

Proximity Detection with Single-Antenna IoT Devices

Timothy J. Pierson, Travis Peters, Ronald Peterson, David Kotz
Department of Computer Science, Dartmouth College, Hanover, NH, USA 03755

ABSTRACT

Providing secure communications between wireless devices that encounter each other on an ad-hoc basis is a challenge that has not yet been fully addressed. In these cases, close physical proximity among devices that have never shared a secret key is sometimes used as a basis of trust; devices in close proximity are deemed trustworthy while more distant devices are viewed as potential adversaries. Because radio waves are invisible, however, a user may believe a wireless device is communicating with a nearby device when in fact the user's device is communicating with a distant adversary. Researchers have previously proposed methods for multi-antenna devices to ascertain physical proximity with other devices, but devices with a single antenna, such as those commonly used in the Internet of Things, cannot take advantage of these techniques.

We present theoretical and practical evaluation of a method called SNAP – SiNgle Antenna Proximity – that allows a single-antenna Wi-Fi device to quickly determine proximity with another Wi-Fi device. Our proximity detection technique leverages the repeating nature Wi-Fi's preamble and the behavior of a signal in a transmitting antenna's *near-field* region to detect proximity with high probability; SNAP never falsely declares proximity at ranges longer than 14 cm.

ACM Reference Format:

Timothy J. Pierson, Travis Peters, Ronald Peterson, David Kotz. 2019. Proximity Detection with Single-Antenna IoT Devices. In *The 25th Annual International Conference on Mobile Computing and Networking (MobiCom '19)*, October 21–25, 2019, Los Cabos, Mexico. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3300061.3300120>

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 ACM 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.

MobiCom '19, October 21–25, 2019, Los Cabos, Mexico

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6169-9/19/10...\$15.00

<https://doi.org/10.1145/3300061.3300120>

1 INTRODUCTION

People and the devices they wear or carry may soon encounter dozens, possibly hundreds, of new devices each day if predictions about the Internet of Things (IoT) come true. These IoT devices are envisioned to share data and actuator control information among themselves, and some of that information may be privacy sensitive or have security implications. This situation suggests that devices that have never met, nor shared a secret, must somehow have a way to securely communicate that is consistent with user intent.

Providing secure, user-intended communications between devices that encounter each other on an ad-hoc basis is a challenge that has not yet been fully addressed [13]. The main difficulty is that the newly discovered devices do not have a common point of trust. In these situations, researchers have previously proposed using physical proximity as one basis of trust [8, 19, 21, 24, 25, 28, 30]. The idea is that a user can express intent to introduce two devices by bringing them within a few centimeters of each other, at least temporarily, and then taking an action such as pressing a button. Adversaries, however, are assumed to be unable to come into such close proximity (e.g., an adversary does not break into a home to gain close physical access to devices). The physical proximity between legitimate devices then forms a basis of trust (whether the devices have been compromised is out of scope) and nearby devices can then bootstrap a secure connection among themselves. A distant adversary, however, may attempt to trick a user's devices into accepting a malicious payload by impersonating a nearby legitimate device.

Several techniques have been proposed to combat such impersonation attacks. Often these techniques rely on short-range out-of-band communication where devices use a secret channel for communication that is impervious to observation or interference by an adversary. These methods frequently require additional hardware such as accelerometers [38], light sensors [23], or specialized radio frequency (RF) devices such as NFC [1]. The required out-of-band hardware may not be present on some devices and these approaches often require complex processing that exceeds the capabilities of many

embedded devices. Other approaches to thwarting distant adversaries use in-band RF, but rely on multiple antennas to simultaneously measure signal strength to determine proximity [8, 28] or to locate a device in three dimensions [22, 35]. Single-antenna IoT devices with limited hardware that follow standard communication protocols, however, cannot use these techniques.

In this paper we present a novel method called SNAP – SiNgle Antenna Proximity – that enables a single-antenna device following the standard Wi-Fi protocol to quickly determine when it is in close proximity to a transmitting antenna. Our technique leverages the repeating nature of Wi-Fi’s preamble and the characteristics of a transmitting antenna’s *near field* region (i.e., the region physically close to the antenna) to detect proximity with a transmitter. When a receiving device is physically close to a transmitter, near-field effects will cause repeated portions of the preamble to differ in phase and amplitude, whereas when the device is far from the transmitter, the repeated portions of the preamble will be received with a consistent phase and amplitude. We use the presence or absence of phase and amplitude mismatches to determine whether the single-antenna device is near a transmitter.

Contributions

We present the following contributions in this paper:

- a novel method for a single-antenna device to quickly and accurately determine when it is in close proximity with a transmitting device;
- a reference Wi-Fi implementation that performs the same frame decoding steps *any* Wi-Fi device must perform; and
- an experimental evaluation of the technique using several popular types of antennas.

2 WI-FI PREAMBLE

In this section we briefly describe the Wi-Fi preamble, focusing on the repeating portions of the Long Training Field (LTF). We show that when a receiver is far from a transmitter, even though the channel changes the transmitted signal, the repeated portions of the LTF are changed consistently and are received with matching phase and amplitude. In Section 3 we show that when the receiver is in the transmitter’s near field, this consistency does not hold, allowing SNAP to determine proximity with the transmitter.

2.1 PHY layer preamble format

Every Orthogonal Frequency Division Multiplexing (OFDM) Wi-Fi frame begins with a physical (PHY) layer preamble




Figure 1: Wi-Fi OFDM PHY preamble format. Each Wi-Fi frame begins with a PHY layer preamble consisting of a Short Training Field (STF) and a Long Training Field (LTF). The STF and LTF are followed by a Signal field and then the frame’s data [17].

to aid in synchronizing the transmitter and receiver.¹ The format of the PHY layer preamble is shown in Figure 1 and consists of a Short Training Field (STF) followed by a Long Training Field. These fields are followed by a Signal Field and then the Wi-Fi frame’s data.

The STF consists of 10 identical short training symbols (denoted t_1 through t_{10} in Figure 1) where each STF symbol is sampled 16 times, for a total of 160 samples. The STF is used by the receiver for frame detection, automatic gain control, coarse frequency offset estimation, and rough symbol timing synchronization. We discuss the STF in more detail in Section 4 below.

The Long Training Field follows the Short Training Field and is used by the receiver for fine frequency correction and channel estimation. Important to our proximity detection technique, the LTF contains two identical 64-sample symbols. We discuss the LTF in Section 2.2.

A Signal field encoded with Binary Phase Shift Keying (BPSK) follows the LTF. It provides information about the rest of the Wi-Fi frame including the number of bytes and the encoding scheme used on the frame’s data.

Data carried by the Wi-Fi frame comes after the Signal field. Each OFDM data symbol consists of a 16-sample guard interval (denoted GI in Figure 1) and 64 samples carrying the actual data.

2.2 Long Training Field

The LTF consists of a 32-sample guard interval (denoted GI2 because it is twice as long as other guard intervals in the frame) followed by two identical 64-sample OFDM symbols that are denoted T_1 and T_2 in Figure 1. The guard interval together with T_1 and T_2 make a total of 160 samples in the LTF. Because T_1 and T_2 are identical, the phase and amplitude of sample i in symbol T_1 matches the phase and amplitude of sample $i + 64$ in T_2 ,

¹Here we focus on 20 MHz wide channels but the technique could easily be extended for wider channels.




