar, "Irshield: A countermeasure against adversarial physical-layer wireless sensing," in *43rd IEEE Symposium on Security and Privacy, SP'22*. IEEE, 2022, pp. 1705–1721.

[10] J. Shenoy, Z. Liu, B. Tao, Z. Kabelac, and D. Vasisht, "Rf-protect: privacy against device-free human tracking," in *Proceedings of the 2022 ACM Conference on Special Interest Group on Data Communication, SIGCOMM'22*. ACM, 2022, pp. 588–600.

[11] calder security, "Things burglars do before committing to the break in," 2024, https://www.caldersecurity.co.uk/things-burglars-do-before-committing-to-the-break-in//.

[12] S. W. R. Centre, "Stalking while staying at home," 2024, https://www.scottishwomensrightscentre.org.uk/covid-stalking-home/.

[13] S. L. . Key, "he psychology of home security: Understanding burglars' tactics and how to outsmart them," 2024, https://surelockkey.com/blog/the-psychology-of-home-security/.

[14] T. Independent, "Convicted burglars advise homeowners on how to protect themselves from break-ins, theft and robbery," 2017, https://www.independent.co.uk/news/uk/crime/burglary-advice-convicts-homeowners-protection-christmas-security-alarms-robbery-break-ins-theft-a8081671.html/.

[15] B. Xie, M. Cui, D. Ganesan, and J. Xiong, "Wall matters: Rethinking the effect of wall for wireless sensing," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT*, vol. 7, no. 4, pp. 190:1–190:22, 2023.

[16] X. Wang, K. Niu, J. Xiong, B. Qian, Z. Yao, T. Lou, and D. Zhang, "Placement matters: Understanding the effects of device placement for wifi sensing," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT*, vol. 6, no. 1, pp. 32:1–32:25, 2022.

[17] F. Zhang, Z. Chang, K. Niu, J. Xiong, B. Jin, Q. Lv, and D. Zhang, "Exploring lora for long-range through-wall sensing," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT'20*, vol. 4, no. 2, pp. 68:1–68:27, 2020.

[18] A. F. Molisch, *Wireless communications*. John Wiley & Sons, 2012, vol. 34.

[19] C. Wu, X. Huang, J. Huang, and G. Xing, "Enabling ubiquitous wifi sensing with beamforming reports," in *Proceedings of the ACM SIGCOMM 2023 Conference, SIGCOMM'23*. ACM, 2023, pp. 20–32.

[20] R. Xiao, J. Liu, J. Han, and K. Ren, "Onefi: One-shot recognition for unseen gesture via COTS wifi," in *The 19th ACM Conference on Embedded Networked Sensor Systems, SenSys'21*. ACM, 2021, pp. 206–219.

[21] S. Ding, Z. Chen, T. Zheng, and J. Luo, "Rf-net: a unified meta-learning framework for rf-enabled one-shot human activity recognition," in *The 18th ACM Conference on Embedded Networked Sensor Systems, SenSys'20*. ACM, 2020, pp. 517–530.

[22] J. Liu, W. Li, T. Gu, R. Gao, B. Chen, F. Zhang, D. Wu, and D. Zhang, "Towards a dynamic fresnel zone model to wifi-based human activity recognition," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT'23*, vol. 7, no. 2, pp. 65:1–65:24, 2023.

[23] X. Li, D. Zhang, Q. Lv, J. Xiong, S. Li, Y. Zhang, and H. Mei, "Indotrack: Device-free indoor human tracking with commodity wi-fi," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT'17*, vol. 1, no. 3, pp. 72:1–72:22, 2017.

[24] M. Kotaru, K. R. Joshi, D. Bharadia, and S. Katti, "Spotfi: Decimeter level localization using wifi," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM'15*. ACM, 2015, pp. 269–282.

[25] J. Xiong and K. Jamieson, "Arraytrack: A fine-grained indoor location system," in *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation, NSDI'13*. USENIX Association, 2013, pp. 71–84.

