



RF-Protect: Privacy against Device-Free Human Tracking

Jayanth Shenoy[†], Zikun Liu[†], Bill Tao[†], Zachary Kabelac[‡], Deepak Vasisht[†]

[†]University of Illinois Urbana-Champaign, [‡]Analytical Space

ABSTRACT

The advent of radio sensing that works through walls & obstacles challenges the notion of indoor privacy. An eavesdropper can deploy such sensing to snoop on their neighbors and a smart sensor embedded with such sensing capabilities can perform large scale behavioral and health data mining. We present RF-Protect, a new framework that enables privacy by injecting fake humans in the sensed data. RF-Protect consists of a novel hardware reflector design that modifies radio waves to create reflections at arbitrary locations in the environment and a new generative mechanism to create realistic human trajectories. RF-Protect’s design doesn’t require any high bandwidth hardware or physical motion. We implement RF-Protect using commodity hardware and validate its ability to generate fake human trajectories.

CCS CONCEPTS

• **Security and privacy** → *Mobile and wireless security*; • **Computer systems organization** → *Embedded and cyber-physical systems*.

KEYWORDS

Wireless, Embedded Systems, Privacy, Sensing, IoT, RF

ACM Reference Format:

Jayanth Shenoy[†], Zikun Liu[†], Bill Tao[†], Zachary Kabelac[‡], Deepak Vasisht[†]. 2022. RF-Protect: Privacy against Device-Free Human Tracking. In *ACM SIGCOMM 2022 Conference (SIGCOMM ’22)*, August 22–26, 2022, Amsterdam, Netherlands. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3544216.3544256>

1 INTRODUCTION

In the last decade, both academia and industry have relied on FMCW¹-radar based radio-frequency (RF) sensors to enable through-wall human tracking. These sensors capture reflections from human bodies to track occupancy of rooms [5], motion patterns of occupants [4, 5], their daily activities [41], and their health metrics [8, 34]. Recently, Google has incorporated high frequency FMCW-based sensing into their smart home devices [29, 30], and Amazon received an FCC waiver [9] to conduct testing for the same.

While RF sensing systems enable many new applications like personalized healthcare and smart homes, they are fraught with

¹FMCW:Frequency Modulated Carrier Wave

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGCOMM ’22, August 22–26, 2022, Amsterdam, Netherlands

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9420-8/22/08...\$15.00

<https://doi.org/10.1145/3544216.3544256>

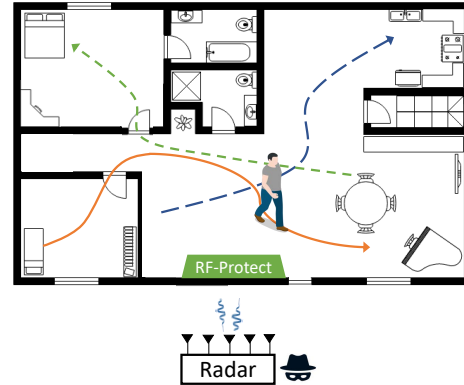


Figure 1: RF tracking systems can monitor an occupant’s motion patterns (solid line) from *outside* the building. RF-Protect creates ‘ghost’ reflections (dashed lines above) that resemble human reflections and corrupt eavesdropper’s information.

serious privacy risks. Such tracking works without requiring users to wear a device and has a range of multiple rooms. Moreover, low-frequency RF sensing systems have the ability to penetrate through walls and obstacles, making them an even greater privacy threat. Using RF sensing, a neighbor or an eavesdropper can track occupancy in a home, when occupants go to bed, take a bath, wake up, and other daily activities simply by deploying a sensor next to a wall. A thief could use such sensing to identify empty houses. Similarly, smart sensors, such as the ones commonly deployed directly inside homes, can mine mobility and activity tracking data at scale and gain insights into people’s behavioral distribution. These insights are then vulnerable to various kinds of abuse, like targeted advertisements or health insurance pricing.

To make this worse, there is no easy way today for users to avoid such tracking. Traditional approaches to counter such tracking rely on jamming [48] the radio spectrum. However, through-wall tracking systems operate at low power in bands shared by other users (e.g. for Wi-Fi [4, 5]). Thus, jamming will disable communication on user devices as well. Moreover, jamming spectrum for non-defense applications is largely illegal. Another alternative is to convert homes into Faraday cages with metal walls (or tin foil walls) that prevent wireless signals to propagate from one home to another. This also limits other signals of interest (like cellular signals) to operate in homes. Fundamentally, disabling such tracking is challenging because human bodies naturally reflect radio signals and such reflections cannot be ‘turned off’.

We present RF-Protect, a hardware-software system design that enables end users to counter unauthorized tracking. We note that this privacy problem closely resembles privacy in the context of web-based user tracking [12, 15, 44, 46, 49]. One cannot hide the queries and clicks they make on web-pages from web-servers.

Therefore, web-based privacy systems inject fake clicks and queries to fool the tracking entities. Inspired by this approach, we build a new privacy primitive that injects fake human-like reflections in the environment and disables accurate inferences by FMCW-based RF sensing systems. Our system doesn't make humans disappear, but fills up a home with fake humans that prevent accurate inferences about the state of a home (see Fig. 1). RF-Protect consists of: (a) a new hardware reflector, deployed in the environment, that generates fake customizable reflections in the environment, and (b) a new algorithmic framework that controls the reflector to ensure that these fake reflections mimic real humans, i.e. they walk and breathe like humans in different activity settings. RF-Protect can be deployed on a wall in the home and inject fake information in RF-based sensing without requiring any physical motion.

Our approach has three key advantages. First, RF-Protect significantly limits inferences on sensed data. By injecting fake human reflections, our system ensures that inferences like house occupancy, occupant sleep status, health metrics, etc. will reveal incorrect results. We show, in Sec. 7, that RF-Protect disables accurate inferences both at the instance level (e.g. is someone home now) and at the distribution level (e.g. what's the distribution of home occupancy for this household). Second, our defense is hard to detect without side-channel information because the reflectors don't transmit any signal of their own. Finally, our design does not interfere with legitimate sensors. Our hardware reflector can communicate the 'fake' information injected into the system to a legitimate tracking device authorized by the user. The legitimate device can remove the fake reflections and get access to real tracking results.

To understand why developing RF-Protect is challenging, recall that FMCW radars use GigaHertz of bandwidth to accurately measure human reflections. In such cases, creating static reflections is easy – a piece of metal works as a reflector. However, static reflectors are already abundant in the environment and RF sensing systems are designed to filter them out (e.g. by background subtraction or doppler shift filtering). Therefore, we must design a reflector that can create dynamic reflections that appear to emerge out of arbitrary locations. This is challenging due to the large bandwidth of the incident signal. A reflector would have to lock onto the GHz-wide FMCW signal, sample at GHz bandwidth, and modify the signal before relaying this signal back. Even then, an eavesdropper can simply weed out such reflections if they do not resemble human behavior across time. RF-Protect solves these challenges through the following innovations:

Generating Customizable Reflections: FMCW systems operate on the principle of time-of-flight, i.e. they measure the time it takes for a transmitted signal to reflect off the human and arrive at the receiver. To introduce reflections that can move over time, without using physical motion, our reflector must be able to introduce nanosecond-level delays in the reflected signal. Introducing delays at such fine-grained control requires complex, expensive radio hardware that operates at GHz bandwidth. We observe that in FMCW, the time delay and frequency are linearly correlated, and hence, we can introduce such delays in the reflections using small frequency shifts (kHz-scale). Interestingly, we can introduce such frequency shifts even without having to sample the signal, by simply switching the reflector on and off at the required frequency.

By changing the frequency of reflector switching, we can create reflections at different distances from the reflector. RF-Protect's hardware reflector builds on this idea by incorporating an array of switched antennas, allowing us to spoof reflections from different directions and create arbitrary trajectories in a 2-D space.

Mimicking Human Reflections: If RF-Protect's injected trajectories do not resemble realistic human motion, it is easy for a smart eavesdropper to consider them noise and remove them from the sensed data. Therefore, we need to create trajectories that accurately resemble human behavior. One way to achieve this is to have the reflection move in a fixed trajectory – e.g. a circle, which is fairly easy to generate. However, the eavesdropper can analyze their received data and realize that it isn't realistic for a human to move repeatedly along a circle. In general, we need to create trajectories that resemble the *distribution* of real human motion. To achieve this, we leverage inspiration from recent research in machine learning that generates new examples from a given distribution using Generative Adversarial Networks (GANs). We use a conditional GAN architecture to generate trajectories that mimic human motion and are indistinguishable by the eavesdropper. Finally, to complete this design, RF-Protect adds phase shifts to the signal to resemble human breathing to avoid identification.

