Benchmarking the Physical Layer of Wireless Cards using Software-Defined Radios

Liangxiao Xin

xlx@bu.edu

Boston University

Boston, MA

Stefan Gvozdenovic

tesla@bu.edu

Boston University

Boston, MA

Johannes K Becker

jkbecker@bu.edu

Boston University

Boston, MA

David Starobinski

staro@bu.edu

Boston University

Boston, MA

ABSTRACT

Many performance characteristics of wireless devices are fundamentally influenced by their vendor-specific physical layer implementation. Yet, characterizing the physical layer behavior of wireless devices usually requires complex testbeds with expensive equipment, making such behavior inaccessible and opaque to the end user. In this work, we propose and implement a new testbed architecture for software-defined radio-based wireless device performance benchmarking. The testbed is capable of accessing and measuring physical layer protocol features of real wireless devices. The testbed further allows tight control of timing events, at a microsecond time granularity. Using the testbed, we measure the receiver sensitivity and signal capture behavior of Wi-Fi devices from different vendors. We identify marked differences in their performance, including a variation of as much as 20 dB in their receiver sensitivity. We further assess the response of the devices to truncated packets and show that this procedure can be employed to fingerprint the devices.

CCS CONCEPTS

• **Networks** → **Network performance analysis**; Mobile and wireless security; *Wireless local area networks*; • **Hardware** → *Analog, mixed-signal and radio frequency test*.

KEYWORDS

Testbed, Wi-Fi, device fingerprinting, signal synthesis, interference, capture effect.

ACM Reference Format:

Liangxiao Xin, Johannes K Becker, Stefan Gvozdenovic, and David Starobinski. 2019. Benchmarking the Physical Layer of Wireless Cards using Software-Defined Radios. In *22nd Int'l ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '19)*, November 25–29, 2019, Miami Beach, FL, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3345768.3355907>

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

MSWiM '19, November 25–29, 2019, Miami Beach, FL, USA

© 2019 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-6904-6/19/11...\$15.00

<https://doi.org/10.1145/3345768.3355907>

1 INTRODUCTION

With the explosion of wireless device adoption, the problems of Wi-Fi channel congestion and resilience to interference are becoming more acute than ever, especially in densely populated areas. New Wi-Fi specifications such as 802.11ax (Wi-Fi 6) aim to mitigate this problem by supporting existing as well as anticipated additional unlicensed spectra (such as the new 3.5 GHz spectrum [25] and the expanded 6 GHz spectrum [9]) to avoid congestion. However, the large and growing number of legacy Wi-Fi devices means that performance bottlenecks on the given spectrum cannot be avoided. Hence, ensuring high performance despite channel congestion and interference is essential.

Wi-Fi devices are commodity hardware on a product level. Yet, subtle manufacturer-specific physical layer implementations can result in substantial performance differences that are opaque to end users. Benchmarking Wi-Fi performance and investigating behavior resulting from complex real-world situations, such as hidden nodes, currently require expensive physical setups in anechoic chambers under high time synchronization constraints. Specialized test equipment vendors offer wireless device testing equipment consisting of specialized hardware and software modules [15, 17, 23], which have to be integrated by trained specialists to perform as intended, and require considerable capital investment to procure.

To address this problem, we propose in this work a novel testbed architecture for physical layer benchmarking that consists of a simple setup made from cost-effective components. The key novelty of this architecture resides in emulating parts of the channel environment (including interference from other users) within a Software-Defined Radio (SDR)-based toolchain. The testbed reduces the complexity and expense required to conduct high-precision physical layer performance benchmarking, while leveraging the precise time synchronization and parameter control within the SDR to enable consistent and reproducible testing results.

We demonstrate the testbed capabilities by comparing the behavior and performance of Wi-Fi cards from four different manufacturers under precisely controlled physical layer testing conditions. First, we show that the cards exhibit noticeable differences in their receiver sensitivity (i.e., the lowest power level at which they can detect and demodulate RF signals). Next, we subject the devices to precisely time- and power-controlled collisions to assess their response to perturbed signals, thus demonstrating their different signal capture behavior. Finally, we show how device types can be

fingerprinted based on chipset-specific implementations. In particular, our results indicate distinct device responses to truncated (non-standard) packets that the testbed allows us to craft.

In summary, this paper makes the following contributions:

- (1) We propose an experimental testbed architecture for generating precisely timed traffic on the physical layer, subjecting real network devices to reproducible test conditions. The testbed can generate one or multiple packets at different power levels, emulate wireless interference and signal collisions on SDR hardware, and transmit the resulting composite signal to the device under test (DUT).
- (2) We demonstrate key features of the experimental testbed by measuring the devices' sensitivity and packet loss rate under different signal gains, and subjecting real Wi-Fi devices to packet collisions with high-fidelity control of timing and signal-to-interference ratio (SIR) parameters.
- (3) We show that it is possible to fingerprint different Wi-Fi devices based on their distinct sensitivity curves and different response to the capture effect and truncated packets.

