Practical Evaluation of Passive COTS Eavesdropping in 802.11b/n/ac WLAN

Daniele Antonioli[1] 0000-0002-9342-3920, Sandra Siby[2] 0000-0002-9481-0826, and Nils Ole Tippenhauer[1] 0000-0001-8424-2602

[1] Singapore University of Technology and Design (SUTD), Singapore
{daniele_antonioli,nils_tippenhauer}@sutd.edu.sg
[2] Ecole polytechnique federale de Lausanne (EPFL), Switzerland
sandra.siby@epfl.ch

Abstract. In this work, we compare the performance of a passive eavesdropper in 802.11b/n/ac WLAN networks. In particular, we investigate the downlink of 802.11 networks in infrastructure mode (e. g. from an access point to a terminal) using Commercial-Of-The-Shelf (COTS) devices. Recent 802.11n/ac amendments introduced several physical and link layer features, such as MIMO, spatial diversity, and frame aggregation, to increase the throughput and the capacity of the channel. Several information theoretical studies state that some of those 802.11n/ac features (e. g. beamforming) should provide a degradation of performance for a passive eavesdropper. However, the real impact of those features has not yet been analyzed in a practical context and experimentally evaluated. We present a theoretical discussion and a statistical analysis (using path loss models) to estimate the effects of such features on a passive eavesdropper in 802.11n/ac, using 802.11b as a baseline. We use Signal-to-Noise-Ratio (SNR) and Packet-Error-Rate (PER) as our main metrics. We compute lower and upper bounds for the expected SNR difference between 802.11b and 802.11n/ac using high-level wireless channel characteristics. We show that the PER in 802.11n/ac increases up to 98% (compared to 802.11b) at a distance of 20 meters between the sender and the eavesdropper. To obtain a PER of 0.5 in 802.11n/ac, the attacker's maximal distance is reduced by up to 129.5 m compared to 802.11b. We perform an extensive set of experiments, using COTS devices in an indoor office environment, to verify our theoretical estimations. The experimental results validate our predicted effects and show that every amendment add extra resiliency against passive COTS eavesdropping.

Keywords: WLAN, 802.11, Eavesdropping, MIMO, Beamforming

1 Introduction

In the last decade, wireless network communication has grown tremendously mainly due to standards such as UMTS (3G) and LTE (4G) for cellular networks and IEEE 802.11 (WLAN) for wireless networks. Cisco estimated that in 2017, 68% of all Internet traffic will be generated by wireless devices [5]. As a result,

it can be expected that a majority of sensitive communication services, such as mobile banking and online payments will involve wireless networks. Indeed, it is paramount to secure the broadcast wireless channel against eavesdroppers to protect the confidentiality and integrity of the information.

In this work, we present a theoretical discussion, a numerical analysis (using path loss models), and a practical evaluation of passive eavesdropping attacks targeting several 802.11 (WLAN) networks. Recent 802.11n/ac amendments introduced interesting physical layer and link layer features such as Multiple-Input-Multiple-Output (MIMO), spatial diversity (e. g. CSD, TxBF, STBC) , spatial multiplexing (e. g. MU-TxBF), dual-band antennas[3] and frame aggregation [14]. It is believed that some of those features, that were developed mainly to increase the robustness and throughput of the channel might also *degrade* the performance of a passive eavesdropper. We would like to investigate this claim and experimentally measure whether this degradation happens or not in practice in a simple but yet realistic scenario (e. g. eavesdropping WLAN networks with COTS devices).

Several theoretical discussions have already been presented about passive and active eavesdropping in the wireless channel. The seminal work by Wyner [31] started the wiretap channel research track that has been extended to Gaussian [16], fading [10], and MIMO [20] channels. This set of papers studies asymptotic conditions that very rarely happen in practice. Recently, special attention was given to MIMO and beamforming as a defense mechanism against passive eavesdropping [25, 32, 22]. However, those works do not focus on 802.11 and they consider only a subset of the 802.11 features. There are also some alternative techniques already proposed against passive eavesdropping including multi-user cooperative diversity and the use of artificial noise [8, 33, 19]. However, those techniques are neither listed in any 802.11 standards nor implemented in any COTS device.

In this paper, we investigate the disadvantages that a passive eavesdropper has to face when attacking the downlink of an 802.11n/ac (MIMO) network versus an 802.11b (SISO) network. We focus on 802.11 networks in infrastructure mode (e. g. an access point connecting several laptops to the Internet) that use Commercial-Of-The-Shelf (COTS) devices. In particular, we compare *three* of the most widely used 802.11 amendments: b, n, and ac. We look at the downlink (e. g. traffic from the access point to the terminals) because it is the link that supports most of the advanced features of 802.11n/ac (e. g. spatial diversity and spatial multiplexing). We use 802.11b as a baseline. Our attacker model choice is explained in detail in Section 3.1, and a brief discussion about a stronger attacker model is presented in Section 4.5.

In our theoretical discussion, we estimate lower and upper bounds for the expected Signal-to-Noise-Ratio (SNR) disadvantage of an eavesdropper in 802.11n and ac compared to 802.11b. We numerically derive the expected Packet-Error-Rate (PER) of the intended receiver and the eavesdropper with respect to their distances to the sender. Finally, we present an 802.11b/n/ac downlink empirical

[3] In this work we always use the word *antennas* rather than *antennae*.

evaluation using COTS devices. After the experiments, we are able to confirm that in 802.11n/ac networks, the PER of the eavesdropper increases with respect to her distance to the sender, given a minimum distance between the attacker and the intended receiver.

We summarize our contributions as follows:

- We derive the theoretically expected eavesdropper's SNR disadvantage (in dB), for attacks using COTS radios, in 802.11b/n/ac downlinks.
- We discuss how the theoretical SNR disadvantage translates to practical constraints (e. g. reduced range, higher PER) for the attacker.
- We perform a series of experiments to validate that the expected disadvantage is experienced in practice and that its effects were correctly predicted.

The structure of this work is as follows: in Section 2, we provide the basic background about a fading wireless channel, the 802.11 standard, and three wireless communications metrics. In Section 3, we present the system and attacker model, we compare passive eavesdropping on the downlink of 802.11b (SISO) and 802.11n/ac (MISO) channels and we estimate the SNR and PER disadvantages for a passive eavesdropper in 802.11n/ac. In Section 4, we present our results from a series of practical eavesdropping experiments that validate our predicted disadvantages. We summarize related work in Section 5, and conclude our paper in Section 6.

2 Background

We now provide a summary of the important concepts used in this work: the fading wireless channel, the 802.11b/n/ac amendments, and three wireless communication metrics (SNR, BER, and PER). For additional details, we refer to influential books such as [23, 9].

2.1 The Fading Wireless Channel

The progression of wireless communication systems evolved around two main metrics: *robustness* and *throughput*. Those metrics are severely influenced by channel fading. Fading can be described as a random process affecting the quality of the transmitted wireless signal, by means of attenuation and distortion over time and frequency. There are three additive phenomena contributing to fading: path loss, shadowing, and multipath.