We note that RF-Protect does not need to know the location of the eavesdropper. We believe that RF-Protect significantly raises the bar for privacy in passive tracking systems. We envision RF-Protect's design will be integrated into future smart surfaces to allow for privacy protections in addition to the intended communication benefits of these metasurfaces.

We build RF-Protect and test it against a state-of-the-art FMCW radar. RF-Protect's reflector can accurately replicate arbitrary trajectories, achieving a median error of 13 cm (home environment) and 24 cm (office environment) between its intended trajectory and the trajectory observed at the radar. We also show that RF-Protect's GAN creates trajectories that follow the same distribution as human trajectories. Finally, our experiments demonstrate RF-Protect's ability to allow a legitimate radar to decode human trajectories by communicating the injected reflection data. Our work, in RF-Protect, makes the following contributions:

- To the best of our knowledge, we build the first system to inject realistic human reflections in an FMCW-based sensing system.
- We propose a new system design to inject customizable 2D reflections for an indoor FMCW-based radar design.
- We build a new conditional GAN-framework for creating realistic human trajectories.
- We provide an information theoretic analysis that quantifies the privacy protection offered by RF-Protect.

2 THREAT MODEL

We target FMCW-based wireless sensors because they are one of the most robust and popular RF-based sensing systems today. They have been used to sense human motion [5, 17] and human health [7], as well as vehicular sensing [11, 14, 20]. In industry, the popular radars designed by Texas Instruments [24, 25] for both indoor and outdoor sensing use FMCW. Similarly, Google's smart devices [29] use FMCW for sensing.




Figure 2: FMCW-sensors measure delay between transmitted (blue solid) and reflected (green dashed) chirps.

Setup: We assume an adversarial eavesdropper interested in monitoring the motion or other activity patterns of a victim in an enclosed space (home, conference room, etc.). However, the eavesdropper has no physical or visible (e.g. deployed cameras) access to the space. An eavesdropper may be able to place the device in or near the space (e.g. along a wall of the neighboring apartment or embedded in a smart device). Fig. 1 illustrates this with the adversary being located at the red box outside the buildings.

Eavesdropper Hardware: To accomplish their goal, the eavesdropper uses FMCW-radar to measure distance to the target humans. Following standard design practices, we assume the eavesdropper uses a 1-D antenna-array (either digital beamforming or phased array) to calculate the direction from the radar to the tracked human. Therefore, the eavesdropper can precisely track the location of victims in a 2-D space. Finally, our work does not deal with an eavesdropper that violates FCC regulations on spectrum occupancy. We believe this assumption is reasonable because (a) smart device manufacturers need FCC certification, and (b) for a one-off attacker, they are unlikely to have the expertise to build custom FMCW radars that operate at uncommon frequencies.

Eavesdropper Algorithms: Our eavesdropper is capable of deploying mobility models that process the FMCW signals to track the trajectory of the victim and their health metrics. This may include machine learning or statistical approaches like Kalman Filters. They also deploy algorithms to isolate human trajectories from random motion (e.g. fans).

3 PRIMER ON FMCW RADARS

FMCW radars transmit a chirp signal (spanning several GHz), receive its reflections from the environment, and measure the time delay between transmission and reflection. These radars can operate in both line-of-sight and non-line-of-sight conditions, depending on their frequency.

Distance Measurement: To measure time delay, the radar transmits a chirp signal (Fig. 2), which has a frequency linearly dependent on time (say with slope, sl). The radar mixes the transmitted signal with the measured received signal. The resulting signal is a narrow-band signal at frequency, Δf , which is the difference of frequency between the transmitted and received signals. Due to the linear relationship of time and frequency, we can identify the time delay between transmission and reception using: $Delay = \frac{\Delta f}{sl}$

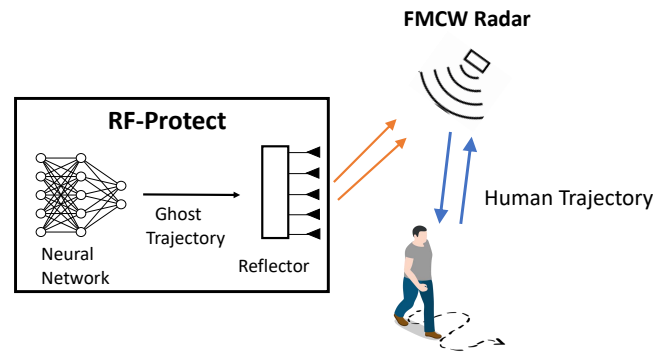


Figure 3: RF-Protect System Overview: A neural network generates ghost trajectories and creates them through the tag, misleading the FMCW Radar.

Since the radio wave travels at the speed of light, we can convert the round trip delay estimate into distance:

$$Distance = C \times \frac{Delay}{2} = C \times \frac{\Delta f}{2sl} \quad (1)$$

where C is the speed of light. If there are multiple reflections, they can be separated using the Fourier transform at a resolution of $\frac{C}{2B}$, where B is the chirp bandwidth [5].

Converting Distance to Location: RF-sensing systems, typically, use antenna arrays comprising of 4 to 8 antennas to estimate the angle at which the reflection is received. Specifically, the eavesdropper can compute $P(\theta)$, the portion of signal arriving along direction θ using the equation:

$$P(\theta) = \left| \sum_{k=1}^N h_k e^{\left(\frac{-j2\pi kd \cos \theta}{\lambda}\right)} \right|^2 \quad (2)$$

where h_k is the signal at the k th antenna, N is the number of antennas in the array, d is the antenna spacing, and λ is the signal wavelength.

Addressing Static Reflectors: The reflections measured by such sensing systems consist of reflections from other static objects in the environment (furniture, walls, etc.), in addition to human reflections. They reject static reflections using background subtraction, for example, by subtracting successive measurements.

4 SYSTEM OVERVIEW

The core technical challenge in preventing FMCW-based tracking is that they rely on human reflections, as opposed to device-tracking like smartphones. Smartphones can be turned off, but human reflections cannot be. An ideal privacy system will offer invisibility to humans. However, such invisibility is extremely challenging to achieve for electromagnetic waves, as discussed before. Approaches such as jamming and Faraday cages prevent other electromagnetic waves as well (like cellular signals and Wi-Fi).

In RF-Protect, we envision a new method to fool the adversary, by making them hallucinate realistic fake trajectories. This corrupts their sensed information, thus making them see humans when none exist, obtain incorrect data about people count in homes, and even observe fake breathing and sleeping activities. This privacy is not perfect, e.g. we cannot make humans disappear. However,

we believe that this system truly raises the bar for privacy against such sensing systems. We quantify the extent of privacy offered by our design in Sec. 7. Our design is motivated by fake data injection attacks in web-privacy – the webpages can still track your actions, but by injecting fake data (e.g. clicks and queries), the user can limit what useful information the webpage can infer about them.

An overview of our design is presented in Fig. 3. As shown, RF-Protect has two components: (a) a hardware reflector that is deployed in the environment and can create reflections at different locations (Sec. 5), and (b) a neural network architecture (Sec. 6) that generates trajectories that resemble human motion and feeds these trajectories to the hardware reflector. A user would ideally deploy the RF-Protect reflector close to vulnerable walls (e.g. with an adjacent neighbor) so that the reflector can be in the reflection range of a radar deployed along that wall. If there are multiple vulnerable walls, a user may deploy multiple such reflectors.

5 CREATING DYNAMIC REFLECTIONS

Our first goal is to create reflections that appear to emerge from arbitrary locations. A simple approach would use a high bandwidth (GHz-wide) receiver to receive the FMCW signal, do signal processing to synchronize itself with the radar transmissions, and then transmit chirps of its own with varying delay. However, this approach requires complex hardware that can sample at GHz bandwidths, needs continuous synchronization with the eavesdropper, and is easy to detect. Past work [27] has shown that the eavesdropper can detect (and later circumvent) such defense mechanisms by simply turning the eavesdropping radar off. Since the defender operates with a lag due to the processing, it would continue to transmit and give itself away.

In contrast, we aim to design a simple reflector that does not need complex GHz-scale sampling, creates true reflections, and yet spoofs the location information present in the reflections. To this end, we build a hardware-software system that leverages the structure of the incident FMCW signal to modify the delay and direction information sensed by the radar. We discuss our design below.

