TECHNISCHE
UNIVERSITÄT
DARMSTADT

PERFORMANCE AND SECURITY ENHANCEMENTS IN PRACTICAL MILLIMETER-WAVE COMMUNICATION SYSTEMS

Vom Fachbereich Informatik der
Technischen Universität Darmstadt genehmigte

DISSERTATION

zur Erlangung des akademischen Grades eines
Doktor-Ingenieurs (Dr.-Ing.)

von

DANIEL STEINMETZER, M. SC.

geboren am 2. Februar 1987 in Northeim.

Erstreferent: Prof. Dr.-Ing. Matthias Hollick
Korreferent: Dr. rer. nat. Joerg Widmer

Tag der Einreichung: 17. Dezember 2018
Tag der Disputation: 28. Januar 2019

Darmstadt, 2019
Hochschulkennziffer D17

Daniel Steinmetzer, *Performance and Security Enhancements in Practical Millimeter-Wave Communication Systems*, Dissertation, TU Darmstadt, 2019.

Fachgebiet Sichere Mobile Netze
Fachbereich Informatik
Technische Universität Darmstadt
Jahr der Veröffentlichung: 2019
Tag der mündlichen Prüfung: 28. Januar 2019
URN: urn:nbn:de:tuda-tuprints-83253
URL: https://tuprints.ulb.tu-darmstadt.de/id/eprint/8325

To my beloved parents for their endless support.

ABSTRACT

Millimeter-wave (mm-wave) communication systems achieve extremely high data rates and provide interference-free transmissions. to overcome high attenuations, they employ directional antennas that focus their energy in the intended direction. Transmissions can be steered such that signals only propagate within a specific area-of-interest. Although these advantages are well-known, they are not yet available in practical networks. IEEE 802.11ad, the recent standard for communications in the unlicensed 60 GHz band, exploits a subset of the directional propagation effects only. Despite the large available spectrum, it does not outperform other developments in the prevalent sub-6 GHz bands. This underutilization of directional communications causes unnecessary performance limitations and leaves a false sense of security. For example, standard compliant beam training is very time consuming. It uses suboptimal beam patterns, and is unprotected against malicious behaviors. Furthermore, no suitable research platform exists to validate protocols in realistic environments. To address these challenges, we develop a holistic evaluation framework and enhance the performance and security in practical mm-wave communication systems. Besides signal propagation analyses and environment simulations, our framework enables practical testbed experiments with off-the-shelf devices. We provide full access to a tri-band router's operating system, modify the beam training operation in the Wi-Fi firmware, and create arbitrary beam patterns with the integrated antenna array. This novel approach allows us to implement custom algorithms such as a compressive sector selection that reduces the beam training overhead by a factor of 2.3. By aligning the receive beam, our adaptive beam switching algorithm mitigates interference from lateral directions and achieves throughput gains of up to 60 %. With adaptive beam optimization, we estimate the current channel conditions and generate directional beams that implicitly exploit potential reflections in the environment. These beams increase the received signal strength by about 4.4 dB. While intercepting a directional link is assumed to be challenging, our experimental studies show that reflections on small-scale objects are sufficient to enable eavesdropping from afar. Additionally, we practically demonstrate that injecting forged feedback in the beam training enables Man-in-the-Middle attacks. With only 7.3 % overhead, our authentication scheme protects against this beam stealing and enforces responses to be only accepted from legitimate devices. By making beam training more efficient, effective, and reliable, our contributions finally enable practical applications of highly directional transmissions.

ZUSAMMENFASSUNG

Kommunikationssysteme auf Basis von Millimeterwellen ermöglichen interferenzfreie Übertragungen mit besonders hohen Datenraten. Signale können gerichtet ausgesehendet werden um der hohen Freiraumdämpfung entgegenzuwirken, so dass sie nur an bestimmten Orten empfangen werden. Durch diese direktionale Kommunikation ergeben sich Vorteile die derzeit noch keine Anwendung in praktischen WLAN Systemen finden. Der aktuelle Standard für Kommunikationen im lizenzfreien 60 GHz Bereich macht von den Möglichkeiten der Richtfunktechnik nur eingeschränkt Gebrauch. Trotz des großen verfügbaren Frequenzspektrums werden nur geringfügig höhere Datenraten als mit üblichen Verfahren im 2 - 6 GHz Bereich erreicht. Insbesondere das standardmäßig verwendete Trainingsverfahren für die Antennenausrichtung führt zu Einschränkungen der Performanz und Sicherheit solcher Systeme. Die zeitaufwändige Suche führt zu keinen optimalen Ergebnissen und ist anfällig für Manipulationen. Erschwerend kommt hinzu, dass es derzeit keine geeigneten Werkzeuge gibt, um die vielfältigen Eigenschaften in realistischen Umgebungen analysieren zu können. Im Rahmen dieser Arbeit stellen wir daher ein umfangreiches Evaluationssystem zur Verfügung. Neben der Analyse und Simulation spezifischer Signalausbreitung, ermöglicht unser System praktische Experimente. Um beliebige Trainingsalgorithmen und Antennenkonfigurationen verwenden zu können, modifizieren wir die proprietäre WLAN Firmware in handelsüblichen Geräten. Dieser neuartige Ansatz ermöglicht es uns u.a. ein kompressives Verfahren zu implementieren, welches den Trainingsaufwand um einen Faktor von 2,3 reduziert. Anpassungen der Signalausrichtung auf der Empfangsseite können Störungen unterdrücken und steigern damit den Datendurchsatz um 60 %. Ferner können wir durch eine genaue Kanalschätzung die Antennenausrichtung optimieren und die Signalstärke am Empfänger um durchschnittlich 4,4 dB erhöhen. Zwar erschweren direktionale Übertragungen das Abhören der Kommunikation, jedoch können kleinste Objekte in der Umgebung zu starken Reflektionen führen. In einem praktischen Aufbau belegen wir, dass diese ausreichen um übertragende Informationen auszulesen. Ferner können manipulierte Rückmeldungen im Trainingsverfahren sogenannte MITM Angriffe ermöglichen. Um diesen entgegenzuwirken, stellen wir ein Authentifizierungsverfahren vor, welches mit 7,3 % Mehraufwand nur legitime Antennenausrichtungen zulässt. Durch die Beiträge in dieser Arbeit verbessern wir die Effizienz, Effektivität und Zuverlässigkeit von auf Millimeterwellen basierten Kommunikationssystemen und ermöglichen somit dessen praktische Anwendung.

ACKNOWLEDGMENTS

Special thanks for guiding and supporting me over the years writing this thesis goes to my supervisor Prof. Matthias Hollick. His ideas and feedback were absolutely invaluable for my doctoral studies.

In addition, I would especially like to thank my co-supervisor Dr. Joerg Widmer not only for accepting my request to review this work but also for our long-term collaboration.

Further, I thank my other collaborators and the students I supervised for their hard work that lead to particular contributions in this thesis.

I thank my colleagues, our secretary, and the technical staff for the great working environment and joyful atmosphere.

Finally, I thank the German Research Foundation (DFG), the German Federal Ministry of Education and Research (BMBF), and the Federal State of Hesse for funding my work through various projects.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

A/D	Analog-to-Digital
AGC	Automatic Gain Control
AP	Access Point

BER	Bit Error Rate
BI	Beacon Interval
BPSK	Binary Phase-Shift Keying
CDF	Cumulative Distribution Function
CE	Channel Estimation
CIR	Channel Impulse Response
CSI	Channel State Information
D/A	Digital-to-Analog
DMG	Directional Multi-Gigabit
DoS	Denial-of-Service
ECDH	Elliptic-curve Diffie-Hellman
FPGA	Field-Programmable Gate Array
FSPL	Free-Space Path Loss
GCC	GNU Compiler Collection
I/O	Input/Output
IF	Intermediate Frequency
IQ	In-phase and Quadrature
LEDE	Linux Embedded Development Environment
MCS	Modulation and Coding Scheme
MIMO	Multiple-Input and Multiple-Output
MITM	Man-in-the-Middle
PCB	Printed Circuit Board
PDP	Power Delay Profile
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RFIC	Radio Frequency Integrated Circuit

RMS	Root Mean Square
RPC	Remote Procedure Call
RSSI	Received Signal Strength Indicator
SBR	Shooting and Bouncing Rays
SDR	Software-Defined Radio
SINR	Signal-to-Interference-plus-Noise Ratio
SIR	Signal-to-Interference Ratio
SNR	Signal-to-Noise Ratio
SoC	System-on-a-Chip
SSH	Secure Shell
SSID	Service Set Identification
STF	Short Training Field
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
URI	Unique Resource Identifier
WARP	Wireless Open Access Research Platform
WMI	Wireless Module Interface

PREVIOUSLY PUBLISHED MATERIAL

This thesis includes material of previously published conference papers and articles. Following the regulations of the Computer Science department at Technische Universität Darmstadt, I list below the affected parts that include verbatim and rephrased fragments from these publications.

CHAPTER 2:

- Section 2.1 revises Section 5 of "Eavesdropping with Periscopes: Experimental Security Analysis of Highly Directional Millimeter Waves" [Ste+15] and Section 2 of "mmTrace: Modeling Millimeter-wave Indoor Propagation with Image-based Ray-tracing" [SCH16b].
- Section 2.3 revises Section 8 of "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices" [Pal+18a], Section 8 of "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices" [Ste+17b], and Section 2 of "Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks" [Ste+17a].
- Section 2.4 revises Section 1 of "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks" [Ste+18] and Section 5 of "Eavesdropping with Periscopes: Experimental Security Analysis of Highly Directional Millimeter Waves" [Ste+15].
- Section 2.5 revises Section 8 of "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices" [Ste+17b].

CHAPTER 3:

- Section 3.1 revises Section 3 of "Eavesdropping with Periscopes: Experimental Security Analysis of Highly Directional Millimeter Waves" [Ste+15].
- Section 3.2 revises Section 3 and 4 of "mmTrace: Modeling Millimeter-wave Indoor Propagation with Image-based Ray-tracing" [SCH16b].
- Section 3.3 revises Section 5 of "mmTrace: Modeling Millimeter-wave Indoor Propagation with Image-based Ray-tracing"

[SCH16b] and summarizes Section 2 of "Exploring Millimeter-Wave Network Scenarios with Ray-tracing based Simulations in mmTrace" [SCH16a].

CHAPTER 4:

- Section 4.1 and Section 4.2 revise Section 3 of "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices" [Ste+17b].
- Section 4.3 revises Section 5 of "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices" [Pal+18a].
- Section 4.4 revises Section 1, 2, and 4 of "TPy: A Lightweight Framework for Agile Distributed Network Experiments" [SSH18].
- Section 4.5 revises "A Practical IEEE 802.11ad Research Platform: The Hidden Potential of Off-the-Shelf Devices" [SWH18].

CHAPTER 5:

- Section 5.1, 5.2, and Section 5.3 revise Section 1, 2, 5, 6, 7, and 9 of "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices" [Ste+17b].

CHAPTER 6:

- Section 6.1, 6.2, and 6.3 revise Section 1, 3, 4, and 5 of "Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks" [Ste+17a].

CHAPTER 7:

- Section 7.1, 7.2, and Section 7.3 revise Section 1, 2, 3, 4, 7, and 9 of "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices" [Pal+18a], as well as Section 1 and 2 of "Demo: Channel Estimation and Custom Beamforming on the 60 GHz TP-Link Talon AD7200 Router" [Pal+18c].

CHAPTER 8:

- Section 8.1, 8.2, and 8.3 revise Section 1, 2, 4, and 6 of "Eavesdropping with Periscopes: Experimental Security Analysis of Highly Directional Millimeter Waves" [Ste+15].

CHAPTER 9:

- Section 9.1 revises Section 1 and 2 of "Beam-Stealing: Intercepting the Sector Sweep to Launch Man-in-the-Middle Attacks on Wireless IEEE 802.11ad Networks" [SYH18] and Section 1 and 2 of "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks" [Ste+18].
- Section 9.2 revises Section 3, 4, and 5 of "Beam-Stealing: Intercepting the Sector Sweep to Launch Man-in-the-Middle Attacks on Wireless IEEE 802.11ad Networks" [SYH18].
- Section 9.3 revises Section 3 and 4 of "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks" [Ste+18].
- Section 9.4 revises Section 7 and 8 of "Beam-Stealing: Intercepting the Sector Sweep to Launch Man-in-the-Middle Attacks on Wireless IEEE 802.11ad Networks" [SYH18], as well as Section 5 and 6 of "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks" [Ste+18].

CHAPTER 10:

- Section 10.1 and Section 10.2 revise Section 4 of "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices" [Ste+17b].

CHAPTER 12:

- Section 12.1 summarizes "Pseudo Lateration: Millimeter-wave Localization using a Single RF Chain" [Che+17] and "Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points" [Bie+18].

COLLABORATIONS

Systematically investigating a technical research topic and engineering the required tools is a demanding and interdisciplinary process. Most achievements could never evolve without collaborations in which colleagues and international partners integrated their intellectual forces. When working in teams, accounting particular contributions and components of the resulting publications to individual collaborators becomes almost impossible. This situation also applies to several contents of this thesis, which arise from collaborations, thus, covering joint contributions. Many of these collaborations persisted even longer than the research projects and became a long-term strategic partnership. In our previous publications, all authors contributed by discussing ideas and debating on results throughout the whole project duration. Each of them has particular strengths that sometimes appear invisible. For this reason, I explicitly state and acknowledge—where possible—the contributions of my collaborators in the following.

FRAMEWORK. The channel sounding platform in Section 3.1 originated during my research visit at RICE University in Houston, Texas, USA. Based on existing hardware components, I developed the concept and assembled the system. J. Chen supported me performing the measurements in various scenarios, while J. Classen, E. Knighly, and M. Hollick contributed their ideas and comments during numerous discussions. Our environment simulation tool in Section 3.2 was joint work with my colleagues J. Classen and M. Hollick. I provided the concept, implemented the simulation tool and achieved valuable support from both of them. The testbed experimentation platform in Chapter 4 grew steadily over time. In our initial work, M. Schulz and D. Wegemer contributed their expertise on the Nexmon firmware patching framework and integrated the required changes to support the specific processor architecture. D. Wegemer supported me in analyzing the firmware layout of the Wi-Fi chip and the router's default system image. Porting the OpenWrt environment to the specific hardware of the router and developing the patches to control the beam training was my contribution. Additionally, I obtained access to the sector configurations and reconstructed the antenna weighting network. J. Widmer and M. Hollick added valuable comments and suggestions for possible applications. Besides, I started to develop our experiment automation tool to control multiple Wi-Fi routers in parallel. M. Stute later adopted this approach for his purposes and jointly we extended it for general applicability.

