PhyCloak: Obfuscating Sensing from Communication Signals

Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora,
The Ohio State University

ISBN 978-1-931971-29-4

PhyCloak: Obfuscating Sensing from Communication Signals

Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan and Anish Arora
Department of Computer Science and Engineering
The Ohio State University
{*qiaoyu, zhouwe, kannan, anish*}*@cse.ohio-state.edu*
zhang.4746@buckeyemail.osu.edu

ABSTRACT

Recognition of human activities and gestures using pre-existing WiFi signals has been shown to be feasible in recent studies. Given the pervasiveness of WiFi signals, this emerging sort of sensing poses a serious privacy threat. This paper is the first to counter the threat of unwanted or even malicious communication based sensing: it proposes a blackbox sensor obfuscation technique PhyCloak which distorts only the physical information in the communication signal that leaks privacy. The data in the communication signal is preserved and, in fact, the throughput of the link is increased with careful design. Moreover, the design allows coupling of the PhyCloak module with legitimate sensors, so that their sensing is preserved, while that of illegitimate sensors is obfuscated. The effectiveness of the design is validated via a prototype implementation on an SDR platform.

1 Introduction

A new form of threat has emerged recently that leaks private information about the whereabouts and activities of physical targets merely by observing the ongoing wireless communications in the scene. Broadly speaking, as a wireless signal gets reflected off of people and other objects in the scene, information about them is leaked to eavesdroppers by computational analysis of the signal distortions. Increasingly, researchers have been demonstrating proofs of concept where not only people presence but also fine-grain information about their locations and even breathing, lip movement or keystrokes is leaked [18, 28, 30, 1, 24]—all from observing communication signals that are widely prevalent in our homes. While the upside is that legitimate users can detect these physical "signatures" simply using existing signals, a burglar can also detect that there are no people in a house, a passerby can decipher key presses without leaving a trace [8], and a neighbor can snoop on the activities in our homes [30]. There is little doubt that several of these privacy exploits will in due course be realized robustly and commoditized for broad use. And, given the pervasive nature of wireless communications, the privacy implications of such attacks will undoubtedly be of major social importance.

It is thus timely and important to develop suitable counter-measures for this type of privacy leakage. We take the first step at tackling this problem by proposing a solution to address a single-antenna eavesdropping sensor. At first glance, it might appear that an obvious way to prevent or deter the privacy leakage is to simply jam the signals [21, 11]. However, jamming is an overkill for this problem, as the protection we wish lies in physical and not in the logical (data) layer. Jamming distorts the information of both layers, therefore it hurts the channel capacity of the network. In contrast to jamming, our approach is to distort the physical information that is environmentally superimposed on the signal as opposed to the data itself. *To make clear the distinction between these two forms of signal distortion, we refer to the latter as signal obfuscation*.

To avoid any modification of existing receivers, we need to build an obfuscator (Ox) that works independently from a receiver (Rx) and can yet deter privacy leakage against a single-antenna eavesdropper. At the same time, Ox should not hurt the ongoing reception at the intended receiver. In addition, given the diversity of the design of RF based sensors and invisibility of eavesdroppers, it is not reasonable to assume Ox that uses a specific obfuscation approach against a specific Eve. Thus, our goal is to build a black-box solution which distorts only the privacy sensitive information while not affecting the logical information. We design Ox by answering the two questions below:

1. How to distort physical information regardless of the RF-sensing mechanism? To answer this question, let us first examine what kind of physical information is contained in RF signals. Assume the received signal at a reflector is $s(t)$, then the received signal $r(t)$

reflected by the reflector can be expressed as follow: $r(t) = a \times s(t) \times e^{j2\pi(f_c+\Delta f)(t+\Delta t)}$, where a is the amplitude gain, f_c is the carrier frequency, Δf is the Doppler shift caused by a reflector that moves at a constant speed relative to the receiver, and Δt is the delay due to transmission over the path. Here, we can see that the reflector modifies the reflected copies by controlling three orthogonal components: amplitude gain a, delay Δt and Doppler shift Δf. All the features exploited by single-antenna RF based sensors are created by these three degrees of freedom (DoFs). Hence, if an Ox distorts the three orthogonal bases respectively, any features that reveal physical information are distorted too.

2. *How to preserve logical information (data communication)*? As the previous observation suggests, Ox needs to change the 3 degrees of freedom (DoFs) of a signal in order to deter eavesdropping of physically sensed features. Note that in a wireless environment, signals traverse through many paths and experience Doppler shifts: These effects are similar to dynamic multipath reflections. Thus, Ox can be a relay node that introduces dynamically changing multipath components of the communication signal. In other words, Ox receives the incoming communication signal, manipulates the signals and forwards them back to the environment. To a legitimate receiver, this forwarded signal will simply look like a multipath component of the signal from the legitimate transmitter (Tx). Commercial off-the-shelf (COTS) Rx is capable of tolerating and even exploiting multipath reflections to decode data. Thus, a carefully designed Ox can distort sensing and still preserve communication.

Challenges: PhyCloak works as a full-duplex amplify-and-forward (A&F) relay at logic layer, and an Ox at physical layer by distorting the 3 DoFs. While the solution may appear at first blush to be a simple instance of full-duplex A&F forwarder [6, 3], there are key challenges that arise from this design that need to be resolved.

1. *Online self-channel estimation with an ongoing external transmission*: Online self-channel estimation is needed for an Ox as it works in an environment where the channel is varying as a result of target movement, gestures and activities. When we combine the Ox module with a legitimate sensor the self-channel variation becomes more significant due to the moving object close to the sensor. Therefore an Ox has to transmit training symbols to acquire channel estimation every channel coherence interval ($\sim$100ms). But a complication arises that the training needs to co-exist with ongoing data transmission. A straightforward way to overcome this problem is to adopt medium access control (MAC), however, that would introduce contention and hurt throughput of legitimate data transmission given the frequent self-channel updates.

2. *Effectiveness of obfuscating physical information*: No work has been done in validating a full-duplex A&F forwarder's capability of controlling physical information contained in the forwarded copy. In addition, the effectiveness of superposing an Ox's distorted signal and a target's reflected signal in obfuscating an eavesdropping sensor has yet to be shown.

Contributions: We propose PhyCloak to protect privacy information from unwanted or even malicious sensing with no modification to existing wireless infrastructures. In this work, we make the following contributions:

1. To our knowledge, we are the first to address the potential threats due to the recent development of communication-based sensing.

2. We propose PhyCloak, the first full-duplex forwarder-based solution that hides physical information superimposed by the channel via adding interference in a 3-dimensional orthogonal basis so that illegitimate sensing is disabled and meanwhile data transmission is not affected (and even improved). We go further and add the capability to spoof human gestures to further confuse illegitimate sensors.

3. We propose an alternative online self-channel estimation scheme that is contention-free and operates in the presence of an ongoing transmission. By doing so we also allow for legitimate sensing by integrating the sensor with our obfuscator.

4. We build a prototype PhyCloak on PXIe-1082, an SDR platform. Experimental results (Section 5.3) on a state-of-the-art sensor show that PhyCloak successfully obfuscates illegitimate sensing, enables legitimate sensing and improves overall throughput of data transmission. Gesture spoofing to the same type of sensor is also proved to be feasible.

2 Related Work

RF sensing from communications has been of great interest in the last few years, as it allows data signals to be exploited to infer remarkable details about the physical world. Although the primary purpose of the communication signals is to carry logical information, concepts of radar analysis [14, 5, 23, 16, 15, 25, 27, 22, 19, 26, 10, 13, 17] are adapted to extract these details. There are however several challenges in the adaptation since communication signal is defined particularly for carrying data. For example, radar systems control their resolution by specially encoding their transmitting signals, say in the form of Frequency-Modulated Carrier Waves (FMCW) for spectrum sweeping, but when sensing from RF communication a similar sort of transmitter cooperation typically cannot be leveraged. As another example,

Existing Work	Feature Basis	Device	Sensing Task
WiSEE: Pu et al. [24]	Doppler Shift	USRP-N210	Gesture recognition
Wi-Vi: Adib and Katabi [1]	Phase	USRP-N210	Gesture based communication,tracking
E-eyes: Wang et al. [30]	RSSI, CSI	COTS 802.11n devices	Activity classification
Gonzalez-Ruiz et al. [12]	RSSI	IEEE 802.11g wireless card	Obstacle mapping
Wang et al. [29]	Phase, CSI	COTS 802.11ac devices	Activity classification
WiKey: Ali et al. [2]	CSI	COTS 802.11n devices	Key recognition
RSA: Zhu et al. [32]	RSS	HXI Gigalink 6451 60GHz radios	Object imaging

Table 1: Summary of recent SISO sensing systems

sophisticated radar signal processing techniques, say creating a synthetic aperture using a *large* number of antennas, cannot be implemented directly in communication systems due to resource limitations.

