An Experimental Study of CSI Management to Preserve Location Privacy

Marco Cominelli
DII, University of Brescia

Felix Kosterhon
SEEMOO, TU Darmstadt

Francesco Gringoli
DII, University of Brescia

Renato Lo Cigno
DII, University of Brescia

Arash Asadi
WISE/SEEMOO, TU Darmstadt

ABSTRACT

Passive device-free localization of a person exploiting the Channel State Information (CSI) from Wi-Fi signals is quickly becoming a reality. While this capability would enable new applications and services, it also raises concerns about citizens' privacy. In this work, we propose a carefully-crafted obfuscating technique against one of such CSI-based localization methods. In particular, we modify the transmitted I/Q samples by leveraging an irreversible randomized sequence. I/Q symbol manipulation at the transmitter distorts the *location-specific* information in the CSI while preserving communication, so that an attacker can no longer derive information on user's location. We test this technique against a Neural Network (NN)-based localization system and show that the randomization of the CSI makes undesired localization practically unfeasible. Both the localization system and the randomization CSI management are implemented in real devices. The experimental results obtained in our laboratory show that the considered localization method (first proposed in an MSc thesis) works smoothly regardless of the environment, and that adding random information to the CSI mess up the localization, thus providing the community with a system that preserve location privacy and communication performance at the same time.

CCS CONCEPTS

• **Networks** → **Network protocols**; **Link-layer protocols**; • **Security and privacy** → *Human and societal aspects of security and privacy*; **Privacy protections**;

KEYWORDS

localization, privacy, channel state information, neural networks, Wi-Fi, randomization, experiments and measures

ACM Reference Format:
Marco Cominelli, Felix Kosterhon, Francesco Gringoli, Renato Lo Cigno, and Arash Asadi. 2020. An Experimental Study of CSI Management to Preserve Location Privacy. In *14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH'20), September 21, 2020, London, United Kingdom.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3411276.3412187

1 INTRODUCTION

The theme of precise localization of devices or people has long been of great interest from both a research and an industrial perspective, particularly indoor positioning, as GPS-based systems cannot work. One specific field of indoor positioning is the localization of "bodies"—human beings, but also other physical objects in the ambient—without these bodies being fitted with an active or passive communication device. Clearly, camera-based localization or anti-intrusion systems (e.g., radar- of lidar-based) are part of this latter field. However, we focus on systems based on Wi-Fi for the contribution of this paper, as wireless communications signals are less detectable by users and such systems can take advantage of the widely deployed Wi-Fi infrastructure.

Of interest is localization based on the CSI [11, 16], whose variations can be correlated to changes in the physical environment. Pioneer in this field of research was [2], where the authors disclosed the possibility of using advanced MIMO technologies combined with signal processing typical of radar systems with just three antennas and a Software Defined Radio (SDR) module to reveal the position, and even gestures, of a person behind a wall. Clearly such a system poses huge privacy concerns, as the localization can be obtained with proper devices even outside the room or building the person is in: A person moving within the room changes the signal propagation in the environment, which is in turn reflected on the CSI. This exposes the person to the risk of being tracked without being aware or having the possibility to avoid it. A possible countermeasure against Wi-Fi sensing attacks has been implemented in [10] for preventing gesture recognition; however, the system proposed relies on an additional component acting as a relay that must be placed in the environment. Moreover, obfuscation performance strongly depends on the position of such device. To date, preventing CSI-based localization is not possible, and the only solution to stop such attacks consists in jamming or disabling all Wi-Fi communication in the vicinity, which is not desirable.

A question arises: is it possible to explicitly alter the CSI to preserve privacy without hampering communications? This paper tries to give an initial answer to this question with two related contributions: *i)* present a measurement-based study that confirms the feasibility of passive localization of people based on CSI analysis; *ii)* disclose a CSI randomization technique that prevent unauthorized localization while maintaining high communication performance.

WiNTECH'20, September 21, 2020, London, United Kingdom

ACM ISBN 978-1-4503-8082-9/20/09…$15.00
https://doi.org/10.1145/3411276.3412187

2 CSI-BASED LOCALIZATION

Localization using Wi-Fi signals has a long literature story, starting from Received Signal Strength Indicator (RSSI) fingerprinting techniques, to mixed and fusion methods, and we refer the interested reader to a recent survey [6] for an overview. We are interested in techniques based on the analysis of CSI that, starting a decade ago, have emerged as the most powerful technique for Wi-Fi-based indoor localization [17, 18]. In particular we concentrate on methodologies based on NN and Machine Learning (ML)/Deep Learning (DL) in general [1, 11, 12, 15, 16].

We exploit the CSI collector and extraction technique presented in [5, 13], which works on many different devices, like Asus 4x4 Access Points. Once extracted, the CSI is fed to a NN-based system as described in detail in the MSc Thesis [8] and summarized briefly in Section 2.2; Section 2.1 discusses the essentials of OFDM that are needed to understand both the NN design and the randomization technique discussed in Section 3.

2.1 OFDM Transmissions

We focus on a 80 MHz OFDM transmission with a single spatial stream, i.e., transmitted by a single antenna. This type of modulation is described in the Very-High-Throughput (VHT) part of the standard; the corresponding physical level is called VHT-PHY: while being an extension of the legacy 20 MHz physical level, called OFDM-PHY, it adds many features including both Single-User (SU) and Multi-User (MU) MIMO capabilities, enhanced modulation rates and up to 8 spatial streams. We report here only details useful to understand the randomisation procedure presented in Section 3 and the rational of the NN design, and we refer the interested readers to the standard [4] and classic literature as [9].

