

Non Intrusive Wi-Fi CSI Obfuscation Against Active Localization Attacks

Marco Cominelli, Francesco Gringoli, Renato Lo Cigno
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 privacy in work spaces, at home, and maybe even outdoors, even if outdoors experiments proving the feasibility are still not available. 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 only 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 looks promising; they show feasibility of location obfuscation with high efficiency and excellent preservation of communication performance, paving the road for further research and improved users privacy.

I. INTRODUCTION AND BACKGROUND

The estimation of the propagation channel through the so-called Channel State Information (CSI) is 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]. In parallel with the transmission performance boost, another interesting use of CSI emerged: CSI-based localization. The attention to CSI-based localization was brought by early works [2]–[5] nearly 10 years ago, immediately proving that CSI-based localization techniques has the ability to outperform traditional Received Signal Strength Indicator (RSSI)-based techniques.

After these initial works, the topic flourished, with proposals to identify activities and gestures [6]–[8], health and medical applications [9] or even “hear” people [10], and many other works and papers that is not possible to cite here.

Recent years witnessed the explosion of Machine Learning (ML) and Artificial Intelligence (AI) methodologies applied to the topic [11]–[16], which achieve astounding results using different classification or analysis techniques, often involving Deep Learning or Reinforcement Learning.

This work has been partially funded at the University of Brescia by the European Commission under the Horizon 2020 Orchestration and Reconfiguration Control Architecture – ORCA project (grant no. 732174) Open Call 3 “Experimental analysis of CSI based anti-sensing techniques – CSI-MURDER” experiment.

What none of these works have ever discussed, is how ethical or intrusive the CSI-based localization can be. What characterize most of these techniques compared with others, e.g., based on the Time of Flight of frames, is that they can be used without consent from the localized person, which is a clear violation of privacy, not to mention the security problems that can arise with the ability of some attacker to tell if and how many people are inside a room, house, or laboratory.

The technology is particularly invasive because the attack can be both passive or active 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 normally have at home or at work; attackers do not need to control such transmitters, they only have to place a receiver somewhere in the same room for precise localization or even outside if the goal is detecting presence or approximate location. In an active attack, instead, attackers control both a transmitter and a receiver: they have more freedom, and power, in the design of the attack. In case of active attack, the attacker may decide to use a different technology for location sensing, but in this case the attack is very easy to detect on-air and prevention by simple jamming is easy, not to mention that the non-Wi-Fi devices can be easily tracked and removed.

This paper introduces CSI-based localization fundamentals, the related work, 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 *active attacks*, i.e., those where the attacker controls at least one transmitter and one receiver. In these scenario, in fact, manipulating the signals at the (legitimate) transmitter is not sufficient to prevent the attack, and manipulating the signal transmitted by the attacker’s device is clearly not possible. One may argue that an active attack is detectable, as it implies on-air traffic which is not “legitimate.” 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 really control what are the legitimate Wi-Fi users. Even in office environments it is very difficult to imagine that the victim can identify an attack, if not for else because the victim can be an employee and the attacker the employer, who wants to control his/her employees beyond what legislation permits.

II. CSI-BASED LOCALIZATION




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.

Whatever the technology used to extract location information from the CSI, this information must be present in the signal itself. It is embedded in the signal during propagation and carries pieces of information on people 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 consecutive received frames is remarkably constant in the same location, while it changes significantly when moving from one location to another. Repeating the experiment at different times shows that there is a time-based variation, but still the CSI carries enough location-specific information to allow a proper algorithm to infer the person 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 important.

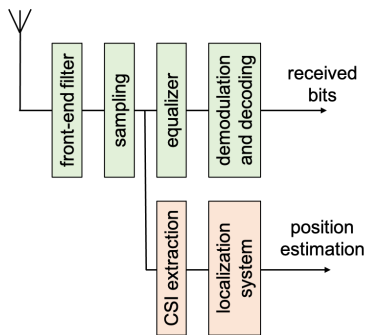


Figure 2: 802.11 receiver modified to infer people location; first the localization system is trained with a person standing in positions of interest building a reference set, during the attack the infers the position of the person classifying CSI data on the reference set.

