

Adversarial WiFi Sensing using a Single Smartphone

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. Many of these invisible transmissions reflect off our bodies, carrying off information about our location, movement, and other physiological properties. While a boon to 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 ambient WiFi signals around us pose real risks to our personal privacy. We identify a passive adversarial sensing attack, where bad actors using a single smartphone can silently localize and track individuals in their home or office from outside walls, by just listening to ambient WiFi signals. We experimentally validate this attack in 11 real-world locations, and show user tracking with high accuracy. Finally, we propose and evaluate defenses including geo-fencing, rate limiting, and signal obfuscation by WiFi access points.

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 of many frequencies, from mid-frequency signals like cellular and WiFi, to very high frequencies in the millimeterwave range. 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 [42, 47, 87, 70, 71, 40].

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. Could we be continuously leaking information about our locations and movements, even when we are not carrying any (trackable) devices?

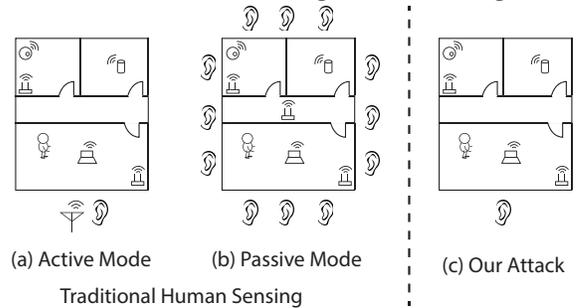


Figure 1: *Traditional human sensing designs either (a) relies on active transmissions by attacker devices, or (b) deploys many sniffer devices. (c) Our attack uses just a single sniffer device (with a single antenna).*

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. Specifically, we consider the malicious task of *human sensing through walls*, where the attacker seeks to localize and track users in their home or office from outside walls, even when the victims do not carry any WiFi devices. Take for example the scenario of thieves looking to break into an office building, either to steal documents or to gain physical access to sensitive data. Being able to identify and track the location of any employees or security personnel gives 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, the attack must function for all human targets, regardless of whether they carry any WiFi devices.

There are a number of technical tools that would enable such an attack in practice (summarized in §9). One straightforward “active sensing” approach is to deploy one of the several existing RF-based human sensing systems (*e.g.* [30, 77, 80]). These techniques require the attacker to actively transmit RF signals and measure the reflections of the target(s) (see Figure 1a)¹. Active RF emissions make it easy to

¹We notice that several existing works [81, 27, 53] have used the term “passive sensing” to refer to device-free sensing scenario, *i.e.* where the target users do not carry any RF devices, but the sensing device is actively

detect and locate the attacker. A stealthy alternative would be to use passive WiFi radar (e.g. [18]), where multiple synchronized WiFi receivers (equipped with directional antennas) coordinate to listen to ambient WiFi signals and extract their doppler shifts to locate human users. This attack can also be done by replacing the group of WiFi radar devices with laptops (each with three antennas) densely placed around the target area [13] (Figure 1b). Finally, some localization techniques rely on fingerprinting their targets ahead of time [51, 70, 12, 20], which clearly does not apply here.

Yet none of these systems can help us answer the key question of interest: *how much information are we already leaking from ambient WiFi reflections, and how easy is it for an attacker to learn about our location and behaviors?* User location and tracking using active transmissions is a well understood problem. We also know that given enough passive listeners, they can cooperate to detect movement in a target area. But what about a *single, passive* attacker device without carefully tuned specialized equipment or help from lots of cooperating devices? *What can it learn about us by just passively monitoring the ambient RF signals and reflections around us?*

We believe that passive user localization and tracking by a single attacker is already possible using today’s commodity smartphones. Our approach to *adversarial user sensing* leverages two recent developments in WiFi networks:

- Our homes and offices are filled with ambient WiFi transmissions from near-ubiquitous deployment of WiFi IoT devices like security cameras, home assistants, media centers and access points (APs).
- Recent work [3, 57] shows that smartphones with Broadcom WiFi chipsets² can accurately capture detailed propagation behavior of any ambient WiFi signals, in the form of the amplitude of channel state information (A-CSI). A-CSI provides a microscopic view of the signal fluctuation caused by human movements, enabling their detection. In the past, CSI could only be captured by a device who actively communicates with a carefully configured target transmitter [1, 74].

Turning Ambient WiFi Signals into Tripwires. Leveraging these two factors, we develop a passive human sensing attack by placing a single commodity smartphone (as a sniffer) outside of the victim’s house/office, which just listens to ambient WiFi traffic (Figure 1c). The key driver of our attack is a novel A-CSI model that converts ambient WiFi signals into *invisible tripwires* radiating around WiFi devices, which silently monitor users in their own home/office. When a user makes a move (e.g. sitting down, waving her hands, opening/closing a door, walking) near a WiFi device x , our sniffer

transmitting signals. We consider these to be “active sensing,” since the sensing device must transmit signals to generate reflection signals.

²Broadcom’s WiFi chipsets are the most common WiFi chipsets used on mobile devices [4].

immediately detects the movement and localizes the victim by observing x ’s A-CSI values. This calculation requires information about the locations of WiFi devices in the victim’s home/office, which our system first estimates from the same ambient signal captured by the same sniffer.

In this paper, we describe our efforts to understand the feasibility, challenges, and defenses surrounding the topic of passive adversarial sensing. In short, our key contributions can be summarized as follows.

- We identify the passive, adversarial human sensing attack using a single smartphone, and design a new A-CSI based technique for adversarial sensing.
- We build a complete prototype of an attacker system, and demonstrate that the attack is not only highly accurate (detecting and localizing users to an individual room with 99% accuracy), but also highly general (effective in 11 different physical settings, including both office buildings and residential apartments).
- We propose and evaluate three potential defenses, geofencing WiFi signals, rate limiting WiFi signals, and signal obfuscation. Our results show that only signal obfuscation by WiFi APs is both practical and effective.

Our work shows that using just a single smartphone as a passive sniffer, an external attacker can already silently localize and track our movements inside our own home/office. While not able to identify fine-grained features like movement speed and type (which require active sensing and/or fingerprinting), our attack does raise an alarming concern on the “vulnerability” of us reflecting ambient RF signals.

Broader Implications. Recent research efforts have aggressively pursued the use of RF sensing as a way to measure and understand ourselves and our environment. Yet few have considered the adversarial aspects of these RF sensing technologies, and what risks they pose to unsuspecting targets. We believe passive human sensing using commodity smartphones is just one of many practical attacks made possible by inventive uses of RF signals. We hope our work brings more attention from the security community to an understudied topic with potentially numerous technical and social challenges.

2 Attack Model

We begin with a definition of our attack model. Our goal is to study a single practical and concrete instance of adversarial sensing attacks. We make basic assumptions about the attacker’s resources and capabilities, in order to determine what a minimally equipped attacker can accomplish. As shown in Figure 1(c), our attacker places a sniffer outside a target building (residential or office), passively sniffs ambient WiFi transmissions, and uses captured signals to locate and track human movements within the building. We make no assumptions about human targets being tracked.

We make the following assumptions about the adversary.

- We assume the adversary does not have physical or remote access to WiFi devices in the target building, thus these WiFi devices are secure. We assume the attacker can physically move outside the target area, either outside exterior walls or along public corridors.
- To avoid detection, the attacker only performs passive WiFi sniffing, and avoids bulky or specialized hardware, *e.g.* directional antenna, antenna array, and USRP [7]. Instead, they prefer commodity smartphones with a single built-in antenna (one for 2.4GHz).
- The adversary has access to rough floorplans of the target building. These are generally publicly available thanks to real estate websites and apps, *e.g.*, Zillow and Redfin.

We intentionally choose a resource-limited attacker to understand the ease with which adversarial sensing attacks can take place. Less resource requirements imply that the attack can be successful in a wider range of adversarial scenarios.

3 Attack Methodology

Before diving into details of our attack, we first present the intuition behind it and an overview of the attack process.

3.1 What We Get from Ambient WiFi Signals

Our attack leverages the ubiquity of ambient WiFi emissions today to locate and track users. Whether it is routers, laptops, media sticks, or new IoT devices like voice assistants, cameras, and smart appliances, WiFi devices reside in every room at our homes and work. They also constantly flood their surroundings with wireless signals³. WiFi sniffers outside the room or building can passively listen to these signals without risking detection. Since WiFi packets do not encrypt source and destination MAC addresses, the sniffer can isolate packets of each transmitting device even under MAC randomization [56, 61, 43] (details in §8). In our attack we refer to these WiFi transmitters as *anchors*.