Many techniques have been developed and demonstrated to address the above mentioned challenges for diverse sensing tasks including motion tracking [1], activity/gesture recognition [24, 30, 29], and obstacle/object mapping/imaging [32, 12], and even minor motions like keystrokes recognition [8, 2] and lip reading [28]. One idea is to use one antenna to emulate an antenna array in the presence of human movement. By tracking the angle of the reflected signal from the target (human) [1], the system is able to track the motion of the target as a form of inverse synthetic aperture radar (ISAR). Ubicarse [18] exploits the idea of circular synthetic aperture radar (SAR), in which the system rotates a single antenna so as to emulate a circular antenna array. As SAR does not require the target to be in motion, unlike the case of ISAR, Ubicarse proposes a method of using a handheld device to create circular antenna array to perform localization. To overcome any imprecision in the circle created by the rotation, it refines the formulation of SAR by using the relative trajectory between two receive antennas. Some other techniques characterize signatures corresponding to the channel variation caused by human activities. E-eyes [30] shows that temporal RSS and CSI features, which are available in COTS devices, can be used in activity classification, albeit this requires relatively heavy training. WiSee [24] proposes a method to extract Doppler shifts from OFDM symbols by applying a large FFT over repeated symbols, and gesture recognition is then shown to be possible from the extracted Doppler shifts. Another interesting technique used by communication based sensors maps obstacles/objects [12, 20]. The Tx-Rx pairs detect the presence of obstacles via wireless measurements and thereby co-operatively draw the indoor obstacle map.

As our protection system is single-input-single-output (SISO), we focus on breaking any SISO illegitimate sensing system in this work. Although SISO sensing systems use diverse techniques exemplified in Table 1, they all leverage a subset of the 3 DoFs discussed in Section 1. Since PhyCloak provides a generic tool to obfuscate in all these three dimensions, it can protect against any SISO sensor.

In contrast, for a multi-antenna sensing system, there is an additional DoF—the relative placement of antennas—that yields other types of information like angle of arrival (AoA) and time difference of arrival (TDoA). Nevertheless, by rotating PhyCloak's transmit antenna or extending our framework to a multi-antenna protection system, we would have the freedom to also obfuscate the fourth dimension provided by a multi-antenna sensing system.

3 Overview

3.1 Threat Model

Assume there is an adversary who is interested in inferring physical information from a SISO wireless communication channel. The adversary may be active or passive, i.e., it can transmit itself or just exploit ongoing wireless transmissions. In both cases, we assume that the adversary uses a single-antenna receiver to sniff the wireless transmission. In general, the design and implementation of adversarial sensing is unknown to the protection system designer.

Note that some types of sensing require a training phase to tune recognition patterns with respect to the environment of interest. To protect against stronger adversaries, we assume that the adversary is well trained for the environment at hand. The details of this training, whether it occurs concurrently with the training of a legitimate sensor or is based on some historical knowledge, are outside the scope of our interest here.

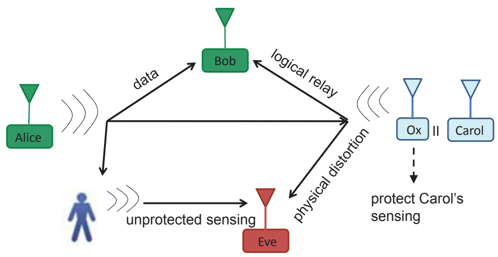

Figure 1: 4 single-input-single-output (SISO) nodes exist in the system: Alice, Bob, Carol and Eve: Alice and Bob perform data transmission and reception; Eve performs illegitimate sensing by exploiting Alice's transmission; Carol also performs sensing, but her obfuscator module forwards the received signal in a way that distorts physical information but preserves logical information

3.2 System and Goals

Our protection system comprises 4 SISO nodes as shown in Figure 1: Alice (data transmitter), Bob (data receiver), Carol (legitimate sensor) and Eve (illegitimate sensor). Both Alice and Bob can be controlled by Eve, thus Carol does not assume that Alice and Bob are honest.

Goals: 3 tasks co-exist in the network: data transmission between Alice and Bob, illegitimate sensing at Eve and legitimate sensing at Carol. By adding Ox to Carol with no cooperation from any of the other nodes, the protection system must satisfy the following three goals:

1. Obfuscate Eve's sensing.
2. Preserve Carol's sensing.
3. Not degrade the throughput of the link between Alice and Bob, nor introduce extra computation at Alice and Bob; i.e., Alice's and Bob's behaviors stay unaltered when Ox operates.

3.3 Three Degrees of Freedom

Usually a forwarder relays the signal directly, but in the context of an Ox a forwarder can do far more. In fact, a forwarder can be viewed as a special type of reflector; in theory, whatever change a natural reflector can induce on a signal, a forwarder can induce likewise. We begin by examining how a reflector changes the signal.

Letting the received signal at a reflector be $s(t)$, the received signal $r(t)$ that it reflects can be expressed as

$$r(t) = a \times s(t) \times e^{j2\pi(f_c+\Delta f)(t+\Delta t)} \tag{1}$$

where a is the amplitude gain due to reflection and propagation, f_c is the carrier frequency, Δf is the Doppler shift caused by a reflector that moves at a constant speed relative to the receiver, and Δt is the delay due to propagation over the path. We see that a reflector modifies signals by changing three components: a, Δf and Δt. Namely reflectors enjoy three DoFs when modifying signals.

We examine what kind of signal processing is needed at the Ox to effect similar changes in the signal being forwarded. Rewrite Equation 1 into the following form:

$$r(t) = a \times s(t) \times e^{j2\pi\Delta f t} \times e^{j2\pi(f_c+\Delta f)\Delta t} \times e^{j2\pi f_c t} \tag{2}$$

Amplitude gain a: It is clear that if a forwarder receives $s(t)$ from the source, then by amplifying the samples with different levels, a can be easily changed.

Doppler shift Δf: To emulate a Doppler shift of Δf, a forwarder can rotate the nth received sample by $2\pi n\Delta f\overline{\Delta t}$, where $\overline{\Delta t}$ = sampling interval.

Delay Δt: A delay of Δt can be introduced by simply delaying the to-be-forwarded signals in either the digital domain or the analog domain at the forwarder. A problem with delaying signals in the digital domain is that digital delays are discrete and do not match the speed of human movement. For example, if an ADC works with a sampling rate 100MHz, then the minimum delay that can be introduced in digital domain is 10ns, which corresponds to a distance of 3m. Controlling analog delay while feasible, however requires effort in modifying existing SDR platforms. Our solution then is to rotate the to-be-forwarded samples by a fixed phase $2\pi(f_c+\Delta f)\Delta t$ in the digital domain, which matches the expected delay of Δt. In our NI PXIe platform, this calculation can be made in two clock cycles ($\frac{1}{\text{ADC sampling rate}}$).

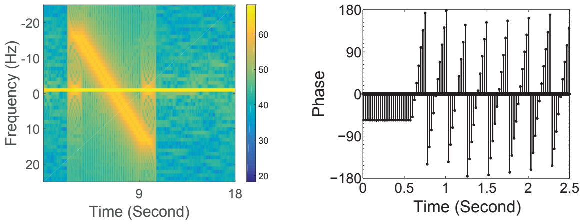

(a) By multiplying the nth to-be-forwarded sample with $2\pi n\Delta f\overline{\Delta t}$, and changing Δf from 20Hz to -20Hz, the Doppler shift profile at the receiver is as expected

(b) By rotating the to-be-forwarded signals with a certain phase which changes by 36° every 30ms at the forwarder, the phase of the signal changes $\sim 36°$ every 30ms

Figure 2: Expected Doppler shift and phases are generated at a forwarder

Figure 2(a) depicts the Doppler shift profile of the received signals that are sent by a forwarder who keeps changing the to-be-forwarded samples' Doppler shift from 20Hz to -20Hz according to the above algorithm. Similarly, from Figure 2(b) we can see that by multiplying the to-be-forwarded samples with a phase φ which increases 0.2π every 30ms at the forwarder, the phase of the received samples changes by $\sim 0.2\pi$ every 30ms. These results show that a forwarder can predictably control Doppler shift and phase.

4 Design

Figure 3 shows a simplified block diagram of our system PhyCloak. The physical distortion is introduced after self-interference cancellation and then the distorted signal is then forwarded to the transmit antenna.