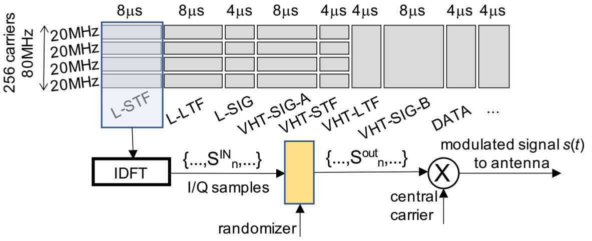

Figure 1: Format of an OFDM frame: initial symbols are known and some are used to infer the CSI at the receiver; the orange block is where randomization is introduced in Section 3.

OFDM divides the transmitted data over 256 equally spaced subchannels and keeps carriers orthogonal by construction: it builds the signal in the frequency domain mapping data to each carrier with the appropriate modulation, and then generates an OFDM symbol with the corresponding time-domain I/Q samples z_n with an Inverse Discrete Fourier Transform (IDFT), as shown in Fig. 1. This operation is repeated until the entire frame is transmitted: in the figure this corresponds to the vertical IDFT block moving left to right, producing a train of OFDM symbols $s(t)$ after multiplication by the carrier that defines the Wi-Fi channel. To help the receiver to decode the signal, the first symbols carry constant (known) content and information about the encoding process. Figure 1 reports these symbols, showing also the first three legacy symbols (L-) that can be decoded by OFDM-PHY receivers located in the corresponding 20 MHz channels that build up the 80 MHz VHT channel. They are used for setting Automatic Gain Control (AGC) and estimating time and frequency offsets (Legacy Short Training Field, L-STF); for fine frequency tuning and for collecting CSI measures of the corresponding 20 MHz channel (Legacy Long Training Field, L-LTF); and for *understanding* the duration of the remaining part of the frame (Legacy Signal, L-SIG), so that legacy receivers can set the channel as busy until then. After the L- symbols there are the VHT headers: the VHT Signal A (VHT-SIG-A) carries information required to interpret VHT data (it is still transmitted as 20 MHz symbols); the VHT-STF is used to improve AGC estimation; the VHT-LTF is used to collect the 80 MHz CSI data; and the VHT-SIG-B is only used for MU-MIMO transmission, but it is always present. After the preambles DATA symbols are transmitted.

At the receiver, operations are executed in the opposite order: I/Q samples are collected from the incoming signal and transposed in the frequency domain with a Discrete Fourier Transform (DFT). Information extracted from header symbols are used to recover the clock, reduce the carrier frequency offset (L-STF, L-LTF), and equalize the DATA symbols before decoding their content (L-LTF and VHT-LTF with the help of the extracted CSI). CSI is fundamental to enable the high throughput of modern Wi-Fi allowing fast and precise equalization, but it can also be used to infer what happened to the signal during the propagation: i.e., it is possible to identify a precise *condition* of the environment, and based on this, for instance, to derive the position of a person in a room.

2.2 A NN-based Solution

The work in [8], inspired by [3], focused on a passive, device-free system to localize a person in a room without her/him being aware of the ongoing tracking process. The system extracts features (I/Q or Amplitude and Phase) of the CSI and feed them to a supervised NN, rather than a model based approach, due to the lack of knowledge about the inner structure of the CSI and the dependencies of the different sub-carriers. In general, this approach consists of two phases: *i)* an offline stage (training) in which so-called *fingerprints* (data associated with specific locations) are collected and stored in a database, and *ii)* an online phase (operation) when the CSI samples (amplitude and phase or I/Q) are fed to the NN that returns the estimated location. Before feeding the CSI measurements into the neural network, different pre-processing steps are necessary. Among them, we identify two steps that are particularly important and worth mentioning.

Selection of Subcarriers: Not all subcarriers are used for localization. As we are focusing on VHT-PHY transmissions, we discard all pilot subcarriers as they are used by the chipset for producing the CSI itself: should we receive a multi-stream transmission, the amplitude of the pilots would be random (as reported in [5]) so we decided to discard them. We also remove subcarriers that are not used for transmission according to the 802.11ac standard.

Normalization: Using both the amplitude and the phase information, there are several issues pointed out in [15]. The lack of synchronization regarding the central frequency leads to a Carrier-Frequency Offset (CFO), while the mismatch in the time-domain introduces a Sampling-Frequency Offset (SFO) generated by the Analog-to-Digital Converter (ADC). All these errors are caused by imperfections of the hardware. After applying further processing steps as detailed in [8], the network can be trained with the CSI measurements.

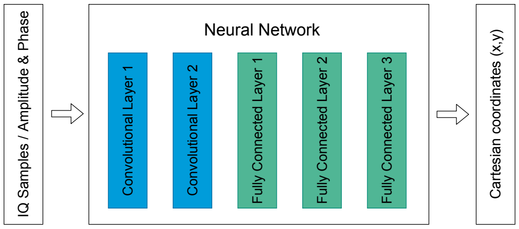

Figure 2: Architecture of the chosen NN

The network design, visualized in Fig. 2, uses two convolutional layers to extract meaningful features from the CSI data; the two layers build more complex descriptors that are essential for the localization process. Moreover, it exploits the spatial closeness of adjacent frequencies, as a convolutional layer considers multiple values at the same time. In cascade to the convolutional layers there are three fully-connected layers that combine the extracted features to obtain the Cartesian coordinates directly. The number of layers, as well as the hidden neurons for each layer, were originally obtained heuristically with experiments starting with only one layer and increasing the number of layers as well as the number of hidden neurons until the network was able to describe the relationship between the CSI and the corresponding position of the user. The loss function uses the Euclidean distance to rate the performance of the network, as it provides an intuitive understanding of the error. We choose the common Rectified Linear Unit (ReLU) as activation function and the Adaptive Momentum Estimation (ADAM) [7] algorithm to adjust the weights based on the corresponding labels.

