AntiSense: Standard-Compliant CSI Obfuscation Against Unauthorized Wi-Fi Sensing

Marco Cominelli, Francesco Gringoli, Renato Lo Cigno

Department of Information Engineering (DII) — University of Brescia, Italy

Abstract

Channel State Information (CSI)-based localization with 802.11 has been proven feasible in multiple scenarios and is becoming a serious threat to people's privacy in workplaces, at home, and maybe even outdoors. Countering unauthorized localization without hampering communications is a non-trivial task, although some very recent works suggest that it is feasible with marginal modification of the 802.11 transmission chain, but this requires modifying 802.11 devices. Furthermore, if the attacker controls two devices and not just a receiver, transmission side signal manipulation cannot help. This work explores the possibility of countering CSI based localization with an active device that, instead of jamming signals to avoid that a malicious receiver exploits CSI information to locate a person, superimpose on frames a copy of the same frame signal whose goal is not destroying reception as in jamming, but only obfuscate the location-relevant information carried by the CSI. A prototype implementation and early results look promising; they show the feasibility of location obfuscation with high efficiency and excellent preservation of communication performance, and indicate that the technique works both against *passive* attacks, where the attacker controls only a receiver, and *active* ones, where he/she controls both a transmitter and a receiver. These results pave the road for further research on smart spaces that preserve users' privacy with a technical solution and not only via legal prescriptions.

Keywords: CSI-based Wi-Fi Localization, Smart Spaces, Privacy Protection, Location Obfuscation

1. Introduction and Background

Wi-Fi sensing is a broad topic that is receiving attention from both academia and industry. Exploring the environment and getting information on it through electromagnetic waves is surely not novel, but doing it as a side-task of wireless communications is indeed novel and compelling. Exploiting Wi-Fi signals to do so is a natural choice given the ubiquity of Wi-Fi communications, and it is indeed enabled quite naturally by the same technology supporting Wi-Fi evolution: Advanced channel sounding. The estimation of the propagation channel through the so-called CSI is indeed one of the enabling mechanisms to support multi-gigabit throughput in 802.11 systems. The development of new CSI-based equalization techniques inside next-generation 802.11be (branded Wi-Fi 7 by the Wi-Fi Alliance) allows up to 16 spatial streams and a data rate of 46 Gbit/s [1]. As the goal of equalization moves from simple compensation of the channel distortion to a complex operation that resembles more the multi-reflection and multi-refraction analysis of a synthetic aperture radar, the idea of using this information for ambient sounding so far carried out with radars emerged, sometimes calling this operation *channel charting* [2, 3], especially when it is done with unsupervised techniques. Among all possible applications, the one that received more attention is localization. The attention to CSI-based localization was brought by early works [4, 5, 6, 7] nearly ten years ago, immediately proving that CSI-based localization techniques can outperform traditional Received Signal Strength Indicator (RSSI)-based techniques.

After these initial works, the topic flourished, with proposals exploiting massive Multiple-Input Multiple-Output (MIMO) [8, 9] or Bayesian estimators [10], or broadening the scope of positioning, for instance to identify activities and gestures [11, 12, 13], for health and medical applications [14] or even to "hear" people [15]. Many other works, flavors and papers exist, interesting but not strictly related to our contribution, as they relate to applications of CSI-based localization and not to fundamental techniques to perform it.

Recent years witnessed the explosion of Machine Learning (ML) and Artificial Intelligence (AI) methodologies applied to the topic [9, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25], which achieve astounding results using different classification or analysis techniques, often involving Deep Learning or Reinforcement Learning.

What none of these works have ever discussed are the following questions:

- How ethical or intrusive is CSI-based localization?
- Is it possible to prevent unauthorized use of CSI-based localization?
- What is the cost in terms of communication performance we have to pay to prevent CSI-based localization?

These are, instead, precisely the question our work addresses. The answer to the first one is clear: CSI-based localization is potentially highly intrusive, and tracking people without their consent is unethical. Furthermore, we should also consider the security problems that can arise when an attacker can tell if and how many people are inside a room, house, or laboratory, where they stand, or if they move. Thus the challenge of our contribution is to give some answers to the last two questions.

The technology is particularly invasive because the attack can be both passive or active, or even combined, and the victim is completely unaware of the attack: she/he does not need to wear any device to be located and has no means to detect the attack. In a passive attack, attackers capture frames transmitted by sources in well-known positions, like almost all the Access Points (APs) we usually have at home or work. Attackers do not need to control such transmitters; they only have to place a receiver somewhere in the same room or just outside it for precise, Cartesian localization or classification-based positioning. In an active attack, instead, the attacker controls both a transmitter and one or more receivers. The attacker has more freedom and power in the attack: the transmitter and receiver can be placed in strategic positions, but in general he has to place these devices outside the room where the victim stands to avoid easy spotting.

This paper overviews CSI-based localization fundamentals to make the contribution self-contained, analyzes the works that tackle the same privacy-preserving problem, and sketches the general principles of localization privacy protection based on the obfuscation of the location information carried by the CSI. The core contribution of this work is the design, implementation, and analysis of an obfuscation technique based on the injection in the channel of artificial signal reflections that can prevent *both passive and active attacks*. One may argue that an active attack is detectable, as it implies "illegitimate" on-air traffic. The observation is valid but of limited use: who cares when yet another Service Set Identifier (SSID) appears at home? Not to mention public spaces, where nobody can control who the legitimate Wi-Fi users are. Even in office environments, it is tough to imagine that the victim can identify an attack because the victim can be an employee and the attacker the employer who wants to control his/her employees beyond what legislation permits. Alternatively, it could be a double-dealing worker who installs a pair of devices in an office, lab, or room to monitor the position and movements of fellow workers.

We already tackled the problem of localization obfuscation for *passive attacks* in [26, 27, 28], as discussed in detail in Sect. 3. The technique adopted in those works, albeit apparently similar to the one presented here, is based on the manipulation of the CSI at the transmitter, and cannot clearly cope with active attacks, where the attacker controls also the transmitter. This work, extending [29], deals instead with both *passive and active attacks* exploiting an external signal reflector.

2. CSI-based Localization

No matter the technology used to extract location information from the CSI, this information must be present in the signal itself. This information is embedded in the signal during propagation and carries pieces of information on people's presence and location because a human body absorbs, scatters, and reflects Wi-Fi signals. Fig. 1 reports the amplitude and unwrapped phase of 100 frames collected with a person standing in two different locations in our lab in Brescia. The exact location is irrelevant, but it is clear that the amplitude of consecutively received frames is remarkably constant in the same location, while it significantly changes when moving from one location to another. Repeating

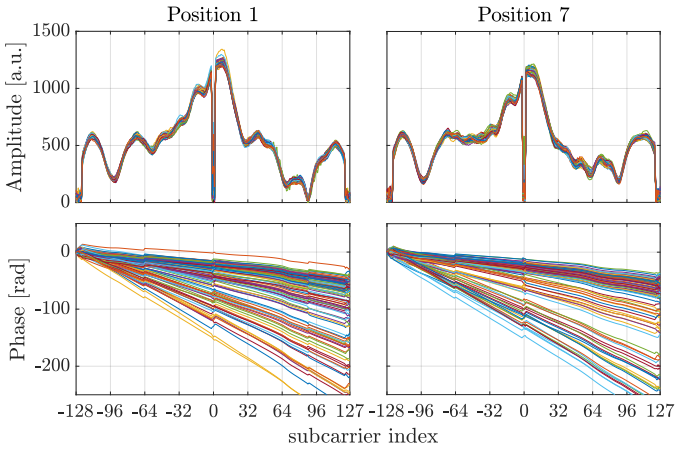

Figure 1: Amplitude and unwrapped phase of the CSI collected from 100 frames with a person standing in two different locations in our lab in Brescia.

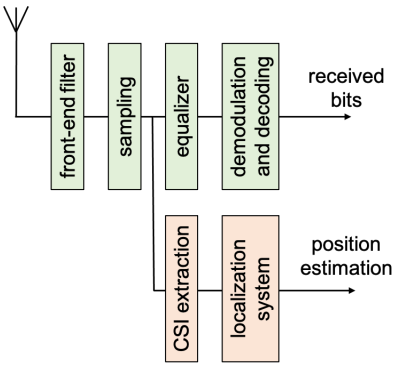

Figure 2: 802.11 modified receiver to infer people location; first the localization system is trained with a person standing in positions of interest building a reference set, while during the attack the localization system infers the position where the person can be classifying CSI data on the reference set.

the experiment at different times shows a time-based variation, but still, the CSI carries enough location-specific information to allow a proper algorithm to infer the person's location. It is clear that both the amplitude and phase are affected, although the linear variation of the phase with the carrier frequency has nothing to do with localization, and only phase jumps are relevant.

The transmission technique has an essential importance in CSI manipulation, and the structure of Wi-Fi frames, their generation, and filtering at both the transmitter and the receiver are fundamental to understand localization techniques fully. Fig. 2 sketches the diagram of a single antenna receiver modified to retrieve information on people's localization. At the receiver, after sampling the incoming signal, samples are duplicated. The standard data path goes through the equalizer that compensates the channel distortions and then to the demodulation and decoding blocks that yield the frame bits if decoding is successful. Instead, the duplicated samples enter the localization system that, exploiting the same CSI used by the equalizer, estimates the person's location. The CSI is implicitly carried by the training sequences at the beginning of the frame, and in particular by

the Long Training Sequence (LTS), whose bits and structure are known, allowing the equalizer to compute the channel frequency response, and the localization system to use this information to fingerprint the person's position.