The rest of this paper is organized as follows. In Section 2, we discuss related work. In Section 3, we describe our testbed architecture and our experimental setup. In Section 4, we discuss the experimental results. Finally, we conclude the paper and discuss future work in Section 5.

2 RELATED WORK

In this section, we provide an overview of previous work related to wireless testbeds and benchmarking, as well as theoretical and experimental analysis of the capture effect in Wi-Fi.

2.1 Benchmarking and Testbeds

Nychis et al. [16] propose an SDR-based platform that achieves precise packet timing by pre-loading a packet from the host to the FPGA and triggering its transmission based on the FPGA main clock on the USRP instead of the host clock (general purpose processor). Subsequent works aiming to satisfy the real-time requirements of wireless protocols follow this "split functionality" approach as well [2, 6], delegating the most real-time constrained functions within the protocol to customized FPGA modules. These workarounds are required to overcome processing, queuing, and bus transfer delays, which can add up to hundreds of microseconds [24]. Our testbed is significantly simpler, since it requires no FPGA modifications. Moreover, overlapping frames are added in software so that their offset is not affected by the host-radio hardware latency.

Park et al. [20] propose a wired testbed where signal of interest and interferers are generated on separate USRPs which are combined with a power combiner. While that testbed uses a sync cable to synchronize the two USRPs, our testbed generates both signals on the same device and therefore the same clock, precluding any frequency offset/drift errors.

In [11], Khorov et al. present a Wi-Fi testbed for investigating the capture effect. The testbed generates two data streams on the application layer, and processes them in parallel USRP transmission chains before sending them out over two antennas, generating a packet collision over-the-air. The offset between the frames is set by assigning each frame a different number and duration of backoff

slots. The two transmitters are synchronized with a common local oscillator. However, our testbed only requires a single transmission chain for multiple colliding packets, and does not require additional synchronization mechanisms.

Our work differs from the related works in the following aspects. First, we provide a cost effective (single USRP) experimental testbed that allows fine control of transmission frame parameters such as power, delay offset between frames, modulation, and frequency channel. As such, we are able to intentionally generate precise collision scenarios of interest instead of relying on a large volume of collision-producing traffic and subsequent filtering of suitable collisions in post-processing [20]. Furthermore, our testbed can easily compare multiple Wi-Fi devices directly, and without requiring calibration. This allows us to reveal differences in manufacturer implementation of the physical layer. Although we showcase the testbed with Wi-Fi devices, this methodology can be applied to devices implementing other protocols. This opens the door for device co-existence testing with multiple protocol stacks easily implemented in GNU Radio, similar to Liu et al. [13].

2.2 Capture Effect

The *capture effect* describes a scenario in which a Wi-Fi receiver receives multiple transmissions at once, and can properly decode the stronger frame despite the signals overlapping. This effect is highly time dependent. The physical layer (PHY) state machine in 802.11 starts by detecting a signal preamble, and – after successfully receiving metadata on the demodulation type and decoding rate of the signal – subsequently decoding the contained symbols into received data. If a stronger signal arrives at just the right time, it may supersede the existing signal on the receiver. Note that overlapping signals can occur in several practical situations, for instance if two nodes transmit (or re-transmit) packets at the same back-off slot time [1] or in a hidden node scenario [21, 31, 32] when two transmitting nodes cannot sense each other.

Traditional analytical models for IEEE 802.11 performance analysis do not take the capture effect into consideration. For instance, Bianchi's Markov chain model [1] and its refined models [10, 14, 22, 28] simply regard a packet collision as a packet loss. The work in [7] analyzes the performance of multi-hop 802.11 networks, under a full capture model (i.e., the stronger signal always captures the channel) and a limited capture model (i.e., the stronger signal captures the channel only if it comes first). In our paper, we show that none of the tested Wi-Fi devices behaves in full accordance with either one of these models.

Other work, such as Chatzimisios et al. [4] and Daneshgaran et al. [5], propose analytical models to calculate packet loss based on the bit error rate (BER). However, those analytical results are only verified in simulation environments and do not consider the additional complexities arising from physical layer implementation in real hardware.

Experimental studies on IEEE 802.11 networks consider the physical layer behavior of Wi-Fi devices. Ware et al. [29] demonstrate that the channel is always captured by the packet having the strongest SIR in hidden node scenarios. This capture behavior can cause unfairness issues within Wi-Fi networks, despite the use of request to send/ clear to send (RTS/CTS). However, the SIR is

Figure 1: Testbed architecture. The SDR and the device under test (DUT) are placed in a shielded test enclosure and controlled from dedicated hosts on the outside.

the only parameter studied in that work. In this paper, we consider additional parameter such as packet arrival time and different chipsets.

The work by Ganu et al. [8] evaluates the capture effect using the ORBIT indoor wireless testbed [19] in a scenario with no hidden nodes. Their experimental results show that the capture effect significantly reduces throughput fairness: When two stations transmit packets to the same receiver, the transmitter with weaker received signal strength indication (RSSI) has higher packet loss probability and longer backoff delays, resulting in negative impact on its throughput. However, they do not test the capture effect in a hidden node scenario. In this paper, we evaluate the capture effect in situations when the transmitters could be hidden nodes with respect to each other (i.e., there is a significant delay between the starts of overlapping frames). Furthermore, we do not require an expensive and complex setup to generate precisely timed signal collisions.

