© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at https://doi.org/10.1109/TWC.2024.3485477.

Physical-Layer Privacy via Randomized Beamforming Against Adversarial Wi-Fi Sensing: Analysis, Implementation, and Evaluation

Marco Cominelli, *Member, IEEE*, Shaghayegh Shahcheraghi, Jakob Link, Matthias Hollick, *Member, IEEE*, Federico Cerutti, *Member, IEEE*, Francesco Gringoli, *Senior Member, IEEE*, and Arash Asadi, *Senior Member, IEEE*

***Abstract*—Wi-Fi sensing applications have achieved remarkable results over the last decade, offering accurate device-free localization and gesture recognition capabilities. Indeed, Wi-Fi sensing has quickly become a critical field of research for future communication systems under the paradigm known as joint communication and sensing. However, device-free wireless sensing can also be exploited for malign purposes against unaware victims, and the omnipresence of Wi-Fi transceivers poses a significant threat to people's privacy. Therefore, it is essential to develop functional solutions that can effectively thwart wireless sensing. All the current attempts to hinder illegitimate wireless sensing rely on specialized hardware deployed in the environment, but their cost and complexity can undermine widespread deployment. In this paper, we explore the possibility of using native capabilities of Wi-Fi systems, namely beamforming, to thwart wireless sensing. To this end, we propose for the first time a solution that enables complete control over the beamforming in commercial Wi-Fi devices. On top of that, we build BeamDancer, which randomizes beamforming vectors to inhibit channel fingerprinting. We empirically demonstrate the effectiveness of the proposed solution against three different wireless sensing techniques, both data-driven and model-based, while preserving almost entirely the legitimate Wi-Fi traffic at the same time.**

***Index Terms*—Physical-layer privacy, Wi-Fi sensing, beamforming, channel state information, deep learning**

I. Introduction

DEVICE-FREE Wi-Fi sensing exploits wireless communication signals for a wide range of applications, from the detection of human *motions* [1] to *emotions* [2]. The term "device-free" implies that the sensed target does not have to carry any radio transceiver; indeed, most sensing applications rely on the physical interaction of the communication signals with the human body and the surrounding environment. Many works already proved that wireless sensing based on Wi-Fi signals can be highly accurate for a variety of use cases, including user tracking with decimeter-level accuracy [3], human presence detection [4] or activity recognition [5], fine-grained gesture recognition (e.g., hands or lips movement [6, 7]) or vital sign monitoring [8]. If implemented on commodity wireless communication devices, such capabilities can catalyze progress in critical areas such as patient and elderly monitoring [9], device-free localization of human and robots [10], and camera-less movement monitoring [11]. However, when leveraging ubiquitous communication technologies like Wi-Fi, device-free sensing applications also threaten users' privacy because they are incredibly accurate and potentially undetectable. Figure 1 shows a scenario in which Wi-Fi sensing can be used to undermine people's privacy. In this scenario, a malicious user secretly exploits Wi-Fi signals in the wild to gain unauthorized—possibly sensitive—information about a target person, who has no clue she is being monitored.

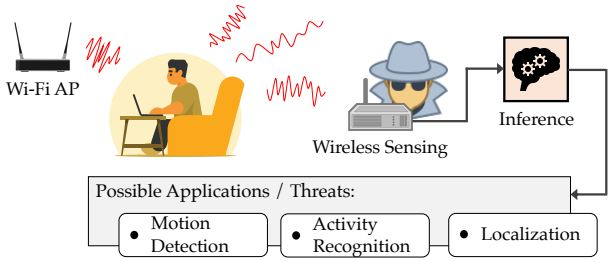

Fig. 1: Adversarial wireless sensing scenario. An eavesdropper exploits the physical properties of the received Wi-Fi signals to gain unauthorized information about the surroundings.

Today, while novel wireless sensing techniques are thriving, research on the security and privacy aspects still lags behind. It is interesting to observe that some state-of-the-art frameworks can already run wireless sensing operations on commercial off-the-shelf (COTS) devices [12]; however, to our knowledge, no countermeasures are currently available on commodity devices. So far, systems that protect users against illegitimate wireless sensing must rely on expensive additional hardware. In this work, instead, we present for the first time a system that can be potentially implemented into the vast majority of COTS devices to enhance users' privacy with no downsides from the communication perspective.

Motivation. Although wireless sensing is generally deemed less privacy-invasive than camera-based systems, it also rep-

M. Cominelli is with the DEIB, Politecnico di Milano, 20133 Milano, Italy. E-mail: marco.cominelli@polimi.it

S. Shahcheraghi is with the Ohio State University, Columbus, OH 43210, United States. E-mail: shahcheraghi.1@osu.edu

J. Link, M. Hollick are with the Computer Science Department, Technische Universität Darmstadt, 64289 Darmstadt, Germany. E-mails: {jlink, mhollick}@seemoo.tu-darmstadt.de

F. Cerutti and F. Gringoli are with the DII, University of Brescia, 25123 Brescia, Italy. E-mails: {federico.cerutti, francesco.gringoli}@unibs.it

A. Asadi is with the EEMCS Faculty, TU Delft, 2628 Delft, Netherlands. E-mail: a.asadi@tudelft.nl

This work was conducted while M. Cominelli was with the University of Brescia, and S. Shahcheraghi and A. Asadi were with TU Darmstadt.

resents an entirely novel attack surface against privacy that should not be underestimated. It has been shown in the literature that adversaries can obtain potentially compromising information of unsuspecting users by just inspecting the Wi-Fi frames sniffed over the air. Sample applications include: burglars monitoring people's living habits in their homes [13]; hand-gesture inference to unlock smartphones [14, 15]; vital signs monitoring to reveal health conditions [16]; illegitimate tracking and identification of people, e.g., via gait analysis [17, 18]; industrial espionage, e.g., manufacturers obtaining information about even active sections of the plant, or hotels measuring the room occupancy of their competitors; at scale, the same capabilities could be used by intelligence agencies for surveillance.

These attacks are easy enough to mount when the adversary can (i) access the raw channel state information (CSI), which is readily available on many commercial devices, or (ii) further process the CSI via well-established parameter estimation methods for Doppler, angle of arrival (AoA) and time of flight (ToF) calculation. It is the availability of inexpensive hardware and accessibility of well-established machine learning (ML) and signal processing methods that make the attacks possible, and their consequences are vast and hard to predict.

Challenges and Novelty. Preventing Wi-Fi sensing privacy attacks can be a daunting effort. Unlike attacks exploiting radars and cameras, an adversary performing wireless sensing can listen to the Wi-Fi communications passively without transmitting any signal, thus remaining fully hidden. Sensing attacks can be mounted even outside the victim's premises, through walls, or using drones flying by. Given the pervasiveness of wireless communication in our daily lives, classical defense mechanisms such as jamming and geofencing are, unfortunately, impractical. On the one hand, jamming has the undesired side effect of severely hindering Wi-Fi communications. On the other hand, trying to set up geofences, i.e., to physically secure the perimeter of the Wi-Fi deployment, does not scale with dense deployments, since defining disjoint geofences can be challenging if not impossible. Optimal solutions should allow wireless devices to communicate but stop adversaries from illegitimate sensing without imposing strong additional constraints. Furthermore, to facilitate broad adoption at a consumer level, such solutions must be cost-effective for the users.

While privacy complications of wireless sensing have already been identified [19, 20], the literature on privacy-oriented signal processing is still limited. Prior work against illegitimate *device-based* sensing only prevents the adversaries from localizing and tracking wireless devices carried by a victim [19, 21–24]. Although crucial, these methods cannot protect the users against novel device-free wireless sensing methods. Only a few works proposed countermeasures against illegitimate *device-free* sensing so far [25–28], as we discuss in detail in Section VII. In [13, 25, 26], the authors rely on secondary wireless transceivers to impede illegitimate sensing. Similarly, the works of [27, 28] propose to leverage reconfigurable reflecting surfaces [27] or radio elements [28] as a passive secondary device to impact the propagation of the signal in the environment, thus avoiding illegitimate sensing. Using secondary devices can be effective; however, such solutions exhibit higher costs for the users—especially if full-duplex custom transceivers are needed—and must be synchronized with the legitimate transmitter, complicating their operation in everyday scenarios. On the contrary, this work aims to show that privacy can be achieved at the physical layer cost-effectively using commodity devices without introducing additional components.

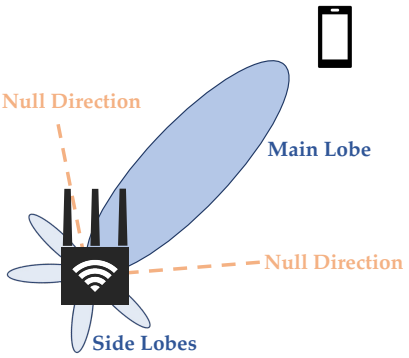

Fig. 2: Beamforming enables high directivity by combining constructive and destructive interference in different directions.

Main contributions. To our knowledge, `BeamDancer` is the first defense mechanism built on top of commodity devices against adversarial Wi-Fi sensing. The first goal—and the main novelty aspect—of this work is addressing the technological challenges that arise when implementing a privacy-preserving solution on COTS devices. In addition, we present a methodically simple yet practical approach to protect users' privacy against adversarial Wi-Fi sensing.

To implement `BeamDancer`, we first develop and validate a tool to access and modify the beamforming configuration of COTS devices, which we name `Beamer` (Section IV). Then, we evaluate `BeamDancer` experimentally against both data-driven and model-based sensing techniques, including (i) a supervised ML classifier for localization applications; (ii) an unsupervised ML clustering algorithm for localization; and (iii) a model-based speed estimation technique based on micro-Doppler analysis. Our experiments prove `BeamDancer` can thwart localization accuracy and distort the velocity estimation entirely with a limited impact on the quality of the communication link. Our analysis indicates that `BeamDancer` maintains on average more than 96% of the original packet delivery ratio in line-of-sight (LoS) scenarios and only slightly degrades the performance in most non-line-of-sight (NLoS) scenarios.

