Adversarial WiFi Sensing

Yanzi Zhu[†], Zhujun Xiao[‡], Yuxin Chen[‡], Zhijing Li[†], Max Liu[‡],
Ben Y. Zhao[‡] and Haitao Zheng[‡]

[†]University of California, Santa Barbara [‡]University of Chicago

{yanzi, zhijing}@cs.ucsb.edu {zhujunxiao, yxchen, maxliu, ravenben, htzheng}@cs.uchicago.edu

ABSTRACT

Wireless devices are everywhere, at home, at the office, and on the street. Devices are bombarding us with transmissions across a wide range of RF frequencies from a few kilohertz to terahertz. Many of these invisible transmissions pass through our bodies, while others reflect off, carrying information about our location, movement, and other physiological properties. While they are a boon to medical professionals with carefully calibrated instruments, they may also be revealing private data about us to potential attackers nearby.

In this paper, we examine the problem of adversarial WiFi sensing, and consider whether our wireless reflections pose a real risk to our personal privacy. We identify an adversarial localization attack, where bad actors using smartphones can localize and track individuals in their home or office from outside walls, by leveraging reflections of ambient WiFi transmissions. We experimentally validate this attack in 11 real-world locations, and show user tracking with high accuracy. Finally, we propose and evaluate defenses ranging from geofencing to rate limiting and signal obfuscation.

1 INTRODUCTION

Advances in wireless technology over the last decade have made wireless devices ubiquitous in our homes, offices, and outdoor settings, covering nearly all areas where urban populations reside today. These devices inundate our surroundings with invisible RF signals covering a wide range of frequencies, from low frequency signals like GPS, AM/FM, and WiFi, to very high frequencies in the millimeter wave or terahertz range (*e.g.*, for 5G cellular picocells). While some signals pass harmlessly through our bodies, others bounce off of our bodies, giving professionals with specialized equipment information about our emotional states, heart rates, or even postures [27, 28, 32, 53, 54, 64].

But are we unknowingly revealing too much about ourselves and our actions? While we live and move in areas densely covered by wireless signals, we remain largely oblivious to the amount of information our bodies divulge on a continuous basis. But how much information are we revealing through wireless reflections, and what sensitive information can be learned this way by bad actors?

In this work, we consider these questions under the umbrella of *adversarial WiFi sensing*, where adversaries leverage reflections from ubiquitous WiFi signals to enable potentially malicious applications. While research in the wireless domain has shown that specialized equipment can extract precise information about us and our bodies, we limit ourselves in this initial work to a less powerful adversary: a mobile adversary without specialized RF transmitters or receivers who is only able to passively observe wireless signals using today's commodity devices.

We believe that, by leveraging statistical data mining techniques, even a weak adversary armed with only passive off-the-shelf WiFi receivers can perform invasive localization attacks against unsuspecting targets. Take for example the scenario of thieves looking to break-in to an office building, either to steal documents or to gain physical access to sensitive data. While it might be too conspicuous for them to bring along bulky specialized equipment, thieves might use commodity WiFi receivers to identify the location of any employees or security personnel, giving them a huge advantage in avoiding detection. Similarly, bad actors could track the location and movements of the occupants of a house, as a precursor to burglary or other crimes. In both cases, attackers would take advantage of near-ubiquitous WiFi transmissions (digital assistants, WiFi access points), to passively locate and track moving users.

In this paper, we study the general problem of adversarial WiFi sensing as a threat to location privacy, and empirically validate the feasibility of adversarial location attacks using passive WiFi sensing. Our work focuses on adversarial localization and tracking of users, unlike prior work on applications that sense users' gestures [35], emotions [62], or physical condition [8]. Our attack scenarios call for passive inference using commodity hardware, since active RF transmitters are obtrusive and easy to detect.

Our work makes four primary contributions:

- First, we identify the risks to location privacy from human body's blockage and reflections of ambient WiFi signals.
- Second, we propose a two-step algorithm for adversarial localization and tracking of unsuspecting moving targets across rooms.

- Third, we implement a prototype of the attacker system on commodity smartphones, and show (using real-world measurements) that the attack is feasible and accurate in 11 different settings, including both office buildings and residential apartments.
- Finally, we propose and evaluate three different possible defenses, including geo-fencing WiFi signals, rate limiting WiFi signals, and signal obfuscation.

There is a growing set of literature on how wireless signals can be used to sense and assess ourselves and our surroundings. Yet few have considered the adversarial aspects of RF sensing technologies, and what risks they pose to unsuspecting targets. We believe target localization and tracking is just one of many attacks made possible by inventive uses of RF signals, and we hope our work helps bring more attention from the security community to a topic with potentially numerous technical and social challenges.

2 ATTACK MODEL

We begin with a description of our attack model. While the space of adversarial sensing attacks is potentially quite large, we are interested in exploring potential attacks possible with a weaker adversary with limited resources. We consider an adversarial sensing attack, where the adversary sniffs WiFi transmissions outside of a target building (home or office), and uses captured signals to monitor the target's location and movements within the building. Our attack does not require the target to carry or use any RF devices or transmitters, instead using ambient WiFi transmissions generated by nearby WiFi devices.

We make the following assumptions about the adversary.

- The attacker does not have physical or remote access to WiFi devices in the target building. We assume such devices are secure and cannot be remotely compromised.
- To avoid detection, the attacker only performs passive WiFi sniffing, and does not actively transmit any RF signals, *i.e.* the attack is completely passive.
- The attacker is limited to compact, light-weight, off-the-shelf WiFi sniffing hardware. Given the nature of the attack, adversaries are unlikely to have access to bulky, expensive devices like directional antenna, antenna array, and USRP [5]. All of our experiments were performed using commodity smartphones with a single built-in WiFi antenna. The necessary functionality can also be implemented on a low-cost Raspberry Pi.
- We assume the adversary can physically move outside the victim's residence area, outside exterior walls or along public corridors inside buildings. To avoid detection, attacks should be performed quickly, within a matter of minutes. We also assume the adversary can place a small number (one or two) of compact sniffing devices hidden outside the target area.
- The adversary has access to rough floorplans of the target building or home. These are generally publicly available thanks to real state websites and apps, *e.g.*, Zillow, Redfin, Realtor.com.

3 OVERVIEW OF AN AMBIENT RF LOCALIZATION ATTACK

Before discussing the details of our localization attack using ambient WiFi signals, we first give a broad overview of the attack and its intuition. Significant prior work showed that by analyzing reflections of radio frequency (RF) emissions, software systems can "sense" users at various levels of granularity. The key enabler of this attack is the ubiquity of ambient RF emissions today. Whether it is routers, laptops, media sticks, or new IoT devices like voice assistants, cameras, door bells, smart appliances, and light switches, WiFi devices are in every room in our homes and offices, and constantly broadcasting wireless signals. Observations on these ambient transmissions are sufficient to provide the information necessary to sense and track users.

In our proposed attack, to accurately localize and track a user, an adversary first analyzes ambient WiFi emissions to determine the locations of static WiFi transmitters inside the building. We call these transmitters *anchor devices*. Their WiFi signals effectively create a dense net of "invisible" tripwires inside each room, using which the adversary can monitor user presence and movements inside the home/office.

We describe the sequence of actions below, then detail the key challenges and solutions in the following sections.

Step 1: Localizing Anchors inside the Target Building. The key idea is to leverage the inherent correlation between the received signal strength (RSS) of the sniffed WiFi packets and the distance from the anchor to the sniffer, and estimate the anchor location from RSS observed at various locations. To do so, the adversary walks and performs a brief measurement outside the target's location (along a public corridor inside an office building, or outside a house), using a standard sniffer device to passively listen to transmissions from WiFi devices inside the rooms (Figure 1). Since WiFi packets do not encrypt source and destination MAC addresses, the adversary can collect packets for each WiFi device and even infer their device type[1] from the packet header of the sniffed packets. The adversary then uses the RSS values of these sniffed packets, measured along the walking

[1] Existing works show that one can infer the device type from the Organizational Unique Identifier (OUI) field of the MAC address [39] and/or the traffic pattern [44], even under MAC randomization [29].

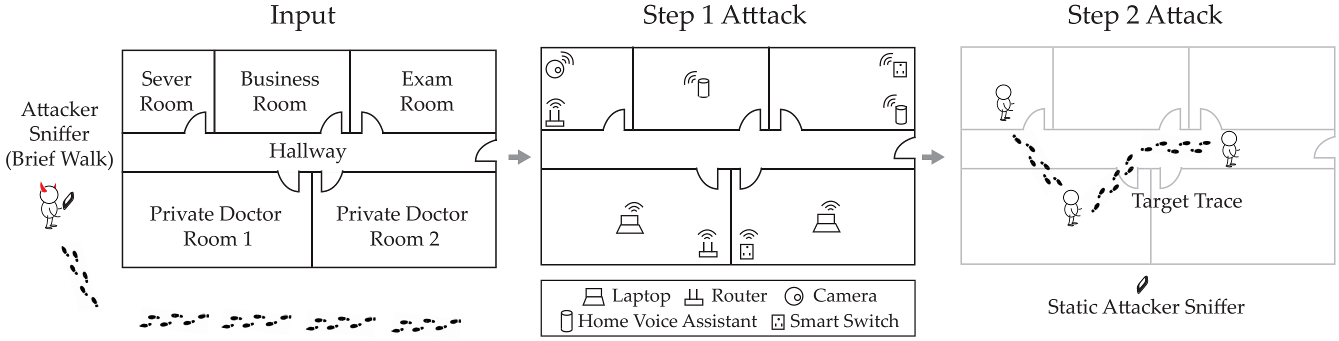

Figure 1: Illustration of our attack scenarios in a doctor's office.

path, to localize each corresponding WiFi device. Resourceful attackers could use robots or drones to carry out the measurements.