Figure 2: Time domain amplitude of the Long Training Field. In the time domain, sample i in T_1 matches sample $i + 64$ in T_2 in phase and amplitude. Here we highlight how sample 16 in T_1 matches sample 80 ($16 + 64$) in T_2 .

where $i = 0 \dots 63$. This relationship between samples is shown in Figure 2.

The time-domain samples can be converted into an equivalent frequency-domain representing by taking a *Discrete Fourier Transform (DFT)*, which is nearly always implemented in real hardware with a *Fast Fourier Transform (FFT)*. Wi-Fi receivers perform a 64-point FFT over the received time-domain samples to transform the time-domain samples into the frequency domain. The FFT operation yields 64 complex numbers representing the phase and amplitude of 64 subcarriers, indexed from -32 to +31 [17]. Figure 3 shows T_1 and T_2 represented in the frequency domain. Provided samples in the time domain in T_1 match corresponding samples in T_2 at the receiver, the phases and amplitudes of each subcarrier after an FFT of the samples in T_1 will also match the phases and amplitudes of each subcarrier after an FFT of the samples in T_2 . If the samples in the time domain do not match, however, the phases and amplitudes of the subcarriers will also not match.

2.3 Channel State Information

The channel between the transmitter and receiver will modify the transmitted signal because the signal takes multiple paths while in flight, reflecting off or passing through objects in the environment. These multi-path signals add up constructively or destructively at the receiver and the result is that the samples will not be received with the same phase and amplitude with which

##	Re	Im	##	Re	Im	##	Re	Im	##	Re	Im
-32	0.000	0.000	-16	1.000	0.000	0	0.000	0.000	16	1.000	0.000
-31	0.000	0.000	-15	1.000	0.000	1	1.000	0.000	17	-1.000	0.000
-30	0.000	0.000	-14	1.000	0.000	2	-1.000	0.000	18	-1.000	0.000
-29	0.000	0.000	-13	1.000	0.000	3	-1.000	0.000	19	1.000	0.000
-28	0.000	0.000	-12	1.000	0.000	4	1.000	0.000	20	-1.000	0.000
-27	0.000	0.000	-11	-1.000	0.000	5	1.000	0.000	21	1.000	0.000
-26	1.000	0.000	-10	-1.000	0.000	6	-1.000	0.000	22	-1.000	0.000
-25	1.000	0.000	-9	1.000	0.000	7	1.000	0.000	23	1.000	0.000
-24	-1.000	0.000	-8	1.000	0.000	8	-1.000	0.000	24	1.000	0.000
-23	-1.000	0.000	-7	-1.000	0.000	9	1.000	0.000	25	1.000	0.000
-22	1.000	0.000	-6	1.000	0.000	10	-1.000	0.000	26	1.000	0.000
-21	1.000	0.000	-5	-1.000	0.000	11	-1.000	0.000	27	0.000	0.000
-20	-1.000	0.000	-4	1.000	0.000	12	-1.000	0.000	28	0.000	0.000
-19	1.000	0.000	-3	1.000	0.000	13	-1.000	0.000	29	0.000	0.000
-18	-1.000	0.000	-2	1.000	0.000	14	-1.000	0.000	30	0.000	0.000
-17	1.000	0.000	-1	1.000	0.000	15	1.000	0.000	31	0.000	0.000

Figure 3: Frequency domain representation of T_1 and T_2 in the LTF. Re is the real component and Im is the imaginary component of the complex number representing the phase and amplitude of each subcarrier [17].

they were transmitted. This signal change suggests the possibility that samples in T_1 may not have the same phase and amplitude as the corresponding sample in T_2 when the signal is received. We see next, however, that those samples will match (except for random noise) when the receiver is not in the transmitter’s near-field region.

The channel between the transmitter and receiver can be mathematically expressed as [33]

$$\mathbf{y}[i] = \mathbf{H}\mathbf{x}[i] + \mathbf{w}[i] \quad (1)$$

where $\mathbf{y}[i]$ is the i^{th} received sample, \mathbf{H} is the channel matrix representing the changes to the signal caused by the channel, $\mathbf{x}[i]$ is i^{th} the transmitted sample, and $\mathbf{w}[i]$ is noise received with sample i . In a static environment (e.g., no moving objects), \mathbf{H} is time invariant and causes the same shift in phase and amplitude for all samples in \mathbf{x} because all transmitted samples take the same multipaths from sender to receiver. Neglecting noise, the result is that sample $\mathbf{y}[i]$ still matches sample $\mathbf{y}[i + 64]$ in phase and amplitude, even though they no longer match $\mathbf{x}[i]$ due to the effects of \mathbf{H} .

This phase and amplitude change in the received sample compared with the transmitted sample is normal for wireless communication and is one of the reasons why Wi-Fi uses a preamble. The phase and amplitude of the preamble samples are pre-defined by the Wi-Fi specification and are known to both the transmitter and receiver. The transmitter sends the preamble at the pre-defined phase and amplitude and the receiver uses these known phase and amplitude values in the STF to detect the start of the frame and apply a coarse frequency

correction. Next it uses the LTF to synchronize symbol timing and apply fine frequency correction. Finally, because each subcarrier may be impacted differently by the channel, the receiver performs an FFT of the received time-domain signal to independently measure the phase and amplitude of each frequency-domain subcarrier in the LTF. The receiver computes the difference from the known transmitted phases and amplitudes for each subcarrier (see Figure 3) and the received phases and amplitudes to estimate the channel’s impact on each subcarrier. This estimate is called *Channel State Information* or CSI. The receiver uses this estimate from the LTF to correct for the channel’s effects.

2.4 Coherence time

Above we consider an environment with no moving objects and we see in Equation (1) that \mathbf{H} is time invariant so corresponding samples in T_1 and T_2 will be received with identical phase and amplitude (except for noise). In the real world, however, the transmitter, receiver, or other objects may be moving and that movement may impact the signal. A channel is said to be *coherent* if it is stable over a particular time interval. We can calculate the needed coherence time, T_c , for the corresponding portions of the preamble. If the channel is coherent over T_c then the corresponding samples will be received with the same phase and amplitude.

Wi-Fi samples at 20 MHz, meaning it takes 20 million samples per second. The time for one sample, T_s , is then $1/(20,000,000 \text{ samples/second})$ which equates to 50 ns. T_1 and T_2 are a total of 128 samples long, and because we are interested in how T_1 matches T_2 , we require a coherence time of $6.4 \mu\text{s}$ ($50 \text{ ns/sample} \times 128 \text{ samples} = 6.4 \mu\text{s}$). In this case, if the channel is stable over $6.4 \mu\text{s}$, then T_1 will match T_2 (aside from noise).