For each anchor, an attack sniffer can extract two metrics from its “ambient” WiFi signals (even if they are encrypted):

- *Amplitude of CSI (A-CSI)* measures the signal strength on each of the many frequency subcarriers of a transmission. Because human movements change the dynamics of multipath signal propagation between the anchor and the sniffer, A-CSI values observed at the sniffer will also fluctuate over time. This fine-grained feature provides a microscopic view of human movements.
- *Received Signal Strength (RSS)* is the signal strength measured from a transmission, but aggregated over all the frequency subcarriers. Such aggregation makes it relatively insensitive to user movements.

It is important to note that from ambient signals, a passive sniffer with a single antenna is unable to extract advanced

³We list a summary of common WiFi devices seen at homes and offices and their traffic patterns later in Table 1.

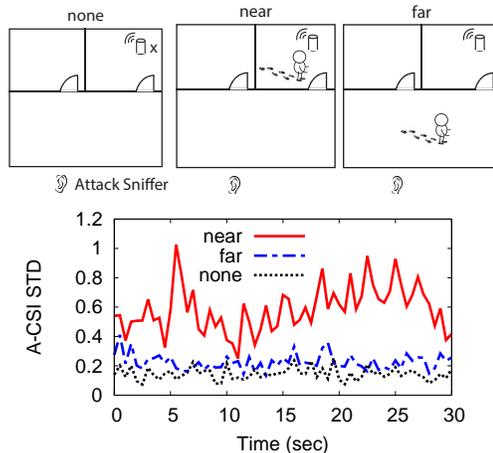


Figure 2: The attack sniffer monitors anchor X 's signal by its A-CSI STD. When a target moves near X , a subset of signal paths from X to the sniffer are affected, leading to larger STD. When the target is away from X , it has less impact on the signal propagation from X to the sniffer, so A-CSI STD is less. When no one is present, A-CSI STD is even smaller.

features like phase of CSI and Angle of Arrival (AoA)⁴ that many existing sensing designs rely on. Estimation of CSI phase requires the sniffer to actively synchronize with the transmitter, while AoA estimation requires the sniffer to have two or more antennas [75]. These physical limitations rule out the feasibility of using most techniques introduced by prior sensing work (*e.g.* [80, 30, 11]).

3.2 Turning Ambient Signals into Tripwires

In an office/home setting, human users are never completely stationary. Whether it is playing games, walking, opening doors, sitting down, or standing up, their natural movements will disturb the signal propagation of nearby WiFi transmitters (anchors), creating variations in their A-CSI values seen at the sniffer. Figure 2 illustrates this trend using two sample traces of A-CSI standard deviation (STD) averaged over multiple sub-carriers, of an anchor x , when a human user is in proximity to x and when there is no human presence.

Furthermore, the degree of disturbance depends heavily on the user's distance to the anchor. When the user is away from x , her activities will produce much less disturbance to the signal propagation from x to the sniffer. Figure 2 shows another sample trace of x 's A-CSI STD when the user is in a different room, and the values are much smaller.

Together, these observations show that a hidden adversarial sniffer can detect human movements *near* any anchor by just monitoring its signals, and locate human movements over time from the temporal sequence of the “triggered” anchors. That is, by capturing A-CSI STD of ambient WiFi signals sent by different anchors, the attacker effectively turns

⁴AoA measures an incoming signal's propagation direction.

these signals to a dense net of *invisible tripwires radiating around each anchor*, and uses them to locate targets.

Our proposed use of A-CSI for detecting and localizing targets differs from existing works on A-CSI based activity detection [46, 70]. The latter matches observed A-CSI patterns to pre-trained ones (*i.e.* fingerprints defined a priori) to recognize specific gestures (also infeasible under our attack scenario), but cannot localize the human target.

3.3 Locating Anchor Devices

To set up these “tripwires,” the attacker needs to know the location of each anchor device. Again, this is achieved using the same, single, passive smartphone with a single antenna.

We take an existing passive device localization design, which estimates the anchor location from RSS observed at various locations. Specifically, the adversary simply holds the sniffer and performs a brief measurement by walking outside the target building, either along a public corridor inside an office building or outside a house. A short walk with the sniffer will capture RSS of all the (observable) anchor devices. It then fits the RSS measurement into a model to estimate the anchor location. Resourceful attackers could even use robots or drones to carry out the measurements.

The key contribution of our work is not in the localization algorithm, but a statistical algorithm to identify proper RSS data samples as input to the localization algorithm. This is because in practice, the adversary has little control on the available walking path and the propagation environment, thus the spatial RSS measurements will contain bias, noise and even human errors. To boost localization accuracy, we develop a consistency-based data sifting algorithm that selects proper RSS samples for localization.

3.4 Overview of the Attack Process

With the above in mind, we now present the high-level overview of the attack (shown in Figure 3).

- **Bootstrapping:** The attacker takes a brief walk around the target building, and sniffs ambient WiFi signals to discover and locate WiFi anchor devices inside the building.
- **Continuous sensing:** The attacker hides the same sniffer at a static location outside the target area. The sniffer continuously monitors ambient WiFi signals, and uses them to locate and track human presence and movements inside. The sniffer also monitors each detected anchor, and any relocation of an anchor will trigger its removal from the anchor list, and possibly another *bootstrapping* phase to (re)locate the anchors.

Next, we describe our design of the *continuous sensing* phase in §4 and the *bootstrapping* phase in §5.

4 Continuous Human Sensing

During continuous sensing, the attacker sniffer seeks to detect, locate and track human presence by converting ambi-

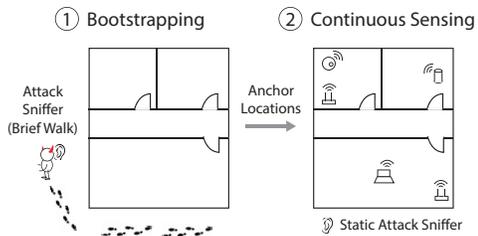


Figure 3: Our attack process includes a bootstrapping phase and a continuous sensing phase.

ent WiFi transmissions into a dense net of tripwires radiating around many anchors in the target area. As discussed in §3.2, the key insight is that human movements near an anchor x increase the A-CSI variance of x ’s signals seen at the sniffer. The key novelty of our work is to recognize and analyze such variation across multiple anchors, and use them to detect, locate human presence to their individual rooms and track their movements over time.

Next we describe our attack design assuming that the attacker has exercised the bootstrapping phase to locate anchor devices, and has a rough floorplan of the target area.

4.1 Detecting Human Movements

Natural user movements will disturb propagation of ambient WiFi signals. Such disturbance can be observed by the attacker sniffer despite being outside of the building. The disturbance is highly visible in terms of A-CSI but not RSS. This is because movements lead to (extra) multi-path fading and cause signal variations in each narrowband sub-carrier of the sniffed signal [69, 70]. RSS is the received signal strength aggregated over all the sub-carriers and the aggregation hides the variations of each sub-carrier.

As examples, Figure 4 plots several 30-second samples of an anchor x ’s A-CSI STD (averaged across the sub-carriers), for scenarios of no human presence, a user sitting down, opening/closing the door, and walking. Compared to the case without any human presence, user movements lead to much higher A-CSI variations.

For our attack, we developed a smartphone app that extracts A-CSI from passively sniffed WiFi signals (details in §6). For signals sniffed from each anchor, we calculate the A-CSI STD per sub-carrier using a 5-second moving window, and then average over all the qualified⁵ sub-carriers. In the rest of the paper, we refer to this value as A-CSI STD.

Threshold for Human Movements. The attacker sniffer monitors, for each anchor, its present value of A-CSI STD to detect human presence. When it goes beyond a threshold $\theta_{presence}$, the attacker declares the presence of human users.

⁵We exclude from our calculation the sub-carriers whose amplitude values remain very high and the standard deviation is less than 0.1. We think this is likely because these sub-carriers operate at very high transmit power, and thus are insensitive to user movement due to saturation.

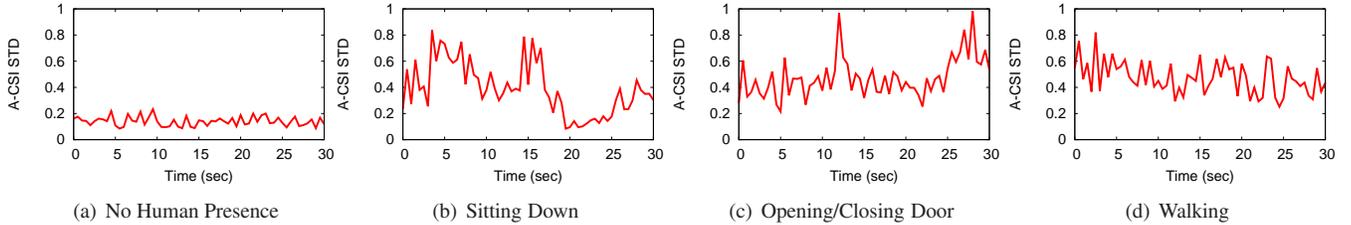


Figure 4: Sample traces of average A-CSI STD for scenarios of no human presence, a user sitting down in a chair, opening/closing the door, and walking.