Path loss is a large-scale fading event due to the propagation nature of the electromagnetic waves (that are carrying the useful signal). There are different path loss models according to the system parameters and the channel environment. For example, in the Free Space Path Loss (FSPL) model the transmitted power decays quadratically with the distance from the transmitter to the receiver. Shadowing is another large-scale fading event due to the presence of obstacles between the transmitter and the receiver. There are different ways to model shadowing such as using a log-normal random variable. Multipath is a small-scale fading phenomenon that takes into account constrictive and/or

Table 1: *Relevant 802.11b/n/ac physical layer specifications. f_c is the carrier frequency, λ is the wavelength, s_{dr} is the theoretical maximum throughput of the channel, n_S is the number of maximum independent data streams, TxBF indicates support for single-user (SU) or multi-user (MU) transmit-beamforming, d_i and d_o are the expected ranges for indoor and outdoor communications.*

	Technology	Modulation	f_c [GHz]	λ [cm]	s_{dr} [Mbit/s]	n_S	TxBF	d_i	d_o
b	SISO	DSSS	2.4	12.5	11	N/A	N/A	35	140
n	SU-MIMO	OFDM	2.4, 5	12.5, 6	135	4	SU	70	250
ac	MU-MIMO	OFDM	5	6	780	8	MU	35	N/A

destructive interference at the receiver between direct, reflected and scattered electromagnetic waves.

There are two well-known fading models that take into account all three fading aspects: *Rayleigh fading* for non-line-of-sight (NLOS) environments, and *Rician fading* for line-of-sight (LOS) environments. In both cases, each channel coefficient h is modeled with a complex random number. Each channel coefficient is providing random attenuation (change in amplitude) and distortion (change in phase). In the Rayleigh fading model, the real and imaginary parts of h are modeled with independent identically-distributed (IID) Gaussian random variables with 0 mean and equal variances and the amplitude of h is Rayleigh distributed. In the (more generic) Rician fading model, the amplitude of h is Rice distributed.

2.2 IEEE 802.11 Standard (WLAN)

802.11 is a family of IEEE standards that regulates wireless local area networks (WLAN) [7]. The standards define the physical layer (PHY), and the link layer specifications. An example of physical layer specification is the modulation and coding scheme (MCS) table that lists the supported modulation types, spatial streams, coding rates, bandwidths and data rates of a given PHY. An example of link layer specification is the medium access control (MAC) protocol that governs how the nodes share the wireless medium.

Table 1 lists some relevant physical layer specifications for 802.11b, n, and ac [14]. 802.11b uses Single-Input-Single-Output (SISO) scheme with direct-sequence spread spectrum (DSSS) modulation techniques. In contrast, 802.11n and 802.11ac are Multiple-Input-Multiple-Output (MIMO) schemes, based on orthogonal frequency division multiplexing (OFDM) modulation techniques. Single user MIMO is supported by 802.11n, while 802.11ac supports multi-user MIMO. The major advantage in terms of throughput and robustness of the channel from b to n/ac is given by the usage of multiple radios and antennas that allows transmitting different independent symbol at the same time (spatial multiplexing) or the same symbol on multiple antennas at the same time (spatial diversity). In particular, 802.11n/ac support transmit-beamforming (TxBF) at

the downlink for single user (n) and multiple users (ac). By using TxBF, an access point can optimize the transmission of the symbols to a device located in a particular region of space, given an estimate of the condition of the downlink channel. For a more detailed comparison among the three 802.11 amendments please refer to [13, 21].

2.3 Wireless Communications Metrics

The following list summarizes the three wireless communication metrics used in our paper.

- The *Signal-to-Noise-Ratio (SNR)* is the ratio between the power of the useful signal denoted with P, and the noise power σ^2. It is typically expressed in decibel dB, and it convertible from logarithmic to linear scale using: $10 \log_{10} \mathrm{SNR} = \mathrm{SNR}_{\mathrm{dB}}$.
- The *Bit-Error-Rate (BER)* is the expected probability of error while decoding 1-bit at the receiver. The BER is not an exact quantity. It can be modeled and estimated according to different factors such as the modulation/coding schemes, the fading model and the number of antennas. Typically, 10^{-6} is considered a reasonable BER value, i. e. 1-bit error per Mbit.
- The *Packet-Error-Rate (PER)* is directly proportional to the BER, and it is computed as: $\mathrm{PER} = 1 - (1 - \mathrm{BER})^N$, where N is the average packet size in bits. In this work, we assume that one or more bit errors in a packet will lead to an incorrect link layer checksum. Packets with an incorrect checksum are not acknowledged by the (legitimate) receiver, and retransmitted by the sender.

3 Passive 802.11 Downlink Eavesdropping

We start this section introducing the system and attacker models. Then we present a theoretical discussion and a numerical analysis (based on 802.11 path loss models) to estimate the SNR and PER disadvantages of a passive eavesdropper in an 802.11n/ac (MISO) downlink, compared to an 802.11b (SISO) downlink.

3.1 System and Attacker Model

Our system model focuses on the *downlink* of indoor 802.11b/n/ac networks in infrastructure mode (e. g. access point that communicates with several wireless terminals), using Commercial-Of-The-Shelf (COTS) devices. The access point is equipped with multiple antennas. The intended receiver and the attacker are equipped either with a single or multiple antennas according to the scenario. We are looking at the ratio of packets that the attacker successfully eavesdrop on the physical layer and we are agnostic to any encryption scheme used at the link layer or above. Attacks on those schemes are possible, but out of the scope of this work [3, 26]).

The attacker is assumed to be *equipotent* to the intended receiver in terms of hardware and software capabilities. In particular, both use COTS devices, with a similar chipset, driver, feature set, and maximum throughput. With COTS devices we refer to wireless radios either built into laptops, smartphones, access point or USB dongles. We do not consider an attacker equipped with a software-defined-radio (SDR) or similar devices. We focus on a *passive* eavesdropper who wants to capture the downlink packets in real-time using her wireless card in monitor mode. We are not considering an attacker who is recording and post-processing the traffic offline. We assume an attacker that is static and we evaluate her eavesdropping performance at different distances from the sender. If the sender is using beamforming, we assume that the attacker is outside the beamforming region.

The effectiveness of the attacker is assessed from the Signal-to-Noise-Ratio (SNR) and the Packet-Error-Rate (PER) at her receiver. We chose PER as metric because we are mainly interested in the relative performance of eavesdropping on 802.11b vs. 802.11n/ac. As our passive attacker is unable to request retransmissions, the only chance to recover from bit errors would be to find the offending bit(s) and correct it using a checksum (possibly by brute force). We note that such corrections are expected to have significant cost for increasing number of flipped bits, and that the number of flipped bits is expected to quickly increase with distance. We plan to further investigate this in future work.

Without loss of generality and to simplify our discussion, we are considering an attacker focused on eavesdropping the downlink channel of one pair of transmitter and intended receiver. We understand that our attacker model is relatively weak (e. g. a single attacker, no SDR), however, given the lack of related experimental work and the number of involved moving parts, we decided to start with a simple scenario that is easy to evaluate (e. g. worst-case scenario for the passive eavesdropper). We look forward to investigate more complex attacker models in future work.