2.5 Moving objects

Moving objects can potentially cause a mismatch by changing the length of the signal’s path as it travels from transmitter to receiver. The length of the path affects the phase and amplitude of the signal according to the following formula [33]

$$\mathbf{H} = \sum_{p=1}^P a_p e^{-j2\pi d_p/\lambda} \quad (2)$$

where a_p is the attenuation of the signal along the path p , d_p is the length of path p , and λ is the wavelength. The length of path p may change as the transmitter, receiver, or multipath-inducing objects move. To cause a significant change in the signal between corresponding samples, however, the movement would need to cause a change in

path length of more than one-quarter wavelength (and one-half wavelength to cause maximum change) [33]. In Wi-Fi’s 2.4 GHz band, the wavelength λ is approximately 12 cm, suggesting that an object would need to move approximately $\lambda/4 \approx 3 \text{ cm}$ in $6.4 \mu\text{s}$ to significantly impact the phase and amplitude between corresponding LTF samples. This translates to a speed of over 17,000 km/hour (and roughly twice this speed for Wi-Fi’s 5 GHz band). Given the extraordinary speed an object would need to be moving to cause a substantial change in path length in the short coherence time needed for the preamble, we eliminate changing path lengths as a possible explanation for corresponding LTF samples to have different phases and amplitudes.

2.6 Implication

The implication of this section is that the channel between the receiver and a transmitter will not cause a significant difference in the phase and amplitude between repeating portions of the LTF, even in the presence of moving objects, but this section implicitly assumes that the transmitter is far from the receiver. We find in the next section that near-field effects *can* cause differences in the corresponding LTF samples when the transmitter is near the receiver. We use those differences to detect when the receiver is near the transmitter. If those differences are not present, we infer that the receiver is far from the transmitter.

3 NEAR FIELD

The area around a transmitting antenna is generally classified into three different regions: (1) the reactive near-field is closest to the transmitting antenna, (2) the radiating near-field begins after the reactive near-field, and (3) the far-field begins after the radiating near-field and extends to infinity. These regions are shown in Figure 4. The boundaries between regions are not sharp, but instead transition gradually.

Using the orientation depicted in Figure 5, where the antenna is aligned vertically with the z axis, a signal’s magnetic fields \mathbf{H} relative to each axis² are determined by the following formulas [4]

$$H_r = H_\theta = 0 \quad (3a)$$

$$H_\phi = j \frac{k I_0 l_t \sin \theta}{4\pi r} \left[1 + \frac{1}{jkr} \right] e^{-jkr} \quad (3b)$$

²Here \mathbf{H} refers to the magnetic field, whereas previously the same symbol referred to CSI. Unfortunately, this overloading is common in the literature.




Figure 4: Regions surrounding a transmitting antenna. A transmitting antenna with length l_t has three surrounding regions: (1) the reactive near-field, (2) the radiating near-field, and (3) the far-field. D is the length of the transmitting antenna l_t plus the length of the receiving antenna, l_r [4].

and the electric fields \mathbf{E} are determined by

$$E_r = \eta \frac{I_0 l_t \cos \theta}{2\pi r^2} \left[1 + \frac{1}{jkr} \right] e^{-jkr} \quad (4a)$$

$$E_\theta = j\eta \frac{k I_0 l_t \sin \theta}{4\pi r} \left[1 + \frac{1}{jkr} - \frac{1}{(kr)^2} \right] e^{-jkr} \quad (4b)$$

$$E_\phi = 0 \quad (4c)$$

where $j = \sqrt{-1}$, $k = 2\pi/\lambda$ is the wavenumber, I_0 is current applied to the transmitter, l_t is the length of the transmitting antenna, $\eta = 120\pi$ is the intrinsic impedance of free space, θ is the vertical angle between the transmitter and receiver, ϕ is the horizontal angle between the transmitter and receiver, and r is the distance extending radially from the transmitter.

3.1 Reactive near-field region

The reactive near-field region is the region closest to the transmitting antenna, where $kr < 1$ (or equivalently, where $r < \lambda/2\pi$). In this region the reactive (e.g. non-radiating) field dominates and there is a high content of non-propagating stored energy. Here the wavefront is not spherical because the electric and magnetic fields are not yet aligned, and in addition to the radiated




Figure 5: Antenna orientation. To provide a common reference, the transmitting antenna is typically assumed to be aligned with the vertical (z) axis as shown. Redrawn from Balanis [4].

energy described by the first term in brackets in Equations (3b), (4a), and (4b), there is a great deal of stored, non-propagating energy because the second and third terms inside the brackets dominate at close range.

With real antennas, the reactive near-field region is commonly estimated to extend from the surface of the antenna to roughly R_1 , defined as [4]

$$R_1 = 0.62\sqrt{D^3/\lambda} \quad (5)$$

where $D = l_t + l_r$ is the combined length of the transmitting antenna, l_t , and the receiving antenna, l_r , and λ is the signal wavelength. With Wi-Fi 2.4 GHz band, and quarter-wavelength dipole antennas, this region extends to roughly 2.7 cm from the transmitter. In Wi-Fi's 5 GHz band this region extends to roughly 1.1 cm.

3.2 Radiating near-field (Fresnel) region

Sometimes referred to as the *Fresnel* or *intermediate field*, the radiating near-field region is the area between the reactive near-field and far-field regions. In this region $kr > 1$ and the electric and magnetic fields are predominantly in phase, but the wavefront is still not yet spherical as it is in the far-field region. Examining Equations (3b) and (4a) we see that, unlike in the reactive near field, the first term in the brackets (1) begins to dominate the second term ($\frac{1}{jkr}$) because kr is greater than one. Likewise, in Equation (4b), the first term in the brackets (1) begins to dominate the second ($\frac{1}{jkr}$)

and third terms ($\frac{1}{(kr)^2}$). Because of the increasing value of kr compared with the reactive near-field region, the energy in the radiating near field is largely real, that is, radiated energy.

We can estimate the average power of the signal, W , using the following equation [4]:

$$\mathbf{W} = \frac{1}{2}(\mathbf{E} \times \mathbf{H}^*) \quad (6)$$

where $*$ denotes complex conjugate and \mathbf{E} and \mathbf{H} are determined using Equations (3) and (4). W can be decomposed into its radial, W_r , and vertical, W_θ components as follows [4]:

$$W_r = \frac{\eta}{8} \left| \frac{I_0 l_t}{\lambda} \right|^2 \frac{\sin^2 \theta}{r^2} \left[1 - j \frac{1}{(kr)^3} \right] \quad (7a)$$

$$W_\theta = j\eta \frac{k|I_0 l_t|^2 \cos \theta \sin \theta}{16\pi^2 r^3} \left[1 + \frac{1}{(kr)^2} \right]. \quad (7b)$$

Mapping the power of each of these components, we see in Figure 6 for Wi-Fi's 2.4 GHz band with quarter-wavelength antennas, at distances larger than roughly 5 cm the W_θ component begins to dominate the W_r component as it does in the far field. At distances closer than about 5 cm the radial component is stronger than the vertical component. This relative strength suggests the power pulses inward and outward near the transmitter, whereas at greater distances, the radial component dies out and vertical component takes over. This vertical component domination is indicative of signals in the far-field region, whereas radial component domination is indicative of signals in the radiating near-field region.