II. Background

A. Beamforming

Beamforming is commonly utilized in antenna arrays and multiple-input multiple-output (MIMO) systems for directional signal transmission or reception [29]. As shown in Fig. 2, the radiation pattern is dynamically configured to focus transmission and reception only towards certain directions by appropriately combining the signals at different antenna array elements. As a result, the capability of *steering* the main

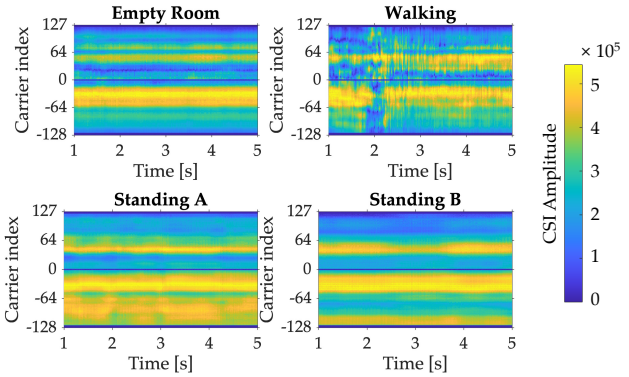

Fig. 3: The CSI captures the physical properties of the wireless channel in time and frequency; hence, the resulting CSI depends on the activity or the location of the subject (A and B denote two separate places in the same room). CSI amplitude values are dimensionless quantities extracted directly from the Wi-Fi chipset.

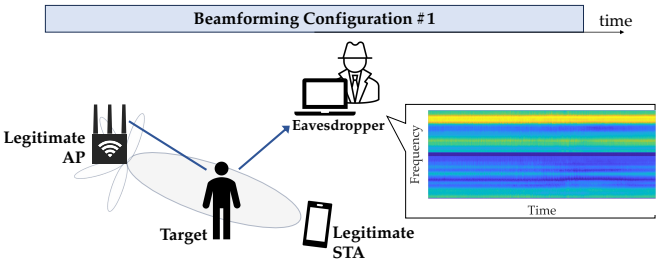

Fig. 4: Attack scenario without `BeamDancer`; the adversary has access to "clean" CSI data and can make sensible guesses.

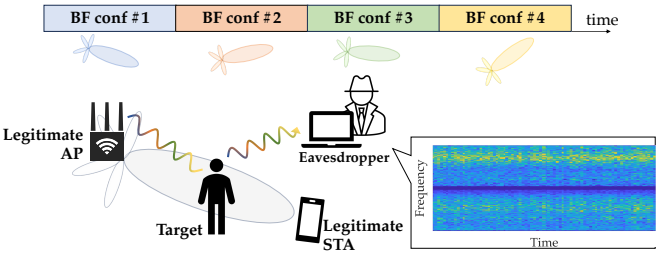

Fig. 5: Attack scenario when `BeamDancer` is deployed; some environment features are still present, but many details in the CSI are lost.

lobe of the radiation pattern generally offers better signal-to-noise ratio (SNR) with respect to omnidirectional transmissions. Beamforming was introduced in Wi-Fi more than ten years ago, with the High Throughput (HT) amendment to the IEEE 802.11n standard. Nowadays, it is a common feature in many commodity Wi-Fi devices, where the transmitter (*beamformer*) applies an appropriate steering matrix to the transmit signal—depending on the wireless channel's physical properties—to facilitate the reception at the receiving node (*beamformee*). There are two methods to compute the steering matrix. The first method, called *implicit beamforming*, represents a naive application of the beamforming technique. The beamformer estimates the wireless channel's properties by opportunistically exploiting the CSI of the frames received from the beamformee. It is based on the assumption that the wireless channel is reciprocal. Still, the steering matrix computed this way is often sub-optimal because of the physical differences in the transmit and receive signal processing chains. The second method, called *explicit beamforming*, can achieve better performance but requires more complex coordination between the beamformer and the beamformee which keep exchanging sounding packets and reports (cf., Section IV-A). Explicit beamforming is the only mechanism standardized in modern Wi-Fi, and hence the one we focus on in this work.

B. Wi-Fi Sensing

Wi-Fi sensing entails the analysis of the received wireless communication signals to obtain information about the surrounding physical environment. In essence, the physical properties of the received communication signals are affected by several factors, including the multipath propagation from the transmitter due to the varying scattering and reflections caused by people and objects in the environment. The overall effect of these factors is expressed by the frequency response H of the time-varying wireless channel [30], which can be defined as a function of the frequency f and time t by

$$H(f;t) = \sum_{i=1}^{N} a_i(t) e^{-j2\pi f \tau_i(t)}, \tag{1}$$

where N is the number of multipath components, while $a_i(t)$ and $\tau_i(t)$ are the complex attenuation factor and the propagation delay for the i-th path, respectively.

The CSI computed at the Wi-Fi receiver estimates the channel's frequency response in Eq. (1) by using pilot signals. Figure 3 illustrates that the CSI is a *fingerprint* of the wireless channel that inherently captures the varying interaction of the communication signals with the surroundings. While some systems use the raw CSI amplitude and phase to perform sensing [31], some others derive different physical quantities by pre-processing the CSI to extract for example the AoA or the Doppler shift [32]. Generally, the former systems use ML techniques (often neural networks) to classify the CSI as corresponding to some target events (activity, location, etc.); the latter systems, instead, rely on traditional signal processing to detect and quantify channel variations. In our experimental evaluation in Section V we consider both types of wireless sensing applications to show the effectiveness of our physical-layer privacy-preserving solution.

III. Threat Model and Proposed Countermeasure

We consider a general indoor environment with one adversary eavesdropping on the Wi-Fi signals and exploiting the CSI from the sniffed Wi-Fi frames to perform illegitimate sensing, e.g., to estimate a person's location or speed. This scenario is schematically depicted in Fig. 4, where we assume a legitimate transmitter is placed in a fixed (either known or unknown) location, like an access point (AP) that provides

Wi-Fi connectivity inside the room. The AP has multiple antennas and supports the explicit beamforming procedure towards connected stations as defined in the IEEE 802.11ac standard. Our assumption looks reasonable because multi-antenna APs are common in residential and office spaces today; their position rarely changes, and they are often even fixed on the walls or the ceiling. In our model, the adversary is mounting a passive privacy attack; therefore, to sense the environment, the adversary has to rely on the existing Wi-Fi traffic only. We assume the traffic volume is sufficiently large for the adversary to obtain the CSI with enough resolution to carry on the desired sensing activity. Finally, we assume that the legitimate station and the eavesdropping device are also in a fixed location. We intentionally consider a high traffic rate and static devices to evaluate the worst-case scenario from the defender's perspective.

The primary privacy goal in this scenario is to prevent the adversary from succeeding in his sensing effort using the eavesdropped CSI. When `BeamDancer` is deployed on the AP, the beamforming configuration of the AP keeps changing, as shown in Fig. 5. Ensuring that slight modifications to the beamforming configuration do not disrupt the communication performance is paramount. We shall verify this later in Section V, when we evaluate the privacy-preserving capabilities of the proposed technique and its impact on communication performance. Furthermore, by continuously changing the beamforming, the channel's effect appears to be ever-changing to the adversarial eavesdropper, nullifying any attempt to fingerprint the channel. Finally, `BeamDancer` must work without prior knowledge of the physical environment and without using additional specialized hardware.

IV. System Implementation and Validation

We design `BeamDancer` to thwart wireless sensing applications by leveraging only native capabilities of commodity Wi-Fi hardware. The core idea of the proposed approach is to continuously modify the beamforming configuration of a MIMO Wi-Fi system, like the legitimate AP in our threat model (Section III). `BeamDancer` builds upon `Beamer`, a tool we developed for this work to create arbitrary beamforming configurations on the beamformer by exploiting the standard IEEE 802.11ac beamforming mechanism.

In this section, after a brief review of the *explicit beamforming* procedure standardized by the Very High Throughput (VHT) amendment in IEEE 802.11ac (Section IV-A), we show how `Beamer` exploits this procedure to provide total control over the beamforming configuration (Section IV-B). We validate our tool by verifying that all the different parameters affect beamforming according to what we expect from the theory (Section IV-C). Finally, we introduce some beamforming manipulation ideas (Section IV-D) that will be employed in the empirical evaluation of `BeamDancer`.

A. VHT Explicit Beamforming Procedure

We briefly review the VHT explicit beamforming procedure using a minimal example with an AP acting as the beamformer and a station (STA) as the beamformee. We denote with N_t and N_r the number of antennas of the AP and the STA, respectively. For each OFDM subcarrier[1], the signal $\mathbf{r} \in \mathbb{C}^{N_r}$ received at the STA depends on the channel matrix $\mathbf{H}$ and the transmitted signal $\mathbf{x} \in \mathbb{C}^{N_t}$ according to $\mathbf{r} = \mathbf{H}\mathbf{x} + \mathbf{n}$, where $\mathbf{n} \in \mathbb{C}^{N_r}$ represents zero-mean additive Gaussian noise.

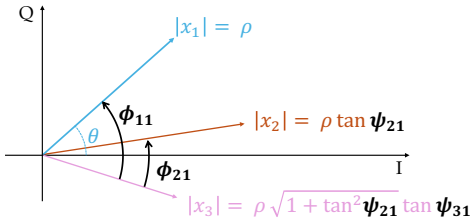

Fig. 6: The angles in Eqs. (2) to (5) express how the AP should transmit the beamformed signals from its $N_t = 3$ antennas.

The ultimate goal of the entire procedure is to determine the best way to split the transmit signal $\mathbf{x}$ across the N_t antennas to maximize the SNR at the STA, i.e., maximize $|\mathbf{H}\mathbf{x}|^2$. The VHT beamforming procedure starts with the AP sending a *Null Data Packet Announcement* (NDP-A), followed by an empty frame (*Null Data Packet*, NDP) which is used for channel sounding. The NDP is a standard pre-coded frame that allows the STA to estimate the channel between each transmit N_t and receive N_r antennas. Since $\mathbf{x}$ is known for the NDP, the STA can estimate the channel matrix $\mathbf{H}$; however, since this matrix can become very large, the information in $\mathbf{H}$ is transmitted back to the AP within a *Compressed Beamforming Report* based on singular value decomposition (SVD) of $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$.

