Spatial-Domain Wireless Jamming with Reconfigurable Intelligent Surfaces

Philipp Mackensen[1,*], Paul Staat[1,†], Stefan Roth[*], Aydin Sezgin[*], Christof Paar[†] and Veelasha Moonsamy[*]
[*]Ruhr University Bochum, Bochum, Germany
[†]Max Planck Institute for Security and Privacy, Bochum, Germany
E-Mail: {philipp.mackensen, stefan.roth-k21, aydin.sezgin, veelasha.moonsamy}@rub.de,
{paul.staat, christof.paar}@mpi-sp.org

***Abstract*—Wireless communication infrastructure is a cornerstone of modern digital society, yet it remains vulnerable to the persistent threat of wireless jamming. Attackers can easily create radio interference to overshadow legitimate signals, leading to denial of service. The broadcast nature of radio signal propagation makes such attacks possible in the first place, but at the same time poses a challenge for the attacker: The jamming signal does not only reach the victim device but also other neighboring devices, preventing precise attack targeting.**

In this work, we solve this challenge by leveraging the emerging reconfigurable intelligent surface (RIS) technology, for the first time, for precise delivery of jamming signals. In particular, we propose a novel approach that allows for environment-adaptive spatial control of wireless jamming signals, granting a new degree of freedom to perform jamming attacks. We explore this novel method with extensive experimentation and demonstrate that our approach can disable the wireless communication of one or multiple victim devices while leaving neighboring devices unaffected. Notably, our method extends to challenging scenarios where wireless devices are very close to each other: We demonstrate complete denial-of-service of a Wi-Fi device while a second device located at a distance as close as $5\,\mathrm{mm}$ remains unaffected, sustaining wireless communication at a data rate of $25\,\mathrm{Mbit/s}$. Lastly, we conclude by proposing potential countermeasures to thwart RIS-based spatial domain wireless jamming attacks.

I. Introduction

Wireless communication systems are ubiquitous and seamlessly provide connectivity to the smart and interconnected devices that permanently surround us. In our modern daily lives, we frequently use instant messaging, media streaming, health monitoring, and home automation – all of which rely on wireless systems and their constant availability. However, wireless systems utilize a broadcast medium that is open to everyone, inherently exposing a large attack surface. One particular critical threat is *wireless jamming*, which allows malicious actors to perform denial of service attacks with minimal effort. In a classical jamming attack, the adversary transmits an interfering signal that overshadows the desired signal, preventing a victim receiver from correctly decoding it. Crucially, loss of connectivity impacts the functionality of wireless devices and can thus have potentially far-reaching consequences, such as in smart grids, smart transportation, and healthcare systems. Recent media reports underscore the real-world threat potential of jamming attacks, *e.g.*, criminals disabling smart home security systems [57, 6] and preventing cars from locking [7].

This basic attack principle has previously been studied by a large body of research: For instance, the attacker can leverage various jamming waveforms, such as noise or replayed victim signals [23], and vary the attack timing, jamming constantly [72] or only at certain times [52]. As evident from the many existing attack strategies [35, 72, 45, 43], wireless jamming has been incrementally refined and became increasingly sophisticated. One particular example for this is the case of selective jamming attacks.

To illustrate a potential attack scenario, consider an adversary attempting to sabotage a complex automated manufacturing process. Distributed actuators might take orders from several previous processing stages that have to be executed in a timely fashion, risking manufacturing failure otherwise. Here, the adversary could use selective jamming to simulate local loss of connectivity on a single actuator but not the entire plant which would likely trigger some emergency response.

So far, the only means to realize such a selective jamming attack is via so-called reactive jamming. Here, the attacker analyzes all wireless traffic in real-time to decide on-the-fly whether to send a jamming signal [52, 46, 3], relying on the existence of meaningful protocol-level information not protected by cryptographic primitives. In our manufacturing plant example, selective disruption of the actuator would require the attacker to receive and identify *every* packet directed to the recipient before sending a jamming signal. This restricts the attacker positioning rather close to the victim. Other downsides of this approach are that it can be mitigated by fully disguising packet destinations and the attack realization being rather complex and cumbersome.

In light of these aspects, we are interested in novel attack strategies resolving the aforementioned shortcomings. Clearly, the ideal solution would be to physically inject a proactive jamming signal directly and only into the victim device which, however, is not possible due to the wireless nature

[1]These authors contributed equally to this work.

Network and Distributed System Security (NDSS) Symposium 2025
24-28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-8-3
https://dx.doi.org/10.14722/ndss.2025.240440
www.ndss-symposium.org

of jamming and the inevitable broadcast behavior of radio signal propagation to other, non-target devices. Thus, we aim to answer the following research question:

How can we physically target and jam one device while keeping others operational?

We solve this challenge by means of a reconfigurable intelligent surface (RIS) to devise the first selective jamming mechanism based on taming random wireless radio wave propagation effects. Using RIS-based environment-adaptive wireless channel control, allowing to maximize and minimize wireless signals on specific locations [27], the attacker gains spatial control over their wireless jamming signals. This opens the door to precise jamming signal delivery towards a target device, disrupting any legitimate signal reception, while leaving other, non-target devices, untouched. Other than reactive jamming, this is a true physical-layer selection mechanism, allowing realization independent of protocol-level information. Moreover, the attacker only initially needs to detect signals from considered devices, removing the need for any real-time monitoring and reaction to ongoing transmissions.

In this work, we experimentally evaluate RIS-based spatially selective jamming attacks against Wi-Fi communication, showing that it is possible to target one or multiple devices while keeping non-target devices operational. To accomplish this, we exploit that considered devices transmit signals, allowing the attacker to passively adapt to the scene. Apart from the attack's core mechanism, we study crucial real-world aspects such as the attack's robustness against environmental factors. We additionally verify the effectiveness of our attack in real-world wireless networks, where mechanisms that could counteract the attack are at play, *e.g.*, adaptive rate control of Wi-Fi networks. We show that RIS-based selective jamming even works despite extreme proximity of devices, *e.g.*, $5\,\mathrm{mm}$, and investigate the underlying physical mechanisms. Finally, we perform comparison experiments with a directional antenna, showing significant of our RIS-based approach.

In summary, our work makes the following key contributions:

- We propose the first true physical-layer selective targeting mechanism for wireless jamming, enabling environment-adaptive attacks in the spatial domain.
- We present an attack realization based on RIS, using passive eavesdropping to determine an appropriate RIS configuration which is the key to deliver jamming jamming signals towards targeted devices while avoiding non-target devices.
- We present a comprehensive experimental evaluation with commodity Wi-Fi devices, environmental changes, and an in-depth analysis of the physical properties of our jamming attack.

II. Technical Background

In this section, we introduce the necessary background on wireless jamming and RISs.

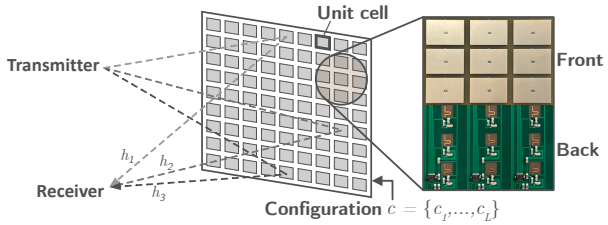

Fig. 1: Illustration of the RIS operation principle along with photos of the RIS hardware implementation [24], where the configuration vector c determines the radiation behavior.

A. Wireless Jamming

Wireless communication quality of service, *e.g.*, reliability and data throughput, is determined by the quality of received signals. An example is weak signal reception which yields a low signal-to-noise ratio (SNR) and consequently, increases the probability of bit errors. Similarly, when multiple radios transmit simultaneously, the receiver observes the superposition of multiple signals, *i.e.*, the desired signal with additional interference, again degrading performance. A *wireless jamming* attacker exploits this mechanism by deliberately sending strong interfering signals. At the victim receiver, they overshadow legitimate signals, increase the bit error probability, and eventually lead to complete denial of service [45, 21, 35].

There exist several types of jamming attacks, differing in terms of, *e.g.*, specifics of the attack target, the used interfering signal waveform, or the level of cognition [21]. The jamming signal can comprise of, *e.g.*, noise, single or multi-tone carriers, or valid waveforms carrying digital information. While noise jamming reduces the SNR, valid signals can enhance attack effectiveness, *e.g.*, by keeping the receiver busy [22]. Attackers may constantly or reactively transmit the jamming signal, *i.e.*, upon detecting a certain trigger [18].

B. Reconfigurable Intelligent Surfaces

An RIS is an engineered surface to digitally control reflections of radio waves, enabling *smart radio environments*. It is worth noting that RISs are likely to become pervasive, as they hold the potential to complement future wireless networks such as 6G [25, 12, 63]. Here, the propagation medium is considered as a degree of freedom to optimize wireless communication by redirecting radio waves in certain directions [34], *e.g.*, to improve signal coverage and eliminate dead zones, to enhance energy efficiency and data throughput [36], and building low-complexity base stations [5].

An RIS does not actively generate its own signals but passively reflects existing ambient signals. For this, it utilizes L identical unit-cell reflector elements arranged on a planar surface, as shown in Figure 1. Importantly, the reflection coefficient of each reflector is separately tunable to shift the reflection phase. Typically, an RIS is realized as a printed circuit board (PCB) with printed microstrip reflectors, enabling very low-cost implementation. To reduce complexity, many RISs use $1\,\mathrm{bit}$ control [48], *i.e.*, to select between two reflection

phases $0°$ and $180°$, corresponding to the reflection coefficients $+1$ and -1. This allows the control circuitry to directly interface with digital logic signals from, *e.g.*, a microcontroller. The technology is still under development [48] which is why RISs are currently not widely used in practice. At the time of writing, first implementations are being made commercially available [20, 40] and field trials are being carried out [42].

1) Finding RIS configurations: To realize a desired RIS reflection behavior between a sender and a receiver, an appropriate RIS configuration is required that matches the radio environment. For example, to maximize signal power at a receiver, the RIS configuration is used to make all signal components traveling via the RIS add coherently with other non-RIS signal components. Usually, such an RIS configuration cannot be be blindly synthesized due to the complexity of scene-dependent and hard-to-predict radio wave propagation effects in conjunction with the vast configuration space of the RIS. That is, an L-element binary-tunable RIS has 2^L possible configurations. Therefore, RIS configurations are often determined based on iterative optimization algorithms, involving measurement feedback to assess how a particular RIS configuration influences the wireless channel [17, 75, 65, 27, 42].

III. Preliminaries

A. Threat Model

1) Attack Scenario: We consider a typical wireless network scenario where a number of wireless devices are deployed and connected to an access point (AP). At least temporarily, the devices are stationary and do not change location. We assume the devices communicate with the AP using Wi-Fi, but the following analysis holds for any time-division duplex (TDD) communication protocol. Additionally, we assume reciprocity of wireless channels, meaning that for a pair of devices, the same radio propagation effects occur, regardless of the communication direction. Finally, we assume that the wireless signals are subject to multipath propagation due to typical propagation phenomena, *e.g.*, reflection and scattering, as commonly found in indoor and urban environments.

2) Attacker Model: We consider a physical-layer wireless jamming attacker who generates radio interference with the goal of disrupting the wireless communication of a set of victim devices. The attacker aims to perform selective jamming, meaning they aim to disrupt only a subset of devices while leaving others unaffected.

The attacker is capable of transmitting and receiving radio signals towards and from considered devices, *e.g.*, by using an ordinary radio transceiver comparable to the hardware of the considered devices. We assume that the attacker utilizes a single antenna to either transmit or receive signals. Additionally, the attacker employs an RIS next to their antenna, which they can configure arbitrarily.

The attacker is external to the wireless network of the considered devices and cannot read encrypted payload information. However, the attacker can estimate the received signal strength indicator (RSSI) and distinguish signals originating from different devices. Finally, the attacker can choose an arbitrary position to launch their attack without knowing the exact location of the considered devices.