The transmission technique has clearly a huge 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 fully understand localization techniques. Fig. 2 sketches the diagram of a single antenna receiver modified to retrieve information on people's localization. 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. The duplicated samples, instead, enter the localization system that, exploiting the same CSI used by the equalizer, yields an estimation of the person's location. The CSI is 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 received more attention recently are based on Neural Networks (NNs) that are trained with someone standing in a known position and then, during the attack, estimate 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 key design decision regards the transmissions. The localization system can be *passive*, i.e., it exploits the data packets normally sent by users, or it can be *active*, i.e., it uses frames that are sent by a device specifically to perform the localization.

A passive attack is easier to mount 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 otherwise the change in CSI determined by the transmitter movement completely confuses the classification at the receiver. In most cases this is not a problem, since APs are fixed and they also generate the largest amount of traffic. On the other hand, as we have shown in [17], it is possible to obfuscate the information on localization carried by the CSI by properly manipulating the transmitted frames, so that protection against passive localization is achievable.

An active localization attack, instead, uses both a transmitter and a receiver, so it is somewhat more complex, it is detectable because there are frames on-air that do not belong to a legitimate Basic Service Set (BSS), but there is no way to confuse the localization technique by manipulating the transmitted frames, as the transmitter is not controlled by legitimate users, but by the attacker. In this work, we concentrate on 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 to have more than one receiver that work coordinately to improve the localization accuracy.

The localization technique adopted in this work relies on a Convolutional Neural Network (CNN) to perform the classification task. The design of the CNN is inspired by the work in [14] and refined in [18]. Based on these two works,

within the CSI-MURDER project¹ we have developed an efficient implementation that has good localization efficiency and properties as we discussed in [17]. A high-level representation of its architecture is shown in Fig. 3.

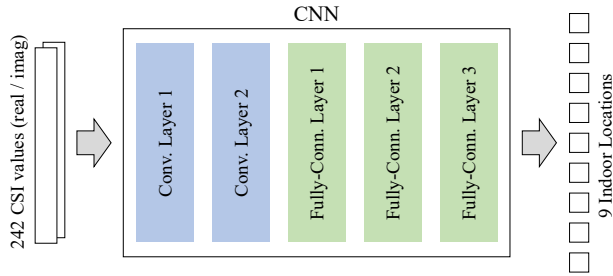


Figure 3: Architecture of the CNN used by our localization system.

One CSI data is extracted from each 802.11 frame correctly decoded at the receiver, each one being an array of complex values (the IQ samples) computed at the receiver to estimate the frequency response of the channel. We consider 802.11ac frames transmitted on 80 MHz channels in the 5 GHz band; therefore, each raw CSI point consists of 256 complex values. During the preprocessing phase, we remove data relative to all the subcarriers that are suppressed by the modulation for communication purposes, since they do not carry any information. The input of the CNN is thus a 242×2 matrix. The first two convolutional layers of the CNN shown in Fig. 3 are used to 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 number of classes can be changed as it is needed by the localization task without modifying the other layers of the CNN, its range and scope is rather flexible, from simple binary decisions (e.g., right hand side or left hand side of the room), to fine classification of arbitrary positions, not necessarily in a regular grid, to estimate (x, y) coordinates in a Cartesian space provided that the training grid is sufficiently dense, as we did in [17] where we explored localization obfuscation against passive attacks. All the layers but the last (which uses a softmax function) use a common 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.

III. RELATED WORK

A review of CSI based localization techniques is out of the scope of this paper; we overview only works whose target is localization obfuscation or localization privacy protection.

Our previous work [17] deals with localization obfuscation against passive attacks. Albeit the scope of the two works is the same, the methodology is rather different. In a passive attack, localization is based on the CSI of frames transmitted by APs, and countermeasures can be based on the manipulation

of frames at the transmitter: no additional devices are needed and the paper presents a simple proof-of-concept showing that proper manipulation at the transmitter can obfuscate the actual position of a person.