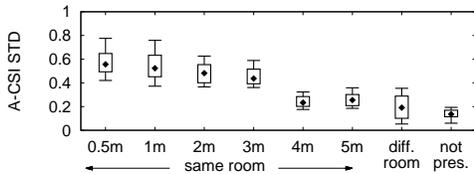


Figure 5: An anchor’s A-CSI STD depends on its distance to the target in the same room, and becomes much smaller when the target is in a different room or is completely absent.

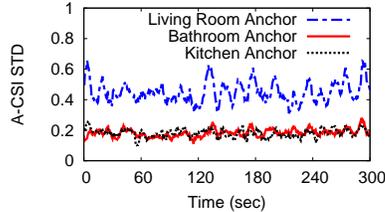


Figure 6: A-CSI STD of anchors in the livingroom, bathroom and kitchen, when the user is in the livingroom (sitting, walking, playing video games).

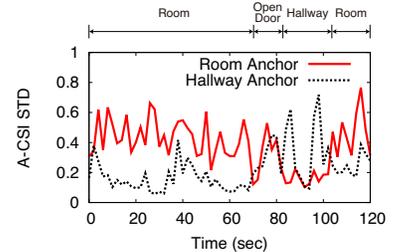


Figure 7: A-CSI STD of anchors in a room and the hallway, where the user walks out to the hallway briefly and walks back to the room.

Choosing $\theta_{presence}$ is similar to finding a threshold for “outliers.” Since human movement is relatively sparse in time and space, anchors are mostly “idle.” Using this intuition, the attacker first records A-CSI STDs of multiple anchors for a period of time (e.g., hours). It then applies a widely used statistical “outlier” measure MAD⁶ [24, 55] to derive $\theta_{presence}$. In our experiments $\theta_{presence} = 0.27$ applies to the sniffers across all test scenes.

4.2 Localizing Targets

After detecting any human presence from A-CSI STD, the attacker also needs to localize the target, e.g. identifying the room she is in. For this we leverage another observation on A-CSI STD. As shown by Figure 2, our hypothesis is that the degree of A-CSI variation would correlate with the distance between the target and the anchor. The closer the target is to an anchor, the more impact she would produce on signals propagated from the anchor to the sniffer.

A-CSI STD vs. Anchor-Target Distance. We build a ray tracing model to explore the correlation between A-CSI observed at the sniffer and the distance between the target and the anchor. The detailed model is in the Appendix. The key intuition is that signals sent by an anchor x will take multiple paths to reach the sniffer; a target is “bigger” when she is closer to x , affecting more signal paths from x to the sniffer. Thus the moving target would create larger temporal varia-

tions in x ’s A-CSI values, leading to a higher A-CSI STD.

We validate our hypothesis empirically using real-world A-CSI measurements of different anchor devices and different test scenes (more details in §7). Figure 5 plots the quantile distribution of the A-CSI STD as a function of the distance between the target and the anchor, when they are in the same room. These results indicate a general tendency of A-CSI STD degrading with the anchor-target distance. Furthermore, we also show the results when no target is present and when the anchor is not in the same room as the target. The STD distribution of the anchors who are close to the target ($< 3m$) is well separated from those of the anchors that are not in the same room with the target.

These observations indicate a general trend that if the number of anchors in a room is sufficiently large (e.g. 4 anchors in a room of common size $25m^2$), any user movement in the room should be “picked up” by at least one anchor in the room who displays large STD values (or peaks). These peaks will be larger than the STD values of the anchors in other rooms. While it is hard to configure a threshold for the peaks (which could be environment specific), we instead identify the right anchor by comparing the STD values of all the anchors in different rooms.

Comparing A-CSI STD across Rooms. Consider the following examples. Figure 6 shows the traces of A-CSI STD for three anchors, located in the living room, bathroom, and kitchen of an apartment, where the target stays in the living room. In this case, the anchor in the living room always has the largest A-CSI STD, thus locating the target to the living room. Another example is Figure 7 with two anchors, one

⁶ We assume the “idle” measurements follow a Gaussian distribution and set MAD scale factor to 1.4826. Then at 99.5% confidence, $\theta_{presence} = 3 * MAD - median(data)$.

in the room and one in the hallway. The target in a room walks towards the hallway, enters the hallway at $t=80s$, and walks back into the room around $t=100s$. In this case, the two rooms are connected so the user movement in the room also triggers the anchor in the hallway. Yet the anchor with higher STD is always in the user’s current room. Furthermore, at $t=80s$, we see the transition of peaks from the room anchor to the hallway anchor, indicating that they are results of movements from the same user.

Assigning Targets to Rooms. With the above in mind, we design the following rules for assigning targets to rooms. Our design assumes that the anchor density in each room is sufficiently high, which the attacker can recognize after running bootstrapping to localize anchors to their rooms.

Case 1: Target in only one room. We start from the simple case where at any given time, at most one room in the target area is occupied, *e.g.*, a single security guard patrolling a company at night, or a user in her apartment. In this case, the attacker first uses $\theta_{presence}$ to identify the set of anchors triggered. Of all the anchors triggered, it picks the one with the largest A-CSI STD and declares that the user is in the room of this anchor.

Formally, let $\theta_x(t)$ be the A-CSI STD of anchor x seen by the sniffer at time t . Our basic rule of room assignment is

$$Room_{human}(t) = Room_{anchor}(\operatorname{argmax}_{x, \theta_x(t) > \theta_{presence}} \theta_x(t)).$$

Case 2: Targets in multiple rooms. We consider more general cases where multiple rooms are occupied. The attacker first calculates for each room i and time t , the room-level A-CSI STD as $\theta^i(t) = \max_{x \in room\ i} \theta_x(t)$. It then applies two rules on the room-level STD traces to determine whether multiple rooms are occupied by different users (rather than a single user moving from one room to another), and to identify the occupied rooms.

- *Temporal Rules:* In general, users in different rooms act asynchronously, producing peak values in each $\theta^i(t)$ in different time segments. At a given time, we will likely see a single peak across all the rooms. Across time, there are no immediate transitions between peaks of different rooms, since they are not triggered by the same user.
- *Spatial Rules:* A single user cannot be in two separate rooms at the same time. At time t , if the attacker does observe multiple STD peaks at different rooms, and these rooms are well separated, *e.g.* with another room in between, then these peaks are caused by different users rather than a single user traveling from one room to another. Similarly, if the floor plan indicates that a user can only travel from room A to room C via room B, and room B’s anchors are not triggered upon detecting subsequent peaks in room A and C’s anchors, then room A and C are occupied by different users.

For cases where the above rules do not apply, we conservatively treat the rooms with large peaks as occupied.

Tracking Targets. Our continuous sensing generates a set of sequential events in terms of (movement time, occupied room), which can be used to generate user trajectories (details omitted for brevity). Since the attacker cannot recognize each specific user, the tracking design works well if the number of users in the target area is small.

4.3 Discussion

Can External Pedestrian Interfere with the Sniffer?

Pedestrians who move outside of the target building near the attack sniffer could also create A-CSI variations, leading to false detection of user presence in the target home/office. Interestingly, such event can be detected because movements near the sniffer will create sudden, simultaneous A-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.

How to Place the Sniffer? The sniffer should be placed where it can capture A-CSI signals from the detected anchors. 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. This helps reduce the chances that user movements in one room trigger anchors in a different room. This is feasible in practice because the attacker has the rough floor plan of the target area.

Limitations. The attacker is unable to recognize a specific user or an activity, *e.g.*, distinguishing between walking and waving hands. Doing so requires extensive knowledge on the activities and A-CSI patterns for each user and anchor pair, which is infeasible under our attack scenario.

5 Bootstrapping: Locating Anchor Devices

During the bootstrapping phase, the attacker uses the same passive sniffer to identify and localize static anchors inside the victim’s home/office. There are many device localization proposals, from those using active communications with the target device [9, 82], those using multiple APs equipped with multiple antennas [75], to those fingerprinting each possible device location a priori [12]. But as the attacker stays passive and only has a single smartphone with a single built-in antenna, RSS-based passive device localization [37, 40] becomes the most feasible candidate. Specifically, with a brief walk outside of the victim’s home/office, the adversary measures RSS of ambient WiFi signals at multiple locations along the trajectory. These spatial RSS values, together with the trajectory, are fed into a RSS-based propagation model to estimate the locations of the anchors.