[26] K. Qian, C. Wu, Y. Zhang, G. Zhang, Z. Yang, and Y. Liu, "Widar2.0: Passive human tracking with a single wi-fi link," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys'18*. ACM, 2018, pp. 350–361.

[27] Y. Chae, Z. Lin, K. Bae, S. M. Kim, and P. Pathak, "mmcomb: High-speed mmwave commodity wifi backscatter," in *21st USENIX Symposium on Networked Systems Design and Implementation, NSDI'24*. USENIX Association, 2024.

[28] "IEEE standard for information technology - telecommuni-cations and information exchange between systems - local and metropolitan area networks–specific requirements - part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications," *IEEE Std 802.11-2020*, pp. 1–4379, 2021.

[29] F. Gringoli, M. Schulz, J. Link, and M. Hollick, "Free your CSI: A channel state information extraction platform for modern wi-fi chipsets," in *Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization, WiNTECH@MobiCom'19*. ACM, 2019, pp. 21–28.

[30] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Tool release: gathering 802.11n traces with channel state information," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 1, p. 53, 2011.

[31] Y. Xie, Z. Li, and M. Li, "Precise power delay profiling with commodity wifi," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, MobiCom'15*. ACM, 2015, pp. 53–64.

[32] Z. Li, B. Chen, X. Chen, H. Li, C. Xu, F. Lin, C. X. Lu, K. Ren, and W. Xu, "Spiralspy: Exploring a stealthy and practical covert channel to attack air-gapped computing devices via mmwave sensing," in *29th Annual Network and Distributed System Security Symposium, NDSS'22*. The Internet Society, 2022.

[33] A. Abedi and D. Vasisht, "Non-cooperative wi-fi localization & its privacy implications," in *The 28th Annual International Conference on Mobile Computing and Networking, MobiCom'22*. ACM, 2022, pp. 570–582.

[34] C. Li, M. Xu, Y. Du, L. Liu, C. Shi, Y. Wang, H. Liu, and Y. Chen, "Practical adversarial attack on wifi sensing through unnoticeable communication packet perturbation," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking, MobiCom'24*. ACM, 2024, pp. 373–387.

[35] Wireshark, "Wireshark," 2024, https://www.wireshark.org/.

[36] E. Yi, D. Wu, J. Xiong, F. Zhang, K. Niu, W. Li, and D. Zhang, "Bfmsense: Wifi sensing using beamforming feedback matrix," in *21st USENIX Symposium on Networked Systems Design and Implementation, NSDI'24*. USENIX Association, 2024.

[37] J. Wilson and N. Patwari, "Radio tomographic imaging with wireless networks," *IEEE Trans. Mob. Comput.*, vol. 9, no. 5, pp. 621–632, 2010.

[38] C. Wu, J. Xu, Z. Yang, N. D. Lane, and Z. Yin, "Gain without pain: Accurate wifi-based localization using fingerprint spatial gradient," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., IMWUT'17*, vol. 1, no. 2, pp. 29:1–29:19, 2017.

[39] X. Li, S. Li, D. Zhang, J. Xiong, Y. Wang, and H. Mei, "Dynamic-music: accurate device-free indoor localization," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp'2016*. ACM, 2016, pp. 196–207.

[40] X. Li, D. Zhang, J. Xiong, Y. Zhang, S. Li, Y. Wang, and H. Mei, "Training-free human vitality monitoring using commodity wi-fi devices," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 3, pp. 121:1–121:25, 2018.

[41] G. Chi, Z. Yang, J. Xu, C. Wu, J. Zhang, J. Liang, and Y. Liu, "Wi-drone: wi-fi-based 6-dof tracking for indoor drone flight control," in *The 20th Annual International Conference on Mobile Systems, Applications and Services, MobiSys'22*. ACM, 2022, pp. 56–68.

[42] K. Qian, C. Wu, Z. Yang, Y. Liu, and K. Jamieson, "Widar: Decimeter-level passive tracking via velocity monitoring with commodity wi-fi," in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc'17*. ACM, 2017, pp. 6:1–6:10.