Let us focus on the case treated in this work, which is N_r=1 and N_t=3, implying that the channel matrix $\mathbf{H}$ is a 1×3 row vector. Applying the SVD to $\mathbf{H}$, we have that $\mathbf{U} = 1$ and $\mathbf{\Sigma}$ is a row vector with only one non-zero component, corresponding to the unique non-zero eigenvalue of $\mathbf{H}^*\mathbf{H}$. Then, the first column of $\mathbf{V}$ is the complex vector $\mathbf{v} = [v_1, v_2, v_3]^\intercal$ that maximizes $|\mathbf{H}\mathbf{x}|^2$ if $\mathbf{x} = \mathbf{v}$. The vector $\mathbf{v}$ can be conveniently computed by the STA as $\mathbf{H}^*(\sqrt{\mathbf{H}\mathbf{H}^*})^{-1}$ starting from the NDP. Rather than focusing on the specific magnitude ρ and phase θ of each of the components $v_i = \rho_i \cdot e^{j\theta_i}$ (i=1, 2, 3), the beamforming procedure is designed to express the relation between the three transmitted signals. In particular, the Wi-Fi standard prescribes how to encode two components (e.g., v_2 and v_3) as functions of the third component (e.g., v_1), using Givens rotations defined by the following angles:

$$\phi_{11} = \theta_1 - \theta_3 \tag{2}$$

$$\phi_{21} = \theta_2 - \theta_3 \tag{3}$$

$$\psi_{21} = \arctan \frac{\rho_2}{\rho_1} \tag{4}$$

$$\psi_{31} = \arctan \sqrt{\frac{\rho_3^2}{\rho_1^2 + \rho_2^2}}\,. \tag{5}$$

These four angles establish a precise geometrical relation between the signals transmitted by the AP when beamforming is applied, which is illustrated in Fig. 6 and holds for any $\mathbf{x}$

[1]We recall IEEE 802.11ac uses Orthogonal Frequency-Division Multiplexing (OFDM) transmissions. The number of OFDM subcarriers depends on the signals's bandwidth. For instance, 80-MHz frames have 256 subcarriers.

Fig. 7: Schematic representation of `Beamer`. Custom reports crafted on the STA control the AP's steering matrix.

transmitted over the same channel $\mathbf{H}$. For example, taking x_1 as reference, the generic beamformed signal $\mathbf{x}$ sent by the AP over one single OFDM subcarrier can be described by

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \rho e^{j\theta} \cdot \begin{bmatrix} 1 \\ \tan\psi_{21} e^{j(\phi_{21}-\phi_{11})} \\ \sqrt{1+\tan^2\psi_{21}}\tan\psi_{31} e^{-j\phi_{11}} \end{bmatrix}. \quad (6)$$

Finally, the four angles computed for each OFDM subcarrier are quantized and packed into the *Compressed Beamforming Report* frame sent back to the AP. Each ϕ angle has support $[0, 2\pi)$ and is quantized using 6 bits, while each ψ angle has support $[0, \frac{\pi}{2})$ and is quantized using 4 bits.

B. Arbitrary Report Injection

To modify the beamforming configuration of a COTS Wi-Fi device, one would typically have to tamper directly with the Wi-Fi chipset of the AP that is transmitting beamformed frames to a connected STA. Unfortunately, we did not find an addressable *beamforming configuration memory* (specific to each STA) in any of the chipsets we analyzed. Thus, we chose a different approach instead and proceeded to modify the compressed report on the STA to *inject* arbitrary beamforming configurations on the AP. We notice that this implementation choice makes the tool very flexible because, in this way, it is possible to arbitrarily customize the steering matrix of any AP, without bothering about the dissimilarities between different APs. After analyzing several Wi-Fi chipsets and the corresponding toolchains, we targeted Broadcom chipsets, since it is possible to modify the low-level behavior of some of the most common chipsets using the Nexmon[2] framework. We implemented the tool on an Asus RT-AC86U router; however, even if we only focus on this specific implementation in the following, the tool also works on Google Nexus 5 smartphones and Raspberry Pi 4B and 5 computers.

In Fig. 7 we summarize the proposed design: the device on the left acts as a STA connected to the AP on the right. By modifying the beamforming report on the STA, we can inject custom reports and configure the steering matrix of the AP with any arbitrary content. Broadcom Wi-Fi chipsets, including the `BCM4365E` embedded in the Asus routers, are equipped with a D11 microcontroller managing all time-critical operations [33], including the generation of the compressed beamforming report exactly $10\,\mu$s after the reception of the NDP. The report's content is fetched from an internal memory whose content is directly updated by the hardware upon reception of the NDP. Since it is impossible to alter this memory content, we modified the D11 microcode

[2]https://nexmon.org, Nexmon: The C-based Firmware Patching Framework

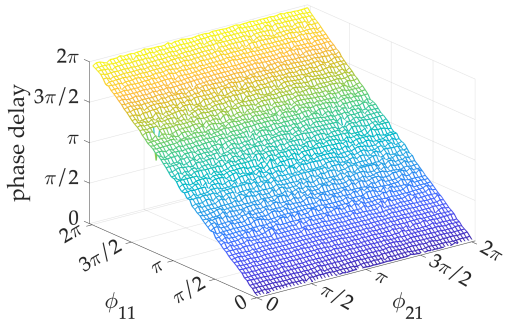

Fig. 8: The phase difference measured between the signals v_1 and v_2 as a function of ϕ_{11} and ϕ_{21} agrees with Eq. (2).

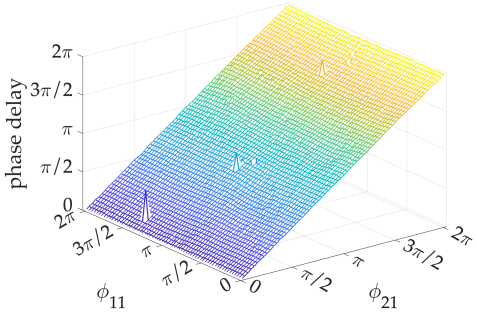

Fig. 9: The phase difference measured between the signals v_2 and v_3 as a function of ϕ_{11} and ϕ_{21} agrees with Eq. (3).

using the Nexmon binary patching tool to load the report's content from another memory region, named Buffer Memory or Template RAM (i.e., the memory also containing the Wi-Fi beacons) [34]. The `Beamer` tool can store arbitrary compressed reports into the Template RAM and read the actual beamforming memory content to add arbitrary modifications dynamically on top of it.

C. System Analysis and Validation

Before deploying our system in a real testbed, we analyze the results obtained by modifying the beamforming report to ensure they match with what we would expect from the theory. In particular, we verify that the arbitrary modification of the four angles ϕ_{11}, ϕ_{21}, ψ_{21} and ϕ_{31}, which are expressed in the beamforming report for each OFDM subcarrier, does change indeed the signal transmitted by the beamformer according to Eqs. (2) to (5). To this end, we connect with coaxial cables the three external antennas of the beamformer *under test* to another device, namely an Asus RT-AX86U, that is used to collect the CSI and measure the signals transmitted by the beamformer.

It is easy to understand and visualize the effect of the first two angles ϕ_{11} and ϕ_{21}. From Eqs. (2) and (3), we see that these two angles describe the phase difference between the signals transmitted by the beamformer's antennas. First, we verify with an exhaustive search over all the possible ϕ_{11} and ϕ_{21} that the amplitude does not change. Then, we measure the phase delay between the different antennas. Figure 8 shows that the phase difference (averaged over all the subcarriers) between the antennas corresponding to signals v_1 and v_2

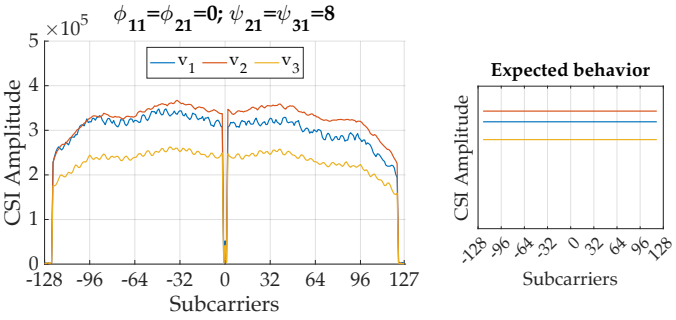

Fig. 10: CSI amplitude measured at the three antennas with the angles ψ_{21} and ψ_{31} constant for all the subcarriers. Amplitude values are dimensionless quantities extracted directly from the Wi-Fi chipset.

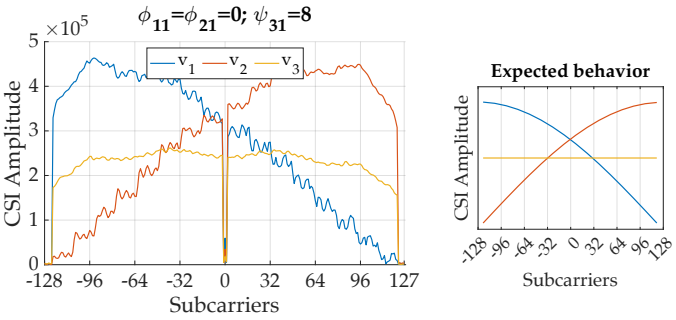

Fig. 11: CSI amplitude measured at the three antennas with the angle ψ_{21} increasing towards the higher subcarriers (the same value of ψ_{21} is assigned to groups of 16 contiguous subcarriers). Amplitude values are dimensionless quantities extracted directly from the Wi-Fi chipset.

depends only on the angle ϕ_{11}, ranging from 0 to 2π. This agrees with the behavior described by Eq. (2). Analogously, in Fig. 9 we observe that the phase difference between the signals v_2 and v_3 depends only on the angle ϕ_{21} as described in Eq. (3). We then verify that also the signals' amplitudes satisfy Eq. (4) and Eq. (5). To this end, we report in Fig. 10 (left) the CSI captured from the three transmitting antennas when $\phi_{11} = \phi_{21} = 0$ and $\psi_{21} = \psi_{31} = \pi/2$, which correspond to the quantized value 8. We measure the amplitudes of v_1, v_2, and v_3 at the center of the spectrum, and use these values as a reference, represented in Fig. 10 (right). Finally, we apply an arbitrary beamforming configuration where we assign in succession all the possible ψ_{21} discretized values (from 0 to 15) to 16 groups of 16 contiguous subcarriers, in steps of $\pi/16$. The measured CSI shown on the left side of Fig. 11 agrees with the theoretical values obtained in this situation, represented on the right side. This experiment confirms the theoretical model and demonstrates that the beamforming can be arbitrarily configured on each subcarrier.

D. `BeamDancer` Algorithm

One crucial feature of the `Beamer` tool is that it decouples the injection of a custom beamforming report from the creation of the beamforming report itself. So far, we have reviewed how `Beamer` takes advantage of the standard VHT explicit beamforming procedure to inject on the AP a custom beamforming configuration crafted on the STA, and we have empirically verified that manipulating the angles yields the results expected from theory. Now, the core idea of `BeamDancer` is to leverage the injection capabilities offered by `Beamer` to keep changing the AP's beamforming vector. The ultimate goal is to hinder the adversarial eavesdropper's sensing capabilities by creating an ever-changing channel response that cannot be fingerprinted. A naive solution to tamper with the beamforming configuration would be to assign arbitrary values to the four angles in Eqs. (2) to (5) for each one of the OFDM subcarriers every time the STA crafts a new beamforming report. For instance, their values could be drawn uniformly over their dynamic range, which is 6 bits for each ϕ angle and 4 bits for each ψ angle. However, this solution could be detrimental to the communication system's performance, as generating completely random beamforming configurations might also hinder legitimate communications. In general, it is safe to assume that a more principled approach is required, and we would like to have fine-grained control over the possible beamforming configurations.