Lee et al. [12] design a testbed based on Atheros Wi-Fi cards and carry out a measurement study on the capture effect with hidden node scenario in IEEE 802.11a networks. They reveal the conditions under which the capture effect takes place, such as packet arrival timing, signal-to-interference ratio (SIR), and bit rate. Furthermore, they show that the packet preamble is more vulnerable to interference than the payload. However, this testbed consists of several independent Wi-Fi nodes, acting as sender, interferer, receiver, and sniffers. As a result, time synchronization between the nodes drifts over time, and other parameters like SIR cannot be precisely controlled. Our testbed allows for full control over all relevant parameters while requiring fewer devices and no complex topology and device manipulation in order to obtain precise results.

Finally, all aforementioned papers except [11] focus on evaluating the behavior of a single type of Wi-Fi card (chipset). In contrast, we compare the behaviors of multiple cards and show that they vary significantly.

3 TESTBED AND EXPERIMENTAL SET-UP

3.1 Testbed

The proposed testbed emulates one or multiple transmission signals on a single host and sends the resulting signal with a USRP to real wireless devices, where reception statistics are collected. Thus, the testbed allows us to emulate physical layer signal collisions and

allows fine-grained control of the parameters of the transmitting frames and of the channel, such as gain (attenuation), offset between frames, modulation, and channel frequency.

3.1.1 Hardware. The hardware setup of the testbed involves a transmitting host and a receiving host, and can be set up on a simple lab desk (see Figure 1(a)), whereas other wireless testbeds such as the ORBIT require extremely complex hardware configurations [18].

As shown in Figure 1(b), the transmitter consists of an Ettus USRP B200 SDR board connected to a host PC¹ via USB, and the receiver consists of a separate host PC configured with the appropriate USB- or PCIe-based network card (i.e., the device under test (DUT)). We use a RF cable to connect the USRP to the DUT. The cable has configurable attenuation to emulate signal loss on the transmission path. The SDR and the DUT are placed in a shielded enclosure to eliminate other interference sources.

3.1.2 Software. The software stack of our testbed consists of GNU Radio for signal generation, and the packet analyzer tcpdump [26] for collecting receiver data. On the transmitter side, we periodically generate Wi-Fi packets, using the gr-ieee802-11 library [3]. We emulate channel environment characteristics, such as relative signal strength, packet collision, and interference, directly on the transmitting host.

As shown in Figure 1(c), complex samples of signal and interference packets are summed up before transmission. Their transmission power gain as well as their delay relative to each other can be precisely controlled since they are both generated and added together on a symbol-level in software on the host (i.e., in GNU Radio) and transmitted with a single USRP. This setup ensures time synchronization in a much more straightforward way compared to setups with multiple physical transmitters. The two competing packets (signal of interest and interferer) are sent out with different MAC addresses to allow for easy packet statistics collection on the receiver side.

On the receiver side, a Wi-Fi card under test is connected to a separate host PC to receive Wi-Fi packets from the USRP. The card is set to monitor mode and data is collected via tcpdump. We then

¹Dell Precision Tower 5810 XCTO Base (CPU: Intel Xeon Processor E5-1607 v3 3.10 GHz $\times$ 4, RAM: 15.6 GB).

Table 1: Tested Wi-Fi cards.

Make	Model	Interface	Protocols	Chipset
Atheros	AR5B22	Mini PCIe	a/b/g/n	Atheros AR9462
TP-Link	TL-WN722N N150	USB	b/g/n	Atheros AR9271
Panda Wireless	PAU06 300Mbps N	USB	b/g/n	Ralink RT5372
AmazonBasics	Wi-Fi 11N USB Adapter - 300 Mbps	USB	b/g/n	Realtek RTL8192EU

count the number of received signal packets and compare it to the number of packets transmitted to obtain the packet loss statistics under each configuration.

3.2 Experimental Setup

We next describe the experiments performed using the testbed, including experimental setup, parameters, and performance metrics.

3.2.1 Devices under Test (DUTs). Our objective is to benchmark Wi-Fi cards with USB and PCIe-based interfaces, as shown in Table 1. All tested devices are popular, commodity devices using different Wi-Fi chipsets.

3.2.2 Parameters. The experiments take advantage of the high degree of parameter control that the testbed offers. In particular, we control the following parameters:

Delay offset (Δt), defined as the difference between the start time of the signal packet and the start time of the interference packet. Note that if the signal packet starts before the interference packet, the delay offset is negative. In the experiments, the delay offset is varied in steps of 1 μ s.

Signal and interference gains, which can be controlled directly within the transmission flowgraph.

Signal-to-interference ratio (SIR), which is the ratio of the strength of the signal packet to the strength of the interference packet in dB. Precise control of the SIR allows for reproducibility in experiments related to packet collisions.