2.3 Attacker Model

As depicted in Fig. 3 (a), an attacker wants to infer the location of a person in a room, e.g., an employee being kept under surveillance in a laboratory. We assume the presence of a common Wi-Fi AP providing Internet access in the laboratory. The attacker (e.g., the employer) has positioned a hidden Wi-Fi receiver—in our case a second AP, but in general any device capable of extracting CSI—in the laboratory and uses the NN-based localization system described before. In our specific setup, visualized in Fig. 3 (b), the receiver RX and the transmitter TX are on the opposite side of the room.

We make the following assumptions regarding the attacker model: *i)* the attacker is able to train the localization system, which only requires collecting some measurements of reference positions; and *ii)* the attacker can only access the receiver and retrieve CSI from it. It's clear that this attack can be easily replicated in hotels and multi-room environments as well as in private homes.

2.4 Learning Insight

The hardware used, the number of antennas at the transmitter and receiver and many other details may influence the precision of the localization; however, in this paper we want to focus on the elementary reasons that allow the localization and the minimal countermeasures that prevent localization. For this reason we limit the discussion to the use of one single TX-RX chain.

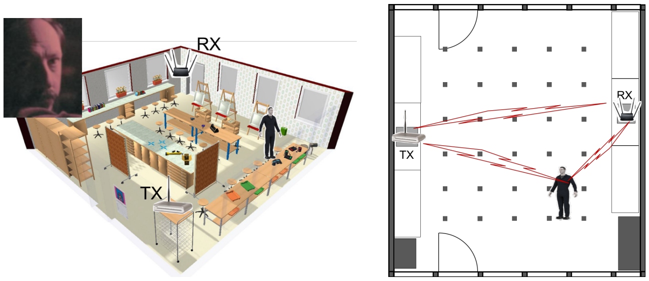

(a) Scenario **(b) Schematic setup**

Figure 3: The location attack scenario (a) mapped to our laboratory setup (b).

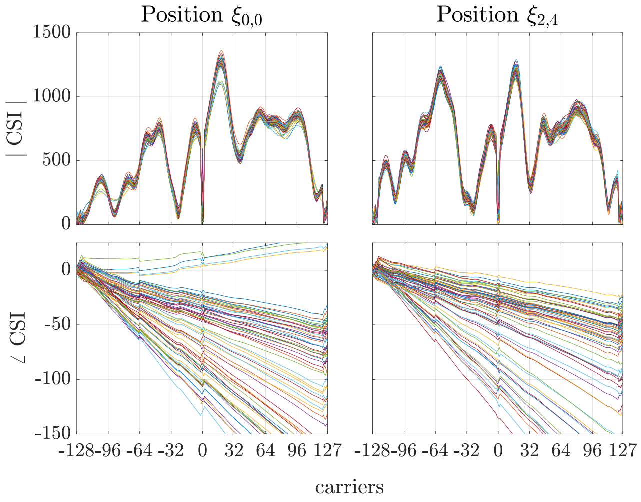

Figure 4: Plots of the amplitude (upper row, Broadcom 4365 802.11ac chipset units) and phase (lower row, radiants) versus the carrier number with the person in two different training spots ξ (see Section 4).

Figure 4 reports the amplitude and phase of the OFDM signal collected on 70 packets at the receiver with a person in two different locations. The details of the experiments are described in Section 4; what is important for our discussion is that the NN is fed exactly with these quantities, so that whatever the NN learns it must be present here. It is clear that the characteristics are remarkably constant across different packets for the same position, so that learning is feasible. Interesting features that are visible are the peaks and notches in the amplitude, and the phase jumps; however, the notch around carrier 0 and the phase jumps seem to be independent from the position of the person (result confirmed by nearly all positions), so that they are probably useless in position estimation. The other peaks and notches, instead, have positions in the spectrum that clearly depend on the person's location, which is probably what the NN learns.

We can try to disrupt the NN ability to learn the position of a person from CSI information by "fiddling around" with the CSI, so that it does not reflect exactly the electromagnetic fingerprint of the environment (thus revealing the position of the person in a deterministic way), but it also contains “deceit features” that confuse the learning and decision process of the NN.

3 CSI RANDOMIZATION

It seems a sound approach to apply a random distortion to the transmitted CSI in such a way that a receiver can still equalize the channel—i.e., not destroying the communication—while localization efforts based on CSI characteristics are hampered.

With reference to Fig. 1, we decided to apply the appropriate distortion at the output of the IDFT block, right before the DAC and the front end, in the orange block named “randomizer” in the figure. This position may be sub-optimal w.r.t. other positions earlier in the transmission chain, but it makes the technique easily understandable and it can also be implemented outside chipsets, allowing the realization of specialized, privacy-preserving devices without the need to develop a new chipset from scratch. With “randomization” in our context we refer to a manipulation of the transmitted signal so that additional peaks, notches or phase jumps appear randomly in the CSI. This disturbance must not be “white,” looking like a memoryless random process, because otherwise the NN would very simply filter it out, selecting only the location-dependent features. In this paper we use a very simple strategy changing the random pattern every 100 packets (corresponding to roughly a second with the packet rate we use), leaving a theoretical analysis of the optimal changing process for future work. Furthermore, the disturbance must not be too distorting to avoid hampering the communication performance. The actual manipulation is different depending on the disturbance that we introduce as shown in Fig. 5, where $s(\cdot)$ is the signal at the output of the Inverse Fast Fourier Transform (IFFT).

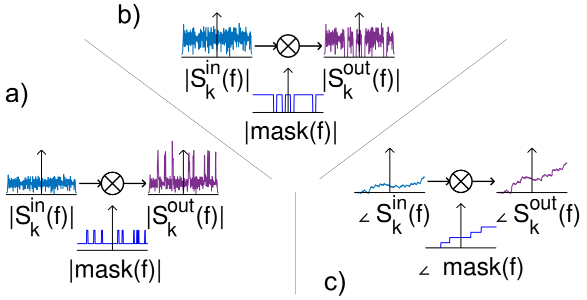

Figure 5: The tree different manipulations we experiment to randomize the position: a) adding random peaks in the channel response; b) introducing random notches; c) introducing random phase jumps.