TABLE I: Terminology Overview

Symbol	Description
$\mathcal{D}$	Set of all devices
D_i	i^{th} device in $\mathcal{D}$
N	Number of total devices in $\mathcal{D}$
$\mathcal{T}$	Subset of target devices
$\mathcal{N}$	Subset of non-target devices
L	Number of IRS elements
c	RIS configuration vector
c_l	Reflection coefficient of the l^{th} RIS element
H_R^T	Channel gain from sender T to receiver R
$H_R^T(c)$	RIS-controlled channel gain
$h_l^{D_i}$	l^{th} RIS sub-channel to device D_i
X_T	Signal from a sender T
J	Jamming signal from the attacker
W	White Gaussian noise

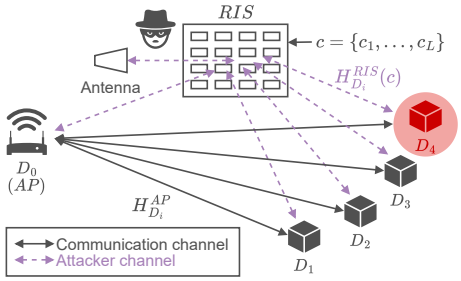

Fig. 2: Illustration of the assumed system model.

B. System Model

In this section, we establish the system model to formally describe the attack scenario and the involved parties from a signal perspective. For the reader's convenience, we summarize the used symbols in Table I.

We assume that the attacker faces a total of N wireless devices from the set $\mathcal{D} = \{D_0, \ldots, D_{N-1}\}$, *e.g.*, forming a wireless network where one device is an AP that the others connect to. The devices $\{D_1, \ldots, D_{N-1}\}$ seek to extract correct digital information from the legitimate signal X_{AP} received from the AP D_0. The attacker seeks to disrupt the wireless communication of K devices, forming the subset $\mathcal{T} \subseteq \mathcal{D}$, while leaving the remaining devices in the subset $\mathcal{N} = \mathcal{D} \setminus \mathcal{T}$ unaffected. Figure 2 illustrates an exemplary scenario with five devices $\mathcal{D} = \{D_0,\ D_1,\ D_2,\ D_3,\ D_4\}$, where the attacker would like to jam $\mathcal{T} = \{D_4\}$, while keeping the remaining devices $\mathcal{N} = \{D_0,\ D_1,\ D_2,\ D_3\}$ operational.

To achieve their goal, the attacker transmits a jamming signal J to overshadow the legitimate signal from the AP X_{AP}. Both signals X_{AP} and J are subject to radio propagation effects, described by the complex channel gains between the respective transmitter and the device D_i, denoted as $H_{D_i}^{AP}$ from the AP to D_i, and $H_{D_i}^{RIS}(c)$ from the attacker to D_i. Thus, the device D_i observes the total received signal

$$Y_{D_i} = H_{D_i}^{AP} X_{AP} + H_{D_i}^{RIS}(c)\ J + W, \tag{1}$$

where W is additive white Gaussian noise. Note that the attacker's channel is reconfigurable by means of the RIS

configuration vector c. In line with the literature [5], we model this channel as the superposition of L sub-channels $h_l^{D_i}$ between the attacker's antenna and the device D_i via the l^{th} RIS element, each of which is multiplied with the selected reflection coefficient c_l of the respective RIS reflector element:

$$H_{D_i}^{RIS}(c) = \sum_{l=1}^{L} h_l^{D_i}\, c_l. \tag{2}$$

From Equation 1, we can formulate the jamming-to-signal ratio (JSR) of each device as

$$\mathrm{JSR_{D_i}} = \frac{|H_{D_i}^{RIS}(c)J|^2}{|H_{D_i}^{AP} X_{AP}|^2}, \tag{3}$$

which is a key metric to assess the success of jamming attacks [45, 43]. With increasing JSR, the probability that the respective radio receiver will be disturbed increases.

IV. RIS-based Selective Jamming Attack Strategy

With the established system model in mind, we now proceed to elaborate the attacker's strategy in order to meet their two principal goals: (i) rendering target devices inoperative while (ii) keeping non-target devices operational. For this, the attacker must maximize $\mathrm{JSR_{D_i}}$ for the target devices and minimize it for the non-target devices. The classical approach to meet goal (i) is to increase the power of the jamming signal J. However, this strategy does not address goal (ii) and carries the risk of also jamming non-target devices. In this work, we resolve this issue by means of an RIS, leveraging for the first time RIS-based wireless channel control to optimize an active jamming attack. In particular, the attacker leverages the RIS configuration c to adapt their wireless channel gains $H_{D_i}^{RIS}(c)$ and control the delivery of J towards each device. In other words, by applying an appropriate configuration c to their RIS, the attacker can selectively increase or decrease the channel gains towards the considered devices in order to control the respective JSR and thereby control the effect of the jamming. Thus, the attacker faces the following multivariate optimization problem:

$$\max_{c} |H_{d\in\mathcal{T}}^{RIS}(c)|, \tag{4}$$

$$\min_{c} |H_{d\in\mathcal{N}}^{RIS}(c)|. \tag{5}$$

To find an appropriate c that meets these goals, the attacker must observe $H_{RIS}^{D_i}(c)$ (see Section II-B) – ideally by measuring the jamming signal strength arriving at each considered device. However, this clearly is not possible as the devices $\mathcal{D}$ do not cooperate with the attacker. To solve this, we leverage *channel reciprocity*, where the wireless channels from the attacker to the considered devices and vice versa are identical, *i.e.*, it holds that $H_{RIS}^{D_i}(c) \approx H_{D_i}^{RIS}(c)$ [62]. Consequently, to assess whether a particular RIS configuration c meets the channel optimization, the attacker can eavesdrop on the considered devices and measure $H_{RIS}^{D_i}(c)$.

In summary, the RIS-based jamming attack is a two-step procedure as illustrated in Figure 3. First, the attacker passively determines an appropriate RIS configuration by eavesdropping the radio communication signals from the considered devices (see Figure 3a). Using this configuration, the attacker then actively transmits the jamming signal J which disrupts target devices while non-target devices remain operational (see Figure 3b).

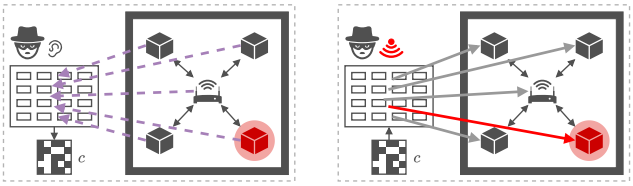

(a) (b)

Fig. 3: Two-step attack strategy of the jamming attack using the RIS. (a) Step 1: Passive channel optimization. (b) Step 2: Active wireless jamming attack.

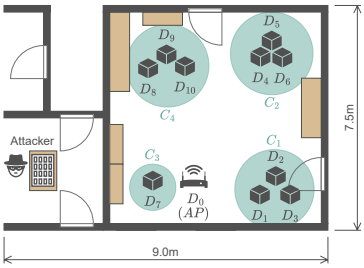

Fig. 4: Floorplan of the environment used for experiments, indicating the positions of all wireless devices and the attacker.

V. Experimental Setup

1) Wireless Environment: We conduct our experiments in an ordinary office environment where we make use of an area of approximate size $9.0\,\mathrm{m} \times 7.5\,\mathrm{m}$. A floor plan is depicted in Figure 4, indicating the position of the attacker and the wireless devices $\{D_0, \ldots, D_{10}\}$. The devices are arranged in four clusters spread across the room.

To demonstrate and evaluate the attack, we utilize commodity off-the-shelf Wi-Fi devices. For the AP device D_0, we utilize a TP-Link N750 router running OpenWrt to provide an IEEE 802.11n Wi-Fi network with 20 MHz bandwidth on channel 112, corresponding to a channel center frequency of 5560 MHz. The router is capable of packet injection using lorcon [69]. For the wireless devices $\{D_1, \ldots, D_{10}\}$, we use ten Raspberry Pi 4 Model B and leverage nexmon [53] to optionally put their Wi-Fi chipset into monitor mode.

The devices D_1 to D_{10} ping the AP (D_0) to trigger wireless traffic. Per default, the AP is part of the set of non-target devices $\mathcal{N}$.

2) Attacker Setup: On the attacker side, we use the following hardware setup. To realize eavesdropping, we employ a Raspberry Pi 4 Model B. Again, we use nexmon [53] to put the Wi-Fi chipset into monitor mode, allowing us to obtain

the RSSI value and MAC address for each received Wi-Fi packet, even in the case of frame errors. In order to connect an external antenna to the Raspberry Pi, we disconnected the onboard PCB antenna and added a coaxial connector.

For the active jamming, we utilize standard IEEE 802.11n Wi-Fi signals (20 MHz bandwidth with modulation and coding scheme (MCS) set to 1), containing 25 randomized payload bytes. For convenient signal generation, we use a Signal Hound VSG60 vector signal generator, allowing to transmit signals with a 100 % duty cycle and precisely controlled signal power. However, we stress that jamming signal generation can likewise be realized with ordinary Wi-Fi devices [52].

We employ an interline PANEL 14 directional antenna and use a Mini-Circuits USB-2SP4T-63H radio frequency (RF) switch to either connect the antenna to the Raspberry Pi for eavesdropping or to the signal generator for jamming. The antenna is directed towards the attacker's RIS. Figure 21 in the Appendix A shows a photo of the setup. The RIS is based on the open-source design of Heinrichs *et al.* [24] and consists of three standard FR4 PCBs. It has $L = 768$ unit-cell reflector elements with binary phase control, optimized to operate in the 5 GHz Wi-Fi frequency range. The elements can be programmed via USB to select the phase of c_l to either be $0°$ (state '0') or $180°$ (state '1'). For further technical details, we refer to [24].

3) RIS Optimization: To determine an RIS configuration that solves the optimization problem formulated in Equation 4 and Equation 5, we employ the greedy genetic optimization algorithm put forward by Tewes *et al.* [65]. The algorithm stores a sorted table, where a cost function f is evaluated for a set of B initially random RIS configurations. The cost function first aggregates the RSSI values from the devices in the respective set, using weighted combinations of the mean and minimum for $\mathcal{T}$, and mean and maximum for $\mathcal{N}$. We weight the mean with 0.3 and the extreme values with 0.7, to emphasize the worst-performing devices in the respective sets stronger during optimization. Finally, we take the signed squared difference of both aggregate values as the result of f. Based on RIS-element-wise empirical probabilities for maximizing the cost function within the table, a new RIS configuration is generated and evaluated to update the table with every algorithm step. For our experiments, we set B to 100 and run the algorithm for 10 000 steps.

4) Evaluation Metrics: In the remainder of the paper, we use the following evaluation metrics:

- *RSSI values*: The Raspberry Pis we use provide estimates of the received signal strength in dBm with a resolution of 1 dB for every received packet. We use RSSI for channel measurement, *e.g.*, to estimate $|H_{RIS}^{D_i}(c)|$ as needed to optimize the attacker's RIS.
- *JSR*: We evaluate $\mathrm{JSR}_{\mathrm{D_i}}$, *cf.* Equation 3 for each device, using RSSI values corresponding to signals received by the devices $\{D_1, \ldots, D_{10}\}$ from the attacker and the AP.
- *Packet rates*: For evaluation of the attacker's selective jamming capabilities, we leverage packet injection on the AP D_0 to transmit Wi-Fi packets with MCS 6 at a constant rate of 100 packets per second. On the devices $\{D_1, \ldots, D_{10}\}$, we measure the successfully received Wi-Fi packets from the AP per second. If a particular device is affected by the attacker's jamming signal J, the received packet rate will be reduced. For the measurements, the devices operate in monitor mode, granting a clear view on the bare jamming effects, independent of adaptive Wi-Fi mechanisms such as rate control, retransmissions, or even disconnections.
- *Data throughput*: We also evaluate the effect of the selective jamming in a standard Wi-Fi network, *i.e.*, where devices do not use monitor mode but operate as a station. Here, $\{D_1, \ldots, D_{10}\}$ are connected to the AP D_0 and we assess the effect of the attacker's jamming signal J by measuring the data throughput from the AP towards a particular device in Mbit/s by means of *iperf3* [16].

VI. Attack Evaluation

After introducing the attack strategy and our experimental setup, we evaluate RIS-based spatially selective jamming attacks in several real-world scenarios. First, we investigate selective jamming of a single device to demonstrate the scheme's feasibility. Then, we target multiple devices to show its scalability. We also assess the robustness of the RIS optimization against environmental changes. In addition, we validate the attack's effectiveness in a fully-fledged Wi-Fi network and study devices in extreme proximity. Finally, we compare the performance of spatially selective jamming using a directional antenna versus the RIS.