Finally, we present the notation used in our paper. The access point is referred as Alice (the transmitter), the victim as Bob (the intended receiver), and the attacker as Eve (passive eavesdropper). We will use A, B, and E subscripts to identify quantities related to Alice, Bob, and Eve respectively. We use x to denote Alice's transmitted symbol, h for complex channel coefficients, and n for the noise at a specific receiver. The relative distances between Alice, Bob, and Eve are written as: d_{AB}, d_{BE}, d_{AE}. Alice is equipped with L antennas and L radios.

3.2 SISO and MISO Channels Eavesdropping

In this section, we analyze and compare two different eavesdropping scenario: i) 802.11b SISO downlink, ii) 802.11n/ac MISO downlink. and we derive two essential conclusions about passive eavesdropping in SISO vs. MIMO 802.11 downlinks.

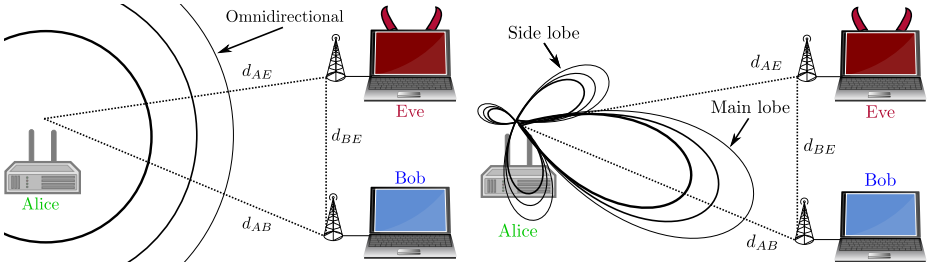

(a) *Omnidirectional radiation ($L = 1$). Eve's success depends on d_{AE}.*

(b) *Transmit-beamforming ($L > 1$). Eve's success depends also on d_{BE} and L.*

Fig. 1: *802.11b SISO (left) vs. 802.11 n/ac MISO (right) passive eavesdropping. Bob and Eve have one antenna. Dashed lines represent distances. Black circles and lobes represent omnidirectional and directional transmission ranges. Circles and lobes decreasing thickness represent the transmission power decay with respect to distance from the transmitter. Both channels are affected by random noise and fading.*

802.11b SISO downlink. Figure 1a shows Eve trying to intercept the communication from Alice to Bob in an 802.11b SISO network. We can represent the signals received by Eve and Bob as:

$$y_E = x \cdot h_E + n_E \tag{1}$$

$$y_B = x \cdot h_B + n_B \tag{2}$$

Intuitively, it is possible to represent Alice's two-dimensional transmission coverage with concentric circles. In free space, the greater is the distance from the transmitter the higher is the transmitted power decay. While one might assume that every receiver inside these circles will be "in range" and receive all transmissions by Alice, this is not the case in practice. If circles are shown around transmitters, their radius commonly refers to a distance in which the average received signal strength is above a certain threshold. However, due to random deep fading (mostly due to multipath), the instantaneous received power will constantly vary. In other words, it is possible to "miss transmissions" while being in the outer circle, or even receive transmissions just outside the outer circle. In this case, Eve's success rate depends on her distance to Alice (d_{AE}) regardless of her distance to Bob (d_{BE}), and random channel characteristics. The SISO wireless channel is providing some sort of resiliency against eavesdropping that an attacker can compensate with other means (eg: increase receiver sensitivity, use directional antenna).

802.11n/ac MISO downlink. Figure 1b shows Eve attempting to intercept the communication from Alice to Bob in an 802.11n/ac MISO network. Alice is equipped with L antennas and uses transmit-beamforming. In this scenario,

beamforming has been theoretically proven to provide resiliency against passive eavesdropping [12]. The received signals by Eve and Bob are as follows:

$$y_E = x \cdot g_E + n_E \tag{3}$$

$$y_B = x \cdot g_B + n_B \tag{4}$$

We can derive two benefits in terms of eavesdropping resiliency, one from g_B, and one from g_E. $\|g_B\|^2$ is defined as the *beamforming gain* and it is modeled by a Chi-squared random variable, with parameter $2L$ (being the sum of squared IID standard Gaussian random variables). Indeed, if $L = 2$ (Alice is using two antennas), then Bob's received signal will be the sum of two signals with independent fading paths. The correspondent beamforming gain is computed as:

$$\|g_B\|^2 = \|h_{B1}\|^2 + \|h_{B2}\|^2 \tag{5}$$

and this quantity is certainly greater (or equal) to $\|h_{B1}\|^2$ and $\|h_{B2}\|^2$. The net result is a better SNR at Bob's receiver with respect to the SISO case.

The second benefit arising from transmit-beamforming is encapsulated by g_E. Eve's ability to eavesdrop depends on two more factors with respect to the SISO case. Firstly, her distance from Bob (d_{BE}), and secondly the number of antennas used by Alice (L). This is a consequence of transmit-beamforming employed by Alice (the beamformer) towards Bob (the beamformee). Figure 1b shows Alice beamforming in the direction of Bob (e. g. inside the main lobe) while Eve is outside the main and the side lobes. This results in a smaller SNR at her receiver compared to the one of Figure 1a (given the same relative distances). Even if we decrease the distance between Eve and Alice, the disadvantage will still hold until Eve is outside the beamforming region. Furthermore, Eve's SNR will be inversely proportional to L because the more antennas are used by Alice to beamform, the more Alice can focus the beam towards a narrower but longer region in space [29].

3.3 Eavesdropper's Theoretical SNR Disadvantage in 802.11n/ac

In the previous section we argued that MISO beamforming from Alice to Bob will degrade Eve's eavesdropping performance according to d_{AE}, d_{BE}, and L. In this section, we will quantify the expected disadvantage of Eve in an 802.11n/ac network compared to an 802.11b network. We will estimate upper and lower bounds for the SNR at Eve's receiver with respect to L. We will provide numerical results for $L = 4$ to match the experimental setup of Section 4.1. We note that the bounds we are providing are not supposed to be *strict*—the actual SNR disadvantage will depend on many factors. Nevertheless, we compute the bounds based on the modeling assumptions to provide an intuition about the theoretically expected disadvantage.

Upper Bound. We start comparing high-level wireless channel characteristics of SISO and MISO channels. Table 2 lists the closed-form expressions for the

Table 2: *SNR and BER of 802.11b (SISO) and 802.11n/ac (MISO transmit-beamforming with L antenna) using BPSK modulation scheme.*

Metric	SISO	MISO Beamforming	
SNR	$\|h\|^2 \frac{P}{\sigma^2}$	$\|g\|^2 \frac{P}{\sigma^2}$	
BER	$\frac{1}{2}(1-\lambda)$	$\left(\frac{1-\lambda}{2}\right)^L \sum_{i=0}^{L-1} \binom{L+i-1}{i} \left(\frac{1+\lambda}{2}\right)^i$	$\lambda = \sqrt{\frac{\text{SNR}}{2+\text{SNR}}}$
DO	1	L	

SNR and the BER of SISO and MISO networks using BPSK modulation scheme. In general, we note that the number of antennas deployed by Alice (L) is playing a central role. If we fix the expected BER to 10^{-6}, then we can compute the minimum SNR for the SISO (57 dB) and the MISO case with $L = 4$ (16 dB). There is a notable difference in SNR of 41 dB between the SISO and the MISO cases. We use 41 dB as an upper bound for the SNR disadvantage of Eve with respect to Bob.