Finding High-Quality RSS Measurements. When deploying this approach to our attacker system, we found that the localization accuracy depends heavily on the “quality” of the RSS measurements. Ideally, these measurements should contain little noise, align with a propagation model, and

cover a wide range of values to minimize fitting bias. Yet in reality, they contain bias, noises and even human errors, leading to inaccurate localization outcomes. Instead of searching for a new/improved localization design, we focus on designing a statistical data sifting algorithm to identify proper RSS data samples as input to the localization algorithm.

In the following, we present the RSS-based passive device localization in §5.1, our proposed statistical data sifting in §5.2, as well as two enhancements in §5.3 where the attack sniffer identifies static anchors and their floor levels.

5.1 Passive Device Localization

Why RSS but not A-CSI? The localization uses RSS measurements of ambient signals rather than other advanced metrics like A-CSI, CSI or Angle of Arrival (AoA) [63, 31]. This is due to two reasons. *First*, accurate CSI-based localization relies on multiple antennas and the phase component of CSI to derive AoA [39]. Our attacker sniffer only has one antenna, and cannot estimate phase of CSI accurately due to lack of synchronization with the transmitter. Recent work [31] estimates AoA from A-CSI but only if the sniffer and targets are in complete line-of-sight, *i.e.*, no walls. *Second*, as shown in §4, A-CSI is sensitive to nearby target movements. As the adversary has no knowledge of the target status during bootstrapping, it cannot rely on A-CSI for localization. In comparison, RSS is robust against target movements and a more reliable metric for localization.

Localization via RSS Model Fitting. RSS model fitting [37, 40] 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 [59]. The detailed calculation is listed in the Appendix. We also experimented with other passive RF localization methods, including (weighted) centroid [17], gradient [25], and ecolocation [79]. They perform worse and require many more spatial RSS measurements.

RSS model fitting requires the walking trajectory, which can be recorded using 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 per trace and has minimum impact on localization.

5.2 Consistency-based Data Sifting

A straightforward solution is to filter out “bad” measurements using de-noising methods, ranging from the traditional Kalman filter [19], wavelet filter [65] to the newly proposed feature clustering algorithm that remove bad measurement rounds [40]. We found that these methods are insufficient under our attack scenarios because the propagation environment is highly complex and unknown to the adversary, mak-

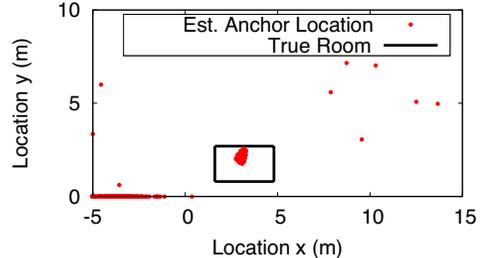


Figure 8: *Localization results from our Monte Carlo sampling. Each red dot is the estimated anchor location from a sample; the rectangle marks the room of the anchor.*

ing it hard to distinguish between noise and natural propagation effect. Features used by [40] 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 [40] will locate the device to a wrong room.

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 [58], *consistent* fitting results from many random sampling of RSS measurements, if exist, can reveal true signal propagation behaviors and produce high-fidelity localization results.

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.

Data Sifting via Monte Carlo Sampling. Given a round of RSS measurements \mathbb{R} , we apply the Monte Carlo method [2] to randomly sample a subset (80%) of \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 clusters among these N results. If they form many small clusters with different room placements, then \mathbb{R} is inconsistent 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 8 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.

Consistency Check across Measurement Rounds. When multiple rounds of sniffing measurements are avail-

able, 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 a room placement accuracy of 92.6% on average (across the 11 test environments).

5.3 Attack Enhancement

We also develop two enhancements that use RSS measurements to check the mobility status (static vs. mobile) of an anchor and its floor level. Our attack only uses static anchors.

Detecting Stationary Anchors. RSS of a stationary transmitter, when captured by a stationary sniffer, should stay relatively stable. When a static device relocates to a different location, its RSS seen by the sniffer will also change.

Before making spatial RSS measurements in bootstrapping, the adversary first places the sniffer statically to record RSS of ambient signals and identify static devices. Later in the continuous sensing phase, the static sniffer also monitors each anchor from its RSS. Upon detecting significant RSS variations for an anchor, the attacker either removes it from the anchor list or run bootstrapping to relocate anchors.

Floor Level Signal Isolation. For buildings with multiple floors, our attack needs to know the floor level of each anchor. Our 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, our commodity sniffer cannot measure AoA because it only has a single antenna.

We show that coarse (vertical) AoA estimation can be achieved by adding a simple, and compact smartphone case to our sniffer, emulating a directional antenna. For our attack, we place a simple cone object of size $8\text{cm} \times 6\text{cm} \times 7\text{cm}$ on top of the smartphone (sniffer) and wrap it with aluminum foils. Now the smartphone sniffer can only capture WiFi signals through the cone. The adversary, standing or sitting in a car, rotates the sniffer while it records the RSS of ambient signals and the phone angle (via the built-in gyroscope). The estimated AoA of a transmitter is the direction that maximizes its RSS. The adversary then infers the floor level by comparing the estimated AoA value to the projected AoA values for different floor levels (derived from 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, while our AoA approximation is sufficient for floor level detection, it is too coarse-grained for anchor localization.

6 Attack Implementation

We prototype our attacker system using a commodity smartphone as the sniffer. We implement the bootstrapping and continuous sensing modules as two Android apps. 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 WiFi firmware from Nexmon [57] to perform passive sniffing. For spatial RSS measurements (during bootstrapping), we use the built-in IMU sensors (accelerometer and gyroscope) 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.

Passive Sniffing of A-CSI. Previously, CSI can only be captured when the WiFi receiver is actively communicating with the transmitter [23]. Our attack leverages a recent development of WiFi firmware [57] to capture A-CSI while operating in the passive sniffing mode.

Our implementation addresses two artifacts in A-CSI measurements caused by the firmware. *First*, the firmware reports each A-CSI 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 A-CSI values at a limited speed, up to 8–11 packets per second. Thus our app subsamples sniffed packets based on this rate limit. Despite these limitations, our prototype sniffer is able to capture sufficient A-CSI samples to successfully launch the attack.

Computation and Energy Cost. One strength of our attack is its simplicity. For our current smartphone prototype, the bootstrapping app runs 1000 rounds of Monte Carlo sampling and model fitting, which finishes in less than 25s per anchor. Our continuous sensing app takes less than 1s to compute average A-CSI standard deviation. These two apps consume 4.18 watts (bootstrapping) and 2.1 watts (continuous sensing), respectively. Using Nexus 5 (2300mAh battery) this enables 4.1 hours of continuous sensing. Currently our apps are not optimized for energy, which we leave to future work.

7 Evaluation

We evaluate our proposed attack using experiments in typical office buildings and apartments. We first describe our experiment setup and scenarios, and then present our evaluation on individual attack phases (bootstrapping and continuous sensing), followed by an end-to-end attack evaluation.

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

	Device Type	Exact Product	Packet Per Second (pps), Idle	Packet Per Second, Active
Static	Cameras (without Motion Detection)	AHD Security Camera	N/A	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	N/A	124*
	Smartphones	Samsung/Google/Apple Phones	≥ 0.5	≥ 6

Table 1: Summary of anchor devices used by our experiments. *Emulated by mounting a camera on a robotic car.

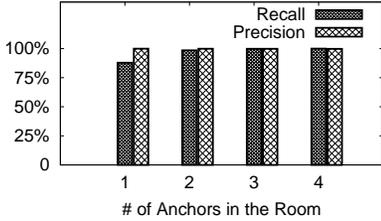


Figure 9: Accuracy of user movement localization, as a function of the number of anchors per room.

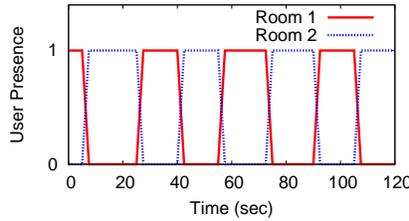


Figure 10: The attacker can track user movement between rooms in real-time

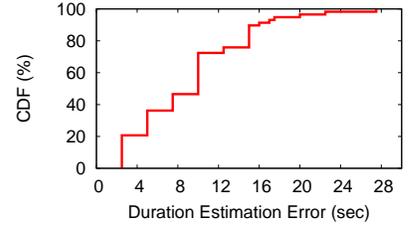


Figure 11: CDF of error in movement duration estimation.

wall materials except for concrete⁷. 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 3 in the Appendix.