A countermeasure against Wi-Fi sensing attacks has been implemented in [19] to prevent gesture recognition; similarly to our proposal, this system relies on an additional component acting as a relay that must be placed in the environment, and 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 superimpose (sum) a different signal (possibly noise) with the goal of making the frame useless, thus also killing communication capabilities. This said the techniques proposed are different and [19] is focused on gesture recognition rather than localization.

A closely related work published recently [20] manipulates the CSI with the goal of avoiding device radiometric fingerprinting and helping in the prevention of impersonation attacks. The goal of the paper is not localization; however the technique used are similar to those we use in this work, and clearly, in case a person also holds a Wi-Fi device the double attack identifying the device and the location of the person is a possibility.

Finally, exploiting a reactive jamming device that selectively kills frames that belong to the localization attack can be conceived and techniques like [21], [22] can be adapted to the scope. Actually, 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 the approach requires to know that a localization attack is under way, and the jamming device must recognize the illegitimate traffic and try to kill those frames only, while our approach is transparent, as it does not affect the reception of frames, so that the obfuscating device can be active at any time on any frame.

IV. ACTIVE CSI BASED ATTACKS AND OBFUSCATION

An active localization attack is based on the control, by the attacker, of one transmitter and at least one receiver. The attacker has thus full control of the transmission chain and the only way to interfere with his/her intrusion is by actively mingling transmitted frames on the channel.

A. Attack Model

Fig. 4 depicts the attack model. A person is standing in a room, can be an office or home or anywhere, and the attacker aims at collecting information on the person 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. II and has the ability to access the room at some time to train the CNN. After training the CNN the attacker can simply configure the system to send frames periodically and collect the estimation of the person position when she/he is in the room. The attacker can possibly use more than one receiver, presumably improving the localization performance by correlating the position estimation

¹Further details on this project, the software produced and so forth can be found at <https://ans.unibs.it/projects/csi-murder/>.




Figure 4: In active localization, the attacker controls both the transmitter and at least one receiver; the location privacy of the victim can be preserved only if an active device is able to obfuscate the CSI of frames.

by all the receivers. This possibility is however beyond the scope of this paper and it is left for future work.

B. Active Obfuscation Requirements

To protect people privacy in presence of an active attack, which can be difficult to detect as we discussed at the beginning of the paper, the only possibility is to install an active device that randomly changes the channel response “reflecting” the incoming signal with a properly 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 it is impossible for the localization system to identify the position of the person based on the CSI fingerprint, because this latter contains too much random information to allow identifying the features that allow the position classification. 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.

Similarly to an attacker controlling more than one receiver, also the reflectors can be more than one, possibly enhancing the obfuscation performance; however, this analysis is left for future research.

V. ACTIVE CSI RANDOMIZATION

In a real anti-localization system the obfuscator can be a repeater that mimics a reflective surface, or, in a bit more futuristic scenario it can be a Reflective Intelligent Surface [23] changing its properties under the control of a proper obfuscation function. The goal of the obfuscator is to add one more “reflecting path” into the propagation environment, and to manage this artificial path to confuse the localization algorithm. At the same time the channel distortion must remain plausible, meaning that it should allow the equalizer at the receiver




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.

to correctly compensate the distortion so that frames can be received without reducing the communication performance.

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 after 10s 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 on. It is clear that the obfuscator actually alters significantly the propagation environment (the blue lines are very different in the two plots) and after 10s the red lines tell a different propagation story, and we can conjecture that any localization technique will have a hard time in fingerprinting and classifying positions.

Ideally, the outcome of the obfuscation at the receiver should be indistinguishable from a standard channel response, both in the distribution of attenuation and phase jumps, and in the time correlation. It is still not clear if this is fully achievable, also because there is a lack of experimental studies that properly characterize the stochastic properties of the channel response.

VI. IMPLEMENTATION

The implementation of the obfuscator/reflector in hardware is unfortunately beyond the possibilities of our lab; let alone realizing a controllable reflective intelligent surface. Moreover, such an expensive endeavor is justified only if the proposed technique works and is tamper proof; thus to realize our proof-of-concept implementation we resort to Software Defined Radio (SDR) devices and a little ‘trick.’