With `BeamDancer`, we explore for the first time the effects of controlled manipulation of beamforming on wireless sensing using COTS devices. For every OFDM subcarrier, the `BeamDancer` algorithm adds a random value r to the true angles ϕ_{11}, ϕ_{21}, ψ_{21} and ψ_{31} in the beamforming report to obtain randomized beamforming angles $\widetilde{\phi}_{11}$, $\widetilde{\phi}_{21}$, $\widetilde{\psi}_{21}$ and $\widetilde{\psi}_{31}$. It is important to notice that `BeamDancer` works with the quantized and discretized values of the beamforming angles; hence, in the following, we will treat the angles as integer values that can range in the interval $[-2^{k-1}, 2^{k-1}-1]$, depending on the number of quantization bits k. In our implementation, the r values are drawn from a discrete uniform distribution described by a single parameter $\mathcal{R}$ that defines the support of the distribution as a fraction of the angles' dynamic range. For example, if $\mathcal{R} = 50$, the r values are taken from a uniform distribution with support equal to the 50% of the angles' range, i.e., $[-16, 15]$ for the angles ϕ_{11} and ϕ_{21} and $[-4, 3]$ for the angles ψ_{21} and ψ_{31}. Now, let us denote α as a generic angle (either ϕ or ψ) and $\widetilde{\alpha}$ the corresponding randomized angle. Then, $\widetilde{\alpha}$ can be simply assigned the value $\alpha + r$ if it falls within the dynamic range of the angle. Some extra care is needed, however, when the resulting $\widetilde{\alpha}$ falls outside of the range. In particular, in case of an overflow (underflow), we wrap the increment r around the maximum (minimum) value instead of just performing a modulo operation to avoid generating randomized angles that are too far from the true angle according to the desired amount of randomization. The resulting operation for a generic angle α is the following equation:

$$\widetilde{\alpha} = \begin{cases} \alpha + r & \text{if } -2^{k-1} \leq \alpha + r \leq 2^{k-1} - 1 \\ 2^{k-1} - 1 - (\alpha + r) & \text{if } \alpha + r > 2^{k-1} - 1 \\ -2^{k-1} - (\alpha + r) & \text{if } \alpha + r < -2^{k-1} \end{cases} \tag{7}$$

and the corresponding pseudocode for the different angles is listed in Algorithm 1.

To summarize, `BeamDancer` algorithm randomizes each

beamforming angle around its actual value according to a parameter that specifies the *randomization intensity* as a percentage of the angle's dynamic range. For the experimental analysis in Section V, the control over the *randomization intensity* will be fundamental to assess how `BeamDancer` impacts on Wi-Fi sensing and its effect on legitimate communications. Thanks to the versatility of the `Beamer` tool, more sophisticated randomization techniques might be considered in different application scenarios. However, this work intends to illustrate the practical design and evaluation method of a novel technique to preserve privacy at the physical layer via arbitrary beamforming manipulation. We leave the analysis of more advanced beamforming manipulation algorithms as future work.

Algorithm 1 `BeamDancer`

Input: Angles $(\phi_{11}, \phi_{21}, \psi_{21}, \psi_{31})$ and $\mathcal{R} \in [0, 100]$
Output: Randomized angles $(\widetilde{\phi}_{11}, \widetilde{\phi}_{21}, \widetilde{\psi}_{21}, \widetilde{\psi}_{31})$

for each ϕ **in** $[\phi_{11}, \phi_{21}]$:
  $r \leftarrow \text{Uniform}[-\frac{32 \cdot \mathcal{R}}{100}, \frac{32 \cdot \mathcal{R}}{100} - 1]$
  if $(\phi + r) > 31$
    $\widetilde{\phi} = 63 - (\phi + r)$
  else if $(\phi + r) < -32$
    $\widetilde{\phi} = -64 - (\phi + r)$
  else
    $\widetilde{\phi} \leftarrow (\phi + r)$
  end if
end for
for each ψ **in** $[\psi_{21}, \psi_{31}]$:
  $r \leftarrow \text{Uniform}[-\frac{8 \cdot \mathcal{R}}{100}, \frac{8 \cdot \mathcal{R}}{100} - 1]$
  if $(\psi + r) > 7$
    $\widetilde{\psi} = 15 - (\psi + r)$
  else if $(\psi + r) < -8$
    $\widetilde{\psi} = -16 - (\psi + r)$
  else
    $\widetilde{\psi} \leftarrow (\psi + r)$
  end if
end for
return $(\widetilde{\phi}_{11}, \widetilde{\phi}_{21}, \widetilde{\psi}_{21}, \widetilde{\psi}_{31})$

E. Theoretical and empirical validation of `BeamDancer`

Our adversarial model (Section III) considers an attacker that senses the environment utilizing eavesdropped CSI data. `BeamDancer` builds upon the assumption that spoiling the information gathered by the adversary can hinder the capability to mount any sensing attack. Specifically, `BeamDancer` works by modifying the statistical properties of the information utilized by the adversary, decreasing the correlation between specific configurations of the channel and the CSI measured by the adversary. Modeling such modifications in non-trivial environments is extremely complex; in this work, their effect on the adversarial sensing accuracy will be evaluated experimentally in Section V. Still, we can use numerical analysis to estimate the theoretical impact of `BeamDancer` on legitimate communications. Let us consider an arbitrary choice of the beamforming vector, corresponding to specific angles in Eqs. (2) to (5) that maximize the signal amplitude along a certain direction. These angles are the ones the standard beamforming procedure would select if the STA is positioned along that direction. We can then assess how the amplitude of the received signal changes when `BeamDancer` is applied, i.e., when we select different angles in the vicinity of the correct ones. By numerically testing all the possible randomizations and averaging the results, we can calculate how much `BeamDancer` lowers on average the Received Signal Strength (RSS) with respect to the optimal configuration selected through the vanilla beamforming procedure. The results of such numerical analysis are reported in Fig. 12 for different randomization intensities, identified by the parameter $\mathcal{R}$ (cf. Algorithm 1), where we take the optimal beamforming case as reference (no randomization, $\mathcal{R} = 0$). It is mundane to observe that the randomization directly affects the RSS, lowering the link budget as the randomization intensity increases. In particular, we verify that selecting a completely random configuration of the four beamforming angles ($\mathcal{R} = 100$) lowers the RSS by almost 4 dB on average.

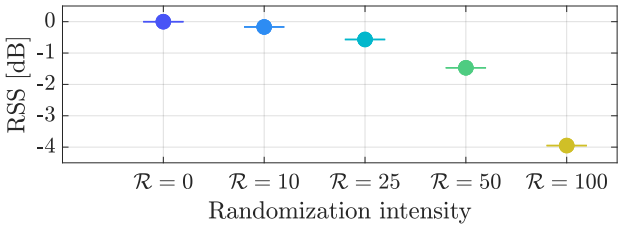

Fig. 12: Numerical results showing the variation in the RSS as a function of the $\mathcal{R}$ parameter of `BeamDancer`. Complete randomization of the angles lowers the RSS by about 4 dB.

V. Experimental Evaluation

We evaluate the effectiveness of `BeamDancer` against Wi-Fi sensing by running multiple experiments in indoor environments using commodity devices. First, we provide a general description of the experimental setup (Section V-A); then, we analyze the performance of `BeamDancer` against different sensing applications:[3] a device-free localization system using deep learning (DL), either supervised (Section V-B) and unsupervised (Section V-C), and a model-based human speed estimation system (Section V-D), based on traditional signal processing techniques. Finally, we show that `BeamDancer` has minimal impact on legitimate Wi-Fi throughput (Section V-E).

A. Setup

Figure 13 depicts the room layout where the first experiments are conducted. The room is approximately 46 m^2 with large windows on the north-south walls and rich multipath characteristics induced by tables, cabinets, and other metallic frames along its perimeter. The node with the label `Tx` represents the main AP in the room, i.e., the beamformer

[3]We pledge to make the tool and the data publicly available upon notification of acceptance at https://github.com/seemoo-lab/beamdancer.

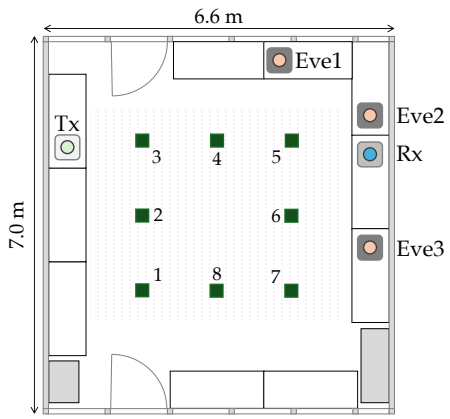

Fig. 13: Layout of the experimental testbed and devices' location. `Tx` is the beamformer, `Rx` is the beamformee, and the `Eve` nodes are the four eavesdroppers.

according to our threat model (Section III), which is the only node required to have multiple antennas in this scenario. In our experiments, `Tx` is an Asus RT-AC86U router, which supports the transmission of beamformed 802.11ac frames with a bandwidth of up to 80 MHz. The STA that takes the role of the beamformee is a Nexus 5 smartphone, identified with the label `Rx` in Fig. 13. In a real scenario, `Rx` is the legitimate receiver and the beamforming vector is usually optimized for this station. As discussed in the previous sections, we use `BeamDancer` to modify the beamforming from `Tx` by crafting specific beamforming reports on `Rx`.

We then deploy four adversarial eavesdroppers, labeled `Eve1` to `Eve4`, in different locations as shown in Fig. 13. The `Eve` nodes are four Asus RT-AX86U routers that extract the CSI of beamformed Wi-Fi frames sent by `Tx` towards the beamformer (co-located with `Eve2` in the layout) using the AX-CSI extractor [35]. According to the considered threat model, the four `Eve` are not supposed to work together. Yet, deploying multiple `Eve` nodes at different locations helps provide some statistics about the expected sensing performance in the environment. Indeed, the CSI measured by the 4 `Eves` changes significantly because the receivers are spaced by many wavelengths λ and experience completely different multipaths.