5.1 Distance Spoofing

Recall from Sec. 3, an FMCW sensor measures the distance to a human based on the time-of-flight of the reflected signals off the victim’s body. Specifically, the sensor mixes the received signal with the transmitted signal, and uses the resulting signal to compute the time-delay, and hence the distance to a reflector. Now, consider Eq. 1. This equation identifies the frequency of the mixed signal as the difference of the instantaneous frequency between the transmitted and received signal, and then maps this frequency difference (Δf) to distance. If we deploy a simple reflector (e.g. an antenna or a piece of metal), R , in the environment, it will cause a corresponding frequency shift, Δf_R . This f_R corresponds to the distance between the reflector and the FMCW sensor.

Since the distance between the reflector and the sensor is static, Δf_R is constant across time. As a result, when the sensor performs background subtraction, this reflection will be eliminated. At this




Figure 4: RF-Protect leverages geometry of radar beams to spoof different angles with respect to the radar.

level, this reflector behaves like any other reflector in the environment like TV, fridge, etc. which the FMCW sensors don’t intend to track and eliminate through the process of background subtraction. However, if we can find a way to vary Δf_R across time in a simple way, we can create dynamic reflections that vary in distance over time.

Our idea is that we can induce a small frequency shift in the radar reflection using simple hardware, leading to a time-varying Δf_R . In fact, we can just turn the reflector on or off at a switching frequency of f_{switch} to create a reflected signal at $\Delta f_R + f_{switch}$. This operation is roughly equivalent to mixing the incident FMCW signal with a wave of frequency f_{switch} , but does not require high-frequency components like mixers, nor does it suffer from attenuation caused by such devices. Formally, the switching operation leads to the FMCW sensor observing a distance, d' that is:

$$d' = C \times \frac{\Delta f_R + f_{switch}}{sl} \quad (3)$$

Before we conclude, we note a few points below:

- Our design aims to add additional delay to the signal, not decrease it. This is because in our deployment scenario, the reflector is deployed on the boundary walls next to the radar. Therefore, any reflections from inside the house must experience extra distance with respect to the FMCW sensor (see Fig. 2).
- Note that Eq. 3 relies on the value sl of the slope of the FMCW chirp. In practice, the considerations of bandwidth, speed of human motion, and FCC regulations limit the slope to a small range. For example, a signal that sweeps too fast will experience aliasing and low SNR, and a signal that sweeps too slow will be disrupted by human motion. Within such variations, RF-Protect’s design works well. The variation in slope ends up scaling the distance spoofed by RF-Protect to be higher or lower, but the structure of motion largely remains the same. Future work may rely on a design where a small narrowband receiver scans the signals for some time and identifies the slope. For public systems, like systems embedded in smart sensors, this information will be publicly available anyway.
- Since we use a square wave (on-off switching), it can cause harmonics in the reflections, leading to additional reflections at $-f_{switch}$, $2f_{switch}$, $3f_{switch}$, etc. The negative harmonics occur behind the radar or outside the home walls and can be removed by using single sideband modulation like [50] if needed. We observe that the higher harmonics are typically much weaker than human motion.

5.2 Direction Spoofing

The distance spoofing step above allows us to emulate motion along a straight line joining the radar and the RF-Protect antenna. As the

distance increases, the radar sees the reflector farther away along that line. Now, we discuss how we can simulate two-dimensional trajectories using this primitive.

As discussed in Sec. 2, we consider an FMCW sensor using an antenna array for angle estimation. Such antenna arrays are standard in various commercial [24–26] and academic FMCW systems [20, 51]. Antenna arrays, depending on their configuration, use two types of beamforming methods for angle estimation – (a) analog beamforming – the FMCW radar electronically steers the antenna array beam using phased arrays and measures the signal reflections along each angle, and (b) digital beamforming – the FMCW radar processes the signal along each antenna in software to isolate the signal along different angles. It is very challenging to spoof the angle of reflection in both these systems. In analog beamforming, if your reflector is not along a specific direction, it won't reflect when the beam is pointed in that direction. In digital beamforming, since each antenna is omni-directional, the signal reflections will arrive at the radar, but to manipulate such reflections, one needs knowledge of the wireless channel between the tag and the radar. This is near impossible to obtain without cooperation with the FMCW radar.

We take a different approach. We leverage the simplicity of our reflector design to create a physical spread in space. Observe that an FMCW antenna array (either phased array or multi-receiver) comprising K antennas approximately divides the angular space in front of it into K sectors. More specifically, the antenna array cannot separate paths that arrive at angular separation of less than $\frac{\pi}{K}$. Such paths will merge with each other from the perspective of the FMCW sensor. We leverage this observation to create smooth angular motion in our reflected signal.

We place a switched antenna array along the wall of the protected area (an alternative design could use multiple single antenna reflectors), as shown in Fig. 4. Each antenna acts as a potential reflector and therefore can emulate motion along the lines shown in the figure. We simply switch between the different antennas as needed to spoof the required angle. For example, if we place K_R antennas along a wall, then we get K_R possible directions to project a human motion. In practice, this gives us a grid of 2D locations where we can simulate human motion. Note that, we do not perform beamforming in our design, we are simply switching between multiple antennas. Therefore, our design is simple and does not require multiple RF chains, phase shifters, or complex synchronization. The flexibility of our approach allows us to distribute the antennas in space and create smooth angular motion with respect to the FMCW sensor.

This design has several advantages. First, it works with both analog and digital beamforming because the antenna reflector is physically present along a given direction. Second, it doesn't require complex synchronization between the multiple antennas or reflectors. Tens of milliseconds of synchronization is sufficient, as opposed to nanosecond-level synchronization required for antenna arrays. Third, we can also use the multiple antennas in a single environment to generate multiple phantoms. Finally, note that the eavesdropper is at an unknown location. However, our

design is naturally resilient to unknown eavesdropper location. If the eavesdropper changes their location, the projected trajectory just undergoes an angular shift (as shown in Fig. 4).

We observe that RF-Protect's angle spoofing mechanism functions ideally when the reflector is deployed against a wall. This observation holds true for several reasons. First, boundary deployment of RF-Protect will always ensure that the eavesdropper is closer to the reflector than the victim and maximizes the range of angles that the reflector can spoof. Boundary deployment of RF-Protect will also improve the likelihood that the reflected signal is strong enough to reach the eavesdropper. Note that as the distance between the reflector and the eavesdropper increases, we observe a trade-off between the range of angles that can be spoofed and the resolution of the spoofed angles (leading to more continuous angular reflections). Increased distance between RF-Protect and the eavesdropper will also degrade the quality of the radar's sensing as the reflections received at the radar from both a victim and RF-Protect will be significantly weaker when they propagate over a longer distance. As a result, we envision that RF-Protect's hardware reflector can become integrated as a component in emerging smart surface technologies. RF-Protect can potentially play a vital role in augmenting the privacy capabilities that smart surfaces can provide.

Finally, note that the number of RF-Protect antennas, K_R , needs to be of the same order as the number of antennas on the radar K . Since the area of interest typically doesn't cover the entire span of the radar, even if $K_R \sim K$, we end up emulating a continuous trajectory.

5.3 RF-Protect Schematic

A schematic of RF-Protect's design is shown in Fig. 5. The reflector receives the radar signal, amplifies it, and passes it through a switch operating at switching frequency, f_{switch} . A microcontroller can manipulate f_{switch} over time to create random variation in distance sensed by the radar signal. Note that for typical radar design parameters, this frequency shift corresponds to tens to hundred kHz which suffices in creating a home-level distance shift. Therefore, it can be created using a low-power low-cost setup.

Given a trajectory τ , RF-Protect maps it to a sequence of antennas and frequency shifts and spoofs this sequence using the RF-Protect reflector. We note that the trajectory sensed by the FMCW sensor may be rotated or scaled due to unknown parameters like location of the radar sensor and the slope of the FMCW chirp.

Spoofing Breathing: Finally, note that our design includes a phase shifter to enable the capability of emulating human breathing motion. We encourage the reader to refer to past work in breath monitoring [6] for details, but at a high level, when humans are static, their chest motion can be used to identify breathing period. Since this motion is miniscule compared to other human motions like walking, this motion manifests itself better in the phase of the signal. By changing the phase of the reflected signal using the phase shifter, we are able to replicate the human breathing behavior.