A. Single-Target Jamming

The first scenario we evaluate is jamming of a single target device. For this, we first optimize the attacker's RIS and then transmit the jamming signal while evaluating how it affects each considered device $\{D_1, \ldots, D_{10}\}$. Serving as a blueprint for the subsequent scenarios, we provide a detailed outline of the attacker's action, covering RIS optimization, active jamming and JSR analysis.

1) RIS Optimization and Active Jamming: The attacker's first step is to find an appropriate configuration for the RIS to maximize the received RSSI from the targeted device while minimizing it for all other devices. For this, the attacker uses their Raspberry Pi to eavesdrop on the signals transmitted by the considered devices and estimates the magnitude of the channels $H_{RIS}^{D_i}(c)$ from the obtained RSSI values. Once RSSI values from all devices have been opportunistically collected, the attacker can perform a step of the RIS optimization algorithm. In our experiments, the algorithm runtime for 10 000 steps is about 5 minutes.

In our initial experiment, we want to jam device D_1 and thus first need to find an RIS configuration that targets device D_1 while excluding the remaining devices. Figure 5a depicts the RIS optimization process, clearly showing an improvement in the channel quality between the RIS and the target device D_1 while it degrades for the remaining devices as the optimization algorithm progresses. Eventually, the RSSI for D_1 converges

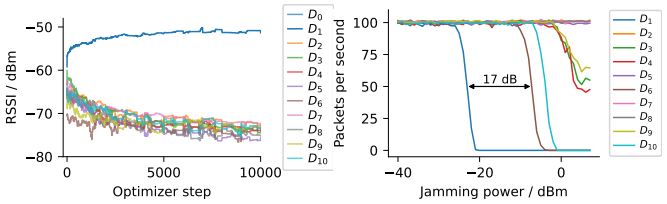

(a) (b)

Fig. 5: (a) RIS optimization process for targeting device D_1 while excluding all other devices. (b) Measurement of packet reception rates on the considered devices over jamming signal power using the previously optimized RIS configuration.

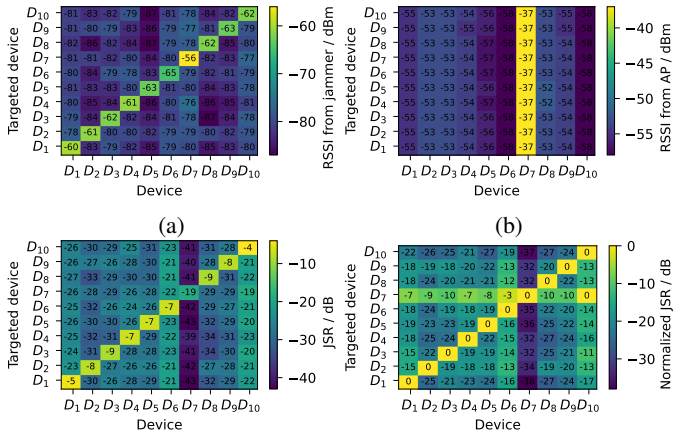

(a) (b) (c) (d)

Fig. 6: (a) RSSI from the jammer as observed for each device. (b) RSSI from the AP. (c) JSR of the RSSI from the jammer and the AP. (d) Normalization of the JSR.

to $-50\,\mathrm{dBm}$, while the RSSI for the non-targeted devices reach levels around $-75\,\mathrm{dBm}$, confirming the effectiveness of the optimization algorithm.

Using the RIS configuration resulting from the optimization algorithm, the attacker switches from eavesdropping to actively sending a jamming signal. Figure 5b shows the packet reception rates from the AP for each device over the jamming signal power. Here, we can see that jamming with a signal power of $-21\,\mathrm{dBm}$ completely disrupts the reception of D_1, while the reception rates of all other devices are unaffected. Moreover, the attacker has a jamming signal power margin of $17\,\mathrm{dB}$ until any other device (D_6) is disrupted.

2) Jamming Success of the Attacker: Thus far, we have demonstrated jamming of D_1 without affecting the other devices. However, the attacker can likewise leverage the RIS optimization to target any other device in the environment. We repeat the previously outlined optimization process for the remaining devices $\{D_2, \ldots, D_{10}\}$ and investigate the overall attack success by studying the JSR.

As discussed in Section IV, the JSR assesses the attacker's success, relating the jamming power at a specific device to the desired signal power from the AP. To evaluate the JSR, we measure the RSSI values from the jammer and the AP (both

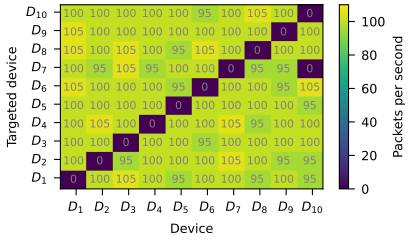

Fig. 7: Packet receive rate for each device and for each single-target configuration D_1 to D_{10}.

sending Wi-Fi signals with constant transmission power) on each device for each optimized RIS configuration. Figure 6a and Figure 6b present the results, visualizing the effect of the respective RIS configuration (per column) as observed on each device (per row). The distinct diagonal entries in Figure 6a show that the attacker succeeds in focusing their signals on the intended devices while achieving rejection towards others, while Figure 6b depicts each device's signal reception from the AP. Here, the RIS does not affect the channels $H_{D_i}^{AP}$ between the AP and the devices. Additionally, we observe the effects of distance-dependent path loss, as the devices with the smallest and largest distance to the jammer and AP (D_7 and D_6) experience the highest and lowest signal strengths, respectively. Similarly, D_7 receives signals from the AP strongest as it is only $1\,\mathrm{m}$ apart.

Using the RSSI measurements from both the jammer and the AP, we derive the JSRs by taking their difference (since RSSI values are logarithmic, this follows Equation 3). The result is shown in Figure 6c. The key observation here is that the JSR values of the targeted devices on the diagonal entries stand out as desired. Furthermore, the variation of the per-target JSR indicates that jamming of, *e.g.*, D_{10} is more efficient than jamming D_7, which has a robust legitimate signal reception.

To highlight the attack effect on non-target devices, we additionally show the row-wise normalized JSR in Figure 6d. This highlights the different legitimate channel conditions: D_7 consistently shows the lowest non-target JSR due to its strong signal reception from the AP, whereas D_6 and D_{10} exhibit higher non-target JSRs because they receive weaker signals from the AP. Finally, the higher relative jamming signal power required for D_7 reduces the signal rejection margins towards the other devices such that the attacker fails to exclude D_{10}. Nonetheless, the attacker achieves JSR reductions of at least $20\,\mathrm{dB}$ in more than $50\,\%$ of the cases ($16\,\mathrm{dB}$ in $90\,\%$ and $24\,\mathrm{dB}$ in $25\,\%$). We refer the reader to Appendix C for an evaluation of the normalized JSR during the RIS optimization.

Now, while using just enough jamming signal power to disrupt the respective target device, we perform the active jamming attack against each device. Following the same evaluation rationale as before, Figure 7 shows the packet reception rates of each device while sending the jamming signal with the corresponding RIS configuration applied. After demonstrating selective jamming of D_1 before, the clearly visible diagonal entries show that the attacker is capable of

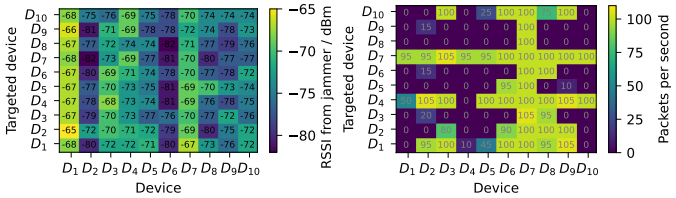

(a) (b)

Fig. 8: Attack performance with randomized RIS configurations. (a) Received RSSI from the attacker for each device. (b) Resulting packet reception rates during jamming.

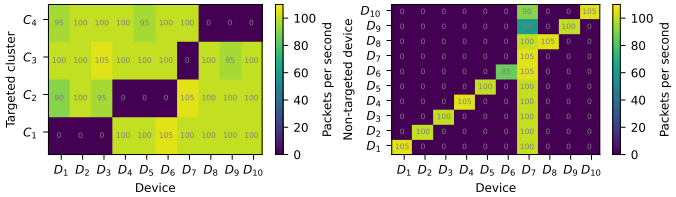

(a) (b)

Fig. 9: (a) Packet reception rates for jamming the four device clusters (see Figure 4). (b) Packet reception rates for targeting all but one device.

selectively jamming all considered devices without affecting the other devices. However, we also recognize the effect of insufficient jamming signal rejection, as discussed before. That is, when jamming D_7, the attacker also disrupts the packet reception of D_{10}, being in line with the previous JSR result shown in Figure 6d.

Overall, our results demonstrate that RIS-based jamming enables precise physical attack targeting without affecting neighboring devices, even through-the-wall from another room. Please note the reported attack performance clearly stems from the attacker's optimized RIS. That is, when using random configurations for the RIS, neither the RSSI from the attacker nor the packet reception rates under jamming are concentrated on a particular device as evident from Figure 8.

B. Multi-Target Jamming

The previous results highlight the effectiveness of RIS-based spatially selective jamming in single-target scenarios. In the following, we extend this scenario and investigate two additional scenarios where the attacker wants to jam not just one, but multiple devices simultaneously.

1) Device Clusters: In the first multi-device scenario, the attacker seeks to disrupt the device clusters depicted in Figure 4. Thus, we repeat the previous experiments but this time specify the devices belonging to the clusters C_1 to C_4 as target devices during the RIS optimization. Subsequently using the resulting four RIS configurations for active jamming, Figure 9a shows the packet receive rates of all devices, analogous to Figure 7. Here, we can see that attacker succeeds to disrupt the devices belonging to the respective clusters, while the remaining devices again remain fully operational. Note that the RIS configuration to target cluster C_3 (comprising only of device D_7) this time sufficiently reduces the jamming signal towards device D_{10}, preventing the unintended non-target jamming previously observed in Figure 7. That is because the greedy optimization algorithm does not necessarily always converge to the same RIS configuration, given that the algorithm is randomly initialized and guided based on noisy measurements. Still, the overall conclusion from this experiment is that the attack approach is extensible to selectively jam even multiple devices simultaneously.

2) Single Exclusion: After jamming multiple devices, we now aim to jam *every* device except one the attacker would like to keep operational. Building on the previous insight that the attacker can leverage the RIS to deliver the jamming signal to multiple devices, we now seek to push this approach even further to jam *every* device except one that the attacker would like to keep operational. Thus, we repeat the RIS optimization, but now specify $\mathcal{T} = \mathcal{D} \setminus \{D_0, D_i\}$ where $i \geq 1$. Like before, we then perform active jamming with each optimized RIS configuration and plot the packet reception rates of all devices in Figure 9b. Here, quite opposite to the single-targeting scenario of Figure 7, we can see that the attacker indeed succeeds to disrupt all receivers except D_7 while keeping one non-targeted device operational. This experiment confirms that jamming of a broader set of devices is possible.

The inability to jam D_7 is due to its strong reception of the legitimate AP signal *cf.* Figure 6b. Successful jamming requires the attacker to deliver sufficient signal power to D_7 to overshadow the legitimate signal. However, in the present scenario, the attacker splits their jamming power among nine targets, reducing the power efficiency. Thus, we expect the jamming power at each device to be reduced by a factor $1/9 \mathrel{\widehat{=}} -10\log_{10}(9) \approx 9.5$ dB. This matches our observations, as the jamming power arriving at D_7 was $-56\,\mathrm{dBm}$ in the single-target scenario, *cf.* Figure 6a, while in the current experiment, it was at most $-65\,\mathrm{dBm}$. Therefore, the attacker lacks sufficient jamming power to disrupt D_7. To address this, the attacker could amplify their jamming signal or prioritize channel maximization towards D_7 during RIS optimization.

C. Effect of Environmental Variation

The attacker relies on the assumption that the optimized RIS configuration is valid during the subsequent active jamming. In the following, we investigate the validity of this assumption, evaluating the robustness of the proposed scheme against environmental variation.