PERFORMANCE. Performing an efficient beam training with compressive sector selection in Chapter 5 was a joint work with J. Widmer and M. Hollick. The idea arose in various discussions. With the features of our practical testbed framework, I designed the algorithm and integrated it into the sector sweep operation. In addition, I took care of the measurements and evaluations. Another collaboration on my initiative with A. Loch, A. García-García, J. Widmer, and M. Hollick yielded our interference mitigation approach in Chapter 6. I performed the first analyses, measurements, and simulations, and then designed the beam switching protocol. A. Loch and A. García-García ran additional measurements and evaluated the protocol in real-world scenarios. Together with A. Loch, I wrote up and presented our results. Our beam optimization approach in Chapter 7 was a collaboration with J. Palacios, A. Loch, J. Widmer, and M. Hollick. After jointly discussing the idea and goals, J. Palacios took care of the mathematical model, the algorithm design, and conducted the evaluation experiments. I contributed by analyzing the antenna and provided the capabilities to change the beam configuration. My achievements enable to extract the signal properties from individual antenna elements with different phase shifts. The presentation of our results was jointly prepared with A. Loch.

SECURITY. Our investigation of reflections and whether eavesdropping is possible from distant locations in Chapter 8 was the first application of our channel sounding platform. The idea evolved in several meetings. Together with J. Chen, I carried-out the measurements at RICE University to characterize reflections of various objects. Based on the results, I derived the security characteristics and finished the write-up. During the whole process, we received valuable support from J. Classen, E. Knightly, and M. Hollick. Our beam stealing attack in Chapter 9 is based on the master thesis of Y. Yuan that M. Hollick and I supervised. I started with the idea of stealing beams with forged feedback in the sector level sweep and offered this thesis. Y. Yuan implemented the attack based on our testbed experimentation framework and conducted the experiments. Afterward, I revised the presentation of our results for publication. The authentication scheme that protects the sector level sweep against forged feedback in Section 9.3 was developed as part of the master thesis of S. Ahmad. Following my idea to authenticate the feedback, he designed the first protocol, provided the simulations, and performed the measurements. Later, I amended the protocol design and revised the write-up that leads to our publication. N. Anagnostopoulos, S. Katzenbeisser, and M. Hollick co-supervised the thesis and consistently integrated beneficial suggestions.

APPLICATIONS. The localization scheme using a single Access Point (AP) in Section 3.3 was the achievement of J. Chen in collaboration with J. Classen, M. Hollick, E. Knightly and myself. I contributed the technical configuration, set up the hardware for the experiments (see Chapter 3), and supported the practical measurements. The work for the second localization system with multiple APs in Section 3.3 was done by G. Bielsa, J. Palacios, A. Loch, P. Casari, and J. Widmer at IMDEA Networks Institute in Madrid, Spain. I provided the platform and framework along the capabilities to extract the required parameters from the utilized hardware devices. Thus, my contribution facilitated the experiments on device localization. The algorithm itself, as well as the measurements, were developed and carried-out under the lead of G. Bielsa.

Part I

INTRODUCTION

We first introduce our goals and contributions in Chapter 1 and then present background information and related work in Chapter 2.

1 INTRODUCTION

Wireless communication technology has become pervasive with billions of mobile devices worldwide. Due to the widespread use of mobile phones and other personal gadgets, the demands on ubiquitous connectivity raise. The wireless traffic annually increases by about 46 % [Cis18] and soon reach a level that cannot be handled by current technologies anymore. Engineers address these demands by developing new approaches to increase the spectrum utilization and coding efficiency. Still, the available spectrum is almost saturated and leaves little space for improvements.

The wireless communication spectrum is almost saturated.

Millimeter-wave (mm-wave) communication systems operate at high frequencies and offer new opportunities to address the aforementioned demands. The unlicensed 60 GHz band has an available spectrum twelve times as large as those in both the prevalent 2.4 GHz and 5 GHz Wi-Fi systems. Additionally, next-generation cellular systems may utilize more than 20 GHz of spectrum in the range of 28 GHz to 100 GHz [Rap+14]. This massive bandwidth available solves the limited spectrum bottleneck and enables data rates of multiple Gigabit per second.

Millimeter-wave communications provide a huge bandwidth.

Transmitting signals at high frequencies causes additional attenuations due to free-space propagation losses and atmospheric absorptions [Rap+14]. In contrast to common communication technologies, mm-wave signals are less affected by scattering and diffractions but strongly impaired by blockage. Only the line-of-sight or direct reflections allow constructing a communication path. These quasi-optical propagation effects lead to naturally decreased communication distances and bounded coverage. To overcome these impairments, mm-wave systems employ steerable directional antennas that focus their signal power in the intended direction of communication. They achieve a strong signal gain inside the designated beam, which significantly drops in any other direction. With this directionality applied, mm-wave communications not only improve wireless performance but also enable emerging applications. For example, exploiting the specific propagation effects facilitates a precise device localization, gesture recognition, or "information showers" with location based services [Rap+14]. Since the size of antennas is in the order of millimeters, inter-chip communications and cable replacements [RMG11] are conceivable. Vehicular applications already utilize 77 GHz transmissions for automotive radars that may soon be adopted for signaling and data transfer [KGH15]. In all these scenarios, narrow beams facilitate spatial reusability with low interference among concurrent

High frequencies induce different propagation effects.

Directional transmissions enable emerging applications.

transmissions, thus increasing the overall performance in ultra-dense networks. However, due to multiple reasons and current hardware imperfections these great advantages are retained from practical use and market availability. In the following, we state these issues and motivate our work in Section 1.1. We define current challenges as well as our goals in Section 1.2 and list our contributions in Section 1.3. Finally, Section 1.4 provides an overview on the remainder of this thesis.

1.1 MOTIVATION

Current systems underutilize the advantages of directional communication.

Current mm-wave devices and enrolled applications only tap a slight fraction of the advantages of directional communications. For example, IEEE 802.11ad, the first Wi-Fi standard for the 60 GHz band, achieves data rates that are not significantly higher than those in common sub-6 GHz communication systems such as IEEE 802.11ac (Wi-Fi 5) and IEEE 802.11ax (Wi-Fi 6). Despite the smaller spectrum, the latter provides data rates of more than 1 Gbps which still meets our current demands. Moreover, the specific propagation effects of mm-waves make it challenging to replace existing communication infrastructure. Due to the high directionality, network deployments of base stations and Access Points (APs) must consider environmental effects and potential link outages. Indoor scenarios with mm-wave devices typically require to place multiple APs in a single room to achieve coverage as in sub-6 GHz Wi-Fi systems. These effects demand a fundamental rethinking of wireless networking and induce a new transition from omnidirectional to directional communications.

Communication at high frequencies is technically challenging.

Designing wireless communication systems that operate at very high frequencies such as 60 GHz is technically challenging. For instance, electrical components must meet tiny tolerances, and phased antenna arrays are hard to manufacture [Won+14; Niu+15]. Phase noise plays a significant role [TCK15; Dha15] and typically impedes the signal decoding. System designers often resort to sub-optimal solutions to address these challenges efficiently. As a result, current systems provide proper communication links but underutilize the advantages of directionality in mm-wave communications.

Research platforms with lower layer access are not yet available.

To enhance the performance of wireless communication, researchers typically use designated development systems to evaluate their protocol designs. Despite that consumer-grade devices already use IEEE 802.11ad, no existing research platform supports the mm-wave frequency bands. Available off-the-shelf devices lack suitable features to control the beam training and other lower layer parameters. The firmware of the Wi-Fi chip in such devices is a black box [Nit+15a] and hinders users from prototyping new protocols. Low-layer access is possible only within the limitations of the provided interface [LBW16]. Hence, many researchers either limit themselves to the-

oretical approaches [MRM16; RVM12b; RVM12a] or use simplified measurements to verify their solutions. Due to the lack of wide-spread hardware, they base their work on prototyping platforms with directional horn antennas [Sur+15; Sur+16; HK16] or custom antenna arrays [Ras+17]. Both exhibit different behavior than any practical device. The low-cost components integrated with the latter cause imperfections and do not achieve the precision of laboratory equipment.

Attack possibilities in practical scenarios differ those in ideal environments.

Most mm-wave applications implicitly assume a strong directionality. Signals should be only receivable within a bounded coverage area of the designated beam. Indeed, the IEEE 802.11ad standard specifies beamwidths as small as three degree [Nit+14] to compensate for the high path loss and realize links at average Wi-Fi-scale distances. Under ideal circumstances, mm-wave communication systems are often asserted to be inherently resilient to eavesdropping and other attacks. Signal interception is assumed infeasible if the attacker is forced to locate itself within several degrees off the path between the transmitter and the receiver [Fre13; Yan+15]. This assumption might hold under laboratory conditions with perfect directionality but leaves a false sense of security for practical deployments. In particular, most development systems feature completely different radiation characteristics than practical deployments with integrated antennas arrays and do not encounter environmental propagation effects. Moreover, the underlying beam training protocols such as the sector level sweep in IEEE 802.11ad [Nit+14; IEE14] are unprotected against malicious behavior. Finding the optimal antenna steering typically takes place before any secure channel establishes.

1.2 CHALLENGES AND GOALS

We derive challenges of applying direction communications in practical systems.

In the previous section, we identify the underutilization of directionality as the primary cause of low adaptation of mm-wave communications. To fully exploit the very specific spatial propagation effects and turn emerging application scenarios into practice, we derive five sub-problems and identify the most critical challenges in mm-wave communications. Based on this, we define individual sub-goals as well as the main goal of our work. Figure 1.1 gives an overview of this separation.

Strong directionality isolates concurrent transmissions and enables spatial reuse.

SPATIAL ISOLATION. Early adopters of mm-wave communication claim to achieve high directionality with pencil beams and provide strong antenna gains in the intended direction with low interference on other communications. However, we find that practical phased antenna arrays feature an imperfect directionality and expose a significant amount of side lobes in their radiation patterns. These side lobes might cause interference with transmissions of other devices in proximity and, in addition to that, compensate the directional ad-

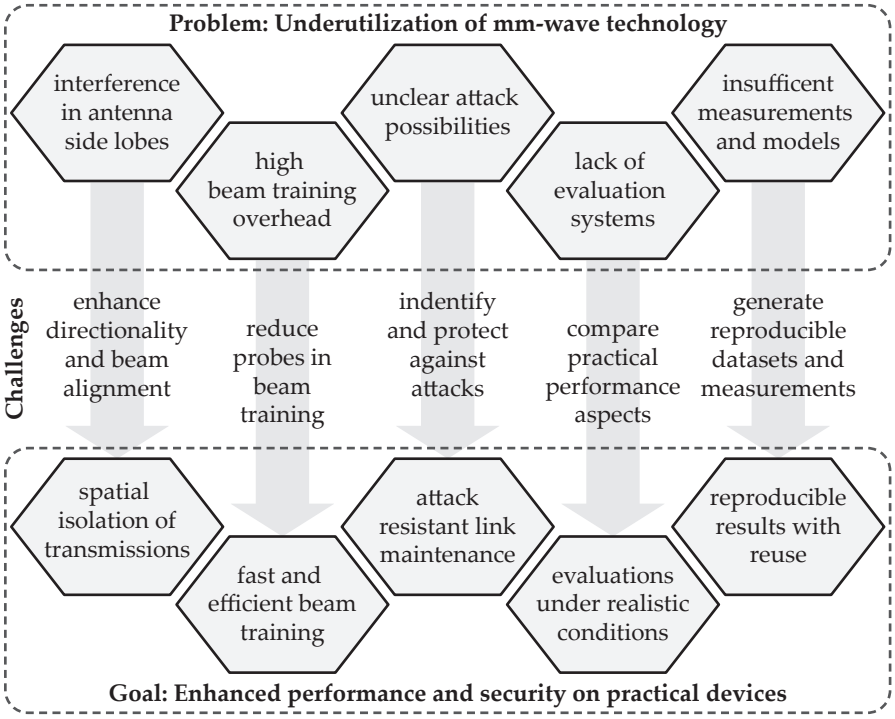

Figure 1.1: Research problems, challenges, and goals.

vantages. Such aspects must be taken into account when aiming for spatial isolation on concurrent transmissions. A strong directionality is required to spatially isolate these transmissions, which is particularly challenging with cost-efficient small size phased antenna arrays as utilized in common off-the-shelf devices. We aim at maximizing the signal gain towards the intended direction while minimizing the side lobe level. When side lobes are unavoidable, we tend to align the side lobes among multiple transceivers to improve the spatial separation. Our goal is to spatially isolate concurrent transmissions between different devices to enhance spatial reuse.

Reducing the number of probing frames increases the efficiency of beam training.

EFFICIENT BEAM TRAINING. Mm-wave communications enforce beam steering to focus the transmitted signals in particular directions. In order to find the optimal steering parameters, beam training techniques sense the current channel conditions by sending periodic probing frames in different directions. This approach, as performed by state-of-the-art beam training techniques (e.g., the sector level sweep in IEEE 802.11ad) has a high overhead. Probing a large set of narrow beams causes devices to spend a long time with frequent training. Besides improving the frame format, as suggested in IEEE 802.11ay, the number of required probing frames during the training sequence should be minimal. Our goal is to find the optimal beam fast and efficiently while achieving an accuracy as an exhaustive search.

ATTACK RESISTANCE. Featuring a high directivity and being susceptible to blockage by objects, mm-waves are often assumed to be hard to intercept. However, it is unclear how non-ideal radiation characteristics of phased array antennas and low-layer beam training protocols affect malicious attacks. The challenge is to identify specific attack strategies that target the directionality by considering environmental reflections and vulnerabilities in the beam training itself. Mm-wave communications require amendments in lower layer protocols to handle the directionality and perform the beam training. By analyzing and identifying possible attack vectors, we aim at protecting these protocol extensions against adversarial behavior. Our goal is to obtain an understanding of the security aspects of directional communications and achieve attack-resistant link maintenance to prevent attackers from tampering with the directionality.

Attack resistance requires to identify emerging attack vectors first.

PRACTICAL EVALUATION. One of the significant challenges in developing novel communication protocols and applications is the unavailability of proper development and evaluation systems. An extensive evaluation of mm-wave communications with available systems is nearly impossible. On the one hand, current Software-Defined Radio (SDR) based evaluation and development systems lack sophisticated implementations above the link layer. Due to high expenses, they are only affordable for single link setups. Consumer grade off-the-shelf devices, on the other hand, typically come as black-boxes and lack access to the lower layers of the communication stack. To address all aspects of mm-wave propagation, we need a holistic full-stack evaluation. The envisioned system provides access to all relevant communication layers. It allows for optimizing the beam training mentioned above and directly comparing results with those of a standard reference implementation. Moreover, an evaluation should also consider realistic deployments and not be limited to single link scenarios. In this regard, our goal is to set up a practical evaluation testbed with standard compliance under realistic conditions.

Directional communications require a holistic evaluation in practical environments.