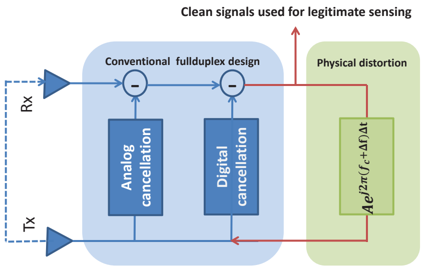
Clean signals used for legitimate sensing
Conventional fullduplex design
Physical distortion
Rx
Tx
-
-
Analog cancellation
Digital cancellation
$Ae^{j2\pi(f_c+\Delta f)\Delta t}$

Figure 3: High-level block diagram of PhyCloak

4.1 Online Maintenance of Self-Channel Estimates

As mentioned earlier, PhyCloak is a full-duplex system that needs to cancel self-interference to operate. However, human movement close to the full-duplex radio changes the self channel and affects cancellation. Figure 4 illustrates this phenomenon as it depicts the power of the residual noise after cancellation over time when a human target walks around the fulld-uplex radio. The full-duplex radio re-estimates the channel every 1s. We see that if we set the residual threshold to -95dBm, which is 5 dB above the maximum digital cancellation capability (noise = -100dBm), the channel estimation works fine only for a short duration (∼100ms) after each channel estimation update. This observation implies that frequent self-channel re-tuning (∼100ms) is required.

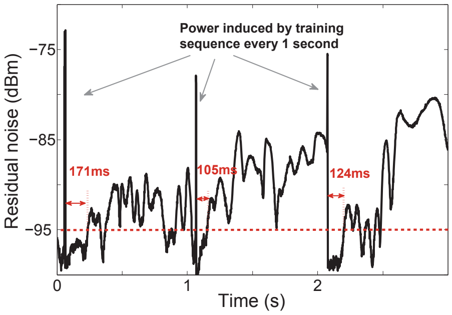

Figure 4: With human movement going on, the self-interference cancellatoin works fine only for a short duration (∼100ms)

A complication, however, arises when an update is attempted during an ongoing external transmission: the external transmission may distort self-channel estimation while the transmission that helps with self-channel estimation may interfere with external data reception. There are two straightforward solutions to this problem: 1) using MAC; 2) exploiting the silent period defined by wireless protocols, like short inter-frame space (SIFS) in WiFi. The former hurts the throughput of data transmission and moreover interrupted external transmission degrades coupling legitimate sensors with the Ox. And in addition, both of the solutions require a big effort to design careful adaptation to various wireless communication protocols.

We therefore propose a self-channel estimation algorithm for PhyCloak that addresses this complication. It uses two main elements: 1) oversampling and differential to get rid of any ongoing external transmission, and 2) a special training sequence that yields minimum interference to external transmissions.

4.1.1 Self-channel estimation with and without external interference

Before we describe our self-channel estimation algorithm, let us first see the impact of training with and without external interference. Assume $A = \{a_{-m}, a_{-m+1}, \ldots, a_m\}$ is the transmitted training sequence, $B = \{b_0, b_1, \ldots, b_m\}$ is the received sample sequence, and $H = \{h_0, h_1, \ldots, h_m\}$ is the channel coefficient vector in time domain with $m+1$ taps. Therefore, we have

$$\begin{Bmatrix} b_0 \\ b_1 \\ \cdots \\ b_m \end{Bmatrix} = \begin{Bmatrix} a_0 & \cdots & a_{-m} \\ a_1 & \cdots & a_{-m+1} \\ \cdots & \cdots & \cdots \\ a_m & \cdots & a_0 \end{Bmatrix} \times \begin{Bmatrix} h_0 \\ h_1 \\ \cdots \\ h_m \end{Bmatrix} \tag{3}$$

In the presence of external transmission, B becomes:

$$\begin{Bmatrix} b_0 \\ b_1 \\ \cdots \\ b_m \end{Bmatrix} = \begin{Bmatrix} a_0 & \cdots & a_{-m} \\ a_1 & \cdots & a_{-m+1} \\ \cdots & \cdots & \cdots \\ a_m & \cdots & a_0 \end{Bmatrix} \times \begin{Bmatrix} h_0 \\ h_1 \\ \cdots \\ h_m \end{Bmatrix} + \begin{Bmatrix} s_0 & \cdots & s_{-m} \\ s_1 & \cdots & s_{-m+1} \\ \cdots & \cdots & \cdots \\ s_m & \cdots & s_0 \end{Bmatrix} \times \begin{Bmatrix} h'_0 \\ h'_1 \\ \cdots \\ h'_m \end{Bmatrix} \tag{4}$$

where $S = \{s_{-m}, s_{-m+1}, \ldots, s_i, \ldots, s_m\}$ is the external transmitted sample sequence, and $H' = \{h'_0, h'_1, \ldots, h'_m\}$ is the channel coefficient vector which corresponds to the channel between the transmit antenna of the external device and the receive antenna of the Ox.

4.1.2 Oversampling and differential to get rid of external interference

To overcome the external interference in Equation 4, which is unknown to PhyCloak, we exploit oversampling. Say PhyCloak samples at a rate $2m$ times higher than the sampling rate of the external transmitter, it follows that approximately $s_{-m} = \ldots = s_m$. So

$$\begin{Bmatrix} s_0 & \cdots & s_{-m} \\ s_1 & \cdots & s_{-m+1} \\ \cdots & \cdots & \cdots \\ s_m & \cdots & s_0 \end{Bmatrix} \times \begin{Bmatrix} h'_0 \\ h'_1 \\ \cdots \\ h'_m \end{Bmatrix} = \begin{Bmatrix} s_0 \times (h'_0 + \ldots + h'_m) \\ s_0 \times (h'_0 + \ldots + h'_m) \\ \cdots \\ s_0 \times (h'_0 + \ldots + h'_m) \end{Bmatrix} \tag{5}$$

Therefore, by differential we have

$$\begin{Bmatrix} b_1 - b_0 \\ b_2 - b_1 \\ \cdots \\ b_m - b_{m-1} \end{Bmatrix} = \begin{Bmatrix} a_1 - a_0 & \cdots & a_{-m+1} - a_{-m} \\ a_2 - a_1 & \cdots & a_{-m+2} - a_{-m+1} \\ \cdots & \cdots & \cdots \\ a_m - a_{m-1} & \cdots & a_1 - a_0 \end{Bmatrix} \times \begin{Bmatrix} h_0 \\ h_1 \\ \cdots \\ h_m \end{Bmatrix} \tag{6}$$

It may appear that we have already been able to get rid of external interference, however, $\mathbf{A}$ is an $m \times (m+1)$ matrix, so the rank of A is less than $m+1$. This means that we can get only a unique solution for at most m

of the $m+1$ unknowns contained in $\mathbf{H}$, where $\mathbf{H} = \{h_0, h_1, \ldots, h_m\}^T$ and

$$\mathbf{A} = \left\{ \begin{array}{ccc} a_1 - a_0 & \ldots & a_{-m+1} - a_{-m} \\ a_2 - a_1 & \ldots & a_{-m+2} - a_{-m+1} \\ \ldots & \ldots & \ldots \\ a_m - a_{m-1} & \ldots & a_1 - a_0 \end{array} \right\} \quad (7)$$

4.1.3 A special training sequence

To ensure that Equation 6 has a unique solution for $\{h_0, h_1, \ldots, h_{m-1}\}^T$, we leverage a special training sequence, namely a square wave, which is shown in Figure 5(a). As shown in Figure 5(b), the fundamental frequency of the square wave is the square wave frequency, and its odd harmonics are decreasing in size. To be more specific, for a square wave over a period consisting of N samples with B MHz sample rate, the frequency components are at $1f$, $3f$,..., $(2i+1)f$, ... with decreasing amplitude, where $f = \frac{B}{N}$MHz.

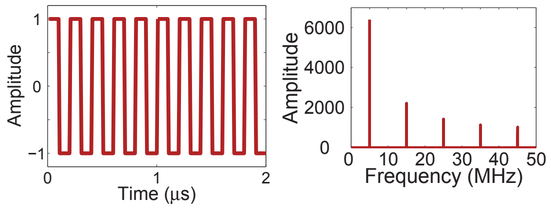

(a) Training sequence in time domain

(b) Training sequence in frequency domain

Figure 5: Training sequence

The rationale for using this training sequence is twofold: First, the square wave has a unique solution to $\{h_0, h_1, \ldots, h_{m-1}\}^T$ as long as $a_{-m} = a_{-m+1} = \ldots = a_0 = a_1 + c = \ldots = a_m + c$, where c is a non-zero constant. And second, the spikes it produces in the frequency domain are sparse. For example, with $B = 100$MHz and $N = 16$, the space between neighboring spikes is 12.5MHz. Such sparse spikes are tolerable in wireless systems. For example, in a 20MHz WiFi band using OFDM, as claimed by Flashback [9], existing WiFi systems have a relatively large SNR margin. And because the interference of any such spike is constrained to at most one subcarrier, the loss of a few bits does not significantly affect decoding, as successful packet transmissions always respect SNR margins.