The localization techniques that have recently received more attention are based on neural networks (NNs) trained with someone standing in known positions and then, during the attack, determine the position of a person based on the training fingerprints. Given a localization technique, the *system* that implements it can follow several design lines. One critical design decision regards the transmissions. The localization system can be *passive*, i.e., it exploits the data packets usually sent by users, or it can be *active*, i.e., it uses frames that are sent by a device specifically to perform the localization.

In some sense, a passive system is more accessible as only a specialized receiver is needed to perform the localization. However, the frames used for localization must come from a transmitter in a fixed location (not necessarily known) because the change in the CSI determined by the moving transmitter will taint the collected fingerprints. This is not a problem in most cases since APs are fixed and generate most of the traffic; thus, the localization device only needs to filter frames transmitted by the AP to achieve its goal. On the other hand, as we have shown in [26, 27], it is possible to obfuscate the information on localization carried by the CSI by properly manipulating the transmitted frames.

An active localization system, instead, requires an attacker to use both a transmitter and a receiver. While this is somewhat more complex and detectable because frames on-air that do not belong to a legitimate Basic Service Set (BSS), there is no way to hinder the localization by manipulating the transmitted frames, as the transmitter is not controlled by legitimate users but by the attacker himself.

In this work, we consider both passive and active localization systems based on fingerprinting and a single transmitter-receiver pair. We do not consider localization techniques based on the angle of arrival, and we do not consider the possibility of having more than one receiver that works coordinately to improve the localization accuracy. All of these topics are extremely interesting, but they are outside the scope of this paper.

2.1. Localization Adopted

The localization technique adopted in this work relies on a Convolutional Neural Network (CNN) to perform a classification task. The design of the CNN is inspired by the work in [19] and refined in [30] and it is not a contribution of this paper, thus we refer the interested reader to the original works. Based on these two works, we have developed an efficient implementation within the CSI-MURDER project[1] with good localization efficiency and properties. The CNN-based localization system can work both with active and passive attacks, as long as the trasmitter—either controlled by the attacker or by some other user depending on the case—stays in a fixed position. In both cases (active and passive scenario), the attack is mounted in two phases: first, the attacker trains the CNN using data collected with the help of another person standing in the target positions; then, when the victim enters the room, the attacker can use the trained model to associate the received CSI to the victim's position. In the remaining part of this section, we briefly present the main features of the localization framework. The interested reader can find a more complete presentation of the properties of our localization system in [26], and as already mentioned in the documentation of the software published on-line.

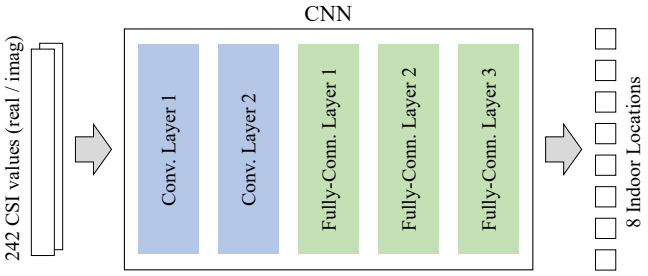

Figure 3: Architecture of the CNN used by our localization system.

In Fig. 3 we show a high-level representation of the CNN architecture. One CSI data point is extracted from each 802.11 frame correctly decoded at the receiver. Each CSI data point is an array of complex values (the IQ samples) computed at the receiver to estimate the channel's frequency response. In this work, we consider 802.11ac frames transmitted on 80 MHz channels; therefore, each raw CSI consists of 256 complex values. During the preprocessing phase, we remove the subcarriers at the edges of the spectrum as well as the three central ones because they are all suppressed by the modulator and cannot carry any information about the propagation channel. The input of the CNN is thus a 242×2 matrix. The first two convolutional layers of the CNN shown in Fig. 3 extract complex features from the input data by exploiting the similarity of adjacent frequencies. In cascade to the convolutional layers, there are three fully-connected layers. The output of the last layer corresponds to a choice among one of the possible classes, i.e., positions. The range and scope of this network are relatively flexible: the number of classes to predict can be changed as needed without modifying the other layers of the CNN and it will still have good fingerprinting performance (clearly, we must train the network every time we use a new dataset). All the layers but the last (which uses a softmax function) use a standard Rectified Linear Unit (ReLU) activation function. The Adaptive Momentum Estimation (ADAM) algorithm is used to adjust the weights of the CNN during the training phase.

As shown also in [26, 27, 30], the CNN structure can extract features from raw data CSI to perform many different localization tasks in indoor environments: from coarse to finer classification in specific locations to, finally, fine-grained training that enables (x, y) coordinates estimation. However, localization becomes more fragile as the precision requirement is increased; thus, to validate the obfuscation technique in this paper we decided to rely on a robust 8-location-points classification.

[1] Further details on this project, the software produced an so forth can be found at `https://ans.unibs.it/projects/csi-murder/`. The software is going to be released as FOSS, so that results can be replicated also in other scenarios.

3. Related Work

A review of CSI-based localization techniques and systems, beyond the overview in Sect. 2, is out of the scope of this paper. We focus instead only on works whose goal is localization obfuscation or localization privacy protection, a topic, as already highlighted, that received so far less attention than it deserves.

The first countermeasure against Wi-Fi sensing attacks has been implemented in [31] for preventing gesture recognition; similarly to our proposal, this system relies on an additional component acting as a relay placed in the environment. We also inherited from this work the term *obfuscation* with the meaning of distorting the information imprinted on a frame by the environment, contrasted to the more common *jamming* that instead superimposes a different signal (possibly noise) with the goal of making the frame useless, thus also killing communication capabilities. However, the artificial reflection techniques proposed here are different, and [31] is focused on gesture recognition rather than localization. Furthermore, in our opinion, [31] opens a research field rather than writing *"The End"* on it, and our work adds novel insight into this fascinating topic.

The idea in common with [31] is the potential use of dynamically changing reflective surfaces or, in general, of active devices that randomly change the electromagnetic (EM) environment. The technology to obtain this random behavior of the EM environment is outside the scope of this work and ranges from reconfigurable intelligent surfaces to metasurfaces and to more traditional fast relays able to retransmit a signal with slight random delays. We only mention here two works that, even if they do not explicitly mention location privacy, are in some way closer to our work. The first one presents a simple yet effective device implementing a passive reflector with different delays corresponding to 0, $\frac{\lambda}{4}$, and $\frac{\lambda}{2}$ additional paths, which can be selected randomly [32]. Though conceptually very broad, the device is in some sense tailored for Wi-Fi, which connects this work to our contribution. The authors are more concerned with the challenge of using their proposal to enhance communications. However, we observe that several of these devices that act randomly and without coordination—adding random delay on a per-frame basis—can well be building blocks of the obfuscation system we propose. As we show at the end of this paper, obfuscation based on intelligent reflecting/relaying does not hamper the communication performance, and we claim that a privacy-preserving ambient is part of a Smart Space, further linking our work to [32]. The second one takes a *networking perspective* to smart EM spaces, called here programmable wireless environment [33], and it is related to our work mainly because, as we do, abstract from physical layer details and focuses on potential goals and services. The work is extensive, but the authors explicitly mention the goal of privacy protection, even if referred to eavesdropping rather than localization.

In previous works [26, 27, 28], we have focused on passive attacks only, and albeit the final goal is the same, the methodology adopted in those works is different because the CSI manipulation is implemented directly at the transmitter as a pre-distortion of the frames. No additional devices are needed, and the countermeasures can be implemented in the AP alone, as it is the only device of a BSS that we can consider fixed in a given position. In [26], we presented a simple proof-of-concept showing that proper manipulation at the transmitter can obfuscate a person's actual position. The work in [27] extends the contribution with a deeper analysis in which we have defined the localization problem as a more straightforward classification problem rather than positioning in the Cartesian coordinate space. Still, the obfuscation technique was proven to remain very effective. Playing the devil's advocate, in [28] we conjectured that a multi-point multi-receiver attack would be far more powerful and far more challenging to counter. While the former claim is valid, the latter turned out to be false, as a proper pre-distortion at the transmitter completely obfuscates the victim position even if the attacker controls five different receivers in different positions. In [28], we also introduced the idea of manipulating the signal according to a Markov random process that introduces memory in the random distortion and makes the distortion more similar to the one introduced by a person's motion or by erratic changes in the EM environment. The work we present in this paper, instead, extends the work in [29] to make a complete archival contribution. It focuses on obfuscation induced by additional time-varying reflections of the frames on-air and can thus counter both active and passive attacks with a single methodology.

The authors of [34] manipulate the CSI to avoiding device radiometric fingerprinting to help preventing impersonation attacks. Their paper's goal is not directly location obfuscation; however, the techniques used are similar to those we use in this work, and we do not exclude that, in case a person holds a Wi-Fi device, the double attack identifying the device and the location of the person is feasible.

Finally, we can conceive reactive jamming devices that "kill" the frames used for localization attacks, adapting techniques like [35, 36] to our scope. To the best of our knowledge, this has never been proposed in the literature, so it is difficult to state how effective it can be. Moreover, this approach would require to know that a localization attack is underway, and the jamming device must recognize the illegitimate traffic and try to kill those frames only. By contrast, our approach is transparent, as it does not affect the reception of frames significantly, indeed it can even improve it, so that the obfuscating device can be active at any time on any frame.