[43] W. Wang, A. X. Liu, M. Shahzad, K. Ling, and S. Lu, "Understanding and modeling of wifi signal based human activity recognition," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, MobiCom'15*. ACM, 2015, pp. 65–76.

[44] K. Qian, C. Wu, Z. Zhou, Y. Zheng, Z. Yang, and Y. Liu, "Inferring motion direction using commodity wi-fi for interactive exergames," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI'17*. ACM, 2017, pp. 1961–1972.

[45] Y. Zhang, Y. Zheng, K. Qian, G. Zhang, Y. Liu, C. Wu, and Z. Yang, "Widar3.0: Zero-effort cross-domain gesture recognition with wi-fi," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 11, pp. 8671–8688, 2022.

[46] Dell, "Xps 13," 2016, https://dl.dell.com/manuals/all-products/esuprt_laptop/esuprt_xps_laptop/xps-13-9350-laptop_reference%20guide_en-us.pdf.

[47] Intel, "Intel® wi-fi 6e ax210," 2024, https://www.intel.com/content/www/us/en/products/sku/204836/intel-wifi-6e-ax210-gig/specifications.html.

[48] KimiNewt, "pyshark," 2024, https://github.com/KimiNewt/pyshark/.

[49] Broadcom, "Broadcom inc. – connecting everything," 2024, https://www.broadcom.com/.

[50] MediaTek, "Mediatek–powering the brands you love–incredible inside," 2024, https://www.mediatek.com/.

[51] Qualcomm, "Qualcomm: Intelligent computing everywhere," 2024, https://www.qualcomm.com/.

[52] Z. Jiang, T. H. Luan, X. Ren, D. Lv, H. Hao, J. Wang, K. Zhao, W. Xi, Y. Xu, and R. Li, "Eliminating the barriers: Demystifying wi-fi baseband design and introducing the picoscenes wi-fi sensing platform," *IEEE Internet Things J.*, vol. 9, no. 6, pp. 4476–4496, 2022.

[53] J. Hu, T. Zheng, Z. Chen, H. Wang, and J. Luo, "Muse-fi: Contactless muti-person sensing exploiting near-field wi-fi channel variation," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking, MobiCom'23*. ACM, 2023, pp. 75:1–75:15.

[54] A. Abedi and O. Abari, "Wifi says "hi!" back to strangers!" in *The 19th ACM Workshop on Hot Topics in Networks, HotNets'20*. ACM, 2020, pp. 132–138.

[55] R. S. Ayyalasomayajula, A. Arun, C. Wu, S. Rajagopalan, S. Ganesaraman, A. Seetharaman, I. K. Jain, and D. Bharadia, "Locap: Autonomous millimeter accurate mapping of wifi infrastructure," in *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI'20*. USENIX Association, 2020, pp. 1115–1129.

[56] Y. He, J. Liu, M. Li, G. Yu, J. Han, and K. Ren, "Sencom: Integrated sensing and communication with practical wifi," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking, MobiCom'23*. ACM, 2023, pp. 60:1–60:16.

[57] C. Han, K. Wu, Y. Wang, and L. M. Ni, "Wifall: Device-free fall detection by wireless networks," in *2014 IEEE Conference on Computer Communications, INFOCOM'14*. IEEE, 2014, pp. 271–279.

[58] F. Zhang, D. Zhang, J. Xiong, H. Wang, K. Niu, B. Jin, and Y. Wang, "From fresnel diffraction model to fine-grained human respiration sensing with commodity wi-fi devices," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 1, pp. 53:1–53:23, 2018.

[59] R. H. Venkatnarayan, G. Page, and M. Shahzad, "Multi-user gesture recognition using wifi," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys'18*. ACM, 2018, pp. 401–413.

[60] S. Kato, T. Murakami, T. Fujihashi, T. Watanabe, and S. Saruwatari, "CBR-ACE: counting human exercise using wi-fi beamforming reports," *J. Inf. Process.*, vol. 30, pp. 66–74, 2022.