We apply the disturbance filtering the sequence of I/Q samples in the frequency domain. Knowing the modulation format (VHT-PHY at 80 MHz with Long Guard Interval) we can collect the samples coming out from the IDFT block that build each symbol. Therefore, since we know the structure of each symbol, we can *i)* easily invert the encoding process and go back to the coefficients assigned to each of the 256 pilots; *ii)* multiply them by the corresponding values of the filter at their frequency; and *iii)* recreate the corresponding “tampered” OFDM symbols.

We tested several different types of filters to explore the effect on localization given by basic *shapes* (like peaks or notches), and the number of such repeating shapes. In all cases we built the spectral response of the filter starting from the identity filter in the frequency domain, i.e., a sequence of unitary samples. With reference to Fig. 5, we tested our randomization technique by applying: a) a random number of peaks (between 5 and 10), each of uniform random width between 2 and 6 samples; the amplitude is constant equal to 8; b) a random number of notches: similar to the previous case, but nulling instead of amplifying the selected points; c) randomly placed phase jumps: all points after randomly selected positions accumulate a phase delay of $\pi/2$, with the number of positions selected randomly from 2 to 6.

From a preliminary performance analysis it turned out that the effect of notches and phase jumps on the randomization of the estimated position is negligible, while interesting results were obtained for the filter with random peaks. Localization results for one sample point under different types of randomization are shown in Fig. 6; it is clear from this example that even if the localization precision is affected by applying random phase shifts to the CSI (b), an attacker can still make a sensible guess about the user's location. Only the introduction of randomly-placed peaks in the CSI (d) effectively degraded localization performance; therefore, we carry on the analysis using this filter.

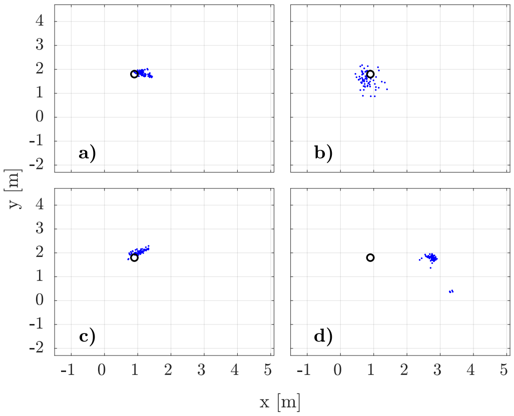

Figure 6: Localization results for one target point (black circle) under different conditions: a) without CSI modification; b) with selective phase shifting; c) with randomly-placed notch filters; d) with randomly-placed spikes.

4 EXPERIMENTAL SETUP AND METRICS

The initial experimentation is carried out in a Laboratory of the ANS[1] group at the University of Brescia; Appendix A reports all the implementation details needed to replicate our experiments. Figure 7 reports a plan of the laboratory with the position of the Tx and Rx devices and the places where a human being stood to train the neural network performing the CSI analysis. The goal of the experiment is twofold. First, we want to validate the results presented in [8] to guarantee that they can be reproduced with an independent implementation; next we want to measure the

[1] The Advanced Networking Systems group https://ans.unibs.it/ is one of the research groups in telecommunications at the University of Brescia

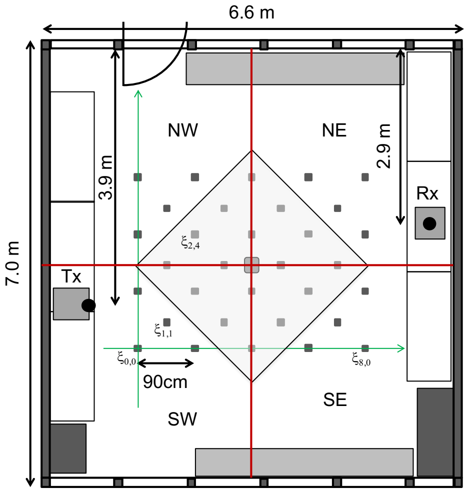

Figure 7: Layout of the experimental setup at the University of Brescia, the square dots on the floor are the training spots (ξ) for the neural network; the space is divided into four quadrants SW–NE for the sake of clarity.

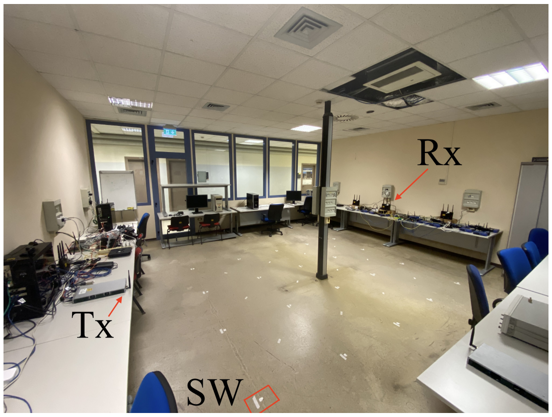

Figure 8: Photo of the laboratory described in Fig. 7, the pole in the middle with electrical outlets clearly creates a complex electromagnetic environment; the transmitter is labeled with 'Tx' on the west side while the receiver used in the one in the middle on the east side of the lab. Some of the ξ spots are visible on the floor (white dots) with the coordinate origin one marked 'SW'.

effectiveness of the randomization technique presented in Section 3 for obfuscating the actual position of a person.