4. Localization Attack Models

A malicious user with the ability to overhear Wi-Fi traffic can perform two different types of localization attacks. The classification is based on the capability of the attacker to control transmissions: if she/he can control a transmitter with a fixed position, the attack is *active*; otherwise, if he must rely on standard, legitimate traffic, the attack is *passive*. Fig. 4 depicts the scenario in both cases: Red arrows refer to the active attack, and blue ones to the passive one. Solid lines define the control of devices, dashed the transmitted frames, and dotted ones the obfuscator replica. The attacker can eventually use more than one receiver, presumably improving the localization performance by correlating the position estimation by all the receivers. This possibility was explored in [28] for passive attacks only and with

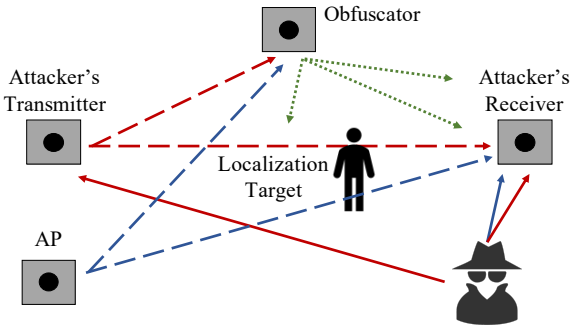

Figure 4: In an active attack, the attacker controls both the transmitter and at least one receiver (red arrows); in a passive attack he/she controls only the receivers and rely on traffic sent by an AP (blue arrows). In both cases the location privacy of the victim can be protected if an active device is able to obfuscate the CSI of frames (green dotted arrows indicating a reflected or relayed frame.

a different obfuscation technique, but it is beyond the scope of this paper.

4.1. Active Attack's Model

During an active localization attack, the attacker controls one transmitter and at least one receiver; therefore, the attacker has complete control over the transmission chain, and the only way to interfere with his/her intrusion is by actively mingling transmitted frames on the channel. Red arrows in Fig. 4 depict this attack model. A person is standing in a room (can be an office or home or anywhere) and the attacker aims to collect information on the person's position. To achieve this goal, the attacker has installed a standard Wi-Fi transmitter and a modified receiver. The receiver implements the localization technique described in Sect. 2 and can access the room at some time to train the CNN. After training the CNN, the attacker can configure the system to send frames periodically and estimate the person's position when she/he is in the room.

4.2. Passive Attack's Model

Passive attacks instead are characterized by the collection of frames sent by a legitimate AP on which the attacker has no control, as indicated by the blue arrows in Fig. 4. The attacker cannot control which frames to send and when, and training of the localization system may be more complex. We have studied this scenario in previous works [26, 27, 28], adopting a radically different obfuscation technique based on the pre-distortion of transmitted frames at the AP. However, active attacks could not be prevented with such a technique. Here we are interested in understanding if one single obfuscation method can prevent unauthorized localization in both attack models.

5. Obfuscation Principles and Requirements

As discussed in Sect. 2, CSI-based localization exploits the information carried by CSI on the EM environment. The goal of an obfuscator is blurring this information without destroying communication capabilities. We are interested in designing an active device that can prevent unauthorized localization against all types of attacks. The device must be able to randomly change the channel response "reflecting" the incoming signal with an adequately designed amplification, delay, and phase distortion. The key idea is that this device acts as an additional feature of the propagation environment, changing it in such a way that the localization system cannot identify the position of the person based on the CSI fingerprint because this latter contains too much random information to allow identifying the features of a target position. This device is the *obfuscator* in Fig. 4, but since it conceptually reflects the Wi-Fi signal, we also call it the *reflector* throughout the paper.

The obfuscator cannot operate only on non-legitimate frames, simply because the reflection delay must be well below a single symbol duration, and it is impossible to read the Medium Access Control (MAC) addresses before reflection. Since the CSI information is embedded in preambles, stopping the reflection when MAC addresses are available would be detrimental to frame reception as it is equivalent to have a channel coherence time shorter than a frame.

Like an attacker controlling more than one receiver, the reflectors can also be more than one, possibly enhancing the obfuscation performance. This possibility points toward the idea of Smart Spaces, where the EM environment is active and participates both in enhancing the performance of communication and in protecting users against intrusions.

A good obfuscation system must meet three fundamental requirements: *i)* it does not hamper communication performance (ideally, with an active device, it should even enhance it); *ii)* it alters the signal in ways that are compatible with people's movements; *iii)* Its random behavior cannot be reversed in a reasonable amount of time.

The last two requirements are needed to guarantee that even sophisticated analysis cannot filter out the obfuscation distortion so that privacy is protected almost surely. In other words, the attacker should not get any information that is significantly better than a random guess.

6. Randomized Reflection Strategies

In a real anti-localization system, the obfuscator can be a repeater that mimics a reflective surface or, in a more futuristic scenario, a reflective intelligent surface (or smart space [37, 33]) changing its properties under the control of a proper obfuscation function. The goal of the obfuscator is to add one or more "reflecting paths" into the propagation environment in such a way that the behavior of the channel that embeds the information on the victim's location in the CSI becomes blurred and time-varying, confusing the localization system. At the same time, the channel distortion must remain plausible, meaning that it should allow the equalizer at the receiver to correctly compensate the distortion so that legitimate frames can be received without reducing the communication capabilities.

To fix the ideas on what an active obfuscation shall achieve, consider Fig. 5. On the left-hand side, there are 100 CSI amplitude samples collected as reference (in blue) and 100 collected

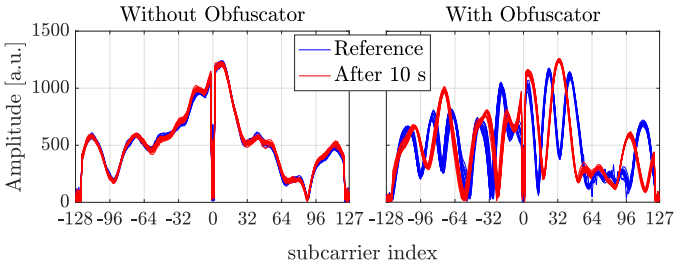

Figure 5: Effect of the active obfuscation on the CSI. In both cases the victim is standing still in one position; however, when the obfuscator is actively relaying the received signal, the channel conditions appear to change over time.

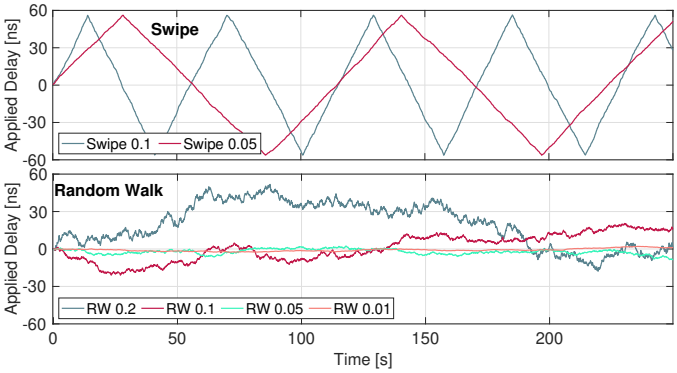

Figure 6: Processes driving the delay added by an active obfuscator, the processes are centered around zero, then a proper offset τ_{off} is added depending on the use cases. The upper plot shows a random delay swipe with $\delta = 0.1$ adn 0.05 respectively; the lower plot shows the outcome of a Markov-Uniform process with four different δ parameters.

after 10 s when a person is standing still in a given position and there is no obfuscation. On the right-hand side, instead, we repeated the same experiment with the obfuscator turned on. It is clear that the obfuscator significantly alters the propagation environment (the blue lines are very different in the two plots) and after 10 s the red lines tell a different propagation story, or at least they *mimic* it. In any case, we can conjecture that any localization technique will have a hard time in fingerprinting and classifying positions.[2]

Ideally, the outcome of the obfuscation at the receiver should be indistinguishable from a standard channel response, both as a distribution of attenuation and phase jumps in frequency and as a correlation in time. However, it is not clear if this is entirely achievable because there is a lack of experimental studies that adequately characterize the stochastic properties of the channel response.

In our preliminary contribution at WONS [29], we used a simple implementation where the delay added by the reflector is a swipe with some randomness between a minimum and a maximum, as described by Eq. (1). It is clear that such an implementation does not meet the requirement *iii)* defined in Sect. 5, as a prolonged observation with appropriate processing can reveal such a periodic behavior.

A better approach would be to use a delay that mimics a random walk of the person in the room. However, understanding what delay corresponds to each position of the person in the room is generally unfeasible, as it would require the characterization of every space we want to protect and the tuning of the obfuscation algorithm. A solution that appears feasible and robust, and that may mimic movements of people in a room, is a random walk of the delay itself, which can be efficiently implemented as a Markov random process, as defined by Eq. (2), whose output is the delay to be added to the next reflected or repeated frame.