Lower Bound. For the lower bound of Eve's SNR disadvantage, we use a standard formula to compute the beamforming gain in a MISO channel where Alice is using Cyclic Delay Diversity (CDD) with L antennas [17]. In this case, the beamforming gain in dB can be computed as follows:

$$\|g\|^2 = 10\log_{10}(L) \quad dB \tag{6}$$

Assuming a COTS access point with 4 antennas and a single receiving antenna, Bob's beamforming gain is 6 dB. As Eve's COTS radio will not benefit from the beamforming gain (being outside the main lobe) Eve's SNR disadvantage lower bound is thus 6 dB with respect to Bob.

Summary. We estimate that an 802.11n/ac downlink that is using transmit-beamforming with four antennas provides an reduction in the SNR of a passive eavesdropper (outside the main lobe, using a COTS receiver) that is bounded *between 6 dB and 41 dB*. The reduction in SNR at Eve's receiver depends on a deterministic and measurable factors: d_{AE} (distance between Alice and Eve) and L (number of antennas used by the Alice). We note that Eve's SNR variation depends also on channel (Rayleigh) fading, however this factor is not considered in our discussion because it randomly affects both Bob and Eve, providing no deterministic disadvantage to Eve. Given this theoretically expected disadvantage, the question now is: "*How does the eavesdropper SNR disadvantage translate to practical constraints on 802.11 passive eavesdroppers?*"

3.4 Numerical Path Loss Analysis

In this section, we present a numerical analysis using three indoor path loss models for 802.11 networks. The models includes both the 2.4 and 5 GHz bands

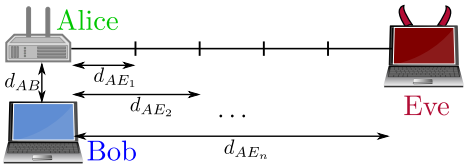

Fig. 2: *Setup used for our numerical analysis and for the experiments: Bob is at a fixed distance away from Alice, Alice is sending 802.11 traffic and Eve is passively eavesdropping from different (stationary) distances on a line perpendicular to Bob.*

and they are taken from [23]. We now describe their relevant parameters. In particular, d_{BP} is defined as the breakpoint distance between the transmitter and the receiver and it determines the cutoff span between LOS and NLOS channel condition. σ_{SF} represents the standard deviation in dB of the log-normal random variable that models the shadowing term of the path loss. s_{PL} represents the path loss slope before and after d_{BP}. Comma-separated values in the following list indicate values before and after the breakpoint distance:

- **Model B**: Residential (e. g. intra-room, room-to-room).
  - $d_{BP} = 5$ m
  - $\sigma_{SF} = 3$, 4 dB
  - $s_{PL} = 2$, 3.5
- **Model D**: Office (e. g. large conference room, sea of cubes).
  - $d_{BP} = 10$ m
  - $\sigma_{SF} = 3$, 5 dB
  - $s_{PL} = 2$, 3.5
- **Model E**: Large office (e. g. multi-storey building).
  - $d_{BP} = 20$ m
  - $\sigma_{SF} = 3$, 6 dB
  - $s_{PL} = 2$, 3.5

Figure 2 shows the setup used for our numerical analysis and for the experiments. Bob is placed at a fixed distance away from Alice, Eve is placed at different (stationary) distances d_i from Alice, and Alice is constantly sending traffic to Bob. In a two-dimensional plane, Bob and Eve distance vectors are perpendicular to avoid Eve being in the main lobe when Alice is using transmit-beamforming. We note that in an indoor environment multipath is playing a major role than visual of RF line-of-sight conditions that is why we decided to keep altitude and angle constant and vary only the distance between Alice and Eve [6].

The path loss model function L_P is constructed considering the sum of a free-space loss component (L_{FS}), a shadowing log-normal component due to obstacles (S_F), and a post breakpoint distance component. All terms vary according to the distance d between the transmitter and the receiver. We used the following equations from [23]:

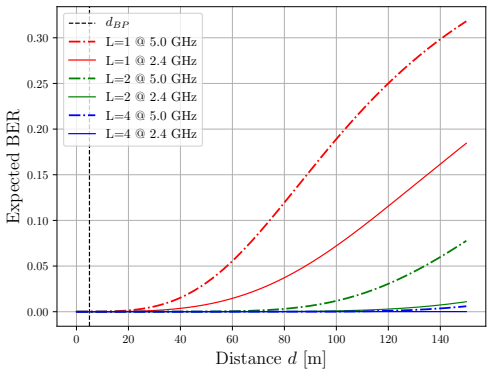

Fig. 3: *802.11n Model B (Residential) expected BER estimation using BPSK. Red lines represent Eve. Green and Blue lines represent Bob when L=2 and L=4.*

$$L_P(d) = \begin{cases} L_{FS}(d) + S_F(d) & \text{if } d \leq d_{BP} \\ L_{FS}(d_{BP}) + S_F(d) + 35\log_{10}\left(\frac{d}{d_{BP}}\right) & \text{otherwise} \end{cases} \tag{7}$$

$$L_{FS}(d) = 20\log_{10}(d) + 20\log_{10}(f) - 147.5 \tag{8}$$

$$S_F(d) = \frac{1}{\sqrt{2\pi}\sigma_{SF}} \exp\left(-\frac{d^2}{2\sigma_{SF}^2}\right) \tag{9}$$

Figure 3 and Figure 4 shows the predicted BER and PER for model B (Residential) vs. distance between the transmitter and the receiver. Solid lines represent results for 2.4 GHz and dash-dotted lines represent results for 5.0 GHz. Red lines represent Eve's expected BER and PER. The other lines represent Bob's expected BER and PER when Alice is using transmit beamforming with two (green lines) and four (blue lines) antennas. If we focus on the solid lines of Figure 4, then we note that a distance between Alice and Eve d_{AE} of 12.5 m is sufficient to drop Eve's expected PER from 0 to 0.5 (50% chance of decoding). Furthermore a d_{AE} of 20 m is sufficient to increase Eve's PER to 0.98 (0.2% chance of decoding). On the other hand, a d_{AB} of 142 m is required to experience a PER of 0.5 at Bob's receiver when Alice is using four antennas (L=4).