6 CREATING REALISTIC TRAJECTORIES

RF-Protect's reflector can generate arbitrary fake trajectories. However, what are the fake trajectories that it should generate? A simple

²In theory, the angular resolution is evenly distributed in the $\cos \theta$ space not the angle (θ) space, but we approximate this for ease of exposition.

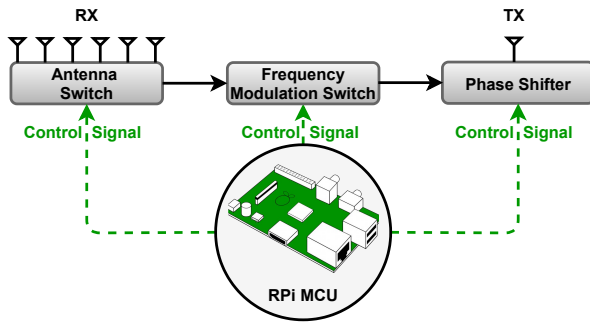


Figure 5: Schematic for RF-Protect Reflector: RF-Protect uses a simple switch and phase shifter to spoof varying distances and breathing motions.

option is to do a fixed human trajectory. However, a smart eavesdropper can easily filter this motion out by observing that such repetitive motion is not realistic for a human. As described in the threat model, the eavesdropper can also filter out naturally oscillating kinetic reflectors in the environment as deterrents, such as ceiling fans. Another option would be to do a random or noisy trajectory. However, this approach would also lead to an unrealistic phantom since human motions typically have certain characteristics that need to be captured (e.g. smoothness and continuity). Such an approach may also lead to the eavesdropper thinking that our injected signal is interference or noise, which may cause the eavesdropper to counter attack by modifying radar parameters such as frequency.

More fundamentally, as long as the distribution of spoofed trajectories is not identical to the distribution of human trajectories, there exists a classifier which can identify real vs fake trajectories with high probability. This is especially true for adversaries (e.g. smart home devices) learning the distribution of human behavior. Ideally, we want the distribution of our fake trajectories to match the distribution of human trajectories so that any classifier trying to separate the two will fail roughly half the time.

Our Design: To generate diverse trajectories that resemble human motion, we propose a conditional Generative Adversarial Network [31] architecture to generate synthetic trajectories. GANs are popular machine learning [19] tools that are used to generate *new samples* from the same distribution as the input samples. GAN architectures formulate the problem using a game-theoretic approach wherein an adversary tries to disambiguate generated samples from real data. This accurately mimics our scenario and will allow us to generate fake human trajectories that a smart eavesdropper won't be able to disambiguate from real human trajectories.

We use the conditional GAN (cGAN) variant for two reasons. First, vanilla GAN architectures suffer from model collapse problem [36], i.e., the output of the GAN is independent of the input. Second, we can tweak the conditional GAN architecture to generate trajectories of different ranges (e.g. purposeful walking, ambling, etc.). RF-Protect's cGAN architecture is shown in Fig. 6.

Dataset Collection and Classification: We collect a set of human trajectory data to train our cGAN in a large office space. We ask participants to move at will. Our experiment spans 2 hours and we construct a dataset with 7000 traces, each approximately 10

seconds long. Each trace has 50 two dimensional data points (x and y coordinates). We classify the dataset into five classes based on ranges of motion. We use the range values as input to the cGAN to enable a coarse yet effective control of the distribution of generated trajectories.

Generator Architecture: The generator aims to create new trajectories that resemble real trajectories. In the generator neural network, we input a Gaussian noise sample z to sample different trajectories. The range labels n pick the type of trajectory. Then, z and n (after embedding) are concatenated and sent into the fully connected layer. We then use a two-layer Long Short-Term Memory (LSTM) network [23] to generate the consecutive points to form a trajectory. We set the dropout probability in LSTM to be 0.5 and the hidden size to be 512. We feed the output of the LSTM into another fully connected network for reshaping.

Discriminator Architecture: The goal of the discriminator is to identify fake trajectories from real trajectories. Our cGAN samples real and fake trajectories and challenges the discriminator to classify them as such. As shown in Fig. 6, the network comprises a fully connected layer, followed by a Bidirectional LSTM [37] with hidden size 512 and dropout probability 0.5. The final output is sent into a fully connected network for reshaping followed by the Sigmoid function to identify the trajectory score, i.e. the likelihood of the trajectory being real.

Loss Function: The generator aims at learning a distribution p_g over data x . It creates a mapping from the prior noise distribution $p_z(z)$ to the data space as $G_{\theta_g}(z | n)$, conditioned on the range label n . The discriminator, $D_{\theta_d}(x | n)$, gives scalar outputs which represents the probability that x comes from the real dataset rather than p_g if it's conditioned on n . We use the following standard loss function [31] for the training:

$$\min_G \max_D V(D, G) = E_{x \sim p_{\text{data}}(x)} \left[\log \left(D_{\theta_d}(x | n) \right) \right] + E_{z \sim p_z(z)} \left[\log \left(1 - D_{\theta_d} \left(G_{\theta_g}(z | n) \right) \right) \right] \quad (4)$$

where θ_g and θ_d are the parameters for our generator and discriminator respectively.

7 PRIVACY PROTECTION ANALYSIS

How does the addition of fake humans limit the ability of an eavesdropper to gain meaningful information about the human motion in the environment? We examine this question for multiple applications below.

Occupancy Status: The simplest question an eavesdropper may be interested in is occupancy detection, i.e. is someone at home? RF-Protect can easily ensure that this question always returns a positive response. So, solving this problem is trivial for RF-Protect. This also applies for occupancy-related activity questions, e.g., when is the home occupied or when does someone go to the bathroom? RF-Protect can disable such inferences by spoofing multiple visits where only a subset of which are true visits. This is in line with past work in web-privacy where injection of fake click or query data removes user-specific references.

Breath Monitoring: With RF-Protect deployed in an environment, an eavesdropper will sense multiple different breaths, some real and

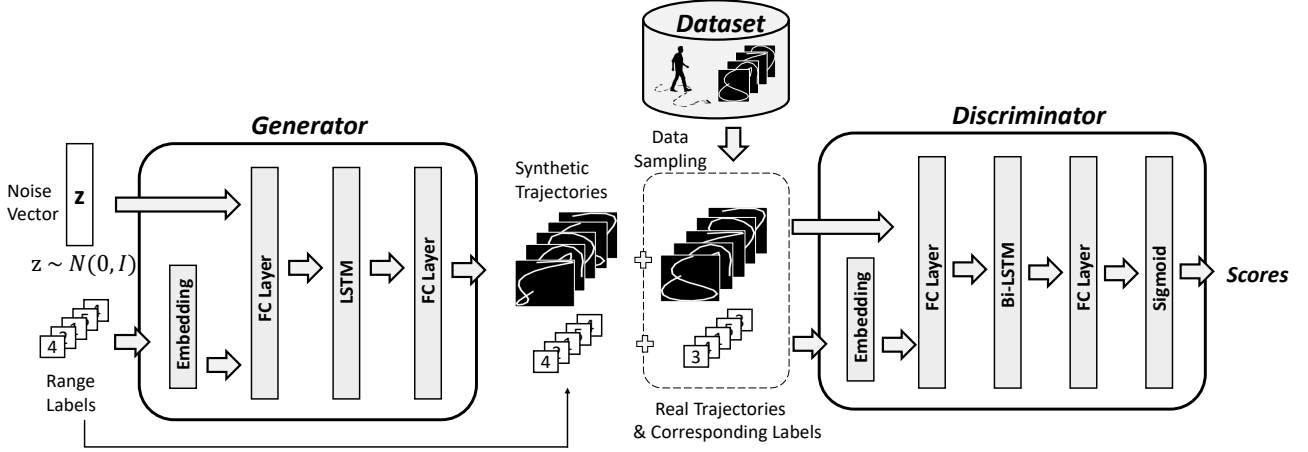


Figure 6: Trajectory GAN Structure: Generator G takes a noise vector and range labels as input and gives synthetic trajectories. Discriminator D takes the synthetic trajectories and real trajectories along with their labels as input to get the scores.




Figure 7: Mutual Information, $I(X, Z)$ between observed distribution & measured distribution decreases with the number of fake human reflections (M).

some fake. However, note that the eavesdropper has no information about real or fake trajectories. Therefore, given M fake breathing patterns and N real ones, even if the eavesdropper knows about RF-Protect’s deployment, they can at best make a random guess. The random guess will be correct with probability $\frac{N}{M+N}$.

Occupant Counting: Similar to occupancy status, RF-Protect spoofs fake humans and disrupts occupant counting. For example, if there are two people at a home, RF-Protect can have the eavesdropper sense four people.