The difference between the two approaches is clear looking at Fig. 6 where the upper plot shows the output of a simple random swipe (Eq. (1)) with two different values of the maximum delay step δ while the bottom one reports the output of the Markov-Uniform random process (Eq. (2)) with four different values of δ. The processes are centered around zero; then, in the implementation, a constant τ_{off} is added depending on the scenario to make the output compatible with the actual scenario. For instance, if the scenario is based on an active relay or reflector, then there is a minimum positive delay corresponding to the propagation time plus the minimum time required by the device to relay/reflect the signal. Let us define the minimum and the maximum admitted delays $\tau_{\min}$ and $\tau_{\max}$ respectively, so that for any frame the actual delay introduced is $\tau_{\min} \leq \tau \leq \tau_{\max}$. Setting these two limits and τ_{off} is fundamental to maintain the additional delay added by the reflector within bounds that are coherent with the ambient to be protected, as already commented. Since the movement of a person in a room is based on time and not on transmitted frames, we also define a time interval Δt, used to pilot the random delay evolution easily; $\tau(i)$ means the delay added at the i-th time interval. For further random behavior, also Δt, $\tau_{\min}$, and $\tau_{\max}$ can be random variables, but we have not explored this option in our implementation.

A random delay swipe, including τ_{off}, is defined as

$$\tau(i) = \tau_{\text{off}} + \tau(i-1) + I(\tau) \cdot \delta U_{[0,1]} \tag{1}$$

where $I(\tau)$ is an indication function that takes the value $+1$ if τ is increasing and -1 if it is decreasing, $U_{[0,1]}$ is a uniform random variable with support $[0, 1]$, and δ is the maximum allowed difference between additional delays in adjacent time intervals. Switching between increasing and decreasing behavior happens when τ reaches $\tau_{\min}$ and $\tau_{\max}$ respectively.

The Markov-Uniform delay is instead computed as

$$\tau = [\tau_{\text{off}} + \tau(i-1) + \delta U_{[-1,1]}]_{\tau_{\min}}^{\tau_{\max}} \tag{2}$$

where $U_{[-1,1]}$ is a uniform random variable with support $[-1, 1]$, and $[\cdot]_{\tau_{\min}}^{\tau_{\max}}$ indicates the clipping between $\tau_{\min}$ and $\tau_{\max}$. Other types of Markovian processes can be used for obfuscation, and the analysis to identify the one that guarantees optimal obfuscation is an interesting future work, with the possibility of finding a theoretical optimum that minimizes, or even nullify, the location information carried by the CSI.

[2]The rationale of the approach can be better appreciated dynamically, with a video that is not possible to include in a paper. We have realized this video and is available from our website at `https://ans.unibs.it/projects/csi-reflector/`.

In the case of the simple random swipe of Eq. (1), δ controls the average period of the swipe, which is $T_s = \frac{2(\tau_{\max} - \tau_{\min})}{\delta}$. In the Markov-Uniform case (Eq. (2)), instead, δ controls the probability that the process is clipped. The plots in Fig. 6 reports the realizations of both the swipe and the Markov random delays for different values of δ. In all the experiments with the swipe method we report results only for $\delta = 0.1$ only, as we have verified that this value has little influence on the results, hence the parameter is not repeated every time. For the Markov case, instead, we report results for the four values $\delta = 0.01, 0.05, 0.1, 0.2$ because we want to analyze the impact of the delay variation amplitude on both the obfuscation performance and the communication performance. Of particular interest are the two extreme values $\delta = 0.01$ and 0.2, because from the bottom plot of Fig. 6 one might argue that $\delta = 0.01$ has too small variations to obfuscate the location while $\delta = 0.2$ might be too "noisy" to guarantee good communication performance.

7. Implementation

Depending on the considered attack scenario, we can envision the transmitter as a device on which the attacker has complete control (active attack) or no control at all (passive attack). In the former case, we cannot predict the attacker's transmission and either an active reflector of a signal relay seem to be the only viable options to counter this type of attack.

Implementing the obfuscation mechanism in hardware, which implied realizing a dedicated chip, is unfortunately beyond the possibilities of our lab, let alone realizing a controllable reflective intelligent surface. Moreover, we believe that such an expensive endeavor is only justified once it is clear that the proposed technique works and is tamper-proof. Thus, we resort to software-defined radio (SDR) devices and a little "trick" to realize our proof-of-concept implementation.

Our setup consists of two SDRs—namely two Ettus USRP N300, one for the transmitter and one for the obfuscator—and a commercial AP (Asus RT-AC86U) used as the receiver. The SDR transmitter keeps sending 802.11 frames generated using the Matlab WLAN Toolbox at a constant rate of approximately one frame every 10 ms. The receiver encloses a Broadcom chipset from which we can extract the CSI data points using the tools provided by the Nexmon project [38]. Finally, the localization system works offline on the memorized CSI data points. In general, there are no strict real-time requirements to identify a person's position. In any case, once the CNN has been trained, the analysis of the CSI made by the localization system is swift: according to our tests, an Intel Core i7 clocked at 4.4 GHz takes as little as 60 μs to process the CSI data extracted from a single frame and estimate the related victim's position.

Implementing a real-time 802.11 signal relay in software is doable but still tricky: the latency introduced by typical SDR systems cannot meet the strict timing requirements, and some kind of hardware-accelerated processing is necessary. For this reason, we resort to a gimmick, as we show in Fig. 7. The two SDRs—one playing as the transmitter and the other as the obfuscator—are synchronized through a common clock source.

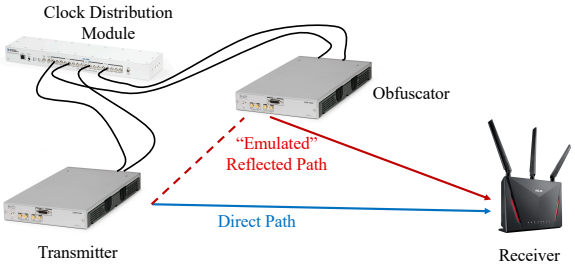

Figure 7: Schematic representation of the experimental setup; the obfuscator acts as a configurable reflector.

This common clock source is provided in our case by an Ettus Octoclock-G and consists of both a 10 MHz reference signal and a 1 Hz separate signal that allow almost perfect synchronization. Once the two SDRs are synchronized, the obfuscator can emulate the effects of a reflected path by re-transmitting the original signal with a delay, as discussed in Sect. 6.

The easiest way to apply a time delay to the signal transmitted by the obfuscator is to shift the sequence of IQ samples transmitted by the obfuscator by a certain amount of samples. However, despite being simple, this solution has a substantial limitation. Since the radio transmits samples at a fixed rate, the available bandwidth determines the time between two consecutive samples and the maximum granularity of the delay. In our implementation, the transmission rate of the N300 SDRs is 125 MSample/s, which corresponds to a sampling period of 8 ns; therefore, the minimum delay corresponding to a shift of the sequence by one sample is equivalent to a path difference of approximately 2.4 m. Moreover, this method can only emulate propagation delays multiples of such quantity, which would be an inconvenient limitation for our obfuscation system.

A better solution is to process the sequence of IQ samples in the frequency domain. Given the digital signal $x[n]$ and assuming that all the conditions on proper sampling are satisfied, we apply the Discrete Fourier Transform (DFT) to get its representation in the frequency domain $X[k]$ (Eq. (3)). Then, we modulate the digital frequencies by a complex exponential as in Eq. (4) to obtain $X_d[k]$, which is the frequency domain representation of the delayed digital signal $x_d[n]$ obtained applying the Inverse DFT (Eq. (5)). The effect of these operations is to produce a new sequence of samples $x_d[n]$ representing a signal that is a copy of $x[n]$ delayed by a generic value τ. In this case, τ can also be a fraction of the sampling period, which allows an arbitrary resolution when tuning the delay introduced by the obfuscator.

$$X[k] = \sum_{n=0}^{N-1} x[n] \cdot e^{-j\frac{2\pi}{N}kn},\ k = \{0, ..., N-1\} \tag{3}$$

$$X_d[k] = X[k] \cdot e^{-j\frac{2\pi}{N}k\tau} \tag{4}$$

$$x_d[n] = \frac{1}{N} \sum_{k=0}^{N-1} X_d[k] \cdot e^{j\frac{2\pi}{N}kn},\ n = \{0, ..., N-1\} \tag{5}$$

As already discussed in Sect. 6, the delay τ changes over time and can vary between a minimum value $\tau_{\min}$ and a maximum value $\tau_{\max}$ as defined in Eqs. (1) and (2). When we consider active attacks, the obfuscator is somehow reflecting the incoming

Wi-Fi signals; hence the reflected signal in Fig. 7 will always be transmitted *after* the original signal. Therefore, for active attacks emulation, we have arbitrarily chosen an offset delay $\tau_{\text{off}} = 88\,\text{ns}$ when generating the delay processes reported in Fig. 6. The applied delay τ is updated every $\Delta t = 100\,\text{ms}$—following either Eq. (1) or Eq. (2)—and it can range between $\tau_{\text{min}} = 32\,\text{ns}$ and $\tau_{\text{max}} = 144\,\text{ns}$. All these values are arbitrary, but in principle they can be tuned based on the expected performance of the emulated system. It is interesting to notice that the time delays we are considering are so tiny (tens of nanoseconds) that they do not affect the Wi-Fi MAC layer, but still have a considerable effect at changing the physical properties of the communication channel.

From a technical perspective, the implementation of the system does not change whether we are considering an active or passive attack scenario, since in both cases the obfuscator can act as a reflector following the same rules. However, to further explore the possibilities of our obfuscation mechanism, we would like to consider a different implementation that can only work when dealing with passive attack scenarios. Let us imagine that the transmitter and the obfuscator can cooperate in obfuscating the CSI; for instance, they can be two radio frontends driven by a single controller. We notice that in such case the obfuscator can even transmit the signal *before* the transmitter once the two devices are synchronized, i.e., τ can take negative values. For this reason, when discussing about the passive attack scenario in the following parts of this paper, we have chosen to set $\tau_{\text{off}} = 0\,\text{ns}$, $\tau_{\text{min}} = -56\,\text{ns}$ and $\tau_{\text{max}} = 56\,\text{ns}$ to explore the effect of this configuration.