Step 2: Continuous Target Monitoring. Next, the adversary secretly places a stationary WiFi sniffer outside of the victim's home/office to continuously monitor WiFi transmissions. Using the detected WiFi devices as anchors, the adversary extracts subtle variations from their signals to identify and track how the target moves across individual rooms inside. The key insight here is that when moving, the target user will block or reflect WiFi signals sent by nearby anchor devices in the same room, triggering variations in signals captured by the adversary. So from the variations in the sniffed signals, the adversary can infer the target's location based on the locations of the "triggered" anchor devices.

The above attack is easy to launch in practice, requiring minimum physical effort from the adversary. Yet achieving accurate localization and tracking is challenging, because the adversary only uses passive sniffing to prioritize stealth. Existing works in indoor localization and user tracking have achieved accurate results, *e.g.*, [20, 37], but all require active communication and synchronization between the target devices and the adversary, which is false under our scenario.

Furthermore, since the adversary can only observe WiFi signals "behind-the-wall", the captured signals are aggregated from multiple paths, and are thus highly complex and hard to model. A successful attack design requires robust localization algorithms that address such complexity and the lack of ground truth reference points and measurements to calibrate the propagation model used in localization.

In the following sections, we describe our solutions to achieve an accurate device localization and the tracking of users, despite the complex propagation environments and unique hardware constraints. We first discuss our design for accurate anchor device localization in §4, then techniques for passive user detection and tracking in §5.

4 LOCATING ANCHOR DEVICES

The first step of the attack involves identifying and localizing stationary "anchor devices" inside the building. Their locations are used to help localize and track users in the step 2 attack. The key component of this attack is a brief spatial measurement of RSS (by the adversary walking outside of the building for a few minutes), and a localization algorithm that uses these RSS values to estimate the room locations of anchor devices (§4.1). We also present two enhancements where the adversary uses a static sniffer (also used by the step 2 attack) to identify static anchor devices and detect their floor levels (§4.2).

4.1 Room-level Device Localization

With a brief walk outside of the target's location area, the adversary makes spatial RSS measurements of neighboring WiFi devices at multiple locations along the trajectory. These values, together with the trajectory, are feed into a localization algorithm to estimate the locations of anchor devices.

Why RSS ? The localization uses simple spatial measurements of RSS rather than other advanced signal metrics like Angle of Arrival (AoA) [21, 46] or the phase component of Channel State Information (CSI). This is because commodity WiFi sniffers (smartphones, Raspberry Pis) cannot report these advanced metrics when operating in the passive mode. The sniffers are only equipped with a single antenna and thus cannot report AoA. The attacker can add more antennas, but needs a bulky antenna array of more than 1.5m in width to perform room-level localization with reasonable accuracy. Similarly, CSI-based localization relies on multiple antennas and the phase component of the CSI to derive AoA [26]. But the attacker sniffer is unable to obtain an accurate phase estimation due to the lack of synchronization with the transmitter. Recent work [21] estimates AoA from CSI amplitudes but only when the sniffer and targets are in complete line-of-sight, *i.e.*, no walls.

Overview. Due to the constraint of passive sniffing and the complexity of propagation environments, achieving accurate localization using limited RSS measurements is challenging. We address this challenge by combining traditional passive localization with statistical data mining. Our proposed localization design includes two layers. The base layer applies *RSS model fitting* to infer the transmitter location from a given input of RSS values. The data layer sifts through all the RSS measurements to identify the input to the base layer that leads to meaningful and consistent localization results. We now discuss these two layers in more details.

Base Layer: RSS Model Fitting. RSS model fitting [25, 27] is widely used for passive transmitter localization. Leveraging the correlation between RSS and signal propagation distance, it fits the captured RSS values into a propagation model to estimate the transmitter location. For our attack design, we use the log distance path loss model [42]. The detailed calculation is listed in the Appendix. We also experimented with other passive RF localization methods, including (weighted) centroid [13], gradient [17], and ecolocation [57]. They perform worse and require many more spatial RSS measurements.

RSS model fitting requires the walking trajectory during RSS measurements, which can be recorded using the IMU sensors (*e.g.*, the built-in accelerometer and orientation sensor on smartphones). For our attack, we built a smartphone app to record the trajectory and the RSS values simultaneously. The tracking error is less than $1m$ within each trace and has minimum impact on the localization result.

Data Layer: Consistency-based Data Sifting. With RSS model fitting, the localization accuracy depends heavily on the "quality" of the RSS measurements. Ideally, these measurements should contain little noise, align with the propagation model, and cover a wide range of values to minimize fitting bias. Yet in reality, the adversary has little control on the available walking path and the complex propagation environment. The resulting RSS measurements inevitably contain bias, noises and even human errors, leading undesirable localization outcomes.

A straightforward solution is to filter out "bad" measurements using de-noising methods, ranging from the traditional Kalman filter [14], wavelet filter [49] to the newly proposed feature clustering algorithm that remove bad measurement rounds [27]. We found that these methods are insufficient under our attack scenarios because the propagation environment is highly complex and unknown to the adversary, making it hard to distinguish between noise and natural propagation effect. Features used by [27] to identify bad measurement rounds are too coarse-grained to effectively control localization accuracy. Our experiments in §7 show that more than half of the good measurement rounds identified by [27] will locate the device to a wrong room.

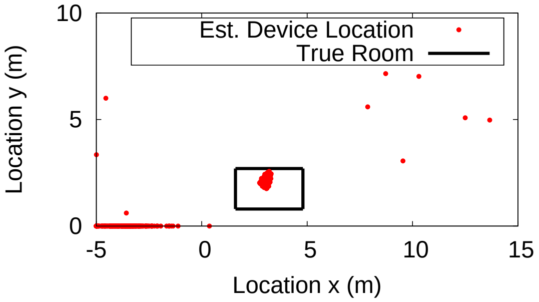

Figure 2: Localization results from our Monte Carlo sampling. Each red dot is the estimated anchor location from a sample; the rectangle marks the room of the anchor.

Instead, we propose *consistency-based data sifting* to identify proper data samples that will be used for model fitting. Our hypothesis is that, by the law of large numbers [41], *consistent* fitting results from many random sampling of RSS measurements, if exist, can reveal true signal propagation behavior and produce high-fidelity localization result.

Based on this hypothesis, we introduce two rounds of data sifting, one within each measurement round and one across different rounds. A measurement round represents RSS measurements collected during a single walk on the corridor.

(1) Data Sifting via Monte Carlo Sampling. Given a round of RSS measurements $\mathbb{R}$, we apply the Monte Carlo method [1] to randomly sample $\mathbb{R}$ as the input to the model fitting. This is repeated by $N = 1000$ times, producing N localization results. Using standard clustering algorithms like K-means, we find natural correlation clusters among these N results. If they form many small clusters with different room-level results, then $\mathbb{R}$ displays inconsistency and cannot be used for localization. If a dominant cluster exists and its averaged fitting mean square error (MSE) is less than those of the other clusters, then $\mathbb{R}$ can be used for localization.

Figure 2 plots an example result of the Monte Carlo sampling on a single round of RSS measurements. The sampling process produced a single, dominant cluster, while the rest of the result data points are widely scattered.

In this case, we consider the dominant cluster, compute the room location of each data point, and use them to compute the statistical distribution of the device's room location, *i.e.* the probability of the device being in each room. In the current design, we simply choose the room with the highest probability as the location of the device. A more advanced design could leverage statistical patterns of the clusters to refine localization decision. We leave this to future work.

(2) Consistency Check across Measurement Rounds. When multiple rounds of sniffing measurements are available, the adversary can also perform consistency check across

them. If the localization result (room-level estimate) is consistent across multiple rounds of measurements, then the result is confident. Overall, we found that consistency check across 4 rounds of measurements is sufficient to achieve room-level localization of 92.6% accuracy on average.

4.2 Attack Enhancement

The above attack design assumes all the sniffed WiFi devices are stationary and the target building has only one level. In practice, in an office or home environment, some WiFi devices (*e.g.*, laptops, smartphones and tablets, telepresence robots), can move around the building/room. Furthermore, when the target building has multiple floors, the adversary needs to group the detected WiFi devices by their floor levels in order to properly localize the devices on the target floor.

We now discuss two attack enhancements that uses RSS measurements from a static sniffer to identify static anchors and detect their floor levels. This static sniffer is also used by the second stage of the attack (§5).

Detecting Stationary WiFi Devices. The intuition is that the RSS value of a stationary WiFi device, captured by a stationary sniffer, should stay relatively stable, while those of mobile WiFi devices will fluctuate over time. Figure 3 plots three sample RSS traces observed in our experiments, corresponding to a static device, a moving robot device, and a static device being relocated to a different location. We see that the degree of RSS variations is a reliable indicator of a device's mobility status.