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 receiver. The SDR transmitter keeps sending 802.11 frames generated using the Matlab WLAN Toolbox at a constant rate of one frame every 10 ms. The receiver includes a Broadcom chipset from which we extract the CSI data points using the tools provided by the Nexmon project [24]. The localization system is implemented off-line working on the memorized CSI data points, since there is in general no real-time requirement in the identification of a person’s position. In any case the analysis of the CSI by the CNN of the localization system is extremely fast: according to our tests, an Intel Core i7 clocked at 4.4GHz takes as only as




Figure 6: Schematic representation of the experimental setup; the obfuscator acts as an active, configurable reflector.

60 μ s for processing the CSI data extracted from every single frame.

Implementing a real-time 802.11 signal relay in software is hard: the latency introduced by typical SDR systems cannot meet the strict timing requirements and in general some kind of hardware-accelerated processing is necessary. In case of hardware implementation, the delay, which is in the order of nano seconds to tens of nano seconds, is easily achievable working in the analog domain. For this reason, we resort to a gimmick in order to implement our system that is shown in Fig. 6. The two SDRs—one playing as the attacker’s transmitter and the other one as the obfuscator—are synchronized by means a common clock source. 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 (Pulse Per Second (PPS)) that together allow almost perfect synchronization. Once the two SDRs are synchronized, the obfuscator can emulate the effects of a reflected path in different ways.

The easiest way to emulate a reflected path is to apply a delay to the signal transmitted by the obfuscator, i.e., to shift the sequence of IQ samples transmitted by the obfuscator by a certain amount of samples. Despite being simple, this solution has a substantial limitation: samples are transmitted by the radio at a fixed rate; therefore, the time between two consecutive samples is completely determined by the available bandwidth. In our implementation, the transmission rate of the N300 SDRs is 125 MSample/s, which corresponds to a sampling period of 8 ns. This means that the minimum delay that corresponds 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 that are multiple of this quantity, which would be an inconvenient limitation in the implementation of our proof-of-concept.

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. (1)). Then, we modulate the digital frequencies by a complex exponential as in Eq. (2) 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. (3)). 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 quantity Δt , expressed in number of samples. In this case Δt can also be a fraction, which allows an arbitrary resolution when tuning the delay introduced by the obfuscator.

$$X[k] = \sum_{n=0}^{N-1} x[n] \cdot e^{-j \frac{2\pi}{N} kn}, k = \{0, \dots, N-1\} \quad (1)$$

$$X_d[k] = X[k] \cdot e^{-j \frac{2\pi}{N} k \Delta t} \quad (2)$$

$$x_d[n] = \frac{1}{N} \sum_{k=0}^{N-1} X_d[k] \cdot e^{j \frac{2\pi}{N} kn}, n = \{0, \dots, N-1\} \quad (3)$$

In our implementation, the delay Δt introduced by the obfuscator can change between a minimum value t_0 (which is manually set depending on the relative position of transmitter and obfuscator) and a maximum value t_1 that is set to $t_0 + 120$ ns. This arbitrary value corresponds to a path of the signal that is approximately 36 m longer than the shortest possible one generated by the obfuscator.

Early tests showed that it is possible to obfuscate the position of the victim by letting the artificial delay Δt vary uniformly within the given range. However, our goal is to obfuscate the position of the person and a clever attacker might identify the obfuscation if the CSI is too random. For this reason, we devise a simple algorithm that changes the delay by a random quantity, but with a Markovian correlation on the previous frame, so that to a simple statistical analysis it would look like a moving person. Obviously, the obfuscator must retain the delay applied previously when applying a new delay. In our implementation, we update the value of the delay Δt every 100 ms according to Eq. (4). The quantity $U_{[0,1]}$ is a uniform random variable taking values between 0 and 1. This specific value of the increment makes the total length of the multipath component introduced by the obfuscator change with an average speed of 1.2 m/s. The sign of the increment becomes negative when Δt reaches the maximum value t_1 and returns positive when Δt falls below t_0 . In this way, Δt continuously swings between t_0 and t_1 with an average round-trip time of 30 s.