3.2.3 Signal Gain and SIR. In order to achieve desired signal and interference gains and SIR, we adjust the amplitudes of the signal and interferer samples before they are summed up in GNU Radio.

Specifically, a wireless signal s can be represented as a sequence of discrete complex samples, with the n^{th} sample denoted by $s[n]$. We denote the transmission power gain of signal s by G_s . The (normalized) power of signal s is

$$P_s(G_s) = \frac{1}{N} \sum_{n=0}^{N-1} |G_s s[n]|^2. \quad (1)$$

The parameter G_s allows us to control the gain of the signal. Therefore, converting to dB units, we have

$$P_s(G_s) \text{ (dB)} = 20 \log_{10}(G_s) + P_{\text{USRP}}, \quad (2)$$

with the first term in the right hand side representing the *signal gain* (in dB), and the second term representing the transmission power offset of the USRP. We stress that the signal gain G_s is a relative quantity that is not calibrated to a specific output transmission power (i.e., one needs to estimate P_{USRP} if one wishes to know the actual transmission power P_s).

Figure 2: IEEE 802.11a/g packet format.

Note that Equation (2) is only applicable in the linear region of the transmitter's RF power amplifier. A too large value for G_s will eventually saturate the output power P_s to its maximum rated output power. Conversely, a too low value for G_s will flatten the output power at the noise floor.

Next, if we consider a desired signal s and interference signal i , we can express the signal-to-interference ratio (SIR) as

$$\text{SIR} = P_s - P_i = 20 \log_{10}\left(\frac{G_s}{G_i}\right), \quad (3)$$

where P_i is the interference power and G_i is the interference gain. In this paper, we use Equation (3) to calculate the SIR (e.g., setting $G_s = 1.0$ and $G_i = 0.1$ results in a SIR of -20 dB). G_s and G_i are chosen within the linear region of the transmitter's RF power amplifier where Equation (2) holds.

3.2.4 Experiments. In the experiments conducted in this paper, the signal packets consist of 200 byte-long IEEE 802.11g packets transmitted at 6 Mbit/s. The generated packets have payload containing random contents. The results are averaged over a larger number of packets (e.g., 100 or 1000).

Each packet contains both a preamble and a data payload (see Figure 2). Therefore, the duration of each packet is 328 μ s, whereby the duration of the preamble is always 20 μ s and the duration of the data is 308 μ s. The preamble consists of a 2-symbol (or 8 μ s) short training field. The following long training field (of the same length) is used for channel estimation, fine frequency offset estimation, and fine symbol timing offset estimation [27]. Finally, the third part of the packet preamble (the SIGNAL field) lasts 4 μ s and encodes the packet length and bit rate.

Using this configuration, we conduct the following experiments and measure the corresponding packet loss statistics:

- (1) **Receiver sensitivity** experiments measure and compare how devices react to different transmission power levels. We increase the signal gain G_s from -80 dB to 0 dB in steps of 4 dB. At each step, we transmit 1000 packets and record packet loss statistics. The RF cable has a 60 dB attenuation to protect the DUT. In this experiment, no interference packet is added.

Figure 3: Packet transmissions for the capture effect.

- (2) **Capture effect** experiments investigate packet loss during packet transmissions, as illustrated in Figure 3. Each experiment generates two packets: one packet defined as the *signal packet* and another packet defined as the *interference packet*. We craft precisely-timed packet collisions and measure whether the DUT experiences the capture effect, i.e., captures the signal packet despite the presence of an interference packet. We subject the DUTs to a range of colliding transmissions, varying Δt in increments of 1 μs . We transmit 1000 packets for each setting and record packet loss statistics. We further distinguish between the following three cases:

- Preamble capture effect*: The signal packet starts before or during the preamble of the interference packet.
- Body capture effect*: The signal packet starts during the frame (body) of the interference packet.
- Trailer capture effect*: The signal packet starts near the end of the interference packet.

Note that all the packet reception statistics reported in this paper pertain to signal packets. Interference packets are only used for emulating collisions.

- (3) **Truncated Packet Fingerprinting** experiments aim to characterize different devices based on their behavior in the presence of a specially crafted collision. We create an interference packet that contains a preamble, but no data afterwards. This truncated packet collides with a regular signal packet. We investigate how long it takes for a device to recover from such a bogus packet, i.e., at what time after the end of bogus packet can a valid packet be received again. We vary the delay offset from $-5 \mu s$ to $335 \mu s$ to capture packet loss statistics across the full length of a signal packet.

4 EXPERIMENTAL RESULTS

In this section, we detail the results of our experiments for each of the four DUTs listed in Table 1.

4.1 Receiver Sensitivity

In our first experiment, we evaluate DUT performance in terms of their receiver sensitivity. Specifically, we measure the packet loss ratio as a function of the transmission power gain G_s .