3.5 Eavesdropping Analysis Summary

In this section, we argued that in 802.11n/ac downlink a passive eavesdropper (Eve) using a COTS radio will have a disadvantage in terms of SNR compared to an eavesdropper in an 802.11b downlink. This disadvantage is due to different

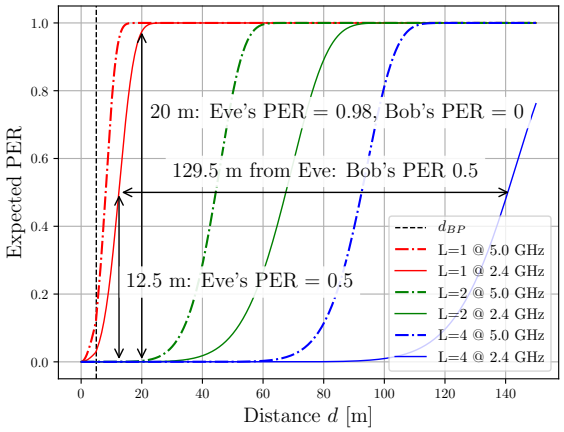

Fig. 4: *802.11n Model B (Residential) expected PER estimation using BPSK. Red lines represent Eve. Green and Blue lines represent Bob when L=2 and L=4.*

features provided by recent 802.11n/ac such as MIMO, and spatial diversity. This disadvantage can be expressed in an SNR decrease at the eavesdropper receiver of 6-41 dB (depending on the chosen scenario). We also express this disadvantage in terms of the distance that the eavesdropper has to be closer to the sender to achieve the same PER as a legitimate receiver, which can reach up to 129.5 m. In contrast, there is no such distance disadvantage for the eavesdropper in 802.11b. Furthermore, we can express the disadvantage in terms of PER at the eavesdropper receiver compared to her distance from the transmitter (d_{AE}). For example, if d_{AE} is 12.5 m, then the PER of Eve is increased to 50%, and if d_{AE} is 20 m, then the PER of Eve is increased to 98%.

4 Experimental Validation

In this section, we present an experimental evaluation of COTS passive eavesdropping in 802.11b/n/ac downlink networks. The presented results are in line with the theoretical estimations from Section 3.

4.1 Experimental Methodology

We focus our experiments on SNR and PER measurements at Eve's receiver using the setup presented in Figure 2. We keep a ninety-degree angle between Bob and Eve to reduce the variability of the results. As we assume the attacker to be outside the beamformed region, we do not expect an impact of the angle on our measurements. We vary the distance from Bob and Eve (d_{BE}) while keeping the distance from Alice to Bob (d_{AB}) constant. Table 3 lists the parameters that

Table 3: *Parameters used for the experiments.*

Parameter	Value(s)	Description
P_A [dBm]	23	Alice's transmitted power
N_0 [dBm]	-91	Mean noise power at the receivers
$Ch_{b/n/ac}$	11, 11, 36	Channels used for 802.11 b/n/ac
d_{AB} [m]	2	Fixed distance from Alice to Bob
$\boldsymbol{d}_{AE}$ [m]	$[2.5, 5.0, \ldots, 20]$	Eight distances from Alice to Eve

we fix for our experiments with a short description. As stated in Section 3.1 we are not using link-layer encryption (which does not influence our measurements). Figure 5 shows the layout of the indoor office environment where we conducted the experiments.

The setup consists of an open access point (Alice) and a laptop (Bob) associated to it. The access point is a Linksys WRT3200 ACM device, equipped with four antennas and supporting 802.11a/b/g/n/ac. We installed the OpenWrt [28] operating system on the access point to have more configuration options at our disposal. For the 802.11b/n experiments (at 2.4 GHz), Bob's laptop runs Ubuntu 16.04 and has a TP-Link TL-WN722N wireless adapter. The adapter has a single antenna and supports 802.11b/g/n. Eve's laptop runs Ubuntu 16.04, and it uses the same TP-Link TL-WN722N wireless adapter. Eve's adapter is not associated with the access point and it tries to record the traffic from Alice and Bob, in monitor mode using `tcpdump`. The eavesdropper listens on the same channel as Alice and Bob and we use channel for 802.11b and n experiments.

For the 802.11ac experiments (at 5 GHz), Bob's laptop runs Ubuntu 16.04 and uses an Asus USB-AC68 wireless adapter, and the access point uses channel 36. The adapter has a 3x4:3 antenna configuration and supports 802.11a/b/g/n/ac. Eve's laptop is a MacBook Pro with an inbuilt adapter with 3x3:3 configuration that supports 802.11a/b/g/n/ac. We use a different adapter for Eve because the

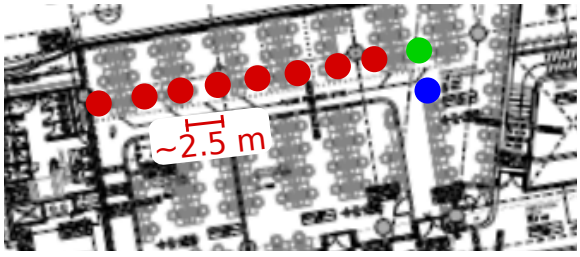

Fig. 5: *The layout of the indoor office environment where we conducted the experiments. The green and blue dots indicate the location of Alice and Bob. The red dots indicate the positions of Eve.*

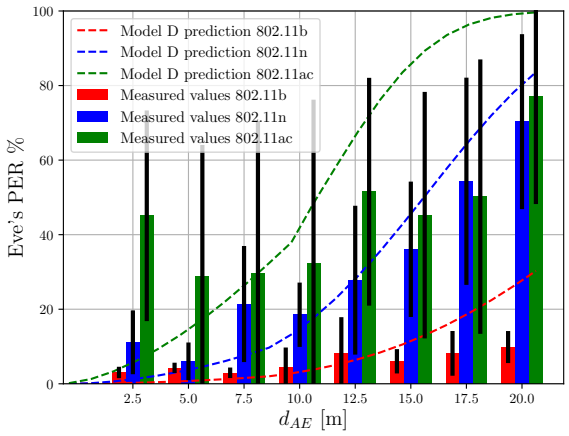

Fig. 6: *Eve's measured PER (bars) vs. Model D predicted PER (dashed lines).*

Asus adapter could not be put into monitor mode due to some issues with its driver. The other parameters remain the same as in the 802.11b/n experiments.

For all the experiments, we vary Eve's distance from Bob and we obtain pcap traces of the packets transferred from Alice to Bob. The distance between Alice and Bob (d_{AB}) is fixed at 2 m. We used `iperf` to generate UDP downlink traffic. We decided to use UDP to avoid retransmissions at the transport layer. PER is computed based on the number of received UDP packets with valid UDP checksums. We acknowledge that this approach slightly underestimates the actual PER, as packets with a valid UDP checksum but incorrect FCS might be included in this calculation. The (maximum) transmission power of Alice is set at 23 dBm. From the experiments, we are able to obtain traces from Eve at d_{AE} between 2.5 m and 20 m, using increments of 2.5 m. We do not change the orientation of Eve with respect to Alice in our tests to better compare the results. All the devices have the same fixed elevation, without a visual line-of-sight path between them. The information about the recorded traffic is obtained from the 802.11 PHY radiotap headers. In the subsequent section we will compare the experimental results with our estimations from the path loss model D (Office). Figure 9 shows the predicted BER and PER curves at Eve's receiver (red curves), and at Bob's receiver when Alice is using transmit-beamforming with two (green curves) and four antennas (blue curves).