Thus during the step 1 attack, the adversary places a *static* sniffer outside the target area to record RSS values[2] of neighboring WiFi devices. Any device with small temporal variance in RSS is marked as static.

Later in our step 2 attack, the adversary also uses the same system (the static sniffer) to detect any relocation of the anchor devices. Once detecting any sudden variation, *e.g.*, the sudden RSS drop at 15s in Figure 3c, the adversary will either stop using the corresponding device as anchors or repeat the localization measurements to update its location.

Floor Level Signal Isolation. The floor level detection leverages the physical geometry of signal propagation: RF signals emitted by devices on different floor levels arrive at the sniffer in different (vertical) directions. If the sniffer can identify the incoming angle of the WiFi signal, *i.e.*, angle of arrival (AoA), we can infer the floor level. However, as discussed earlier, commodity sniffers cannot measure AoA.

We address this issue by adding a compact smartphone case to our sniffer, emulating a directional antenna. As shown by Figure 4, we place a simple cone object of size 8cm × 6cm ×7cm on top of the smartphone, and wrap the smartphone sniffer with aluminum foils. Now the sniffer can only capture WiFi signals through the cone. The adversary, standing or sitting, rotates the sniffer while it records the WiFi signal (RSS) and the phone angle (via the built-in gyroscope). The estimated AoA is the direction that maximizes the RSS value. The adversary then infers the floor level by comparing the estimated AoA value to the projected AoA values for different floor levels (calculated based on the floor plan).

We validate our design by the adversary staying on the first floor and measuring the AoA of the WiFi devices on the first and the second floor of a building (ground truth AoA of 0° and 25°, respectively). The measured AoAs for these devices are 5° and 32°, respectively, which are widely separated. This indicates that the devices are on different floors, proving the effectiveness of the floor detection.

Finally, we note that while the above AoA measurements are sufficient for floor level detection, they are too coarse-grained to perform anchor localization.

5 TRACKING MOVING TARGETS

Using located static WiFi devices as anchors, our step 2 attack uses a static sniffer to detect and monitor human targets over time, especially when they do not carry any WiFi devices. The key insight is that in an office/home setting, human users are never completely stationary. Whether it is typing on keyboards, waving hands, opening doors, sitting down, or standing up, their natural movements will disturb the signal propagation of the nearby anchor devices, creating subtle signal variations observable at the sniffer. Therefore, the adversary can detect human presence by detecting subtle variations among the sniffed signals of the anchor devices, and track movements over time from the temporal sequence of the "triggered" anchors.

In the following, we discuss our proposed design to extract signal variations triggered by target movements, followed by the detailed tracking algorithm. We then discuss several practical implications of our design.

5.1 Detecting Signal Variations

When it comes to extracting the (instantaneous) signal variations caused by the target movements, RSS is no longer a reliable metric. This is because movements near an anchor lead to fast channel fading and thus independent, subtle signal variations in each (narrowband) sub-carrier of the sniffed signals [52, 53]. RSS, on the other hand, is the received signal strength averaged over all the sub-carriers in a wide frequency band. The averaging removes the subtle variations in each sub-carrier.

Instead, our attack uses the amplitude Channel State Information (CSI) extracted from the sniffed signals, which measures the signal amplitude on each individual sub-carrier.

[2] We consider RSS values from packets of the same packet type, since different packet types may be transmitted with different transmit powers, causing variations in RSS even if the device is static.

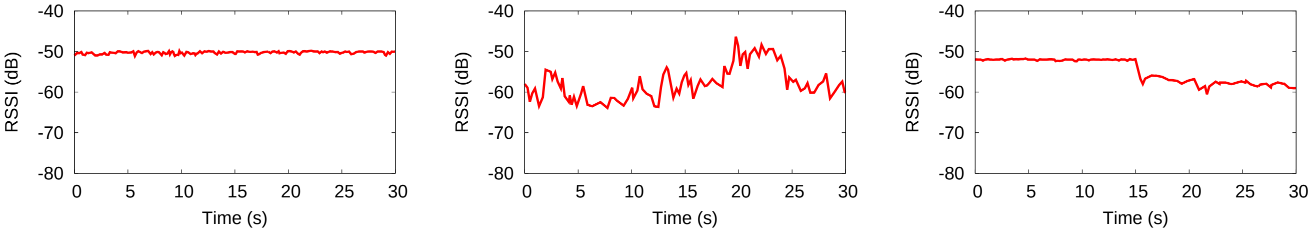

(a) RSS trace of a static device

(b) RSS trace of a mobile device

(c) RSS trace of a static device being relocated

Figure 3: Using RSS trace to distinguish between static and mobile devices.

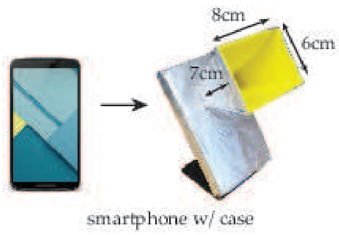

Figure 4: Our phone case prototype for floor level detection, which when rotated by the adversary, performs coarse-grained AoA measurements.

This fine-grained feature provides a microscopic view of the signal fluctuations caused by target movements. In general, WiFi devices can only report CSI values for active communications [16]. For our attack, we developed a smartphone app on the sniffer that is able to extract the CSI amplitudes from passively sniffed WiFi signals. We present the detailed implementation later in §6.

To detect the signal variations of a given anchor device, we process, for each sub-carrier, its observed CSI amplitude value over time. We take a window-based approach, detecting movements at a 5-second time granularity. Given a 5-second segment of the CSI observation, we first calculate the standard deviation θ_i of the amplitude at sub-carrier i, then average θ_i across all the sub-carriers. If the averaged θ goes beyond a threshold, a user movement is detected. We then move the window forward by 2.5 seconds, and compute θ for the next 5-second segment. In this process, we exclude from our calculation the sub-carriers whose amplitude values remain very high. Due to WiFi link adaptation, these sub-carriers operate at very high transmit power, and thus are less sensitive to user movement in proximity.

As examples, Figure 5 plots the CSI amplitude at a single sub-carrier, for scenarios of no human presence, a user sitting down on a chair, opening/closing the door, typing on keyboard, waving hand and walking. Compared with the case without any human in presence, one can clearly observe signal variations as a result of user movements, confirming that they can serve as a reliable indicator of the user presence and movement in a room.

There are a number of factors that could affect the fidelity of the signal variation detection, which we discuss next.

Trigger range of anchors. The amount of signal variations depends on the distance between the target and the anchor. The closer the target moves around the anchor, the more impact the target produces on the sniffed signal. Our experiments in §7 show that the triggering range of a typical WiFi devices is around 3 meters (9.8 feet). For a standard room in offices and homes (of size 10 feet× 15 feet), 2–3 anchors can cover the entire room.

Sniffer placements. When the target area has multiple rooms, the sniffer should be placed at locations where direct propagation paths from anchor devices to the sniffer do not align with each other, *i.e.* sharing the same AoA. This is feasible in practice because the attacker has the rough floor plan of the target area. In this way, user movements in a room do not trigger any anchors in a neighboring room. Finally, while our attack design and experiments only used a single sniffer, resourceful attackers can place 1-2 more sniffers to obtain a more complete view of the target.

Interference from pedestrians near the sniffer. External movements near the sniffer (*i.e.* the receiver) will create CSI variations at the sniffed signals, leading to false detection of user presence in the target rooms. This is particularly true when the sniffer is placed in indoor hallways. On the other hand, such event can be detected because movements near the sniffer will create sudden, simultaneous CSI variations and reduced RSS values at *all* the anchors (or at least the majority of them). When detecting such pattern, the attacker can mark the corresponding sniffer data as uncertain.

5.2 Tracking Targets over Time

After identifying a set of "triggered" anchor devices, the adversary identifies the presence of the human targets and tracks their movements over time. Without any prior knowledge of the targets and their movement patterns, our signal measurements are unable to recognize each individual user or the

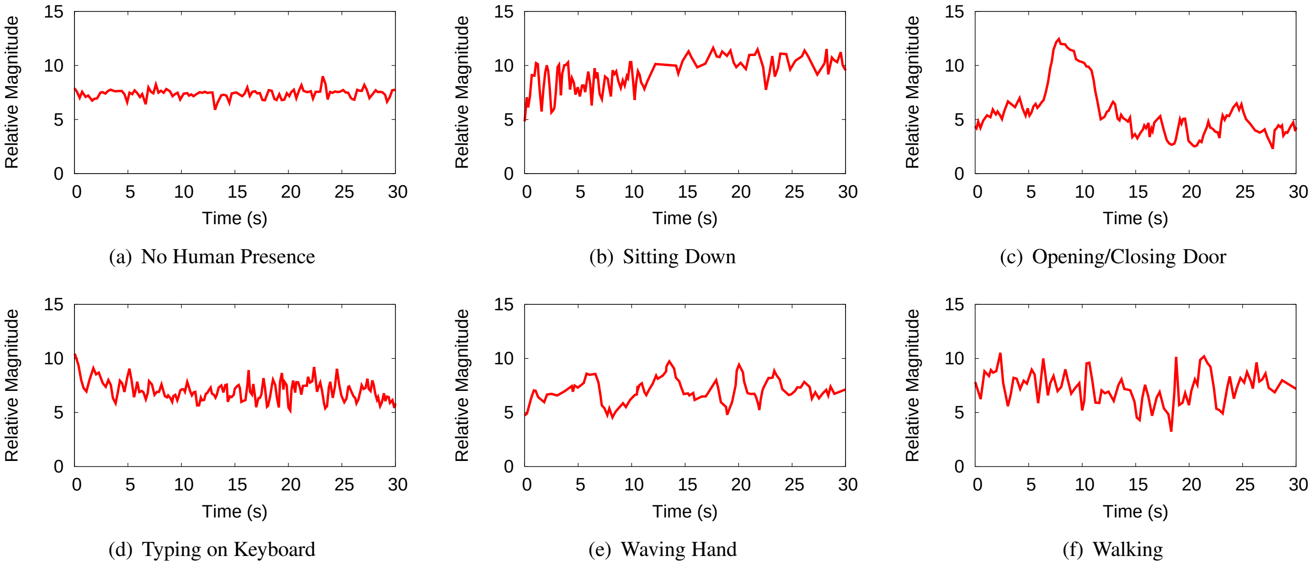

(a) No Human Presence

(b) Sitting Down

(c) Opening/Closing Door

(d) Typing on Keyboard

(e) Waving Hand

(f) Walking

Figure 5: Examples of CSI signal amplitude variations (at a single sub-carrier) caused by user movements.