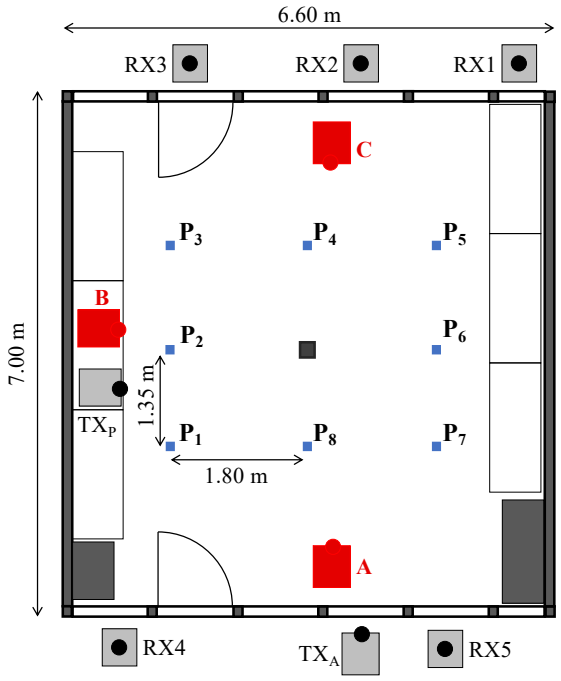

Figure 8: Plan of the lab in which localization experiments are performed. The small square dots represent the target locations of the victim. The red 'shadows' labeled A, B, C are the locations of the obfuscator in different scenarios. Five different positions for the receiver are also considered, while the transmitter is inside (TX_P) for the passive attacks and outside (TX_A) for the active attacks.

8. Scenario and Measures

We carried out the experiments in a laboratory of the ANS[3] group at the University of Brescia. The plan of the laboratory with the positions of the transmitting and receiving nodes is shown in Fig. 8. Here, we distinguish between the cases of active and passive attacks. The transmitter is outside the room when the attack is active (TX_A), and the attacker controls it. In the case of passive attacks, the transmitter is instead inside the room (TX_P), and the attacker cannot control it. All the receivers, controlled by the attacker, are placed outside the room on two opposite sides. The five different positions are useful to explore whether relative positions of the transmitter and receiver have an influence on the localization and/or the obfuscation. For instance, RX5 should have abysmal performance in localization during the active attack, as it is too close to the transmitter (TX_A) for the EM environment inside the room to have a meaningful effect on the received CSI.

We assume that the victim is standing inside the room in one of eight possible spots, indicated by small blue squares in Fig. 8 and enumerated from 1 to 8. The square at the center of the room indicates a metallic pole with an electrical cabinet; thus, it is a position where a person cannot stand, and it introduces additional EM complexity beyond walls, cabinets, tables, and chairs (not shown in the figure).

As indicated in Fig. 8, we consider three different positions of the obfuscator, which have a different relative positioning w.r.t. the transmitter depending on the attack scenario. If the attack is active and the transmitter is TX_A, then the obfuscator is: A) in front of the transmitter; B) at a 45° angle from TX_A and the receiver RX2; and C) in front of RX2, on the line of sight with TX_A. If the attack is passive and the transmitter is TX_P, then the obfuscator is: A) at a 45° angle from TX_P and on the line of sight with RX5; B) right on the side of the transmitter TX_P; and C) in front of RX2, at 45° degrees from TX_P.

The three positions of the obfuscator and the five positions of the receivers cover many different configurations, both for the case of active and passive attacks. At first glance, the position of RX4 and RX5 can seem "weird" in case of active attacks, and one may think that with the receiver outside the room on the same side the localization system cannot work, but this is not always the case as results will show. Overall the setup consists of 15 possible configurations (three positions of the obfuscator times five positions of the receiver) for each attack type (active and passive). For each configuration, we use 5 different random delay patterns (one random swipe and four Markov processes). Considering also the measures when the obfuscator is off, overall the experiments consist of $(15 + 1) \times 5 \times 2 = 160$ different configurations, giving a good "coverage" of different layouts

[3]The Advanced Networking Systems (ANS) group is a research groups in telecommunications at the Department of Information Engineering of the University of Brescia

and scenarios.

We collect a few thousand samples (i.e., CSI data points associated with one 802.11 frame) for each target position and experiment; the exact number of samples can vary depending on several other parameters, but this is irrelevant. Overall, we use 8000 samples for the training phase and 8000 samples for the testing phase, i.e., 2000 samples for each target position. Training and testing samples are collected from two different experiments with the same setup separated by several minutes to make the setup more realistic. Given the impossibility of leaving the experiments mounted overnight, we cannot say if training on one day and testing on another one gives good localization results or not, but we deem that obfuscation will always work.

To assess the validity of the proposed CSI randomization technique, we compare the classification accuracy of the localization system when the obfuscator is on (in the three different positions) and when it is turned off for all five receivers. Training and testing are always performed with the same obfuscation setup, which is the most favorable case for the attacker and the most challenging one for the obfuscator.

8.1. Scenario Validity and Extension

In this paper we consider one single scenario, inside the laboratory of our group. The reason is twofold. First, and trivial, we only have one laboratory, and since experiments are rather long and intrusive, we could not occupy offices, corridors, meeting rooms, etc. Second, the scenario we crafted is extremely favorable to the attacker (recall that our contribution is the *obfuscation*, not the localization technique that we inherit from the state of the art). In the CSI-MURDER experiments, which partially funded this work too, presented also in [26, 27], we have shown that a less sophisticated obfuscation technique still works in a completely different environment, i.e., a very large room full of reflective metallic objects in the *imec w-iLab.t* in Ghent, Belgium. Indeed, considering a completely different scenario would be an additional challenge from the attacker's point of view but does not change the main ideas at the core of the proposed obfuscation technique.

Extending this work to other scenarios and topologies is conceptually trivial, but it does not add insight in the problem. We think that a large-scale measurement campaign may be due if and when localization obfuscation techniques will have found a sound theoretical framework, and hopefully found their way into standards to protect people while offering them novel services.

9. Experimental Results

The amount of data collected prevents the presentation of all the results in detail, so we try to highlight the properties of the obfuscation system, selecting the more meaningful ones.[4]

[4]All the data we collected, the software of the implementation, and all the results are available through our web site at
`https://ans.unibs.it/projects/csi-murder/`
If the paper gets accepted for publication, we will make the data available as Open Data for future research and include some of the results omitted here (which are repetitive, thus inappropriate for a scientific paper) in an extended Technical Report.

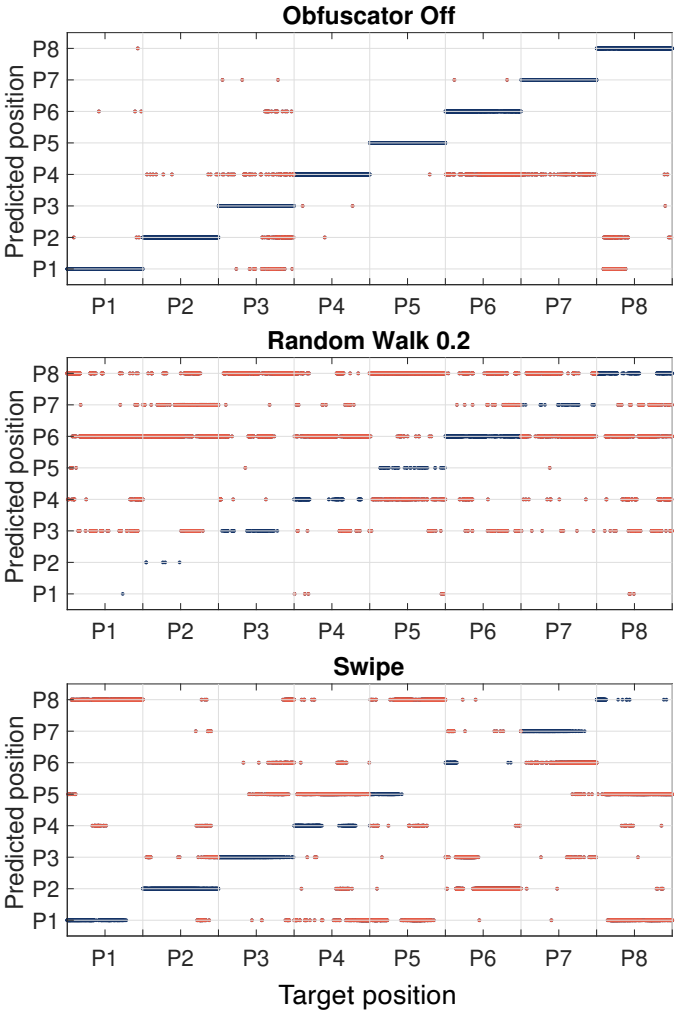

Figure 9: Active attack, RX1, Obfuscator in C). Correct and misclassified location estimates, indicated with blue and red dots respectively. Top: obfuscator off; Middle: Markov process with $\delta = 0.2$; Bottom: Swipe.

We split the results between the obfuscation performance and the communication performance in the following two subsections.