1) Long-Term Stability: For our previous experiments, the attacker's RIS was optimized shortly before obtaining JSR and jamming packet reception rate results. However, it is not yet clear whether we can expect the RIS optimization to be long-term stable, which would be a desirable property for attack practicality. To investigate this aspect, we perform RIS optimization with $\mathcal{T} = \{D_1\}$ and monitor the resulting JSR values of each device for a duration of 24 hours. Figure 10 shows the JSRs normalized to the initial JSR measurement of D_1 as a time series. The key observation is that the

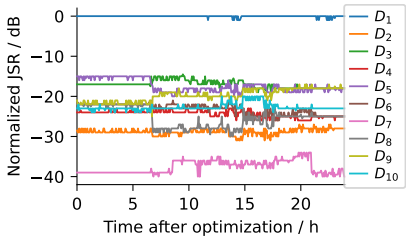

Fig. 10: Time stability of optimized JSR values over 24 h.

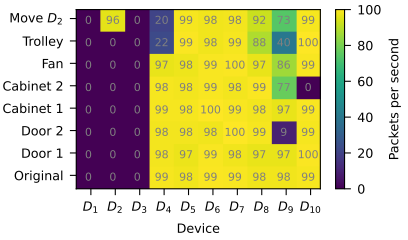

Fig. 12: Packet reception rates during jamming with incremental environmental changes.

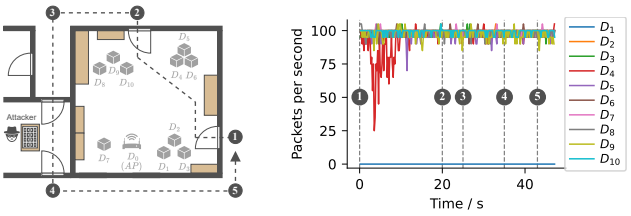

(a) (b)

Fig. 11: Effects of human motion in the experimental environment. (a) Path layout for the motion including checkpoints. (b) Measured time series for each device.

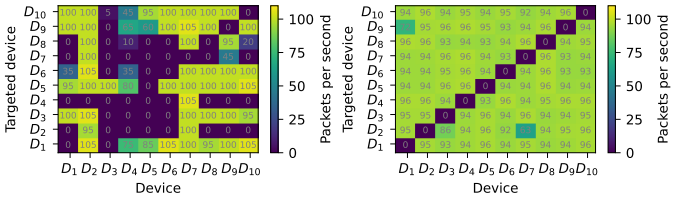

(a) (b)

Fig. 13: Effect of device-repositioning: Selective jamming performance after rearranging the devices in a grid across the room (a) before and (b) after rerunning RIS optimization.

JSRs remain largely stable with sustained focusing of D_1 and rejection of the other devices, showing that the optimized RIS configuration stays valid after an entire day has passed. However, we do observe some variation in the JSRs of the non-target devices, starting 7 hours after the RIS optimization at midnight. Since our experiments took place in an ordinary, actively used office environment, we attribute this to office activity, *e.g.* individuals walking around. We investigate the effect of human activity on the attacker's jamming performance more systematically in the next experiment.

2) Human Motion: In this experiment, we again perform single-target jamming with $\mathcal{T} = \{D_1\}$ while an individual walks through the environment, passing by the considered devices and the attacker's RIS. At the same time, we record the packet receive rate of each device for 45 s. Figure 11a shows the walking path with annotations matching the receive rate time series shown in Figure 11b. The first thing to note is that the reception of D_1 remains completely suspended as desired, regardless of the individual. Likewise, the reception of the non-targeted devices is mostly unaffected at approx. 100 packets per second. Still, we can see that when the individual is within the room, the device D_4 is temporarily affected by the jamming as evident from the reduced packet reception. That is, the individual potentially affects the wireless channels $H_{D_4}^{RIS}(c)$ and $H_{D_4}^{AP}$, which caused an increased JSR.

3) Changes to the Environment: Next, we perform multi-target jamming with $\mathcal{T} = \{D_1, D_2, D_3\}$ (cluster C_1) and incrementally change the room where the devices are located. In particular, we open its two doors, introduce additional items (two $60 \times 60 \times 43$ cm rolling office cabinets, a pedestal standing fan, and a 60×60 cm flat platform trolley), and finally move the device D_2 by two centimeters. Figure 12 shows the packet receive rates of all devices for each environmental state. Without changes in the environment (labeled as 'Original'), the attacker achieves their goal to jam the target devices while the others remain unaffected. Then, as we introduce more changes to the environment, we can see that especially the non-target devices D_4, D_8, and D_9 become more affected by the jamming, indicating an increased JSR. As in the previous experiment, jamming of the targeted devices is largely robust against environmental variation. However, when finally moving the targeted device D_2, it is no longer disrupted.

4) Different Device Positions: In the previous experiment, we have seen that relocating the device D_2 caused jamming of that device to become ineffective. This observation underscores the attack's desired dependency on the device location. To further study the effect of relocating devices, we change the device positioning in clusters (see Figure 4) to a grid, uniformly distributing the device across the room. Then, we repeat the single-target experiment outlined in Section VI-A with the originally optimized RIS configurations. The results are shown in Figure 13a. We now see that the attacker – using outdated RIS configurations – fails to selectively jam the respective devices. However, after renewing the RIS configurations, the attacker is again capable of selectively jamming each device, as can be seen in Figure 13b. Apart from the clear spatial dependence of the attack, this result also highlights the attacker's ability to adapt to different scenes.

D. Attack Performance in a Wi-Fi Network

Thus far, we studied RIS-based selective jamming by means of the JSR and packet reception rates. For this, the devices operated in monitor mode while observing packets with a fixed MCS setting. This allowed us to evaluate the attack

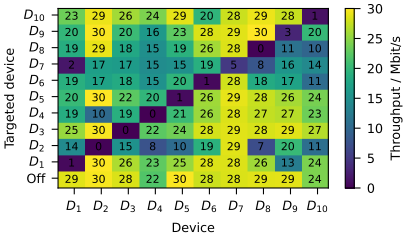

Fig. 14: Wi-Fi throughput for each device and for each single-target configuration D_1 to D_{10}.

performance independently of the behavior of custom rate adaption algorithms employed in fully operational Wi-Fi networks. Here, the physical-layer transmission speed is not fixed but is adapted to the wireless channel conditions [71, 32]. The sender monitors whether transmissions with higher data rates were successful and otherwise reduces the data rate which in turn increases the link robustness. This is controlled via the MCS setting [67], denoting a standardized combination of modulation scheme and error coding rate.

In the context of jamming, rate adaption can be viewed as a countermeasure where the victim party adaptively enhances their jamming resiliency. Switching to a lower MCS setting allows the victim to cope with reduced link quality or, put differently, with a higher JSR. For example, for single-antenna communication over $20\,\mathrm{MHz}$ bandwidth, Wi-Fi receiver sensitivity improves by approx. $18\,\mathrm{dB}$ when switching from MCS 7 to MCS 0 [44], yet reducing data rate by a factor of 10 [67].

To assess whether Wi-Fi rate adaption could diminish the attacker's selective jamming success and to further investigate the attack's real-world potential, we repeat the single-target experiment from Section VI-A. However, we measure the data throughput on each device $\{D_1, \ldots, D_{10}\}$ in the Wi-Fi network of the AP. In particular, we use iperf3 to transfer a $30\,\mathrm{Mbit/s}$ UDP datastream towards each device while the attacker transmits their jamming signal to disrupt one device. Figure 14 shows the resulting throughput measurements and additionally indicates the throughput without an attack in the first row. Importantly, we again observe clearly distinct diagonal entries, showing that the throughput on the targeted devices is reduced to (nearly) $0\,\mathrm{MBit/s}$.

Other than in Figure 7, non-target devices are slightly affected by the jamming. The reason for this is the rate adaption of the targeted device which reacts to the jamming by switching to a more robust MCS that can withstand a higher JSR. In turn, to achieve a JSR that disrupts the target reception, the attacker must expend more jamming power. However, this may exhaust the JSR reduction provided by RIS optimization, causing a non-target device to be affected and switching to a more robust MCS, sacrificing some data rate. Still, the throughput of the non-target devices remains above $22\,\mathrm{Mbit/s}$ in $50\,\%$ of the cases ($11\,\mathrm{Mbit/s}$ in $90\,\%$ and $27.5\,\mathrm{Mbit/s}$ in $25\,\%$).

Consequently, while the attacker does not completely prevent the non-target devices from being affected, the observed throughput degradation is within acceptable limits and can likely be further improved by refining the RIS optimization. The results demonstrate the feasibility of the attack even when the victim devices use adaptive rate control, posing a significant threat to real-world Wi-Fi networks.

E. A Detailed Look in the Spatial Domain

Thus far we have investigated RIS-based spatially selective jamming of devices distributed across an entire room. Now, we seek to explore the attack on a smaller scale, *i.e.*, when devices are in close proximity.

1) Selective Jamming of Devices at Sub-Wavelength Distance: Previously, we have seen that it is possible to jam single devices although being in close proximity to others within clusters, *cf.* Figure 4. Intuitively, due to the inherent spatial correlation of electromagnetic fields, one would expect that if one device is disrupted by the attacker's jamming, then another very close-by device would likewise be affected. However, interestingly, we found that RIS-based spatially selective jamming even works when device antenna separation is deeply in the sub-wavelength region.

In our experiment, we consider the devices D_5 and D_6 (and the AP D_0) which we place very close to each other in two geometrical configurations. In the first (see Figure 15a), we place the devices directly above each other to minimize their antenna distance while being in the same orientation. In the second (see Figure 15c), we place the devices facing each other to minimize their antenna distance regardless of the orientation, being approx. $5\,\mathrm{mm}$. For both scenarios, the attacker optimizes their RIS to (i) target D_5, (ii) target D_6, and (iii) target both. Like before, we use iperf3 to measure the Wi-Fi throughput on both devices while the attacker conducts their jamming attack.

The resulting throughput measurements for both device placements are shown as time series in Figure 15b and Figure 15d. The first thing to note is that without jamming, both devices initially have data rates of around $25\,\mathrm{Mbit/s}$. Then, as the attacker starts to jam with the first RIS configuration, the data rate of D_5 drops to zero while D_6 remains completely unaffected as evident from the unaltered throughput. The attacker then switches to the next RIS configuration, alternating the attack target. Now, the throughput of D_5 is restored to the level without the attack while the throughput of D_6 is close to zero. Finally, as attacker switches to the third RIS configuration, the throughput of both devices drops to (nearly) zero. Please note that – after activating the jamming – the attacker achieves this result by merely reconfiguring the RIS configuration. Importantly, this experiment highlights the attacker's ability to dynamically change the targeted device.

2) Further Analysis of Close-By Antennas: We now investigate the mechanics behind the previous result. First of all, theoretical limits of separating wireless channels in the spatial domain are rooted in the correlation of multipath components at different locations, *i.e.*, the correlations of the L subchannels via the RIS from the attacker antenna towards D_5 and D_6, $h_l^{D_5}$ and $h_l^{D_6}$. As shown by Clarke [13], the correlation

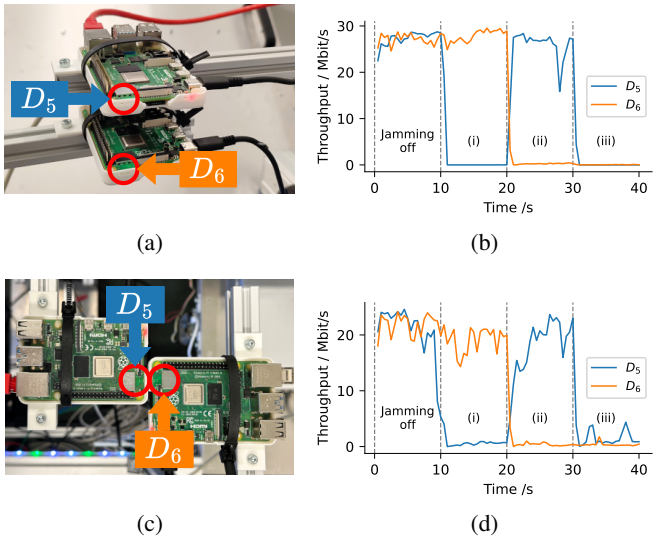

Fig. 15: Real-world spatial jamming attack demonstration against Wi-Fi communication. (a) Parallel aligned antennas. (b) Measured Wi-Fi datarates for parallel antennas. (c) Facing antennas. (d) Measured Wi-Fi datarates for facing antennas.