REPRODUCIBILITY. Although many researchers investigate mm-wave communications, only a few models and extensive data sets that reveal the diverging communication paradigms of directionality are publicly available. Most of these models cover very specific scenarios [Mal10] without general applicability. Due to the lack of alternatives, many research groups engineer custom prototyping systems from scratch and use those to perform measurements in small office environments to base their studies. We find that this approach has a limited reproducibility as it requires to assemble custom hardware components to re-build the systems that are often poorly documented. Moreover, this approach is an unnecessary repetition of existing work. Publicly available data sets with extensive measurements and traces in

Reproducible measurements and results enable continued use.

realistic scenarios would allow researchers to validate their theses and approaches without the necessity to create yet another prototyping system. To this end, we emphasize the release of tools and measurements to build on existing results. In other words, our goal is to generate reproducible results that allow for continued use.

Our goal is to enhance the performance and security of practical mm-wave communication systems.

MAIN GOAL. The main goal of this thesis is to enhance the performance and the security of practical mm-wave communication systems. To achieve this, we follow a holistic approach that includes a profound experimental evaluation and tackle the challenges of spatial signal propagation in high-frequency bands as mentioned above. The diverse aspects require a cross-layer consideration as the typical isolation of layers in the network stack becomes inappropriate. For example, beam training should consider the physical appearance of the antenna and its steering capabilities. Moreover, the antenna steering should adapt to the current environment and network aspects. By jointly achieving spatial isolation, efficient beam training, and attack resistance, we provide the foundation to enable emerging applications that fully exploit the spatial sparsity in mm-wave signal propagation. With our focus on practical evaluation and reproducibility, we enable other researchers to benefit from our findings and continue to study related aspects.

1.3 CONTRIBUTIONS

To address the challenges above in mm-wave communications, we *(1)* propose a practical evaluation framework, *(2)* introduce performance enhancements, *(3)* identify and address security aspects, and *(4)* provide practical measurements and traces. Our contributions are described in the following.

1.3.1 *Framework*

Our first contribution is a holistic framework that allows to analyze the signal propagation and to conduct experiments with customized and standard compliant protocols on commercial off-the-shelf devices in large-scale deployments.

We develop a channel sounding platform and a ray tracing based simulation environment.

SIGNAL PROPAGATION ANALYSIS. Current mm-wave indoor propagation analysis techniques rely on common channel models to predict the signal propagation which has limited options when it comes to more than one transmitter and receiver. Available models are rare and often limited to simple applications in a conference room, living room, or cubicle. Measurements are only available for concrete scenarios, and experimental hardware with mm-wave transceivers is expensive. To address this issue, we develop a channel sounding platform that brings the advantages of SDRs to mm-wave applications and enables

transmission of arbitrary waveforms over wireless links at 60 GHz. It supports variable data modulation schemes and adjustable frame formats to implement custom protocols. Still, channel sounding does not scale with large deployments and multiple transceivers. We present a fast deterministic image-based ray-tracing simulation framework for mm-wave propagation. It supports developing mm-wave specific protocols and, in contrast to conventional statistical models, deals with multiple transceivers. Our framework is written in MATLAB and computes the channel impulse response by taking into account the antenna radiation characteristic and alignment. The strengths of our simulation constitute signal variations at different receivers and interference of multiple transmitters which are crucial in dense deployments.

Our testbed experimentation platform consists of off-the-shelf devices.

TESTBED EXPERIMENTATION. Even though consumer-grade mm-wave devices are commercially available, no suitable research testbed exists that fully complies with IEEE 802.11ad by now. We present our practical experimentation and evaluation platform that bases on off-the-shelf devices. Using a commodity tri-band wireless router, we achieve standard compliance with IEEE 802.11ad. Through software adjustments, we provide full system access to the router's operating system and interface the protocol execution in the 60 GHz Wi-Fi chip. By patching the binary firmware that is running in the chip, we obtain access to the beam training configuration and implement custom beam patterns on the phased antenna array. Doing so, we bare the full potential of the integrated hardware components and provide a low-cost alternative to expensive mm-wave evaluation systems. Using this hardware, we investigate the performance and security in IEEE 802.11ad communications.

Firmware modifications integrate new beam training features.

1.3.2 *Performance*

State-of-the-art beam training causes a high overhead.

Achieving data-rates of multiple Gbps in mm-wave communication systems requires efficient and accurate beam training algorithms. To find the steering direction on IEEE 802.11ad compatible devices, state-of-the-art approaches sweep through a set of predefined antenna sectors with generic beam patterns. This approach is robust and straightforward but causes a high training overhead and provides only sub-optimal antenna beams. The antenna modules typically deployed in such devices are capable of generating much more precise antenna beams. We provide sophisticated approaches and optimize the beam training for commercial off-the-shelf devices. In particular, we *(1)* increase the efficiency with compressive sector selection, *(2)* mitigate lateral interference in antenna side lobes, and *(3)* improve the directional antenna gain with adaptive codebook optimization.

Compressive sector selection reduces the training time.

COMPRESSIVE SECTOR SELECTION. To improve the efficiency of beam training on IEEE 802.11ad devices, we adopt compressive path tracking for the sector selection in off-the-shelf devices. In contrast to existing solutions, our compressive sector selection tolerates the imperfections of low-cost hardware, tracks beam directions in spherical dimensions, and does not rely on pseudo-random beams. We select the best sector based on measured radiation patterns and sweep only through a subset of probing sectors.

We steer beams away from interference directions.

MITIGATING LATERAL INTERFERENCE. To minimize the interference from lateral directions, we present an adaptive beam switching mechanism. Our mechanism steers receive beams to maximize the signal gain but also to minimize interference via both their main and side lobes. Doing so, we enable efficient parallel operation of incompatible standards such as WiGig and IEEE 802.11ad.

Optimized beams adapt to the channel and maximize the signal strength.

ADAPTIVE BEAM OPTIMIZATION. To improve the accuracy of selected beam patterns, we adaptively adjust the sector codebook to optimize the transmit beam patterns for the current channel. Isolating the phase and magnitude from individual antenna elements in the array, we extract the full Channel State Information (CSI) and dynamically compute a transmit beam pattern that maximizes the signal strength at the receiver. Thereby, we automatically exploit reflectors in the environment and improve the received signal quality.

1.3.3 *Security*

Our security consideration includes *(1)* the feasibility of eavesdropping on reflections, and *(2)* attack vectors in the beam training that allow attackers to tamper with the beam selection.

Reflections facilitate eavesdropping from distant locations.

EAVESDROPPING ON REFLECTIONS. Featuring a high directivity and being susceptible to blockage by objects, mm-waves are often assumed to be hard to intercept. However, mm-waves reflect on a variety of objects which might have a significant impact on the security. With practical experiments using our framework, we reveal that small-scale objects within the beam cause reflections and enable eavesdropping from remote locations. In contrast to large blocking obstacles, we consider objects that are sufficiently small not to impede the communication between the intended transceivers. Still, they are sufficiently large to enable an eavesdropper to decode a reflected signal. With this approach, we practically demonstrate the vast impact that inconspicuous objects have on mm-wave security.

BEAM STEALING. The low layer amendments in IEEE 802.11ad to handle directional communications lack proper security mechanisms and disclose unprecedented attack possibilities. Distant attackers might tamper with the beam training and literally 'steal' the beam from other devices. We investigate the threat of such beam stealing attacks that intercept the sector level sweep. Injecting forged feedback forces victims to steer their signals towards the attacker's location. Using our framework, we evaluate the impacts on eavesdropping and acting as a Man-in-the-Middle (MITM). To protect against this kind of attack, we discuss possible detection metrics. We extend the prevalent sector level sweep with an authentication scheme to ensure that devices only accept the feedback from their intended peers. Thereby, we emphasize the threat of beam stealing on mm-wave networks and propose an authentication to counterfeit such an attack.

Attackers may tamper with the beam training.

We propose an authentication scheme to protect the beam selection.

1.3.4 *Measurements*

To complement our findings and allow other researchers to build on our work, we conduct measurements campaigns and process our results for integration in future work. In particular, we provide *(1)* the antenna radiation patterns of predefined and custom beams along with the array factor of individual antenna elements and *(2)* performance traces in practical deployments.

ANTENNA RADIATION PATTERNS. Using our testbed system, we precisely measure the predefined beam patterns of commercial off-the-shelf hardware in an anechoic environment. Our results show irregular shapes that lead to significant side lobe levels but cover the whole area around the device. Moreover, we measure the antenna array factor that reveals the phase difference of individual antenna elements in various directions. Our measurements allow to design and shape custom beam patterns with arbitrary directionality and implement them for communication on the devices in practical environments. To illustrate the possible gains in directionality, we also measure exemplarily generated beam patterns that expose a strong main-lobe.

Measurements in an anechoic chamber reveal the radiation patterns.

PRACTICAL DEPLOYMENT. The practical impact that is achievable with optimized beam pattern enables new applications. To highlight the performance differences of generic beam patterns and highly directional beams, we perform an extensive measurement campaign in an atrium with 28 distributed devices. Between these devices, we measure the signal strength and the CSI and investigate the throughput that is achieved with different beam configurations. Operating multiple links in parallel, we reveal the limitations of spatial-reuse. With directional beams, side lobe suppression, and multi-cast patterns, our experiment highlights the advantages of beamforming on off-the-

We perform practical measurements with 28 devices in an open atrium.

shelf mm-wave hardware. Our obtained dataset allows to analyze the impacts of custom beam steering on the network performance in post-processing. We reprocess our data-set for easy integration in custom studies.

1.4 OUTLINE

The contents of this thesis are separated in different parts.

The remainder of this work is structured as follows and illustrated in Figure 1.2. In Chapter 2, we state background information on mm-wave communications and summarize related work. Part II covers our evaluation framework for signal propagation analysis in Chapter 3 and practical testbed experimentation in Chapter 4. Part III describes performance improvements. We propose our compressive sector selection in Chapter 5 and mitigate lateral interference in Chapter 6. Our adaptive beam optimization approach is described in Chapter 7. Part IV addresses security aspects. We practically investigate eavesdropping on reflected signals in Chapter 8. Chapter 9 considers beam stealing attacks and proposes corresponding countermeasures. Our measurements and traces in Part V cover the antenna radiation patterns in Chapter 10 and practical deployments in Chapter 11. In Part VI, we discuss our findings. Chapter 12 states specific applications and directions for future work. Finally, Chapter 13 concludes this thesis.

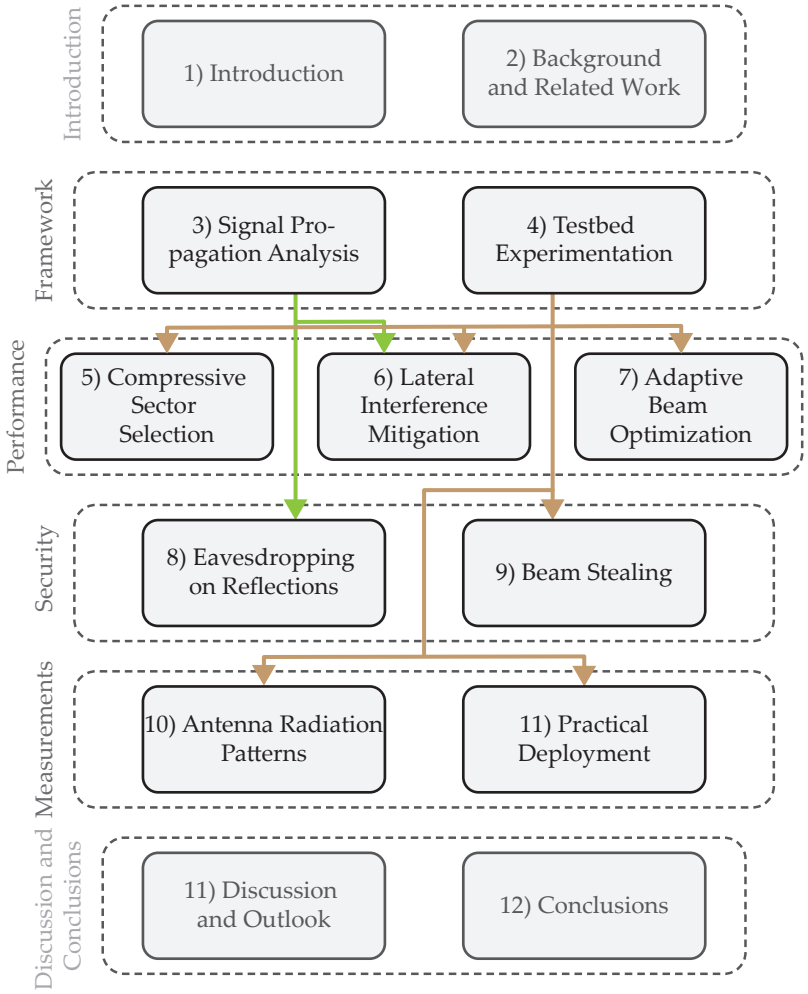

Figure 1.2: Structure and overview of contributions and chapters in this thesis. Annotations indicate the use of framework components.

2

BACKGROUND AND RELATED WORK

In this chapter, we provide background information on mm-wave communication systems and state related work. Section 2.1 describes the characteristics of the mm-wave channel, measurements, and propagation models. Details on the current Wi-Fi standardization within IEEE 802.11ad/ay are provided in Section 2.2. Recent performance improvements for beam training and maintenance are described in Section 2.3. In Section 2.4, we review related security aspects and, finally, summarize existing testbed and evaluation systems in Section 2.5.

2.1 CHANNEL CHARACTERISTICS

The channel characteristics of mm-wave communication systems are fundamentally different from those in the prevalent sub-6 GHz bands. In the following, we state recent work on channel measurements and propagation models that encounter this aspect.

2.1.1 *Measurements and Analyses*

Mm-waves have different propagation characteristics than lower frequency signals.

Early measurements of mm-waves in the 60 GHz band are taken by Smulders and Wagemans in [SW92]. On short distances, weather effects and atmospheric absorptions have little effect [MMI96], but human blockage can be severe [Sin+09]. While small objects cause diffraction [Jac+12; Kle+12], its effects are less relevant for mm-waves than in lower frequencies [Mal+10b]. Several works analyzed how mm-wave signals reflect on building materials and indoor structures [LLH94; Ahm+09; Sat+97]. Their findings imply that reflections should not be neglected in indoor environments. Most of these reflections enable communication via indirect line-of-sight paths. Ansari et al. [Ans+15] confirm that mm-wave transmissions are feasible when the reflections are direct. Polarization of signals is analyzed by Maltsev et al. in [Mal+10c] and turns out to be important due to low multipath effects. Interference is a strong limiting factor in mm-wave networks [SF15]. Especially consumer grade antenna arrays exhibit several side lobes [Nit+15a] that potentially distort concurrent transmissions. All these works outline the fundamentally different aspects of mm-waves that need to be considered in realistic propagation analyses.

2.1.2 *Propagation Models*

Simulations of mm-wave propagation can be performed with statistical or deterministic models. While the former use measured statistics from particular scenarios, the latter draws on physical theories.

Statistical models enable fast computations but are limited to scenarios of the underlying measurements.