Inside each test scene, we either reused existing WiFi devices or deployed our own WiFi devices to generate ambient WiFi signals. 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 1 summarizes these devices and their traffic patterns during idle and active periods. Even when idle, these devices periodically transmit packets. The packet rate varies from 0.5 packet per second (pps) to more than 100 pps. These devices were naturally 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 focus on the 2.4GHz WiFi band due to its dominant coverage. We also tested 5GHz WiFi and did not observe notable difference except its shorter coverage.

For the bootstrapping phase, 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 0.5–2 minutes in time. We also changed the target’s WiFi device placements and repeated the experiments. In total, we collected more than 3000 RSS measurement traces, with more than 121,000 location-RSS tuples.

⁷Our attack does not work when the wall separating the targets and the adversary is made of concrete, which blocks WiFi signals completely.

For the continuous sensing phase, we hid the static sniffer behind plants or at the corners (on the ground) outside of the target building within 1m to the building wall. We asked volunteers to carry out normal activities in each test scene (one moving person per room at any given time), and collected more than 12 hours of CSI entries. The volunteers were aware of the goals and results of the attack but not the specific techniques. We also experimented with different types of target activities and movements.

7.2 Accuracy of Continuous Human Sensing

To test the performance of our continuous sensing phase, we assume that the attacker knows the exact room where each anchor device resides. Our evaluation focuses on answering the following questions on attack effectiveness.

Q1: Can the attacker detect & localize targets?

Our results were based on all the CSI recordings across our test scenes. For each room with at least 1 anchor device, we studied the decision made by the attacker in terms of whether a human user is present in the room or not, and compared this result to the actual human presence. Here we divided time into 5s slots, and for each slot we calculated the average A-CSI STD for each anchor device, and used them to identify the room occupancy in the target area. We then calculated the precision and recall values across all the time slots and rooms. Recall measures the detection rate upon a target’s actual presence in a room, and precision measures the fraction of the actual presence in a room among all detections. To an attacker, a high recall is relative more important since a low recall value means the attacker will miss some human presences in a room.

Figure 9 plots the recall and precision values for rooms with different number of anchor devices. As expected, the

recall value depends on the number of anchor devices per room, 87.8%, 98.5% and 99.8% with 1, 2, and 3 anchor devices in the room, respectively, while the precision remains >99% for all the cases. With only one anchor in a room, the recall is lower, because the user could be further away from the anchor, and her movement introduces less observable impact on the anchor’s A-CSI, leading to possible misses. With more anchor devices in the room, the detection coverage increases quickly.

Q2: Can the attacker track a moving human target?

We first consider cases where a user travels back and forth between two connecting rooms in the building (room 1 and 2) and each room has two anchor devices. Specifically, a user walks in one direction about 25 seconds, turns around to walk in another direction and repeats. Figure 10 shows the detected user occupancy of the two rooms, where our detection is highly responsive to (rapid) human movements.

We also consider all the A-CSI traces and look at the duration of individual movement events estimated by the attacker. We compare these estimations to the ground truth. Figure 11 plots the CDF of the duration estimation error, where for 80% of the cases, the error is less than 16 seconds.

Q3: If an anchor device transmits infrequently, does it “help” the attack?

So far, our results assume that the anchor devices send packets at no less than 11pps⁸. But in reality, certain WiFi devices are often in the idle state, e.g. home voice assistants, and transmit packets infrequently. To study the impact of anchor packet rates, we take the CSI traces of WiFi security cameras (w/o motion detection) and sub-sample them to produce desired packet rates. Our experiments show that when an anchor operates at its full rate (an equivalent CSI rate of 11pps), the recall is 88.5%, which reduces to 58.4% at 2pps, and 31% at 0.5pps. The precision remains >99%. This means that each low-rate anchor still helps the attacker identify and locate targets. When a room consists of multiple low-rate anchors, the attacker should take the *union* of the detection results produced by these anchors to improve the attack recall value.

7.3 Evaluation of Bootstrapping

For bootstrapping (where the attacker locates anchors), we consider two performance metrics: *absolute localization error* which is the physical distance between the ground-truth location and the attacker-estimated location, and *room placement accuracy* which is 1 if the attacker always find the exact room the anchor is at, and 0 if the attacker always places the anchor to a wrong room. Figure 12 plots, for each test scene, the quantile distribution of the absolute localization error and the room placement accuracy.

Q4: Does data sifting improve anchor localization?

⁸ As discussed in §6, the firmware reports CSI in an equivalent packet rate of 8-11pps.

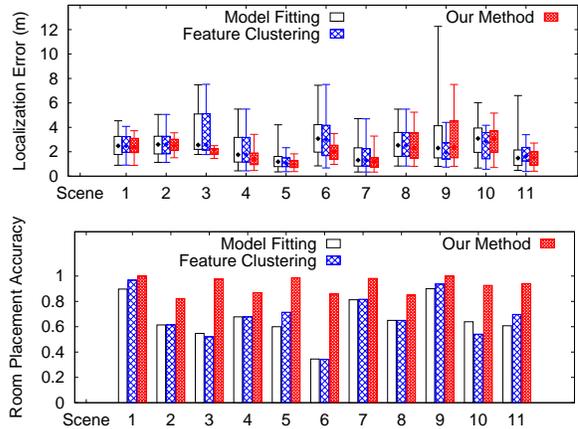


Figure 12: Bootstrapping performance: anchor localization accuracy in terms of absolute localization error (m) and room placement accuracy, for each of the 11 test scenes.

We compare the localization performance of RSS model fitting with and without our proposed data sifting, and when applying feature clustering based data filtering [40]. We make two key observations. First, blindly feeding RSS measurements into model fitting leads to a considerable amount of localization errors and room placement errors. In 5 out of the 11 test scenes, this baseline solution places more than 40% of anchor devices in the wrong room. Second, our data sifting significantly boosts the localization accuracy and room placement accuracy. For more than 90% of the cases, an anchor is placed at the right room.

Using fine-grained data sampling rather than coarse features, our sifting design also outperforms the feature clustering-based filtering [40]. In scene 8, 9, and 10, our design produces similar (and even larger) absolute localization error but higher room placement accuracy. This is because our design directly accounts for the room placement consistency, rather than raw localization errors. Smaller absolute localization error does not always translate into higher room placement accuracy.

Q5: What kinds of anchors are hard to localize?

It is difficult to localize anchors placed at room boundaries, e.g. those directly plugged into wall outlets. These boundary anchors do create a dominant Monte Carlo cluster, but the data points in the cluster map to either of the two neighboring rooms. Currently, we make a simple binary decision by choosing the room with the higher probability, which might not be accurate. As future work, we plan to improve our design by marking these devices as “boundary” anchors and treat them with caution during the continuous sensing phase.

Q6: Are idle anchors hard to localize?

The localization performance is insensitive to the device type and transmission rate. All the 31 devices we have tested always transmit packets at 0.5pps and above. The RSS measurements are relatively time insensitive and thus can be ag-

gregated 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 (and room placement) accuracy.

7.4 End-to-End Attack Evaluation

Finally, we evaluate the end-to-end performance of our attack, combining both bootstrapping and continuous sensing phases. Since the goal of our attack is to recognize and track human user’s presence and movement, we again use recall and precision as the key performance metrics. An effective attack should have both high recall and high precision values, indicating a high rate of detection and low rate of false alarms. Lower values in these two metrics can be the result of misplaced anchors during bootstrapping, or errors in localizing users during continuous sensing.

Table 2 lists the precision and recall values for detecting and localizing human users to their individual rooms (per room). We also vary the number of WiFi devices per room to examine its impact on the success rate of the attack.

		# of WiFi Devices Per Room			
		1	2	3	4
Per Room	Recall	81.67%	96.65%	99.39%	99.89%
	Precision	91.84%	87.37%	83.10%	79.35%
Per Area	Recall	87.84%	98.53%	99.82%	99.98%
	Precision	99.93%	99.88%	99.82%	99.77%

Table 2: End-to-end performance of our attack

With more than 2 WiFi devices in a regular room, our attack can detect more than 99% of user presence and movement in each room tested. The tradeoff is slightly lower precision values, because the probability of assigning a WiFi anchor to the wrong room also increases. On the other hand, if one can “relax” the requirement of detecting user activity in each individual room to detecting in the target area (per area result in the table), then our attack can achieve very high recall and precision (>99.77%). Here, a potential improvement to our attack is to perform movement detection using a carefully chosen subset of anchors with more confident room assignments. We leave this optimization as future work.