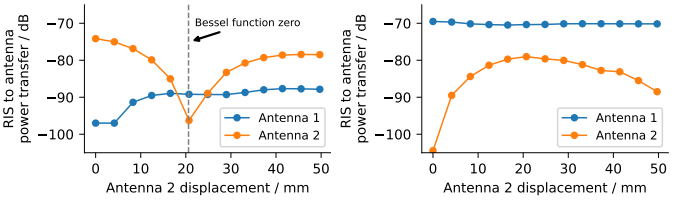

Fig. 16: Channel effects when repositioning an antenna in close proximity of another one. RIS optimized with Antenna 2 at 0 mm displacement to (a) minimize and maximize and (b) maximize and minimize the respective channels.

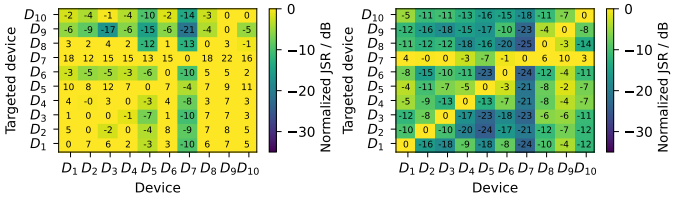

Fig. 17: Normalized JSR for single-target jamming with other devices being hidden before (a) and after (b) optimization.

as a function of the distance d can be described using the Bessel function of the first kind. Given the attack scenario where one channel is maximized, the smallest distance from the maximized location to a minimum could be approximated as the first zero point of the Bessel function, given by $2.4048\frac{c}{2\pi f} \approx 20.6\,\mathrm{mm}$, where c is the speed of light and f is the signal frequency ($5560\,\mathrm{MHz}$ in our experiments). However, this number is considerably higher than the $5\,\mathrm{mm}$ device separation from Figure 15c. One reason for this is because two antennas in real-world scenarios will seldom exhibit the exact same radiation patterns. This effect is due to differences in relative orientation (such as in Figure 15c) as well as differences in the relative environment, *e.g.*, objects in the antenna nearfield. Notably, the latter also includes the case that antennas get into each other's proximity: So called mutual coupling effects distort the individual antenna radiation patterns and therefore reduce spatial correlation effects [60, 15]. In consequence, the aforementioned effects would allow the attacker, for example, to exploit that one device's antenna might exhibit a high sensitivity in an angular direction where the other does not.

To assess the influence of the device antennas, we conducted additional experiments using a Keysight P9372A vector network analyzer (VNA) for high-accuracy wireless channel measurements. At the position of D_8, we place two antennas of the same type ('Antenna 1' and 'Antenna 2') directly next to each other and measure the wireless channels between them and the attacker's antenna (via the RIS). We optimize the RIS to maximize the channel towards Antenna 1 and minimize the channel towards Antenna 2 and vice versa. Then, for both cases, we move Antenna 2 in steps of $4\,\mathrm{mm}$ away from Antenna 1. To ensure accurate and repeatable antenna positioning, we use a 3D-printed positioning fixture. A photo of the setup is shown in Figure 22d in Appendix A.

Figure 16 shows the magnitude channel measurement results for both antennas over the tested displacements of Antenna 2. Here, we can see that the channel of Antenna 2 clearly decorrelates with its displacement. Notably, the measurement is in good agreement with theory, where the signal power is strongly reduced at the zero point of the Bessel function, as indicated by the dashed line Figure 16a. Furthermore, it is also evident that the measurement results of the fixed Antenna 1 are affected by moving Antenna 2. Given the stronger relative impact of small channel variations, this behavior is more pronounced when Antenna 1 is minimized (power increases by approx. $8\,\mathrm{dB}$) than when it is maximized (power reduced by approx. $0.5\,\mathrm{dB}$). Finally, these results confirm that the attacker can take advantage of antenna coupling effects to selectively target devices in close proximity configurations.

3) Jamming Impact on Hidden Devices: Thus far, we have focused on devices detectable via their wireless transmissions. However, some devices may remain passvie and only act as receivers without transmitting. Next, we examine the effects of attacks on these *hidden devices*, where the attacker cannot optimize their wireless channel. To align with our system model from Section III-B, we define a subset of devices, denoted as $\mathcal{H} \subseteq \mathcal{N}$, which is ignored for the optimization problem in Equation 5, effectively treating them as hidden. We replicated the single-target experiment from Section VI-A2 but consider $\mathcal{H} = \{D_1, \dots D_{10}\} \setminus \mathcal{T}$, leaving only the AP D_0 effective within the non-target devices $\mathcal{N}$. Thus, we optimize the RIS to maximize the channel of the respective target while

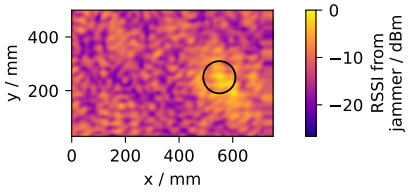

Fig. 18: Spatial distribution of normalized attacker signal RSSI values. During RIS optimization, the targeted device was placed within the black circle.

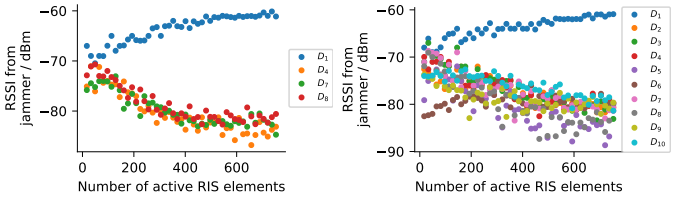

(a) (b)

Fig. 19: RIS optimization results over increasing number of active RIS elements for (a) $K = 5$ and (b) $K = 11$ devices.

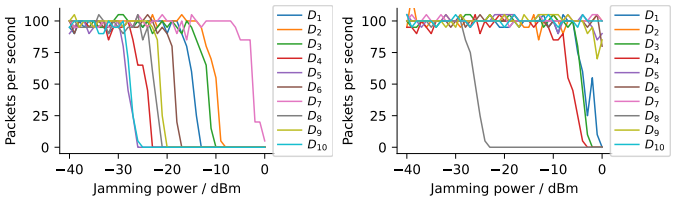

(a) (b)

Fig. 20: Comparison between a directional antenna and the RIS: (a) Directional antenna pointed towards D_8 and (b) RIS jamming, optimized towards with $\mathcal{T} = \{D_8\}$.

minimizing the channel of the AP and ignoring all others. Using the resulting RIS configurations, we then measure the JSR at each device. We show the normalized JSR with respect to each target device prior to and after the optimization process in Figure 17. Initially, most hidden devices experience a JSR at least as high as the respective target device. Crucially, after the optimization, we observe a JSR concentration on the diagonal, indicating that the attacker achieves selective jamming, despite the devices being hidden. The reason for this is that the channel towards the respective target is maximized while the channel towards the hidden devices is not, effectively reducing the jamming interference at the hidden devices. However, the JSRs at the hidden non-target devices are higher than when they are not hidden, *cf.* Figure 6, which is due to the lack of the additional minimization.

The previous result can be explained by the RIS maximization resulting in a focal point around the targeted device, as previously described by Kaina *et al.* [27]. This is different from classical beamforming, which rather affects an area and not a particular spot. To validate this, we mounted device D_5 on a precision dual-axis Cartesian robot, optimized the RIS with $\mathcal{T} = \{D_5\}$, and then measured the RSSI of the jamming signals at D_5 while re-locating the device. Figure 18 shows the resulting jamming signal distribution, measured in $10\,\mathrm{mm}$ steps within an area of size 75 cm $\times$ 50 cm around the initial device position at $(x, y) = (550, 250)$ mm, normalized to the initial position. At positions at least $6\,\mathrm{cm}$ away from the initial position (indicated by the black circle), the attacker signal power is at least 5 dB and on average 13 dB lower. The lesson from this experiment is that one can expect a jamming signal reduction at passive hidden devices, *i.e.*, without explicitly enforcing channel minimization during RIS optimization.

F. Further Evaluation of the RIS

1) Effect of Surface Size: A relevant factor for the attacker's ability to target and exclude devices is the physical size of the RIS [70], motivating the following experiment. To simulate variations in RIS size, we vary the number of active RIS elements. Specifically, we perform RIS optimization using a randomly selected subset of the total available RIS elements, while the other elements are configured random but remain fixed. In this way, we optimize the surface for a single-targeting scenario with $\mathcal{T} = \{D_1\}$ while varying the number of active RIS elements from 16 to 768. We plot the resulting attacker signal strength at each device over the number of active RIS elements used for optimization for $K = 5$ and $K = 11$ total devices in Figure 19.

Here, we see that the attacker fails to separate their RSSI values on targeted and non-targeted devices when using less than approx. 100 active elements during the RIS optimization. Another observation is that the resulting jammer power levels do not significantly improve beyond approx. 500 elements, potentially opening the door utilizing smaller RISs. However, the performance saturation might also be attributed to the particular parametrization of the greedy optimization algorithm, calling for further investigation. Nonetheless, a larger RIS improves the attacker's signal control.

Additionally, this experiment also highlights the effect of increasing the number of considered devices K. Arguably, as the attacker has to consider more devices, the complexity of the RIS channel optimization problem increases. Thus, the optimization algorithm must balance the performance among all devices in the corresponding sets $\mathcal{T}$ and $\mathcal{N}$, potentially sacrificing the optimization quality of one particular device in favor of another.

2) Comparison against Directional-Antenna Jamming: Alternatively to the RIS, the attacker could attempt to use a directional antenna pointed at a targeted device. To investigate whether such an approach could be effective, we perform another jamming experiment where the attacker points an elboxRF TetraAnt 5 19 20 RSLL directional antenna with $19\,\mathrm{dBi}$ gain towards the targeted device $\mathcal{T} = \{D_8\}$.

Figure 20 shows the packet receive rates of all devices over the jamming signal power both for the directional antenna and for the RIS. First of all, in both cases the attacker succeeds to disrupt the reception of D_8 completely. However, with

the directional antenna, the required jamming power is $3\,\mathrm{dB}$ higher than with the RIS. When using the directional antenna, the attacker also jams the non-target devices D_5, D_{10}, D_4, and D_9. Furthermore, the jamming power margin until affecting another device (D_6) is only $4\,\mathrm{dB}$. In clear contrast, with the RIS, the attacker succeeds to only jam the device D_8 while achieving a jamming power margin of $20\,\mathrm{dB}$ before another device (D_4) is affected.

This result shows that the RIS significantly outperforms the directional antenna when considering power efficiency and device selectivity. Directional antennas are designed to radiate a single beam under ideal (free-space) conditions, requiring mechanical adjustment while flexibility is limited: Target devices need to be within the beam width and non-target devices must be sufficiently separated. In contrast, the RIS allows fully-electronic scene adaption, combining many different propagation paths, allowing focusing and nulling of energy at one or multiple targets devices, regardless of whether they are very close or further apart from each other.

VII. Discussion

In this section, we discuss the experimental setup and our results, limitations of the attack, reason about potential countermeasures, and provide directions for future research.

A. Experimental Setup and Results

a) Wireless Devices: We designed the experimental setup to explore the RIS for wireless jamming attacks in a realistic scenario, yet made some simplifications to aid experimentation. While the device population comprises of off-the-shelf Wi-Fi devices located in an ordinary office environment, we made the devices regularly ping the AP to trigger wireless traffic required for the attacker's RIS optimization. In a real-world scenario, the attacker has to contend with the available traffic. However, to tackle potential low transmission rates, the attacker could provoke transmissions by injecting fake packets that devices often respond to with acknowledgements [1, 2]. Furthermore, to evaluate the physical-layer mechanism underlying the devised attack scheme while avoiding potential bias due to vendor-specific adaptive device behavior, we relied on monitor mode packet reception for several experiments. To this end, our Wi-Fi throughput measurements demonstrate that the scheme still works when the devices communicate within a fully-fledged Wi-Fi network.

b) Attacker Setup: All components for our attacker implementation are either available commercially or open-source, promoting reproducibility of the results. Regarding Wi-Fi signal reception and transmission, the attack can in principle be mounted merely using commodity wireless devices offering monitor mode and packet injection, *e.g.*, using low-level Wi-Fi chipset firmware control [53]. We utilized a dedicated external antenna with the attacker's Raspberry Pi to illuminate the RIS. However, in principle, the attack can be mounted with any antenna, provided the signal reaches the RIS.