STATISTICAL CHANNEL MODELS. Statistical models allow for fast computation. They typically utilize statistics obtained from physical measurements in concrete scenarios. A seminal statistical model was investigated by Saleh and Valenzuela in [SV87] and assumes that multi-path components of the channel arrive in clusters. Amplitudes of the signal within each cluster follow a Rayleigh or Rician distribution. In either case, the phases of received signals are uniformly distributed. The channel models for IEEE 802.11 [IEE04] adopt statistical predictions to Wi-Fi-based Multiple-Input and Multiple-Output (MIMO) communication systems in the 2.4 GHz band. Their MATLAB implementation [Mat] provides a fundamental design and analysis tool for low layer Wi-Fi aspects. The indoor simulation tool SIRCIM of Rappaport, Seidel, and Takamizawa [RST91] predicts statistical channel impulse responses in line-of-sight and non-line-of-sight scenarios. It is designed for mm-waves at 60 GHz but also works for lower frequencies. Their model produces realistic multi-path channel characteristics with high precision. A complete statistical channel model for indoor environments is established in [Mal+10b], confirming the quasi-optical nature of mm-waves and negligibility of diffraction. The model in [Mal10] specifically predicts the channel statistics for IEEE 802.11ad systems. This approach is similar to earlier IEEE 802.11 models [IEE04] but also considers directional antennas, variable beamwidth, and antenna orientation. Its implementation [ML10] makes these insights directly accessible for a variety of MATLAB simulations. Unfortunately, statistical propagation models are limited to scenarios similar to those in the underlying measurements. Detailed empirical analyses are required to simulate the channel in complex environments.

Ray tracing is a practical approximation of quasi-optical characteristics.

DETERMINISTIC PROPAGATION MODELS. Deterministic models are based on the theory of electromagnetic wave propagation that could be solved with Maxwell's equations. Unfortunately, this approach is impractical because of the high computational requirements. However, as mm-wave propagation behaves quasi-optically [Mal+09], ray tracing becomes a practical approximation. Ray tracing is typically applied for image generation in computer graphics and effectively models reflections and attenuations. In wave propagation analytics, it predicts the received signal strength at a specific position in a given environment.

Ray tracing can be divided into Shooting and Bouncing Rays (SBR) and image-based methods. SBR methods intuitively launch rays in

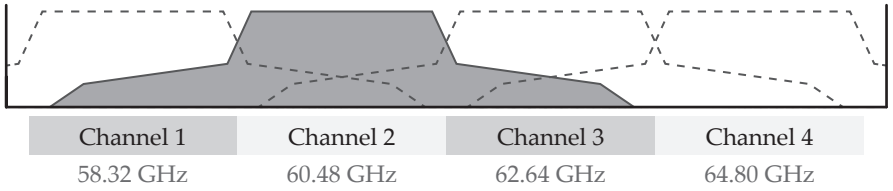

Figure 2.1: Available IEEE 802.11ad channels in the 60 GHz band.

many directions. They are reflected by objects and eventually hit the receiver. With a high number of rays, SBR methods achieve a high precision and in addition to that encounter diffraction and scattering. However, only a small fraction of computed rays contribute to the received signal; most of them never reach the receiver. In contrast, image-based ray tracing only considers direct reflections between the transmitter and the receiver. Receivers' locations are mirrored on the objects surfaces to derive the reflection paths.

Image-based ray tracing exposes only small inaccuracies.

The feasibility of mm-wave ray tracing to identify clusters of arrivals is verified by Neekzad et al. in [Nee+07]. Comparisons with measurements in this work indicate that ray tracing provides precise results for scatter-free line-of-sight scenarios. Their approach exhibits slight inaccuracies only with heavy scattering in non-line-of-sight paths. Fortune et al. [For+95] propose an SBR-based ray tracing tool for 2 GHz and 900 MHz bands. Measurements of Xu, Kukshya, and Rappaport in [XKR02] confirm that image-based ray tracing can determine the majority of multi-path components in mm-wave transmissions. Calculations of first- and second-order reflections are sufficient for line-of-sight applications [XKR02]. A hybrid approach leveraging both, SBR- and image-based methods, is proposed by Peter, Keusgen, and Felbecker in [PKF07]. Their approach is verified with practical measurements and reveals only small deficits in time domain predictions. In all these works, image-based ray tracing provides an efficient tool for coarse-grained channel modeling.

2.2 WI-FI STANDARDIZATION

IEEE 802.11ad specified mm-wave communication for Wi-Fi application.

IEEE 802.11ad [Nit+14; IEE14] standardizes the usage of mm-wave communication for Wi-Fi networks in the unlicensed spectrum between 57 GHz and 66 GHz. The standard was ratified in 2012 and specifies amendments for wireless IEEE 802.11 networks to operate at such extremely high frequencies with directional links. As illustrated in Figure 2.1, it splits the available spectrum into four channels separated by 2.16 GHz each. Since operating at high frequencies requires narrows beams to compensate the high attenuation, the standard integrates a couple of new extensions. In the following, we highlight the most critical specifications in IEEE 802.11ad and summarize the antenna model, frame format, and beam training.

MCS	MODULATION	CODING RATE	SENSITIVITY	DATA RATE
1	π/2-BPSK	1/2	−68 dBm	385.0 Mbps
2	π/2-BPSK	1/2	−66 dBm	770.0 Mbps
3	π/2-BPSK	5/8	−65 dBm	962.5 Mbps
4	π/2-BPSK	3/4	−64 dBm	1155.0 Mbps
5	π/2-BPSK	13/16	−62 dBm	1251.3 Mbps
6	π/2-QPSK	1/2	−63 dBm	1540.0 Mbps
7	π/2-QPSK	5/8	−62 dBm	1925.0 Mbps
8	π/2-QPSK	3/4	−61 dBm	2310.0 Mbps
9	π/2-QPSK	13/16	−59 dBm	2502.5 Mbps
10	π/2-16-QAM	1/2	−55 dBm	3080.0 Mbps
11	π/2-16-QAM	5/8	−54 dBm	3850.0 Mbps
12	π/2-16-QAM	3/4	−53 dBm	4620.0 Mbps

Table 2.1: Supported single-carrier modulations and coding schemes in IEEE 802.11ad [IEE14].

2.2.1 *Antenna Model*

Directional antennas can be steered in different sectors.

Operating at high frequencies requires directional antennas that are steerable in different directions. IEEE 802.11ad abstracts this aspect and separates the coverage of an antenna into sectors that cover only a specific area in one angular direction. Each of these sectors is defined by a steering direction and a beamwidth. At maximum 128 sectors are supported. Thus, sectors can be as narrow as 2.8° to achieve 360° coverage. Moreover, IEEE 802.11ad introduces quasi-omnidirectional sectors with beams that exhibit low gain variations over a wide angular range. Such beams are applied whenever the direction of communication is unknown. Since omnidirectional communication is inefficient at high frequencies, at least one of the transceivers should select an antenna sector with a directional beam.

2.2.2 *Frame Format and Modulation Schemes*

IEEE 802.11ad introduces a new DMG frame format for directional transmissions.

Communicating over directional links requires new physical layer frame formats. The IEEE 802.11ad standard specifies three different Directional Multi-Gigabit (DMG) frame formats for control, single-carrier, and low-power transmissions. In the initial version, it specified an additional frame format for multi-carrier transmissions that meanwhile became obsolete. Only control and single-carrier frames and mandatory to implement, the low-power mode is optional. Control frames are used for signaling and apply the most robust Modulation and Coding Scheme (MCS). Other frames carry an arbitrary payload and support different data rates. The standard specifies a bandwidth

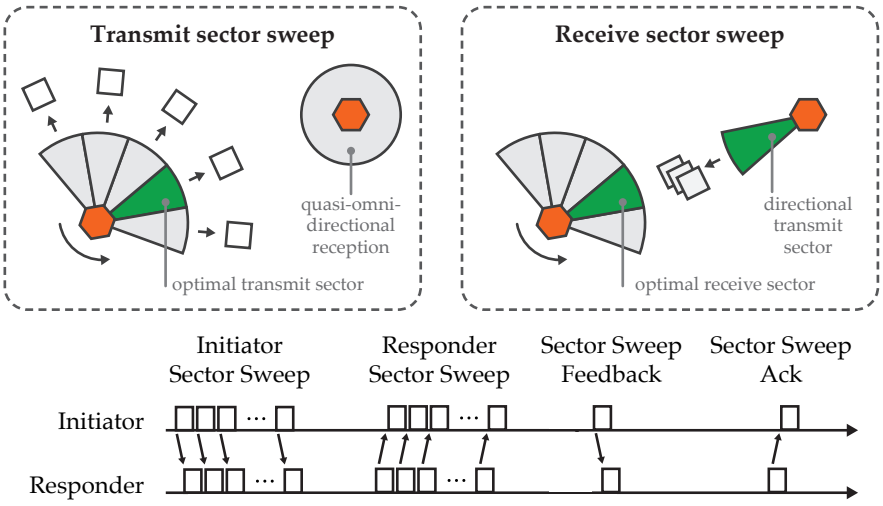

Figure 2.2: Mutual beam training in the IEEE 802.11ad sector level sweep between an initiator and responder.

of 1760 MHz and defines different MCS that provide single-carrier data rates from 385 Mbps to 4620 Mbps. Table 2.1 summarizes the supported combinations of modulations and coding rates. Each frame consists of a preamble, header, data, and an optional training field. The preamble uses a Short Training Field (STF) for frame detection and Automatic Gain Control (AGC). An additional Channel Estimation (CE) field enables receivers to estimate the channel and correct frequency offsets. Both fields are constructed mathematically from Golay sequences and achieve good auto-correlation characteristics. The header contains general information such as the used modulation, the data length, and a checksum. It has a length of 64 bit for single-carrier and 40 bit for control frames. The data field encodes a payload with variable length. The training field is used to adapt the beam configuration during transmission.

Single-carrier transmissions achieve data rates of up to 4620 Mbps.

2.2.3 *Beam Training*

The IEEE 802.11ad standard proposes a two-phase beam training approach [Nit+14]. First, the sector level sweep determines an initial coarse-grained antenna sector configuration. Then, the beam refinement phase fine-tunes the selected sectors to obtain a well-aligned pair of directional beams.

Beam training is performed with a sector level sweep and a beam refinement.

The sector level sweep performs a beam training with predefined sectors. In obtains the steering direction for the antenna by sweeping through all available sectors while sending probing frames. The training is performed mutually between an initiator and responder, where both either determine their best transmit or receive sectors. A responder sweep directly follows that of an initiator. The typical sequential operation is illustrated in Figure 2.2. In the transmit sector

The sector level sweep probes predefined antenna sectors.

sweep, as depicted on the left in Figure 2.2, devices send probing frames on different sectors while the receiver is listening in quasi-omnidirectional mode. Each frame is marked with a specific ID such that the best sector can be identified. During the receive sector sweep, the transmitter uses a fixed directional beam—the one that performed best during the transmit sweep—while the receiver sweeps through different receive sectors as shown on the right in Figure 2.2. The optimum Signal-to-Noise Ratios (SNRs) and, in case of a transmit sweep, the IDs of the best sectors are reported to the other device. The feedback from the initiator is carried in all frames during the responder sweep, while the responder feedback is only transmitted once in a specific feedback frame at the end of the sequence. Finally, the sweep acknowledgment completes the sector training and negotiates parameters of an additional beam refinement.

Beam refinement continuously adjusts the antenna steerings.

The beam refinement phase optimizes the antenna weights and fine-tunes the beam alignment between transceivers. This operation is independent of the sectors mentioned above that have a specific geometrical shape. Using the training fields specified in the DMG frame format, devices evaluate alternative beam configurations during regular data transmission. Doing so, devices continuously adapt their antenna configuration to compensate for mobility and small channel distortions. Devices need to change their antenna configuration rapidly as multiple settings are probed within a single frame. However, due to the complexity, beam refinement is optional to implement on standard compliant devices. Most of them only train their transmit sectors and use a quasi-omnidirectional beam for receiving.

2.2.4 *Future Enhancements*

IEEE 802.11ay introduces several enhancements for efficient beam training and a new frame format.

For a low number of sectors, the sweep completes in a reasonable amount of time. However, since its complexity increases linearly with the number of probed sectors, it leads to high overhead in case of frequent beam sweeping. Current devices keep the number of sectors low and prefer wide sectors to handle the training overhead. Optimizations are likely to appear as part of the upcoming IEEE 802.11ay standard [Gha+17]. It introduces several refinements to make beam training more efficient. Besides a short frame format for fastened sectors sweeps [EC16], it defines an enhanced-DMG frame format. Additional extensions cover support for multiple antennas which enable directional MIMO over multiple spatial streams. Furthermore, throughput enhancements are possible by channel bonding that combines the bandwidth of multiple channels for a single transmission. Besides these advantages, researchers develop novel ideas for beam management not limited to Wi-Fi networks as described in the next section.

2.3 BEAM MANAGEMENT

Optimal algorithms for sector selection and beam maintenance in mm-wave communication systems are heavily discussed in the research community. In the following, we present selected approaches for efficient training and adaptive tracking.

Hierarchical beam structures reduce the training complexity.

EFFICIENT BEAM TRAINING. The simplest beam training method is to probe all possible steering combinations which cause a high overhead. Several protocols and improvements have been proposed that allow devices to steer their antennas more efficiently. Phased array antennas that are used in many devices achieve beamforming by weighting individual antenna elements. Combining the characteristics of these elements enables new training concepts [KS16]. Junyi Wang et al. [Jun+09] found that the optimal number of probing beams must be only twice the number of antenna elements to efficiently estimate the path direction [Jun+09]. Hierarchical codebook structures of beams with different beamwidths allow to refine the beam training accuracy iteratively [Hur+13; Alk+14a; Alk+14b; NZL17; AH16]. For instance, a hierarchical beam-search [Hur+13] probes sectors with wide beams first and then continues with sub-sectors of narrower beams. Organizing beam patterns in a tree hierarchy achieves logarithmic search complexity. However, it introduces an additional communication overhead due to the required feedback for each of the multiple probing rounds.

Compressive sensing approaches derive the optimal antenna steering from a few random probes.

Compressive sensing based path tracking protocols as proposed in [MRM16; RVM12a; RVM12b] achieve the same complexity but only require a single training round. They derive the CSI from magnitude and phase measurements of pseudo-random probing beams. As phase information are seldom accessible on practical systems, Rasekh et al. [Ras+17] propose a variant that only relies on non-coherent signal strength measurements. Despite lacking phase information, Rasekh et al.'s approach achieves high accuracy for line-of-sight and direct reflection paths but cannot distinguish multi-path components.

Environmental changes and mobility require adaptive approaches.