Occupant Distribution: Finally, some eavesdroppers may not be interested in instantaneous answers (e.g. how many people are at home now), but in distribution of occupancy (e.g. how does the number of people vary through the day). Therefore, even with RF-Protect deployed, they may extract information about the distribution of occupant count (e.g. more people are at home in the evening). Similar to past work [49], we examine this question from an information theoretic perspective.

Assume that the random variable X denotes the number of real humans. Similarly, Y is the number of phantom humans generated by RF-Protect. We model X and Y as binomial random variables, i.e. $X \sim \text{Bin}(N, p)$ and $Y \sim \text{Bin}(M, q)$, where Bin denotes the binomial

distribution, p is the probability of a single human moving, q is the probability of a single reflector generating a phantom, N is the maximum occupancy of the environment and M is the maximum number of phantom humans. In RF-Protect, we can control both q and M .

Then, the number of humans seen by the adversary is $Z = X + Y$. Although the adversary’s wireless sensor measures the distribution Z , the goal of the adversary is to get information about the true distribution of occupants X . The *Mutual information* $I(X, Z)$ is a measure of the amount of information that can be inferred from X after observing the variable Z [16]. Naturally, we ask – how much information about X is leaked through Z ? Mathematically,

$$I(X, Z) = \sum_{x \in X} \sum_{z \in Z} P_{X,Z}(x, z) \log \frac{P_{X,Z}(x, z)}{P_X(x)P_Z(z)} \quad (5)$$

We derive the joint distribution, $P_{X,Z}(x, z)$, using the conditional distribution $P_{Z|X}$. Since $Z = X + Y$ and X and Y are independent, the conditional distribution $P_{Z|X}(z|x)$ is simply the probability distribution of Y . We derive the marginal distribution $P_Z(z)$ by summing $P_{X,Z}(x, z)$ across $x \in X$. Lastly, solving for $I(X, Z)$ yields the following:

$$I(X, Z) = \sum_{x=0}^N \sum_{z=x}^{N+M} P_{Z|X}(z|x) P_X(x) \log \frac{P_{Z|X}(z|x) P_X(x)}{P_X(x) P_Z(z)} \quad (6)$$

Fig. 7 plots the variation of $I(X, Z)$ for different values of M for a home with 4 occupants and $p = 0.2$ (a higher estimate for moving humans in a space). First, note that when $q = 0$ and $q = 1$, the mutual information is high. When we set $q = 1$, the reflector is always on, so the distribution of true occupant count and observed count is perfectly in sync (the real values are incorrect though). However, when q is close to 0.5, the mutual information is much lower. This information gets worse as we get the ability to spoof more humans, which is possible in our setup given the multiple antennas and the ability to deploy multiple reflectors.

8 DISCUSSION AND LIMITATIONS

We discuss some aspects of RF-Protect’s design below:

- **Side Channel Information:** An eavesdropper can identify RF-Protect’s defense if they have side channel information about the home. For example, if they visit the house temporarily, verify who is in the home, and later correlate this information with the output for the FMCW radar. However, this information will be futile after a short duration of time as additional trajectories are added and/or removed.
- **Changing Radar Parameters:** We note that RF-Protect’s trajectories are not invariant to FMCW parameters and location. Within standard operating range, change in FMCW slope and radar position leads to scaling and rotation. However, such trajectories would still appear to be genuine human trajectories. Finally, FMCW radars may operate in multiple frequency bands. However, the higher frequency systems that operate above 10 GHz don’t easily penetrate walls. RF-Protect’s design is relatively wideband and can sustain shifts in frequencies in sub 10-GHz range.
- **Incorporating Floor Plan Information:** A current limitation with RF-Protect is that the phantom trajectories we generate using the cGAN can move anywhere within the 2D space that the hardware reflector can cover. However, if the eavesdropper has prior knowledge about the floor plan of the building, some of RF-Protect’s spoof trajectories may unintentionally "walk through walls." In future work, we plan to incorporate floor plan knowledge into the cGAN’s trajectory generation process such that the trajectories will rather move around walls rather than through them. The high-level idea here would be to add the floor plan bounds as a conditional input to the cGAN and modify the loss function to penalize trajectories generated that move through walls of the floor plan.
- **Radars with 2D Antenna Arrays:** RF-Protect currently assumes the eavesdropper uses a 1D antenna array. However, the eavesdropper can potentially use a 2D antenna array to localize motion in a 3D space. In response, we can extend RF-Protect by adding a 2D antenna panel that can spoof trajectories in both the elevation and azimuth directions. We envision using simultaneous antenna transmissions with the 2D panel can help generate 3D phantoms that will address the 2D antenna array radar threat model.
- **Simultaneously Reflecting from Multiple Antennas:** RF-Protect currently performs angular spoofing by reflecting on a single antenna from the switched array deployed on the hardware reflector. Therefore, this design spoofs discrete angle values. To improve RF-Protect’s angle spoofing ability, we can reflect using multiple antennas on the reflector. Incorporating this idea will enable RF-Protect to emulate finer, more continuous trajectories from the angular spoofing perspective. This would also make our design more robust against radars that use more advanced software processing techniques to achieve a finer granularity of angular resolutions than the standard π/k separation mentioned in section 5.2.
- **Radar Cross Section:** An eavesdropper can analyze the radar cross section to potentially determine whether a reflection is an actual human or a ghost trajectory. The radar cross section is

typically influenced by factors such as the size and material of the reflector. As future work, we plan to incorporate reflective power variation into RF-Protect’s GAN architecture to mimic human distribution.

9 EXPERIMENTAL SETUP

9.1 FMCW Radar Design

Hardware: We build a custom FMCW radar prototype operating over the same frequency band (and similar configuration) as [4–6]. We use a Texas Instruments LMX2492EVM chirp generator [1] to generate our FMCW waveform. The generated chirp sweeps from 6 – 7 GHz over a span of 500 μ s. This band is designated by the FCC for civilian use of the spectrum [3]. To obtain angle, we also implement an antenna array on the radar comprising of seven antennas.

Processing Pipeline: We mix the received reflection signal with a copy of the original chirp to extract the *beat signal* along each antenna. We call this matrix of seven beats a frame. We obtain range and angle from this frame using the method described in Sec. 3. We subtract successive frames to remove static reflectors. Upon completing this process, we obtain the power profiles across distance and angles for frames across time, where peaks in the profiles represent human motion as shown in Fig. 10a. As past work has noted [4, 5], the peaks can be sporadic with intermittent noise. Therefore, we perform smoothing over time and peak rejection to extract human trajectories, as is standard in radar processing. Artifacts on our experiments are available and open sourced.³

9.2 RF-Protect Reflector Design

Fig 5 provides an overview of the RF-Protect hardware reflector design. We use 6 directional antennas as the angle spoofing panel to receive the eavesdropper’s FMCW transmitted signal. The antenna separation used in our experiments is roughly 20cm. Based on the the angle we wish to spoof, we control the antenna panel using an SP8T RF switch implemented by EV1HMC345ALP3 [2]. Afterwards, this signal is sent through another switch that modulates the frequency of the signal for distance spoofing. We amplify this signal using an LNA and send it out through a TX antenna back at the eavesdropper radar. The control logic for the switch is implemented on a Raspberry Pi microcontroller.

Conditional GAN Training: We implement our trajectory GAN using the PyTorch deep learning framework. The learning rates of the generator and discriminator are 0.0001 and 0.0002 respectively. We use the Adam optimizer with mini-batches of size 128. The training process takes 5 hours on an Nvidia RTX 1080Ti GPU. RF-Protect doesn’t need per-location training, so the model needs to be trained only once.