In each experiment, we consider different `BeamDancer` configurations. The scenario Clean represents the typical scenario considered in our threat model, with one AP doing beamforming according to the standard and no countermeasures in place. In the scenario Fixed, we study what happens when the beamforming is fixed to an arbitrary value and does not change over time. This can be achieved by letting the `Beamer` tool always inject the same constant report. Then, we consider different possibilities for the `BeamDancer` scenarios, all identified by the label Rand`R`. The number `R` represents the amount of randomization (as a percentage of the angles' dynamic range) applied to the actual values of the angles in the beamforming report. For instance, Rand25 means that `BeamDancer` is randomizing the beamforming angles *of every subcarrier* according to a uniform distribution that spans over 25% of the dynamic range of the corresponding angle.

TABLE I: MLP Architecture for Device-free Localization.

Layers	Size	Activation
Input	256x2	-
Flatten	-	-
3x Dense	64	ReLU
Output	8	Softmax

We recall that ϕ angles can take 64 values (6 bits) while ψ angles can take 16 values (4 bits); hence, Rand25 implies a randomization of ± 8 on ϕ angles and of ± 2 on ψ angles. The same idea applies to Rand10, Rand50, and Rand100; of course, Rand100 is equivalent to complete randomization.

B. Device-free Localization

In the first experiment, we assess the impact of different `BeamDancer` configurations on a device-free localization system. To evaluate the privacy-preserving performance of `BeamDancer`, we implement a localization system using a multi-layer perceptron (MLP) trained to map a single CSI into the corresponding user's location. The model proposed here is simple but has two advantages: first, it allows us to focus on the effects of `BeamDancer`; second, it can be used to define a baseline for future sensing (and anti-sensing) techniques.

Localization System Architecture. The MLP takes as input one CSI vector at a time—one complex vector of 256 elements—and estimates where the person is among the eight target locations shown in Fig. 13. The architecture of the MLP is shown in Table I. For every considered scenario, the MLP is first re-trained from scratch for 20 epochs using a dataset of 6000 CSI collected with the person standing in each target position. Then, the MLP is tested against another dataset collected shortly after the training set under the same experimental conditions.

Experimental Results. We report in Fig. 14 the results of the localization experiment, showing that beamforming manipulation significantly affects wireless sensing. It is clear that increasing the randomization of the beamforming angles has the intended effect of impairing device-free localization to the point that, when the beams are randomly directed over time, the proposed system can only estimate the position with random guesses. It is interesting to notice that the location of the eavesdropper in the room affects the localization performance only marginally, with `Eve4` performing slightly better than the other receivers but following the same general trend. As expected, the Clean configuration (no modifications to beamforming) has the highest localization accuracy. Moreover, we notice that small randomization of the beamforming angles (Rand10) does not affect the localization performance. The results for the Fixed scenario, for which we freeze the beamforming to one specific configuration, are worth a final comment. In this case, since the presence of the person in the wireless channel does not alter the beamforming, variations in the channel are caused only by the human body reflections and scattering, which have a noticeable but weaker effect on the CSI.

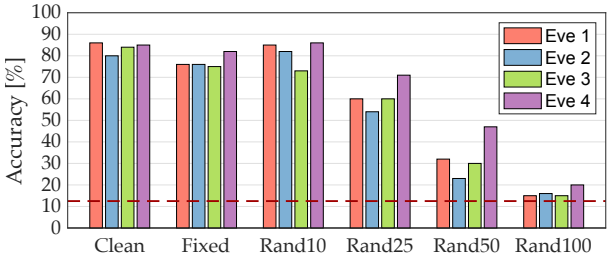

Fig. 14: Average localization accuracy for different `BeamDancer` configurations at the four eavesdroppers over the eight positions reported in Fig. 13. The dashed line indicates the expected accuracy of a random guess.

TABLE II: VAE Architecture for Device-free Localization.

Layers	Nodes/(Kernel Size)	Activation
3x Conv2D	$(3,5) \times 16$, stride (2,4)	ReLU
Flatten	-	-
3x Dense	32	ReLU
Latent Space	*2*	
Dense	256	ReLU
Reshape	(4,4,16)	-
3x Conv2D$^\intercal$	$(3,5) \times 16$, stride (2,4)	ReLU

C. Unsupervised Device-free Localization

To further analyze the impact of `BeamDancer` on data-driven sensing techniques and on device-free localization in particular, we consider another problem that is closely related to the previous one. This time, instead of providing the ML system with a ground truth from which to learn the correct mapping between the CSI and the corresponding user's position, we examine the possibility of discerning among the different target locations with an unsupervised approach, i.e., solely based on the properties of the CSI itself.

To this end, we implement a variational autoencoder (VAE) that (i) encodes short sequences of CSI into a latent space of uncorrelated bivariate Gaussian distributions with means (μ_1, μ_2) and standard deviations (σ_1, σ_2); and (ii) in the decoder part, tries to reconstruct the original input CSI. Unlike traditional autoencoders, VAEs build upon the causal assumption that the latent distribution generates the observations. Once trained, the VAE can be used both for generating new samples that are CSI-alike or to compress a given CSI into a more compact representation.

Localization System Architecture. The encoder takes a short sequence of CSI as input, maps it into a latent space of bivariate Gaussian distributions, and then reconstructs the input data trying to minimize the reconstruction error and the KL divergence of the approximate to the true posterior. The architecture of the VAE is reported in Table II. For every considered scenario, the VAE is trained for 20 epochs using a dataset of 6000 CSI collected with the person standing in each target position. We then proceed to visually inspect the mapping of different CSI data into the VAE's latent space.

Experimental Results. We recall that, in this work, we use the VAE as an *unsupervised* ML technique that separates CSI data into different clusters without any information about the ground truth. Figure 15 shows a visualization of the latent space of the VAE trained using the CSI data collected by `Eve1` with different `BeamDancer` configurations. Every input CSI is mapped into a point whose x-y coordinates represent the expected values (μ_1, μ_2) of the corresponding bivariate Gaussian distribution. Here, we omit the representation of the standard deviations. Experimental results show that when `BeamDancer` is off (Clean case), we can distinguish eight clusters corresponding to the eight target positions in Fig. 13. Even though some clusters overlap, they are still fairly separable, and a simple classification algorithm (e.g., kNN) would yield high classification accuracy on the target position. However, as the obfuscation increases, different clusters overlap more until they eventually collapse into a single point (Rand100). This result is a visual proof of why data-driven sensing methods fail when tampering with the beamforming configuration. Moreover, the graphs in Fig. 15 confirm what we already saw in Fig. 14: while for Rand10 it is still possible to discern different activities, localization in the Rand100 case would be practically impossible with this method.

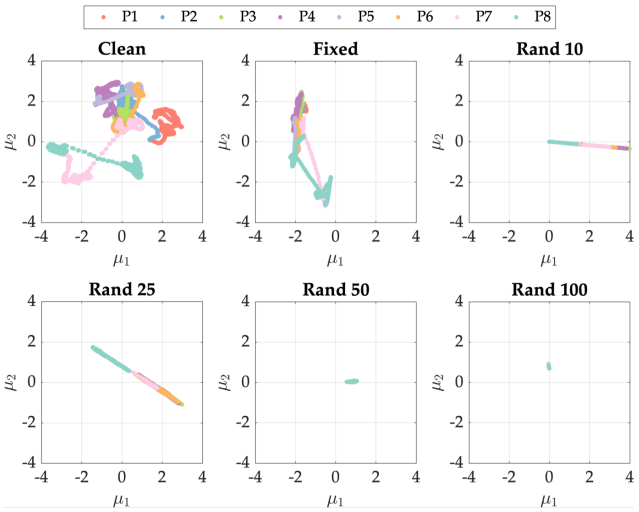

Fig. 15: Expected values of the latent variables into which every CSI from `Eve1` is mapped. Different locations use different colors. As the beamforming randomization increases, separating different CSI clusters becomes harder.

D. Device-free Speed Estimation

We now turn our attention to a third type of wireless sensing application that instead builds upon a physical model of the interaction of Wi-Fi signals with the surrounding environment to estimate the speed of a moving person. The extraction and the analysis of the micro-Doppler effect from the CSI to estimate the speed of a target object is today a common technique in wireless sensing and a fundamental building block for many activity recognition frameworks.

In this subsection, we briefly elaborate on the algorithm we used to extract the micro-Doppler effect from the CSI [36], and then we present the results obtained in our experiment.

Speed Estimation. We use the classic MUSIC algorithm to estimate frequency shifts from the incoming signals that reach

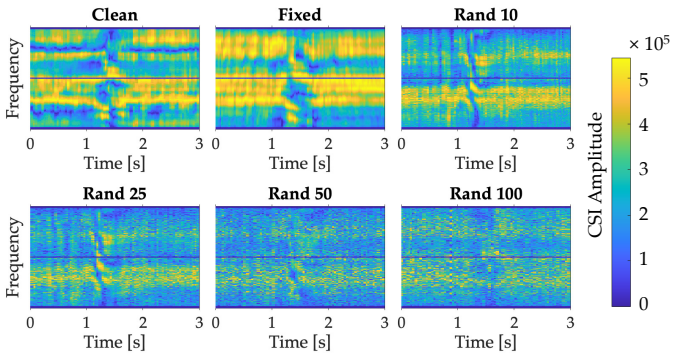

Fig. 16: Effect on the CSI of a person moving across the LoS between `Tx` and `Eve2` in the speed estimation experiment. The movement "fades away" as the randomization increases. CSI amplitude values are dimensionless quantities extracted directly from the Wi-Fi chipset.

the antennas of an `Eve` device. The underlying assumption to this CSI-based speed estimation approach is that when a person walks, the paths of the signals reflected from his body change, and the rate of change of every path directly defines the Doppler frequency shift that can be used for tracking the moving person. The delay introduced by this movement appears as a small frequency shift in the CSI that is *captured* within the coefficient $a_n(t)$ in Eq. (1).