ADAPTIVE BEAM TRACKING. Efficient training alone is insufficient to handle mobility and blockage. Adaptive beam maintenance ensures connectivity in scenarios with changing environmental conditions. Duan et al. [Dua+15] demonstrate the feasibility of device tracking and obtain the angle-of-departure and angle-of-arrival of communication links. Algorithms such as *BeamSpy*, proposed by Sur et al. [Sur+16], improve resilience to blockage by instantaneously predicting the availability of alternative paths. Zhou, Zhang, and Ma [ZZM17] consider the mobility of devices and correlate channel profiles at nearby locations to reconstruct the path when devices start to move. Their proposed protocol continuously realigns the links of mobile devices

without explicit channel probing. To reduce the risk of disconnections due to mobility, Haider and Knightly [HK16] adaptively control data-rates and beamwidths during communications. Their approach allows to widen the beam and lower the data rate before an outage occurs. Likewise, Sur et al. [Sur+15] study the possibility of switching to a wider sector when a blockage is detected to recover a communication link. In [WZZ17], Wei, Zhou, and Zhang sense the environment and trace the line-of-sight and non-line-of-sight paths between devices. By applying common ray tracing techniques, their approach determines the most significant reflections in the paths. It predicts the quality of different links and optimizes the deployment of APs. Inserting additional pilot tones, Araújo et al. [Ara+14] increase the channel estimation accuracy with low overhead. Tracking link directions without any probing overhead is possible by explicitly generating multi-lobe beam patterns as shown by Loch et al. [Loc+17]. Their mechanism detects if devices move out of the beam's coverage area and adjust the steering accordingly. The approaches in [An+09; Gen+10] suggest learning mechanisms to detect a suitable non-line-of-sight path when the direct path is blocked. Sur et al. [Sur+16] instantaneously predict the signal quality and discover an alternative beam when the primary link fails. In cases where the receiver is unreachable through any reflections, an indirect connection could be established through intermediate devices [Sin+09; YDX15].

Lower frequency systems can provide a coarse steering direction.

ASSISTED BEAM STEERING. Beam steering algorithms also benefit from multi-band connectivity. Wi-Fi connections in the 2.4 and 5 GHz bands can assist mm-wave link adaption [Sur+17]. Angular estimations of the communication direction in lower frequencies already provide a coarse antenna steering direction for the mm-wave beam [Nit+15b]. Ali and Heath [AH17] follow a similar approach combined with compressed sensing techniques to keep track of the optimal beam steering direction.

Using multiple antennas with individual RF-chains enables spatial multiplexing.

MULTI-ANTENNA SYSTEMS. In multi-antenna scenarios, additional performance gains are expected. MIMO systems drive multiple antenna arrays with individual Radio Frequency (RF) signals and transmit independently in multiple directions. Hybrid beamforming facilitates such scenarios and enables parallel channel measurements to increase the estimation accuracy. Approaches such as [MG09; SR14; SR15; Aya+12a; Aya+12b] serve different users at the same time and improve spatial reusability. Structured compressive sensing approaches such as [GDW16; Gao+16] estimate the wireless channel of massive MIMO systems with few probes efficiently. Choi [Cho15] choose beams for multiple users with low interference. Theoretical considerations of the advantages of compressed sensing in multi-path sparsity have been formalized by Bajwa et al. [Baj+10]. The coding based beamforming

scheme from Tsang, Poon, and Addepalli in [TPA11] allows in-packet training and continuously switches the beams during transmissions. Unfortunately, at the time of writing, no available device supports multiple RF-chains or hybrid beamforming in practice.

2.4 SECURITY ASPECTS

Despite of high directionality, mm-waves expose the same challenges as any other wireless communication system.

Typical mm-wave applications implicitly assume that signals only propagate within a bounded area. The spatial propagation effects of mm-waves make eavesdropping and other attacks more challenging than in lower frequency bands. The wire-like connections are harder to intercept than omnidirectional transmissions [Dai+13; KHK17; Zhu+17; Zhu+16]. An attacker must typically reside in the coverage of the beam to infer the transmission. Nevertheless, mm-wave communication systems operate on a shared medium that guarantees neither confidentiality, integrity, nor availability. They are exposed to the same security challenges as any other wireless communication system.

Lower layer protection schemes ensure that signals can only be decoded at specific locations.

Lower-layer protection schemes have been proposed to secure mm-wave communications. With many antennas and different carrier frequencies, the secrecy rate of transmitted signals can be increased to impede eavesdropping [Zhu+17; Zhu+16]. Antenna subset modulation [VLH13; Elt+16] protects mm-wave communications on the physical layer against attackers at a single position [RGH15]. This approach randomizes the antenna selection and symbol modulation to cause fluctuations in the antenna's side lobe intensity. While the intended receiver maintains a constant signal gain, potential eavesdroppers are distorted by random signal variations and, thereby, prevented from decoding. Forman and Young [FY10] extract symmetric keys from an mm-wave channel impulse response. Their experiments show that moving objects and people in the environment cause significant changes in the channel which increase the key quality. Unfortunately, none of these protection schemes protects existing IEEE 802.11ad systems. Applying common lower-layer protection schemes to mm-wave communications has a low impact when the antennas are perfectly aligned. Thus, existing schemes from lower frequencies cannot be directly applied. Novel mechanisms to protect mm-wave communication based on the unique propagation characteristics are necessary [Yan+15].

2.5 TESTBED SYSTEMS

Many researchers build their own evaluation platforms.

While millimeter-wave research is well advanced, wide-spread practical testbed and evaluation systems are rare. Available IEEE 802.11ad modules expose a single RF-chain, provide limited control, and only allow to extract a few low-layer parameters [LBW16; Nit+15a]. As a result, researchers tend to create evaluation platforms that are equipped

with directional horn antennas [HK16; Sur+15; HK18; Sur+16] or custom phased antenna arrays [Ras+17; Zha+16]. Due to the high amount of hardware customization, such platforms exhibit different behavior than commercial devices. They typically lack IEEE 802.11ad compliant implementations and provide limited features. Since most of these systems are poorly documented, they additionally complicate the reproducibility of results. Commercial prototyping platforms, such that used in X60 by Saha et al. [Sah+17], allow evaluating different protocols on the lower layers but are hardly affordable for large deployments.

Implementing custom protocols in the Wi-Fi firmware is only possible on specific chips.

The firmware that is running in some specific sub-6 GHz Wi-Fi modules can be modified to support custom MAC-layer protocols [Ber+16; Ber+14; SWH17] or extract the CSI [Hal+11]. Unfortunately, such approaches require community projects and a detailed understanding of the wireless chip. Even for common Wi-Fi systems that have been around for decades, only a few chipsets to date provide accessible CSI. Manufacturers of commodity Wi-Fi equipment typically keep their systems closed. As of now, no such access is available for any IEEE 802.11ad device.

Recent mm-wave devices only provide limited features.

Recent work has shown that some low-layer statistics can be obtained from wireless docking stations operating at 60 GHz [LBW16; Nit+15a] but the features are strongly limited. In [Sur+17; WZ17], the authors perform practical evaluations with commodity routers and modify the device driver to capture extended statistics. Still, all works on commodity devices lack encompassing access to low-layer parameters, which is required to evaluate advanced beam training protocols.

Part II

FRAMEWORK

In this part of the thesis, we propose our framework and methodology to investigate practical aspects of mm-wave communications. Chapter 3 describes our channel sounding platform and environment simulation for signal propagation analysis. In Chapter 4, we present our practical testbed evaluation platform that is based on off-the-shelf IEEE 802.11ad devices with custom software modifications.

3 SIGNAL PROPAGATION ANALYSIS

We develop a channel sounding platform and a simulation tool.

In mm-wave communications, the signal propagation effects fundamentally differ from those in prevalent wireless communication technologies at sub-6 GHz frequencies. As directional links are required, reflections and blockage in the environment have a substantial impact on the received signal quality. To acquire a profound understanding of these effects in practical setups, we propose our signal propagation analysis framework that consists of *(1)* a channel sounding platform and *(2)* an environment simulation tool. Our channel sounding platform, described in Section 3.1, allows to transmit arbitrary signals and to measure the received signal strength with 60 GHz transceivers. To investigate deployments with multiple devices in complex environments, we propose our simulation tool in Section 3.2. It simulates wireless communication channels for arbitrary environments and allows to assess the impacts of blockage, reflections, and various antenna characteristics. In Section 3.3, we describe specific applications such as signal strength mapping, interfering transmissions, and common network problems to highlight the benefits of our system. Our channel sounding platform, as well as our simulations, are used throughout the contributions of this thesis to obtain an understanding of propagation effects under specific conditions.

3.1 CHANNEL SOUNDING PLATFORM

SDRs enable rapid prototyping of wireless communications.

Software-Defined Radios (SDRs) have significant advantages in prototyping wireless communication systems. They typically use Field-Programmable Gate Arrays (FPGAs) that are connected to radio frontends to transmit and receive arbitrary signals. The FPGA processes the digital signal with high reconfigurability and flexibility. SDRs that use this architecture efficiently allow evaluating custom protocols and signal processing techniques. For common wireless communications standards, such as IEEE 802.11a/b/g/n, predefined reference designs [WARP+] exist, which support interoperability with off-the-shelf devices. Unfortunately, this is not the case for mm-wave communication systems. Existing solutions such as [Sah+17] are expensive and limited in configurability. In our mm-wave channel sounding platform, we combine the advantages of a wide-spread SDR platform with commercial 60 GHz radio transceivers. We enable the transmission and the reception of arbitrary waveforms over wireless links at 60 GHz in a well known and easy-to-use SDR environment. Our platform provides full access to the physical layer and supports adjustable data

Our channel sounding platform brings SDR capabilities to mm-wave communications.

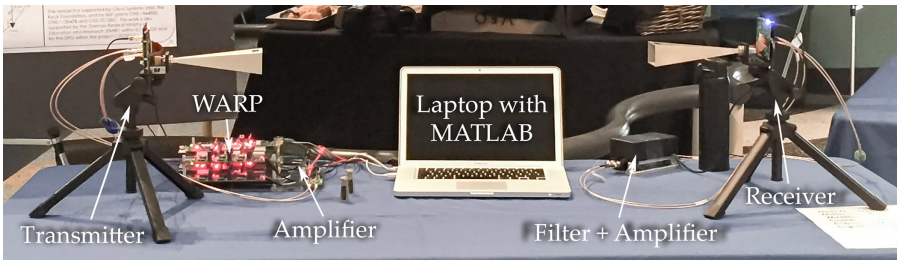

Figure 3.1: Channel sounding platform with 60 GHz transceivers and SDRs.

modulation schemes and frame formats. The hardware of our platform is composed of an SDR, commercial 60 GHz radio transceivers and custom interconnections. A typical setup is shown in Figure 3.1. In the following, we describe the hardware and the operation of our platform.

3.1.1 *Hardware Setup*

Our platform uses WARP SDRs and commercial 60 GHz radio transceivers.

The hardware components in our platform consist of SDRs, 60 GHz radio transceivers, and custom interconnecting circuits as illustrated in Figure 3.2. We use a pair of SDRs from the Wireless Open Access Research Platform (WARP) [WARP] and 60 GHz radio transceivers from the Pasternack/VubIQ 60 GHz development system [Pas]. To match the Input/Output (I/O) specifications of both systems, additional circuits filter and amplify the baseband signal. MATLAB scripts generate the transmitted and process the received signals. The detailed architecture and assembly of the components in our channel sounding platform is described as follows.

WARP SDRs process raw IQ samples of analog baseband signals.

WARP SDR. WARP [WARP] is an SDR platform for rapid prototyping of wireless communication systems, developed at RICE University in Houston and maintained and distributed by *Mango Communications, Inc.* It provides excellent benefits for research on wireless communication systems in the 2.4 GHz and 5.0 GHz bands with direct access to the physical layer of the communication stack. Thereby, it allows generating arbitrary signals and to raw sample processing. Samples are represented by complex numbers with In-phase and Quadrature (IQ) components. A WARP SDR instantiates each transceiver in our system. To adapt the WARP to the 60 GHz band, we use extension boards[1] with Digital-to-Analog (D/A) and Analog-to-Digital (A/D) converters to provide in- and outputs for analog baseband signals. Moreover, we utilize the WARPLab framework that enables direct interaction and control of the WARP SDR from a single MATLAB instance via Ethernet.

1 The WARP Analog Board v1.1 with baseband in- and outputs is available at https://warpproject.org/trac/wiki/HardwareUsersGuides/AnalogBoard_v1.1.

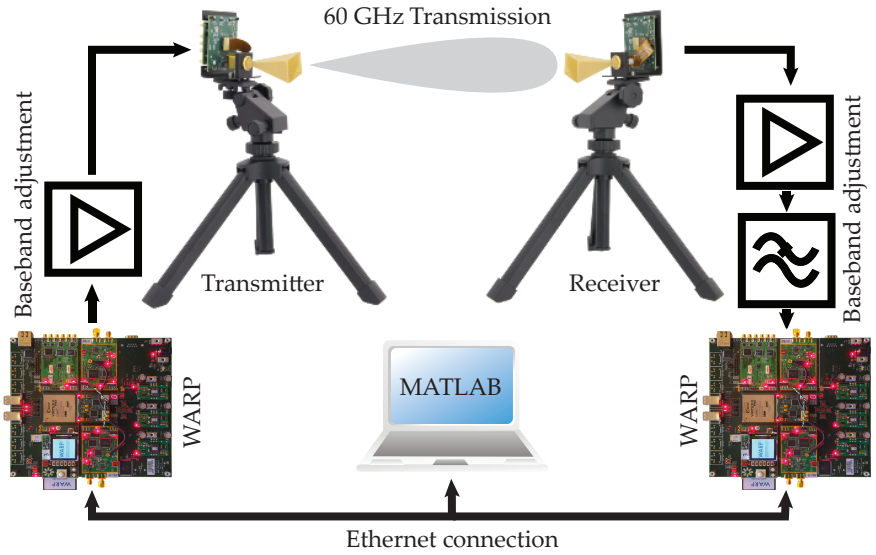

Figure 3.2: Hardware setup showing the interconnection of WARP SDRs and mm-wave transceivers.

WARPLab is a pre-configured firmware image that runs on the FPGA for buffered sample processing. A MATLAB script creates a single carrier signal, transfers it into the buffers of the transmitting SDR, and triggers the transmission. Other transceivers store the received signals in their buffers, such that they can be processed in off-line analyses. Since WARP is designed for common Wi-Fi systems, it samples the signal at 40 MHz, which is sufficient for basic encodings and channel sounding at 60 GHz as well.

The radio transceivers use horn antennas with different beam widths.

RADIO TRANSCEIVERS. The Pasternack/VubIQ 60 GHz development system consists of a transmitter and a receiver with exchangeable horn antennas and baseband I/O connectors. They internally convert the baseband signal to an Intermediate Frequency (IF) first and then to a Radio Frequency (RF) in the 60 GHz band. All mixers for the signal conversion can be configured to use an internal or external clock source. Waveguide connectors allow mounting external horn antennas, which are available with different beam widths. Additional amplifiers, attenuators, and filters enable fine-grained control of the analog signal processing chain. Throughout our experiments, we use the horn antenna with the narrowest available beamwidth of seven degrees and share the clock signal among both transceivers. To mechanically steer the signal, we mount the transmitter on an electronically controllable rotation-head. This setup enables us to orient the transmitting antenna into particular directions and investigate the impacts of antenna misalignment. Unfortunately, the I/O specifications of the transceivers do not match those of the WARP and vice versa. Thus, additional signal adjustments are required.