4.1.4 The training procedure

Training is performed as follows: PhyCloak samples at a rate n times higher than that of external transmission. A training sequence which is the concatenation of consecutive 1s and -1s is sent during training. The received samples corresponding to the transition points (1 to -1 or vice versa) are used to calculate the channel coefficients. More specifically, the received sample b_0 which corresponds to the point right before the transition occurs is equal to $h_0 + \cdots + h_m$, and the next received sample b_1 is equal to $-h_0 + \cdots + h_m$. Thus, we can compute $h_0 = (b_0 - b1)/2$. The rest of the channel coefficients are calculated in a similar way. One concern is whether the desired oversampling rate can be supported. Take 802.11g as an instance, which has the smallest bandwidth (20MHz) among WiFi standards. If training were to require a 20X oversampling rate, we would need a platform that supports 400MHz sampling rate, which is very expensive. We figure out that, however, a 4X oversampling rate is sufficient to eliminate the effect of an external transmission of 802.11g. The reason is that the delay spread of non-ultra-wideband transmission in an indoor setting does not expand more than 3 taps.

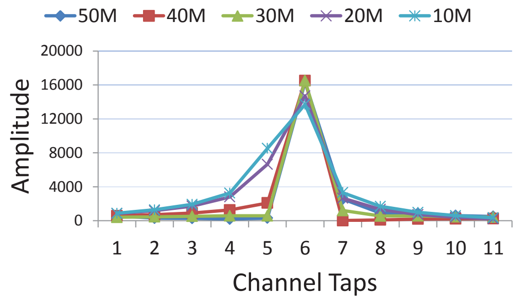

Figure 6: Channel coefficients measured at different sampling rates

To understand that, we need to know the fact that power delay profile is decided by two factors: multipath propagation and inter-symbol-interference (ISI). Let us study them one by one. First is the multipath propagation. For a 20MHz radio, one tap corresponds to $\frac{3 \times 10^8 \text{m/s}}{20\text{MHz}} = 15$m. So the fourth tap corresponds to a 60-meter reflective path. The power conveyed by the 60-meter reflective path is significantly smaller than that conveyed by the short ($\sim$10cm) line-of-sight path between the co-located transmitting and receiving antennas. Second, due to ISI each received sample is affected by not only the intended transmitted symbol, but also its two neighboring symbols. Therefore the delay spread expands across 3 taps. Figure 6 plots the channel estimation of the self channel under different sampling rates in the same environment. We see that in all cases, the main energy is always spread across 3 taps. So as long as we can accurately estimate the three dominant taps in non-ultra-wideband, we can achieve good cancellation performance. That implies we need the external interference to be stable during the reception of at least four consecutive samples at the transition point of the training sequence so as to get the three main taps by differential. Namely 4X oversampling is required.

Note that 4X oversampling does not guarantee the re-

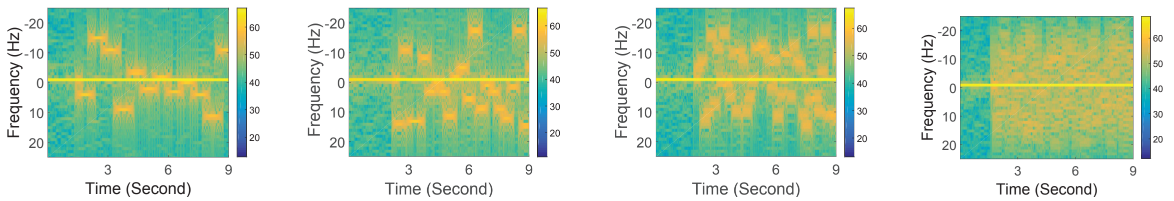

(a) 1s (b) 0.5s (c) 0.125s (d) 0.05s

Figure 7: The granularity of the spectral decreases as the Doppler shifts change from 1s to 0.05s

ception of the desired 4 samples happen in the duration of one external interference sample. But we can leverage the interference reduction provided by averaging over multiple transition points, and partially accurate estimation of the channel taps, and still achieve good performance. Even lower oversampling rate (2X/3X) also performs well according to the experiment (see Section 5.2).

4.2 Obfuscation of Patterns in 3 DoFs

To motivate how we obfuscate patterns in the three DoFs, let us first examine the result of superposing a signal via one path with an obfuscated version via another path.

Assume we have two paths: one with $\{a_1, \Delta f_1, \Delta t_1\}$, and the other via the Ox with $\{a_2, \Delta f_2, \Delta t_2\}$. The superposition of the signals through these two paths is given by the following formula:

$$\begin{aligned}\hat{r}(t) =& a_1 \times s(t) \times e^{j2\pi(f_c+\Delta f_1)(t+\Delta t_1)} \\ &+ a_2 \times s(t) \times e^{j2\pi(f_c+\Delta f_2)(t+\Delta t_2)}\end{aligned} \tag{8}$$

Now, is superposing an obfuscated signal sufficient for hiding the original triplet $\{a_1, \Delta f_1, \Delta t_1\}$? The answer is partially yes: The amplitudes and delays are instantaneously covered in the superposed signal, but the respective Doppler shifts remain distinguishable after superposition. So, a and Δt can be hidden instantly by randomly changing amplitude and delay of the signal by the Ox.[1] To see why Doppler shifts are distinct even after superposition, consider the frequency response of the received signals:

$$\begin{aligned}R(f) =& \int \hat{r}(t) e^{-2\pi j f t} dt \\ =& \int (a_1 \times s(t) \times e^{j2\pi(f_c+\Delta f_1)(t+\Delta t_1)}) e^{-j2\pi f t} dt \\ &+ \int (a_2 \times s(t) \times e^{j2\pi(f_c+\Delta f_2)(t+\Delta t_2)}) e^{-j2\pi f t} dt \\ =& a_1 e^{j2\pi(f_c+\Delta f_1)\Delta t_1} S(f - fc - \Delta f_1) \\ &+ a_2 e^{j2\pi(f_c+\Delta f_2)\Delta t_2} S(f - fc - \Delta f_2)\end{aligned} \tag{9}$$

where $S(f)$ is the frequency response of $s(t)$. In an OFDM system, we can see two frequency components that are shifted by Δf_1 and Δf_2 around the subcarrier f.

4.2.1 Doppler shift obfuscation

As amplitude and delay can be instantly changed by superposition with an obfuscated signal, patterns that rely only on amplitude and delay can be hidden by Ox, by randomly changing them on a per packet basis. At first glance, it may appear that this scheme cannot be made to work for patterns that rely on Doppler shift, but it turns out the scheme can be made to work for Doppler shift, assuming the moments of change are carefully chosen.

The rationale for choosing the moments of change is based on the fact that a t-second observation in the time domain leads to $1/t$ Hz granularity in the frequency domain. To choose the appropriate Δf at $1/t$ Hz granularity, there is an implicit requirement that the Δf needs to last for at least t seconds. Therefore, if the forwarder changes its Δf every t seconds while the other copy's Δf does not change, an observer would still only see $1/t$ Hz granularity. Since human movements typically result in -20Hz to 20Hz Doppler shifts in the 2.4GHz band, a Doppler shift of the forwarded copy that changes every 0.1s creates sufficient confusion at an observer. Figure 7 shows that when the Doppler shifts of the transmitted signals are varied from every 1s to every 0.05s, the spectral seen by an observer with 1s observation interval have progressively finer granularity, to the point where a time-frequency pattern gets hidden.

4.2.2 Effect of superposing with randomly changing obfuscated signals

The basic idea of PhyCloak then is to superpose signals from the target with naturally changing $\{a, \Delta f, \Delta\phi\}$ with the obfuscated signals with randomly changing $\{a, \Delta f, \Delta\phi\}$. More specifically, as analyzed above, PhyCloak changes the value of the triple every 0.1s. We illustrate the blackbox effect of obfuscation experimentally using two state-of-the-art sensors, WiSee [24] and Wi-Vi [1], which we implemented. WiSee performs gesture recognition by extracting Doppler shifts from OFDM symbols, whereas Wi-Vi uses ISAR to track the angle of human motion with respect to the receive antenna of the sensor.

For the case of obfuscating Doppler shift patterns, Figure 8 shows the superposition of a signal with the syn-

[1] In theory for a high sampling rate receiver, delays might be separable in the brief prefix that arrives before the obsfuscated signal arrives, but how much information a sensor can accurately extract from the brief clean prefix is questionable.