The Euclidean coordinate system origin is set on the South-West (SW) corner of the grid as shown by the thin green axes; to make explanations easy we assume the lab is oriented with the north to the top. The training spots $\xi_{x,y}$ are numbered starting from the axis origin, so that $\xi_{0,0}$ is in the origin and $\xi_{8,6}$ is the one in the North-East (NE) corner.

Let us focus on the scenario where the goal of the attacker, i.e., the person who is trying to illegally track the position of someone, is to know in front of which working desk somebody working in the laboratory is passing his time, for instance to determine the fraction of his work time dedicated to different tasks, an act contrary to labor legislation in most countries, at least in Europe. To this end, the map in Fig. 7 is divided into four sectors: NW, NE, SE, SW, as shown with the red thick lines, while Fig. 8 presents a photo of the lab, which clearly shows the realism of the setup. The shaded square of 2 m edge at the center of the room is not considered for the localization purposes, as it is clearly an area where a person would not normally stay, but simply transit moving between the quadrants of the lab. As a side note, consider how simple it is to setup such an attack: the presence of an AP in a laboratory is very likely, a small sniffing device can be hidden easily, the training can be done when nobody else is present; given all this, then the attacker can very easily tell how much time the person spends in which part of the lab.

Within the scenario sketched above, we selected two different metrics to measure the localization performance of the methodology described in Section 2 and the effects of CSI randomization introduced in Section 3 to protect the privacy of people.

The first metric is a Euclidean distance measure to verify and validate the methodologies under analysis. However, the classical mean square error of the distance is not appropriate for our goals. The NN outputs a (x, y) position in a plane (2D), while a human body occupies a fairly vast space in 3D, so that it is indeed not possible to define the distance between the body and the (x, y) estimate. Call ρ a radius around the point estimate (x, y) of the NN, so that the circle of radius ρ and center (x, y) can be considered the projection of the human body on the 2D plane. If not otherwise stated, in this paper we can consider $\rho = 0.25$ m. Furthermore, recall that the NN is not used as a classifier, but it is trained to output a position in the Cartesian coordinates of Fig. 7, so that there is no decision process (e.g., assign the position to the nearest training spot ξ) in the NN.

Given the coordinate estimate (x, y) as computed by the NN, and the coordinates (x_c, y_c) of the ξ where the person stands, we construct a localization reliability $\mathcal{L}_\mathrm{R}$ index as follows

$$\mathcal{L}_\mathrm{R} = \frac{1}{N_l}\sum_{i=1}^{N_l} \mathcal{I}_d(i)\ ;\ \mathcal{I}_d(i) = \begin{cases} 1 & \text{if } d_i < \rho \\ 0.5 & \text{if } \rho \le d_i < 2\rho \\ 0.25 & \text{if } 2\rho \le d_i < 3\rho \\ 0 & \text{otherwise} \end{cases} \qquad (1)$$

where $d_i = \sqrt{(x - x_c)^2 + (y - y_c)^2}$ is the Euclidean distance between the position estimate and the coordinates of the ξ where the person was when the i-th sample is taken, and N_l the number of position estimates collected.

Clearly $\mathcal{L}_\mathrm{R} \in [0, 1]$, converges to one when all position estimates are within ρ from the true position and converges to zero when all estimates are three times ρ from the true position. As a useful comparison to understand the reliability of localization we can use the metric in Eq. (1) assuming the location is simply a random point in the portion of the laboratory where a person can reasonably stay, i.e., the lab minus the border where tables and furniture are. If we exclude 0.8 m around the wall, then the useful area of the lab is $A_u = (6.6 - 0.8) \times (7.0 - 0.8) \simeq 36\,\mathrm{m}^2$, thus randomly placing the location of a person inside this area yields

$$\mathcal{L}_\mathrm{R}^\mathrm{rand} = \min\left(1, \frac{15\pi\rho^2}{4A_u}\right) \qquad (2)$$

as a function of ρ, neglecting the border effect[2].

The second metric is instead focused on the privacy breaching scenario we sketched above, and measures the capability of the system to actually localize a person with high reliability, but with relaxed precision. The person does not stand in a specific location, but stays, possibly moving slowly, in one of the four quarters (NW, NE, SE, SW) of the lab. To this second end we use the probability defined in Eq. (3), where P_l is the empirical probability that the attacker successfully infer the position of the person in the lab quarter where he actually is, N_l is the number of position estimations collected, including those that infer the position in the shaded area that are considered "wrong" ($\mathcal{I}_l(i) = 0$), and $\mathcal{I}_l(i)$ is an indication function that tells if the i-th location estimation is correct ($\mathcal{I}_l(i) = 1$) or not ($\mathcal{I}_l(i) = 0$).

$$P_l = \frac{1}{N_l} \sum_{i=1}^{N_l} \mathcal{I}_l(i) \tag{3}$$

It is immediate to notice that if the location samples i are taken at fixed intervals, i.e., sending packets at constant intervals, and P_l is close to one, then the time spent by the person in the different lab corners is revealed to the attacker.

5 INITIAL RESULTS

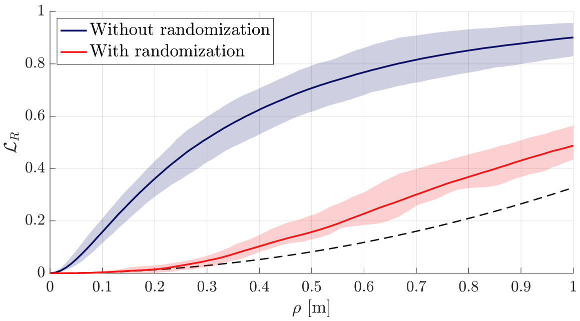

Figure 9: Localization performance according to metric $\mathcal{L}_R$ for ρ ranging from 0 to 1. Solid lines report the average result, while the shaded areas are the envelope of all measures including using different antennas and positions of the receiver. The dashed line represents the theoretical result for uniformely distributed random guesses.