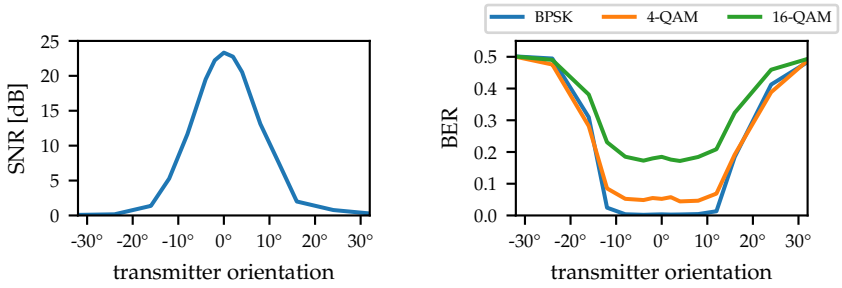

Figure 3.3: Average SNR and BER with our channel sounding platform using BPSK and QAM encodings and a 7° antenna. The transceivers are placed at 2 m distance while the transmitter takes different orientations.

Signal adjustments are required to match the I/O specifications.

INTERCONNECTIONS. The output signals of the WARP are biased and more powerful than the input specifications of the 60 GHz radio transceivers allow. By adjusting the signal levels, we protect the hardware and optimally use the dynamic range. At the transmitter side, our adjustments convert the single-ended output of the WARP to a differential signal and attenuate the magnitude. At the receiver, the signal is amplified again to obtain the maximum resolution at the WARP's A/D converters. An additional low-pass filter suppresses mirror frequencies and prevents aliasing effects. For all these adjustments, we utilize common differential operational amplifiers ('op-amps') and passive electronic components mounted on custom circuit boards. Doing so, we achieve interoperability between the radio transceivers and the WARP and, thus, create our mm-wave SDR for channel sounding and signal propagation analysis.

3.1.2 *Platform Operation*

Our implementation uses single-carrier transmissions with BPSK *and* QAM *encodings.*

Our implementation incorporates basic data encoding and decoding. It supports single-carrier transmissions with Binary Phase-Shift Keying (BPSK) or Quadrature Amplitude Modulation (QAM). Transmitted frames contain a pre-known BPSK encoded preamble for delay detection and channel equalization. Receivers compute a signal strength indicator from the received preamble and determine the Bit Error Rate (BER) by decoding all data symbols in the frame. Figure 3.3 shows the achievable Signal-to-Noise Ratio (SNR) and BER as experimentally obtained by placing the transceivers 2 m apart. In this setup, we first measured the noise floor at the receiver while keeping the transmitter disabled. The presented SNR is the received signal strength relative to this reference. While rotating the transmitter, we track the SNR during the transmission. Under perfect antenna alignment, our platform provides an SNR of 23.3 dB. The shape of the measurements visualizes the directionality of the transmitted signal and matches the stated

characteristics of the horn antenna with a beamwidth of about seven degrees. Evaluating the transmitted data, we obtain a BER of 0.002 for BPSK encoding. With 4-QAM and 16-QAM the BERs increase to 0.044 and 0.171, respectively. In all encodings, impairments occur for alignment offsets larger than ten degrees. High phase-noise in the transceivers limit the data rates, such that the BERs do not match the beamwidth and saturate at the values stated above. Nevertheless, our setup is suitable for channel sounding of a single link and allows to investigate the propagation effects with reflections and blockage. Simultaneous transmissions between multiple devices and complex environment are considered by our simulation in the next section.

3.2 ENVIRONMENT SIMULATION

Channel sounding is inappropriate to analyze complex environments.

Using ideal horn antennas, our channel sounding platform, described in the previous section, provides highly directional mm-wave transmissions. Signals are steered directly to their destination with low overshoot and enable communication. However, using a channel sounding platform to analyze challenging scenarios is inappropriate. Practical environments expose many different link characteristics that are infeasible to measure one-by-one.

High directionality requires precise propagation models.

Comprehensive propagation models are required to assess the practicality of mm-wave communication in arbitrary setups. The statistical channel model of Maltsev [Mal10] provides a foundation to investigate mm-waves propagation according to the IEEE 802.11ad specifications. It predicts statistical channels between two nodes in well-defined setups, such as a conference or living room, but have shortcomings in complex environments with distinct reflections and blockage. It neither handles specific obstacles nor considers multiple transceivers, which are crucial for most applications. Hence, we need better models to predict mm-wave propagation in such environments and dense networks. Ray tracing techniques can handle multiple transceivers as well as obstacles easily. Indoor channels strongly depend on the precise layout of the environment [Nee+07]. This issue becomes especially critical for signals with short wavelengths, which reflect on most surfaces and behave quasi-optically [XKR02]. Even though efficient techniques exist [PKF07], no practical mm-wave ray tracing tool that focuses on different antenna settings and multiple transceivers is available to date.

We simulate the signal propagation with reflecting and blocking obstacles.

In the following, we propose mmTrace, an mm-wave ray tracing simulation tool implemented in MATLAB. Simulations with mmTrace predict the Channel Impulse Response (CIR) between multiple transceivers in line-of-sight and non-line-of-sight scenarios. By applying the image-based ray tracing approach with different antenna configurations, mmTrace enables the analysis of interference and location-dependent signal characteristics. It supports a variable num-

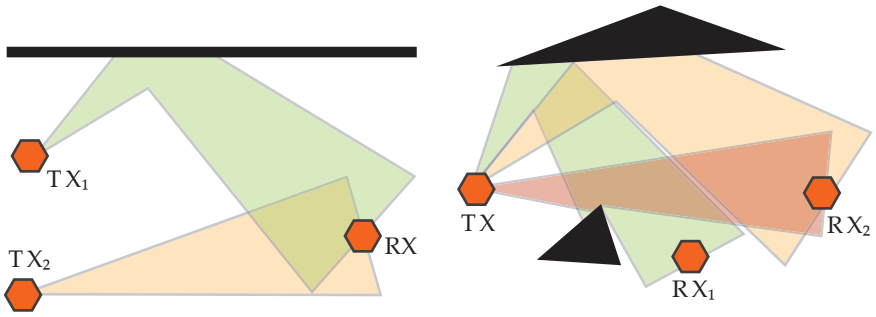

(a) Interfering transmissions (b) Multiple receivers

Figure 3.4: Simulation scenarios with multiple transmitters or receivers.

ber of reflecting and blocking obstacles in various two-dimensional environments. Modeling interfering signals, as illustrated in Figure 3.4a, becomes a simple addition of two signals. Similarities and differences between two receivers that overhear the same signal on independent paths, as sketched in Figure 3.4b, are also measurable. As the effects in such situations highly depend on the antenna orientation and environment, they cannot be obtained from statistics. With mmTrace, we achieve a precision to predict the signal strength and delay spread comparable to that of common statistical models. Additionally, it adapts to various environments and handles multiple transceivers in parallel.

mmTrace traces the paths between transceivers, derives channel models, and characterizes the channels.

The simulation of the channel effects in mmTrace is separated into three stages. First, mmTrace traces the paths between all transceivers in the environment by applying image-based ray tracing. Second, it models the channel for individual paths between the transceivers by taking into account the antenna gain, reflection coefficients, and path attenuation. The third stage, finally, characterizes the derived channel and determines the CIR, and Power Delay Profile (PDP). This independent processing allows for comparison of intermediate results with other models. In the following, we describe the processing stages in detail and validate our approach by comparing our results with those of an existing statistical model.

3.2.1 *Image-based Ray Tracing*

Paths are determined by projecting the transceivers' locations in reflecting surfaces.

Including image-based ray tracing, mmTrace determines the multipath channels between transceivers over multiple reflections. Due to computing direct paths over projections, the image-based ray tracing approach is more efficient than the usually known Shooting and Bouncing Rays (SBR) technique. It considers the position, orientation, and antenna characteristics of transceivers placed in a two-dimensional environment with reflecting and blocking obstacles. Figure 3.5 exemplary considers two transceivers TX and RX, and two reflecting walls

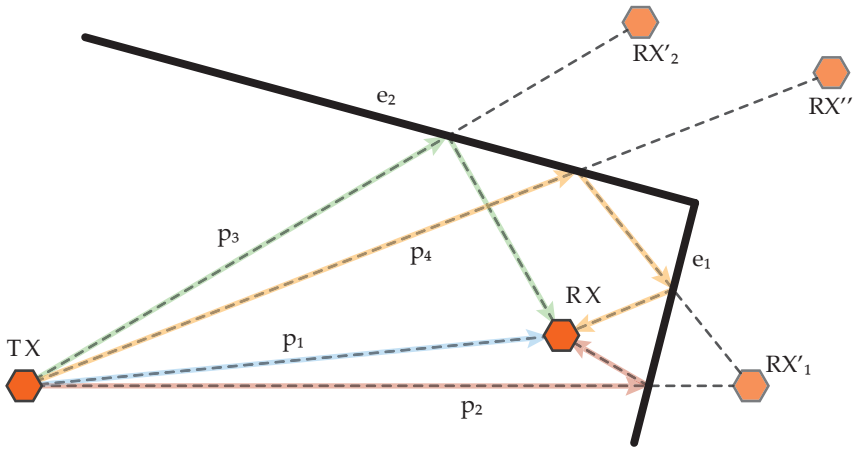

Figure 3.5: Derived paths between a transmitter (TX) and receiver (RX) with image-based ray tracing.

represented by the edges e_1 and e_2. After finding the first order images RX'_1 and RX'_1 of RX in e_1 and e_2, we determine the second-order image of RX'_1 in e_2 resulting in RX″. The paths are constructed using edge intersections towards the images. Paths that miss the reflecting edge are ignored. In doing so, the image-based approach obtains the first order reflection paths p_2 and p_3, and the second-order reflection path p_4 in addition to the line-of-sight path p_1. In mmTrace this is recursively implemented. Polygons represent the walls and obstacles in the room as a set of reflecting edges. In each recursion step, our implementation filters valid reflectors that are not entirely blocked to decrease the computation time. Thus, mmTrace constructs the paths between two transceivers as a pure geometrical problem.

Signal paths are reconstructed geometrically.

3.2.2 *Channel Modeling*

The channel representation is obtained by assigning each path a complex amplitude and delay based on the environmental properties. mmTrace computes the delay as $\tau = d/c_0$, where d is the length of the complete path and c_0 the speed of light. The complex amplitude h_k of path k is compound from multiple properties: *(1)* the antenna gain of the transmitter and the receiver, *(2)* the reflection coefficients for all reflections on the path, and *(3)* the attenuation due to the length of the path. In the following, we describe these properties in detail.

Channels are affected by antenna gains, reflections, and attenuations.

ANTENNA GAINS Antenna radiation patterns describe the gains of a signal that radiates in particular directions. mmTrace implements the three different radiation patterns shown in Figure 3.6a: *(1)* an omnidirectional pattern, *(2)* a directional pattern, and *(3)* a side lobe pattern. The omnidirectional pattern with a constant gain in all directions is the simplest model and typically applied in lower frequency systems. The directional pattern steers the radiation in a particular direction to

We implement directional and side lobe affected antenna patterns.

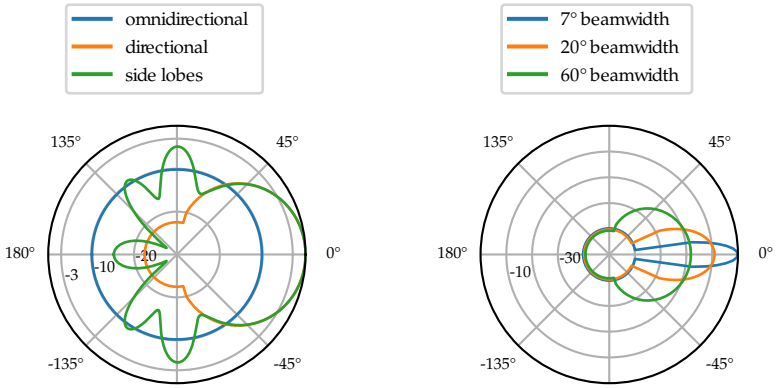

(a) Radiation patterns (b) Impact of beamwidths

Figure 3.6: Radiation patterns implemented in mmTrace.

achieve a very high gain. We implement this as specified by Maslennikov and Lomayev [ML10] in dependency of the beamwidth. Further, mmTrace models the maximum gain according to Maltsev [Mal10]. This approach leads to varying gains with different directionality as illustrated in Figure 3.6b. The side lobe pattern describes an exemplary, non-ideal antenna that exposes a main lobe in the indented direction, a back lobe in the opposite direction, and several side lobes in other directions. We integrate this as the sum of multiple directional radiation patterns with random strength and angular offset. To be independent of specific antenna models, mmTrace additionally supports the integration of custom radiation patterns such as those obtained from physical measurements.

Reflections are characterized by the angle-of-incidence and the reflectors permittivity.

REFLECTION COEFFICIENTS. Assuming that each object has a constant permittivity, we compute reflections using Fresnel's equations [Bas+09] as a function of the permittivity and the angle-of-incidence. The reflection coefficient is the product of the phase shifts and attenuations that affect the signal through all reflections on a path. As Peter, Keusgen, and Felbecker [PKF07] already revealed, making realistic model assumptions and characterizing reflections on particular surfaces is challenging. Providing realistic models that exactly match physical environments is beyond the scope of this work. We propose mmTrace as a framework that supports different reflection characteristics with arbitrary permittivity.

Attenuation is considered as path loss in the open space.

PATH ATTENUATION. Besides the effects of antenna radiation and reflection loss, signals are affected by attenuation and phase shift while propagating through the open space. This attenuation can be described by the Free-Space Path Loss (FSPL). It defines the gain of a received signal dependent on the propagating distance d, and wavelength λ. The phase shift is considered as a rotation of the signal by d/λ in the complex plane.

3.2.3 *Channel Characterization*

A channel is described by its CIR, PDP, *and delay.*

From the channel representation in the previous stage, we extract metrics that describe different characteristics of the channel shown in Table 3.1. In particular, we derive the CIR, which is the most common model to describe channel effects in the time domain. It is the sum of Dirac impulses, representing the delay, and power of all multi-path components. The PDP is the squared CIR magnitude and describes the distribution of the channel power over time. The mean delay describes the average time of arrival of the earliest significant multi-path component, while the Root Mean Square (RMS) delay spread expresses its variation. The cumulative channel power is the sum of the channel powers of all multi-path components.

3.2.4 *Model Validation*

We validate mmTrace in well-defined conference and living room environments.