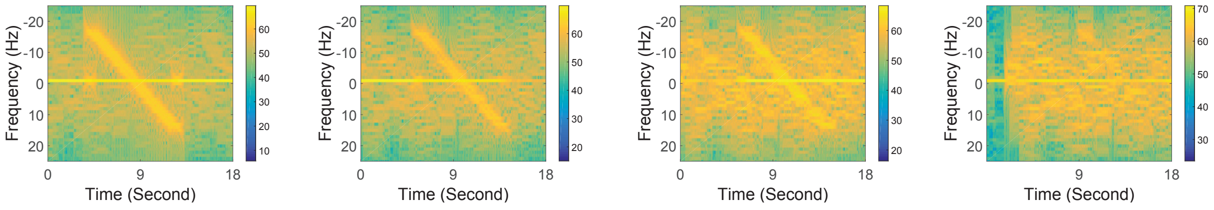

(a) 0dB (b) -3dB (c) -6dB (d) -9dB

Figure 8: The pattern that a WiSee sensor sees in Figures 2(a) is hidden by an obfuscated signal where Doppler shift changes every 0.1 second

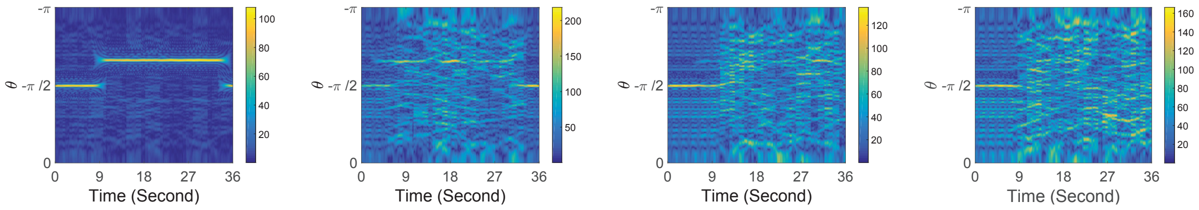

(a) Motion towards a Wi-Vi style sensor with constant angle (b) 0dB (c) -3dB (d) -6dB

Figure 9: The constant angle of human motion (starting from 9th second) that a Wi-Vi style sensor sees in (a) is hidden by an obfuscated signal where phase changes randomly every 0.1 second

thetically generated Doppler shift pattern described in Figure 2(a) and an obfuscated copy of the pattern where Doppler shift changes randomly every 0.1s. We see that pattern of Figure 2(a) is covered by the "noise map" created by the randomly changing copy. As the strength ratio of the former relative to the latter, which we define as signal to obfuscation ratio (SOR), decreases from 0dB to -9dB, the visibility of the artificial pattern decreases.

For the case of obfuscating phase-based patterns, we synthetically emulated a human moving towards the receive antenna of our Wi-Vi style sensor at a constant angle, as shown in Figure 9(a), and then superposed the signal with a randomly obfuscated copy where phase changes every 0.1s. Figure 9 shows that as SOR decreases from 0dB to -6dB, the pattern shown in Figure 9(a) becomes progressively invisible at the Wi-Vi style sensor.

It is worth noting that power passively reflected by human is much smaller compared to that actively forwarded by an Ox that has its own power supply. Therefore 0dB SOR can be readily achieved. To illustrate this point, we can build a simplified power model of our system. In our system Ox's goal is to minimize SOR at Eve with no knowledge of the locations of any of the other parties, so its best strategy is to work at the maximum transmission power. If we assume free-space attenuation, then SOR $\sim \frac{a}{A}(\frac{d_3 d_4}{d_1 d_2})^2$, where a and A are the reflection gains at target and Ox respectively, and d_1,d_2,d_3 and d_4 are the distances as shown in Figure 10(a). Figure 10(b) plots the simulation result of the CDF of SOR when we randomly place Alice, Eve, Ox and human target in a 10m×5m room, with reflection gains being set to -3dB and 10dB respectively. We see that in around 88% cases, SOR is smaller than 0dB.

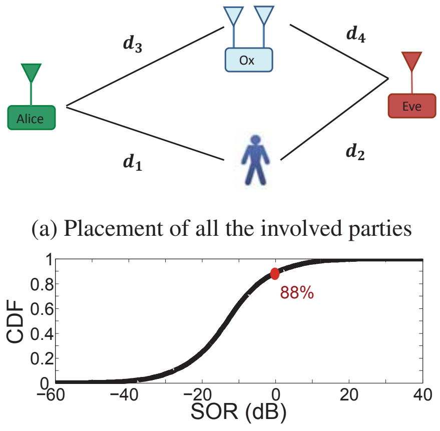

(a) Placement of all the involved parties

(b) SOR distribution

Figure 10: A simplified power model

4.2.3 Security analysis

We believe that our system is robust against a single antenna eavesdropper given certain SOR because of the fact: little information can be extracted from two random signals which occupy similar bands as long as the power of the undesired signal is higher than that of the desired one. In our case, the desired signal is the natural channel variation induced by target, while the undesired one is the artificial channel variation induced by PhyCloak. It is worth noting that as human motion is slow, natural channel variation has a small bandwidth, which is comparable to that of the artificial channel variation that changes ev-

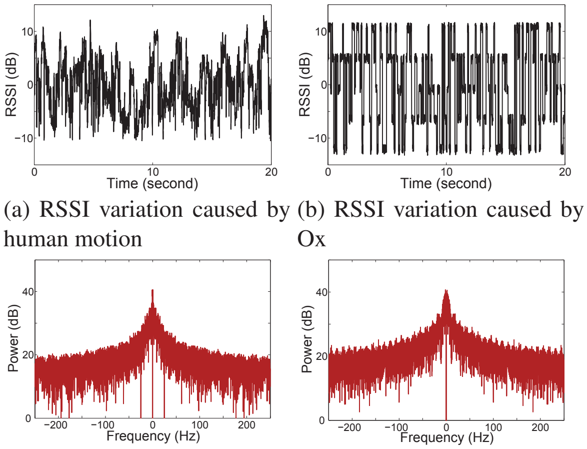

(a) RSSI variation caused by human motion (b) RSSI variation caused by Ox

(c) Power spectrum of the RSSI trace in (a) (d) Power spectrum of the RSSI trace in (b)

Figure 11: The signal (real channel state information) and the noise (artificial channel state information) have similar bandwidths

ery 0.1s.

To illustrate the above point, we compare the RSSI variations induced by human and PhyCloak. Figure 11(a) and 11(b) plot the RSSI changes caused by human movement and PhyCloak respectively, and Figure 11(c) and 11(d) plot the corresponding power spectrums. From the figure, we see that the occupied bandwidths of the two channel state traces are similar.

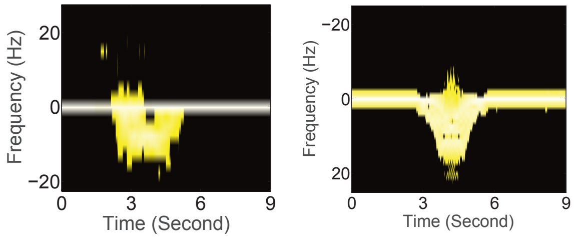

(a) Doppler shifts created by a human gesture (pull) (b) Doppler shifts emulated by an Ox

Figure 12: Spoofing

4.3 Spoofing

According to the above discussion, our design succeeds in obfuscating any RF-based single-antenna sensors by creating false negative results. But an Ox can achieve more than that: it can create false positives also by spoofing changes in the 3 DoFs that are similar to the changes created by a target. By splitting the to-be-forwarded samples into multiple streams, applying different instantiations of the triple $\{a, \Delta f, \Delta t\}$ to them, and forwarding the combination of the processed streams as one stream, an Ox can emulate multiple reflectors corresponding to different parts of the target (say a human body). But unlike the case of false negatives, the effectiveness of creating false positives at a sensor grows as the Ox knows more about the features and algorithms used by the sensor. For example, if an Ox knows a sensor uses the WiSee algorithm [24], it can create a Doppler shift profile accordingly without making an effort to model accurate human movement. Figure 12 depicts the extracted Doppler profile of a human gesture (pull) and that spoofed by an Ox. WiSee segments a Doppler profile into positive and negative parts according to its power distribution and encodes them into 1s and -1s respectively. Since both of the profiles contain positive Doppler shifts of negligible power, they will be encoded as -1s and mapped to the same target by a WiSee sensor.

4.4 PhyCloak

By obfuscating using random physical distortion, an Ox is able to confuse Eve, and by online maintenance of self-channel estimates, Ox is able to output interference-free signals to Carol for legitimate sensing. However, one critical requirement is still not met: preserving the communication throughput in the presence of Ox.