Subjecting all DUTs to test packets with varying signal gain G_s , we obtain the results shown in Figure 4. We can clearly identify and distinguish the receiver sensitivity of different devices with great precision (the 95% confidence interval based on 1000 samples is tight ($\pm 0.47\%$ around the mean), as indicated by the barely visible colored bands around the chart lines).

Interestingly, the devices exhibit markedly different sensitivity. In particular, the Atheros and TP-Link cards first start picking up packets at -60 dB and -56 dB, respectively, whereas the Panda card

Figure 4: Receiver sensitivity of different Wi-Fi cards depending on the transmission power gain G_s .

only starts picking up packets at -36 dB. Being able to distinguish these differences in receiver sensitivity allows us to compare devices regarding their performance in weak signal scenarios, such as strong attenuation occurring in densely developed areas.

We also note that in the range between -28 dB and 0 dB, packets are reliably picked up by all of the devices. In subsequent experiments involving packet collision, we use signal gains in this range, as we need to ensure that packets would have been received correctly if they were transmitted without overlap.

4.2 Capture Effect

We then apply our testbed to investigate the capture effect occurrence in different Wi-Fi devices. Successful capture in the presence of interference depends on different parameters, such as the SIR, and the delay offset Δt .

4.2.1 SIR. We first determine the power and delay conditions under which the capture effect occurs. We vary the SIR from 0 dB to 36 dB by fixing $G_s = 0$ dB and varying the interfering signal gain from 0 dB to -36 dB in steps of 4 dB. We also vary the delay offset from $-1 \mu s \leq \Delta t \leq 10 \mu s$. For each configuration, we generate 100 packets and measure the packet loss of signal packets. Note that we reduce the number of packets for this experiment due to the high number of different SIR and delay offset combinations examined, this after confirming that the confidence intervals remain acceptable: We indeed observe an average 95%-confidence interval of $\pm 1.3\%$ around the mean across all measurements.

Figure 5 shows the packet loss of signal packets at different SIRs and Δt . This graph shows bright spots for all parameter configurations with reliable reception (low packet loss) of the signal packet and darker spots wherever the packet loss is high.

In Figure 5, we observe that the devices behave quite differently, i.e., they experience the capture effect within different boundary conditions. For example, the TP-Link manages to receive the signal packet only if the SIR is above 4 dB, but, independently of the SIR, only up to a delay of $3 \mu s$. In contrast, the Panda Wireless device requires a higher SIR for successful reception, but is capable of receiving the signal up to $8 \mu s$ after the interference packet, while showing a greater variance in its behavior overall.

In general, the data shows that the capture effect requires a certain minimum SIR and gives reason to assume that after a certain Δt , the capture effect does not occur any more – independent of the

Figure 5: Impact of SIR and packet delay on the capture effect in different Wi-Fi cards. Darker shade means higher packet loss.

SIR. This may be due to the receiver already locking on to a signal during the preamble, based in individual vendor implementation.

The Atheros AR5B22 card is an exception to this observation. In Figure 5(a), we observe that the Atheros card stops capturing new packets – independent of the chosen SIR – at 4 μ s, but then resumes capture above a certain SIR threshold. To confirm this finding, we conduct further related experiments in Section 4.2.2.

4.2.2 Delay Offset. The previous experiments showed that after a certain delay offset, the capture effect does no longer occur in several of the devices. We investigate whether this result remains consistent throughout the whole range of possible delay offsets, i.e., for all possible overlaps between interference and signal packets.

In the following, we fix the signal gain $G_s = 1.0$ and $G_i = 0.1$, such that $SIR = 20$ dB. At these settings, both packets would be reliably received if they were sent without overlap. We vary the delay offset Δt and transmit 1000 packets for each configuration, collecting packet loss statistics at the receiver. This time, the range of Δt values considered exceeds the length of a single packet transmission (328 μ s). The goal is to find out whether signal capture behavior occurs when a signal packet starts right after an interference packet.

Figure 6(a) shows the capture effect of different cards for low Δt . We observe that each tested device has a characteristic capture behavior, and transitions to 100% packet loss after a certain delay offset. This result indicates that the capture effect occurs only if the delay offset is small, and implies that the receiver locks on to the packet after it receives the first few bits of a packet. Then, receivers typically cannot detect another packet until the packet transmission

ends. This result shows that the delay offset plays a critical role in the packet loss of the signal packets.

Indeed, this behavior remains consistent until the end of the interference packet. However, as shown in Figure 6(b), we can observe that devices again behave differently after receiving an incoming packet. Some devices exhibit the capture effect shortly before the interference packet ends (at 328 μ s), while others cannot immediately switch to receive the signal packet after the end of the interference packet. We believe this is again due to different physical layer implementations of the standard in the various chipsets.

Coming back to the Atheros AR9642 chipset, we run additional tests on the Atheros AR5B22 card only, varying the SIR from between 16, 24, and 32 dB SIR, and testing the whole range of Δt from the beginning of the interferer preamble at $\Delta t = 0$ μ s until the end of the packet (at $\Delta t = 328$ μ s) in steps of 5 μ s. Indeed, as shown in Figure 7, capture is possible not only during the whole length of the preamble, but along the full length of the interference packet, if the SIR is strong enough. In other words, the Atheros AR9642 chipset seems to implement body capture above a certain SIR^2 . We note that this behavior can be found in the Atheros AR9462 chipset, but not in the AR9271 chipset of the TP-Link device that we tested.