With real antennas, the radiating near-field region is commonly estimated to extend from R_1 to R_2 , where R_2 is defined [4]

$$R_2 = 2D^2/\lambda. \quad (8)$$

With Wi-Fi's 2.4 GHz band and quarter-wavelength dipole antennas, Equation (8) suggests this region extends to approximately 6.2 cm from the transmitter. This estimate roughly matches the results shown in Figure 6 using Equations (7) where the vertical component of the energy begins to dominate as it does in the far field. We note, however, that this boundary is not a sharp distinction between the radiating near-field and the far-field. We see in Figure 6 that the radial and vertical components are nearly equivalent for some distance past this point, but that by 12 cm distance W_θ is roughly three times stronger than W_r .




Figure 6: Power of the radial and vertical components of a signal. Using Wi-Fi's 2.4 GHz band and Equation (7), the vertical component W_θ begins to dominate the radial component W_r at about 5 cm.

3.3 Far-field (Fraunhofer) region

The far field, sometimes referred to as the *Fraunhofer* region, is the area far from the transmitting antenna where $kr \gg 1$. Because kr is large in the far field, several of the terms in Equations (3) and (4) become extremely small and the \mathbf{E} and \mathbf{H} fields can be approximated by the much simpler formulas [4]

$$E_\theta \simeq j\eta \frac{kI_0 l_t e^{-jkr}}{4\pi r} \sin \theta \quad (9a)$$

$$E_r \simeq E_\phi = H_r = H_\theta = 0 \quad (9b)$$

$$H_\phi \simeq j \frac{kI_0 l_t e^{-jkr}}{4\pi r} \sin \theta. \quad (9c)$$

In Equation (9) we see that the electric and magnetic fields are aligned orthogonal to each other (e.g., θ is orthogonal to ϕ), transverse to the direction of propagation, and are in time synchronization. This alignment creates a spherical wavefront with average power given by Equation (6).

3.4 Near-field impact on LTF samples

At ranges closer than roughly R_2 , the overall \mathbf{E} and \mathbf{H} fields are not in phase with respect to time, and because those fields do not have equal magnitude, they form a vector that rotates in time in a plane parallel to the

direction of propagation, rather than the stable orthogonal relationship in the far-field region [4]. Wi-Fi samples taken as the \mathbf{E} and \mathbf{H} fields rotate can result in different phase and amplitude readings between corresponding samples in the LTF. We see in Section 5, as suggested by Figure 6, there is a difference in phase and amplitude in the corresponding LTF samples out to about 12 cm.

4 IMPLEMENTATION

We wanted to test the insight from Section 3 that the phase and amplitude of corresponding Long Training Field samples would change when the receiver is in close proximity to a transmitting antenna (e.g., when the receiver is in the transmitter’s near-field region) and would not change when the receiver is far from the transmitter. To perform these tests, however, we needed access to the raw LTF samples measured by the receiver. All Wi-Fi adapters must evaluate LTF samples to correct for channel effects, so the information required to implement SNAP is available on *any* Wi-Fi adapter, but sample-level data is not provided outside the firmware of commercial Wi-Fi adapters. Because we could not get the granular phase and amplitude data from commercial Wi-Fi adapters (even with the CSI tool frequently used by researchers [16]), we built a custom Wi-Fi receiver using Software Defined Radio (SDR) and GNU Radio. We used our Wi-Fi receiver to evaluate signals sent by four different types of commercial-off-the-shelf (COTS) transmitters commonly found in consumer electronics.

4.1 Receiver

We used an Ettus Research USRP N210 SDR [11] with a UBX40 daughterboard and GNU Radio [32]. The USRP SDR hardware allowed us to receive signals from 10 MHz to 6 GHz, enough bandwidth to cover both the 2.4 and 5 GHz Wi-Fi bands, and GNU Radio is open-source software that processes signals sent by the SDR hardware. Our GNU Radio software implementation closely followed Bloessl et al. [6], but focused on comparing corresponding LTF samples.

We used our custom Wi-Fi receiver software with the USRP SDR hardware and a quarter-wavelength dipole antenna to receive Wi-Fi signals transmitted by several different types of COTS Wi-Fi antennas. Our receiver is shown in Figure 7.

4.2 Transmitters

Because dipole and micropatch antennas are the most common antennas found in consumer electronics [4], we



Figure 7: Receiver. We used a USRP 210N Software Defined Radio with a quarter-wavelength dipole antenna as a Wi-Fi receiver. The USRP was connected to the antenna via a 3 m RF cable and to a laptop via a 2 m ethernet cable (laptop not shown).

focused our testing on those types of transmitting antennas. We tested two different types of dipole antennas and two different types of micropatch antennas.

4.2.1 Dipole transmit antennas. We used both a half-wavelength and a quarter-wavelength dipole antenna connected to an internal Intel Ultimate N WiFi Link 5300 card [18] installed in a Linux laptop to test different transmitting dipole antennas. We attached each dipole antenna to a circular base to hold the antenna stationary and upright during testing. We connected the base to the laptop with a 3 m long RF cable as shown in Figure 8. We tested each type of dipole antenna separately.

4.2.2 Micropatch transmit antennas. We tested a Panda Ultra Wireless N USB Adapter [27] attached to a Linux laptop via a 1 m USB extender cable. We also used a micropatch antenna connected to the Intel 5300 card via a 3 m RF cable.

5 EVALUATION

We transmitted 1,000 Wi-Fi frames from each of the four different types of antennas (two dipole plus two micropatch) using BPSK 1/2 encoding [17] on Wi-Fi channel 1 at distances ranging from 2 cm to 3 m. The Wi-Fi frames were sent on Layer 2 using a Python program written with Scapy [5]. Frames were not acknowledged by our custom Wi-Fi receiver. All tests were conducted in a busy computer science lab bustling with student activity. To accurately measure the preamble deviation at various ranges, all antennas were stationary and oriented vertically. Testing with antennas offset at angles of 45 and



Figure 8: Transmitters. We used a quarter-wavelength and a half-wavelength dipole antenna as well as a micropatch antenna connected via a 3 m cable to an internal Intel Ultimate N WiFi Link 5300 adapter installed in a Linux laptop. We also used a Panda Ultra Wireless N USB adapter connected via a 1 m USB cable.

90 degrees had different signal strength due to polarity, but had similar LTF differences.

In our experiments we dropped frames that were not properly decoded as valid Wi-Fi frames by our receiver. We used a fixed gain, but Automatic Gain Control (AGC) might have allowed our system to drop fewer frames (we did not quantify the number of dropped frames, however, because our goal was not to evaluate the effectiveness of the receiver at decoding frames). Blossel et al. note that for strong signals, such as those from a nearby transmitter, frames can be successfully decoded by a receiver with relatively low gain. A strong signal, received with high gain, can result in clipping if the signal overdrives the analog-to-digital converter. Frames cannot be successfully decoded when clipping occurs. Similarly, weak signals cannot be decoded without additional gain due to high quantization error [7]. Our fixed gain allowed the receiver to avoid these issues and decode signals with high probability over our experimental range.

A commercial Wi-Fi product would likely incorporate AGC. In that case, gain would be automatically adjusted downward for nearby strong signals and upward for weaker, more distant signals. The AGC adjustment normally happens during the Short Training Field (STF) and is stabilized prior to the LTF. Because SNAP relies on the relative difference between repeated portions of the LTF (we only use the STF for frame detection), and because the AGC is stabilized prior to the LTF, the presence of AGC would not significantly change our results. In engineering a future product, the developer could also choose to turn off AGC and configure the radio for a low

fixed gain during the operation of SNAP (because it is only intended to work at close range), and could enable AGC during other modes of operation.