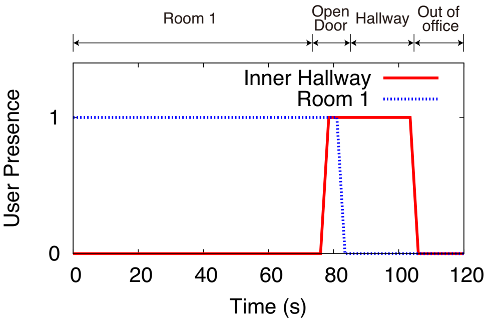

Figure 6: Tracking the movement of a mobile target.

number of users in the target areas. Instead, we detect, for each room, the presence of human users, and connect the sequential events to identify their trajectory. Given a room, the presence of human users is 1 if at least one anchor in the room is triggered, and 0 if no device is triggered.

Now consider an example scenario where a doctor works in his private office, gets up, opens the door, walks out of the office into the hallway and leaves the practice. Figure 6 shows the physical transition in time, and the detected user presence in the room and the hallway, respectively. Since the action of opening door will affect anchors in both the room and the hallway, the user presence values of the two rooms overlap briefly.

Finally, the efficacy of target detection/tracking depends on the localization accuracy of the anchor devices. In fact, our experiments in §7 show that our CSI based user presence detection is highly accurate, achieving a recall value of 87.8%, 98.5% and 99.8% with 1, 2, and 3 anchor devices in the room, respectively, and a precision value of 99.95%. The majority of the end-to-end sensing errors are due to the errors of anchor localization.

5.3 Discussion

Strengths and Limitations. The strength of our attack is its simplicity, generality, and stealthiness: it uses a single low-cost sniffer which does not transmit but *passively* monitors *ambient* transmissions of *many* WiFi devices; it does not require prior knowledge on traffic patterns of WiFi devices (which is required by existing attacks on home privacy [24, 39, 60]); it does not make any assumption on WiFi device placements or user movement/activity patterns.

A limitation of our current design is that the attacker is unable to recognize a specific user or a specific activity, *e.g.*, distinguishing between walking and waving hands. Doing so requires extensive knowledge on the activities and CSI patterns for each user and WiFi device, which is infeasible under our attack scenario. The goal of our attack is to detect the presence and movement of targets, rather than recognizing their detailed activities.

Using CSI for anchor localization (step 1)? A natural question here is why not use CSI for localization in the step 1 attack, as it provides more fine-grained information about signal propagations. This is due to two reasons. *First*, accurate CSI-based localization relies on multiple antennas and the phase component of CSI to derive AoA [26]. The commodity sniffer only has one antenna, and is unable to obtain an accurate phase estimation due to the lack of synchronization with the transmitter. *Second*, CSI amplitude is sensitive to nearby target movements. As the adversary has no knowledge of the target status at the time of the measurements, it cannot rely on CSI amplitude for localization. In comparison, RSS is robust against target movements and thus a more reliable metric for anchor localization.

6 ATTACK IMPLEMENTATION

We implement our attack using smartphones as sniffers, leveraging their built-in IMU sensors (accelerometer and gyroscope) to track walking trajectory and device rotation. Specifically, we use two popular and inexpensive Android phones, Nexus 5 and Nexus 6. They are equipped with the Broadcom WiFi chipset with a single antenna, and a customized WiFi firmware by Nexmon [40] to achieve passive sniffing. The captured RSS values range between $-85dB$ and $-30dB$, and the noise floor is $-85dB$. During the walking signal measurement, we use the built-in IMU sensors to detect and count user stride, and construct the walking trajectory. The RSS measurement is at a much faster rate, and we average the RSS values measured during a single stride.

We implement the RSS/trajectory measurements (step 1) and CSI measurements (step 2) as an Android app. The app runs in the background and consumes the same amount of the energy as that of system standing by.

Passive Sniffing of CSI Amplitudes. Traditionally, CSI can only be captured when the WiFi receiver is actively communicating with the transmitter [16]. Our attack leverages a recent development of WiFi firmware [40] to capture CSI amplitudes while operating in the passive sniffing mode.

Our implementation also addresses two artifacts in CSI measurements caused by the firmware. *First*, the firmware reports each CSI magnitude as a projected value between 0 and 40dB, where the projection factor is unknown. Thus we configure the movement detection threshold accounting for normalization. *Second*, the firmware can only report CSI amplitude values at a limited speed, up to 8–11 packets per second. Thus the app subsamples sniffed packets based on this rate limit. Despite these limitations, our prototype sniffer is able to capture sufficient CSI information to successfully launch the attack in practice.

Computation Cost. One strength of our attack is its simplicity. Currently we post-process all the measurement data in python on a MackBook Pro. Finishing 1000 rounds of Monte Carlo sampling and model fitting takes 1-3 seconds, while CSI based tracking is instantaneous. Thus these computations can be ported to smartphones or Raspberry Pis for real-time processing. We leave this to future work.

7 EVALUATION

In this section, we evaluate our attack design using experiments in typical office buildings and apartments. We start from describing our experiment setup and scenarios, and then present our evaluation on individual Step 1 and Step 2 attacks, followed by the end-to-end attack evaluation.

Sniffer Path	Test Scene	# of Rooms	Mean Room Size (m^2)	# of Devices	# of Building Floors
Indoor Hallway	1	6	14.19	7	10
	2	7	14.60	5	10
	3	8	13.65	3	37
	4	3	14.50	13	15
	5	3	9.51	5	13
	6	6	14.21	15	3
	7	5	16.75	8	3
	8	4	44.39	8	9
	9	2	69.83	4	3
Outdoor Sidewalk	10	2	47.20	4	3
	11	4	12.99	6	2

Table 1: Test scene configuration.

7.1 Experiment Setup

We performed attack experiments at 11 typical offices and apartments that are accessible to us. The owners of each test scene volunteered for our experiments. These test scenes are of different sizes and configurations, and have different wall materials except for concrete[3]. For each test scene, the building has multiple floor levels, but all the rooms of the test scene are on the same level. The walking path available to the adversary also differs across experiments, from indoor corridors to outdoor pathways. We listed the configuration of our test scenes in Table 1.

Inside each test scene, we either reused the existing WiFi devices or deployed our own WiFi devices. These are popular commodity products for smart offices and homes, *e.g.*, wireless security cameras, voice assistants, WiFi routers, and smart switches. In total, we have experimented with 31 WiFi devices, including 6 security cameras and 6 laptops. Table 2 summarizes the devices used in our experiments and their traffic patterns during idle and active periods. Even when idle, these devices periodically transmit packets (data, ACK, QoS maintenance). The packet rate varies from 0.5 packet per second (pps) to more than 100 pps.

All the WiFi devices are placed at locations where they are designed to be: security cameras at room corners, smart switches on the wall outlets, laptops on desks, and WiFi routers in the center of the room for coverage. We experimented with both 2.4GHz and 5GHz for WiFi connectivity and did not observe any notable difference.

To perform our step 1 attack, the adversary holds the sniffer while walking outside the target scene (indoor corridor or outdoor pathway). For each test scene, we collected 50 walking measurements, each of 25–50 meters in length and

[3] Our attack does not work when the wall separating the targets and the adversary is made of concrete, which blocks the majority of the WiFi signals.

	Device Type	Exact Product	Rate (pps), Idle	Rate (pps), Active
Static	Cameras (without Motion Detection)	AHD Security Camera	-	124
	Cameras (with Motion Detection)	Amcrest/Xiaomi IP Camera	≥0.5	108
	Home Voice Assistance	Amazon Echo, Google Home	2	16
	Smart TV (& Sticks)	Chromecast, Apple TV, Roku	6.64	200
	Smart Switches	LifeSmart Plug	≥2.44	≥3.33
	WiFi Router	Xiaomi/Cisco/Asus Routers	28.6	257
Mobile	Surveillance Robot	iPATROL Riley Robot Camera	-*	124*
	Smartphones	Samsung/Google/Apple Phones	≥0.5	≥6

Table 2: Summary of victim devices used by our experiments. We emulate the Robot Surveillance Cameras by mounting a camera on a robotic car.