Although PhyCloak works as a relay at logical layer which can potentially improve the throughput [31], it is not clear that obfuscation would not hurt the decoding process. We find that, however, as long as the change of the triplet $\{a, \Delta f, \Delta \phi\}$ does not happen in the middle of packet transmission, obfuscation is safe with respect to data communication. The reason for this is that from the perspective of a data receiver, the Ox effectively just adds variability to the channel. Since data receivers usually perform channel estimation at the beginning of the received packet, as long as the channel is stable during the reception of the packet, decoding can be successful. We, therefore, refine the design of PhyCloak as follows: PhyCloak switches between two transmitting modes: training and forwarding. In the training phase, the PhyCloak sends the above mentioned training sequence and computes its self-channel estimate according to Section 4.1.4; in the forwarding phase, PhyCloak then performs self-interference cancellation, applies the physical distortion $\{a, \Delta f, \Delta \phi\}$ to the interference-free signal and forwards the distorted signal via the transmit antenna. The PhyCloak randomly chooses an instance of $\{a, \Delta f, \Delta \phi\}$ in the predefined pool and updates the current value when the channel is free and the last update happened more than 0.1s ago. In this way, PhyCloak avoids interfering with the transmission. And in theory, there is still a chance that due to the delay caused by free-channel detection, PhyCloak changes the channel after several samples of a packet has been transmitted, but that chance is quite low. Even if it happens, because PhyCloak only affects a few samples at the beginning, the packet might still be decodable.

5 Validation

We now describe a prototype of PhyCloak that we have built, and our experiments to validate its performance.

5.1 Experimental Setup

Our prototype is based on PXIe 1082 SDR platform. We built the transmitter, receiver, eavesdropping sensor and legitimate sensor on the same platform, which all follow the 802.11g standard, i.e., working at 2.4GHz with a 20MHz band. PhyCloak works at the same center frequency but with a 50MHz sampling rate, about 3 times the rate of an external data transmission, which gives it a reasonable margin to perform self-channel estimation with an ongoing external transmission (see Section 5.2).

PhyCloak contains two RF chains, one for transmitting and one for receiving. Each of the RF chains contains an NI-5791 (FlexRIO RF transceiver equipped with one antenna) for transmitting or receiving and an NI PXIe-7965R (a Xilinx Virtex-5 FPGA) for digital processing. Analog cancellation is implemented according to our earlier design [7, 4]. The self-channel estimation, digital cancellation and physical layer distortion are implemented on the FPGA. The distortion processing introduces a latency of about 100ns. Our experiments were conducted in a 5m×7m lab.

5.2 Self-Interference Cancellation

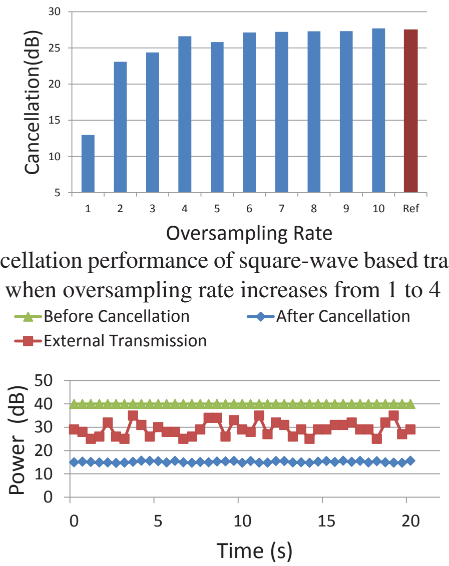

(a) Cancellation performance of square-wave based training increases when oversampling rate increases from 1 to 4

(b) Insensitivity of square-wave based training to external transmission power variation, which is necessary for preserving legitimate amplitude based sensing

Figure 13: Self-interference cancellation performance

We begin with the performance of the digital cancellation of our self-channel estimation algorithm. As discussed in Section 4.1.4, Ox tolerates external interference during self-channel estimation using oversampling. So, we first examine the oversampling rate needed to achieve reasonably accurate self-channel estimates in the presence of external transmission. We let a full-duplex transceiver operate at 50MHz with a 10-tap filter for self-interference (digital) cancellation. Self-channel estimation is obtained by averaging over 128 training rounds, which altogether takes about 20μs.

Figure 13(a) plots the self-interference cancellation performance of our square-wave based training. In the figure, as we fixed the sampling rate of the full-duplex radio (50MHz), different oversampling rates correspond to different external transmission rates with the received power of the external transmissions being the same as that of self-interference signal at Ox's receive antenna[2]. 1X oversampling rate corresponds to the case when the training and data communication use the same sampling rate, in which case square-wave based training and traditional pilot based training would achieve similar performance. We see that the performance of self-interference cancellation of square-wave based training gets better as the oversampling rate increases from 1 to 4, but it stops increasing after 4, and achieves similar performance as that in the case when there is no external transmission going on (indicated by the red bar). It shows that Ox can reliably estimate and cancel self-interference even in the presence of strong external transmission when the oversampling parameter is 4X as supported by our observation in Section 4.1. In addition, 2X and 3X oversampling rates also produce high cancellation as they benefit from two factors: 1) accurate estimation of part of the channel taps, and 2) averaging over multiple transition points. **Takeaway:** *Our oversampling technique makes self-interference cancellation reliable at modest oversampling rates even in the presence of strong ongoing external transmission.*

The analysis above considers external interference sent at a fixed power. To enable legitimate sensing, self-interference cancellation performance needs to be stable even when the received power from external transmission is varying. For example, an unstable self-interference canceler can render an amplitude-based sensor useless since the (varying) residual self-interference will affect the received signal amplitude. Figure 13(b) plots the full-duplex radio's cancellation performance with 3X oversampling rate over time during which the received power from the external transmitter fluctuates. We see that the self-interference cancellation performance of square-wave based training is insensitive to the variation of external interference. **Takeaway:** *Our oversampling technique results in a stable cancellation performance at*

[2] Note that this is a very strong external interference and we choose this setting to show oversampling strategy's performance even under strong external interference.

modest oversampling rates even when the received signal from external transmitter is varying.

5.3 Obfuscation Performance

5.3.1 Obfuscation vs. SOR in 3 DoFs

We first measure the different levels of obfuscation created by PhyCloak by comparing the correlation of the amplitude, phase and Doppler shift with and without the presence of PhyCloak. The transmitter is programmed to send continuous OFDM symbols with QPSK modulation with varying amplitude and phase. An artificial Doppler shift of 10Hz is also added at the transmitter. PhyCloak performs obfuscation by randomly changing the amplitude, phase and Doppler shifts every 0.1s.

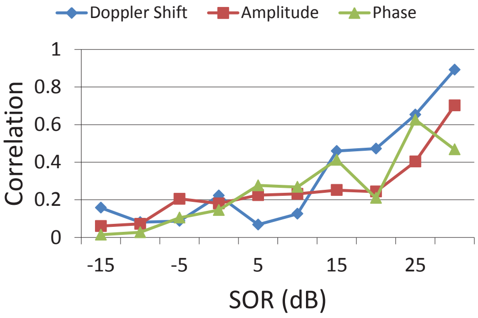

Figure 14: Obfuscation level of each of the three features decreases as SOR (original signal over obfuscation signal) increases

Figure 14 depicts the correlation between the pairs of amplitude, phase, and Doppler shifts at different SORs: Again, SOR is the signal strength ratio of original signal over obfuscation signal (see Section 4.2.2). We see that as SOR increases, the correlation of each pair of the three features increases, i.e., the obfuscation degree decreases. Amplitude sequence pair and phase sequence pair see lower correlation than Doppler shift pair when SOR is high. This is because amplitude and phase are instantaneous quantities, while Doppler is a statistical quantity that is derived from multiple instantaneous samples. But, even for Doppler shift, a 10dB SOR is low enough to hide the patterns contained in signals reflected by targets. It's worth noting that in practice, as PhyCloak is independently powered while the target only passively reflects signals, the desired SOR to successfully obfuscate is readily achieved. **Takeaway:** *PhyCloak effectively obfuscates sensing even at a relatively high SOR.*

As different sensors differ in their robustness to noise, PhyCloak's effectiveness is sensor dependent. While we are unaware of any research on the robustness of the communication-based sensors, we may infer from Figures 8, 9 and 14 that less obfuscation power is needed to confuse a phase or amplitude based sensor as compared to a Doppler shift based sensor. Therefore, we choose to validate the PhyCloak's capability of confusing illegitimate sensing and preserving legitimate sensing in the context of WiSee, which is the state-of-the-art Doppler shift based sensor.

5.3.2 Degradation of illegitimate sensing

We built a Doppler-based sensor in our platform per the method proposed by WiSee [24]. The method consists of two parts: 1) extraction of Doppler shifts from repeated OFDM symbols by applying a large size FFT; and 2) using sequence matching to classify gestures. We note since we could not get to the original WiSee code and some of the details are missing, we implement WiSee with a few adaptations. For example, we randomly map the sequence to the predefined classes with uniform distribution in case the sequence does not match any of the predefined sequence. Our implementation shows a classification accuracy of 93% across 5 gestures in none-line-of-sight (NLoS) setting with the human target 5 feet away from the WiSee sensor, while WiSee reports 94% across 9 gestures. While there is this small discrepancy in replication, the core algorithm is the same and our main goal is to study obfuscation performance.