5.1 Preamble deviation

To measure the difference in phase and amplitude of the matching portions of the LTF, we calculate the total Euclidean distance between the phase and amplitude of subcarriers in T_1 and T_2 as:

$$E_j = \sum_{k=-32}^{31} \left[(\Re(Y_1[k]) - \Re(Y_2[k]))^2 + (\Im(Y_1[k]) - \Im(Y_2[k]))^2 \right]^{\frac{1}{2}} \quad (10)$$

where E_j is the total Euclidean distance between the phase and amplitude of all subcarriers k for frame j , and where Y_1 is the result of an FFT over T_1 and Y_2 is the result of an FFT over T_2 , $\Re(Y_x[k])$ is the real component and $\Im(Y_x[k])$ is the imaginary component of each subcarrier k in Y_x , for $x \in \{1, 2\}$. We call this difference E_j the *preamble deviation* of a frame. If the subcarriers in the two corresponding portions of the LTF are substantially the same, then the preamble deviation will be small. If the subcarriers are different in the two corresponding portions of the LTF, then the preamble deviation will be large.

Figure 9 shows the distance between Y_1 and Y_2 for subcarrier 1 of one frame when the transmitter was located at 6 cm from the receiver and for subcarrier 1 of another frame sent from 30 cm. We see that at 30 cm Y_1 matches Y_2 , but at 6 cm Y_1 does not match Y_2 due to near-field effects discussed in Section 3. Each Wi-Fi symbol, however, is comprised of 64 subcarriers. Figure 10 shows distance between Y_1 and Y_2 for all subcarriers of one frame. We see at 30 cm Y_1 and Y_2 match for all subcarriers, but at 6 cm many subcarriers do not match.

The preamble deviation E_j sums the Euclidean distance between all subcarriers in a frame. Figure 11 shows the distribution of preamble deviations according to Equation (10) for 1,000 Wi-Fi frames received from a transmitting half-wavelength dipole antenna. We see at close range E_j is typically large, but varies due to near-field effects. At long range (greater than about 12 cm) the preamble deviation is small and has much lower variance because the near-field effects have attenuated to near zero as predicted by Equations (3) and (4). For brevity we omit the distribution from other types of transmitting antennas, but they follow a similar pattern, with each having small preamble deviations and low variability beyond about 12 cm.




Figure 9: Constellation diagram showing the distance between Y_1 and Y_2 for subcarrier 1. At long range (30 cm) the distance between Y_1 and Y_2 for subcarrier 1 is small. At close range (6 cm) the difference is large.

All subcarriers of one frame




Figure 10: Constellation diagram showing the distance between Y_1 and Y_2 for all subcarriers of one frame. Y_1 and Y_2 match at long (30 cm) range, but do not match at close (6 cm) range due to near-field effects.




Figure 11: Distribution of preamble deviations E_j for 1,000 Wi-Fi frames transmitted by a half-wavelength antenna at each distance. The red line indicates the median value, the box indicates the 75th and 25th percentile, and the whiskers indicate the maximum and minimum value.

We calculate the average preamble deviation over a number of frames for each antenna type as:

$$A_t = \frac{1}{n} \sum_{j=1}^n E_j \quad (11)$$

where $t \in \{\text{half-wavelength, quarter-wavelength, micro-patch, Panda}\}$ is the type of antenna used to send Wi-Fi frames and $n = 1,000$ is the number of frames received. The average preamble deviation over all 1,000 frames sent at each distance for each antenna type is shown in Figure 12 for distances from 2 cm to 3 m. As predicted in Section 3, at short range we see large preamble deviations and at distances beyond roughly 12 cm, we see small preamble deviations. This relationship holds across all antenna types and suggests that a single-antenna device can monitor the preamble deviation and declare proximity when the preamble deviation rises above a predetermined threshold.

5.2 Proximity detection

We would like a single-antenna device to be able to determine proximity with a transmitting device without help from another source. A simple way to make that proximity determination using the data in Figure 12 would be to set a threshold, τ , where if the preamble deviation for a frame is greater than τ , the single antenna




Figure 12: Average preamble deviation by distance and antenna type using Equation (11). The average preamble deviation over 1,000 Wi-Fi frames is large at close range and small at long range for each antenna type.

device declares proximity, otherwise it does not declare proximity.

If τ were set relatively high, say around 0.2 (indicated by the dashed line in Figure 12), then the single-antenna device would like not falsely declare proximity when the transmitter is far away because the preamble deviations are never over the threshold for any transmitting antenna type at distances over 14 cm. If the single-antenna device uses only one frame to determine proximity, however, it could be the case that the particular frame happens to have a low preamble deviation as indicated by the whiskers in Figure 11 and the single-antenna device would fail to recognize proximity even though it should.

This situation suggests that proximity detection with a single-antenna device may benefit from measuring the preamble deviation from multiple frames before declaring proximity. Instead of relying on the preamble deviation from a single frame, we can use Equation (11) to average the preamble deviation from multiple frames and then compare that average value with threshold τ .

To determine the likelihood of detecting proximity using the average preamble deviation from multiple frames, we created a Monte Carlo simulation where we randomly sampled n frames from the 1,000 Wi-Fi frames we captured at each distance between transmitter and receiver, and then calculated an average preamble deviation over those n frames. The single-antenna device declares proximity if the average is greater than τ .




Figure 13: Likelihood of declaring proximity. The likelihood of a single-antenna device declaring proximity in 1,000 Monte Carlo simulation runs at each distance where proximity was declared if the average preamble deviation for different numbers of frames was greater than $\tau = 0.2$. We see proximity declared with high probability at close range and never declared over 14 cm.

In Figure 13 we show the likelihood of declaring proximity from 1,000 runs of the Monte Carlo simulation that randomly selected $n \in \{1, 2, 5, 10, 20\}$ Wi-Fi frames at each distance with $\tau = 0.2$. The results shown are the average over all antenna types. We see there is a high likelihood of declaring proximity when the transmitter, regardless of antenna type, is within about 9 cm if the receiver uses more than one frame. In fact, we see that using only two frames performs much better than a single frame, consistently predicting proximity at short ranges, and never predicting proximity at ranges over 14 cm. Using more than two frames results in improved detection probability, however, the amount of improvement decreases as the number of frames used increases. Because the single-antenna device is able to accurately determine proximity using a small number of frames exchanged on a single Wi-Fi channel, SNAP can be used by devices during the standard Wi-Fi association process to ensure devices are in close proximity. This differs significantly from non-standard frequency hopping approaches proposed by other researchers [34].

5.3 Future exploration

Our experiments suggest that the preamble deviations caused by the near field of a transmitting antenna allow a single-antenna device to reliably detect proximity at

distances under 9 cm and never falsely declare proximity over 14 cm. We focused our testing on dipole and micropatch antennas because they are the most common antennas used in consumer devices. There are, however, a myriad of other types of antennas: horn, helical, parabolic dish, spiral, loop, spiral, Yagi, and bow tie, to name a few. In future work we plan to examine how the near field affects these other types of antennas. We also used the default transmission power settings on the Wi-Fi devices and have not yet tested with non-standard power settings, but we expect little difference in results. Finally, we note that our experiments were conducted using omni-directional antennas. Directional antennas may have different characteristics. We believe, however, that our work here is a starting point and opens an important new area of research that warrants further investigation.