4.3 Truncated Packet Fingerprinting

Wi-Fi devices implement the physical layer as a state machine, i.e., the receiver has one state to detect the packet preamble and another

²This confirms a recent finding by Khorov et al. [11], who identified the body capture effect in the similar, but not identical, AR9485.

(a) Packet loss ratio at the beginning of an interference packet.

(b) Packet loss ratio at the end of an interference packet.

Figure 6: Packet loss depending on the signal delay offset Δt relative to the beginning of an interference packet at 20 dB SIR. Figure (a) shows packet loss at low Δt , and Figure (b) around the end of the interference packet. Yellow and orange background indicates collision with the preamble and payload of the interference packet, respectively.

Figure 7: The Atheros AR9462 chipset captures new packets even while it is already receiving a packet body, if the SIR is sufficiently high. The graph shows packet loss for different packet delay offsets and SIR.

state to receive the packet frame [30]. We next investigate whether different devices implement such state machines in different ways by examining their response to certain crafted signals.

The experiment setup is the same as in Section 4.2.2, except that the interference is not a valid Wi-Fi packet this time. Instead, we only transmit a preamble and truncate the packet data (MAC frame). Generally, if the signal packet arrives after the end of the interference preamble (without a frame) and experiences loss, such packet loss is not caused by a collision (as there is no data to collide with). Instead, the reason for the packet loss is that the receiver is in a state that does not allow it to capture a new packet.

Figure 8 depicts the results of this experiment. When $\Delta t \leq 20 \mu s$, the interference packet collides with the preamble of the truncated packet. The packet loss ratio jumps to 100% after a few microseconds delay offset, in the same chipset-specific way that we observed in the capture effect experiment. This shows that the truncated packet colliding with the preamble of the signal packet results in the same capture behavior as described in the previous section.

Once Δt exceeds $20 \mu s$, the delay offset is such that the signal preamble would collide with the data field of the interference packet. However, since the truncated interference packets have no data field, there is no data to collide with. Interestingly, the behavior of the DUTs in this scenario varies considerably: Whereas after 30 to 50 microseconds the TP-Link and AmazonBasics cards recover to a state in which they can capture new packets, the Atheros card experiences about 50% packet loss for the whole duration of the non-existent interference packet's data, and the Panda card experiences near total packet loss until the nominal end of the expected packet duration. This demonstrates that the tested cards have widely different state machine implementations, especially regarding the transition from the state of packet preamble detection to the state of packet reception and back.

Probing devices with such specially crafted signals allows for physical layer fingerprinting of the devices based on their chipset implementation. Such way of fingerprinting could be used, for instance, as an additional factor in authentication scenarios in which the physical device identity is critical, confirming or rejecting that communication is coming from the desired device without alerting the application layer.

5 CONCLUSION

In this paper, we present an SDR-based testbed that achieves precise parameter control suitable for wireless device testing. We use the testbed to evaluate a range of Wi-Fi cards regarding different performance aspects. In particular, the receiver sensitivity of the cards varies by as much as 20 dB. We also investigate the capture effect on IEEE 802.11 networks by designing experiments that allow us to capture differences emerging on the scale of microseconds. Thanks to the precise parameter control made possible by the testbed, we provide quantitative analysis on the impact of packet arrival time, SIR and manufacturer-specific implementation on the occurrence of the capture effect. Notably, among the four different Wi-Fi cards tested, capture of the preamble varies by as much as $7 \mu s$ in terms of the delay offset. The experiments further show that some Wi-Fi cards exhibit body capture effects while others do not, thus cross-validating findings from [11].

Our work shows that it is valuable to compare multiple cards at high temporal resolution, because manufacturers differently implement physical layer features that are not precisely defined in the standard. Thus, one should not assume that implementation characteristics of a specific Wi-Fi card are generally applicable to all Wi-Fi cards. This finding is especially important when developing analytical and simulation models of Wi-Fi networks.

Another interesting finding is that two of the tested Wi-Fi chipsets appear to return to the preamble detection state earlier than the standard defines. This specific feature may potentially have performance benefits in congested networks, allowing them to detect PHY preambles more aggressively.

The experimental results of this paper can further serve to fingerprint the tested devices, especially based on their physical layer responses to truncated interference packets. Since these responses are hardwired into the chipset, the fingerprints may be of interest both as an additional authentication factor, as well as for covert device tracking devices (circumventing higher layer anonymization).