8 Defenses

Having demonstrated the effectiveness of our passive sensing attacks, we now explore 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 reducing the amount of WiFi signal leakage to external sniffers or adding inconsistency to WiFi signals could render the attack ineffective.

The Failure of MAC Randomization. The first solution to come to mind would be *MAC address randomization*, a well-known method for protecting mobile devices

against tracking. Since the attack sniffer uses MAC address to isolate signals of each anchor device, MAC randomization can disrupt both bootstrapping and continuous sensing phases. However, recent work has shown that MAC randomization is disabled on most devices (<3% of adoption rate so far) [44] and can be easily broken to reveal the real MAC address [43, 5]. Thus Android 9.0 Pie switches to per-network MAC randomization [6], which does not apply any MAC randomization to static WiFi devices. Thus MAC randomization is not a plausible defense against our attack.

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

8.1 Geofencing WiFi Signals

Geofencing creates a geographical boundary for WiFi signal propagation to significantly reduce or eliminate WiFi signals accessible to the adversary, in terms of the size of the area where the adversarial sniffer can hear signals, and the total number of packets captured. For example, while our experiments in §7 were based on walking traces of 25–50 meters each, geofencing might reduce the area with a signal in our walking trace to 10 meters or less. If we reduce our sniffed packet trace accordingly, the localization error increases significantly. Raw errors more than doubled, and room-level accuracy dropped from 92.6% to 41.15%.

Practical Implications. Geofencing, deployed effectively, 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 almost always undesirable, since it degrades connectivity of WiFi clients. Another alternative is to equip WiFi devices with directional antennas, thus limiting RF emissions in the spatial domain. This approach is undesirable because it not only requires users to upgrade to equipment with higher cost and larger form factors, but also carefully configure their antenna directionality. Finally, the extreme solution is to block RF signals from propagating beyond walls by painting (boundary) walls with electromagnetic shielding paint. This is impractical, since it would also block cellular signals.

A more practical alternative is to customize WiFi signal coverage using 3D fabricated reflectors, proposed recently by [76]. It has limited applicability and considerable complexity, since the reflector configuration depends on both WiFi device placement and details of the environment.

8.2 WiFi Rate Limiting

While geofencing reduces spatial leakage of WiFi signals, rate limiting reduces their temporal volume. When anchors transmit less signals over time, the sniffer will not have sufficient data to compute A-CSI and STD. Results in §7 show that reducing anchor packet rates lowers the recall value.

Practical Implications. Rate limiting is simple to implement, but creates undesirable artifacts to network applications. As shown in Table 1, many WiFi devices in offices and homes, even when idle, transmit beyond 2 packets per second (pps). This makes rate limiting impractical.

8.3 Signal Obfuscation

Our third defense is to add noise to WiFi signals, so the adversary cannot accurately localize anchors or detect user movements. We refer to this to as *signal obfuscation*.

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 [40] 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, a sniffer cannot accurately predict of their locations. Yet this requires tight synchronization and active coordination between devices, without which it is possible for the sniffer to separate data streams.

AP-based Signal Obfuscation. We propose a practical defense where the WiFi access point (AP) actively injects cover traffic for any of its associated WiFi device w that is actively transmitting. As soon as the AP detects a transmission from w , it estimates w 's transmission rate T_w and injects a cover traffic stream with the same T_w , at a randomized power level and with w 's MAC address. If AP limits its cover traffic stream to match w 's throughput, then WiFi's CSMA protocol will randomly interleave packets from the two streams together. In the worst case (T_w is at or higher than available channel throughput), the cover traffic will reduce w 's effective throughput by half (50%). If T_w is less than half of available throughput, then the additional cover traffic will have minimal impact on w 's throughput.

With this defense, the attacker's RSS measurements of anchor w will display fluctuations, tricking the adversary to think that w is moving and making it useless as an anchor. Even if the adversary assumes w is stationary, the noisy RSS measurements will lead to inaccurate anchor placement. It is possible for Monte Carlo sampling (§5) to extract "clean" measurements of w , but the probability is extremely low. More importantly, A-CSI of these anchors will contain sufficient variations, indicating that a user is always present.

The insertion of "fake" packets requires a careful design, so that it disrupts the attack rather than creating obvious "anomalies" or heavily affecting the WiFi network. The AP configures the sequence numbers of fake packets to (partially) interleaved with those of real packets, so that the attacker is unable to separate the two streams based on sequence number and packet arrival time. The AP also needs to continue to periodically adjust its transmit power.

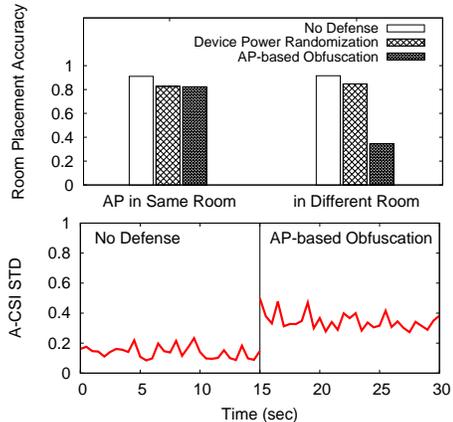


Figure 13: Performance of bootstrapping (top) and continuous sensing (bottom) with and without signal obfuscation.

When evaluating this defense, we assume that the adversary deploys countermeasures by adding an extra stationary sniffer, and applies signal subtraction [40] to remove "injected" signal variations. Figure 13 (top) compares the anchor localization accuracy with 3 schemes: using AP-obfuscation, using only device power randomization, and no defense. For anchors not in the same room as the AP, even with attacker countermeasures, AP-obfuscation reduces anchor localization accuracy from 90% down to 38%. Device-based power randomization is rendered ineffective by the countermeasure. If the AP and w are in the same room, the attacker can still localize the transmissions to the room.

Figure 13 plots the impact on continuous sensing when no user is present. The AP-based defense injects signal variations and confuses the attacker, making it constantly sense the presence of a user, effectively protecting the location/room from our attack.

Practical Implications. This defense can be deployed by today's WiFi APs that support transmit power adaptation with minor changes. The major drawback is the extra consumption of bandwidth and energy at the AP. We note that this defense targets attackers with a single antenna device. Advanced attackers with multiple antennas (> 3) could potentially separate AP signals from device signals by estimating their angle of arrival, making the defense less effective.

9 Related Work

Location Privacy. Whether it is compromising service providers [52], hacking into social networks [26] or smartphone sensors and power meters [48, 45], existing works have identified a wide variety of attacks on location privacy and subsequent defenses [14, 73, 60, 52, 49, 29, 15]. Our work targets a different type of location privacy attack, which tracks presence and movement of targets by monitoring ambient WiFi transmissions.

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 [16]. Prior works use traffic patterns of sniffed signals to infer user status [36, 56, 84], but require accurate identification of each device, knowledge of their transmission behaviors, and can be easily countered by adapting device transmission behaviors. Our attack does not make any of these strong assumptions.

Privacy Invasion from Signal Sniffing. Similar to our attack, existing works develop attacks to locate devices and infer user activities (based on the located room type) using WiFi and ZigBee signals [40, 13] or acoustic signals [47]. Our attack differs from them as follows. [40] focuses solely on locating WiFi cameras using RSS, and applies feature clustering to identify good measurements. Our work develops a different and more effective sifting method to identify good signal measurements, and targets a different problem of locating users who do not carry any WiFi devices.

With a strong assumption that a known router is placed in the center of a home, [13] deploys multiple laptops (each with three antennas) outside to detect human movements using either CSI or RSS measurements. With their design, each laptop can only detect user movements that block the direct path between the laptop and the router. Thus [13] deploys many laptops around the house to detect user movements but still cannot locate them. Our work takes a different methodology: we use a single smartphone (with a single antenna) to monitor ambient transmissions from many devices in the home; our A-CSI STD model also provides accurate room-level localization of user movements.

[47] detects user presence using specially crafted acoustic signals transmitted by devices in their own homes. It requires remote access to these devices, a strong assumption in practice. Our attack leverages ambient WiFi signals and does not require any access to devices in the target area.

Device-Free Human Sensing. Non-adversarial human sensing correlates human movements with wireless signal variations caused by these movements, thus not requiring human to carry any devices. Sensing can be achieved by either active probing or passive snooping.