The time for each RIS optimization step is governed by the time it takes until a signal from each device has been received and therefore depends on the packet transmission rate of Wi-Fi devives which often is higher than 100 packets per second [74]. In our setup, optimization of the RIS with $10\,000$ steps takes approx. 5 minutes to finish. When device positions remain static, the attacker is not limited by optimization time, yet to operate in more dynamic environment, optimization speed becomes more relevant. As discussed in Appendix C, the number of optimization steps can be substantially reduced without heavily sacrificing performance.

For the sake of experimental simplicity, we used device media access control (MAC) addresses to distinguish signals from different devices. A fully payload-agnostic physical-layer attacker would have to rely on physical-layer measures such as radio transmitter fingerprints [26, 56, 58, 14] for this.

While RISs are already commercially marketed [20], there is no broad consumer-level availability. However, the RIS we used is based on the well-documented open-source design by Heinrichs *et al.* [24] who also provide detailed manufacturing information. We estimate that the parts to manufacture the 768-element RIS we used can be purchased for approximately € 750. The way we leverage the RIS resembles an electronically tunable *reflectarray antenna* which was thus far exclusive to high-profile applications such as air surveillance radar [51] or satellites [28]. With the advent of RISs, such technology now becomes accessible even for individuals, extending the toolkit for advanced physical-layer attacks and necessitating a re-evaluation of attacker capabilities.

B. Countermeasures

Wireless jamming itself cannot be prevented due to the broadcast nature of radio wave propagation. Instead, jamming-resistant modulation schemes such as spread-spectrum techniques can be used to enhance robustness against jamming. However, changing the physical-layer modulation scheme is not an option for conventional standard-compliant wireless communication systems such as Wi-Fi. Therefore, in the context of this work, we discuss potential countermeasures geared towards hampering the attacker's RIS optimization.

a) MAC Address Randomization: As a raise-the-bar countermeasure, dynamic randomization of MAC addresses would make it more difficult for the attacker to associate received wireless signals with specific devices. However, the adversary could utilize payload-independent physical-layer properties such as radio frequency fingerprints [19, 56, 26] to distinguish radio signals from different devices.

b) Randomizing Transmit Power: The attacker's RIS optimization during attack preparation relies on RSSI values obtained from eavesdropped wireless signals. Thus, as an ad hoc countermeasure, a victim party could randomize their transmit power to hamper the RIS optimization. However, this would imply a reduced wireless communication quality of service. However, the attacker could also observe fine-grained channel state information (CSI) values which are not affected much by moderate changes of signal power.

c) Randomized Transmit Beamforming: Devices with multiple antennas could employ randomized transmit beam-

forming during their wireless communication. This would yield randomization of the channel $H_{RIS}^{D_i}(c)$ towards the eavesdropping attacker who uses $H_{RIS}^{D_i}(c)$ to optimize their RIS. In consequence, the attacker would be unable to assess whether channel changes stem from the RIS or the victim's transmit beamforming, hampering RIS optimization.

d) Avoiding Channel Reciprocity: A key mechanism underlying the attack is channel reciprocity, allowing the attacker to passively adapt their jamming channel before launching the active attack. Thus, to hamper attack preparation, reciprocal channels should be avoided, *e.g.*, by using sufficiently separated frequencies or antennas for reception and transmission.

e) Attack detection: Since the jamming signal is not fully suppressed at non-target locations, *cf.* our evaluation of attack effects towards hidden devices in Appendix VI-E3, passive wireless receivers can be used to detect the jamming signals, permanently monitoring the wireless environment, *e.g.*, to raise an alarm upon detecting malicious activity.

C. Limitations

We have shown that the RIS enables precise spatial control for targeted wireless jamming. However, as certain preconditions must be met for this, the attack also is subject to limitations. First of all, like in every other jamming attack, sufficient jamming signal power must be delivered to disturb the victim receiver. While our scheme offers fine-grained spatial jamming control at considered device locations, the jamming effect does not completely vanish at other locations.

To passively optimize the jamming channel towards the victim device, the attacker relies on passive eavesdropping and a reciprocal wireless channel. Thus, the attack does not work with wireless systems that rely on non-reciprocal wireless channels, *e.g.*, when transmission and reception employ different signal frequencies or antennas. To overcome this, the attacker could perform active jamming while observing whether the victim's throughput is disrupted to indirectly infer the quality of the jamming channels $H_{D_i}^{RIS}(c)$.

Our attack is geared towards bidirectional wireless communication devices that not only receive but also transmit RF signals which is crucial for the attacker to optimize their jamming channel. Therefore, completely passive radio receivers in unidirectional systems, *e.g.*, media broadcasting or satellite navigation, cannot be targeted as the attacker has no means of optimizing their RIS.

D. Future Work

In this work, we have studied selective targeting of individual receivers on the physical layer. Opposite to that, it would also be possible to target individual transmitters by employing transmitter-reactive jamming. More work is needed to assess the feasibility of such an approach and how it compares to ours. Moreover, the combination of spatial and time-varying jamming techniques provides an interesting opportunity to realize spatio-temporal jamming, *e.g.*, to enhance stealthiness. For instance, time-varying modulation of the RIS during jamming could be used for effective multi-device targeting, changing between a set of single-target RIS configurations. Although our attacker implementation already yields satisfactory results, we believe the hardware setup can be further improved, *e.g.*, realizing the attack with a single-chip wireless transceiver or using different RIS designs, possibly promoting hardware miniaturization. We employed a greedy algorithm from the literature to optimize the RIS. We believe that this process can be further improved, *e.g.*, by studying alternative algorithms, including machine learning-based approaches that might be capable of one-shot synthesis after an initial training. Using more fine-grained CSI channel measurements would likely aid faster and more accurate convergence while enabling spatio-spectral control of jamming signals.

VIII. Related Work

a) Differentiation from Previous Work: Previous research has investigated the adversarial use of RISs for jamming, yet with a clear focus on *passive* attacks. In particular, there are two main approaches: The first, proposed by Lyu *et al.* [38], is to use the RIS to reflect legitimate signals in way that a cancellation signal is formed at the targeted receiver, interfering destructively and thus reducing the received signal power. The second approach, first proposed by Staat *et al.* [59], leverages the RIS to create fast environmental variation which disturbs a targeted Wi-Fi receiver. Both approaches manipulate legitimate signals which is the key difference compared to our work: We *actively* transmit a jamming signal while using the RIS for precise attack targeting.

Karlsson *et al.* [30, 29] proposed to exploit channel reciprocity of TDD communication systems for jamming attacks. However, different from our work, their goal was to enhance the attacker's power efficiency and not selective jamming. Crucially, they consider an attacker employing a massive multiple-input and multiple-output (MIMO) radio instead of a single-antenna radio in conjunction with an RIS as we do.

b) Adversarial RIS Applications: Apart from jamming, the RIS can be used adversarially to, *e.g.*, manipulate radar sensing, as shown by Vennam *et al.* [49] and Chen *et al.* [10]. Zhu *et al.* [73] have shown that the RIS allows attackers to evade wireless sensing-based physical intrusion detection. Other works consider the RIS to facilitate eavesdropping, *e.g.*, Chen *et al.* [9], Chen and Ghasempour [8], and Shaikhanov *et al.* [54]. Finally, Li *et al.* [33] have shown RIS-based jamming of wireless key generation.

c) Jamming Attacks: An early study on the threat of jamming in wireless communication networks is the work of Xu *et al.* [72], covering several attack strategies, including constant random signal jamming, deceptive jamming based on packets with valid encoding, time-pulsed jamming, and reactive jamming. These types are also covered in various survey and overview works on attacks and defenses by, *e.g.*, Mpitziopoulos *et al.* [39], Grover *et al.* [21], Poisel [45], and Lichtman *et al.* [35]. Some of these recognize the attacker's antenna characteristics as a degree of freedom or make distinctions between omnidirectional and directional antennas. However, a concept like our scene-adaptive spatially selective

jamming is not mentioned. Proano and Lazos [46] describe time-domain selective wireless jamming based on real-time packet classification for reactive jamming. Pursuing the same goal, Aras *et al.* [3] describe a packet classification method for LoRaWAN. Reactive jamming has been implemented on smartphones [52] and software-defined radios [68], yet can be counteracted using hiding methods as outlined by Proano *et al.* [47]. Apart from the general threat of jamming, the literature also presents threat analyses for recent cellular systems such as 4G [18] and 5G [4], *e.g.*, discussing the impact of disrupting certain control channels.

A different line of work addresses the detection of jamming attacks [61, 11, 37], more recently also including machine-learning based methods [64, 41]. Other works examine friendly jamming, where the goal is to disrupt potential adversaries, *e.g.*, to achieve confidentiality [55, 31]. However, Tippenhauer *et al.* [66] and Robyns *et al.* [50] have shown that such schemes can be circumvented.

IX. Conclusion

In this paper, we investigated the merits of the RIS technology for active wireless jamming attacks. In particular, we have shown that the RIS enables precise physical-layer attack targeting in the spatial domain, enabling protocol level-agnostic selective jamming. For this, the attacker first determines an RIS configuration by eavesdropping wireless traffic from the victim devices. Then, the attacker uses the RIS to transmit a jamming signal that disrupts the wireless communication of targeted devices while leaving other devices operational. We have demonstrated the effectiveness of the attack under real-world conditions with extensive experimentation using commodity Wi-Fi devices and an open-source RIS. Notably, we found that that it is possible to differentiate between devices that are located only millimeters apart from each other. Overall, our work underscores the threat of wireless jamming attacks and recognizes the adversarial potential of RISs to enhance the landscape of wireless physical-layer attacks.

Acknowledgements

We thank Simon Tewes, Markus Heinrichs, and Rainer Kronberger for providing the RIS prototypes and Harald Elders-Boll for discussions. We thank the anonymous reviewers from both the initial and current versions of this paper for their valuable feedback. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972, RWTÜV Foundation (project number: S0189/10037/2021) and German Federal Office for Information Security (FKZ: Pentest-5GSec - 01MO23025B).

References

[1] Ali Abedi and Omid Abari. WiFi Says "Hi!" Back to Strangers! In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*, pages 132–138, 2020.

[2] Ali Abedi and Deepak Vasisht. Non-Cooperative Wi-Fi Localization & its Privacy Implications. In *Proceedings of the 28th Annual International Conference On Mobile Computing And Networking*, pages 570–582, 2022.

[3] Emekcan Aras, Nicolas Small, Gowri Sankar Ramachandran, Stéphane Delbruel, Wouter Joosen, and Danny Hughes. Selective Jamming of LoRaWAN Using Commodity Hardware. In *Proceedings of the 14th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, pages 363–372, Melbourne VIC Australia, November 2017. ACM.

[4] Youness Arjoune and Saleh Faruque. Smart Jamming Attacks in 5G New Radio: A Review. In *10th Annual Computing and Communication Workshop and Conference (CCWC)*, pages 1010–1015. IEEE, 2020.

[5] Ertugrul Basar, Marco Di Renzo, Julien De Rosny, Merouane Debbah, Mohamed-Slim Alouini, and Rui Zhang. Wireless Communications Through Reconfigurable Intelligent Surfaces. *IEEE Access*, 7:116753–116773, 2019.

[6] N.J. Burkett. Warning About Burglars Using Electronic Devices to Disable WiFi in Homes, 2024. https://abc7ny.com/post/warning-burglars-using-electronic-devices-disable-wifi-homes/14950485/, Accessed: October 19, 2024.

[7] BusinessTech. Remote jamming on the rise in South Africa, 2024. https://businesstech.co.za/news/business-opinion/496845/remote-jamming-on-the-rise-in-south-africa/, Accessed: October 19, 2024.