Much additional work can be performed based on the testbed proposed in this paper, as it provides a flexible platform for any kind of wireless experimentation and is not limited to a specific communication protocol. For instance, one could investigate fingerprinting groups of similar devices as compared to fingerprinting individual devices from separate vendors, as done in this work. Aside from further expansion on performance characterization of Wi-Fi cards (including different 802.11 variants), one could expand the scope of this work to investigate low-layer performance, privacy, and security characteristics of other popular wireless protocols, such as Bluetooth. Future work could also involve expanding the testbed to bidirectional communication testing, which opens up a new range of methods, e.g., fingerprinting based on response delays.

ACKNOWLEDGMENTS

This work is funded in part by NSF under grant CNS-1409053.

REFERENCES

- [1] Giuseppe Bianchi. 2000. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications* 18, 3 (3 2000), 535–547. <https://doi.org/10.1109/49.840210>
- [2] Bastian Bloessl, Andre Puschmann, Christoph Sommer, and Falko Dressler. 2014. Timings matter. In *Proceedings of the 9th ACM international workshop on Wireless network testbeds, experimental evaluation and characterization - WiNTECH '14*.

Figure 8: To fingerprint Wi-Fi chipsets, we generate collisions between signal packets and specially crafted truncated packets containing only a preamble, and measure the DUT's packet loss. Note that no actual signal collision occurs after the end of the preamble, i.e., packet loss at $\Delta t \geq 20 \mu s$ is only a result of the receiver's physical layer state machine implementation.