9.3 Evaluation

We evaluate RF-Protect in two indoor settings – a home environment and an office environment. In each environment, we deploy the eavesdropper radar and RF-Protect reflector in non-line of sight of each other (as would be done in a realistic scenario). Fig. 8b and

³<https://github.com/ConnectedSystemsLab/rf-protect>

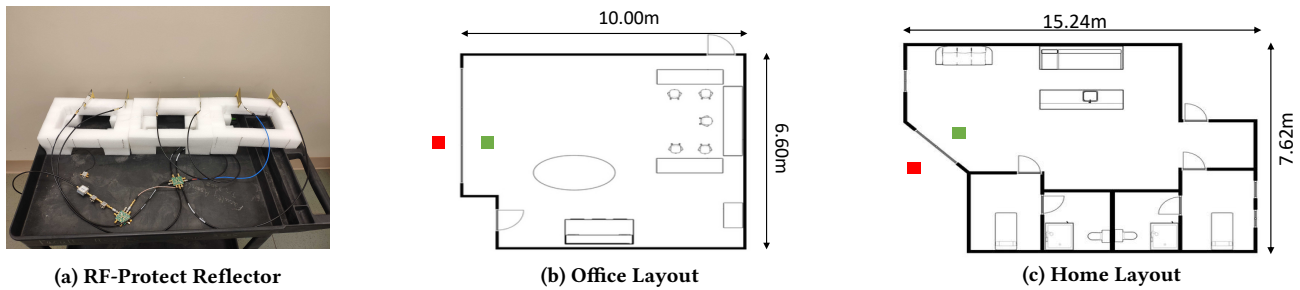


Figure 8: (a) RF-Protect reflector. (b,c) The floor plans of (b) an office and (c) home layout are shown above. The red square denotes the eavesdropper radar and the green square denotes the RF-Protect reflector.




Figure 9: FMCW Radar Localization: Our FMCW Radar prototype can accurately locate users moving along different shapes (representative graphs shown). Blue lines are trajectories detected by our FMCW Radar and red dots are ground truth references.

8c outline the floor plan and device setup locations used for our experiments. These environments have typical multipath and human motion, and span 6.6×10 m (21.65×32.8 ft) and 15.24×7.62 m (50×25 ft) respectively. We spoof a total of 90 different trajectories, 45 for each environment. The distance between the radar and the reflector is approximately 1.2 meters.

10 MICROBENCHMARKS

We present some microbenchmarks for RF-Protect below.

10.1 FMCW Radar Performance

We start by evaluating the performance of our FMCW radar design. We run two different radar measurement experiments in the office environment. In each of the experiments, we have a single human subject walk around in a different trajectory. Fig. 9 demonstrates the results of our experiments. We label the measured ground truth points (red circles) that the subject walked on for each of the experiments. We also overlay the trajectory measured by the radar (blue trace). For these experiments, the radar’s antennas are located at the origin (0,0). The measured radar trajectory closely matches the ground truth labels for all three experiments. The high localization accuracy indicates that our radar is successfully replicating state-of-the-art FMCW radar design.

10.2 Reflector Design

RF-Protect aims to design a hardware reflector that can create dynamic reflections that move through time. As a result, these reflections should not be cancelled out by background subtraction. In Fig. 10b, we plot a range-angle profile sensed by the radar after

background subtraction. The range-angle profile of the phantom produced by RF-Protect is indistinguishable to the profile of the human reflector in Fig. 10a. Observe how the reflection power of RF-Protect is identical to the actual human reflection because the hardware reflector simply reflects the radar signal without adding any additional signal of its own. Moreover, secondary reflections around the phantom appearing due to dynamic multipath are also prevalent in our spoofed range-angle profile.

10.3 Spoofed Trajectory

Fig. 10c illustrates one of the trajectories used in our office room evaluation. We compare the ground truth trajectory created by our cGAN to the trajectory measured by the radar while RF-Protect’s reflector is deployed and actively spoofing the cGAN’s generated trajectory. Observe that the generated trajectory closely follows the expected trajectory. The relative shape of the trajectory also remains intact. The trajectory spans total motion of nearly 20 feet, showing that RF-Protect can generate long trajectories across a home and office.

11 RESULTS

11.1 Accuracy of 2D Spoofing

We evaluate the end to end implementation of RF-Protect by measuring spoofing performance. For both the home and office environments, we use our cGAN to generate 45 trajectories each and then program our reflector to spoof these trajectories. We start by first converting the trajectories into polar-coordinate representation (distance and angles) with the reflector at the origin. The reflector then modifies the eavesdropper radar’s signals based on these generated distance and angle measurements. Recall, the goal of RF-Protect is to spoof the relative trajectory produced by the cGAN rather than the absolute location since the relative trajectory still resembles actual human motion. Therefore, we measure the metrics below modulo translation and rotation of the entire trajectory. Our results are plotted in Fig. 11.

Distance: First, in Fig. 11a, we plot the deviation in the distance (or polar radius) measured by the eavesdropper and our intended spoofed distance. The median distance error is 5.56 cm and 10.19 cm in home and office respectively. Given that the localization resolution of the FMCW radar implementation is 15 cm, the error in our spoofed trajectory is within 1 bin of the range FFT.




Figure 10: Microbenchmarks: RF-Protect generates fake reflections (b) which resemble actual human reflections (a), as observed by an FMCW Radar. (c) RF-Protect uses a sequence of such reflections to create trajectories.




Figure 11: RF-Protect's Accuracy: (a)(b) RF-Protect provides accurate spoofing in angle and distance. (c) Relative Trajectory Error between fake trajectories detected by the radar and the ground truth trajectory.

Angle: In Fig. 11b, we plot the deviation in the angle measured by the eavesdropper and our intended spoofed angle. There are two factors to consider here. The RF-Protect tag is capable of constructing discrete angles only. Second, the FMCW radar is limited in terms of its angular resolution based on the size of its antenna array. Overall, we observe that the median error between the angle measured at the radar and the spoofed angle is 2.05 degrees and 4.94 degrees for the home and office environments respectively.

Location Error: Fig. 11c plots the 2-D location error of our experiments. The median error in spoofing the absolute 2-D location of a reflection is 12.70 cm for home environment and 24.49 cm for office environment. Observe that the errors in the office environment are larger. This is because the FMCW radar performs worse in that environment due to presence of metallic cabinets that cause multipath. We observe high errors in human location in that environment as well. This indicates that RF-Protect's reflections undergo similar effects as humans, and hence appear realistic also in terms of variance to an eavesdropper. Overall, our results indicate that the phantom reflections generated by RF-Protect's hardware reflector are accurate (comparable to the FMCW radar range resolution of 15 cm). These results suggest that RF-Protect's hardware reflector provides a new primitive to generate RF reflections at arbitrary points in the environment.

11.2 cGAN Performance

Qualitative: Now, we ask how realistic are the trajectories generated by the RF-Protect cGAN model. We first plot some representative generated trajectories (solid lines) and real trajectories (dashed lines), as shown in Fig. 12. Qualitatively, both the real and fake trajectories have similar characteristics such as smoothness and continuity.

Quantitative: We employ the Fréchet Inception Distance (FID) to evaluate the cGAN [22]. FID is a standard metric to evaluate the output of Generative Adversarial networks [10, 42, 45]. FID captures the distance between the distance between the distribution of the real data $p_r(\cdot)$ and the distribution of generated data $p_g(\cdot)$. For ease of exposition, we plot a normalized version of FID in Fig. 12(right). Specifically, we take two sets of real human trajectories and compute the FID between them. We divide by this value to compute normalized FIDs. As shown, the real dataset has a normalized FID of 1 and is closely followed by RF-Protect with a normalized FID of 1.229. We compare RF-Protect with these baselines – single trajectory performed by a user repeatedly, uniform linear motion between points, and random motion. Unlike RF-Protect, these trajectories do not capture the distribution of human motion and have normalized scores of 1.867, 2.022, and 3.440 respectively.

User-Study: Following established practices, we also conduct a survey asking humans to distinguish between real-world human trajectories and RF-Protect's generated trajectories. In this experiment, we present to each of 32 participants 5 real trajectories and




Figure 12: (Left) Some samples of real (top row) and generated trajectories (bottom row). (Right) Normalized FID scores show that RF-Protect’s GAN generates close-to-real trajectories compared to other baselines.

#Instances	Real	Fake
Perceived as real	93	89
Perceived as fake	67	71

Table 1: Human study results: RF-Protect’s spoofed trajectories can not be successfully identified by humans.

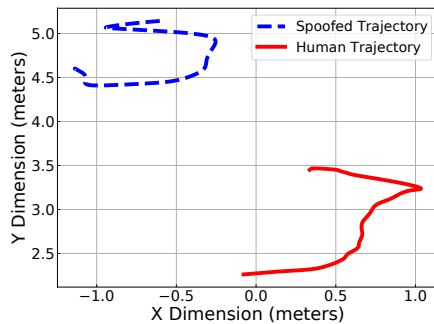


Figure 13: An FMCW sensor can sense both fake and real trajectories. By conveying the RF-Protect-generated trajectory to a legitimate sensor, we enable legitimate sensing, while disabling eavesdroppers.