0.5–2 minutes in time. We also vary the office WiFi device placements and repeat the experiments. In total, we collected more than 3k RSS measurement traces, with more than 121k location-RSS tuples.

For our step 2 attack, we hide the sniffer behind plants or at the corners (on the ground). We ask our lab mates to carry out normal activities in each test scene, and collect more than 12 hours of CSI entries. The lab mates are aware of the goals and results of the tracking but not the specific techniques.

7.2 Effectiveness of Step 1 Attack

To evaluate our Step 1 attack, we process the collected RSS traces to detect and locate stationary WiFi devices in the target areas and compare the result to the ground truth. Figure 7 plots, for each of the 11 test scenes, the average room-level localization accuracy across all the WiFi devices (for which we have ground truth location). We compare the performance of RSS model fitting with and without data sifting, and when applying feature clustering proposed in [27].

We make two key observations from this experiment. First, "blindly" feeding RSS measurements into model fitting leads to considerable amount of localization errors. In 5 out of the 11 test scenes, the adversary places more than 40% of WiFi devices in the wrong room, producing false anchors for the Step 2 attack. Second, our proposed data sifting significantly boosts the localization accuracy. For more than 90% of the cases, a device is placed at the right room. Our design also outperforms the feature clustering-based filtering [27] by using fine-grained data sampling rather than coarse features.

Impact of Anchor Placements. To correctly locate WiFi devices placed at room boundaries, the absolute error (distance to the ground truth location) needs to be very small (a few centimeters) for room-level accuracy. Thus these devices are harder to locate than those placed in the middle of the room. In our experiments, these boundary devices will create a dominant Monte Carlo cluster, but the data points in the cluster will map to either of the two neighboring rooms, and the resulting room probabilities of the two rooms are similar (differed by <20%). Our current design makes a simple binary decision by choosing the room with the higher probability, which could place the device in the wrong room. As future work, we plan to improve our design by marking these devices as "boundary" anchors and treating them differently in the step 2 attack.

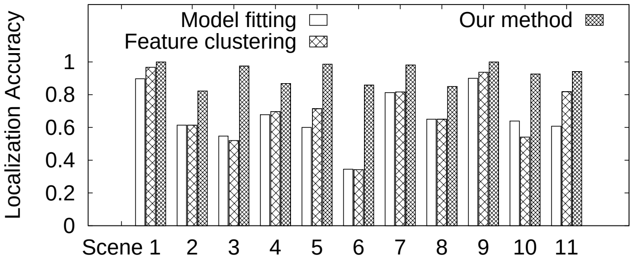

Figure 7: The localization accuracy (room-level) for static WiFi devices, for each of the 11 test scenes.

Impact of Anchor Transmit Rate. We found that the localization performance is insensitive to the device type and transmission rate. For all the 31 devices we have tested, they always transmit packets at 0.5pps and above. The RSS measurements are relatively time insensitive and thus can be aggregated over time. As long as the measurements cover over 20m in distance (space) and sample the RSS values evenly between -75dB and -30dB, we observed no notable difference in localization accuracy.

WiFi Devices outside of Target Area. During our experiments, our sniffer also captured signals from unknown WiFi devices that were placed outside of the target area. Since the adversary has no prior knowledge on the WiFi devices, it will localize these devices as well. We found that our localization design always placed these devices (with strong signals) outside of the target area. Devices with weak signals are automatically filtered out by our data sifting process. Overall, the adversary is able to isolate WiFi devices in the target area from those outside of the area.

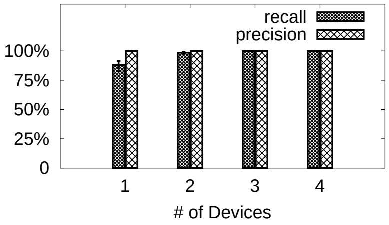

Figure 8: Performance of our user presence/movement detection.

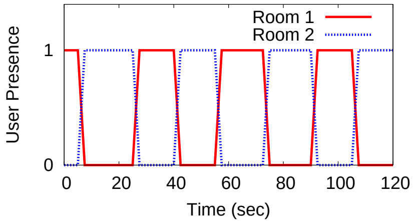

Figure 9: Detecting periodic patterns of a mobile target.

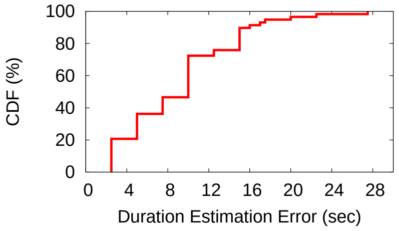

Figure 10: CDF of errors in movement duration estimation.

7.3 Effectiveness of Step 2 Attack

For our Step 2 attack, we experimented with different types of target activities and movements.

Detecting User Presence in a Room. Figure 8 plots the performance of our CSI based user presence/movement detection over 12 hours of CSI recordings across our test scenes. We present the result in terms of precision and recall, as a function of the number of anchor devices in the room. Our attack detector is highly accurate, achieving a recall value of 87.8%, 98.5% and 99.8% with 1, 2, and 3 anchor devices in the room, respectively, and a precision value of 99.95% for all three cases. With only one anchor device, the recall is lower because the user could be further away from the device, thus her movement introduces less observable impact on the sniffed CSI signal. With more anchor devices in the room, the attack coverage increases quickly.

Tracking User Movement across Rooms. Our attack can also track user movements across rooms. To study its efficacy, we first did a controlled experiment where we have two connecting rooms (1 & 2), each with two anchor devices. The sniffer is placed outside behind room 1. We let a human user walk back and forth between the two rooms. Figure 9 shows the detected room occupancy of the two rooms, indicating that our detection is highly responsive to human movements.

Next, using all the recorded CSI traces, we analyze the tracking accuracy by comparing the duration of each detected movement to the ground truth value recorded by the users. Figure 10 plots the CDF of the duration estimation error, where for 80% of the cases, the error is less than 16 seconds.

Trigger Distances of WiFi Devices. As mentioned in §5, each anchor device also has a trigger distance. The closer the user is to an anchor, the more impact she creates on its signal propagation (to the sniffer). To study this effect, we perform controlled experiments in four of our test scenes. Here the sniffer is placed behind a wall 10m away from all the anchor devices. We break the movement pattern into two groups: moving across the direct link from an anchor to the sniffer, and moving on the side of the anchor that affects its reflection paths to the sniffer. Our results show that the first type of movement triggers more signal variations. Overall, the trigger range is about 3m (for 87.8% accuracy). At 5m, the accuracy drops to 40%.

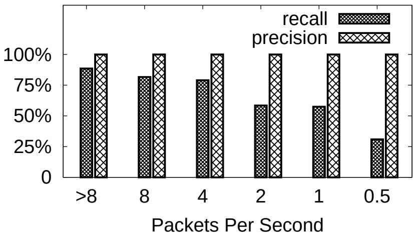

Figure 11: Anchor devices operating at lower packet rates provide less accurate user presence detection (same precision but lower recall).

Impact of Anchor Transmit Rate. Our above results assume that the anchor devices are in the active mode. The reported CSI packet rate is between 8-11pps (due to the firmware limitation discussed in §6). To study the impact of anchor packet rate below 8pps, we sub-sample the CSI traces to emulate low packet rates. Figure 11 shows the detection recall and precision as a function of the packet rate of a WiFi security camera. At its full rate (equivalent CSI rate of 11pps), the recall value is 88.5%, which reduces to 58.4% at 2pps, and 31% at 0.5pps. But the precision is constant at 99.94%. This means that certain WiFi devices, when idle, cannot be used *alone* for user presence detection. But since devices transmit packets at different times, the attacker can aggregate results from multiple anchors to boost detection accuracy.

7.4 End-to-End Attack Evaluation

Finally, we evaluate the end-to-end performance of our attack. Since the goal of our attack is to recognize and track human user's presence and movement in the target area, we consider two end-to-end performance metrics.

Misdetection & Falsealarm. Misdetection refers to cases where the adversary fails to detect a user's presence in a room, either because step 1 mis-locates all the anchor

		# of WiFi Devices Per Room			
		1	2	3	4
Per Room	Misdetection	20.24%	4.21%	0.88%	0.19%
	Falsealarm	7.79%	14.80%	21.30%	27.49%
Per Area	Misdetection	12.16%	1.47%	0.00%	0.00%
	Falsealarm	0.05%	0.05%	0.05%	0.05%

Table 3: End-to-end performance of our attack

devices to a different room or step 2 fails to detect the user's presence. Falsealarm occurs when the adversary falsely reports a nonexistent user presence. This can be the result of anchor mis-location in the step 1 attack, or a false positive event in the step 2 user detection.

Table 3 lists the misdetection and falsealarm rates across all of our experiments, calculated for each individual room in the test scene (per room result) and when treating the entire target area as a single room (per area result). We also vary the number of WiFi devices per room to examine its impact on the attack effectiveness.

We see that with more than 2 WiFi devices in a regular room, our attack can detect more than 99% of the user presence and movement in each room we have tested. The cost is a higher falsealarm rate because the probability of mislocating a WiFi device to a different room also increases. Here a potential improvement to our attack is to intelligently select a subset of high-fidelity anchors for movement detection. We leave this optimization as future work.