We examine the performance of an illegitimate WiSee sensor with obfuscation from a PhyCloak. We conduct two sets of experiments to validate PhyCloak's coverage range and its overall effectiveness under different channel conditions respectively.

Obfuscation coverage: First, we randomly choose 10 pairs of locations to place Tx and Eve, and then place Ox in locations such that the distance d_{TE} between Tx and Eve is equal to the distance d_{TO} between Tx and Ox as shown in Figure 15(a), but the distance d_{EO} between Eve and Ox varies from $0.5d_{TE}$ to $2d_{TE}$. The channels between any two of the three parties are line-of-sight (LoS).[3] A human target performs five gestures drag, push, pull, circle and dodge close to Eve. With no obfuscation, Eve's classification accuracy in this placement is about 90% across the five gestures.

For simplicity, we normalize d_{EO} by d_{TO} (d_{TE}), and plot the classification accuracy against the normalized d_{EO} in Figure 15(b). As we know, the received obfuscation power at Eve from Ox is a function of d_{TO} and d_{OE}, therefore as d_{EO} increases the power ratio of obfuscation over human reflection decreases. From the figure we see that classification accuracy of Eve increases as d_{EO} increases as expected. Note that since we have 5 classes,

[3] WiSEE sensors have a slightly worse performance in LoS (≈ 90%) than NLoS (≈ 93%) as strong direct power from the transmitter hides the information provided by target's reflection. For the next two experiments, we choose LoS instead of NLoS because it makes the placement easier to make sure d_{EO} is the only variable which would change the power ratio of the obfuscation and human reflection.

a classification accuracy of 0.2 means a random guess. PhyCloak can obfuscate Eve near perfectly when d_{EO} is smaller than 0.8, and it totally fails when it is larger than 1.7. **Takeaway:** *The closer Ox is to Eve, the better the achieved obfuscation.*

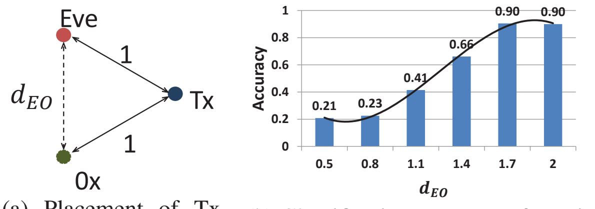

(a) Placement of Tx, Ox and Eve with all three channels LoS

(b) Classification accuracy of Eve in the presence of PhyCloak increases as d_{EO} increases

Figure 15: Eve's classification accuracy vs d_{EO}

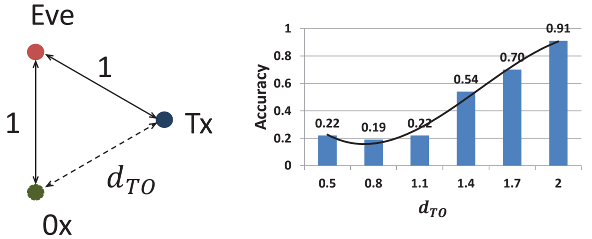

(a) Placement of Tx, Ox and Eve with all three channels LoS.

(b) Classification accuracy of Eve in the presence of PhyCloak increases as d_{TO} increases

Figure 16: Eve's classification accuracy vs d_{TO}

In the second experiment, we make $d_{TE} = d_{OE}$, and vary d_{TO} as shown in Figure 16(a). And again in Figure 16(b), we see that as d_{TO} increases, Eve's classification accuracy increases. **Takeaway:** *the closer Ox is to Tx, the better obfuscation is achieved.*

In other experiments we vary either the human-Eve or human-Ox distance while keeping the power received by Ox and human from Tx stay constant. As these distances respectively reduced, the effectiveness of the sensing and obfuscation respectively increased.

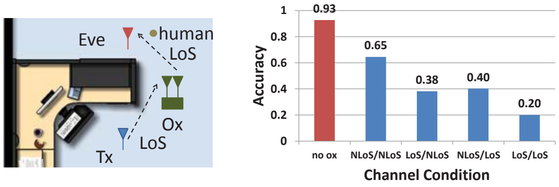

(a) Example of LoS/LoS placement in the lab

(b) Obfuscation degrades somewhat if Tx-Ox or Ox-Eve channels are NLoS

Figure 17: Eve's classification accuracy under different Tx-Ox and Ox-Eve channel conditions

Obfuscation effectiveness under different channel conditions: In addition to the coverage range in LoS setting, we also measure Eve's classification accuracy when channels between the transmitter and obfuscator and the channel between the obfuscator and Eve are under different LoS and NLoS combinations. Intuitively, when both channels are NLoS, Eve receives the least power forwarded by the obfuscator, and therefore, she achieves the best performance. We care about these channel conditions because in some scenarios the transmitter is under control of the adversary, and therefore the adversary may enjoy the freedom to create "good" channels to mitigate PhyCloak's obfuscation.

In the experiment, we make the channel between Tx and Eve NLoS, and the channel between Tx and the human and that between human and Eve LOS, so as to make sure Eve sees high classification accuracy when no obfuscation is going on. The channel between Tx and Ox and the channel between Ox and Eve have four possible channel condition combinations. A human target performs 500 times of the 5 predefined gestures near Eve in each of the four combinations. Figure 17(a) is an example of how we create a channel combination of Los/Los in the lab, where the first LoS refers to the channel condition of the channel between Tx and Ox, while the second refers to that of the channel between Ox and Eve. NLoS channels are created by placing obstacles in the direct propagation paths.

Figure 17(b) depicts Eve's classification accuracy without obfuscator and with obfuscator in four channel combinations. We can see that as expected, Eve sees the highest classification accuracy (65%) in NLoS/NLoS setting among the four channel conditions, but it is still smaller than the case when no obfuscation is happening (93%). Eve sees similar performance in Los/NLoS and NLoS/LoS scenarios as power forwarded by obfuscator in both the settings is similar. **Takeaway:** *although NLoS channel degrades the received power at Eve from Ox, the degradation is not dramatic since there is rich multipath propagation in indoor environment.*

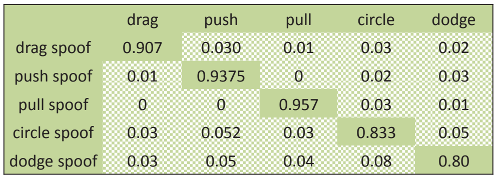

	drag	push	pull	circle	dodge
drag spoof	0.907	0.030	0.01	0.03	0.02
push spoof	0.01	0.9375	0	0.02	0.03
pull spoof	0	0	0.957	0.03	0.01
circle spoof	0.03	0.052	0.03	0.833	0.05
dodge spoof	0.03	0.05	0.04	0.08	0.80

Figure 18: False positives with a spoofing Ox

5.3.3 Feasibility of spoofing

We built a spoofing obfuscator by reverse engineering the five predefined sequences corresponding to the five gesture types that our WiSee sensor recognizes. The basic difference between this spoofing obfuscator and PhyCloak is that the former changes Doppler shift according to the five well-defined gestures, while the latter changes

Doppler shift randomly. The result is shown in Figure 18. **Takeaway:** *the spoofing obfuscator fools a WiSee sensor with a high success rate, averaging 88.69% across the 5 gestures, in the absence of human gesturing*.

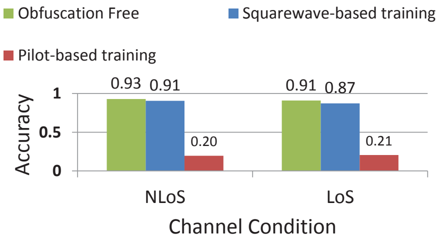

Figure 19: Square wave based training preserves legitimate sensing

5.3.4 Preservation of legitimate sensing

Next, we examine PhyCloak's capability of supporting coupled legitimate sensing. That is, we evaluate whether our self-channel estimation method produces consistent and sufficient self-interference cancellation in a changing environment to preserve legitimate sensing. Figure 19 depicts the legitimate sensor's classification accuracy for three different sensing modes: 1) obfuscation free sensing; 2) legitimate WiSee sensing coupled with a PhyCloak module that uses the proposed square-waved based self-channel estimation; 3) legitimate WiSee sensing coupled with a PhyCloak module that uses traditional pilot based self-channel estimation. We also vary the channel between Tx and the legitimate sensor by placing and removing obstacles. From the figure we see that the WiSee sensor equipped with PhyCloak module that uses square-wave based training achieves comparable performance as obfuscation-free sensing in both LoS and NLoS, while the WiSee sensor equipped with PhyCloak module that uses traditional training fails dramatically. This is because not enough self-interference cancellation is achieved in the presence of external transmissions using extant self-channel estimation techniques. **Takeaway:** *Square wave based training provides sufficient self-interference cancellation to preserve legitimate sensing with external transmission going on*.