Successive CSI samples over time are arranged in a matrix of shape $M \times K$, where M and K are the numbers of samples and orthogonal frequency-division multiplexing (OFDM) subcarriers, respectively. Then, the CSI matrix can also be computed as follows:

$$\text{CSI}_{M\times K} = \boldsymbol{A}_{M\times N}\boldsymbol{X}_{N\times K} + \boldsymbol{W}_{M\times K}\,, \tag{8}$$

where $\boldsymbol{X}$ is the OFDM signal received on K subcarriers from N different paths, $\boldsymbol{W}$ is a matrix of Gaussian noise, and each column of $\boldsymbol{A}$ contains the Doppler velocity of each path. The MUSIC pseudo-spectrum is then constructed as $P_{MUSIC}(v) = [\boldsymbol{a}^*(v)\boldsymbol{U}_n\boldsymbol{U}_n^*\boldsymbol{a}(v)]^{-1}$, where the columns of $\boldsymbol{U}_n$ are the eigenvectors corresponding to near-zero eigenvalues of the received signal covariance matrix, and $\boldsymbol{a}(v)$ is the steering vector. This spectrum has sharp peaks at the Doppler velocity of the incoming signals. We apply this method for some sets of CSI samples collected in a scenario where a person walks perpendicularly to the link between the transmitter and the receiver.

Experimental Results. The proposed method is evaluated with a person walking perpendicularly to the LoS between `Tx` and `Eve2` with nearly constant velocity. In other words, according to the layout in Fig. 13, the target person is moving along the direction north-south. The impact of the movement on the CSI is shown in Fig. 16 under the different randomization schemes. Two scenarios were considered with two different speeds: slow (0.5 m/s) and fast (2 m/s). A preliminary analysis revealed that human speed is almost constant over short intervals of 50 ms. Therefore, we consider a coherent processing interval (CPI) of 50 ms, and for every CPI, we estimate the Doppler speed using the MUSIC algorithm.

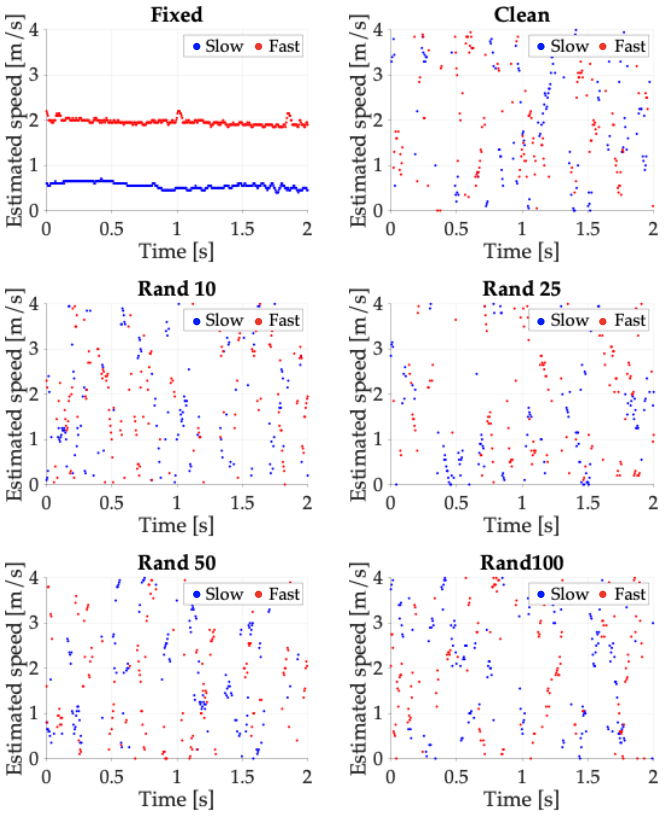

Fig. 17: Even subtle variations in the beamforming configuration can greatly affect the speed estimation results. Indeed, speed estimation is not accurate even in the Clean case, where beamforming changes according to standard procedures.

The speed estimation results shown in Fig. 17 reveal that the estimation is fairly accurate only in the Fixed case. When the beamforming configuration changes over time, then it is impossible to produce accurate estimates. This is true when `BeamDancer` is active, but also when we leave the beamforming on (Clean case). In this sense, `Beamer` is used to force devices to *fix* the beamforming of COTS devices for specific research purposes. Indeed, the original work from which we drew the speed estimation algorithm did not consider transmissions with beamforming [36]. The impact of beamforming can be evaluated by writing the CSI samples collected over each CPI as:

$$\sum_{m=1}^{M} \alpha_m e^{-j2\pi f \frac{v\Delta t_m}{c}} e^{-j2\pi\psi_m} = \sum_{m=1}^{M} \alpha_m e^{-j2\pi\Psi_m} \tag{9}$$

where M is the number of packets in each CPI, and ψ_m is a random phase shift introduced for the m^{th} packet. What Eq. (9) tells us is that, by introducing the random term ψ_m, the phase shift over subsequent CSI samples changes, affecting the Doppler estimates. This is confirmed by the experimental results in Fig. 17, which show that accurate Doppler velocity estimation is not possible when beams are changing over time. Averaging over time the data points reported in Fig. 17 still does not help the adversary, as the average measured speeds are off in all the cases. Moreover, the large variance would make any estimate practically useless.

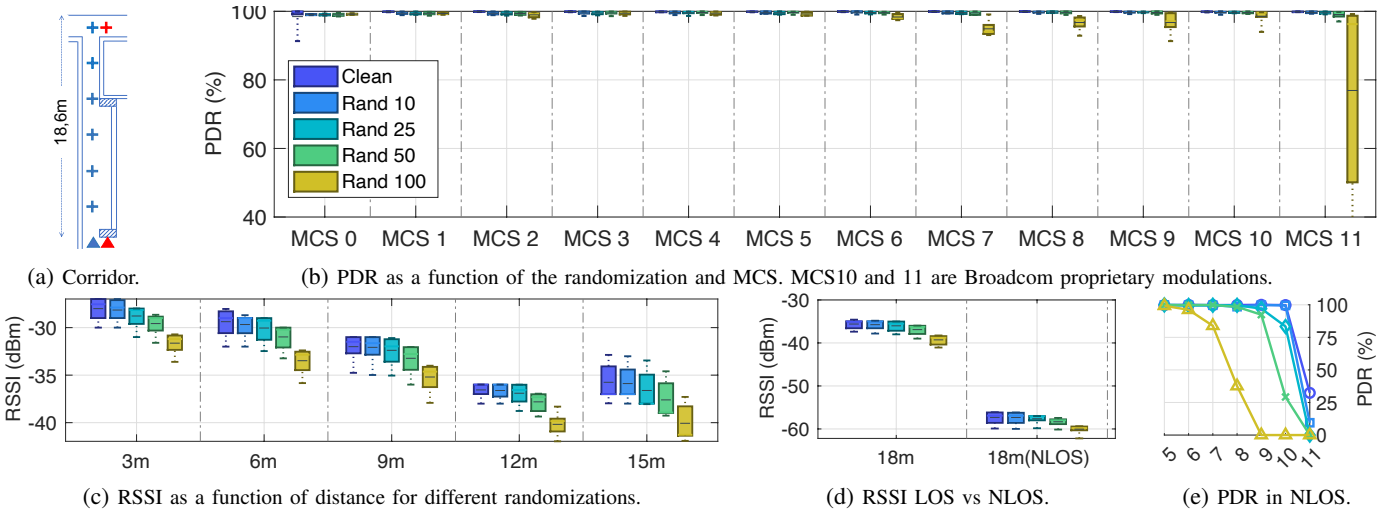

(a) Corridor. (b) PDR as a function of the randomization and MCS. MCS10 and 11 are Broadcom proprietary modulations.

(c) RSSI as a function of distance for different randomizations. (d) RSSI LOS vs NLOS. (e) PDR in NLOS.

Fig. 18: Effects of `BeamDancer` on the downlink performance in the second testbed, designed for the throughput evaluation. (a) Corridor testbed with target points. (b) In LOS, the PDR is affected only at higher MCS when applying Rand100. (c) The RSSI slightly decreases with increasing randomization and distance. (d) The RSSI deteriorates abruptly when switching from LOS to NLOS at similar distances. (e) In NLOS, the PDR drops even with light obfuscation.

E. Impact on throughput

An essential aspect of any privacy-enhancing mechanism is that it should not hinder legitimate communications. We performed a preliminary analysis in the testbed in Fig. 13 and found that the impact of `BeamDancer` on the throughput was negligible. Hence, in this section, we focus on another scenario where we measure and analyze in depth the performance of `BeamDancer` under different and varying conditions.

Performance in Line-of-Sight (LoS). The experiments are performed in the corridor shown in Fig. 18a by placing the AP (blue triangle) and the station (blue crosses) at different distances in steps of 3 m. In Fig. 18b we report the effects of `BeamDancer` on the communication performance by measuring the packet delivery ratio (PDR) when the AP and the connected station are in LoS. Specifically, we disable the automatic AP retransmissions to count the packets received correctly, and we transmit downlink traffic considering all the obfuscation schemes proposed and all the Wi-Fi MCS (including Broadcom-only 1024-QAM MCS10 and MCS11). In this case, the PDR is a more indicative performance indicator than the bit error rate (BER) as some bit errors are implicitly corrected by the coding schemes with no tangible effects for the users. Indeed, the communication performance is only affected when Wi-Fi frames cannot be decoded at the receiver (that is, the PDR drops) because of too many errors.

Figure 18b reports the effect on PDR for separate MCS and obfuscations by averaging over all the six measured distances, from 3m up to 18m. We observe that the effects of all the randomization schemes are almost negligible as the PDR exceeds 99.5% on average, except for Rand100, for which the average PDR is still 96.5%. An interesting observation regarding Fig. 18a is that the "robustness" of the transmissions under the Rand100 scheme depends on the specific modulation and coding rate of the different MCSs. For instance, transmissions with MCS8 (256-QAM, coding rate 3/4) exhibit higher PDR than MCS7 (64-QAM, coding rate 5/6) because the coding rate is more robust. The PDR suddenly drops when Rand100 is applied to MCS11, which is arguably the most complex modulation scheme (1024-QAM, coding rate 5/6). In this case, the boxplot has a very large span indicating that the results strongly depend on the distance. In fact, the boxplot shows the average results for all the measurement points in Fig. 18a, and while the PDR remains high for short distances, it drops to about 60% or even less for larger distances.

Link budget. To show the impact of `BeamDancer` on the link budget, we plot the Received Signal Strength Indicator (RSSI) measured at the receiver as a function of the distance in Fig. 18c. As expected, the RSSI decreases as the distance increases. However, we also point out that the RSSI slightly decreases (about 4 dB for Rand100) when the randomization is stronger. This is because when using `BeamDancer` the steering vector is configured "sub-optimally," and does not steer all the power toward the intended target. The experimental measurements confirm the theoretical results for the RSS anticipated by the numerical analysis in Section IV-E. While it is evident that beamforming randomization can marginally affect legitimate communication processes depending on the randomization intensity, our experiments suggest that there could be a trade-off between the amount of randomization applied and the impact on legitimate transmissions. Nevertheless, we highlight that Rand100 can completely thwart adversarial sensing at the cost of lowering the RSSI by about 4 dB.