To validate our ray tracing based propagation model and ensure that our simulations are correct, we compare simulated channels with those obtained from an existing statistical model. For the latter, we use the MATLAB implementation [ML10] from Maslennikov and Lomayev of Maltsev's statistical model [Mal10]. By neglecting reflections on the floor and ceiling, this model adapts to our two-dimensional environment. We rebuild the empty $4.5\,\mathrm{m} \times 3.0\,\mathrm{m}$ conference room and the $7.0\,\mathrm{m} \times 7.0\,\mathrm{m}$ living room environments from this model in our simulation. In the conference room, transceivers are located randomly on a $2.5\,\mathrm{m} \times 1.0\,\mathrm{m}$ table in the center with a distance of $2\,\mathrm{m}$. The living room represents a typical scenario for wireless entertainment systems, as the receiver is located close to the middle of one wall—a common place for a TV screen. The transmitter takes an arbitrary location, $4\,\mathrm{m}$ away from the receiver. To characterize wall reflections, we use the permittivity of $2.26 - 2.4\mathrm{je}\text{-}4$ as experimentally

CHANNEL CHARACTERISTIC	MATHEMATICAL EXPRESSION
Channel Impulse Response (CIR):	$h(\tau) = \sum_{k=1}^{K} h_k \delta(\tau - \tau_k)$
Power Delay Profile (PDP):	$p(\tau) = \lvert h(\tau) \rvert^2$
mean delay:	$\bar{\tau} = \frac{\sum_{k=1}^{K} \lvert h_k \rvert^2 \tau_k}{\sum_{k=1}^{K} \lvert h_k \rvert^2}$
RMS delay spread:	$\tau_{\mathrm{RMS}} = \sqrt{\frac{\sum_{k=1}^{K} (\tau_k - \bar{tau})^2 \lvert h_k \rvert^2}{\sum_{k=1}^{K} \lvert h_k \rvert^2}}$
cumulative channel power:	$P = \sum_{k=1}^{K} \lvert h_k \rvert^2$

Table 3.1: Applied metrics for channel characterization.

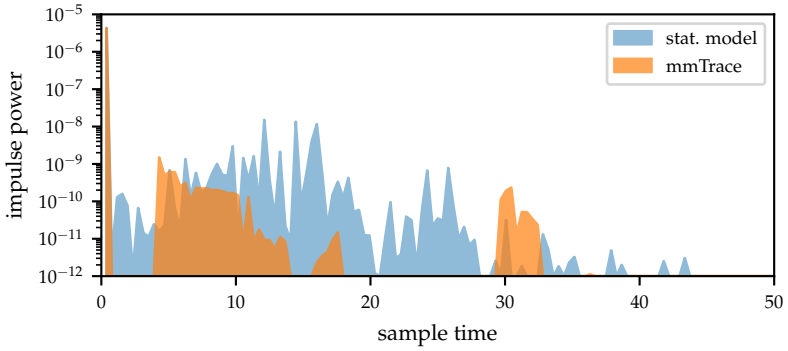

Figure 3.7: Simulated power delay profiles with mmTrace and the statistical model in the conference room scenario at 10 random positions.

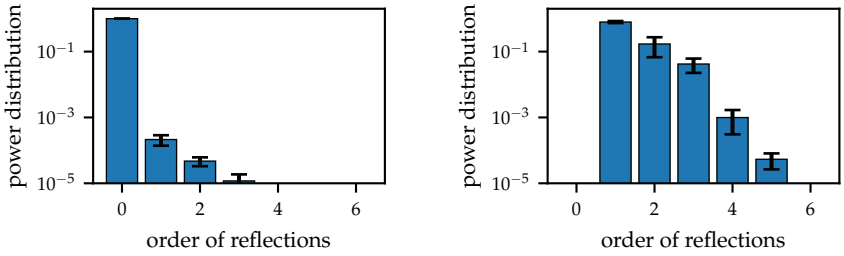

(a) line-of-sight scenario (b) non-line-of-sight scenario

Figure 3.8: Distribution of the RSS in dependency of the order of reflections.

determined by Lu et al. [Lu+14]. All antennas use the directional radiation pattern with a beamwidth of 60° or 20° and take 64 different orientations. In the conference room, we distinguish between line-of-sight and non-line-of-sight scenarios. All simulations are repeated 1000 times and provide the median and 95 % confidence intervals. By characterizing the channels obtained from both, the statistical model and mmTrace, we validate the feasibility of ray tracing for mm-wave channel modeling in the following.

The statistical model exposes smoother impulse due to scattering effects.

POWER DISTRIBUTION. First, we investigate the power distribution in the conference room scenario with antennas of both transceivers directly oriented towards each other. In this setup, we simulate the CIRs and extract the PDPs shown in Figure 3.7. At first sight, both results significantly differ. Only the high power peak without delay perfectly matches. For higher delays, the CIR of mmTrace exhibits only single impulses that correspond to specific reflection paths. The statistical model has smoother impulses which come from clustering effects and scattering that is not encountered in mmTrace. Especially the paths with 30 ns delay show that mmTrace predicts the intensity of channel paths similarly to the statistical models.

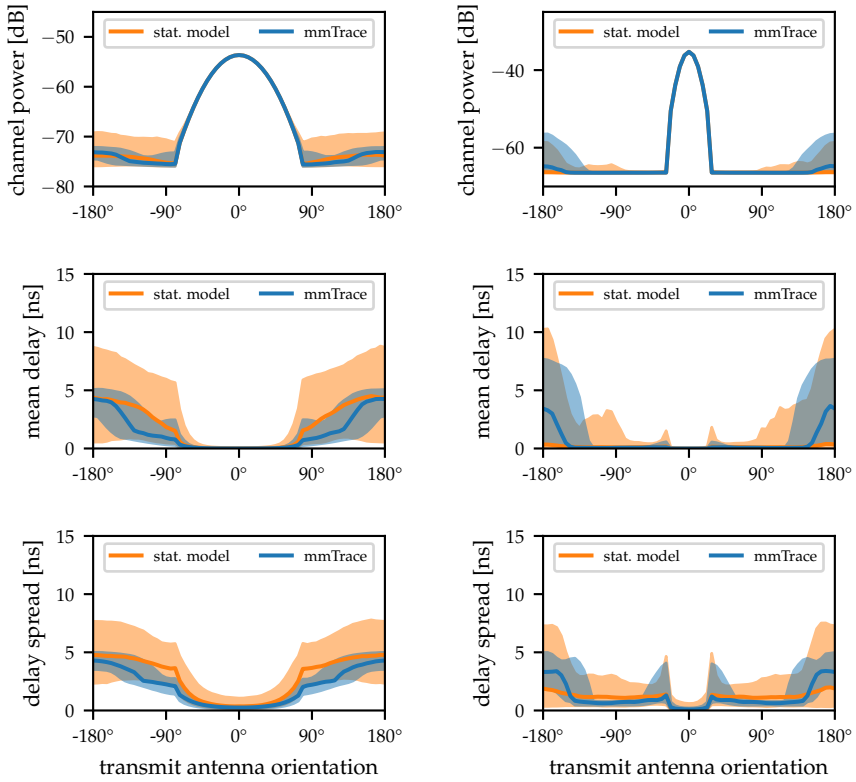

(a) 60° beamwidth (b) 20° beamwidth

Figure 3.9: Channel characteristics with rotating transmitters in the conference room scenario with line-of-sight paths and different beamwidths. The shaded areas indicate the 95 % confidence intervals as obtained in 1000 simulations.

Figure 3.8 shows the distribution of received power over the reflection order in mmTrace. In line-of-sight scenarios (Figure 3.8a), more than 99 % of the received power comes over the direct path. Reflections take a minor part in the power contribution. In non-line-of-sight scenarios (Figure 3.8b), multi-order reflections obtain increased relevance. First order reflections provide 79 % of the received power, while 17 % come over second-order reflections. Reflection paths with a higher order than four barely contribute to the received power. While the statistical model considers up to two reflections, we limit the maximum number of reflections to four in the default configuration.

mmTrace considers fourth order reflections by default.

FEASIBILITY. Analyzing the conference room scenario with line-of-sight path and a beamwidth of 60° as shown in Figure 3.9a confirms the feasibility of mmTrace. The cumulative channel power perfectly matches within the half-power beamwidth. Only for misaligned antennas, the statistical model exposes slightly higher confidence intervals. Instead of statistical probabilities, we use distinct path characteristics

The cumulative channel powers perfectly matches.

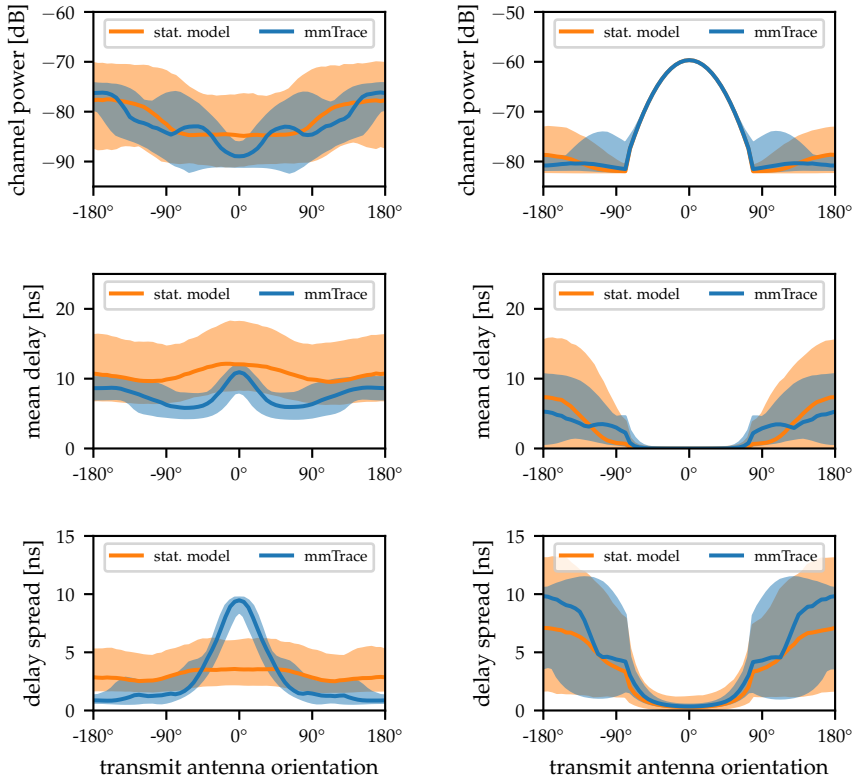

(a) conference room, non-line-of-sight

(b) living room, line-of-sight

Figure 3.10: Channel characteristics with rotating transmitters in the conference and living room scenario with line-of-sight and non-line-of-sight paths and a beamwidth of 60°. The shaded areas indicate the 95 % confidence intervals as obtained in 1000 simulations.

that lead to these observations in all experiments. Moreover, we observe a lower bound for the received signal strength that is caused by a constant side lobe level in the applied radiation pattern. The mean delay and RMS delay spread also match within the beamwidth. For small misalignments, mmTrace slightly underestimates the delay. When the transmitter orientates its antenna to the opposite direction, it approximates the statistical model again. Nevertheless, the differences are marginal.

Narrow beamwidths decrease the mean delay and delay spread.

Narrowing the beamwidth to 20°, as shown in Figure 3.9b, mainly decreases the mean delay and the delay spread in both models. Only for a strong misalignment, a high delay remains recognizable. In this situation, mmTrace slightly overestimates the parameters, but the median still falls into the statistical model's confidence intervals.

BLOCKAGE. In non-line-of-sight scenarios, where the direct path is blocked, we obtain the results shown in Figure 3.10a for a beamwidth of 60°. The channel power is significantly lower than in line-of-sight

scenarios and remains almost independent from the antenna orientation. While the statistical model provides smooth transitions of all three parameters, our results exhibit narrow power peaks that correspond to the concrete reflection paths. For antenna misalignments, mmTrace predicts a lower mean delay and delay spread. With aligned antennas, both parameters significantly grow, which is not observable for the statistical model. Only the delay spread exceeds the confidence intervals of the statistical model.

We identify power peaks for distinct reflection paths.

INCREASED DISTANCES. With increased distances, as in the living room scenario, we identify significantly higher delays than in the conference room scenario as shown in Figure 3.10b for a beamwidth of 60° and line-of-sight paths. This effect comes from the longer reflection distances. Due to the receivers position and the larger room, similarly powered reflections arrive from multiple directions that contribute to a higher delay. In summary, these validation results show that mmTrace successfully adapts to the situation: channel power, mean delay, and RMS delay spread are very close to the statistical channel model.

mmTraces achieves similar results than the statistical model.

3.3 APPLICATION SCENARIOS

In the previous section, we validate mmTrace in comparison to the statistical approach and demonstrate comparable results with only a few discrepancies. Next, we focus on the advantages of mmTrace and exemplary present specific application scenarios that are impossible to investigate with statistical channel models. After describing our simulation environment, we map the achievable signal strength at different positions in a room and consider multiple interfering transceivers. Finally, we discuss the impact of common network challenges on mm-wave communications.

We exemplary analyze specific application scenarios.

3.3.1 *Simulation Environment*

The environment of our following simulations comprises a room of size 5.0 m × 3.0 m with a blocking and reflecting wooden obstacle. We use the permittivities taken from Lu et al. [Lu+14] of $2.26 - 0.24je\text{-}3$ for concrete walls and $2.8 - 4je\text{-}2$ for wooden objects. Two transmitters are placed and oriented in parallel to cause interferences. Received signals are obtained on a grid throughout the room. Unless noted otherwise, all antennas use the directional radiation pattern with a beamwidth of 60°. In total, the simulations cover 540 locations with 64 different antenna orientations. All receivers are assumed to flexibly adjust their antenna orientation and steer it for optimal results. We analyze the signal reception as well as the interference with ideal and more realistic side lobe antenna patterns.

Simulations consider a typical indoor environment.

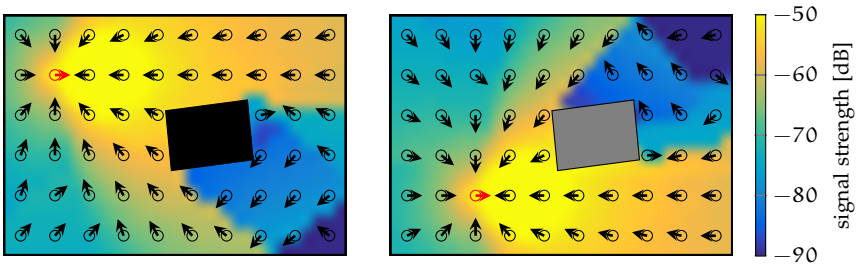

(a) first transmission (b) second transmission

Figure 3.11: Signal strength with ideal radiation patterns and a beamwidth of 60°.