5.4 Throughput Performance

As discussed in Section 4.4, PhyCloak would not hurt the average throughput by virtue of being a relay as long as it avoids parameter changes in the middle of packet transmissions. And, its online training would introduce some interference albeit of small measure. To validate that the net throughput benefit that a data receiver obtains from PhyCloak is not affected but even improved, we measured the throughput performance of a data link with and without PhyCloak in our testbed. We randomly picked 20 location triples to place a data transmitter, a data receiver, and PhyCloak. The data transmitter transmits back-to-back packets continuously, and we can thus see the throughput performance in the worst case where PhyCloak performs parameter updates in the middle of some packets. Figure 20 plots the CDF of the throughput with and without the PhyCloak. **Takeaway:** *The average throughput increases with the help of PhyCloak*.

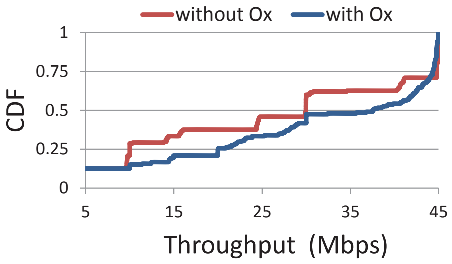

Figure 20: Throughput

6 Conclusion

We have shown that the threat created by recent developments in communication based sensing can be countered in a black-box fashion. PhyCloak obfuscates multi-dimensional physical signatures of human targets. We have empirically validated this for certain state-of-the-art sensors. We have also shown that when white box details of particular sensors can be obtained, PhyCloak can be refined to spoof those sensors. Notably, the methodology not only preserves but in fact improves the link throughput of the ongoing data transmissions, and supports co-existence of legitimate sensors while obfuscating illegitimate sensors.

Looking beyond the scope of the present work, we find that the methodology is readily generalized to protect against sensing of other types of physical targets and their properties, and allows for a network of PhyCloak devices to collaboratively cover a large region, the details of which are topics for future studies. In addition, when we extend our current single-antenna PhyCloak to a multiple-antenna system, how to fully exploit the space diversity provided by the multiple antennas is worth studying.

References

[1] Fadel Adib and Dina Katabi. *See through walls with WiFi!*, volume 43. ACM, 2013.

[2] Kamran Ali, Alex Xiao Liu, Wei Wang, and Muhammad Shahzad. Keystroke recognition using wifi signals. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 90–102. ACM, 2015.

[3] Dinesh Bharadia and Sachin Katti. Fastforward: fast and constructive full duplex relays. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 199–210. ACM, 2014.

[4] Dinesh Bharadia, Emily McMilin, and Sachin Katti. Full duplex radios. In *ACM SIGCOMM Computer Communication Review*, volume 43, pages 375–386. ACM, 2013.

[5] Igal Bilik and Joseph Tabrikian. Radar target classification using doppler signatures of human locomotion models. *IEEE Transactions on Aerospace and Electronic Systems*, 43(4):1510–1522, 2007.

[6] Bo Chen, Yue Qiao, Ouyang Zhang, and Kannan Srinivasan. Airexpress: Enabling seamless in-band wireless multi-hop transmission. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 566–577. ACM, 2015.

[7] Bo Chen, Vivek Yenamandra, and Kannan Srinivasan. Flexradio: Fully flexible radios and networks. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 205–218, 2015.

[8] Bo Chen, Vivek Yenamandra, and Kannan Srinivasan. Tracking keystrokes using WiFi. In *Proceedings of ACM MobiSys*, 2015.

[9] Asaf Cidon, Kanthi Nagaraj, Sachin Katti, and Pramod Viswanath. Flashback: decoupled lightweight wireless control. *ACM SIGCOMM Computer Communication Review*, 42(4):223–234, 2012.

[10] Theodoros Damoulas, Jin He, Rich Bernstein, Carla P Gomes, and Anish Arora. String kernels for complex time-series: Counting targets from sensed movement. In *2014 22nd International Conference on Pattern Recognition (ICPR)*, pages 4429–4434. IEEE, 2014.

[11] Shyamnath Gollakota, Haitham Hassanieh, Benjamin Ransford, Dina Katabi, and Kevin Fu. They can hear your heartbeats: non-invasive security for implantable medical devices. *ACM SIGCOMM Computer Communication Review*, 41(4):2–13, 2011.

[12] Alejandro Gonzalez-Ruiz, Alireza Ghaffarkhah, and Yasamin Mostofi. An integrated framework for obstacle mapping with see-through capabilities using laser and wireless channel measurements. *Sensors Journal, IEEE*, 14(1):25–38, 2014.

[13] Jin He and Anish Arora. A regression-based radar-mote system for people counting. In *International Conference on Pervasive Computing and Communications (PerCom), 2014 IEEE*, pages 95–102, March 2014. doi: 10.1109/PerCom.2014.6813949.

[14] Chih-Wei Huang and Kun-Chou Lee. Application of ica technique to pca based radar target recognition. *Progress In Electromagnetics Research*, 105:157–170, 2010.

[15] Youngwook Kim and Hao Ling. Human activity classification based on micro-doppler signatures using an artificial neural network. In *Antennas and Propagation Society International Symposium, 2008. AP-S 2008. IEEE*, pages 1–4. IEEE, 2008.

[16] Youngwook Kim and Hao Ling. Human activity classification based on micro-doppler signatures using a support vector machine. *IEEE Transactions on Geoscience and Remote Sensing*, 47(5):1328–1337, 2009.

[17] Vinit Kizhakkel, Rajiv Ramnath, and at el. Pulsed doppler radar target recognition based on micro-doppler signatures using wavelet analysis. In *IEEE High Performance Extreme Computing Conference (HPEC)*, September 2014.

[18] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, pages 483–494. ACM, 2014.

[19] Kun-Chou Lee, Jhih-Sian Ou, and Ming-Chung Fang. Application of svd -reduction technique to pca based radar target recognition. *Progress In Electromagnetics Research*, 81:447–459, 2008.

[20] Yasamin Mostofi. Cooperative wireless-based obstacle/object mapping and see-through capabilities in robotic networks. *IEEE Transactions on Mobile Computing*, 12 (5):817–829, 2013.

[21] Rajalakshmi Nandakumar, Krishna Kant Chintalapudi, Venkat Padmanabhan, and Ramarathnam Venkatesan. Dhwani: secure peer-to-peer acoustic NFC. In *ACM SIGCOMM Computer Communication Review*, volume 43, pages 63–74. ACM, 2013.

[22] Jeffrey A Nanzer and Robert L Rogers. Bayesian classification of humans and vehicles using micro-doppler signals from a scanning-beam radar. *Microwave and Wireless Components Letters, IEEE*, 19(5):338–340, 2009.

[23] Byung-Kwon Park, Olga Boric-Lubecke, and Victor M Lubecke. Arctangent demodulation with DC offset compensation in quadrature doppler radar receiver systems. *IEEE Transactions on Microwave Theory and Techniques*, 55(5):1073–1079, 2007.

[24] Qifan Pu, Sidhant Gupta, Shyamnath Gollakota, and Shwetak Patel. Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, pages 27–38. ACM, 2013.

[25] RG Raj, VC Chen, and R Lipps. Analysis of radar human gait signatures. *Signal Processing, IET*, 4(3):234–244, 2010.

[26] Graeme E Smith, Karl Woodbridge, and Chris J Baker. Radar micro-doppler signature classification using dynamic time warping. *IEEE Transactions on Aerospace and Electronic Systems*, 46(3):1078–1096, 2010.

[27] Thayananthan Thayaparan, Sumeet Abrol, Edwin Riseborough, LJ Stankovic, Denis Lamothe, and Grant Duff. Analysis of radar micro-doppler signatures from experimental helicopter and human data. *IET Radar, Sonar & Navigation*, 1(4):289–299, 2007.

[28] Guanhua Wang, Yongpan Zou, Zimu Zhou, Kaishun Wu, and Lionel M Ni. We can hear you with Wi-Fi! In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, pages 593–604. ACM, 2014.

[29] Wei Wang, Alex X Liu, Muhammad Shahzad, Kang Ling, and Sanglu Lu. Understanding and modeling of wifi signal based human activity recognition. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 65–76. ACM, 2015.

[30] Yan Wang, Jian Liu, Yingying Chen, Marco Gruteser, Jie Yang, and Hongbo Liu. E-eyes: device-free location-oriented activity identification using fine-grained WiFi signatures. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, pages 617–628. ACM, 2014.

[31] Yang Yang and Ness B Shroff. Scheduling in wireless networks with full-duplex cut-through transmission. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 2164–2172. IEEE, 2015.

[32] Yanzi Zhu, Yibo Zhu, Ben Y Zhao, and Haitao Zheng. Reusing 60ghz radios for mobile radar imaging. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 103–116. ACM, 2015.