5 fake trajectories that are randomly shuffled. We asked them to identify each trajectory as real or fake. The response is summarized in Table 1. A Pearson χ^2 test on the survey results yields a $\chi^2 = .2$ and $p = .65$. The χ^2 test indicates that the trueness and the trueness perceived by human do *not* have a statistically significant correlation.

Based on our qualitative, quantitative, and human experiments, we conclude that we are able to generate trajectories that are not likely to be distinguishable from real human motion. Note that this indicates to RF-Protect’s critical advantage – an adversary won’t know that they are being deceived without considering side-channel information.

11.3 Enabling Legitimate Sensing

In Fig. 13, we demonstrate the ability of a legitimate sensor to successfully decode trajectories in presence of RF-Protect. Specifically, RF-Protect injects a fake trajectory (blue), while a real human motion happens in a different part of the room. A legitimate FMCW sensor observes both trajectories and can filter the blue trajectory out by communicating with RF-Protect tag. Recall that unlike prior defense mechanisms against wireless sensing such as Faraday cages and jamming which prohibit legitimate sensors in the environment, RF-Protect serves as a reliable defense without disrupting the functionality of legitimate sensors.




Figure 14: Breathing Rate Spoofing Results: RF-Protect can spoof the breathing rate by mimicking the phase of the real human breathing.

11.4 Breathing Rate Spoofing

Prior advanced FMCW radar designs have demonstrated the ability to measure even the smallest movements such as human breathing through chest movement [6]. To accomplish this, these systems analyze periodic variations in received antenna phase to compute the breathing patterns of humans in the environment. We note these past works rely on the underlying breathing signal to infer many other private metrics such as sleep patterns, emotional state, and health profiles. We evaluate RF-Protect’s ability to also spoof these small breathing motions. We accomplish this by using an analog phase shifter to periodically modify the signal passing through the tag, as shown before. Fig. 14 highlights the similarities between actual human breathing phase measured by the radar and the phase measured by the radar from RF-Protect’s spoof breathing. As shown in the figure, RF-Protect’s hardware reflector can accurately mimic human breathing motion. Our ability to spoof human breathing enables us to defend against the large variety of wireless sensors that rely on the breathing signal to infer even more private information.

12 RELATED WORK

FMCW Radar Spoofing: Due to the popularity of FMCW radars for vehicular sensing, past work has focused on failure modes through distance spoofing attacks [28, 32, 33, 43]. In these attacks, the attacker senses the radar signal and replays the signal with some delay. First, note that these attacks require complex circuitry to receive, synchronize with, and replicate high bandwidth FMCW

signal. These attacks require the spoofing device to have tight synchronization with the sensing FMCW radar. Second, these attacks are easy to detect. Unlike RF-Protect they are not real reflections. These spoofing attacks can be detected by periodically turning off the sensing radar [27]. The spoofing device needs some time to stop transmission and to synchronize with new transmissions when the sensing radar turns on. In contrast, RF-Protect generates real reflections that turn off automatically when the radar turns off. Finally, these spoofing mechanisms are limited to spoofing distance, but do not spoof actual location, nor do they create realistic trajectories like RF-Protect.

Location Privacy: There has been a lot of past work [13, 15, 38–40] on enabling privacy for location-based services. This line of work maps location information X to perturbed location $f(X)$, and shares this data with location-based service providers (like mobile applications). The goal is to design the function f such that it achieves the intended utility without leaking too fine-grained location information of individuals. The design of f uses various techniques like obfuscation and data injection. Our work is motivated by this line of research, yet differs in terms of its setting and techniques. RF-Protect focuses on passive sensing using FMCW radars, not the setting where the user willingly shares their location data. Furthermore, RF-Protect does not have access to the raw location that an eavesdropper FMCW radar is sensing. Therefore, we must tweak this location information at the physical layer using a new tag design. Finally, we develop a new approach to generate realistic human trajectories.

Physical Layer Security & Privacy: PhyCloak [35] and RF-Cloak [21] focus on the problem of physical layer security mechanisms for communications [21] and sensing [35]. RF-Cloak focuses on preventing eavesdroppers to read RFID data. PhyCloak is the closest to our work in that it tries to protect privacy sensing using signals meant for communication. However, both RF-Cloak and PhyCloak focus on relatively narrowband communication signals like Wi-Fi. In contrast, RF-Protect provides privacy against wideband FMCW radars. Moreover, their approach scrambles the wireless channel at the eavesdropper. This defense is also easy to detect (and thus, counter – e.g. by deploying additional antennas) for the eavesdropper. RF-Protect presents a different approach – the eavesdropper sees fake human reflections and hence, cannot detect the presence of a defense mechanism. Finally, RF-Protect generates new ‘realistic’ human trajectories, a feature that neither of these works have.

13 CONCLUSION

We present RF-Protect the first smart reflector design that spoofs fake trajectories against eavesdroppers who use FMCW radars to monitor private activities through walls and obstacles. Unlike prior solutions, our system design is simple and does not require complex signal processing. Our experiments show that RF-Protect can accurately spoof human-realistic trajectories. We believe that RF-Protect raises the bar for privacy against through-occlusion sensing systems. We envision RF-Protect being included in future smart surfaces to deliver privacy from different RF sensing modalities.

We expect future work to consider following axes of research:

- **New Sensing Applications:** The systemic and algorithms contributions of RF-Protect are generalizable to other RF sensing applications. For instance, autonomous vehicles use FMCW radars for distance tracking, speed tracking, and even imaging [18, 47]. A RF-Protect reflector deployed outdoors can disrupt such sensing by creating fake reflections and necessitates a discussion about securing these radars against such adversarial attacks.
- **New Sensor Types:** We have focussed on FMCW radars due to their popularity for indoor and outdoor sensing applications. Other kinds of radar like pulsed radars are prone to similar defenses. While we leave such exploration to future work, we note that the primitives of generative trajectories that are similar to humans extend to such radars. However, distance spoofing in such radars need to be achieved through other mechanisms (e.g. by adding a set of delay lines and switching between them)
- **Extension to Smart Devices:** We presented RF-Protect in the context of an out-of-home eavesdropper. However, smart devices may embed such tracking inside the case and sense human motion and health metrics. We envision an extension of RF-Protect designed as a ring around such smart devices, as opposed to a deployment on the wall. This would include formulating the hardware reflector design in different frequency bands, but the principles of distance spoofing and cGAN-based trajectory generation will translate easily to such a design.
- **Extended Threat Model:** RF-Protect assumes that the eavesdropper uses a single radar to sense the private environment. However, it is possible for the eavesdropper to use multiple radars in coordination to determine which reflections are spoofed and which are real. If the eavesdropper deploys multiple radars against all boundaries of the environment, a single RF-Protect reflector would likely not be able to deceive the eavesdropper since the generated trajectories would have a limited angular range from the perspective of at least one of the coordinating radars. Future work can address this updated threat model by identifying the positioning and control of multiple RF-Protect reflectors in an environment that can protect against any configuration of radar deployment.

ETHICS STATEMENT

This work does not raise any ethical issues.

ACKNOWLEDGMENTS

We thank anonymous reviewers and our shepherd, Fadel Adib, for their valuable feedback. We thank Jayden Guan and Mengze Sha for their mentorship and advice on the hardware design. We are grateful to Haitham Hassanieh, Elahe Soltanaghai, and Dinesh Bharadia for providing feedback on early versions of this paper. We also thank Rem Yang for providing feedback on the final version of the paper. Jayanth Shenoy is funded by the NSF Graduate Research Fellowship Program under Grant No. DGE-1746047.