6 LIMITATIONS

Above we see that a single-antenna device is able to reliably detect proximity with a transmitting device if the single-antenna device is in the near field of the transmitter. Our experiments suggest that the single-antenna device can reliably detect proximity out to roughly 9 cm for dipole and micropatch antennas. These results assume, however, that the transmitter is sending properly formed Wi-Fi frames where T_1 in the Long Training Field matches T_2 . It could be the case that a sophisticated adversary transmits a malformed Wi-Fi preamble where T_1 does not match T_2 in an attempt to trick the single-antenna device into falsely declaring proximity. That adversary might pre-compute a mismatched LTF and send those samples with a Software Defined Radio. In this case, a single-antenna device could potentially be tricked into declaring proximity when in fact the transmitter is far away. We propose two ways a single-antenna device might overcome such a sophisticated adversary and we also discuss scenarios where proximity alone is insufficient for trust.

6.1 Help from a trusted device

A single-antenna device might be able to overcome an adversary transmitting malformed preambles if the single-antenna device has a pre-existing trusted relationship with another device known to be far away (perhaps by measuring the preamble of signals from the trusted device), such as a Wi-Fi router. If the single-antenna device detects a preamble deviation greater than τ , it could ask the trusted device to confirm the trusted device sees a matching preamble from the same transmitter. Provided the trusted device is located more than about 18 cm from

the single-antenna device (i.e., two times the effective range of the preamble detection technique to rule out a legitimate transmitter positioned in between the single-antenna device and trusted device), the trusted device will see a matching preamble if the preamble is properly formed and can inform the single-antenna device. If the preamble is malformed, both devices will see the high preamble deviation; with the trusted device’s input, the single-antenna device can conclude an adversary sent the frames with malformed preambles.

6.2 Signal strength

In many uses there will not be a trusted device with which the single-antenna device can confer. In those instances, the single-antenna device can examine the strength of the signal when it detects a high preamble deviation. Signal strength is a notoriously bad indicator of distance, but because signal strength drops with the square of distance, a distant adversary will need to transmit a high-power signal for the single-antenna device to receive it with the same strength as a signal from a legitimate device located a few centimeters away. To prevent the distant adversary from tricking the single-antenna device into believing malformed preambles are legitimate signals from a nearby device, the single-antenna device can measure the signal strength of frames with high preamble deviations and reject frames with a signal strength below a threshold. We see next that even with a high-gain antenna, and an unlawfully high transmit power, an adversary’s signal strength will be well below the signal strength of a legitimate device located nearby, making it possible for the single-antenna device to reject the weaker signals.

Assuming the adversary is in the far-field region, then signal strength will attenuate as it travels from transmitter to receiver as predicted by the well-known Friis equation [33]

$$P_r = P_t + G_t + G_r + 20 \log \left(\frac{\lambda}{4\pi d} \right)^2 \quad (12)$$

where P_r is the power received in dBm, P_t is the power at the surface of the transmitting antenna in dBm, G_t and G_r are the gains of the transmitting and receiving antennas, λ is the frequency of the signal, and d is the distance between the transmitting and receiving antennas.

Using Equation (12), we see that a signal from a legitimate device with $P_t = 27$ dBm, $G_t = 3$ dBi, and located 5 cm away would be received on a single-antenna device with a signal strength of slightly over 19 dBm, assuming $G_r = 3$ dBi. To estimate the amount of power

a distant adversary could need to transmit to achieve the same signal strength at the single-antenna device, we can rewrite Equation (12) as

$$P_t = P_r - G_t - G_r - 20\log\left(\frac{\lambda}{4\pi d}\right)^2. \quad (13)$$

To match the signal strength of a nearby device, Equation (13) suggests that an adversary with a standard 3 dBi antenna located 1 m away would need to transmit a signal at a whopping 53 dBm – well above the capabilities of SDRs. The Ettus Research UBX-40 daughterboard for the N210 SDR, for example, transmits at roughly 20 dBm in the Wi-Fi frequency range. The popular HackRF One [15] tops out at 15 dBm in the Wi-Fi frequency bands. In fact, in the United States the Federal Communications Commission sets the legal limit for transmitters in the 2.4 GHz band at 30 dBm [12]. Nonetheless, we assume that an adversary does not stay within the FCC’s limits and may have a more powerful transmitter. We acknowledge that there is theoretically no limit to how much power the adversary can transmit, but we assume a realistic adversary has some practical bounds on its transmitting capability. For discussion purposes, we consider an adversary capable of transmitting at 36 dBm – four times the limit set by the FCC and more than 32 times the power of the HackRF One. Furthermore, we assume the adversary may use a high-gain antenna to increase the strength of the signal received by the single-antenna device.

Commercially available high-gain Wi-Fi antennas are normally relatively large panel or parabolic dish antennas that provide up to approximately 20 dBi gain. For example, the Altelix 2.4 GHz parabolic dish antenna provides 15 dBi gain and is nearly one-half meter long [2]. While it is possible this antenna, or one like it, could be concealed inside of furniture or possibly behind a low signal-attenuating wall, in many cases a user would notice the presence of the one-half meter long antenna if it is within two meters from the single-antenna device. Additionally, the single-antenna device may be mobile, making it difficult to preposition a high-gain antenna such that it is focused on the mobile device while data is transferred.

Using Equation (13), and assuming $G_r = 3$ dBi on the single-antenna device, in Table 1 we show the amount of power in dBm that an adversary would need to transmit to match the signal strength of a legitimate device transmitting from a few centimeters away. We assume the legitimate device transmits at 27 dBm (e.g., half the FCC limit) from 5 cm away. We assume the adversary is located one to three meters away and has a high gain



Figure 14: Altelix high-gain antenna. This antenna provides 15 dBm gain, but is nearly one-half meter wide.

G_t	Adversary distance (cm)				
	100	150	200	250	300
0	56	60	62	64	66
3	53	57	59	61	63
6	50	54	56	58	60
9	47	51	53	55	57
12	44	48	50	52	54
15	41	45	47	49	51
18	38	42	44	46	48
21	35	39	41	43	45
24	32	36	38	40	42

Table 1: Required adversary transmit power. This table shows the amount of power in dBm that an adversary located a given distance from a single-antenna device using an antenna with gain G_t would need to transmit to have a signal arrive at the single-antenna device with the same signal strength as a legitimate nearby device. Highlighted cells indicate configurations where the adversary would be able to match or exceed the legitimate signal if transmitting at four times the legal limit set by the FCC or below. We see that even with high transmit power and high-gain antenna, in most cases an adversary cannot achieve enough signal strength.

antenna with G_t ranging up to 24 dBi. Highlighted cells indicate configurations where the adversary would be able to match or exceed the legitimate signal if transmitting at four times the legal limit set by the FCC (36 dBm) or below. We see that in most cases, even with the high transmit power and high-gain antenna, the adversary cannot achieve a higher signal strength than a legitimate nearby device transmitting at one-half the maximum level set by the FCC.