Our results also reveal another important aspect of `BeamDancer`; namely, its complete transparency to the normal functioning of the Wi-Fi devices. Comparing the RSSI values in Fig. 18c, we notice that the signal strength remains consistent for distances larger than 12 meters in LoS (about -35 dBm to -40 dBm). This behavior is caused by the power control algorithm of the firmware that automatically regu-

TABLE III: PDR measured in each environment. Results for corridor (LoS) are averaged over all the distances in Fig. 18a.

Environment	Clean [%]	Rand [%] 10	25	50	100
Lab	99.4	99.2	99.2	99.2	97.1
Classroom	99.9	99.4	99.4	99.4	96.7
Corridor (LoS)	99.8	99.6	99.6	99.5	96.5
Corridor (NLoS)	94.2	92.1	89.8	84.8	67.9

lates the TX/RX gain of the radio. During our experiments, we did not disable any firmware operation to validate that `BeamDancer` is fully compliant with the normal functioning of the real device.

LoS vs NLoS. In Fig. 18d, we compare the RSSI in LoS and in NLoS (averaged for all the MCSs) when the station and the AP are 18 m apart. The NLoS case is measured in the same corridor by moving both the AP and the station behind two large concrete pillars (as indicated by the red markers in Fig. 18a). We observe a drop of about 20 dB in the RSSI due to the loss of the LoS, which is expected. However, we observe in Figs. 18d and 18e that `BeamDancer`'s impact on the link budget (both RSSI and PDR) is again minimal. Indeed, the PDR exceeds 98.5% for $\mathrm{MCS} \leq 8$ for all the obfuscation schemes except Rand100.

Different environments. Table III reports the PDR in four different scenarios for different obfuscation schemes, averaged over all the MCSs. These results include (i) the lab where we performed the sensing experiments, (ii) the corridor, and (iii) another classroom whose detailed results are omitted due to space constraints. Table III confirms our previous observations, but it also shows that the system performance remains consistent across different indoor locations. Specifically, in the first three environments, the average PDR is always above 96% even with the maximum randomization. In the NLoS case, things get only a little worse up to Rand50 and then deteriorate more for Rand100, where only $2/3$ of the packets are correctly decoded. Thus, in some generic scenarios without LoS, a tradeoff between performance and security might have to be further considered.

VI. Discussion

In this section, we briefly review the main advantages and some limitations of `BeamDancer`, and discuss some considerations about the future of wireless sensing.

`BeamDancer` does not require specialized hardware. The main goal of this work is to propose a practical solution against adversarial Wi-Fi sensing that does not require custom hardware. The performance of `BeamDancer` is evaluated in terms of sensing accuracy from the perspective of an adversary. It would be interesting to compare its performance against other obfuscators that use secondary devices, even though this is very difficult as they usually require highly specialized equipment. Nonetheless, we conjecture that `BeamDancer` is more scalable and practical than other obfuscators as, in general, their location needs to be optimized to the environment and their operations require coordination with other transmitters.

`BeamDancer` does not require protocol modification. The proposed system does not require any modification to the current beamforming operation. Indeed, `BeamDancer` is completely transparent to the beamforming procedure and can be implemented on any Wi-Fi device without requiring additional coordination or alternative negotiation phases. Implementing `BeamDancer` on a specific device demands tampering with low-level details that are likely to differ among Wi-Fi chipsets and firmware versions. These modifications should be trivial for the device manufacturers, and since they require only very simple modifications (just a few additions and multiplications), we suspect that this method can be easily ported to different devices without impacting performance. Moreover, since `BeamDancer` can be implemented independently on every device, there is no need for any interaction between manufacturers. This is a key point that can facilitate the widespread adoption of similar privacy-preserving techniques against adversarial Wi-Fi sensing.

Robustness against adversarial analysis of the randomization scheme. An adversary could become aware that the beamformed signals are randomized by an algorithm like `BeamDancer` and try to anticipate the specific random configuration. In principle, `BeamDancer` would be compromised only if an adversary has access to the specific beamforming configuration. However, even with such information, the adversary should still decompose and differentiate the signals arriving from different directions, which is a non-trivial task, especially in the sub-6 GHz band where scattering is very rich. We also note that—unlike fingerprinting works like the one in [23] where knowing the distribution of the randomized fingerprints could compromise the identity of the transmitter—knowing the specific randomization values does not help the attacker, and very complex analysis of the channel frequency response in indoor environments is still needed on the adversary's side to mount a successful attack.

Sensing NDP-A and NDP frames. It could be argued that an adversary can perform wireless sensing by exploiting NDP-A, NDP, or the compressed beamforming report only, which indeed are not beamformed. According to our threat model (Section III), it is not possible to use the compressed beamforming report sent by the STA because the STA can potentially move around the room and change the effect of the channel completely. Moreover, the NDP-A and the NDP could be transmitted with various combinations of transmitting antennae, making it impossible to exploit them to perform accurate sensing without knowing the exact contributions of each antenna.

Future Extensions. `BeamDancer` can be an effective research tool for experimental analysis of MIMO beamforming in quite complex setups, i.e., 80 MHz 11ac single stream transmissions on commercial devices. While in this paper we consider a 3x1 setup, the system already supports also 4x1. In the future, we plan to extend the system to support up to 4x4 MIMO setups. This requires a new method for handling long reports, which are fragmented into multiple frames. In particular, the firmware of the beamformee must be modified to compose each fragment on the fly by loading the corresponding angles from a larger memory area. Finally, we

would like porting BeamDancer to 802.11ax, for which CSI extraction is already available [35, 37]. Specifically, 802.11ax is interesting for research on MU-MIMO OFDMA systems where a tool for studying the effect of beamforming on multi-destination transmissions is still highly needed.

VII. Related Work

Today, while research on device-free Wi-Fi sensing is thriving, only a few works are focusing on privacy to restrain sensing when not authorized. Since our contribution is focused on a practical privacy-preserving technique, in this section we discuss the major works that aim to hinder Wi-Fi sensing and examine the main differences from BeamDancer.

To our knowledge, the seminal work in [25] is the first to propose an obfuscation method to avoid illegitimate sensing. A secondary full-duplex device, called *obfuscator* and linked to the legitimate access points, is used to transmit slightly different phase-delayed versions of the signal transmitted by the AP. These "echoes" are shown to effectively disrupt adversarial sensing and can be used to improve the signal to legitimate receivers at the same time. However, the results show that the system's effectiveness strongly depends on the relative positions of the transmitter, the obfuscator, and the eavesdropper. Moreover, since the obfuscator is subject to very stringent timing requirements, the hardware device is emulated using a software-defined radio (SDR). In [26], the authors propose a different approach that does not rely on a full-duplex transceiver and strict coordination with legitimate APs. The work employs a secondary transmitter with a mechanically rotating antenna to direct interference toward the adversaries. Again, depending on the location of the malicious sensor and the obfuscating device, the adversary might have access to the "clean" information. Furthermore, the rotation pattern can be predicted easily. BeamDancer provides a much more effective and generalizable solution to the problem than these works. In particular, BeamDancer does not require an additional specialized device to obfuscate the radio signals and is implemented directly on commercial devices, enabling quick and easy deployment in real application scenarios.

The work in [38] proposes a new technique to ensure privacy at the physical layer through a significant modification of the transmitted Wi-Fi signal. The core idea is to pre-filter every Wi-Fi frame just before transmission, using an artificial channel response that can arbitrarily change over time. The principle of obfuscation is similar to the one presented in this work, i.e., is based on impeding any fingerprinting of the wireless channel. Similarly to our work, the technique in [38] does not require additional specialized hardware. However, crafting artificial channel responses that are still reasonable and do not hamper legitimate communications too much is a very complex problem that is still open. BeamDancer solves this issue by leveraging as much as possible the standard communication procedures established by commercial nodes.

In [39], the authors consider a specific adversarial scenario in which the attacker tries to recognize finger motion in public spaces using Wi-Fi signals. Even in this case, the work proposes a solution that does not require additional hardware. However, the countermeasure against unauthorized sensing relies on creating strong interference on adjacent Wi-Fi channels. While this is effective in hindering sensing, we argue that it also reduces the network's capacity. On the contrary, BeamDancer does not require spurious transmissions and only utilizes the in-band resources of the transmitters.

The works in [27, 28] investigate the topic of physical-layer privacy using a reconfigurable intelligent surface (RIS) instead of radio transceivers. In [28], the authors present a technique to ensure privacy against active adversaries using millimeter-wave radars, which is out of the scope of this work since it does not provide privacy against passive Wi-Fi sensing. In [27], the authors show that a smart surface can be used to modify the magnitude of the signals received by the adversary, thwarting motion detection applications. Still, tampering only with the amplitude cannot prevent eavesdroppers from leveraging phase information, and even with this method the location of the RIS is essential to the scheme's effectiveness.

To summarize, we found that BeamDancer presents two main advantages with respect to other state-of-the-art works: i) no additional hardware is needed to secure Wi-Fi communications from adversarial sensing in public places; ii) no modifications are made to the Wi-Fi protocol, and only minor modifications are required on standard nodes with limited drawbacks in terms of communication performance.

We emphasize that BeamDancer represents a highly practical contribution to the community. While we acknowledge that establishing theoretical foundations is paramount, we did not find any rigorous analytical framework for privacy-preserving signal manipulation in our literature review. A few works so far investigated the problem of estimating the sensing performance using an information-theoretic approach [40]. However, to our knowledge, no work has yet considered this problem from a privacy perspective. We found that modeling the effects of BeamDancer (or other equivalent systems) in non-trivial indoor environments is, in fact, a complex task that would require analytical tools completely different from the ones presented in the present work. Still, this kind of analysis is an extremely interesting research topic that should be considered for future work to identify the theoretical limits of this promising field.

VIII. Conclusion

In this paper, we proposed a practical countermeasure against the physical-layer privacy issues posed by device-free wireless sensing Specifically, we enable for the first time complete control of the beamforming configuration on commodity Wi-Fi devices. Using this capability, we apply straightforward variations to the beamforming vector to effectively obfuscate human localization and motion in indoor environments. Moreover, we show that by using this solution we can preserve almost entirely legitimate Wi-Fi traffic.

While in this work we use Beamer to thwart illegitimate sensing, we believe that this tool can be used in many other areas for low-layer experimentation with COTS devices. To facilitate the reproducibility of the results and to bolster research on privacy-aware wireless sensing, we make the tool publicly available at https://github.com/seemoo-lab/beamdancer.