On the other hand, if one can "relax" the requirement of detecting activity in each individual room to detecting in the target area, then our attack can detect all the activities while maintaining a very low falsealarm rate of 0.05%.

8 DEFENSES

Having demonstrated the effectiveness of these location attacks, we now discuss robust defenses against them. Our key insight for developing defenses is that the effectiveness of the attack depends heavily on both the quantity and quality of the WiFi signals captured by the sniffer. Thus a defense that reduces the amount of WiFi signal leakage to external sniffers or adds inconsistency to WiFi signals could render the attack ineffective.

Why MAC Randomization Fails? An immediate candidate would be *MAC address randomization*, a well-known method for protecting mobile devices from being tracked. Since the attack sniffer uses MAC address to isolate RSS measurements for each anchor device, MAC randomization can break this isolation and disrupt the proposed localization and tracking. However, recent works have shown that the MAC randomization feature is disabled on most devices (less than 3% of adoption rate so far) [30] and can be easily broken to reveal the real MAC address [2, 29]. Thus Android 9.0 Pie switches to per-network MAC randomization [4], where the static WiFi devices do not apply any MAC randomization.

Next, we describe three alternative defenses for reducing the quantity and/or quality of sniffable WiFi signals. We experimentally evaluate their effectiveness against the attack and discuss the strengths and limitations of each defense.

8.1 Geofencing WiFi Signals

Geofencing creates a geographical boundary for WiFi signal propagation, so that WiFi signals can only reach limited areas accessible to the adversary. This effectively reduces the number of packets captured by the adversarial sniffer. While our experiments in §7 were based on walking traces of 25-50 meters each, we found that when shrinking each walking trace to 10 meters or less, the localization error increased significantly. The raw errors more than doubled, and the room-level accuracy droped from 92.6% to 41.15%.

Practical Implications. Geofencing, once deployed, can be very effective against adversarial sensing attacks. But in practice, geofencing is extremely difficult to deploy and configure. The simplest form is reducing the transmit power of the WiFi devices, which is generally infeasible since it degrades the connectivity. A similar approach is equipping WiFi devices with directional antennas, limiting RF emission in the spatial domain. Requiring higher cost and larger form factor, this approach is not applicable to commodity WiFi/IoT devices. The extreme solution is to paint (boundary) walls with electromagnetic field shielding paint, preventing any RF signal from propagating beyond these walls. Doing so also blocks cellular signals, thus is undesirable in practice.

A practical alternative is to customize WiFi signal coverage using 3D fabricated reflectors, proposed recently by [56]. Yet it faces considerable complexity and limited applicability since the reflector configuration depends on the WiFi device placements and the environment settings.

8.2 WiFi Rate Limiting

While geofencing reduces the spatial leakage of WiFi signals, rate limiting reduces the temporal volume. Now each WiFi device transmits less signals over time, the signals captured by the sniffer can become insufficient to execute proper model fitting or signal variation detection. As shown in §7, rate limiting is effective against Step 2 attack that detects user presence/movements by signal variations (Figure 11), but ineffective against Step 1 attack.

Practical Implications. Rate limiting is simple to implement but introduces undesirable artifacts to applications. As shown by Table 2, many of WiFi devices used in offices and

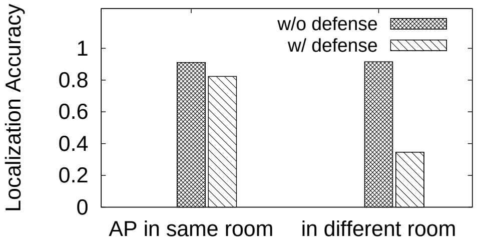

Figure 12: The localization accuracy (room-level) with and without our signal obfuscation defense.

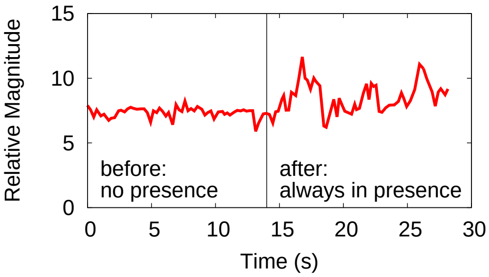

Figure 13: With our defense, the sniffer will detect constant user presence when no one is present.

homes are IoT devices, even during idle, transmit packets beyond 2pps. Thus in many cases, it is impractical to use rate limiting as a continuous defense.

8.3 Signal Obfuscation

Our third defense is to add noises to WiFi signals, so the adversary cannot accurately localize the anchor devices or detect user movements. We refer to this to as *signal obfuscation*. Compared to geofencing and rate limiting, this defense imposes minimum impact on anchor WiFi devices.

Signal obfuscation can take place in both temporal and spatial forms. In temporal obfuscation, WiFi devices change their transmit power (randomly) over time, injecting artificial noises to signals seen by the sniffer. But recent work [27] shows that the adversary can counter this defense by deploying an extra static sniffer to infer the injected signal power changes and remove them from the signal traces. In spatial obfuscation, two WiFi devices transmit via a single MAC address. Since signals come from two physically separated transmitters, the sniffer cannot make accurate prediction of its location. However, this requires tight synchronization and active coordination across the devices and it is still possible for the sniffer to separate the two data streams.

AP-based Signal Obfuscation. In this paper, we propose a practical defense that integrates both temporal and spatial obfuscation. The key idea is to let the WiFi access point (AP) transmit packets upon receiving packets from an associated WiFi device w. Specifically, the AP pretends as w and transmits packets (to itself) at a randomly chosen time and power level. This secretly inserts random noise into the signal traces of w at the sniffer.

The process of inserting "fake" packets requires a careful design, so that it disrupts the attack but creates minimal impact on the WiFi network. We found the following strategy works well under our current experiment scenarios (and leave the systematic design to future work). Upon receiving a packet from any anchor w, with a probability of 20% the AP transmits a random number of packets (between 1 and 20), and adjusts its transmit power every 100 packets. The sequence numbers of the fake packets (partially) overlap with those of (past and future) packets from w, so that the attacker is unable to separate the two packet streams based on sequence number and packet arrival time.

With this defense, the RSS trace of w will display fluctuations, tricking the adversary to think that w is moving and cannot be used as an anchor device. Even if the adversary treats w as stationary, the use of the spatial obfuscation will lead to an inaccurate localization result. It is possible for Monte Carlo sampling (§4) to extract "clean" measurements of w, but the probability is extremely low. Finally, the sniffed CSI traces will contain sufficient signal variations across each sub-carrier, indicating that a user is always present.

We experimented with this defense for cases where the WiFi devices are in the same or different room as the AP. Figure 12 plots the Step 1 anchor localization accuracy for both cases with and without the proposed defense. For both cases, the adversary also deploys an extra stationary sniffer, and applies the same signal subtraction method in [27] to attempt at removing "injected" signal variations in RSS. We see that despite the countermeasure deployed by the adversary, the proposed defense lowers the accuracy of anchor localization. The degradation is particularly visible (from 90% to 38%) for anchor devices not in the same room as the AP.

Figure 13 plots the impact on the Step 2 attack, in terms of a single sub-carrier's CSI amplitude trace before and after applying the defense, when no one is in the room. We see that the defense produces signal variations and confuses the attacker to detect constant human presence.

Practical Implications. The strength of this defense is that it can be quickly deployed by today's WiFi APs that support transmit power adaptation on the fly. No firmware or hardware changes are needed for individual WiFi devices. The major drawback is the extra consumption of WiFi bandwidth and energy at the AP. As future work, we plan to develop more efficient AP obfuscation strategies.

The above defense can be further enhanced by making WiFi device randomly adapt its transmit power at the same time, adding more randomness and inconsistency to the signal traces. This, however, requires commodity WiFi devices to adapt their transmit power on the fly.

9 RELATED WORK

Location Privacy. With geolocation present in more than 90% of apps installed on smartphones [3], we are particularly vulnerable to attacks that reveal our private locations. Whether it is compromising service providers [36], or hacking into social networks [18] or smartphone sensors and power meters [31, 33], existing works have identified a wide variety of attacks on location privacy and subsequent defenses [10, 11, 19, 34, 36, 43, 55].

Our work targets a new type of location privacy attack, where the attacker tracks presence and movement of unsuspecting targets by monitoring ambient WiFi transmissions outside of their building. This new attack does not require access to any services and devices, transmitting any signals, and users carrying any devices.

Privacy Invasion from Traffic Analysis. User presence and activity can change traffic patterns of some WiFi devices, *e.g.*, cameras with motion detection transmit more packets when an active object is present [12]. Prior works use traffic patterns of sniffed signals to infer user status [24, 39, 60]. This approach requires accurate identification of each device, knowledge of their transmission behaviors, and can be easily countered by adapting transmission behaviors. Our proposed attack does not make any of these strong assumptions.

Privacy Invasion from Signal Sniffing. Similar to our attack, existing works seek to locate devices and infer user activities (based on the located room type) using either sniffed WiFi and ZigBee signals [9, 27] or acoustic signals [32]. Focusing solely on locating WiFi cameras, [27] applies feature clustering to identify good measurement rounds. Our work proposes a much more effective method for identifying good data, and targets an advanced goal of inferring the presence and movement of users who do not carry any WiFi device.