$$\Delta t_{new} = \Delta t_{old} \pm 0.1 \cdot U_{[0,1]} \quad (4)$$

This evolution mimics, with some simplifications, what would happen if the reflection is due to a person moving in the room with a random waypoint mobility model and a strolling speed around 0.6 m/s. Clearly more sophisticated patterns can be used, but we kept it simple for the sake of results interpretation.²

VII. SCENARIO AND MEASURES

Our experiments are carried out in a laboratory of the ANS³ group at the University of Brescia. A plan of the laboratory

²Observing the evolution in time of the channel response (at least of its complex envelope amplitude) when the reflector is on and when it is off is extremely insightful. Unfortunately, it is impossible to include a movie in a paper; the interested reader can find some examples of the evolution in time of the channel response on the CSI-MURDER web site (<https://ans.unibs.it/projects/csi-reflector/>).

³The Advanced Networking Systems group is a research group in telecommunication at the University of Brescia




Figure 7: 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.

with the positions of the transmitting and receiving nodes is shown in Fig. 7. The transmitter and all the receivers, controlled by the attacker, are placed outside of the room on two opposite sides. We assume that the victim is standing inside the room in one of the nine possible spots P_1, \dots, P_9 in Fig. 7. Position 9, at the center of the room, is very close to a metallic pole with an electrical cabinet, thus it is a position where the presence of a person may not induce significant variations to the environment from an electromagnetic point of view.

As indicated in Fig. 7, we consider three different positions of the obfuscator: A) in front of the transmitter (TX); B) at a 45° angle from TX and the receiver (RX2); and C) in front of RX2. The three positions of the obfuscator with respect to the other receivers (RX1,3,4,5) assume many different configurations. Indeed, the position of RX4 and RX5 can seem “weird” 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. Overall the setup consists of 15 possible configurations, giving a good “coverage” of different layouts and scenarios, considering also experiments when the obfuscator is off, it makes 20 different experiments.

For each target position and experiment we collect 800 samples (i.e., CSI data points associated to one 802.11 frame) for the training phase and 100 samples for the testing phase. Training and testing samples are collected from two different experiments with the same setup, but not at the same time to make experiments more realistic.

To assess the validity of the proposed CSI randomization technique, we compare the classification accuracy of the

Table I: Confusion matrix for receiver RX2 in two different scenarios. Upper numbers in every cell refer to the results obtained with the obfuscator off, while the lower numbers are obtained with the obfuscator in position A. The last column (in gray) reports the True Positive Ratio, i.e. the number of correct guesses among all the ones produced for a given position. For each target position, we are considering 100 CSI samples.

		Target Position									TPR
		P1	P2	P3	P4	P5	P6	P7	P8	P9	
Predicted Position	P1	74 0	0 0	0 22	0 0	0 0	0 0	0 0	0 0	0 0	100% 0%
	P2	0 0	100 0	0 24	0 0	0 0	0 71	0 0	0 2	0 0	100% 0%
	P3	26 14	0 0	100 0	0 0	6 0	0 0	0 0	0 30	4 82	74% 0%
	P4	0 0	0 5	0 0	100 0	0 1	0 0	0 0	0 1	82 3	55% 0%
	P5	0 0	0 0	0 0	0 94	94 99	0 0	0 0	0 0	0 0	100% 51%
	P6	0 0	0 1	0 0	0 0	0 0	100 0	0 0	0 0	0 0	100% 0%
	P7	0 86	0 94	0 0	0 0	0 0	0 0	100 100	0 67	0 15	100% 28%
	P8	0 0	0 0	0 54	0 6	0 0	0 29	0 0	100 0	14 0	88% 0%
	P9	0 0	0 0	0 0	0 0	0 0	0 0	0 0	0 0	0 0	- -
Average Accuracy										85% 22%	

localization system when the obfuscator is on (in the three different positions) and when it is turned off, and this for all the 5 receivers. When we consider obfuscation, the obfuscator is on also during training, otherwise the training classification would have no meaning during testing and localization is obviously impossible.