REFERENCES

[1] 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)*, ACM, 2015, pp. 65–76.

[2] Y. Gu et al., "EmoSense: Data-driven emotion sensing via off-the-shelf WiFi devices," in *Proceedings of the 2018 IEEE International Conference on Communications (ICC)*, IEEE, 2018, pp. 1–6.

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

[4] E. Soltanaghaei et al., "Robust and practical WiFi human sensing using on-device learning with a domain adaptive model," in *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, ACM, 2020, pp. 150–159.

[5] Y. Ma et al., "Location- and person-independent activity recognition with WiFi, deep neural networks, and reinforcement learning," *ACM Transactions on Internet of Things*, vol. 2, no. 1, Jan. 2021.

[6] H. Li, W. Yang, J. Wang, Y. Xu, and L. Huang, "WiFinger: Talk to your smart devices with finger-grained gesture," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, ACM, 2016, pp. 250–261.

[7] G. Wang, Y. Zou, Z. Zhou, K. Wu, and L. M. Ni, "We can hear you with Wi-Fi!" *IEEE Transactions on Mobile Computing*, vol. 15, no. 11, pp. 2907–2920, 2016.

[8] J. Qiu, P. Zheng, K. Chi, R. Xu, and J. Liu, "Respiration monitoring in high-dynamic environments via combining multiple WiFi channels based on wire direct connection between RX/TX," *IEEE Internet of Things Journal*, vol. 10, no. 2, pp. 1558–1573, 2023.

[9] U. Ha, S. Madani, and F. Adib, "WiStress: Contactless stress monitoring using wireless signals," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 5, no. 3, 2021.

[10] M. Cominelli, F. Kosterhon, F. Gringoli, R. Lo Cigno, and A. Asadi, "IEEE 802.11 CSI randomization to preserve location privacy: An empirical evaluation in different scenarios," *Computer Networks*, vol. 191, 2021.

[11] Y. Xie, J. Xiong, M. Li, and K. Jamieson, "mD-Track: Leveraging multi-dimensionality for passive indoor Wi-Fi tracking," in *Proceedings of the 25th Annual International Conference on Mobile Computing and Networking (MobiCom)*, ACM, 2019.

[12] B. Korany, H. Cai, and Y. Mostofi, "Multiple people identification through walls using off-the-shelf WiFi," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6963–6974, 2021.

[13] Y. Zhu et al., "Et tu Alexa? When commodity WiFi devices turn into adversarial motion sensors," in *Proceedings of the 2020 Network and Distributed System Security Symposium*, Internet Society, 2020.

[14] 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)*, ACM, 2015, pp. 90–102.

[15] M. Li et al., "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)*, ACM, 2016, pp. 1068–1079.

[16] D. Zhang, Y. Hu, Y. Chen, and B. Zeng, "BreathTrack: Tracking indoor human breath status via commodity WiFi," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3899–3911, 2019.

[17] C. Feng, J. Xiong, L. Chang, F. Wang, J. Wang, and D. Fang, "RF-Identity: Non-intrusive person identification based on commodity RFID devices," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 5, no. 1, Mar. 2021.

[18] B. Korany, C. R. Karanam, H. Cai, and Y. Mostofi, "XModal-ID: Using WiFi for through-wall person identification from candidate video footage," in *Proceedings of the 25th Annual International Conference on Mobile Computing and Networking (MobiCom)*, ACM, 2019.

[19] S. Oh, T. Vu, M. Gruteser, and S. Banerjee, "Phantom: Physical layer cooperation for location privacy protection," in *Proceedings of IEEE INFOCOM*, 2012, pp. 3061–3065.

[20] Y. Ma, G. Zhou, and S. Wang, "WiFi sensing with channel state information: A survey," *ACM Computing Surveys*, vol. 52, no. 3, 2019.

[21] S. Taha and X. Shen, "A physical-layer location privacy-preserving scheme for mobile public hotspots in NEMO-based VANETs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1665–1680, 2013.

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

[23] L. F. Abanto-Leon, A. Bäuml, G. H. Sim, M. Hollick, and A. Asadi, "Stay connected, leave no trace: Enhancing security and privacy in WiFi via obfuscating radiometric fingerprints," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 4, no. 3, Nov. 2020.

[24] R. Ayyalasomayajula, A. Arun, W. Sun, and D. Bharadia, "Users are closer than they appear: Protecting user location from WiFi APs," in *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications (HotMobile)*, ACM, 2023, pp. 124–130.

[25] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, and A. Arora, "PhyCloak: Obfuscating sensing from communication signals," in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation (NSDI)*, USENIX Association, 2016, pp. 685–699.

[26] Y. Yao et al., "Aegis: An interference-negligible RF sensing shield," in *Proceedings of IEEE INFOCOM*, 2018, pp. 1718–1726.

[27] P. Staat et al., "IRShield: A countermeasure against adversarial physical-layer wireless sensing," in *Proceedings of the 43rd IEEE Symposium on Security and Privacy (SP)*, 2022, pp. 1705–1721.

[28] J. Shenoy, Z. Liu, B. Tao, Z. Kabelac, and D. Vasisht, "RF-protect: Privacy against device-free human tracking," in *Proceedings of the ACM SIGCOMM 2022 Conference*, ACM, 2022, pp. 588–600.

[29] W. Liu and S. Weiss, *Wideband beamforming: concepts and techniques*. John Wiley & Sons, 2010, ISBN: 0470713925.

[30] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005, ISBN: 9780521845274.

[31] F. Wang, J. Feng, Y. Zhao, X. Zhang, S. Zhang, and J. Han, "Joint activity recognition and indoor localization with WiFi fingerprints," *IEEE Access*, vol. 7, 2019.

[32] Y. Zheng et al., "Zero-effort cross-domain gesture recognition with Wi-Fi," in *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys)*, ACM, 2019, pp. 313–325.

[33] M. Schulz, D. Wegemer, and M. Hollick, "The Nexmon firmware analysis and modification framework: Empowering researchers to enhance Wi-Fi devices," *Computer Communications*, vol. 129, pp. 269–285, 2018.

[34] J. Link, D. Breuer, F. Gringoli, and M. Hollick, "Rolling the D11: An emulation game for the whole BCM43 family," in *Proceedings of the 17th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*, ACM, 2023, pp. 88–95.

[35] F. Gringoli, M. Cominelli, A. Blanco, and J. Widmer, "AX-CSI: Enabling CSI extraction on commercial 802.11ax Wi-Fi platforms," in *Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (WiNTECH)*, ACM, 2022, pp. 46–53.

[36] X. Li et al., "IndoTrack: Device-free indoor human tracking with commodity Wi-Fi," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 3, Sep. 2017.

[37] Z. Jiang et al., "Eliminating the barriers: Demystifying Wi-Fi baseband design and introducing the picoscenes Wi-Fi sensing platform," *IEEE Internet of Things Journal*, vol. 9, no. 6, pp. 4476–4496, 2022.

[38] R. L. Cigno, F. Gringoli, M. Cominelli, and L. Ghiro, "Integrating CSI sensing in wireless networks: Challenges to privacy and countermeasures," *IEEE Network*, vol. 36, no. 4, pp. 174–180, 2022.

[39] J. Zhang et al., "Defeat your enemy hiding behind public WiFi: Wiguard can protect your sensitive information from CSI-based attack," *Applied Sciences*, vol. 8, no. 4, 2018.

[40] M. Ahmadipour, M. Kobayashi, M. Wigger, and G. Caire, "An information-theoretic approach to joint sensing and communication," *IEEE Transactions on Information Theory*, vol. 70, no. 2, pp. 1124–1146, 2024.

Marco Cominelli is a researcher at Politecnico di Milano, Italy. He received his Ph.D. in Information Engineering in 2023 from the University of Brescia, Italy. He has been a visiting research student at the University of Edinburgh and Northeastern University. His research interests include studying device-free wireless sensing techniques and their impact on users' security and privacy.

Shaghayegh Shahcheraghi is a Ph.D. student at The Ohio State University. She was a Marie-Curie Early-Stage Researcher for the EU Horizon 2020 MINTS project at TU Darmstadt, Germany. She obtained her M.Sc. degree in Telecommunication engineering with the focus on signal processing from Politecnico di Milano. She received her B. Sc and M.Sc. degrees in Electrical engineering from Shiraz University. Her research interests include joint communication and sensing, opportunistic navigation, cognitive radio, and software-defined radio.

Jakob Link is a researcher at the Technical University of Darmstadt as part of the Secure Mobile Networking Lab (SEEMOO). He received the master's degree in IT-Security from the Technical University of Darmstadt in 2021. His focus lies on re-purposing of Wireless LAN COTS devices through reverse-engineering and low-level firmware modifications, with the vision of seeing more theoretical and lab-only research coming to practical real-world setups.

Matthias Hollick is a Full Professor of Computer Science at the Technical University of Darmstadt where he is heading the Secure Mobile Networking Lab (SEEMOO) since 2009. Over the last 20 years, he has been contributing practical, award-winning solutions in decentralized wireless/mobile networking as well as network security and privacy. His current research focus is on enhancing ICT resilience, thus empowering humans and society to better deal with crises.

Federico Cerutti is an associate professor at the University of Brescia, Italy. His research is in learning and reasoning with uncertain and sparse data for supporting (cyber-threat) intelligence analysis. He is a Rita Levi-Montalcini Fellowship laureate, an Honorary Senior Lecturer at Cardiff University (UK), a Visiting Fellow at the University of Southampton (UK), and the Chair of the University of Brescia's local branch of the Italian Cybersecurity National Laboratory. Contact him at federico.cerutti@unibs.it

Francesco Gringoli is Full Professor at the University of Brescia, Italy. He received the master's degree in Telecommunications Engineering from the University of Padova, Italy, in 1998 and the Ph.D. degree in Information Engineering from the University of Brescia, Italy, in 2002. His research interests include security assessment, performance evaluation and medium access control in Wireless LANs.

Arash Asadi is an Assistant Professor at Embedded Systems Group at TU Delft where he leads the Wireless Communication and Sensing Lab (WISE). His research is focused on wireless communication and sensing for 6G networks. He is a recipient of several awards, including Athena Young Investigator award from TU Darmstadt and outstanding PhD and master thesis awards from UC3M. Some of his papers on D2D communication have appeared in IEEE COMSOC best reading topics on D2D communication and IEEE COMSOC Tech Focus.