4.2 Comparison between 802.11b/n/ac Networks

For the comparison between 802.11b/n/ac networks, we set a 2.4 GHz carrier for 802.11b/n and a 5 GHz for 802.11ac. To extract the results we capture packets both from Eve and Alice. We measured two parameters—the PER of the passive eavesdropper (percentage of packets that Eve failed to capture),

Table 4: *Results from 802.11n and 802.11ac experiments. d_{AE} is the distance from Alice to Eve in meters. n_r is the total number of runs. μ_p is the average number of UDP packets sent by Alice per run. μ_{PER} and σ_{PER} are the Eve's PER means and standard deviations measured in our experiments for 802.11n (n subscript) and 802.11ac (ac subscript).*

$d_{AE}[m]$	n_r	μ_p	μ_{PER_n}	σ_{PER_n}	$\mu_{PER_{ac}}$	$\sigma_{PER_{ac}}$
2.5	30	894.00	11.13	8.56	45.07	28.25
5.0	30	894.00	6.02	5.06	28.94	35.13
7.5	30	894.00	21.39	15.57	29.64	40.86
10.0	30	894.00	18.52	8.63	32.33	43.88
12.5	30	894.00	27.79	19.97	51.52	30.55
15.0	30	894.00	36.08	18.16	45.23	33.07
17.5	30	894.00	54.33	27.79	50.20	36.80
20.0	30	894.00	70.32	23.46	77.01	28.80

and her SNR. We compute Eve's PER by comparing her pcap traces with the ones from Alice. We compute the SNR by dividing the extracted signal strength values by the average channel noise. We computed the average channel noise using noise measurements from the access point. The channel noise averaged to -91 dBm. We repeat the same experiments with the same distances 30 times and we average the results to obtain mean SNR and PER values, and related errors (standard deviations).

Figure 6 shows Eve's PER measurements and estimated values for d_{AE} varying from 0 m to 20 m. The red/blue/green bars indicate the experimental results for 802.11b/n/ac, respectively. The dotted lines indicate the predicted estimates (from model D). It can be observed that Eve's PER is almost always *increasing* from b to n and from n to ac. In particular, the PER starts to increase significantly for distances greater than 15 meters. While such (relatively small-scale) experiments will hardly produce the exact same results as our theoretical analysis, we observe that the increase in PER that was predicted by us, for even relatively short distances of around 20 m, can be observed in practice. In particular, our D model predicted a PER for Eve in an 802.11n downlink of around 78% for 20 m distance and in our experiments the average PER was around 70%. For convenience, we tabulate in Table 4 the numerical results of Figure 6.

Figure 7 shows Eve's mean SNR varying her distance (d_{AE}) from Alice for 802.11b (red bars), 802.11n (blue bars) and 802.11ac (green bars). It can be observed that Eve's SNR in 802.11n/ac is always *smaller* than in 802.11b—an effect that we assumed to be caused by advanced 802.11n/ac physical and link layer features (such as TxBF).

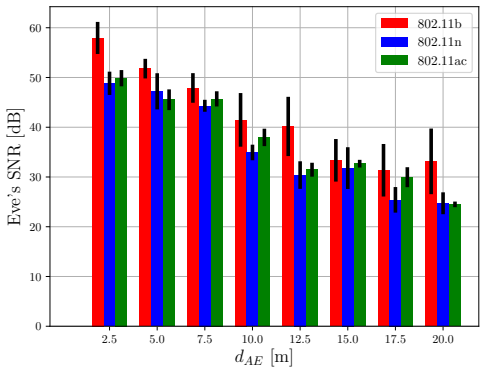

Fig. 7: *Eve's measured SNR with respect to* d_{AE}.

4.3 Bob vs. Eve in 802.11n/ac

We conducted a second set of experiments targeting Bob in order to compare his SNR and PER with respect to Eve's SNR and PER in 802.11n/ac networks. In this case, we increased Bob's distance from Alice. As in the previous experiment, we start from 2.5 m and we end at 20 m, with increments of 2.5 m. Bob is placed at the same location that Eve was placed in the previous experiment. In this scenario, we are expecting that Bob would benefit from 802.11n/ac features (such as TxBF). We are not showing a plot for Bob's PER compared to the one Eve experienced in Figure 6. This is because we observed that Bob's PER is very low (less than 1%), and yet not comparable with Eve's PER. This confirms our assumption that the intended receiver experiences significantly *lower* PER than a passive eavesdropper in 802.11n/ac networks.

Interestingly, as we can see from Figure 8a, the mean SNR of Bob and Eve at various distances are relatively close. In particular, Bob's SNR in 802.11n is always higher than Eve's SNR (as expected). However in the 802.11ac case, we measure a higher SNR for Eve than Bob. We assume that this is an artifact resulting from the fact that Eve's SNR is reported only for successfully received packets.

4.4 Eve's PER and PER Thresholds

It is important to note that even a small decrease in PER could affect a passive eavesdropper depending on the type of exchanged traffic. That is why we decided to analyze Eve's PER compared to different PER thresholds and distances d_{AE}. Table 5 shows the results of our analysis for 802.11b/n/ac. For example, if we fix the threshold to 15%, then Eve's PER in 802.11ac is above the threshold in at least 33% of all cases. The same holds for 802.11n except for the 5m measurement. With regards to 802.11b, fixing the same 15% threshold, we note that

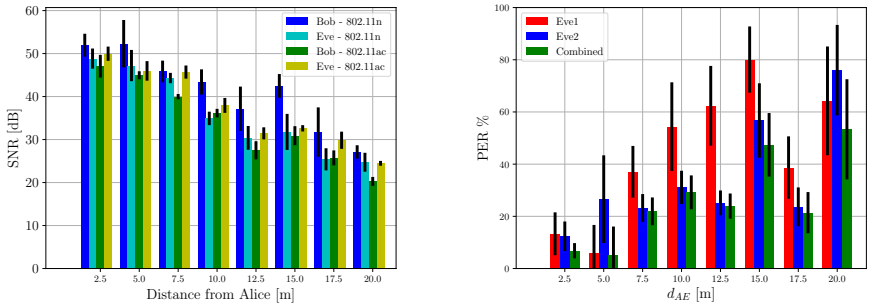

(a) *802.11n and 802.11ac SNR comparison between Bob and Eve at different distances from Alice.*

(b) *802.11n PER of Eve using two COTS radios. The green bars represent combined PER.*

Fig. 8: *Experimental results from Section 4.3 (a) and Section 4.5 (b).*

Eve's PER does not exceed the threshold in more than 16% of all cases. This is another way to confirm our predictions about 802.11n/ac passive eavesdropping.

4.5 Eve with Two COTS Radios in 802.11n

We argued earlier that attackers with COTS radios will not be able to benefit from advanced 802.11n/ac physical layer and link layer features, and discussed an attacker with a single COTS radio. We now discuss a passive eavesdropper with multiple COTS radios in an 802.11n downlink. The attacker aggregates the eavesdropped packets to reduce the number of packets lost (e. g. due to deep fading). In Figure 8b, we show the PER for an attacker with *two* COTS radios. The radios are placed at a distance of 50cm from each other (to avoid mutual coupling). Note that we used a different data set from the previous experimental section, and we repeated this experiment 30 times. It can be observed that such a scheme *reduces* the number of lost packets for the attacker (as expected). However, the PER in the aggregated case is still higher than the 802.11b one, especially at distances greater than 5m. For a threshold PER of 15%, the PER for the aggregated case is higher than the threshold in about 23% of the runs, compared to 6% for 802.11b.