[8] Haoze Chen and Yasaman Ghasempour. Malicious mmWave Reconfigurable Surface: Eavesdropping through Harmonic Steering. In *Proceedings of the 23rd Annual International Workshop on Mobile Computing Systems and Applications*, HotMobile '22, page 54–60, New York, NY, USA, 2022. Association for Computing Machinery.

[9] Haoze Chen, Hooman Saeidi, Suresh Venkatesh, Kaushik Sengupta, and Yasaman Ghasempour. Wavefront Manipulation Attack via Programmable mmWave Metasurfaces: From Theory to Experiments. In *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '23, pages 317–328, New York, NY, USA, June 2023. Association for Computing Machinery.

[10] Xingyu Chen, Zhengxiong Li, Baicheng Chen, Yi Zhu, Chris Xiaoxuan Lu, Zhengyu Peng, Feng Lin, Wenyao Xu, Kui Ren, and Chunming Qiao. MetaWave: Attacking mmWave Sensing with Meta-material-enhanced Tags. In *The 30th Network and Distributed System Security (NDSS) Symposium*, volume 2023, 2023.

[11] Jerry T. Chiang and Yih-Chun Hu. Cross-Layer Jamming Detection and Mitigation in Wireless Broadcast Networks. *IEEE/ACM Transactions on Networking*, 19(1):286–298, 2011.

[12] Mostafa Zaman Chowdhury, Md Shahjalal, Shakil Ahmed, and Yeong Min Jang. 6G Wireless Communication Systems: Applications, Requirements, Technologies, Challenges, and Research Directions. *IEEE Open Journal of the Communications Society*, 1:957–975, 2020.

[13] Richard Hedley Clarke. A Statistical Theory of Mobile-Radio Reception. *Bell System Technical Journal*, 47(6):957–1000, 1968.

[14] Boris Danev, Srdjan Capkun, Ramya Jayaram Masti, and Thomas S. Benjamin. Towards Practical Identification of HF RFID Devices. *ACM Trans. Inf. Syst. Secur.*, 15(2), jul 2012.

[15] Anders Derneryd and Gerhard Kristensson. Signal Correlation Including Antenna Coupling. *Electronics letters*, 40(3):1, 2004.

[16] Jon Dugan, Set Elliott, Bruce A. Mah, Jeff Poskanzer, and Kaustubh Prabhu. iPerf - The TCP, UDP and SCTP Network Bandwidth Measurement Tool, 2024. https://iperf.fr/, Accessed: October 19, 2024.

[17] Biqian Feng, Junyuan Gao, Yongpeng Wu, Wenjun Zhang, Xiang-Gen Xia, and Chengshan Xiao. Optimization Techniques in Reconfigurable Intelligent Surface Aided Networks. *IEEE Wireless Communications*, 28(6):87–93, 2021.

[18] Felix Girke, Fabian Kurtz, Nils Dorsch, and Christian Wietfeld. Towards Resilient 5G: Lessons Learned from Experimental Evaluations of LTE Uplink Jamming. In *IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1–6. IEEE, 2019.

[19] Hadi Givehchian, Nishant Bhaskar, Eliana Rodriguez Herrera, Héctor Rodrigo López Soto, Christian Dameff, Dinesh Bharadia, and Aaron Schulman. Evaluating Physical-Layer Ble Location Tracking Attacks on Mobile Devices. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1690–1704. IEEE, 2022.

[20] Greenerwave. Our technology - Greenerwave. https://greenerwave.com/our-technology/, (Accessed: October 19, 2024).

[21] Kanika Grover, Alvin Lim, and Qing Yang. Jamming and Anti-jamming Techniques in Wireless Networks: A Survey. *International Journal of Ad Hoc and Ubiquitous Computing*, 17(4):197–215, 2014.

[22] Stefan Gvozdenovic, Johannes K Becker, John Mikulskis, and David Starobinski. Truncate after Preamble: PHY-based Starvation Attacks on IoT Networks. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 89–98, 2020.

[23] Wang Hang, Wang Zanji, and Guo Jingbo. Performance of DSSS against Repeater Jamming. In *13th IEEE International Conference on Electronics, Circuits and Systems*, pages 858–861. IEEE, 2006.

[24] Markus Heinrichs, Aydin Sezgin, and Rainer Kronberger. Open Source Reconfigurable Intelligent Surface for the Frequency Range of 5 GHz WiFi. In *IEEE International Symposium On Antennas And Propagation (ISAP)*, pages 1–2. IEEE, 2023.

[25] Wei Jiang, Bin Han, Mohammad Asif Habibi, and Hans Dieter Schotten. The Road Towards 6G: A Comprehensive Survey. *IEEE Open Journal of the Communications Society*, 2:334–366, 2021.

[26] Kyungho Joo, Wonsuk Choi, and Dong Hoon Lee. Hold the Door! Fingerprinting Your Car Key to Prevent Keyless Entry Car Theft. In Dongyan Xu and Ahmad-Reza Sadeghi, editors, *Proceedings 2020 Network and Distributed System Security Symposium*, Reston, VA, 2020. Internet Society.

[27] Nadège Kaina, Matthieu Dupré, Geoffroy Lerosey, and Mathias Fink. Shaping complex microwave fields in reverberating media with binary tunable metasurfaces. *Scientific Reports*, 4(1):6693, 2014.

[28] Majid Karimipour, Nader Komjani, and Iman Aryanian. Shaping Electromagnetic Waves with Flexible and Continuous Control of the Beam Directions Using Holography and Convolution Theorem. *Scientific Reports*, 9(1):11825, 2019.

[29] Marcus Karlsson, Emil Björnson, and Erik G Larsson. Jamming a TDD Point-to-point Link Using Reciprocity-based MIMO. *IEEE Transactions on Information Forensics and Security*, 12(12):2957–2970, 2017.

[30] Marcus Karlsson and Erik G Larsson. Massive MIMO as a Cyber-weapon. In *2014 48th Asilomar Conference on Signals, Systems and Computers*, pages 661–665. IEEE, 2014.

[31] Yu Seung Kim, Patrick Tague, Heejo Lee, and Hyogon Kim. Carving Secure Wi-Fi Zones with Defensive Jamming. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security*, pages 53–54, 2012.

[32] Mathieu Lacage, Mohammad Hossein Manshaei, and Thierry Turletti. IEEE 802.11 Rate Adaptation: A Practical Approach. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pages 126–134, 2004.

[33] Guyue Li, Paul Staat, Haoyu Li, Markus Heinrichs, Christian Zenger, Rainer Kronberger, Harald Elders-Boll, Christof Paar, and Aiqun Hu. RIS-Jamming: Breaking Key Consistency in Channel Reciprocity-based Key Generation, March 2023.

[34] Christos Liaskos, Shuai Nie, Ageliki Tsioliaridou, Andreas Pitsillides, Sotiris Ioannidis, and Ian Akyildiz. A novel communication paradigm for high capacity and security via programmable indoor wireless environments in next generation wireless systems. *Ad Hoc Networks*, 87:1–16, 2019.

[35] Marc Lichtman, Jeffrey D Poston, SaiDhiraj Amuru, Chowdhury Shahriar, T Charles Clancy, R Michael Buehrer, and Jeffrey H Reed. A Communications Jamming Taxonomy. *IEEE Security & Privacy*, 14(1):47–54, 2016.

[36] Yuanwei Liu, Xiao Liu, Xidong Mu, Tianwei Hou, Jiaqi Xu, Marco Di Renzo, and Naofal Al-Dhahir. Reconfigurable intelligent surfaces: Principles and opportunities. *IEEE communications surveys & tutorials*, 23(3):1546–1577, 2021.

[37] Nikita Lyamin, Alexey Vinel, Magnus Jonsson, and Jonathan Loo. Real-Time Detection of Denial-of-Service

Attacks in IEEE 802.11p Vehicular Networks. *IEEE Communications Letters*, 18(1):110–113, 2013.

[38] Bin Lyu, Dinh Thai Hoang, Shimin Gong, Dusit Niyato, and Dong In Kim. IRS-Based Wireless Jamming Attacks: When Jammers Can Attack Without Power. *IEEE Wireless Communications Letters*, 9(10):1663–1667, 2020.

[39] Aristides Mpitziopoulos, Damianos Gavalas, Charalampos Konstantopoulos, and Grammati Pantziou. A Survey on Jamming Attacks and Countermeasures in WSNs. *IEEE Communications Surveys & Tutorials*, 11(4):42–56, 2009.

[40] NovoFlect. NovoFlect - Reconfigurable Intelligent Surfaces. https://novoflect.de/, (Accessed: October 29, 2024).

[41] Jered Pawlak, Yuchen Li, Joshua Price, Matthew Wright, Khair Al Shamaileh, Quamar Niyaz, and Vijay Devabhaktuni. A Machine Learning Approach for Detecting and Classifying Jamming Attacks Against OFDM-based UAVs. In *Proceedings of the 3rd ACM Workshop on Wireless Security and Machine Learning*, pages 1–6, 2021.

[42] Xilong Pei, Haifan Yin, Li Tan, Lin Cao, Zhanpeng Li, Kai Wang, Kun Zhang, and Emil Björnson. RIS-aided Wireless Communications: Prototyping, Adaptive Beamforming, and Indoor/Outdoor Field Trials. *IEEE Transactions on Communications*, 69(12):8627–8640, 2021.

[43] Konstantinos Pelechrinis, Marios Iliofotou, and Srikanth V Krishnamurthy. Denial of Service Attacks in Wireless Networks: The Case of Jammers. *IEEE Communications Surveys & Tutorials*, 13(2):245–257, 2010.

[44] Eldad Perahia and Robert Stacey. *Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n*. Cambridge University Press, 2008.

[45] Richard Poisel. *Modern Communications Jamming Principles and Techniques, Second Edition*. Artech house, 2011.

[46] Alejandro Proano and Loukas Lazos. Selective Jamming Attacks in Wireless Networks. In *IEEE International Conference on Communications*, pages 1–6. IEEE, 2010.

[47] Alejandro Proano and Loukas Lazos. Packet-hiding methods for preventing selective jamming attacks. *IEEE Transactions on dependable and secure computing*, 9(1):101–114, 2011.

[48] Biswarup Rana, Sung-Sil Cho, and Ic-Pyo Hong. Review Paper on Hardware of Reconfigurable Intelligent Surfaces. *IEEE Access*, 2023.

[49] Rohith Reddy Vennam, Ish Kumar Jain, Kshitiz Bansal, Joshua Orozco, Puja Shukla, Aanjhan Ranganathan, and Dinesh Bharadia. mmSpoof: Resilient Spoofing of Automotive Millimeter-wave Radars Using Reflect Array. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 1807–1821, May 2023.

[50] Pieter Robyns, Peter Quax, and Wim Lamotte. PHY-layer Security Is No Alternative to Cryptography. In *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 160–162, Boston Massachusetts, July 2017. ACM.

[51] Javier Rosado-Sanz, María-Pilar Jarabo-Amores, Jean-Yves Dauvignac, David Mata-Moya, Jérôme Lanteri, and Claire Migliaccio. Design and Validation of a Reflectarray Antenna with Optimized Beam for Ground Target Monitoring with a Dvb-s-based Passive Radar. *Sensors*, 21(16):5263, 2021.

[52] Matthias Schulz, Francesco Gringoli, Daniel Steinmetzer, Michael Koch, and Matthias Hollick. Massive Reactive Smartphone-Based Jamming using Arbitrary Waveforms and Adaptive Power Control. In *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 111–121, 2017.

[53] Matthias Schulz, Daniel Wegemer, and Matthias Hollick. Nexmon: The C-based Firmware Patching Framework, 2017.

[54] Zhambyl Shaikhanov, Fahid Hassan, Hichem Guerboukha, Daniel Mittleman, and Edward Knightly. Metasurface-in-the-Middle Attack: From Theory to Experiment. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec '22, pages 257–267, New York, NY, USA, May 2022. Association for Computing Machinery.

[55] Wenbo Shen, Peng Ning, Xiaofan He, and Huaiyu Dai. Ally Friendly Jamming: How to Jam Your Enemy and Maintain Your Own Wireless Connectivity at the Same Time. In *IEEE Symposium on Security and Privacy*, pages 174–188. IEEE, 2013.