6.3 Raising the bar for an adversary

In most practical cases the single-antenna device can use signal strength to determine if a frame with a high preamble deviation was sent by a distant adversary, even if that adversary transmits at unlawfully high power and uses a high-gain antenna. We consider our methods as a way of “raising the bar” that an adversary must overcome. It could be the case that the adversary has an extraordinarily high transmit power, an extremely high-gain antenna, and is within three meters of the single-antenna device. In those cases our method could fail, but our method raises the bar well above the capabilities of off-the-shelf hardware.

6.4 Proximity alone is not sufficient

In many scenarios proximity is necessary, but not sufficient, as an indicator of trust. Indeed, in most cases the user may not want their devices to pair with other devices that get physically close. For example, in a crowded subway people may be packed together tightly. Any devices they wear or carry may then come into unintended proximity with other devices. In those use cases, where devices may encounter untrusted devices, we anticipate that a system using SNAP’s technique would require the user to take an action such as intentionally pushing a button to initiate the proximity detection process, rather than blindly trusting nearby devices. Proximity detection used in conjunction with user intent could help prevent distant adversaries from tricking legitimate devices into accepting malicious frames.

7 RELATED WORK

Exploiting the repeating nature of Wi-Fi OFDM preambles and near-field effects has not been explored in the literature. Work thus far has primarily focused on covertly embedding a small amount of data into the Wi-Fi frame and on device fingerprinting based on PHY layer attributes. Both of these approaches are substantially different from SNAP’s methods and we examine each of them in this section. Additionally, because Near Field Communications (NFC) also uses the near field of a transmitting antenna, we review that approach as well.

7.1 Embedding covert information in Wi-Fi frames

Classen et al. examined covert channels possible within standard Wi-Fi frames [9] to covertly embed information. In particular they analyzed methods to use the Short Training Field with Phase Shift Keying (STF PSK), Carrier Frequency Offset with Frequency Shift Keying (CFO

FSK), using additional subcarriers while still conforming to Wi-Fi standards, and replacing portions of the Cyclic Prefix with covert data. In each of these cases the idea is to embed information in the Wi-Fi frame in a way that a standard receiver would not notice. While these techniques involve the PHY layer and often use the preamble, they do not enable proximity detection.

In related work Rahbari and Krunz proposed a technique they call *P-modulation* to modulate the STF in a standards-compliant manner to include up to eight user-chosen bits in a 20 MHz wide Wi-Fi frame [29]. These bits can be used to inform other devices of the transmitter’s status, possibly eliminating the need for additional control frames. This technique, however, is different from our technique in that we use the repeating nature of the Long Training Field to establish proximity, not to use it to include a small number of indicator bits.

7.2 Device fingerprinting

There have been numerous studies on fingerprinting wireless devices. Many of these approaches focus on PHY layer imperfections resulting from manufacturing such as clock skews [3] and can determine subtle radiometric differences between devices. These differences can often lead to identification of device type, and can sometimes also lead to specific device identification [26, 36].

While identifying specific devices based on their radio characteristics can be useful in some use cases, these techniques do not provide an indication of trust. It could be the case that a single-antenna device is able to distinguish between transmitters based on the (presumably non-spoofed) radiometric signatures of signals it receives, but the single-antenna device does not know if the transmitter is a legitimate nearby device or a distant adversary – it simply knows transmissions came from the same device.

Fingerprinting and our single-antenna proximity detection technique could, however, work together. One solution would be to use our proximity detection technique to know when a user brings a new device near a single-antenna device. The idea is that proximity determines the trustworthiness of the new device (whether the new device is compromised in some manner is out of scope for this work). In addition to detecting proximity, the single-antenna device could record radiometric details about the signal it receives from the new device. If the user later separates the devices, the single-antenna device could examine the radiometric details of the signal to identify data sent by the new device.

The goal of re-identifying a known device could also be accomplished, however, if the single-antenna device

shares a unique key with the new device when they are in close proximity. In future communications the devices could use the key to encrypt or sign messages, even if the two devices are now far apart.

7.3 Near Field Communications

Near Field Communications (NFC) is a short-range communication technique, not a proximity detection method, that uses two loop antennas to transfer data between devices located roughly 10 cm apart. These systems transmit in an ISM (Industrial, Scientific, Medical) frequency band at 13.56 MHz [10] and rely on magnetic induction to transfer data [31]. Induction occurs when a receiver is in the transmitter’s near-field region and does not occur in the transmitter’s far-field region. This near-field requirement implies that data can be securely transferred at short ranges (e.g., within R_1 from Equation (5)), without fear of interception by more distant eavesdroppers. Recent work, however, has shown that NFC communications can be read up to 2.4 m from the transmitter [37], making NFC a questionable choice for secure communications.

Despite the questionable security aspects of NFC, it is becoming increasingly popular on cell phones. Aside from mobile payment systems, this increased popularity, however, has not translated into wide-spread usage. Android Beam, for example, is a feature of the Android operation system initially deployed with version 4.0, Ice Cream Sandwich, that can bootstrap a Bluetooth connection by exchanging keys over NFC [14]. Android Beam initially generated a great deal of excitement but has seen limited adoption. In fact ComputerWorld noted that “... despite the admirable marketing effort, Android Beam never quite worked particularly well ...” [20].

Additionally, because NFC requires specialized radios and antennas, it is not commonly found on IoT-type devices. SNAP, however, accomplishes some of the same tasks NFC was designed to accomplish, but does not require specialized radios or antennas. Instead, SNAP uses the in-band Wi-Fi radio commonly found on such devices.

8 CONCLUSION

Because radio waves are invisible, a user may believe that a wireless device is communicating with a nearby device when in fact the user’s device is communicating with a distant adversary. Knowing that a transmitter is in close physical proximity could eliminate that distant adversary problem. While other researchers have suggested methods for multiple-antenna devices to estimate the location of a transmitter, these techniques are not

available to single-antenna IoT-type devices. In this paper we show that a single-antenna device following the standard Wi-Fi protocol can reliably determine when it is in close proximity to a transmitting device by leveraging the repeating nature of Wi-Fi’s preamble and the physical characteristics of signals in the transmitter’s near-field region.

Our experiments suggest that mismatches in the preamble caused by the near field of a transmitting antenna can allow a single-antenna device to reliably detect proximity at distances under 9 cm while never falsely declaring proximity at ranges greater than 14 cm. This proximity determination can then be used as one basis for trust when new IoT devices are encountered and can ensure that a device is communicating with an intended nearby device, not the distant adversary.

We focused our testing on dipole and micropatch antennas because they are the most common antennas used in consumer devices, but there are a myriad of other types of antennas we have not yet tested. In future work we will examine the near-field effects on these other types of antennas. We believe, however, that our work here is a starting point and opens an important new area of research that warrants further investigation.

9 ACKNOWLEDGEMENTS

This research program is supported by National Science Foundation award number CNS-1329686. The views and conclusions in this document are those of the authors and may not necessarily represent the official policies of NSF.