9.1. Obfuscation Performance

We start the analysis with detailed results, as they allow us to appreciate better and understand the aggregated results presented later. Figs. 9 and 10 present a set of plots that show the accuracy of the localization classification in several different situations. Fig. 9 refers to an active attack, while Fig. 10 to a passive one. Both figures refer to RX1 and the obfuscator in position C) (see Fig. 8). Results for other receivers and obfuscator positions confirm the same discussion and conclusions we draw here. As discussed in Sect. 7, when both the transmitter and the reflector are controlled by the same organization and not by the attacker (passive attack scenario), our prototype implementation allows centering the additional delay around zero. This means that in some cases, the "reflected" copy is transmitted before the original one. Indeed, given the setup based on two identical

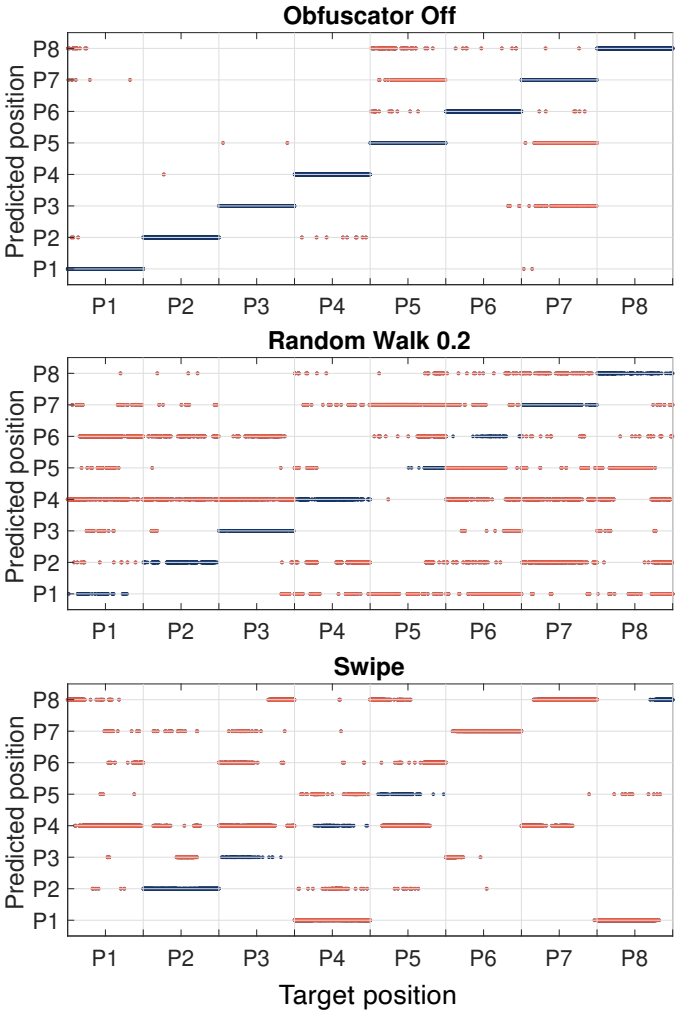

Figure 10: Passive attack, RX1, Obfuscator in C). Correct and misclassified location estimates, indicated with blue and red dots respectively. Top: obfuscator off; Middle: Markov process with $\delta = 0.2$; Bottom: Swipe.

devices, it is impossible to state which one is acting as the transmitter and which one as the obfuscator, as it would really be in a EM Smart Space.

The x-axis identifies the person's true position and the y-axis the position predicted by the localization system; thus, the plots report on the diagonal, as blue dots, the correct estimates, while all the red dots outside the diagonal are misclassifications. For each position, we plot one thousand dots, i.e., position estimates. These plots represent the same information of a confusion matrix: they miss the precision of numbers, but we argue that they are easier to appreciate at a glance. The top plot refers to the benchmark of the localization without obfuscation, and it is clear that localization is very effective, albeit not perfect: the precision is, on average, above 90%.

The middle and bottom plots report the outcome when the obfuscator is active, with the Markov process and with the swipe-based delay, respectively. Obfuscation is effective in both cases, but the Markov process generally ensures a better dispersion of the misclassified position. The dispersion is not uniform, i.e.,

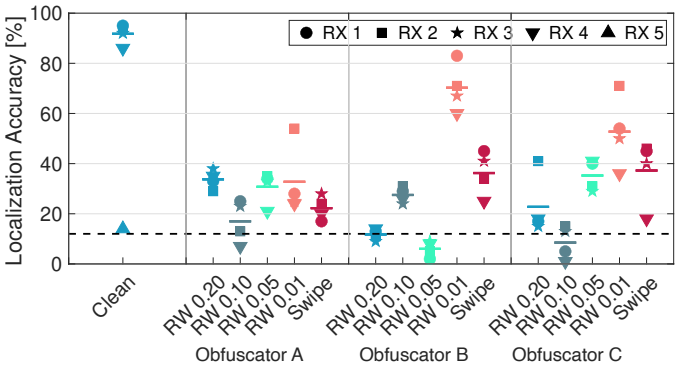

Figure 11: Active attack. Accuracy of the localization system for all the experiments. For each experiments the colored markers are the accuracy at each receiver, while the horizontal line is the average. The dashed horizontal line at 12.5% marks the random guess over 8 positions. RX5 is excluded from these results because of its peculiar position, the relative marker (blue triangle) for the clean experiment shows that even in this case the localization outcome is just a random guess.

when the localization system makes a mistake, the distribution of the mistakes is not evenly distributed on all possible positions, as a perfect obfuscator should do. The reasons are many, but they can all be reconducted to the complexity of the scenario and the behavior of the CNN. In the end, the CNN makes a decision based on the similarity of the CSI received, and since the measures are taken in a relatively short time of a few tens of seconds per position, the randomness introduced has enough correlation to make the CNN decide for some positions with a higher probability. It must be noted that introducing a totally uncorrelated random delay is inappropriate because a "white noise" can be filtered out given enough observations. One last remark regards the simple swipe technique. Apparently, the technique works reasonably well, but the cases of positions P1, P6, and P7 in Fig. 10 highlight a specific behavior that can probably be used to overcome the effects of the obfuscation in the long run. The errors in these cases are almost deterministic. This is most probably due to a coincidence in the added random delay between training and testing: during the training phase for P4, P7, and P8, the added delay of the reflector is very similar to the one added in P1, P6, and P7 during testing. This is a possibility that with the swipe technique is relatively probable, while with the Markov process is almost impossible.

Having analyzed the detailed behavior of the localization system and the obfuscation counter-measures, we can now analyze the overall results. Figs. 11 and 12 report the aggregated results of all the experiments for the active and passive attacks, respectively. First, let us comment on the localization performance when the obfuscator is off, named *clean* in the figures. The average accuracy is around 90%, with minor variations depending on the position of the receiver. Accuracy may vary from one position to another, but the analysis of all the detailed results, as done for a single case in Figs. 9 and 10, indicate that these variations are casual from one experiment to another, most probably due to how the CNN interprets the CSI during training. Indeed, CSI-based localization has never been proven able of

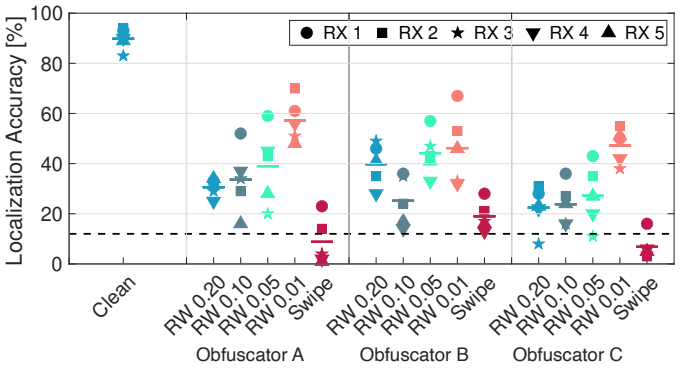

Figure 12: Passive attack. Accuracy of the localization system for all the experiments. For each experiments the colored markers are the accuracy at each receiver, while the horizontal line is the average. The dashed horizontal line at 12.5% marks the random guess over 8 positions.

higher accuracy in complex scenarios; thus, these results are in line with the literature. One note is due for RX5 performance in active attacks. As we expected as a sanity check, in this case, the localization system cannot correctly localize the person due to the "overwhelming" direct path between a transmitter and a receiver so close. The upward triangle for the clean experiment in Fig. 9 shows in this case the localization outcome is just a random guess, so we excluded RX5 for all the active attack experiments.

When the obfuscator is turned on, the localization accuracy drops both for active and passive attacks, which is very interesting as it shows that a smart space with intelligent reflective devices can successfully preserve privacy in the face of different attacks. The performance of the obfuscator is still not ideal, as for very few experiments the output of the localization system is comparable to a random guess. However, for nearly all experiments the average accuracy is below 40%, a value that makes the system hardly usable for any attacker. In both figures, the system based on the Markov process is called RW, meaning that the delay selection process is a Random Walk. Interestingly, even if it is less efficient than the other cases, setting $\delta = 0.01$ (RW 0.01) does significantly disturb the localization system, even if the delay added by the reflector is hardly detectable (cfr. Fig. 6).

From all the results we have it is difficult to draw a general conclusion on what is the best possible location for the obfuscator, or in turn, what is the best place for the attacker to place its nodes knowing where the reflector is, and this is valid both for the active and for the passive attacks. We notice that in the case of a passive attack, the receiver's position seems to have a higher impact on the misclassification (the markers are more spread). However, there is no clear correlation with the actual positions of the receiver and the obfuscator, as the ranking between the receivers change from one configuration to the other.