Our first experiment is focused on evaluating the localization capabilities of an attacker that is using the method described in Section 2.2. Here we compare the performance obtained with and without the randomization procedure; the randomization is applied (or not) to the entire experiment, i.e., both training and testing phases. For both cases, we train the NN with 700 packets for each one of the 32 positions highlighted in Fig. 7 and then we test the localization on a different set of measures consisting of 150 packets per position collected at a different time. We capture CSI data from each of the four antennas available at the three receivers in the lab (visible in the picture of Fig. 8) for a total of 12 CSI feeds. Interestingly, results for the three receivers are similar and it also turns out that training the NN with data from one single antenna or any combination of the four antennas for each of the three receivers does not have any significant impact on the results. Figure 9 reports the average results obtained considering the metric $\mathcal{L}_R$. The solid line is the average for the 32 positions computed by considering all the CSI (average over 12 antennas). We also show the shaded regions between the worst and best performing antenna, obtained again by averaging over the 32 positions. While all the lines increase with ρ, the most interesting cases for localization are the ones with small ρ, e.g., for values between 0.2 and 0.6. In particular, for $\rho = 0.3$, the average $\mathcal{L}_R$ score is above 0.5 for the localization system, but drops below 0.05 when CSI randomization is active. The benefit of using our randomization system is evident from the fact that the curves obtained using randomized CSI are much closer to the black dashed line corresponding to the approximated result for uniformly distributed random guesses.

With respect to the second metric defined in Section 4, which represents the probability of an attacker locating the victim at least in the correct quadrant of the lab, we measured the performance of our system by comparing the value of P_l when the randomization is active or not and when the user is moving outside the shaded area in Fig. 7. Under normal conditions, the NN predicts positions that fall in the correct quadrant with probability 0.66. Despite the result appearing quite low, from Fig. 10 we can clearly identify clusters of points in the correct quadrant that would help the attacker making a more sensible guess. However, we can see from the same Figure that this analysis is significantly hampered when randomization is active: in this case, in fact, the probability that the estimated position falls in the right quarter of the lab drops to 0.30. We recall that random guessing the position of the victim in one of the four quadrants would give a value of 0.25, so we can affirm that our system is successfully defeating the localization mechanism.

For better assessing the performance of the system with and without randomization, we ran a second experiment where we collected CSI data when the user was sitting at four different desks located at each corner of the lab as indicated in Fig. 11. In this experiment the user can slightly move, i.e., by rotating over the chair vertical axis, or by moving arms and hands on the desk. The only constraint is to stay within the circles reported in the Figure. It is clear the while without randomization the NN predicts the positions with very high accuracy (they almost always fall within their circle), when our system is on the NN always fails.

5.1 Impact on Throughput

So far we have discussed only CSI-based localization and the possibility of obfuscating the location information by random manipulation of the OFDM symbols prior to transmission. In our experiments, we only transmitted with the lowest-order modulation and coding scheme (MCS) (i.e., `MCS0`) that uses BPSK. However, it is important to investigate the communication performance for higher-order MCSs because they are more susceptible to channel errors. Hence, we computed the Packet Delivery Ratio for all VHT-PHY MCS transmitted with 80 MHz bandwidth and a single spatial stream:

[2] The border effect, i.e., areas with a positive weight in Eq. (1) that are outside A_u, is marginal for small ρ and underestimate the location reliability as ρ increases, because it counts as a valid position also portions of the lab that are outside A_u, thus $\mathcal{L}_R$ is an actual lower bound for the location reliability, giving a good reference for the location reliability reduction provided by the randomization techniques we propose.

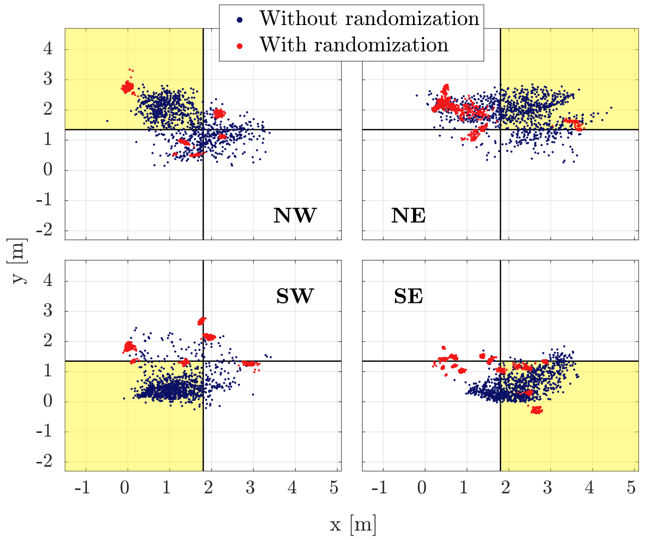

Figure 10: Positioning results considering different quarters of the lab. Each dot represents the position estimated by the attacker for each packet received. We report the estimates performed with and without CSI randomization using orange and purple dots respectively.

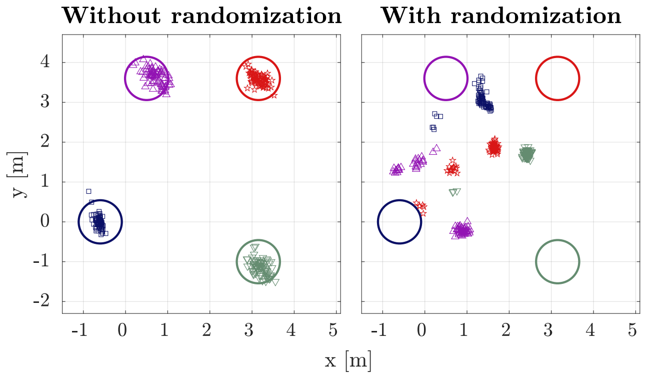

Figure 11: Localization results when the user is working at four different desks placed in the corners of the lab and not moving.