- ACM Press, New York, New York, USA, 57–64. <https://doi.org/10.1145/2643230.2643240>
- [3] Bastian Bloessl, Michele Segata, Christoph Sommer, and Falko Dressler. 2013. An IEEE 802.11 a/g/p OFDM Receiver for GNU Radio. In *Proceedings of the second workshop on Software radio implementation forum*. ACM, 9–16.
 - [4] Periklis Chatzimisios, A.C. Boucouvalas, and Vasileios Vitsas. 2004. Performance analysis of IEEE 802.11 DCF in presence of transmission errors. In *2004 IEEE International Conference on Communications (IEEE Cat. No.04CH37577)*, Vol. 7. IEEE, IEEE, 3854–3858. <https://doi.org/10.1109/ICC.2004.1313274>
 - [5] Fred Daneshgaran, Massimiliano Laddomada, Fabio Mesiti, Marina Mondin, and Massimiliano Zanolò. 2008. Saturation throughput analysis of IEEE 802.11 in the presence of non ideal transmission channel and capture effects. *IEEE Transactions on Communications* 56, 7 (7 2008), 1178–1188. <https://doi.org/10.1109/TCOMM.2008.060397>
 - [6] Paolo Di Francesco, Seamas McGettrick, Uchenna K Anyanwu, James C. O'Sullivan, Allen B MacKenzie, and Luiz A DaSilva. 2015. A Split MAC Approach for SDR Platforms. *IEEE Trans. Comput.* 64, 4 (4 2015), 912–924. <https://doi.org/10.1109/TC.2014.2308197>
 - [7] M. Durvy, O. Dousse, and P. Thiran. 2007. Modeling the 802.11 Protocol Under Different Capture and Sensing Capabilities. In *IEEE INFOCOM 2007 - 26th IEEE International Conference on Computer Communications*. IEEE, 2356–2360. <https://doi.org/10.1109/INFCOM.2007.280>
 - [8] Sachin Ganu, Kishore Ramachandran, Marco Gruteser, Ivan Seskar, and Jing Deng. 2006. Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture. In *Proceedings of the second international workshop on Multi-hop ad hoc networks: from theory to reality - REALMAN '06*. ACM, ACM Press, New York, New York, USA, 7. <https://doi.org/10.1145/1132983.1132986>
 - [9] Neil Grace. 2018. FCC Proposes More Spectrum For Unlicensed Use. <https://www.fcc.gov/document/fcc-proposes-more-spectrum-unlicensed-use>
 - [10] Zoran Hadzi-Velkov and Boris Spasenovski. 2003. Saturation throughput - delay analysis of IEEE 802.11 DCF in fading channel. In *IEEE International Conference on Communications, 2003. ICC '03.*, Vol. 1. IEEE, IEEE, 121–126. <https://doi.org/10.1109/ICC.2003.1204154>
 - [11] Evgeny Khorov, Aleksey Kureev, Ilya Levitsky, and Andrey Lyakhov. 2018. Testbed to Study the Capture Effect: Can We Rely on this Effect in Modern Wi-Fi Networks. In *2018 IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*. IEEE, 1–5. <https://doi.org/10.1109/BlackSeaCom.2018.8433688>
 - [12] Jeongkeun Lee, Wonho Kim, Sung-Ju Lee, Daehyung Jo, Jiho Ryu, Taekyoung Kwon, and Yanghee Choi. 2007. An experimental study on the capture effect in 802.11a networks. In *Proceedings of the the second ACM international workshop on Wireless network testbeds, experimental evaluation and characterization - WinTECH '07*. ACM, ACM Press, New York, New York, USA, 19. <https://doi.org/10.1145/1287767.1287772>
 - [13] Wei Liu, Eli De Poorter, Jeroen Hoebeke, Emmeric Tanghe, Wout Joseph, Pieter Willemen, Michael Mehari, Xianjun Jiao, and Ingrid Moerman. 2017. Assessing the Coexistence of Heterogeneous Wireless Technologies With an SDR-Based Signal Emulator: A Case Study of Wi-Fi and Bluetooth. *IEEE Transactions on Wireless Communications* 16, 3 (2017), 1755–1766. <https://doi.org/10.1109/TWC.2017.2654256>
 - [14] David Malone, Ken Duffy, and Doug Leith. 2007. Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions. *IEEE/ACM Transactions on Networking* 15, 1 (2 2007), 159–172. <https://doi.org/10.1109/TNET.2006.890136>
 - [15] National Instruments. 2019. Simple Solutions to Complex Problems. <http://www.ni.com/en-us/shop.html>
 - [16] George Nychis, Thibaud Hottelier, Zhuocheng Yang, Srinivasan Seshan, and Peter Steenkiste. 2009. Enabling MAC Protocol Implementations on Software-defined Radios. *NSDI'09 Proceedings of the 6th USENIX symposium on Networked systems design and implementation* (2009), 91–105.
 - [17] octoScope Inc. 2019. Wireless Personal Testbeds. <http://octoscope.com/English/Products/Ordering/index.html>
 - [18] ORBIT Lab. 2016. Hardware. <https://www.orbit-lab.org/wiki/Hardware>
 - [19] ORBIT Lab. 2019. Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT). <https://www.orbit-lab.org/>
 - [20] Jin Soo Park, Hyungoo Yoon, and Byung Jun Jang. 2016. SDR-based frequency interference analysis test-bed considering time domain characteristics of interferer. *International Conference on Advanced Communication Technology, ICACCT 2016-March* (2016), 517–521. <https://doi.org/10.1109/ICACCT.2016.7423454>
 - [21] Saikat Ray, David Starobinski, and Jeffrey B. Carruthers. 2005. Performance of wireless networks with hidden nodes: a queueing-theoretic analysis. *Computer Communications* 28, 10 (6 2005), 1179–1192. <https://doi.org/10.1016/j.comcom.2004.07.024>
 - [22] J.W. Robinson and T.S. Randhawa. 2004. Saturation Throughput Analysis of IEEE 802.11e Enhanced Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications* 22, 5 (6 2004), 917–928. <https://doi.org/10.1109/JSAC.2004.826929>
 - [23] Rohde & Schwarz GmbH. 2019. Test Systems & Accessories. https://www.rohde-schwarz.com/us/products/test-and-measurement/wireless-communications-testers-systems/wireless-communication-testers-systems/test-systems-accessories_86246.html
 - [24] Thomas Schmid, Oussama Sekkat, and Mani B. Srivastava. 2007. An experimental study of network performance impact of increased latency in software defined radios. (2007), 59. <https://doi.org/10.1145/1287767.1287779>
 - [25] Cecilia Sulhoff. 2018. FCC Takes Action To Encourage Increased Investment And Deployment In The 3.5 GHz Band. <https://www.fcc.gov/document/fcc-acts-increase-investment-and-deployment-35-ghz-band>
 - [26] TCPDUMP & LIBPCAP. 2019. TCPDUMP & LIBPCAP. <https://www.tcpdump.org/>
 - [27] The Mathworks Inc. 2019. WLAN Packet Structure. <https://www.mathworks.com/help/wlan/ug/wlan-packet-structure.html#buytqq7-12>
 - [28] Ilenia Tinnirello, Giuseppe Bianchi, and Yang Xiao. 2010. Refinements on IEEE 802.11 Distributed Coordination Function Modeling Approaches. *IEEE Transactions on Vehicular Technology* 59, 3 (3 2010), 1055–1067. <https://doi.org/10.1109/TVT.2009.2029118>
 - [29] Christopher Ware, John Judge, Joe Chicharo, and Eryk Dutkiewicz. 2000. Unfairness and capture behaviour in 802.11 adhoc networks. In *2000 IEEE International Conference on Communications. ICC 2000. Global Convergence Through Communications. Conference Record*, Vol. 1. IEEE, IEEE, 159–163. <https://doi.org/10.1109/ICC.2000.853084>
 - [30] Liangxiao Xin and David Starobinski. 2018. Cascading Attacks on Wi-Fi Networks with Weak Interferers. In *Proceedings of the 21st ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems - MSWiM '18*. ACM Press, New York, New York, USA, 255–258. <https://doi.org/10.1145/3242102.3242142>
 - [31] Liangxiao Xin and David Starobinski. 2018. Mitigation of Cascading Denial of Service Attacks on Wi-Fi Networks. In *2018 IEEE Conference on Communications and Network Security (CNS)*. IEEE, Beijing, China, 1–9. <https://doi.org/10.1109/CNS.2018.8433124>
 - [32] Liangxiao Xin, David Starobinski, and Guevara Noubir. 2016. Cascading denial of service attacks on Wi-Fi networks. In *2016 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 91–99. <https://doi.org/10.1109/CNS.2016.7860474>