The swipe method, disregarding its other drawbacks, seems to work particularly well for the passive attack. We think this is indeed accidental, and further experiments would highlight that any one of the many configurations seems to work particularly well, just to change the ranking if a new configuration is tested.

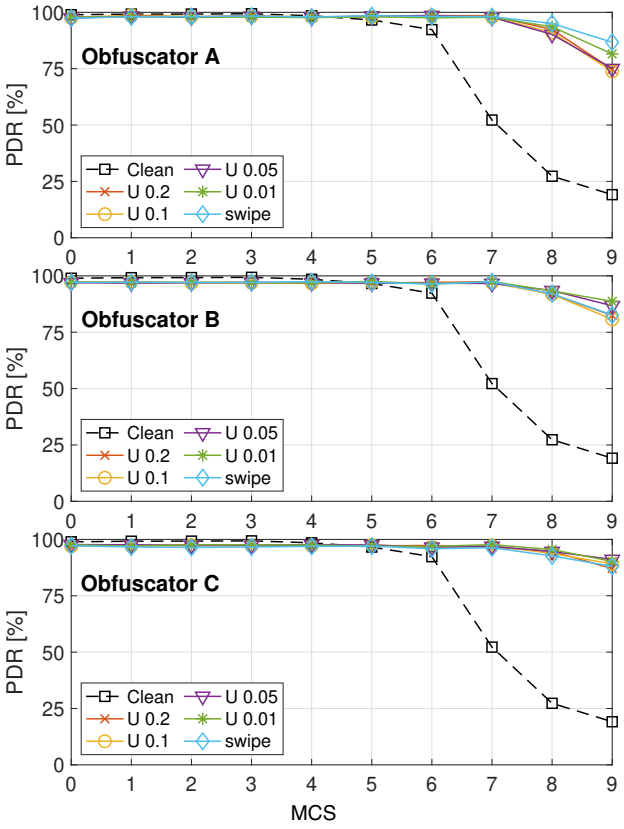

Figure 13: Active attack (transmitter TX_A outside the room). Average packet delivery ratio computed over all the receivers and all the positions of the obfuscator as a function of the Modulation and Coding Scheme (MCS) for all the obfuscation techniques tested and without the obfuscator.

The fundamental observation remains that a simple reflector can protect people's privacy against unauthorized surveillance, opening one more application for EM Smart Spaces, and granting that Wi-Fi-like systems can be safely used also in the future, even with pervasive EM fingerprinting capabilities.

9.2. *Impact on Throughput*

Protecting users' privacy is useless if the service gets destroyed; therefore, we have run experiments to verify that the obfuscating node is not harming the communication throughput between the transmitter and the receiver. To this end, we send 1000 frames from the transmitter and monitor how many of them we correctly decode at the receiver so that we can compute the Packet Delivery Rate (PDR) for different scenarios as reported in Figs. 13 and 14 for the transmitter in position TX_A and TX_P respectively. Notice that now we are now not much concerned with the localization attack, as the interest is in legitimate frames, and we use the same positions of the transmitter and the receivers just for convenience. Moreover, even if the receiver has four receiving chains, we collect these measures with a single antenna like in a Single-Input Single-Output (SISO) system, just like we do for the CSI-based localization.

The PDR is expected to decrease with the MCS order because frames with a higher MCS yield a high throughput, but they are

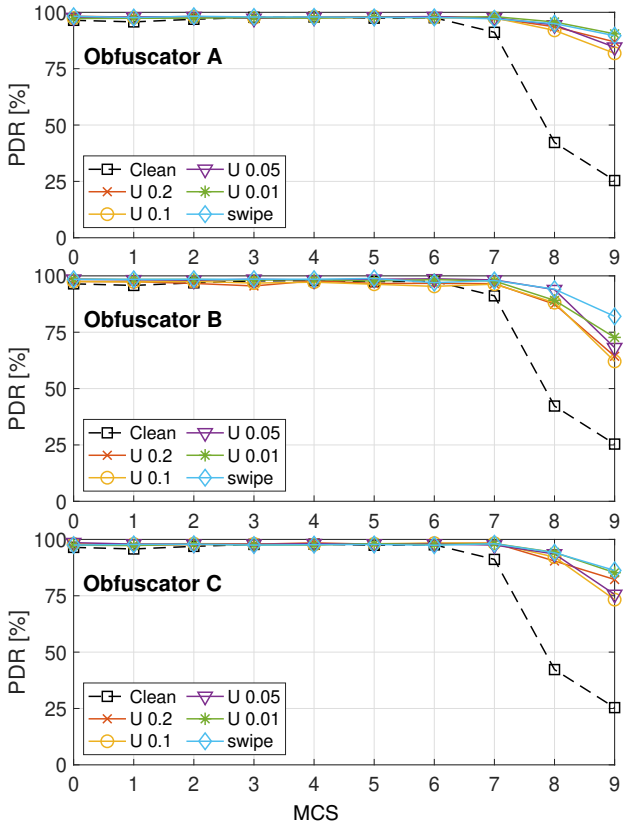

Figure 14: Passive attack (transmitter TX_P inside the room). Average packet delivery ratio computed over all the receivers and all the positions of the obfuscator as a function of the MCS for all the obfuscation techniques tested and without the obfuscator.

more sensitive to noise and interference, especially in a complex environment such as our lab. This expected behavior is entirely confirmed for the clean case (obfuscator off), with the PDR dropping sharply for MCS larger than 6 in Fig. 13 and larger than 7 in Fig. 14. The difference between the two cases can be explained with the additional wall that frames received by RX1, RX2, and RX3 have to traverse. Indeed, for these cases, the PDR is much lower for these receivers than for RX4 and RX5. In general, high order MCS are known to be fragile, meaning that, if there is no clear line-of-sight between transmitter and receiver, then the channel distortions, not completely compensated for by the equalizer, lead to systematic decoding errors, that cannot be corrected by the forward error correction code. Recall that MCS = 8, 9 use a 256QAM modulation, clearly extremely sensitive to distortions, while MCS = 7 uses a 64QAM modulation, but with a 5/6 code ratio, which is not capable of compensating high error rates.

However, the most intriguing result is that the PDR improves when the obfuscator is active for all the transmitter and receiver positions and all the obfuscation techniques. The rationale is that the obfuscator is working as a relay for the frame sent by the transmitter, creating a dominant second path and increasing the overall signal strength; hence, improving the quality of the link between transmitter and receiver. Actually, in our prototype, the signal strength is doubled as each SDR device transmits a copy of the frame. This observation is valid in general for any scenario and layout: injecting two copies of the same frame separated by a sub-symbol delay, doubles (in a stochastic sense, as the instantaneous power depends on reflections) the signal power at the receiver independently from the specific propagation environment, hence improves performance.

It can be exciting to explore the performance when the power of the reflector is different or when there are more reflectors. This observation is extremely interesting because again it hints at creating privacy-preserving Smart Spaces with extremely high communication performance, possibly using more than one reflector. A detailed analysis of the potential communication gains maintaining the overall power injected in the channel, but exploiting intelligent reflective surfaces that protect privacy is very interesting, but goes well beyond the scope of this paper.

10. Discussion and Conclusions

Environment sensing attacks exploiting 802.11 BSSs have been proven feasible by recent works and represent a severe threat to users' privacy, exposing the presence of people in a room and even their precise position within it.

In this work, we have shown that it is possible to counter CSI-based localization with an active device that acts as a relay and forwards the received frames with a random delay instead of jamming malicious signals and killing communications. The device produces continuous variations of the electromagnetic environment (mimicking the movement of a person in a room) to obfuscate the CSI and prevent unauthorized localization.

This work considers the CSI from a single antenna for localization purposes. We believe that using MIMO systems can drastically improve the localization accuracy since a set of other measures (e.g., angle of arrival or time of flight) can concur to give a better location estimate. However, it is not easy to imagine how these techniques can be used for the localization of a person that does not carry any device. It would be interesting to explore whether our obfuscation technique also works in MIMO communications with the localization algorithm optimized for such systems. Indeed, this can open new research paths and foster novel ideas on using EM Smart Spaces to enhance communication performance and peoples' privacy. Maybe 6G networks, or beyond, will go in this direction.

To conclude, we highlight that the random delay introduced by the obfuscator and the distance between the transmitter and the obfuscator itself make our system very different from an "extended" MIMO platform, even when the transmitter and the obfuscator work together as in the passive attack case we analyzed. In a MIMO system, indeed, the phase-delay between closely spaced antenna elements relates only to the carrier phase; in our system, the delay—albeit of the order of ns and even smaller—is a time delay of the entire signal, which is repeated by a device several meters away from the transmitter.

Acknowledgement

This work has been partially funded at the University of Brescia by the European Commission under the Horizon 2020 Or-

chestration and Reconfiguration Control Architecture – ORCA project (grant no. 732174) Open Call 3 “Experimental analysis of CSI based anti-sensing techniques – CSI-MURDER” experiment.

References

[1] E. Khorov, I. Levitsky, I. F. Akyildiz, Current status and directions of ieee 802.11be, the future wi-fi 7, IEEE Access 8 (2020) 88664–88688.

[2] J. Deng, S. Medjkouh, N. Malm, O. Tirkkonen, C. Studer, Multipoint Channel Charting for Wireless Networks, in: 2018 52nd Asilomar Conf. on Signals, Systems, and Computers, 2018, pp. 286–290.