4.6 Summary of 802.11b/n/ac Experiments

Overall, we were able to experimentally confirm our main findings: a) there is a significant increase of the PER of a passive eavesdropper attacking 802.11n and 802.11ac networks (compared to a 802.11b network). In our experiments, the difference was approximately 60% increased PER for 802.11n and 70% increased PER for 802.11ac at 20 m distance. In addition, the PER rises from around 12.5 m onward, similar to our predictions based on the theoretical analysis. We

also confirmed that the PER experienced by the attacker is related to the non-cooperating Alice. In particular, legitimate receivers at the same locations were able to receive traffic with close to zero PER.

5 Related Work

There are several empirical studies for 802.11 networks. Most of them focus on specific link layer [18] or physical layer [27] features. There are also more generic empirical studies, for example about enterprise WLAN [4], intrusion detection [15], denial of service [2] co-existence [11] and signal manipulation [24] Anyway, those studies neither focuses on wireless security nor compares end experimentally evaluate eavesdropping in various 802.11 networks.

An interesting aspect of eavesdropping is to study how to optimally place a set of antennas in a multiple users scenario to obtain the maximum amount of private information. In [30] Wang et al compare co-located vs. distributed eavesdropping schemes performance with respect to Eve's number of antennas and the presence of a guard zone. The de-facto standard countermeasure against eavesdropping (complementary to physical layer security) is cryptography. Several studies were done to secure [1] and break [3, 26] cryptographic systems used by 802.11 such as WEP and WPA.

6 Conclusions

In this work, we investigated the impact of physical and link layer 802.11n/ac features over a passive eavesdropper using COTS devices. We focused on downlink networks in infrastructure mode. We performed a theoretical discussion, a numerical simulation and several experiments comparing the Signal-to-Noise-Ratio and Packet-Error-Rates of the eavesdroppers in 802.11b/n/ac.

We showed that theoretically the eavesdropper's effective SNR is decreased by 6-41 dB in 802.11n/ac networks with four antennas ($L = 4$), which translates to a Packet-Error-Rate increase of up to 98% at a distance of 20 m between sender and eavesdropper. To obtain same Packet-Error-Rates as in a legitimate receiver, the attacker's maximal distance has to be reduced by 129.5 m in the case of 802.11n. In our practical experiments, we showed that the predicted effects occur in practice (although we were not able to exactly reproduce the theoretic predictions). Eve's PER for n was at least 20% higher than for b, and more than 30% for ac (with increasing impact over distances greater than 10m).

We conclude that the evolution of the 802.11 standard actually introduced several physical and link layer features, such as MIMO and spatial diversity, that might degrade the performance of a passive eavesdropper. If properly exploited those features could be used as a part of a defense-in-depth strategy as a complement to well-known eavesdropping defense mechanism (e. g. symmetric and asymmetric cryptography). Nevertheless, we understand that further investigations are necessary to characterize the benefits against stronger attacker models

(e. g. multiple attackers, SDR radios) and in more complex scenarios (e. g. multi-user, multi-access point). We leave this discussion to future work.

References

1. W. A. Arbaugh et al. *Real 802.11 security: Wi-Fi protected access and 802.11 i*. Addison-Wesley Longman Publishing Co., Inc., 2003.
2. M. Bernaschi, F. Ferreri, and L. Valcamonici. Access points vulnerabilities to dos attacks in 802.11 networks. *Wireless Networks*, 2008.
3. N. Borisov, I. Goldberg, and D. Wagner. Intercepting mobile communications: the insecurity of 802.11. In *Proceedings of the 7th annual international conference on Mobile computing and networking*. ACM, 2001.
4. Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *Proc. of Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, 2006.
5. Cisco. Cisco's visual networking index forecast projects nearly half the world's population will be connected to the internet by 2017. https://newsroom.cisco.com/press-release-content?articleId=1197391, 2013.
6. D. D. Coleman and D. A. Westcott. *CWNA: Certified Wireless Network Administrator Official Study Guide: Exam CWNA-106*. Sybex, 2014.
7. B. P. Crow, I. Widjaja, L. G. Kim, and P. T. Sakai. IEEE 802.11 Wireless Local Area Networks. *IEEE Communications Magazine*, 1997.
8. L. Dong, A. P. P. Z. Han, and H. V. Poor. Improving wireless physical layer security via cooperating relays. *IEEE Transactions on Signal Processing*, 2010.
9. A. Goldsmith. *Wireless communications*. Cambridge university press, 2005.
10. P. K. Gopala, L. Lai, and H. El Gamal. On the secrecy capacity of fading channels. *IEEE Transactions on Information Theory*, 2008.
11. R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of rf interference on 802.11 networks. *ACM SIGCOMM Computer Communication Review*, 2007.
12. A. Hero. Secure space-time communication. *IEEE Transactions on Information Theory*, 2003.
13. G. R. Hiertz, D. Denteneer, L. Stibor, Y. Zang, X. P. Costa, and B. Walke. The IEEE 802.11 universe. *IEEE Communications Magazine*, 2010.
14. IEEE. IEEE standard for information technology—telecommunications and information exchange between systems local and metropolitan area networks—specific requirements - part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications. http://standards.ieee.org/getieee802/download/802.11-2016.pdf, 2016.
15. C. Kolias, G. Kambourakis, A. Stavrou, and S. Gritzalis. Intrusion detection in 802.11 networks: empirical evaluation of threats and a public dataset. *IEEE Communications Surveys & Tutorials*, 2016.
16. S. K. Leung-Yan-Cheong and M. E. Hellman. The Gaussian Wire-Tap Channel. *IEEE Transactions on Information Theory*, 1978.
17. S. Martin. Directional Gain of IEEE 802.11 MIMO Devices Employing Cyclic Delay Diversity, 2013.
18. A. Mishra, M. Shin, and W. Arbaugh. An empirical analysis of the ieee 802.11 mac layer handoff process. *ACM SIGCOMM Computer Communication Review*, 2003.