Existing works in the active category deploy a transmitter to continuously send probing signals (either standard RF signals or crafted RADAR signals like FMCW [11, 86]), and deploy receiver(s) to capture signals as they bounce off the targets. Existing designs operate on either time of flight (ToF) [30, 53], frequency shift [11, 86]), CSI phase shifts [69, 67, 80] or RSS [81, 62]. Our attack differs from these works by being passive, not requiring any RF transmission by the attacker device. Also our sensing design operates on A-CSI rather than CSI phase, ToF or RSS.

Works in the passive category sniff existing wireless signals to detect human presence and activities. The majority of existing proposals rely on fingerprinting, *i.e.* mapping any observed signals to a pre-defined fingerprint representing a specific target location and/or activity. A fingerprint can be

based on A-CSI [46, 70], CSI phase [51], RSS [27, 64], or raw signals [72]. Yet fingerprinting requires target cooperation, clearly infeasible under our attack scenario. Another work extracts doppler shift from sniffed signals to detect human presence but cannot locate the target [18]. It also requires a large antenna dish.

Others use non-WiFi/RF signals, *e.g.* RFIDs [85], visible light [83, 38], acoustic [41], to sense human activities. They require control of transmitters inside or outside of the target's home/office, and are infeasible under our attack scenario.

Transmitter Localization. Solutions to this well studied problem can be divided into three categories: fingerprinting, active probing, and passive trilateration. *Fingerprinting* first uses measurements of RSS, CSI or other metrics within the target area to build a database of signal patterns (fingerprints) for each location. It then maps any observed signal pattern to the closest fingerprint to determine the device location [64, 12, 20]. This is clearly infeasible under our attack scenario. *Active probing* exchanges RF or acoustic communications with the target device to measure signal propagation delay, degradation or phase shift in order to compute the distance to the target [8, 50, 82, 9], followed by trilateration.

Finally, *passive localization* often leverages receivers with multiple antennas [10, 63, 31, 75, 35] to estimate signal incoming angle (AoA), and applies triangulation across multiple receivers to derive the target location. Recent work [34] lowers the antenna count to 3, but requires at least two line-of-sight paths between the transmitter and receiver (*i.e.* no walls) and multiple APs. Our work (bootstrapping) differs by using a single smartphone with a single antenna. We use an existing branch of passive localization that fits spatial measurements of RSS to a propagation model [28, 40, 22]. Our key contribution is a data sifting algorithm that identifies good RSS samples as input to the model fitting.

Defense against RF Eavesdropping. Existing works [33, 54, 21, 68] defend against eavesdropping on a transmitter by a jammer transmitting simultaneously, preventing the attacker from decoding packets or estimating AoA. This requires precise synchronization between the transmitter and the jammer [32] or a high-end full-duplex obfuscator [54]. Our defense uses the AP to insert fake packets (rather than transmitting simultaneously), which is simple to deploy and effective against attackers with a single antenna.

10 Conclusion

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

References

- [1] Linux 802.11n csi tool. <https://dhalperi.github.io/linux-80211n-csitool/>, 2011.
- [2] Howto estimate parameter-errors using monte carlo. <http://www-personal.umd.umich.edu/~wiclarks/AstroLab/HOWTOs/NotebookStuff/MonteCarloHOWTO.html>, 2014.
- [3] Nexmon: The c-based firmware patching framework. <https://nexmon.org>, 2017.
- [4] Over the air: Exploiting broadcom’s wi-fi stack. https://googleprojectzero.blogspot.com/2017/04/over-air-exploiting-broadcoms-wi-fi_4.html, 2017.
- [5] 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.
- [6] 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.
- [7] Ettus research products. <https://www.ettus.com/product/>, 2018.
- [8] Wifi indoor positioning. <https://fit-iot.com/web/wifi-indoor-positioning/>, 2018.
- [9] ADIB, F., KABELAC, Z., AND KATABI, D. Multi-person localization via RF body reflections. In *Proc. of NSDI* (2015).
- [10] ADIB, F., AND KATABI, D. See through walls with wifi! In *Proc. of SIGCOMM* (2013).
- [11] 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).
- [12] BAHL, P., AND PADMANABHAN, V. N. Radar: an in-building rf-based user location and tracking system. In *Proc. of INFOCOM* (2000).
- [13] 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).
- [14] BINDSCHAEDLER, V., AND SHOKRI, R. Synthesizing plausible privacy-preserving location traces. In *Proc. of SP* (2016).
- [15] BORDENABE, N. E., CHATZIKOKOLAKIS, K., AND PALAMIDESSI, C. Optimal geo-indistinguishable mechanisms for location privacy. In *Proc. of CCS* (2014).
- [16] CHENG, Y., JI, X., LU, T., AND XU, W. Dewicam: Detecting hidden wireless cameras via smartphones. In *Proc. of Asia CCS* (2018), ACM.
- [17] CHENG, Y.-C., CHAWATHE, Y., LAMARCA, A., AND KRUMM, J. Accuracy characterization for metropolitan-scale wi-fi localization. In *Proc. of MobiSys* (2005).
- [18] CHETTY, K., SMITH, G. E., AND WOODBRIDGE, K. Through-the-wall sensing of personnel using passive bistatic wifi radar at standoff distances. *IEEE Transactions on Geoscience and Remote Sensing* 50, 4 (2012).
- [19] EVENNOU, F., AND MARX, F. Advanced integration of wifi and inertial navigation systems for indoor mobile positioning. *EURASIP J. Appl. Signal Process* 2006 (2006).
- [20] FARID, Z., NORDIN, R., AND ISMAIL, M. Recent advances in wireless indoor localization techniques and system. *Journal of Computer Networks and Communications* 2013 (2013).
- [21] GOLLAKOTA, S., AND KATABI, D. ijam: Jamming oneself for secure wireless communication. Tech. rep., Computer Science and Artificial Intelligence Laboratory Technical Report, 2010.
- [22] GOSWAMI, A., ORTIZ, L. E., AND DAS, S. R. Wigem: A learning-based approach for indoor localization. In *Proc. of CoNEXT* (2011).
- [23] 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).
- [24] HAMPEL, F. R. The influence curve and its role in robust estimation. *Journal of the American Statistical Association* 69, 346 (1974), 383–393.
- [25] 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).
- [26] 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).
- [27] HUANG, H., AND LIN, S. Widet: Wi-fi based device-free passive person detection with deep convolutional neural networks. In *Proc. of MSWIM* (2018).
- [28] JI, Y., BIAZ, S., PANDEY, S., AND AGRAWAL, P. Ariadne: A dynamic indoor signal map construction and localization system. In *Proc. of MobiSys* (2006).
- [29] JIN, X., ZHANG, R., CHEN, Y., LI, T., AND ZHANG, Y. Dpsense: Differentially private crowdsourced spectrum sensing. In *Proc. of CCS* (2016).
- [30] JOSHI, K., BHARADIA, D., KOTARU, M., AND KATTI, S. Wideo: Fine-grained device-free motion tracing using rf backscatter. In *Proc. of NSDI* (2015).
- [31] KARANAM, C. R., KORANY, B., AND MOSTOFI, Y. Magnitude-based angle-of-arrival estimation, localization, and target tracking. In *Proc. of IPSN* (2018).
- [32] KHALEDI, M., KHALEDI, M., KASERA, S. K., AND PATWARI, N. Preserving location privacy in radio networks using a stackelberg game framework. In *Proc. of Q2SWinet* (2016).
- [33] KIM, Y. S., TAGUE, P., LEE, H., AND KIM, H. Carving secure wi-fi zones with defensive jamming. In *Proc. of Asia CCS* (2012).
- [34] KOTARU, M., JOSHI, K., BHARADIA, D., AND KATTI, S. Spotfi: Decimeter level localization using wifi. In *Proc. of SIGCOMM* (2015).