[3] P. Ferrand, A. Decurninge, L. G. Ordoñez, M. Guillaud, Triplet-based wireless channel charting, in: IEEE Global Communications Conf. (GLOBECOM), 2020, pp. 1–6.

[4] K. Chetty, G. Smith, K. Woodbridge, Through-the-Wall Sensing of Personnel Using Passive Bistatic WiFi Radar at Standoff Distances, IEEE Trans. on Geoscience and Remote Sensing 50 (4) (2012) 1218–1226.

[5] F. Adib, D. Katabi, See through walls with WiFi!, in: Conf. of the Special Interest Group on Data Communication (SIGCOMM), ACM, Hong Kong, Aug. 2013, 2013, pp. 75–86.

[6] Z. Yang, Z. Zhou, Y. Liu, From RSSI to CSI: Indoor Localization via Channel Response, ACM Comput. Surv. 46 (2) (2013) 25:1–25:32.

[7] K. Wu, J. Xiao, Y. Yi, D. Chen, X. Luo, L. Ni, CSI-Based Indoor Localization, Trans. Parallel Distrib. Syst. 24 (7) (2013) 1300–1309.

[8] M. Widmaier, M. Arnold, S. Dorner, S. Cammerer, S. ten Brink, Towards Practical Indoor Positioning Based on Massive MIMO Systems, in: 90th IEEE Vehicular Technology Conf. (VTC2019-Fall), 2019, pp. 1–6.

[9] S. D. Bast, A. P. Guevara, S. Pollin, CSI-based Positioning in Massive MIMO systems using Convolutional Neural Networks, in: 91st IEEE Vehicular Technology Conf. (VTC2020-Spring), 2020, pp. 1–5.

[10] S. Shi, S. Sigg, L. Chen, Y. Ji, Accurate location tracking from csi-based passive device-free probabilistic fingerprinting, IEEE Trans. on Vehicular Technology 67 (6) (2018) 5217–5230.

[11] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, H. Liu, E-Eyes: Device-Free Location-Oriented Activity Identification Using Fine-Grained WiFi Signatures, in: Proc. of the ACM 20th Int. Conf. on Mobile Computing and Networking (MobiCom’14), 2014, p. 617–628.

[12] H. Abdelnasser, M. Youssef, K. A. Harras, WiGest: A ubiquitous WiFi-based gesture recognition system, in: IEEE Conference on Computer Communications (INFOCOM), 2015, pp. 1472–1480.

[13] F. Zhang, C. Chen, B. Wang, K. J. R. Liu, WiSpeed: A Statistical Electromagnetic Approach for Device-Free Indoor Speed Estimation, IEEE Internet of Things Jou. 5 (3) (2018) 2163–2177.

[14] Y. Wang, K. Wu, L. M. Ni, WiFall: Device-Free Fall Detection by Wireless Networks, IEEE Trans. on Mobile Computing 16 (2) (2017) 581–594.

[15] G. Wang, Y. Zou, Z. Zhou, K. Wu, L. M. Ni, We Can Hear You with Wi-Fi!, IEEE Transactions on Mobile Computing 15 (11) (2016) 2907–2920.

[16] X. Wang, L. Gao, S. Mao, S. Pandey, CSI-based Fingerprinting for Indoor Localization: A Deep Learning Approach, Trans. Veh. Technol. 66 (1) (2017) 763–776.

[17] G.-S. Wu, P.-H. Tseng, A Deep Neural Network-Based Indoor Positioning Method using Channel State Information, in: Int. Conf. on Computing, Networking and Communications (ICNC), IEEE, Maui, HI, USA, Mar. 2018, 2018, pp. 290–294.

[18] T. F. Sanam, H. Godrich, An Improved CSI Based Device Free Indoor Localization Using Machine Learning Based Classification Approach, in: 26th European Signal Processing Conf. (EUSIPCO), IEEE, Rome, Italy, Sept. 2018, 2018, pp. 2390–2394.

[19] C. Cai, L. Deng, M. Zheng, S. Li, PILC: Passive Indoor Localization Based on Convolutional Neural Networks, in: Ubiquitous Positioning, Indoor Navigation and Location-Based Services (UPINLBS), Wuhan, China, 2018, pp. 1–6.

[20] E. Schmidt, D. Inupakutika, R. Mundlamuri, D. Akopian, Sdr-fi: Deep-learning-based indoor positioning via software-defined radio, IEEE Access 7 (2019) 145784–145797.

[21] M. Abbas, M. Elhamshary, H. Rizk, M. Torki, M. Youssef, WiDeep: WiFi-based Accurate and Robust Indoor Localization System using Deep Learning, in: IEEE Int. Conf. on Pervasive Computing and Communications (PerCom), 2019, pp. 1–10.

[22] A. Foliadis, M. H. C. Garcia, R. A. Stirling-Gallacher, R. S. Thomä, CSI-Based Localization with CNNs Exploiting Phase Information, in: IEEE Wireless Communications and Networking Conf. (WCNC), 2021, pp. 1–6.

[23] Z. Zhou, J. Yu, Z. Yang, W. Gong, MobiFi: Fast Deep-Learning based Localization using Mobile Wifi, in: IEEE Global Communications Conf. (GLOBECOM), 2020, pp. 1–6.

[24] Cerar, Gregor and Švigelj, Aleš and Mohorčič, Mihael and Fortuna, Carolina and Javornik, Tomaž, Improving CSI-based Massive MIMO Indoor Positioning using Convolutional Neural Network, in: Joint European Conf. on Networks and Communications 6G Summit (EuCNC/6G Summit), 2021, pp. 276–281.

[25] X. Wang, X. Wang, S. Mao, Indoor Fingerprinting With Bimodal CSI Tensors: A Deep Residual Sharing Learning Approach, IEEE Internet of Things Jou. 8 (6) (2021) 4498–4513.

[26] M. Cominelli, F. Kosterhon, F. Gringoli, R. Lo Cigno, A. Asadi, An Experimental Study of CSI Management to Preserve Location Privacy, in: 14th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH), London, UK, 2020, pp. 1–8.

[27] M. Cominelli, F. Kosterhon, F. Gringoli, R. Lo Cigno, A. Asadi, IEEE 802.11 CSI randomization to preserve location privacy: An empirical evaluation in different scenarios, Elsevier Computer Networks 191 (22) (2021) 107970.

[28] M. Cominelli, F. Gringoli, R. Lo Cigno, Passive Device-Free Multi-Point CSI Localization and Its Obfuscation with Randomized Filtering, in: 19th IEEE Mediterranean Communication and Computer Networking Conference (MedComNet), 2021, pp. 1–8.

[29] M. Cominelli, F. Gringoli, R. Lo Cigno, Non Intrusive Wi-Fi CSI Obfuscation Against Active Localization Attacks, in: 16th IFIP/IEEE Conf. on Wireless On demand Network Systems and Services (WONS 2021), 2021, pp. 87–94.

[30] F. Kosterhon, Device-Free Indoor Localization: A User-Privacy Perspective, Master’s thesis, Technische Universität Darmstadt, Secure Mobile Networking Lab, Department of Computer Science (April 2020).

[31] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, A. Arora, PhyCloak: Obfuscating Sensing from Communication Signals, in: 13th USENIX Conf. on Networked Systems Design and Implementation (NSDI’16), Santa Clara, CA, USA, 2016, p. 685–699.

[32] A. Welkie, L. Shangguan, J. Gummeson, W. Hu, K. Jamieson, Programmable Radio Environments for Smart Spaces, in: 16th ACM Workshop on Hot Topics in Networks, 2017, p. 36–42.

[33] C. Liaskos, A. Tsioliaridou, S. Nie, A. Pitsillides, S. Ioannidis, I. F. Akyildiz, On the Network-Layer Modeling and Configuration of Programmable Wireless Environments, IEEE/ACM Transactions on Networking 27 (4) (2019) 1696–1713.

[34] Abanto-Leon, Luis F. and Bäuml, Andreas and Sim, Gek Hong (Allyson) and Hollick, Matthias and Asadi, Arash, Stay Connected, Leave no Trace: Enhancing Security and Privacy in WiFi via Obfuscating Radiometric Fingerprints, Proc. ACM Meas. Anal. Comput. Syst. 4 (2020) 44:1–44:31.

[35] M. Schulz, F. Gringoli, D. Steinmetzer, M. Koch, M. Hollick, Massive Reactive Smartphone-Based Jamming Using Arbitrary Waveforms and Adaptive Power Control, in: 10th ACM Conf. on Security and Privacy in Wireless and Mobile Networks (WiSec), 2017, p. 111–121.

[36] D. Nguyen, C. Sahin, B. Shishkin, N. Kandasamy, K. R. Dandekar, A Real-Time and Protocol-Aware Reactive Jamming Framework Built on Software-Defined Radios, in: ACM Workshop on Software Radio Implementation Forum, 2014, p. 15–22.

[37] M. Di Renzo, M. Debbah, D. Phan-Huy, et al., Smart radio environments empowered by reconfigurable AI meta-surfaces: an idea whose time has come, J Wireless Com Network 129 (2019).

[38] F. Gringoli, M. Schulz, J. Link, M. Hollick, Free your CSI: A channel state information extraction platform for modern Wi-Fi chipsets, in: 13th Int. Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH ’19), ACM, Los Cabos, Mexico, Oct. 2019, 2019, pp. 21–28.