[32] detects user presence from specially crafted acoustic signals that are transmitted by devices in the home. This requires remote access to these devices, a strong assumption in practice. Relying on unique properties of acoustic signals, this attack only works on devices equipped with speakers and microphones. Instead, our attack leverages ambient WiFi signals emitted by devices in homes and offices and does not require any access to these devices.

[9] deploys multiple laptops (each with three antennas) outside of a user's home to detect her movements. It makes a very strong assumption where a single (known) router is placed in the center of the home and actively communicates with the attacker laptops. This work can only detect a very limited set of user movements that directly block communications between the router and the laptop. Instead, our work uses a single smartphone sniffer with a single antenna, who *passively* monitors ambient transmissions of *many* WiFi devices. Our design does not make any assumption on device placements or user movements/activities.

Human Activity Detection. Existing works correlate human activities with wireless signal variations extracted from Doppler shift [35], AoA [7], CSI [48, 52, 53, 58], and even RSS [45, 59]. Our work builds on these existing efforts, but differs in two key aspects.

First, existing works are non-adversarial and focus on applying machine learning techniques to map each observed signal pattern to a predefined human activity. These mapping are in general environment specific. Under our attack scenario, the adversary has no prior knowledge of the transmitter, the target, or the ground truth reference on the activities and signal patterns. Our goal is to detect the presence and movement of any target, rather than recognizing their specific activities. Second, our attack uses CSI amplitudes reported by portable commodity WiFi sniffers. Existing works either require specialized, bulky hardware [7, 35], and/or active communications with WiFi devices [48, 52, 53, 58]. [7] also requires custom-made transmitters that emit carefully crafted signals for localization/tracking. All are infeasible under our attack scenario.

Transmitter Localization. Our step 1 attack builds on existing works on RF localization. Our contribution is not to invent a new localization algorithm (we reuse RSS model fitting as our base), but to develop a novel data sifting method to identify suitable RSS data for localization.

Our localization design uses spatial measurements of RSS rather than other advanced signal metrics like AoA [21, 46], phase component of CSI (Phase-CSI), or time of flight (ToF) [6, 20, 37]. These advanced metrics cannot be measured by commodity passive WiFi sniffers. Phase-CSI and ToF require active communications and tight synchronization with the transmitter, while AoA requires large antenna arrays. Recent work [23] lowers the antenna count to 3, but requires at least two line-of-sight paths between the transmitter and receiver. Under our attack scenario, this is clearly impossible. Another work [21] requires the sniffer and targets to be in the same large open space (no walls), while our attack is "behind-the-wall." Finally, radar-based localization [7] transmits crafted signals towards an object and uses reflection signals to infer its location. It requires transmissions by the attacker sniffer which is false under our scenario.

Defense against RF Eavesdropping. Existing works [15, 22, 38, 51] defend against eavesdropping on a RF transmitter by a jammer transmitting simultaneously, preventing the attacker to decode packets. This requires precise synchronization between the two devices or a high-end full-duplex obfuscator. Our defense differs by using the AP to insert fake packets to obfuscate sniffed signals (rather than transmitting simultaneously). Our goal is to disrupt adversarial localization rather than disabling packet decoding.

10 CONCLUSION

Our work brings up an inconvenient truth about wireless transmissions. While greatly improving our everyday life, they also unknowingly reveal information about ourselves and our actions. By designing a simple and powerful attack, we show that bad actors outside of a building can secretly track user presence and movement inside the building by just passively listening to ambient WiFi transmissions (even if they are encrypted). To defend against these attacks, we must control the volume and coverage of WiFi signals, or ask APs to obfuscate signals using cover traffic.

While our attack targets WiFi localization and tracking, our methodology can be generalized to sensing mechanisms at different RF frequencies (*e.g.* UHF, cellular, millimeter wave [63]) and other mediums (acoustic [32], ultrasound [47, 61], visible light, magnetics). Beyond this single attack, we hope to highlight largely overlooked privacy risks from ambient RF (and other) signals around us.

REFERENCES

[1] Howto estimate parameter-errors using monte carlo. http://www-personal.umd.umich.edu/~wiclarks/AstroLab/HOWTOs/NotebookStuff/MonteCarloHOWTO.html, 2014.

[2] Researchers break mac address randomization and track 100% of test devices. https://www.bleepingcomputer.com/news/security/researchers-break-mac-address-randomization-and-track-100-percent-of-test-devices/, 2017.

[3] 7 hot ideas for location-based apps. https://gbksoft.com/blog/5-hot-ideas-for-location-based-apps/, 2018.

[4] Android p feature spotlight: Per-network mac address randomization added as experimental feature. https://www.androidpolice.com/2018/03/08/android-p-feature-spotlight-per-network-mac-address-randomization-added-experimental-feature/, 2018.

[5] Ettus research products. https://www.ettus.com/product/, 2018.

[6] Wifi indoor positioning. https://fit-iot.com/web/wifi-indoor-positioning/, 2018.

[7] ADIB, F., AND KATABI, D. See through walls with wifi! In *Proc. of SIGCOMM* (2013).

[8] ADIB, F., MAO, H., KABELAC, Z., KATABI, D., AND MILLER, R. C. Smart homes that monitor breathing and heart rate. In *Proc. of CHI* (2015).

[9] BANERJEE, A., MAAS, D., BOCCA, M., PATWARI, N., AND KASERA, S. Violating privacy through walls by passive monitoring of radio windows. In *Proc. of WiSec* (2014).

[10] BINDSCHAEDLER, V., AND SHOKRI, R. Synthesizing plausible privacy-preserving location traces. In *Proc. of SP* (2016).

[11] BORDENABE, N. E., CHATZIKOKOLAKIS, K., AND PALAMIDESSI, C. Optimal geo-indistinguishable mechanisms for location privacy. In *Proc. of CCS* (2014).

[12] CHENG, Y., JI, X., LU, T., AND XU, W. Dewicam: Detecting hidden wireless cameras via smartphones. In *Proc. of Asia CCS* (2018), ACM.

[13] CHENG, Y.-C., CHAWATHE, Y., LAMARCA, A., AND KRUMM, J. Accuracy characterization for metropolitan-scale wi-fi localization. In *Proc. of MobiSys* (2005).

[14] EVENNOU, F., AND MARX, F. Advanced integration of wifi and inertial navigation systems for indoor mobile positioning. *EURASIP J. Appl. Signal Process 2006* (2006).

[15] GOLLAKOTA, S., AND KATABI, D. ijam: Jamming oneself for secure wireless communication. Tech. rep., Computer Science and Artificial Intelligence Laboratory Technical Report, 2010.

[16] HALPERIN, D., HU, W., SHETH, A., AND WETHERALL, D. Tool release: Gathering 802.11n traces with channel state information. *ACM SIGCOMM CCR 41*, 1 (2011).

[17] HAN, D., ANDERSEN, D. G., KAMINSKY, M., PAPAGIANNAKI, K., AND SESHAN, S. Access point localization using local signal strength gradient. In *Proc. of PAM* (2009).

[18] HASSAN, W. U., HUSSAIN, S., AND BATES, A. Analysis of privacy protections in fitness tracking social networks-or-you can run, but can you hide? In *Proc. of USENIX Security* (2018).

[19] JIN, X., ZHANG, R., CHEN, Y., LI, T., AND ZHANG, Y. Dpsense: Differentially private crowdsourced spectrum sensing. In *Proc. of CCS* (2016).

[20] JOSHI, K., BHARADIA, D., KOTARU, M., AND KATTI, S. Wideo: Fine-grained device-free motion tracing using rf backscatter. In *Proc. of NSDI* (2015).

[21] KARANAM, C. R., KORANY, B., AND MOSTOFI, Y. Magnitude-based angle-of-arrival estimation, localization, and target tracking. In *Proc. of IPSN* (2018).

[22] KIM, Y. S., TAGUE, P., LEE, H., AND KIM, H. Carving secure wi-fi zones with defensive jamming. In *Proc. of Asia CCS* (2012).

[23] Kotaru, M., Joshi, K., Bharadia, D., and Katti, S. Spotfi: Decimeter level localization using wifi. In *Proc. of SIGCOMM* (2015).

[24] Li, H., He, Y., Sun, L., Cheng, X., and Yu, J. Side-channel information leakage of encrypted video stream in video surveillance systems. In *Proc. of INFOCOM* (2016).

[25] Li, L., Shen, G., Zhao, C., Moscibroda, T., Lin, J.-H., and Zhao, F. Experiencing and handling the diversity in data density and environmental locality in an indoor positioning service. In *Proc. of MobiCom* (2014).

[26] Li, X., Li, S., Zhang, D., Xiong, J., Wang, Y., and Mei, H. Dynamic-music: Accurate device-free indoor localization. In *Proc. of UbiComp* (2016).

[27] Li, Z., Xiao, Z., Zhu, Y., Pattarachanyakul, I., Zhao, B. Y., and Zheng, H. Adversarial localization against wireless cameras. In *Proc. of HotMobile* (2018).

[28] Mare, S., Sorber, J., Shin, M., Cornelius, C., and Kotz, D. Adapt-lite: Privacy-aware, secure, and efficient mhealth sensing. In *Proc. of WPES* (2011).