[61] T. Kanda, T. Sato, H. Awano, S. Kondo, and K. Yamamoto, "Respiratory rate estimation based on wifi frame capture," in *19th IEEE Annual Consumer Communications & Networking Conference, CCNC'22*. IEEE, 2022, pp. 881–884.

[62] J. Hu, H. Wang, T. Zheng, J. Hu, Z. Chen, H. Jiang, and J. Luo, "Password-stealing without hacking: Wi-fi enabled practical keystroke eavesdropping," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS'23*. ACM, 2023, pp. 239–252.

[63] F. Meneghello, M. Rossi, and F. Restuccia, "Deepcsi: Rethinking wi-fi radio fingerprinting through MU-MIMO CSI feedback deep learning," in *42nd IEEE International Conference on Distributed Computing Systems,ICDCS'22*. IEEE, 2022, pp. 1062–1072.

[64] R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, "Packet-level signatures for smart home devices," in *27th Annual Network and Distributed System Security Symposium, NDSS'20*. The Internet Society, 2020.

[65] W. Zhang, Y. Meng, Y. Liu, X. Zhang, Y. Zhang, and H. Zhu, "Homonit: Monitoring smart home apps from encrypted traffic," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS'18*. ACM, 2018, pp. 1074–1088.

[66] A. Acar, H. Fereidooni, T. Abera, A. K. Sikder, M. Miettinen, H. Aksu, M. Conti, A. Sadeghi, and A. S. Uluagac, "Peek-a-boo: i see your smart home activities, even encrypted!" in *13th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec'20*. ACM, 2020, pp. 207–218.

[67] J. Freudiger, "How talkative is your mobile device?: an experimental study of wi-fi probe requests," in *Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks, WiSec'15*. ACM, 2015, pp. 8:1–8:6.

[68] K. Ali, A. X. Liu, W. Wang, and M. Shahzad, "Keystroke recognition using wifi signals," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking, MobiCom'15*. ACM, 2015, pp. 90–102.

[69] M. Li, Y. Meng, J. Liu, H. Zhu, X. Liang, Y. Liu, and N. Ruan, "When CSI meets public wifi: Inferring your mobile phone password via wifi signals," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS'16*. ACM, 2016, pp. 1068–1079.

[70] S. Fang, I. D. Markwood, Y. Liu, S. Zhao, Z. Lu, and H. Zhu, "No training hurdles: Fast training-agnostic attacks to infer your typing," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS'18*. ACM, 2018, pp. 1747–1760.

[71] E. Yang, Q. He, and S. Fang, "WINK: wireless inference of numerical keystrokes via zero-training spatiotemporal analysis," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS'22*. ACM, 2022, pp. 3033–3047.

[72] E. Yang, S. Fang, I. D. Markwood, Y. Liu, S. Zhao, Z. Lu, and H. Zhu, "Wireless training-free keystroke inference attack and defense," *IEEE/ACM Trans. Netw.*, vol. 30, no. 4, pp. 1733–1748, 2022.

[73] Z. Zhang, N. Avazov, J. Liu, B. Khoussainov, X. Li, K. Gai, and L. Zhu, "Wipos: A POS terminal password inference system based on wireless signals," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 7506–7516, 2020.

[74] J. Xiao, H. Li, M. Wu, H. Jin, M. J. Deen, and J. Cao, "A survey on wireless device-free human sensing: Application scenarios, current solutions, and open issues," *ACM Comput. Surv.*, vol. 55, no. 5, pp. 88:1–88:35, 2023.

[75] J. Luo, H. Cao, H. Jiang, Y. Yang, and Z. Chen, "Mimocrypt: Multi-user privacy-preserving wi-fi sensing via MIMO encryption," in *IEEE Symposium on Security and Privacy, SP 2024*. IEEE, 2024, pp. 2812–2830.

[76] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, and A. Arora, "Phycloak: Obfuscating sensing from communication signals," in *13th USENIX Symposium on Networked Systems Design and Implementation, NSDI'16*. USENIX Association, 2016, pp. 685–699.