VIII. EXPERIMENTAL RESULTS

First of all we want to assess the quality of the experimental setup, to avoid biases and results that are peculiar to some specific features of the setup itself. Tab. I reports the confusion matrix of one experiment with the obfuscator in A and RX2. The two numbers in the squares report the absolute number of classifications for the position on the rows when the actual position is the one in the columns. The upper number is with the obfuscator off, while the lower one is with the obfuscator on. The diagonal (green squares) is in practice the accuracy for every position, and the lowest square is the overall average. It is clear that the localization works well when the obfuscator is off, while it is barely better than a random guess when the obfuscator is on.

Position P9 results to be peculiar, and the localization never works in this point. We think the reason is the presence of the pole with the electrical cabinet, so in the remaining of the paper we exclude P9 from the evaluation, as it is indeed a bias due to the specific location.

A specific note is due also for P5 and P7. For these two points the accuracy of the localization system is very high even with the obfuscator on. However, we have to consider that the

Table II: Accuracy of the localization system in the considered scenarios.

Positions	Accuracy [%]																			
	RX1 w. obfuscator				RX2 w. obfuscator				RX3 w. obfuscator				RX4 w. obfuscator				RX5 w. obfuscator			
	off	A	B	C	off	A	B	C	off	A	B	C	off	A	B	C	off	A	B	C
1	100	9	0	6	75	0	0	0	93	0	11	96	0	19	0	0	17	0	0	0
2	100	0	4	0	100	0	10	0	100	41	0	0	91	8	0	0	14	0	0	0
3	100	0	20	95	100	0	62	95	100	0	100	74	98	0	0	14	12	7	0	94
4	94	0	100	0	88	0	99	0	99	2	89	0	100	0	100	0	17	1	86	0
5	97	100	100	66	94	99	100	61	66	98	98	69	100	99	100	49	17	93	91	19
6	80	3	60	93	99	0	0	38	100	2	0	0	97	0	0	0	15	0	0	3
7	89	22	0	0	100	89	1	78	79	76	29	0	99	33	0	0	13	3	0	0
8	100	0	78	100	100	0	0	94	100	5	0	80	100	0	0	84	9	0	0	91
Average	95	17	45	45	95	24	34	46	92	28	41	40	86	20	25	18	14	13	22	26

localization system is a classifier: it must assign a position to each CSI data point, and it seems that in this experiment it has a certain preference for these two locations. Looking at the last column, in fact, that reports the True Positive Rate (TPR), i.e., the ratio between correct guesses and all the guesses in a given position, it is clear that these two positions are selected much more frequently than the others, so that the TPR for these two positions is just 55% and 28% respectively, which is not that high.

We can now present the overall results, which are condensed in Tab. II, even if they refer to 20 different experiments and thousands of collected and analyzed CSI data points. The rows refer to the 8 positions we consider, while the columns are the 20 different experiments. It is clear that when the obfuscator is off, the localization system can efficiently determine the location of the target with high accuracy, above 90% on average when the receiver is on the other side of the room, and above 85% for RX4. Only RX5 is not able to properly localize the person, and this is clearly due to an “overwhelming” direct path between a transmitter and a receiver so close.

When the obfuscator is turned on, the localization accuracy drops. While some configurations seem to work better than others, in all the considered cases the accuracy of the localization system is significantly impacted by the presence of the obfuscating node and the average accuracy varies between 17% and 46%. Recall that a random guess would be 12.5% accuracy. 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. It is clear that RX4 is a less favorable position to infer the location, but it is also clear that even such position works without countermeasures, while it does not with the obfuscator on.

A specific comment is due for RX5, where we observe the surprising behavior of an increased accuracy when the obfuscator is active in positions B and C. Indeed this is not surprising at all, because in this case the reflector injects a signal that is far from the receiver, thus it can be influenced by the person position, adding information in the CSI that is not present when the transmitter is “too close” to the receiver. This

result is very interesting and shows that localization obfuscation may be trickier than expected.