[29] Martin, J., Mayberry, T., Donahue, C., Foppe, L., Brown, L., Riggins, C., Rye, E. C., and Brown, D. A study of MAC address randomization in mobile devices and when it fails. *CoRR abs/1703.02874* (2017).

[30] Matte, C., and Cunche, M. Spread of mac address randomization studied using locally administered mac addresses use historic. *RR-9142, Inria Grenoble Rhône-Alpes* (2017).

[31] Michalevsky, Y., Schulman, A., Veerapandian, G. A., Boneh, D., and Nakibly, G. Powerspy: Location tracking using mobile device power analysis. In *Proc. of USENIX Security* (2015).

[32] Nandakumar, R., Takakuwa, A., Kohno, T., and Gollakota, S. Covertband: Activity information leakage using music. In *Proc. of UbiComp* (2017).

[33] Narain, S., Vo-Huu, T. D., Block, K., and Noubir, G. Inferring user routes and locations using zero-permission mobile sensors. In *Proc. of SP* (2016).

[34] Oya, S., Troncoso, C., and Pérez-González, F. Back to the drawing board: Revisiting the design of optimal location privacy-preserving mechanisms. In *Proc. of CCS* (2017).

[35] Pu, Q., Gupta, S., Gollakota, S., and Patel, S. Whole-home gesture recognition using wireless signals. In *Proc. of MobiCom* (2013).

[36] Puttaswamy, K. P., Wang, S., Steinbauer, T., Agrawal, D., El Abbadi, A., Kruegel, C., and Zhao, B. Y. Preserving location privacy in geosocial applications. *IEEE Transactions on Mobile Computing 13*, 1 (2014).

[37] Qian, K., Wu, C., Zhang, Y., Zhang, G., Yang, Z., and Liu, Y. Widar2.0: Passive human tracking with a single wi-fi link. In *Proc. of MobiSys* (2018).

[38] Qiao, Y., Zhang, O., Zhou, W., Srinivasan, K., and Arora, A. Phycloak: Obfuscating sensing from communication signals. In *Proc. of NSDI* (2016).

[39] Sanchez, I., Satta, R., Fovino, I. N., Baldini, G., Steri, G., Shaw, D., and Ciardulli, A. Privacy leakages in smart home wireless technologies. In *Proc. of ICCST* (2014).

[40] Schulz, M., Link, J., Gringoli, F., and Hollick, M. Shadow wi-fi: Teaching smart- phones to transmit raw signals and to extract channel state information to implement practical covert channels over wi-fi. In *Proc. of MobiSys* (2018).

[41] Sen, P. K., and Singer, J. M., Eds. *Large sample methods in statistics*. Chapman & Hall, Inc., 1989.

[42] Seybold, J. *Introduction to Rf Propagation*. Wiley, 2005.

[43] Shokri, R., Theodorakopoulos, G., Troncoso, C., Hubaux, J.-P., and Le Boudec, J.-Y. Protecting location privacy: optimal strategy against localization attacks. In *Proc. of CCS* (2012).

[44] Siby, S., Maiti, R. R., and Tippenhauer, N. O. Iotscanner: Detecting privacy threats in iot neighborhoods. In *Proc. of IoTPTS* (2017).

[45] Sigg, S., Shi, S., Buesching, F., Ji, Y., and Wolf, L. Leveraging rf-channel fluctuation for activity recognition: Active and passive systems, continuous and rssi-based signal features. In *Proc. of MoMM* (2013).

[46] Soltanaghaei, E., Kalyanaraman, A., and Whitehouse, K. Multipath triangulation: Decimeter-level wifi localization and orientation with a single unaided receiver. In *Proc. of MobiSys* (2018).

[47] Song, L., and Mittal, P. Inaudible voice commands. *CoRR abs/1708.07238* (2017).

[48] Srinivasan, V., Stankovic, J., and Whitehouse, K. Protecting your daily in-home activity information from a wireless snooping attack. In *Proc. of UbiComp* (2008).

[49] Tan, S., and Yang, J. Wifinger: Leveraging commodity wifi for fine-grained finger gesture recognition. In *Proc. of MobiHoc* (2016).

[50] Tsai, M. Path-loss and shadowing (large-scale fading). Tech. rep., 2011.

[51] Wang, T., Liu, Y., Pei, Q., and Hou, T. Location-restricted services access control leveraging pinpoint waveforming. In *Proc. of CCS* (2015).

[52] Wang, W., Liu, A. X., Shahzad, M., Ling, K., and Lu, S. Understanding and modeling of wifi signal based human activity recognition. In *Proc. of MobiCom* (2015).

[53] Wang, Y., Liu, J., Chen, Y., Gruteser, M., Yang, J., and Liu, H. E-eyes: Device-free location-oriented activity identification using fine-grained wifi signatures. In *Proc. of MobiCom* (2014).

[54] Wei, T., Wang, S., Zhou, A., and Zhang, X. Acoustic eavesdropping through wireless vibrometry. In *Proc. of MobiCom* (2015).

[55] Xiao, Y., and Xiong, L. Protecting locations with differential privacy under temporal correlations. In *Proc. of CCS* (2015).

[56] Xiong, X., Chan, J., Yu, E., Kumari, N., Sani, A. A., Zheng, C., and Zhou, X. Customizing indoor wireless coverage via 3d-fabricated reflectors. In *Proc. of BuildSys* (2017).

[57] Yedavalli, K., Krishnamachari, B., Ravula, S., and Srinivasan, B. Ecolocation: a sequence based technique for rf localization in wireless sensor networks. In *Proc. of IPSN* (2005).

[58] Yousefi, S., Narui, H., Dayal, S., Ermon, S., and Valaee, S. A survey on behavior recognition using wifi channel state information. *IEEE Communications Magazine 55* (2017).

[59] Youssef, M., Mah, M., and Agrawala, A. Challenges: device-free passive localization for wireless environments. In *Proc. of MobiCom* (2007).

[60] Zhang, F., He, W., Liu, X., and Bridges, P. G. Inferring users' online activities through traffic analysis. In *Proc. of WiSec* (2011).

[61] Zhang, G., Yan, C., Ji, X., Zhang, T., Zhang, T., and Xu, W. Dolphinattack: Inaudible voice commands. In *Proc. of CCS* (2017).

[62] Zhao, M., Adib, F., and Katabi, D. Emotion recognition using wireless signals. In *Proc. of MobiCom* (2016).

[63] Zhu, Y., Yao, Y., Zhao, B. Y., and Zheng, H. Object recognition and navigation using a single networking device. In *Proc. of Mobisys* (2017).

[64] Zhu, Y., Zhu, Y., Zhao, B. Y., and Zheng, H. Reusing 60GHz radios for mobile radar imaging. In *Proc. of MobiCom* (2015).

11 APPENDIX

11.1 Details on RSS Model Fitting

Our RSS model fitting uses the log distance path loss model, which is shown to be robust in indoor environments [25]. This model captures the relation between the RSS P_i and the sniffer's distance d_i to a WiFi transmitting device (TX) when the attacker sniffer is at a location index i:

$$\begin{aligned} P_i &= (P_{TX} - P_{REF}) - 10\gamma \log_{10}(d_i/d_{REF}) + noise \\ &= P_{TX^o} - 10\gamma \log_{10} d_i + noise \end{aligned} \tag{1}$$

where γ is the path loss component, P_{TX} is the transmit power of the target device TX, P_{REF} is its reference power received at distance d_{REF}, and $P_{TX^o} = P_{TX} - P_{REF} + 10\gamma \log_{10} d_{REF}$. When the attacker detects that the sniffer and the target device TX are on the same floor level (see §4.2), we can approximate d_i by

$$d_i \approx \sqrt{(x_i - x_{TX})^2 + (y_i - y_{TX})^2}$$

where xs and ys are 2D coordinates. If TX is detected to be on a different floor,

$$d_i \approx \sqrt{(x_i - x_{TX})^2 + (y_i - y_{TX})^2 + (z - z_{TX})^2}$$

where z and z_{TX} are vertical heights of the sniffer and the target TX. The attacker will pre-calculate $z - z_{TX}$ using our proposed method in §4.2.

The goal of RSS modeling fitting is to estimate (x_{TX}, y_{TX}) as well as (γ, P_{TX^o}), using spatial measurement of RSS values $\{P_i\}$. The corresponding model fitting is formulated into a least square optimization problem:

$$\begin{aligned} \underset{\hat{x}_{TX}, \hat{y}_{TX}, \hat{P}_{TX^o}, \hat{\gamma}}{\text{minimize}} \quad & \sum_i (P_i - \hat{P}_i)^2, \\ \text{subject to} \quad & (\hat{x}_{TX}, \hat{y}_{TX}) \in \text{Candidate area}, \\ & \hat{P}_{TX^o} \leq 30dB, \\ & \hat{\gamma} \in [2, 6] \end{aligned} \tag{2}$$

The constraint on $\hat{\gamma}$ follows the well-known observations from empirical measurements [50] while the value of $\hat{P}_{TX^o}$ is upper bounded by the maximum transmit power for WiF frequency defined by the FCC.

For the task of anchor location in the Step 1 attack, we also experimented with other types of propagation models. Among them, only a complicated ray-tracing model accounting the floor plan of the target building [56] achieves a marginal gain over the above log distance model. Given its high complexity and computation cost, we did not include it in the final attack. Resourceful attackers can further improve the localization by switching to more sophisticated models.