REFERENCES

- [1] TI LMX2492EVM. <https://www.ti.com/>. Texas Instruments.
- [2] EV1HMC345ALP3 switch. <https://www.analog.com/>. Analog Devices.
- [3] Understanding the fcc regulations for low-power, nonlicensed transmitters. office of engineering and technology federal communications commission, 1993.
- [4] Fadel Adib, Zachary Kabelac, and Dina Katabi. Multi-person localization via rf body reflections. In *Usenix NSDI*, 2015.
- [5] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C. Miller. 3d tracking via body radio reflections. In *Usenix NSDI*, 2014.
- [6] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. Smart homes that monitor breathing and heart rate. In *ACM CHI 2015*.
- [7] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 837–846, 2015.
- [8] Mostafa Alizadeh, George Shaker, João Carlos Martins De Almeida, Plinio Pelegrini Morita, and Safeddin Safavi-Naeini. Remote monitoring of human vital signs using mm-wave fmcw radar. *IEEE Access*, 7:54958–54968, 2019.
- [9] Amazon.com Services, LLC. 20210622 amazon rule 15.255(c)(3) waiver request. [https://ecfsapi.fcc.gov/file/10622234731686/20210622%20Amazon%20Rule%2015.255\(c\)\(3\)%20Waiver%20Request.pdf](https://ecfsapi.fcc.gov/file/10622234731686/20210622%20Amazon%20Rule%2015.255(c)(3)%20Waiver%20Request.pdf), 2021. Accessed: 2021-09-14.
- [10] Siddarth Asokan and Chandra Sekhar Seelamantula. Teaching a gan what not to learn, 2020.
- [11] Kshitiz Bansal, Keshav Rungta, Siyuan Zhu, and Dinesh Bharadia. Pointillism: Accurate 3d bounding box estimation with multi-radars. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems, SenSys '20*, 2020.
- [12] Vincent Bindschaedler and Reza Shokri. Synthesizing plausible privacy-preserving location traces. In *2016 IEEE Symposium on Security and Privacy (SP)*, 2016.
- [13] Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Optimal geo-indistinguishable mechanisms for location privacy. In *ACM SIGSAC Conference on Computer and Communications Security*, ACM CCS, New York, NY, USA, 2014. Association for Computing Machinery.
- [14] Samuele Capobianco, Luca Facheris, Fabrizio Cuccoli, and Simone Marinai. Vehicle classification based on convolutional networks applied to fmcw radar signals. In *Italian Conference for the Traffic Police*, pages 115–128. Springer, 2017.
- [15] Richard Chow and Philippe Golle. Faking contextual data for fun, profit, and privacy. In *ACM Workshop on Privacy in the Electronic Society*, ACM WPES, 2009.
- [16] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006.
- [17] Chuanwei Ding, Hong Hong, Yu Zou, Hui Chu, Xiaohua Zhu, Francesco Fioranelli, Julien Le Kermec, and Changzhi Li. Continuous human motion recognition with a dynamic range-doppler trajectory method based on fmcw radar. *IEEE Transactions on Geoscience and Remote Sensing*, 57(9):6821–6831, 2019.
- [18] Florian Folster, Hermann Rohling, and Urs Lubbert. An automotive radar network based on 77 ghz fmcw sensors. In *IEEE International Radar Conference, 2005.*, pages 871–876. IEEE, 2005.
- [19] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [20] Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. Through fog high-resolution imaging using millimeter wave radar. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [21] Haitham Hassanieh, Jue Wang, Dina Katabi, and Tadayoshi Kohno. Securing rfids by randomizing the modulation and channel. In *USENIX Symposium on Networked Systems Design and Implementation*, NSDI, 2015.
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [23] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [24] Texas Instruments. Awr1642 evaluation module single-chip mmwave sensing solution user's guide (rev. c). https://www.ti.com/lit/ug/swru508c/swru508c.pdf?ts=1631686675007&ref_url=https%253A%252F%252Fwww.ti.com%252Ftool%252FAWR1642BOOST, 2020. Accessed: 2021-09-14.
- [25] Texas Instruments. Iwr1642 evm (iwr1642boost) single-chip mmwave sensing solution user's guide (rev. c). https://www.ti.com/lit/ug/swru521c/swru521c.pdf?ts=1631686792288&ref_url=https%253A%252F%252Fwww.ti.com%252Ftool%252FIWR1642BOOST, 2020. Accessed: 2021-09-14.
- [26] Texas Instruments. 77ghz single chip radar sensor enables automotive body and chassis applications. Technical report, 2021.
- [27] Prateek Kapoor, Ankur Vora, and Kyoung-Don Kang. Detecting and mitigating spoofing attack against an automotive radar. In *IEEE Vehicular Technology Conference (VTC-Fall)*, 2018.
- [28] Rony Komissarov and Avishai Wool. Spoofing attacks against vehicular FMCW radar. *CoRR*, abs/2104.13318, 2021.
- [29] Jaime Lien, Nicholas Gillian, M. Emre Karagozler, Patrick Amihood, Carsten Schwesig, Erik Olson, Hakim Raja, and Ivan Poupyrev. Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Trans. Graph.*, 2016.
- [30] Google LLC. Fcc: Soli request for waiver. <https://www.fcc.gov/document/google-llc-request-waiver-part-15-project-soli>, 2018. Accessed: 2021-09-14.
- [31] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [32] Noriyuki Miura, Tatsuya Machida, Kohei Matsuda, Makoto Nagata, Shoeni Nashimoto, and Daisuke Suzuki. A low-cost replica-based distance-spoofing attack on mmwave fmcw radar. In *ACM Workshop on Attacks and Solutions in Hardware Security Workshop*, ACM ASHES, 2019.
- [33] Shoeni Nashimoto, Daisuke Suzuki, Noriyuki Miura, Tatsuya Machida, Kohei Matsuda, and Makoto Nagata. Low-cost distance-spoofing attack on fmcw radar and its feasibility study on countermeasure. *Journal of Cryptographic Engineering*, 2021.
- [34] Octavian Postolache, Pedro Silva Girão, Rui Neves Madeira, and Gabriela Postolache. Microwave fmcw doppler radar implementation for in-house pervasive health care system. In *2010 IEEE International Workshop on Medical Measurements and Applications*, pages 47–52. IEEE, 2010.
- [35] Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora. Phycloak: Obfuscating sensing from communication signals. In *USENIX Symposium on Networked Systems Design and Implementation*, NSDI, 2016.
- [36] Eitan Richardson and Yair Weiss. On gans and gmms. *arXiv preprint arXiv:1805.12462*, 2018.
- [37] M. Schuster and K.K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [38] Reza Shokri, George Theodorakopoulos, Panos Papadimitratos, Ehsan Kazemi, and Jean-Pierre Hubaux. Hiding in the mobile crowd: Locationprivacy through collaboration. *IEEE Transactions on Dependable and Secure Computing*, 2014.
- [39] Reza Shokri, George Theodorakopoulos, Carmela Troncoso, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. Protecting location privacy: Optimal strategy against localization attacks. In *ACM Conference on Computer and Communications Security*, ACM CCS, 2012.
- [40] Reza Shokri, Carmela Troncoso, Claudia Diaz, Julien Freudiger, and Jean-Pierre Hubaux. Unraveling an old cloak: K-anonymity for location privacy. In *ACM Workshop on Privacy in the Electronic Society*, ACM WPES '10, 2010.
- [41] Aman Shrestha, Haobo Li, Julien Le Kermec, and Francesco Fioranelli. Continuous human activity classification from fmcw radar with bi-lstm networks. *IEEE Sensors Journal*, 20(22):13607–13619, 2020.
- [42] Rajhans Singh, Pavan Turaga, Suren Jayasuriya, Ravi Garg, and Martin W. Braun. Non-parametric priors for generative adversarial networks, 2019.
- [43] Zhi Sun, Sarankumar Balakrishnan, Lu Su, Arupjyoti Bhuyan, Pu Wang, and Chunming Qiao. Who is in control? practical physical layer attack and defense for mmwave-based sensing in autonomous vehicles. *IEEE Transactions on Information Forensics and Security*, 2021.
- [44] Vincent Toubiana, Lakshminarayanan Subramanian, and Helen Nissenbaum. Trackmenot: Enhancing the privacy of web search, 2011.
- [45] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space, 2021.
- [46] Udi Weinsberg, Smriti Bhagat, Stratis Ioannidis, and Nina Taft. Blurme: Inferring and obfuscating user gender based on ratings. In *Proceedings of the Sixth ACM Conference on Recommender Systems*, RecSys '12, 2012.
- [47] Volker Winkler. Range doppler detection for automotive fmcw radars. In *2007 European Radar Conference*, pages 166–169. IEEE, 2007.
- [48] Jing Yang, Xiaoxu Guo, and Yunjie Li. Design of a novel drfm jamming system based on afb-sfb. In *IET International Radar Conference*, 2013.
- [49] Shaozhi Ye, Felix Wu, Raju Pandey, and Hao Chen. Noise injection for search privacy protection. In *2009 International Conference on Computational Science and Engineering*, 2009.
- [50] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. Hitchhike: Practical backscatter using commodity wifi. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, pages 259–271, 2016.
- [51] Mingmin Zhao, Shichao Yue, Dina Katabi, Tommi S Jaakkola, and Matt T Bianchi. Learning sleep stages from radio signals: A conditional adversarial architecture. In *International Conference on Machine Learning*, pages 4100–4109. PMLR, 2017.