[56] Joshua Smailes, Sebastian Köhler, Simon Birnbach, Martin Strohmeier, and Ivan Martinovic. Watch This Space: Securing Satellite Communication through Resilient Transmitter Fingerprinting, September 2023.

[57] Nathan Solis. LAPD Warns Homeowners About a Simple Trick Burglars Use to Disable Home Security Systems, 2024. https://www.latimes.com/california/story/2024-03-06/lapd-warns-about-wi-fi-jamming-methods-for-thieves-to-disable-home-security-systems, Accessed: October 19, 2024.

[58] Naeimeh Soltanieh, Yaser Norouzi, Yang Yang, and Nemai Chandra Karmakar. A Review of Radio Frequency Fingerprinting Techniques. *IEEE Journal of Radio Frequency Identification*, 4(3):222–233, 2020.

[59] Paul Staat, Harald Elders-Boll, Markus Heinrichs, Christian Zenger, and Christof Paar. Mirror, Mirror on the Wall: Wireless Environment Reconfiguration Attacks Based on Fast Software-Controlled Surfaces. In *Proceedings of the ACM on Asia Conference on Computer and Communications Security*, pages 208–221, 2022.

[60] Anders Stjernman. Antenna Mutual Coupling Effects on Correlation, Efficiency and Shannon Capacity. In *2006 First European Conference on Antennas and Propagation*, pages 1–6. IEEE, 2006.

[61] Mario Strasser, Boris Danev, and Srdjan Čapkun. Detection of Reactive Jamming in Sensor Networks. *ACM Transactions on Sensor Networks (TOSN)*, 7(2):1–29,

2010.

[62] Wankai Tang, Xiangyu Chen, Ming Zheng Chen, Jun Yan Dai, Yu Han, Shi Jin, Qiang Cheng, Geoffrey Ye Li, and Tie Jun Cui. On Channel Reciprocity in Reconfigurable Intelligent Surface Assisted Wireless Networks. *IEEE Wireless Communications*, 28(6):94–101, 2021.

[63] Harsh Tataria, Mansoor Shafi, Andreas F Molisch, Mischa Dohler, Henrik Sjöland, and Fredrik Tufvesson. 6G Wireless Systems: Vision, Requirements, Challenges, Insights, and Opportunities. *Proceedings of the IEEE*, 109(7):1166–1199, 2021.

[64] Enrico Testi, Luca Arcangeloni, and Andrea Giorgetti. Machine Learning-Based Jamming Detection and Classification in Wireless Networks. In *Proceedings of the ACM Workshop on Wireless Security and Machine Learning*, pages 39–44, 2023.

[65] Simon Tewes, Markus Heinrichs, Paul Staat, Rainer Kronberger, and Aydin Sezgin. Full-duplex meets Reconfigurable Surfaces: RIS-assisted SIC for Full-Duplex Radios. In *IEEE International Conference on Communications (ICC)*, pages 1106–1111. IEEE, 2022.

[66] Nils Ole Tippenhauer, Luka Malisa, Aanjhan Ranganathan, and Srdjan Capkun. On Limitations of Friendly Jamming for Confidentiality. In *IEEE symposium on security and privacy*, pages 160–173. IEEE, 2013.

[67] François Vergès. MCS Index - MCS Index Table, Modulation and Coding Scheme Index 11n, 11ac, and 11ax. https://mcsindex.com/, Accessed: October 19, 2024.

[68] Matthias Wilhelm, Ivan Martinovic, Jens B. Schmitt, and Vincent Lenders. Short paper: reactive jamming in wireless networks: how realistic is the threat? In *Proceedings of the Fourth ACM Conference on Wireless Network Security*, WiSec '11, page 47–52, New York, NY, USA, 2011. ACM.

[69] Joshua Wright and dragorn. lorcon - Loss Of Radio CONtrol, 2009. https://github.com/kismetwireless/lorcon, Accessed: October 19, 2024.

[70] Qingqing Wu and Rui Zhang. Towards Smart and Reconfigurable Environment: Intelligent Reflecting Surface Aided Wireless Network. *IEEE Communications Magazine*, 58(1):106–112, 2019.

[71] Dong Xia, Jonathan Hart, and Qiang Fu. Evaluation of the Minstrel Rate Adaptation Algorithm in IEEE 802.11 g WLANs. In *2013 IEEE International Conference on Communications (ICC)*, pages 2223–2228. IEEE, 2013.

[72] Wenyuan Xu, Wade Trappe, Yanyong Zhang, and Timothy Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, pages 46–57, 2005.

[73] Yuxuan Zhou, Chenggao Li, Huangxun Chen, and Qian Zhang. RIStealth: Practical and Covert Physical-Layer Attack against WiFi-based Intrusion Detection via Reconfigurable Intelligent Surface. In *21th ACM Conference on Embedded Networked Sensor Systems (SenSys 2023)*, 2023.

[74] Yanzi Zhu, Zhujun Xiao, Yuxin Chen, Zhijing Li, Max Liu, Ben Y. Zhao, and Haitao Zheng. Et Tu Alexa? When Commodity WiFi Devices Turn into Adversarial Motion Sensors. In *27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23-26, 2020*, NDSS '20. Internet Society, 2020.

[75] Yuze Zou, Yusi Long, Shimin Gong, Dinh Thai Hoang, Wei Liu, Wenqing Cheng, and Dusit Niyato. Robust Beamforming Optimization for Self-sustainable Intelligent Reflecting Surface Assisted Wireless Networks. *IEEE Transactions on Cognitive Communications and Networking*, 8(2):856–870, 2021.

APPENDIX

A. Experimental Setup

Complementing the description in Section V, we show a photo of the experimental setup in Figure 21, comprising of the RIS that is illuminated by the attacker's antenna.

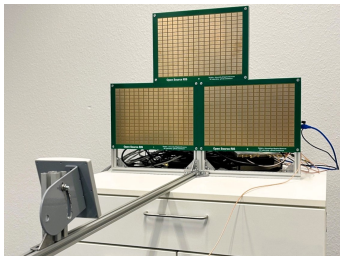

Fig. 21: Attacker setup used for experimental evaluation.

B. Close-By Antennas

We here report additional results from the experiment described in Section VI-E2.

Figure 22a shows the RIS optimization process to achieve maximization of the attacker's channel towards Antenna 1 while minimizing the channel towards Antenna 2. The antennas are placed directly next to each other as shown in Figure 22a and face the attacker's RIS. In this experiment, we used a VNA to gather channel measurements, confirming the observations previously made with Wi-Fi devices.

In the next experiment, we study the effect of antenna positioning and removal. For this, we utilize a 3D-printed positioning fixture as shown in Figure 22d. In Figure 22c, we show the channel measurements for various combinations of antenna locations. On the x-axis, we indicate which antenna is located at which position in Figure 22d. The first value in parentheses corresponds to the left position, the second value to the right position. After the initial RIS optimization, the power transfer towards the two antennas differs by more than $30\,\mathrm{dB}$. Then, when removing Antenna 2, the maximized Antenna 1 is barely affected. In contrast, when removing Antenna 2, the initial $-107\,\mathrm{dB}$ minimization of Antenna 2 deteriorates significantly to around $-80\,\mathrm{dB}$. Similar effects are observed

when exchanging Antenna 1 and Antenna 2. However, with both antennas present, yet swapped, the initial performance is not matched, indicating slight deviations and imperfections regarding equal antenna positioning. However, when finally repeating the initial measurement, the performance again matches the initial values. This experiment clearly highlights the effects of mutual antenna coupling on the attacker's ability to separate the antenna channels despite close proximity.

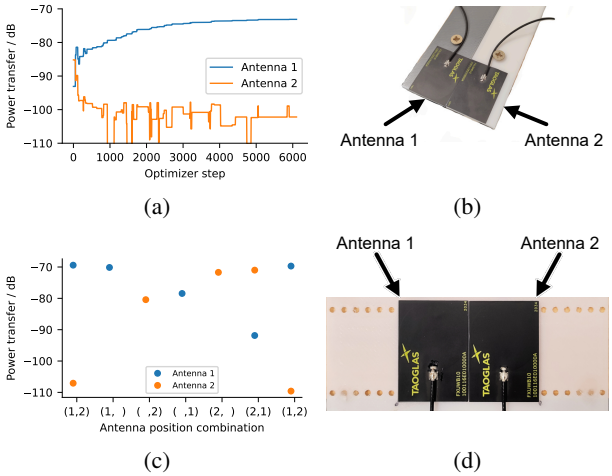

Fig. 22: Optimization progress (a) of antenna arrangement in (b) using VNA measurements. (c) Illustrates effects of removing and reordering antennas, with (1, 2) indicating the arrangement in (d).

C. Behavior of the RIS Optimization

Throughout this work, we conducted numerous experiments involving optimization of the RIS configuration. As outlined in Section V-3, we utilize a greedy heuristic from the literature [65]. To characterize the consistency and course of the optimization, we performed the single-target experiment from Section VI-A2 with $\mathcal{T} = \{D_4\}$, repeated 50 times. For each run and algorithm step, we stored the current best RIS configuration and measured the resulting JSR at the devices.

Evaluating the optimization convergence speed, we treat RIS configurations as 768-bit sequences and calculate the

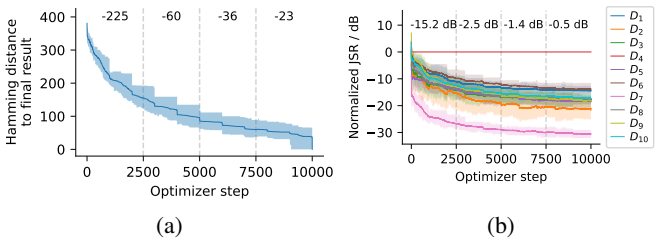

Fig. 23: Optimization results from the RIS over 50 runs targeting D_4: (a) Hamming distances to the final configuration and (b) measured JSR per step. Annotations indicate changes in average Hamming distance and JSR over 2500 steps.

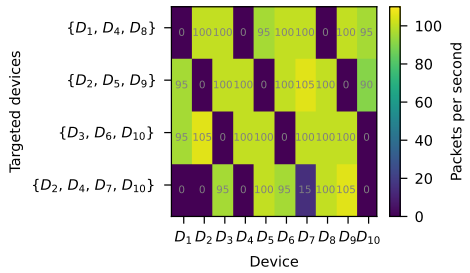

Fig. 24: Additional results highlighting the ability to target multiple devices in different clusters.

Hamming distances of the final optimization results to previous ones of the optimization progress. The average and 5th and 95th percentiles across repetitions over the optimization progress are shown in Figure 23a. Due to the random algorithm initialization, we first observe an average difference of 382 elements, close to the ideal expected value of 384. At the beginning, the RIS configuration quickly evolves towards the final result, as evident by the steep reduction of the Hamming distance by 225 after 2500 steps. After 4573 algorithm steps, the average Hamming distance to the final optimization result is below 100 elements. From the percentiles of the distribution, we see that this behavior is largely consistent across different instantiations of the algorithm, regardless of the random initialization. Please note that the staircase pattern stems from periodic re-evaluation of all B RIS configuration candidates every 1000 steps.

In Figure 23b, we present the corresponding normalized JSR during the optimizer progress, again with the 5th and 95th percentiles across algorithm repetitions. Here, it becomes evident that similar JSR performance is achieved, regardless of the random algorithm initialization and the inherently noisy RSSI measurements. In the plot, we annotate the average JSR reductions after 2500 steps, showing only marginal improvement of 1.9 dB during the last 5000 steps. Thus, we conclude that terminating the optimization early is possible without significantly sacrificing JSR performance, potentially allowing quicker adaption in dynamic environments.

D. Cross-Cluster Multi-Target Jamming

Complementing the results from Section VI-B, we now demonstrate that it is possible to target multiple devices even when these do not belong to the same device clusters. In particular, we conducted a multi-target jamming experiment where we subsequently targeted the devices $\{D_1, D_4, D_8\}$, $\{D_2, D_5, D_9\}$, $\{D_3, D_6, D_{10}\}$, and $\{D_2, D_4, D_7, D_{10}\}$. The resulting packet rates are illustrated in Figure 24, showing successful disruption of one device per cluster. When we include the device D_7 to be targeted, which is closest to the access point and hence demands most jamming signal power to be disrupted, we observed that the D_1 is unintentionally jammed as well.