Table 1: PDR as a function of the MCS, with and without the CSI obfuscation block.

MCS		Rx 1		Rx 2		Rx 3	
		w/o	w	w/o	w	w/o	w
index	Mbit/s	%	%	%	%	%	%
0-BPSK	29.3	95.5	94.8	96.4	95.0	95.7	94.9
1-QPSK	58.5	90.7	92.8	91.7	93.8	91.5	93.3
2-QPSK	87.8	95.6	94.6	96.1	94.6	96.2	94.9
3-16-QAM	117.0	92.4	93.2	92.7	94.0	92.8	94.0
4-16-QAM	175.5	92.2	91.8	92.4	94.3	92.3	94.3
5-64-QAM	234.0	93.2	11.9	94.4	91.2	93.8	80.2
6-64-QAM	263.3	94.0	4.7	93.3	93.4	93.4	65.8
7-64-QAM	292.5	94.3	1.1	95.4	79.8	95.5	40.9
8-256-QAM	351.0	92.4	0.0	93.6	12.0	93.6	0.1
9-256-QAM	390.0	71.0	0.0	94.9	0.2	94.3	0.0

we report in Table 1 the PDR for the three receivers when randomization is off (w/o) and on (w).

It appears from the table that the positions of the three receivers enable acceptable performance for all MCS: only one receiver (Rx 1) suffers a bit with MCS9 without randomization. As easily predictable, only robust MCSs retain acceptable PDR when the randomizing filter is applied. In particular, when the modulation used is sensitive to distortion (i.e., 64- and 256-QAM modulations) the systematic errors introduced by the filter prevent correct decoding of the frame at one receiver (Rx 3) and kills reception at another one (Rx 1). Things get worse when further increasing the MCS: MCS8 and 9 cannot be received at almost any position.

Let's analyze the reasons and discuss how this problem may be overcome, as it is clear that a localization obfuscation method cannot destroy communication capabilities. Eq. (4) describes the mathematical model of the signal at the receiver, where $S(F)$ is the signal spectrum at the output of the IDFT of the 802.11ac transmitter, $M(f)$ is the filtering function used for obfuscation (the 'mask' in Fig. 5), and $H(f)$ is the channel response, including attenuation, distortion and multipath fading.

$$R(f) = S(f) \times M(f) \times H(f) \qquad (4)$$

Equation (4) is an equivalent model of the system we propose, but it's not how we actually perform the obfuscation, as it is done in the digital domain an not in the analog one as described by Eq. (4). $R(f)$ is converted back in the digital domain and passed through the dynamic equalizer driven by the CSI information before it is fed to the digital receiver implementing demodulation and error correction. From Eq. (4) it is clear that to preserve the communication performance $M(f) \times H(f)$ should maintain the properties of an equivalent, physically realizable and admissible (for 802.11ac) channel response $H'(f)$. The theoretical analysis of the properties of a randomizing filter that preserve this property and also obfuscate location is part of the future work of our research teams.

6 DISCUSSION AND FUTURE WORK

New technologies cannot come at the price of reducing people's rights. High performance Wi-Fi communications, which use advanced signal processing techniques to compensate the distortions of the electromagnetic environment, enable tracking the location of people, even if they do not carry any Wi-Fi device with them.

In this paper we have proposed a novel methodology that, by introducing carefully crafted random distortion of the Wi-Fi signal spectrum at the transmitter, prevent inferring the position of a person in a room exploiting the CSI at the receiver. We have shown with a real implementation that the method is feasible and works as intended. At the same time it does not completely destroy communications as, for instance, jamming would do, and this is fundamental to have location-obfuscation techniques widely adopted.

We think that this methodology can be widely adopted, even finding its way into future standards, so that citizens can use Wi-Fi without even bothering that their precise location can be tracked by unauthorized people. The way ahead to achieve this goal however is still quite long, and it includes having a full, formal understanding of the signal manipulations that allow achieving location obfuscation without hampering communications *at all*; extensive experimental campaigns to verify that the technique works with different localization methodologies and different channel bandwidths; analysis of more sophisticated attacks based on the joint analysis of many MIMO channels, and so forth. Finally, we would like to investigate if this same methodology can be applied to obfuscate the position of a device, rather than the position of a person who does not hold or wear a communication device.

ACKNOWLEDGMENTS

This work has been co-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, at TU Darmstadt by the LOEWE initiative (Hesse, Germany) within the emergenCITY center and the German Research Foundation (DFG) in the Collaborative Research Center (CRC) 1053 - MAKI.

REFERENCES

[1] Moustafa Abbas, Moustafa Elhamshary, Hamada Rizk, Marwan Torki, and Moustafa Youssef. 2019. WiDeep: WiFi-based Accurate and Robust Indoor Localization System using Deep Learning. In *Int. Conf. on Pervasive Computing and Communications (PerCom)*. IEEE, Kyoto, Japan, Mar. 2019, 10.

[2] Fadel Adib and Dina Katabi. 2013. See through walls with WiFi!. In *Conf. of the Special Interest Group on Data Communication (SIGCOMM)*. ACM, Hong Kong, Aug. 2013, 75–86.

[3] Chenwei Cai, Li Deng, Mingyang Zheng, and Shufang Li. 2018. PILC: Passive Indoor Localization Based on Convolutional Neural Networks. In *2018 Ubiquitous Positioning, Indoor Navigation and Location-Based Services (UPINLBS)*. IEEE, Wuhan, China, 6.

[4] IEEE Standard for Information technology. 2016. 802.11-2016 - Telecommunications and information exchange between systems Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.