A. Impact on throughput

Protecting users’ privacy is useless if in doing this the service is destroyed, thus we have run an experiment aiming at verifying that the obfuscating node is not harming the throughput of the communication between the transmitter and the receiver. To this end, we send 1000 frame from the transmitter and we 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 Tab. III. The PDR can vary in function of the MCS employed for the communication. Frames with a higher MCS are required to reach a high throughput but they are more sensitive to noise and interference, especially in a complex environment such as our lab. One striking evidence is that the PDR apparently improves when the obfuscator is active inside the room. The rationale of this is that the obfuscator is actually working as a relay for the frame sent by the transmitter creating a dominant second path, but also increasing the overall signal strength, hence improving the quality of the link between transmitter and receiver. This observation is once more extremely interesting, because it hints to the possibility of creating privacy preserving environments with extremely high communication performance, possibly using more than one reflector.

IX. CONCLUSIONS AND FUTURE WORK

Environment sensing attacks exploiting 802.11 BSSs have been proven feasible by recent works and represent a serious 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, instead of jamming malicious signals, acts as a relay and forwards the received frames with a varying delay to the sensing receiver, thus confusing the CSI information in some sense mimicking a continuous variation of the electromagnetic environment. A brief heuristic discussion on how location privacy can be achieved against an active attack where the attacker

Table III: PDR as a function of the MCS, with the obfuscating node placed in different positions, averaged over the five positions of the receiver.

MCS index	Mbit/s	Obfuscator			
		Off [%]	Pos. A [%]	Pos. B [%]	Pos. C [%]
0-BPSK	29.3	99.3	99.8	96.1	99.7
1-QPSK	58.5	99.3	99.4	99.4	99.4
2-QPSK	87.8	99.3	99.5	99.5	99.5
3-16-QAM	117.0	99.3	99.4	99.5	99.4
4-16-QAM	175.5	97.6	98.9	99.2	99.4
5-64-QAM	234.0	69.5	97.7	97.9	98.5
6-64-QAM	263.3	57.8	96.0	97.8	98.2
7-64-QAM	292.5	49.7	94.0	90.6	95.2
8-256-QAM	351.0	27.3	54.2	49.5	71.9
9-256-QAM	390.0	19.1	30.5	38.1	26.2

controls both a transmitter and a receiver is followed by the design of such a privacy protection system. The experimental results obtained with an SDR framework fully support the heuristic derivation, and provide a clear proof-of-concept of the feasibility of location privacy protection coupled with high quality communications.

This work is just an initial step. Experiments with more receivers that correlate their data are needed, a more complete theoretical analysis and understanding of the system is due, as well as the design of future environments where Reflective Intelligent Surfaces may boost communications and privacy for networks beyond 5G.