- [35] KOTARU, M., AND KATTI, S. Position tracking for virtual reality using commodity wifi. In *Proc. of CVPR* (2017).
- [36] 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).
- [37] 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).
- [38] LI, T., LIU, Q., AND ZHOU, X. Practical human sensing in the light. In *Proc. of MobiSys* (2016).
- [39] 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).
- [40] LI, Z., XIAO, Z., ZHU, Y., PATTARACHANYAKUL, I., ZHAO, B. Y., AND ZHENG, H. Adversarial localization against wireless cameras. In *Proc. of HotMobile* (2018).
- [41] MAO, W., HE, J., AND QIU, L. Cat: High-precision acoustic motion tracking. In *Proc. of MobiCom* (2016).
- [42] 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).
- [43] 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).
- [44] 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).
- [45] 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).
- [46] NANDAKUMAR, R., KELLOGG, B., AND GOLLAKOTA, S. Wi-fi gesture recognition on existing devices. *CoRR abs/1411.5394* (2014).
- [47] NANDAKUMAR, R., TAKAKUWA, A., KOHNO, T., AND GOLLAKOTA, S. Covertband: Activity information leakage using music. In *Proc. of UbiComp* (2017).
- [48] 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).
- [49] 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).
- [50] PENG, C., SHEN, G., ZHANG, Y., LI, Y., AND TAN, K. Beepbeep: A high accuracy acoustic ranging system using cots mobile devices. In *Proc. of SenSys* (2007).
- [51] PU, Q., GUPTA, S., GOLLAKOTA, S., AND PATEL, S. Whole-home gesture recognition using wireless signals. In *Proc. of MobiCom* (2013).
- [52] 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).
- [53] 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).
- [54] QIAO, Y., ZHANG, O., ZHOU, W., SRINIVASAN, K., AND ARORA, A. Phycloak: Obfuscating sensing from communication signals. In *Proc. of NSDI* (2016).
- [55] ROUSSEEUW, P. J., AND CROUX, C. Alternatives to the median absolute deviation. *Journal of the American Statistical association* 88, 424 (1993), 1273–1283.
- [56] SANCHEZ, I., SATTÀ, 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).
- [57] 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).
- [58] SEN, P. K., AND SINGER, J. M., Eds. *Large sample methods in statistics*. Chapman & Hall, Inc., 1989.
- [59] SEYBOLD, J. *Introduction to Rf Propagation*. Wiley, 2005.
- [60] 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).
- [61] SIBY, S., MAITI, R. R., AND TIPPENHAUER, N. O. Iotscanner: Detecting privacy threats in iot neighborhoods. In *Proc. of IoTPTS* (2017).
- [62] 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).
- [63] 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).
- [64] SRINIVASAN, V., STANKOVIC, J., AND WHITEHOUSE, K. Protecting your daily in-home activity information from a wireless snooping attack. In *Proc. of UbiComp* (2008).
- [65] TAN, S., AND YANG, J. Wifinger: Leveraging commodity wifi for fine-grained finger gesture recognition. In *Proc. of MobiHoc* (2016).
- [66] TSAI, M. Path-loss and shadowing (large-scale fading). Tech. rep., 2011.
- [67] WANG, J., JIANG, H., XIONG, J., JAMIESON, K., CHEN, X., FANG, D., AND XIE, B. Lifis: Low human-effort, device-free localization with fine-grained subcarrier information. In *Proc. of MobiCom* (2016).
- [68] WANG, T., LIU, Y., PEI, Q., AND HOU, T. Location-restricted services access control leveraging pinpoint waveforming. In *Proc. of CCS* (2015).
- [69] 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).

- [70] 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).
- [71] WEI, T., WANG, S., ZHOU, A., AND ZHANG, X. Acoustic eavesdropping through wireless vibrometry. In *Proc. of MobiCom* (2015).
- [72] XIAO, N., YANG, P., YAN, Y., ZHOU, H., AND LI, X. Motion-fi: Recognizing and counting repetitive motions with passive wireless backscattering. In *Proc. of INFOCOMM* (2018).
- [73] XIAO, Y., AND XIONG, L. Protecting locations with differential privacy under temporal correlations. In *Proc. of CCS* (2015).
- [74] XIE, Y., LI, Z., AND LI, M. Precise power delay profiling with commodity wifi. In *Proc. of MobiCom* (2015).
- [75] XIONG, J., AND JAMIESON, K. Arraytrack: A fine-grained indoor location system. In *Proc. of NSDI* (2013).
- [76] 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).
- [77] YANG, L., CHEN, Y., LI, X.-Y., XIAO, C., LI, M., AND LIU, Y. Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices. In *Proc. of MobiCom* (2014).
- [78] YANG, Z., ZHOU, Z., AND LIU, Y. From rssi to csi: Indoor localization via channel response. *ACM Comput. Surv.* 46, 2 (2013).
- [79] 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).
- [80] 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).
- [81] YOUSSEF, M., MAH, M., AND AGRAWALA, A. Challenges: device-free passive localization for wireless environments. In *Proc. of MobiCom* (2007).
- [82] YOUSSEF, M., YOUSSEF, A., RIEGER, C., SHANKAR, U., AND AGRAWALA, A. Pinpoint: An asynchronous time-based location determination system. In *Proc. of MobiSys* (2006).
- [83] ZHANG, C., AND ZHANG, X. Litell: Robust indoor localization using unmodified light fixtures. In *Proc. of MobiCom* (2016).
- [84] ZHANG, F., HE, W., LIU, X., AND BRIDGES, P. G. Inferring users' online activities through traffic analysis. In *Proc. of WiSec* (2011).
- [85] ZHANG, J., TIAN, G., MARINDRA, A. M. J., IMAM, A., AND ZHAO, A. A review of passive rfid tag antenna-based sensors and systems for structural health monitoring applications. *Sensors* 17 (2017).
- [86] ZHAO, M., ADIB, F., AND KATABI, D. Emotion recognition using wireless signals. In *Proc. of MobiCom* (2016).
- [87] 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 [37]. 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 :

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

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 §5.3), we can approximate d_i by

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

where x_s and y_s 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 floor level detection (§5.3).

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:

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

The constraint on $\hat{\gamma}$ follows the well-known observations from empirical measurements [66] while the value of \hat{P}_{TX^o} is upper bounded by the maximum transmit power for WiFi frequency defined by the FCC.

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 [76] 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.

11.2 Details on Distance Impact on A-CSI

In §4.2, we make a hypothesis about A-CSI STD and the target-anchor distance. When a human user moves around a transmitter (TX), she blocks and diffracts some signal propagation paths from TX to a receiver (RX). When the user is close to TX, the set of paths affected by her movements is larger than that when she is far away from TX. As such the received signals seen by RX will display a larger variation σ^2 as the user gets closer to TX.

To confirm this hypothesis, we build a ray tracing model on signal propagation from TX to RX as follows. Let P_k represent the received signal power at RX on sub-carrier k , measured at wavelength λ_k . P_k can be modeled as an aggregation of multiple (N) signal paths $\mathbb{S} = \{p_i\}_{i=1..N}$ [78]:

$$P_k = \left\| \sum_{i=1}^N p_i(d_i, \Gamma_i, \lambda_k) \right\| \quad (3)$$

where d_i and Γ_i are the propagation distance and reflection/diffraction coefficient of the signal path $i = 1..N$, respectively.

For simplicity, we also assume that the human user moves at a constant speed around TX with a fixed distance r . Similar to a widely used model [28], we consider complete blockage (*i.e.* the coefficient $\Gamma_i = 0$ or 1) and disregard the contribution of signal phase in the above summation.

When the user moves near the TX (at distance r_1), some signal paths $\mathbb{S}(r_1)$ are blocked. We denote the time upon blockage as $[t_1, t_2]$ during time $[t_0, t_3]$, where $t_0 < t_1 < t_2 < t_3$. Then the received signal power in time is:

$$P_k(t, r_1) = \begin{cases} \sum_{i \notin \mathbb{S}(r_1)} p_i(d_i, \Gamma_i, \lambda_k) & \text{if } t \in [t_1, t_2] \\ \sum_i p_i(d_i, \Gamma_i, \lambda_k) & \text{otherwise} \end{cases} \quad (4)$$

From above, we know the mean power $\mu_k(r_1)$ over $[t_1, t_2]$ must follow $\mu_k(r_1) > P_k(t, r_1)$. Similarly, $\mu_k(r_2) > P_k(t, r_2)$ when $t \in [t_1, t_2]$. At distance $r_2 > r_1$, the set of blocked signal paths $|\mathbb{S}(r_2)| < |\mathbb{S}(r_1)|$, thus $P_k(t, r_2) > P_k(t, r_1) > 0$ $t \in [t_1, t_2]$. Therefore,

$$\mu_k(r_2) > \mu_k(r_1) \quad (5)$$

The standard deviation $\sigma_k(r)$ between t_0 and t_3 is defined as:

$$\sigma_k^2(r) = \frac{1}{t_3 - t_0} \int_{t_0}^{t_3} (P_k(t, r) - \mu_k(r))^2 dt \quad (6)$$

From these conditions, we can easily show that $\sigma_k^2(r_2) < \sigma_k^2(r_1)$. And the averaged standard deviation over k sub-carriers follow the same observation.

11.3 Details on Test Scenes

The following table lists the configuration of our test scenes, which include both offices and apartments of different sizes.

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 3: Test scene configuration.