[5] Francesco Gringoli, Matthias Schulz, Jakob Link, and Matthias Hollick. 2019. 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, 21–28.

[6] Xiansheng Guo, Nirwan Ansari, Fangzi Hu, Yuan Shao, Nkrow Raphael Elikplim, and Lin Li. 2020. A Survey on Fusion-Based Indoor Positioning. *Comm. Surveys & Tutorials* 22, 1 (First Quarter 2020), 566–593.

[7] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG]

[8] Felix Kosterhon. April 2020. *Device-Free Indoor Localization: A User-Privacy Perspective.* Master's thesis. Technische Universität Darmstadt, Secure Mobile Networking Lab, Department of Computer Science.

[9] Ramjee Prasad. 2004. *OFDM for Wireless Communications Systems.* Artech House, London, UK.

[10] Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora. 2016. PhyCloak: Obfuscating Sensing from Communication Signals. In *13th Conf. on Networked Systems Design and Implementation.* USENIX Association, Santa Clara, CA, USA, Mar. 2016, 685–699.

[11] Tahsina F. Sanam and Hana Godrich. 2018. An Improved CSI Based Device Free Indoor Localization Using Machine Learning Based Classification Approach. In *26th Eur. Signal Proc. Conf. (EUSIPCO)*. IEEE, Rome, Italy, Sept. 2018, 2390–2394.

[12] Schmidt, Erik and Inupakutika, Devasena and Mundlamuri, Rahul and Akopian, David. 2019. SDR-Fi: Deep-Learning-Based Indoor Positioning via Software-Defined Radio. *IEEE Access* 7 (Oct. 2019), 145784–145797.

[13] Matthias Schulz, Francesco Gringoli, Jakob Link, and Matthias Hollick. 2018. Shadow Wi-Fi: Teaching smartphones to transmit raw signals and to extract channel state information to implement practical covert channels over Wi-Fi. In *Int. Conf. on Mobile Systems, Applications, and Services (MobiSys'18)*. ACM, Munich, Germany, June 2018, 256–268.

[14] Matthias Schulz, Daniel Wegemer, and Matthias Hollick. 2017. *Nexmon: The C-based Firmware Patching Framework.* https://nexmon.org

[15] Xu Wang, Lingjun Gao, and Shiwen Mao. 2016. CSI Phase Fingerprinting for Indoor Localization with a Deep Learning Approach. *IEEE Internet of Things Journal* 3, 6 (Dec. 2016), 1113–1123.

[16] Guan-Sian Wu and Po-Hsuan Tseng. 2018. A Deep Neural Network-Based Indoor Positioning Method using Channel State Information. In *Int. Conf. on Computing, Networking and Comm. (ICNC)*. IEEE, Maui, HI, USA, Mar. 2018, 290–294.

[17] Kaishun Wu, Jiang Xiao, Youwen Yi, Dihu Chen, Xiaonan Luo, and Lionel M. Ni. 2013. CSI-Based Indoor Localization. *IEEE Trans. Parallel Distrib. Syst.* 24, 7 (July 2013), 1300–1309.

[18] Zheng Yang, Zimu Zhou, and Yunhao Liu. 2013. From RSSI to CSI: Indoor Localization via Channel Response. *ACM Comput. Surv.* 46, 2 (Dec. 2013), 25:1–25:32.

A IMPLEMENTATION DETAILS

The localization system is implemented on Commercial Off-The-Shelf (COTS) devices: for the receiver we use the *ASUS RT-AC86U* router; for the transmitter instead, we use an Ettus N300 SDR radio and we use the Matlab Wi-Fi toolbox to generate the frame and apply the randomization procedure.

At the transmitter, a never-ending Matlab loop generates a Wi-Fi frame at each iteration with random payload, and randomizes the corresponding raw signal as explained in Section 3. To this end, the software converts the packet into a vector containing I/Q samples according to the VHT-PHY modulation with one spatial stream and 80 MHz bandwidth. Then, it parses the vector and separates the VHT-PHY symbols. For each of them, and corresponding to its structure, the software may apply a specific procedure to isolate 256 I/Q samples from the symbol and apply the DFT to get back the OFDM coefficients that Matlab assigned to each carrier. Once in the frequency domain, the software multiplies each coefficient of the OFDM spectrum by the value of the randomization filter at the same frequency. Finally, it generates a new sequence of 256 I/Q samples by inverting the OFDM spectrum applying a IDFT: it also recovers the structure of the symbol by either adding the GI or by completing the missing part taking into account the periodicity.

At receivers we replace the firmware that controls the Wi-Fi card the nexmon-csi one [14]. With this software we have access to all the 256 subcarriers of each RX-TX stream. Received packets are saved to a pcap together with the corresponding CSI data that are conveyed to user-space as UDP datagrams. We extract CSI data from such packets and we feed the NN.

In designing the NN, we analyzed the influence of the overall number of training measurements as well as the number and distribution of positions needed to properly train the NN finding out that in general an even distribution of positions in a regular grid lead to a reasonable accuracy.

We collected measurements on different days, to evaluate the system under realistic conditions, and the approach provides a stable accuracy over several days, assuming a setup without external distortions or significant changes in the room furniture placement. Training needs to be done periodically to grant precision. We further observe that neither the characteristics of the tracked person nor the transmission power affects the precision of our system, while changes in the location and orientation of the router heavily influence the localization precision as expected. Further information on this topic like possible improvements, various countermeasures, and the usage of a high-level feature like the Angle-of-Arrival (AoA) can be found in [8] along with more details on the described aspects.

Software Availability. All the software we have developed is available through github at
https://github.com/seemoo-lab/csicloak
where references to datasets for validation are also available. Further information on the CSI-MURDER project-experiment and its evolution is available at https://ans.unibs.it/projects/