[77] Y. Yao, Y. Li, X. Liu, Z. Chi, W. Wang, T. Xie, and T. Zhu, "Aegis: An interference-negligible RF sensing shield," in *2018 IEEE Conference on Computer Communications, INFOCOM'18*. IEEE, 2018, pp. 1718–1726.

APPENDIX A
IMPACT OF MOVING OCCUPANT ON BFI RATE

We investigate how the victim's movement affects the BFI rate across different client background traffic scenarios. We conducted BFI rate measurements in the same manner as in §V-E3, but with a moving occupant. The results, shown in Figure 25(a), indicate that the BFI rate is consistently higher than 4.6 Hz, except for the Dlink AP with an idle client. Moreover, we observe that under the same background traffic conditions, the BFI rate with a moving occupant is consistently higher than in scenarios with no occupant or a stationary occupant, as illustrated in Figure 25(b), with an average increase of 2.3 Hz. This increase occurs because the APs, detecting dynamic changes in the channel, raise their sounding frequency to allow for timely adjustments in beam alignment, thereby making them more susceptible to *LeakyBeam* attack.

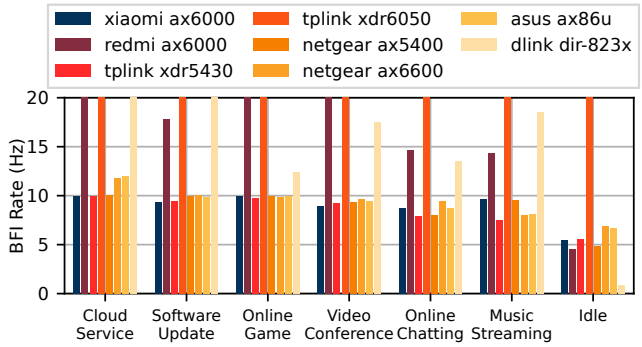

(a)

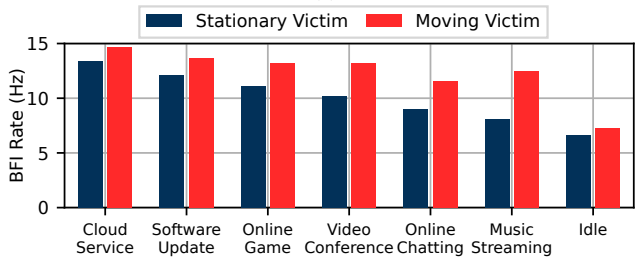

(b)

Fig. 25: Figure depicts (a) BFI rates of different background applications with a moving victim (values > 20 Hz clipped for clarity), and (b) comparison of BFI rates between stationary and moving victims.

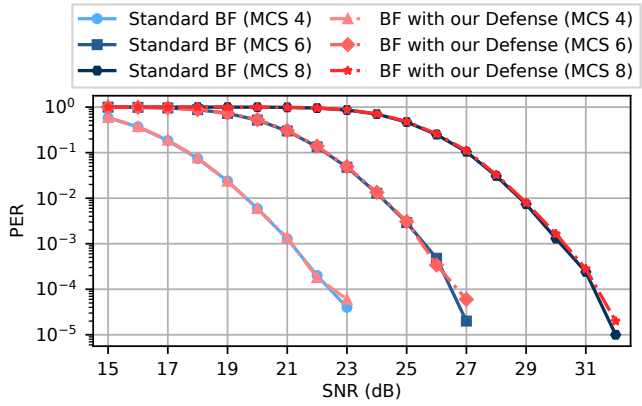

Fig. 26: Figure depicts the impact of *LeakyBeam*'s defense on PER under different modulation schemes. Beamforming (BF) with our defense shows comparable results to standard BF.

Appendix B
Security Analysis on Countermeasure

We provide detailed security analysis on *LeakyBeam*'s countermeasure design by proving the following theorem.

Theorem B.1: If the obfuscation matrix $\mathbf{Q}_k^{obf}$ is chosen uniformly at random from all possible candidates, then defense is secure, i.e.,

$$P(\mathbf{V}_k = V_k | \mathbf{V}_k^{obf} = V_k^{obf}) = P(\mathbf{V}_k = V_k)$$

for all possible V_k and all possible V_k^{obf}, meaning eavesdropping givens no advantage to guess V_k.