19. A. Mukherjee and A. L. Swindlehurst. Robust beamforming for security in MIMO wiretap channels with imperfect csi. *IEEE Transactions on Signal Processing*, 2013.
20. F. Oggier and B. Hassibi. The secrecy capacity of the MIMO wiretap channel. *IEEE Transactions on Information Theory*, 2011.
21. E. H. Ong, J. Kneckt, O. Alanen, Z. Chang, T. Huovinen, and T. Nihtilä. IEEE 802.11 ac: Enhancements for very high throughput WLANs. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2011 IEEE 22nd International Symposium on*. IEEE, 2011.
22. K. P. Peppas, N. C. Sagias, and A. Maras. Physical layer security for multiple-antenna systems: A unified approach. *IEEE Transactions on Communications*, 2016.
23. E. Perahia and R. Stacey. *Next generation wireless LANs: 802.11 n and 802.11 ac*. Cambridge University Press, 2013.
24. C. Pöpper, N. O. Tippenhauer, B. Danev, and S. Čapkun. Investigation of signal and message manipulations on the wireless channel. In *Proc. of the European Symposium on Research in Computer Security*, 2011.
25. V. U. Prabhu and M. R. Rodrigues. On wireless channels with-antenna eavesdroppers: Characterization of the outage probability and-outage secrecy capacity. *IEEE Transactions on Information Forensics and Security*, 2011.
26. P. Robyns, B. Bonné, P. Quax, and W. Lamotte. Exploiting WPA2-enterprise vendor implementation weaknesses through challenge response oracles. In *WiSec*. ACM, 2014.
27. A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker. MOJO: A distributed physical layer anomaly detection system for 802.11 WLANs. In *Proceedings of the 4th international conference on Mobile systems, applications and services*. ACM, 2006.
28. O. D. Team. Openwrt wireless freedom. https://openwrt.org/.
29. B. Van Veen and K. Buckley. Beamforming: A Versatile Approach to Spatial Filtering. *IEEE ASSP Magazine*, 1988.
30. J. Wang, J. Lee, and T. Q. S. Quek. Best antenna placement for eavesdroppers: Distributed or co-located? *IEEE Communications Letters*, Sept 2016.
31. A. D. Wyner. The wiretap channel. *Bell System Technical Journal*, 1975.
32. N. Yang, P. L. Yeoh, M. Elkashlan, R. Schober, and I. B. Collings. Transmit antenna selection for security enhancement in MIMO wiretap channels. *IEEE Transactions on Communications*, 2013.
33. Y. Zou, J. Zhu, X. Wang, and V. C. M. Leung. Improving physical-layer security in wireless communications using diversity techniques. *IEEE Network*, 2015.

A Appendix

Figure 9 shows the result of our BER and PER analysis using model D. Figure 10 shows the result of our BER and PER analysis using model E. Figure 11 shows expected BER and PER for a free-space path-loss model.

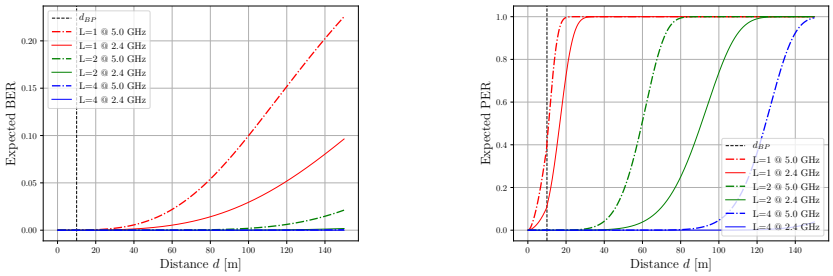

(a) *Expected BER vs. Distance.* **(b)** *Expected PER vs. Distance.*

Fig. 9: *802.11n Model D (office) BER/PER using BPSK. Red lines represent Eve. Green and Blue lines represent Bob when L=2 and L=4.*

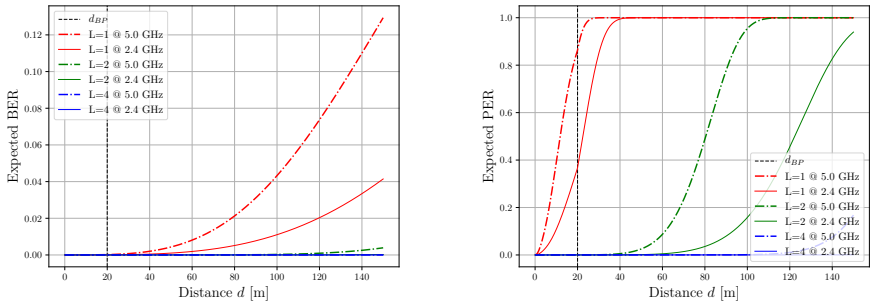

(a) *Expected BER vs. Distance.* **(b)** *Expected PER vs. Distance.*

Fig. 10: *802.11n Model E (Large office) BER/PER using BPSK. Red lines represent Eve. Green and Blue lines represent Bob when L=2 and L=4.*

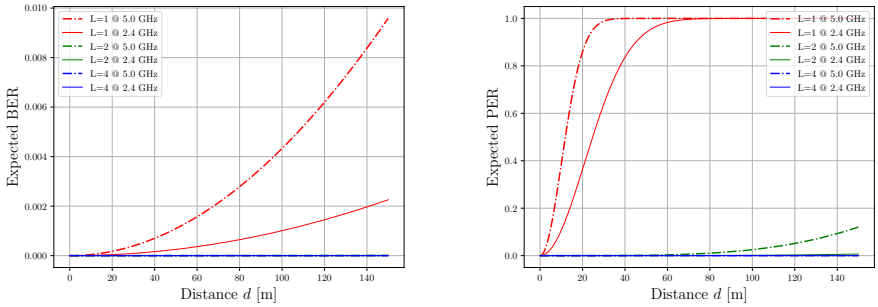

(a) *Expected BER vs. Distance.* **(b)** *Expected PER vs. Distance.*

Fig. 11: *Free Space Path Loss (LOS) BER/PER using BPSK. Red lines represent Eve. Green and Blue lines represent Bob when L=2 and L=4.*

Table 5: *Eve's PER vs. PER Thresholds vs. Distances. Columns represent different distances from Eve to Alice (d_{AE}). Rows represent different PER thresholds. Comma-separated values represent the rounded-down percentage of experimental runs where Eve's PER was above the threshold for 802.11b, n, and ac.*

	5.0 [m]	**7.5 [m]**	**10.0 [m]**	**12.5 [m]**	**15.0 [m]**	**17.5 [m]**
5%	33, 36, 50	10, 100, 33	20, 100, 33	36, 100, 90	43, 100, 80	60, 100, 96
10%	0, 26, 40	0, 73, 33	6, 83, 33	30, 90, 83	16, 96, 70	30, 100, 70
15%	0, 3, 36	0, 56, 33	6, 53, 33	16, 66, 76	0, 90, 63	13, 100, 60
20%	0, 0, 33	0, 43, 33	3, 36, 33	13, 53, 56	0, 76, 56	6, 96, 53
25%	0, 0, 33	0, 30, 33	3, 26, 33	10, 40, 53	0, 66, 56	0, 83, 53
30%	0, 0, 33	0, 20, 33	0, 13, 33	6, 30, 50	0, 60, 43	0, 73, 53
35%	0, 0, 30	0, 13, 30	0, 3, 33	3, 30, 43	0, 56, 43	0, 63, 50
40%	0, 0, 30	0, 10, 30	0, 0, 33	0, 23, 43	0, 40, 43	0, 53, 46
45%	0, 0, 26	0, 10, 30	0, 0, 33	0, 16, 43	0, 26, 43	0, 46, 46
50%	0, 0, 23	0, 6, 26	0, 0, 33	0, 16, 33	0, 16, 36	0, 43, 46