REFERENCES

- [1] NFC Forum, <http://nfc-forum.org>, visited 7/28/2018.
- [2] Altelix LLC. 2.4 GHz 15 dBi Wi-Fi Parabolic Antenna, <http://www.altelix.com>, visited 3/22/2018.
- [3] Chrisil Arackaparambil, Sergey Bratus, Anna Shubina, and David Kotz. On the reliability of wireless fingerprinting using clock skews. In *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*, March 2010.
- [4] Constantine A. Balanis. *Antenna Theory: Analysis and Design*. Wiley, third edition, 2005.
- [5] Philippe Biondi. Scapy, <http://www.secdev.org/projects/scapy>, visited 4/15/2018.
- [6] Bastian Bloessl, Michele Segata, Christoph Sommer, and Falko Dressler. An IEEE 802.11a/g/p OFDM receiver for GNU Radio. In *Proceedings of the Workshop on Software Radio Implementation Forum (SRIF)*, pages 9–16. ACM, 2013.
- [7] Bastian Bloessl, Christoph Sommer, and Falko Dressler. Power matters: Automatic Gain Control for a Software Defined Radio IEEE 802.11 a/g/p receiver. In *IEEE International Conference on Computer Communications Workshop (INFOCOM)*, pages 25–26. IEEE, 2015.
- [8] Liang Cai, Kai Zeng, Hao Chen, and Prasant Mohapatra. Good neighbor: Ad hoc pairing of nearby wireless devices by

- multiple antennas. In *Proceedings of Network and Distributed Systems Security Symposium (NDSS)*, 2011.
- [9] Jiska Classen, Matthias Schulz, and Matthias Hollick. Practical covert channels for WiFi systems. In *Conference on Communications and Network Security (CNS)*, pages 209–217. IEEE, September 2015.
- [10] Vedat Coskun, Busra Ozdenizci, and Kerem Ok. A survey on Near Field Communication (NFC) technology. *Wireless Personal Communications*, 71(3):2259–2294, 2013.
- [11] Ettus Research. USRP N210 Software Defined Radio, <https://www.ettus.com>, visited 3/18/2018.
- [12] Federal Communications Commission. Section 15.247 Operation within the bands 902-928 MHz, 2400-2485.3 MHz, and 5725-5850 MHz, <https://www.fcc.gov/fdsys/pkg/CFR-2013-title47-vol1>, visited 3/22/2018.
- [13] Mikhail Fomichev, Flor Alvarez, Daniel Steinmetzer, Paul Gardner-Stephen, and Matthias Hollick. Survey and systematization of secure device pairing. *IEEE Communications Surveys & Tutorials*, 2017.
- [14] Google. Share content by NFC with Android Beam, <https://support.google.com/nexus/answer/2781895>, visited 3/22/2018.
- [15] Great Scott Gadgets. HackRF One, <https://greatscottgadgets.com>, visited 3/22/2018.
- [16] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. Tool release: Gathering 802.11n traces with Channel State Information. *SIGCOMM Computer Communication Review*, 41(1):53, January 2011.
- [17] Institute of Electrical and Electronics Engineers. 802.11n standard, <http://standards.ieee.org>, visited 7/20/2018.
- [18] Intel. Intel Ultimate N Wi-Fi Link 5300: Product Brief, <http://www.intel.com>, visited 4/15/2018.
- [19] R. Jin, L. Shi, K. Zeng, A. Pande, and P. Mohapatra. Mag-pairing: Pairing smartphones in close proximity using magnetometers. *IEEE Transactions on Information Forensics and Security*, 11(6):1306–1320, June 2016.
- [20] JR Raphael. Android nostalgia: 13 once-trumpeted features that quietly faded away, <https://www.computerworld.com/article/3239864>, visited 3/22/2018.
- [21] Andre Kalamandeen, Adin Scannell, Eyal de Lara, Anmol N. Sheth, and Anthony LaMarca. Ensemble: cooperative proximity-based authentication. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 331–344. ACM, June 2010.
- [22] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. SpotFi: Decimeter level localization using WiFi. In *Proceedings of the Conference on Special Interest Group on Data Communication (SIGCOMM)*, pages 269–282. ACM, 2015.
- [23] Xiaohui Liang, Tianlong Yun, Ronald Peterson, and David Kotz. LightTouch: Securely connecting wearables to ambient displays with user intent. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–9, May 2017.
- [24] Suhas Mathur, Rob Miller, Alexander Varshavsky, Wade Trappe, and Narayan Mandayam. ProxiMate: Proximity-based Secure Pairing using Ambient Wireless Signals. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 211–224. ACM, June 2011.
- [25] Rene Mayrhofer and Hans Gellersen. Shake well before use: Intuitive and secure pairing of mobile devices. *IEEE Transactions on Mobile Computing*, 8(6):792–806, June 2009.
- [26] Christoph Neumann, Olivier Heen, and Stephane Onno. An empirical study of passive 802.11 device fingerprinting. In *International Conference on Distributed Computing Systems Workshops*, pages 593–602. IEEE, June 2012.
- [27] Panda Wireless. Panda Ultra Wireless N USB Wi-Fi adapter, <http://www.pandawireless.com>, visited 4/15/2018.
- [28] Timothy J. Pierson, Xiaohui Liang, Ronald Peterson, and David Kotz. Wanda: securely introducing mobile devices. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–9. IEEE, April 2016.
- [29] Hanif Rahbari and Marwan Krunz. Exploiting frame preamble waveforms to support new physical-layer functions in OFDM-based 802.11 systems. *IEEE Transactions on Wireless Communications*, 16(6):3775–3786, June 2017.
- [30] R. Rawassizadeh, T. J. Pierson, R. Peterson, and D. Kotz. NoCloud: Exploring network disconnection through on-device data analysis. *IEEE Pervasive Computing*, 17(1):64–74, 2018.
- [31] K. Ren, Q. Wang, D. Ma, and X. Jia. Securing emerging short range wireless communications: the state of the art. *IEEE Wireless Communications*, 21(6):153–159, December 2014.
- [32] The GNU Radio Foundation. GNURadio, <https://www.gnuradio.org>, visited 3/22/2018.
- [33] David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge University Press, 2005.
- [34] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single WiFi access point. In *Symposium on Networked Systems Design and Implementation (NSDI)*, volume 16, pages 165–178. USENIX, 2016.
- [35] Ju Wang, Jie Xiong, Hongbo Jiang, Xiaojiang Chen, and Dingyi Fang. D-watch: Embracing bad multipaths for device-free localization with COTS RFID devices. *IEEE Transactions on Networking (TON)*, 25(6):3559–3572, December 2017.
- [36] Qiang Xu, Rong Zheng, Walid Saad, and Zhu Han. Device fingerprinting in wireless networks: Challenges and opportunities. *IEEE Communications Surveys & Tutorials*, 18(1):94–104, January 2016.
- [37] Ruogu Zhou and Guoliang Xing. nShield: A noninvasive NFC security system for mobile devices. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 95–108. ACM, 2014.
- [38] Hongzi Zhu, Jingmei Hu, Shan Chang, and Li Lu. ShakeIn: secure user authentication of smartphones with single-handed shakes. *IEEE Transactions on Mobile Computing*, (10):2901–2912, 2017.