REFERENCES

- [1] E. Khorov, I. Levitsky, and I. F. Akyildiz, "Current Status and Directions of IEEE 802.11be, the Future Wi-Fi 7," *IEEE Access*, vol. 8, pp. 88 664–88 688, May 2020.
- [2] K. Chetty, G. Smith, and K. Woodbridge, "Through-the-Wall Sensing of Personnel Using Passive Bistatic WiFi Radar at Standoff Distances," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 50, no. 4, pp. 1218–1226, 2012.
- [3] F. Adib and D. Katabi, "See through walls with WiFi!" In *Conf. of the Special Interest Group on Data Communication (SIGCOMM)*, Hong Kong, Aug. 2013: ACM, Aug. 2013, pp. 75–86.
- [4] Z. Yang, Z. Zhou, and Y. Liu, "From RSSI to CSI: Indoor Localization via Channel Response," *ACM Comput. Surv.*, vol. 46, no. 2, 25:1–25:32, Dec. 2013.
- [5] K. Wu, J. Xiao, Y. Yi, D. Chen, X. Luo, and L. Ni, "CSI-Based Indoor Localization," *Trans. Parallel Distrib. Syst.*, vol. 24, no. 7, pp. 1300–1309, Jul. 2013.
- [6] Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, and 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)*, Maui, Hawaii, USA, 2014, pp. 617–628.
- [7] H. Abdelnasser, M. Youssef, and K. A. Harras, "WiGest: A ubiquitous WiFi-based gesture recognition system," in *IEEE Conference on Computer Communications (INFOCOM)*, Kowloon, Hong Kong, Apr. 2015, pp. 1472–1480.
- [8] F. Zhang, C. Chen, B. Wang, and K. J. R. Liu, "WiSpeed: A Statistical Electromagnetic Approach for Device-Free Indoor Speed Estimation," *IEEE Internet of Things Jou.*, vol. 5, no. 3, pp. 2163–2177, 2018.
- [9] Y. Wang, K. Wu, and L. M. Ni, "WiFall: Device-Free Fall Detection by Wireless Networks," *IEEE Trans. on Mobile Computing*, vol. 16, no. 2, pp. 581–594, 2017.
- [10] 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.
- [11] X. Wang, L. Gao, S. Mao, and S. Pandey, "CSI-based Fingerprinting for Indoor Localization: A Deep Learning Approach," *Trans. Veh. Technol.*, vol. 66, no. 1, pp. 763–776, Jan. 2017.
- [12] G.-S. Wu and P.-H. Tseng, "A Deep Neural Network-Based Indoor Positioning Method using Channel State Information," in *Int. Conf. on Computing, Networking and Communications (ICNC)*, Maui, HI, USA, Mar. 2018: IEEE, 2018, pp. 290–294.
- [13] T. F. Sanam and H. Godrich, "An Improved CSI Based Device Free Indoor Localization Using Machine Learning Based Classification Approach," in *26th European Signal Processing Conf. (EUSIPCO)*, Rome, Italy, Sept. 2018: IEEE, 2018, pp. 2390–2394.
- [14] C. Cai, L. Deng, M. Zheng, and S. Li, "PILC: Passive Indoor Localization Based on Convolutional Neural Networks," in *2018 Ubiquitous Positioning, Indoor Navigation and Location-Based Services (UPINLBS)*, Wuhan, China: IEEE, Mar. 2018.
- [15] E. Schmidt, D. Inupakutika, R. Mundlamuri, and D. Akopian, "SDR-Fi: Deep-Learning-Based Indoor Positioning via Software-Defined Radio," *IEEE Access*, vol. 7, pp. 145 784–145 797, Oct. 2019.
- [16] M. Abbas, M. Elhamshary, H. Rizk, M. Torki, and M. Youssef, "WiDeep: WiFi-based Accurate and Robust Indoor Localization System using Deep Learning," in *Int. Conf. on Pervasive Computing and Communications (PerCom)*, Kyoto, Japan, Mar. 2019: IEEE, 2019.
- [17] M. Cominelli, F. Kosterhon, F. Gringoli, R. Lo Cigno, and 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, Sep. 2020, pp. 1–8.
- [18] F. Kosterhon, "Device-Free Indoor Localization: A User-Privacy Perspective," M.S. thesis, Technische Universität Darmstadt, Secure Mobile Networking Lab, Department of Computer Science, April 2020.
- [19] Y. Qiao, O. Zhang, W. Zhou, K. Srinivasan, and A. Arora, "PhyCloak: Obfuscating Sensing from Communication Signals," in *13th Conf. on Networked Systems Design and Implementation*, Santa Clara, CA, USA, Mar. 2016: USENIX Association, 2016, pp. 685–699.
- [20] 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.*, vol. 4, 44:1–44:31, 3, Article 44 Dec. 2020.
- [21] M. Schulz, F. Gringoli, D. Steinmetzer, M. Koch, and 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)*, Boston, MS, USA, 2017, pp. 111–121.
- [22] D. Nguyen, C. Sahin, B. Shishkin, N. Kandasamy, and 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*, Chicago, Illinois, USA, 2014, pp. 15–22.
- [23] M. Di Renzo, M. Debbah, D. Phan-Huy, and et al., "Smart radio environments empowered by reconfigurable AI meta-surfaces: an idea whose time has come," *J Wireless Com Network*, vol. 129, 2019.
- [24] F. Gringoli, M. Schulz, J. Link, and 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)*, Los Cabos, Mexico, Oct. 2019: ACM, 2019, pp. 21–28.