Proof: Suppose there are n candidates for the obfuscation matrix, then

$$P(\mathbf{Q}_k^{obf} = Q_k^{obf}) = \frac{1}{n},$$

where Q_k^{obf} is one of the candidate and is an unitary matrix. Since $\mathbf{Q}_k^{obf}$ is chosen independently of $\mathbf{V}_k$, we have

$$\begin{aligned} &P(\mathbf{V}_k^{obf} = V_k^{obf} \cap \mathbf{Q}_k^{obf} = Q_k^{obf}) \\ &= P(\mathbf{V}_k = Q_k^{obf} V_k^{obf} \cap \mathbf{Q}_k^{obf} = Q_k^{obf}) \\ &= P(\mathbf{V}_k = Q_k^{obf} V_k^{obf}) P(\mathbf{Q}_k^{obf} = Q_k^{obf}) \\ &= P(\mathbf{V}_k = Q_k^{obf} V_k^{obf})(1/n), \end{aligned} \tag{B.1}$$

given that $\mathbf{V}_k^{obf} = \mathbf{Q}_k^{obf\dagger} \mathbf{V}_k$.

If Q_k^{obf} runs all possible candidate matrices, then the corresponding $Q_k^{obf} V_k^{obf}$ runs all possible messages:

$$\sum_{Q_k^{obf}} P(\mathbf{V}_k = Q_k^{obf} V_k^{obf}) = 1. \tag{B.2}$$

With Eq. (B.1) and Eq. (B.2), we have:

$$\begin{aligned} P(\mathbf{V}_k^{obf} = V_k^{obf}) &= \sum_{Q_k^{obf}} P(\mathbf{V}_k^{obf} = V_k^{obf} \cap \mathbf{Q}_k^{obf} = Q_k^{obf}) \\ &= \frac{1}{n} \sum_{Q_k^{obf}} P(\mathbf{V}_k = Q_k^{obf} V_k^{obf}) = \frac{1}{n}, \end{aligned} \tag{B.3}$$

The definition of conditional probability and the independence of $\mathbf{Q}_k^{obf}$ and $\mathbf{V}_k$ yield:

$$\begin{aligned} &P(\mathbf{V}_k = V_k | \mathbf{V}_k^{obf} = V_k^{obf}) P(\mathbf{V}_k^{obf} = V_k^{obf}) \\ &= P(\mathbf{V}_k^{obf} = V_k^{obf} \cap \mathbf{V}_k = V_k) \\ &= P(\mathbf{Q}_k^{obf} = Q_k^{obf} \cap \mathbf{V}_k = V_k) \\ &= P(\mathbf{Q}_k^{obf} = Q_k^{obf}) P(\mathbf{V}_k = V_k). \end{aligned} \tag{B.4}$$

Since $P(\mathbf{V}_k^{obf} = V_k^{obf}) = \frac{1}{n} = P(\mathbf{Q}_k^{obf} = Q_k^{obf})$, we can obtain:

$$P(\mathbf{V}_k = V_k | \mathbf{V}_k^{obf} = V_k^{obf}) = P(\mathbf{V}_k = V_k). \tag{B.5}$$

The proof is hence completed.

Appendix C
Impact of *LeakyBeam*'s Defense under Different Modulation Scenarios

We evaluate the impact of *LeakyBeam*'s defense on communication performance across various WiFi modulation scenarios, specifically for MCS indices 4, 6, and 8, corresponding to 16-QAM, 64-QAM, and 256-QAM, respectively. Similar to §VI-C, the simulations utilize MATLAB's WLAN Toolbox, employing LDPC channel coding and the TGax Model-B channel model to introduce realistic fading and noise variations. As depicted in Figure 26, our results demonstrate that *LeakyBeam*'s defense strategy sustains performance comparable to standard beamforming even under higher modulation schemes, aligning with our analysis in §VI-A.