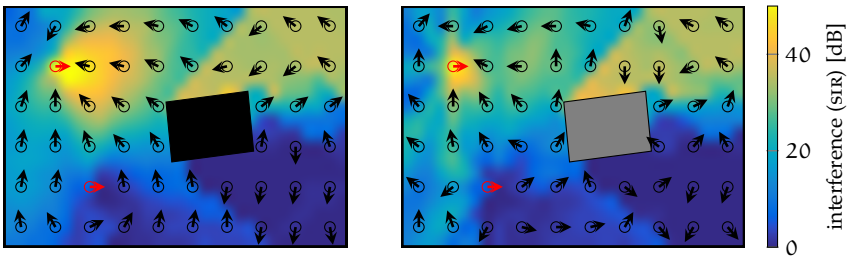

(a) ideal antenna patterns (b) patterns with custom sidelobes

Figure 3.12: Interference of two transmissions with ideal and sidelobe affected radiation patterns.

3.3.2 *Signal Strength Mapping*

mmTrace maps the received signal strength throughout the environment.

Figure 3.11a shows the map of reception capabilities with optimal antenna alignment for the first transmitter which is marked in red. The background color indicates the achievable signal strength while the arrows indicate the optimal antenna orientation. At all positions roughly inside the signal beam, receivers directly steer their antennas towards the transmitter. At locations behind the blocking objects or far out of the beam, they are oriented towards the first-order reflections. We identify very sharp transitions between certain areas, at which small distances between the receivers lead to entirely different orientations. A small location offset can have a significant impact on the channel.

3.3.3 *Interfering Transmissions*

The impacts of interference are location dependent.

For the second transmitter, we observe a similar situation shown in Figure 3.11b. Only the shadow region moves to the other side of the obstacle. The signal of this transmitter represents the interference. The Signal-to-Interference Ratio (SIR) calculates as $\text{SIR} = P_{TX1}/P_{TX2}$, where P_{TX1} and P_{TX2} are the received signal strengths from both transmitters

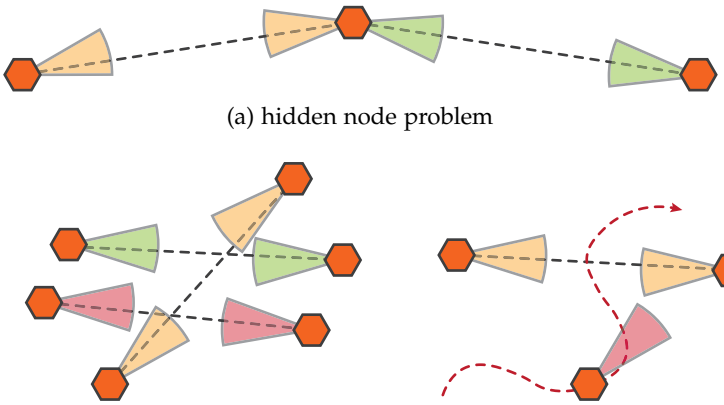

(a) hidden node problem

(b) concurrent transmissions (c) nomadic eavesdropping

Figure 3.13: Illustration of network scenarios with multiple transmitters and receivers.

respectively. The SIRs in optimal antenna orientation are shown in Figure 3.12a. Receivers aim at steering their antennas towards the first transmitter but away from the second one. Even some receivers in line-of-sight of the first transmitter misalign their antennas to suppress the interference. The optimal orientation might change from the non-interfering scenario. Receivers are affected differently: some regions still exhibit a suitable signal quality while receivers at others get strongly distorted.

Interference with antenna side lobes leads to unexpected steering directions.

So far, we only considered the ideal directional radiation pattern with a single lobe. The side lobe pattern defined in Section 3.2.2 leads to different signal strength mappings. In the interference scenario in Figure 3.12b receivers not only try to steer their main lobe towards the first transmitter; they also aim at keeping away their side lobes from the interfering signal. This can have interesting effects even on the receivers behind the object that blocks the interfering signal: receivers tend to orientate their antenna completely away from the transmitter to overcome interference in the side lobes. In the shadow region, antenna orientations appear random. This shows that in case of interference, antenna orientations for best SIR do not necessarily correspond to the line-of-sight direction even though a direct path exists. We revisit this scenario in Chapter 6 and propose an adaptive beam switching mechanism to mitigate interference in antenna side lobes.

3.3.4 *Network Scenarios*

To demonstrate the possibilities of mmTrace, we consider three well-known problems of wireless networks and analyze their impact. In particular, these are *(1)* the hidden node problem with two transmitters that cause interference at a single receiver, *(2)* the problem of concurrent transmissions, where multiple transceiver pairs distort

Finding the optimal antenna alignment can be challenging.

each other, and *(3)* the threat of nomadic eavesdroppers that aim to overhear transmissions from arbitrary positions. Figure 3.13 illustrates these scenarios. Simulations with mmTrace show that the hidden node problem becomes superficial for narrow beams and perpendicular node arrangement. Concurrent transmissions lead to various interference levels and only appear manageable with low beamwidth. The success of nomadic eavesdropping depends on the environment and antenna settings but generally becomes more challenging than in conventional wireless networks. Detailed simulation results under angular alignment dependencies in these network scenarios are provided in [SCH16a]. Our findings empathize the advantages of mm-wave communications but also outline the challenges in finding optimal antenna alignment, which is vital for reliable and efficient communication.

3.4 DISCUSSION AND SUMMARY

Channel sounding allows for analyzing reflections and blockage.

Predicting mm-wave network behavior is challenging as propagation effects are different to those of conventional wireless networks. We propose a signal propagation analysis framework that consists of a channel sounding platform and an environment simulation tool. Our channel sounding platform combines the advantages of SDRs with mm-wave transceivers. It operates in the 60 GHz band and enables to transmit arbitrary signals with modulations of up to 4-QAM. In typical indoor scenarios transmissions with low error rates are possible. This efficiently allows investigating reflections and blockage on single communication links. In larger deployments, statistical channel models are the state-of-the-art to predict propagation effect. However, they reach their limits when dealing with highly directional communication. Steerable antennas and distinct reflections on environmental objects lead to significant channel variations. Such properties that were negligible for omnidirectional communications take an essential role for mm-waves. At high frequencies, ray tracing is suitable to predict the signal propagation. We present an analysis tool that uses image-based ray tracing to simulate channel properties for specific antennas and environments with multiple transceivers. Validation results reveal only low discrepancies with existing statistical models and demonstrate that our simulation provides a valuable benefit for situations in which channel sounding is inappropriate. It allows adapting to different physical scenarios and considers multiple transmissions in parallel.

Simulations reveal environment and location specific propagation effects.

The channel sounding platform proposed in this chapter is applied to measure the signal propagation with environmental reflections in Chapter 8. Additionally, our simulations provide a foundation to understand specific propagation effects. They motivated our investi-

gation of interference in antenna side lobes in Chapter 6 and nomadic eavesdropping in Chapter 8. The evaluations in [Che+17; ZGP18; Loc+17] use our framework as well. To support the community and allow others to benefit from our work, we release the source code of our environment simulation tool on our public project page[2] (see Section A.5).

2 The source code of mmTrace, our environment simulation tools, is available at: https://github.com/seemoo-lab/mmTrace

4

TESTBED EXPERIMENTATION

Our testbed system consists of common off-the-shelf IEEE 802.11ad devices.

Our previously presented channel sounding platform and simulation environment for signal propagation analysis allow understanding how signals traverse environments and identifying the limitations of mm-wave signal propagation. This toolset is beneficial for propagation analytics but cannot evaluate the performance of mm-wave networks in practical scenarios with high-speed data transmission. Until now, no suitable mm-wave research platform exists that fully comply with the IEEE 802.11ad standard. In this section, we present our testbed experimentation framework which enables practical evaluations on off-the-shelf IEEE 802.11ad devices with high configurability. Our framework provides full access to the devices' operating system and their wireless interface drivers. By patching the firmware that is running on the IEEE 802.11ad Wi-Fi chip, we obtain control over the beam training and antenna steering properties. Without any hardware modifications, we turn cheap consumer-grade devices into alternatives to expensive mm-wave evaluation systems.

Custom software modifications provide high configurability.

In the following, we describe the platform architecture in Section 4.1. Section 4.2 details how we achieve accessibility to internal operations. Advanced antenna control is covered in Section 4.3. Section 4.4 describes the automation of large-scale experiments. Finally, we summarize the features and application of our testbed platform in Section 4.5.

4.1 PLATFORM ARCHITECTURE

The Talon AD7200 features a phased antenna array and IEEE 802.11ad compatibility.

Currently, only a few off-the-shelf devices support IEEE 802.11ad. One of these is TP-Links Talon AD7200 tri-band router that contains a Qualcomm QCA9500 60 GHz Wi-Fi module, depicted in Figure 4.1. The QCA9500 features a chip for baseband processing and is connected

Figure 4.1: Talon AD7200 tri-band router with a 32-element antenna array.

to an external phased antenna array. It achieves beam steering by individually adjusting the signal on each antenna element in the array. According to IEEE 802.11ad, the 60 GHz Wi-Fi interface supports single carrier modulations with up-to 16-QAM and achieves a theoretical data rate of 4620 Mbps. With about $250, the hardware of this platform is significantly cheaper than any other mm-wave evaluation system.

Firmware modifications provide control over the beam training and allow implementing custom beam patterns.

Unfortunately, the full control over the hardware's capabilities is encapsulated in the closed source firmware. Manufacturers typically distribute their Wi-Fi interface as black-boxes with limited access to the internal operations. We replace the router's vendor-supplied operating system with a Linux based OpenWrt variant and patch additional features directly into the binary firmware that runs on the Wi-Fi chip. Thus, we open the platform for testbed driven experiments and make the full potential of the hardware accessible to researchers. Adopting the Nexmon firmware patching framework from Schulz, Wegemer, and Hollick [SWH17] to the architecture of the QCA9500 Wi-Fi chip allows extending the beam training and adjusting internal configurations with custom patches written in C. Besides, we obtain control over the phase and gain adjustments of the antenna array which enables us to implement arbitrary beam shapes and instant switching between them during runtime. To automate research experiments and measurements in large deployments, we develop an automation system that controls devices over the network via simple Python scripts.

Our testbed experimentation framework consists of the following mutable components:

- an OpenWrt system image for the Talon AD7200,
- a modified QCA9500 firmware with beam training control,
- an extended interface driver to handle the modifications, and
- an experiment automation system to handle large deployments.

Our testbed platform consists of multiple components.

Figure 4.2 illustrates the application of these components. All components in our framework are independent of each other, can be easily ported to other architectures, and address various purposes. For example, the customized Wi-Fi firmware also runs in other devices than the Talon AD7200 that incorporate a QCA9500 chip. We provide our reconfigurable tools as open-source software to enable the integration of additional features. The build-chain for the OpenWrt system can fetch and install additional software packages. The adapted Nexmon framework for the QCA9500 architecture facilitates various customizations of the firmware that are not limited to the beam training access as done in our work. It allows other researchers to build upon our work and further improve the device's MAC and PHY mechanisms. In the following, we reveal how we achieved the open-access to the closed-source software components.

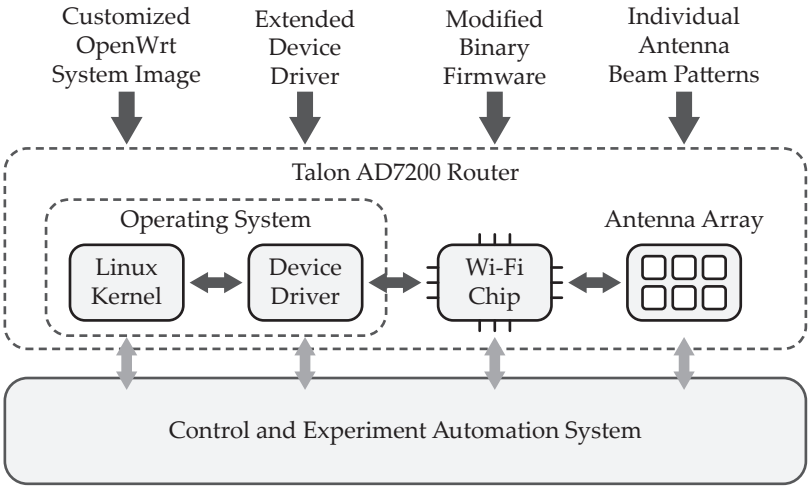

Figure 4.2: Components of our testbed experimentation platform.

4.2 SOFTWARE MODIFICATIONS

Software modifications make the full potential of the hardware accessible.

To make the full potential of the Talon AD7200 accessible, we first open up the router's operating system to control the wireless interface. To bypass the limitations of this interface in the operating system and obtain direct access to the internal operations on the Wi-Fi chip, we jailbreak the proprietary Wi-Fi firmware. Doing so, we integrate binary firmware patches that interface the beam training and extract the signal strength measurements of individual beam patterns. Next, we outline our detailed approach to achieve this accessibility.

4.2.1 *Open System Access*

Porting OpenWrt to the specific hardware architecture provides open access to the operating system.

The router's system image shipped with the Talon AD7200 allows very little access to the 60 GHz interface. Configuration is only possible via a web interface, and modifications are limited to basic parameters for the Access Point (AP) mode. An open-source operating system with full root access to the device gives us more control over the system and, especially, over the 60 GHz interface. Thus, we replace the router's vendor-supplied system image with a Linux based OpenWrt [OpenWrt] variant. Unfortunately, the hardware architecture of the Talon AD7200 is not supported in the official OpenWrt version. Therefore, we reconstruct the device tree definition, which specifies the components of the Talon hardware including the CPU, memory layout, buses, and peripherals. To gather the required information, we enable the serial console on one of our devices by adding a solder bridge between the serial pins and the connection to the main System-on-a-Chip (SoC). Reading the logged information of the bootloader on the serial output, we reconstruct the memory layout as well as the image file format. Fortunately, the Talon AD7200 uses hardware components

similar to those in an Archer C2600, which is already supported by LEDE, a fork of OpenWrt. Consequently, we use this architecture as a reference and adjust the parameters for the specific memory segmentation, I/O pin-configuration, and Ethernet interfaces. As the 60 GHz interface is not present in the Archer C2600, we additionally integrate the latest *wil6210* driver of the Linux Kernel for communicating with the QCA9500 chip. The QCA9500 is accessible on the PCI Express bus, whose controller requires additional adjustments to be handled correctly under load.

The wil6210 *driver and binary firmware files are required to enable the interface.*

Enabling the IEEE 802.11ad Wi-Fi interface requires the *wil6210* driver to load the firmware on the QCA9500 chip. This firmware must be present in the local file system and consists of two binary files, "*wil6210.fw*" and "*wil6210.brd*". The former contains the code and data partitions, whereas the latter carries device-specific configuration and calibration data. For this work, we initially used firmware files extracted from the Windows driver for Acer TravelMate Notebooks that also contain the QCA9500 chip. These firmware files are labeled with the version number "3.3.3.7759". After Qualcomm officially released the firmware binaries in the public linux-firmware repository [LFW] in May 2017, we switched to this variant with the version label "4.1.0.55" that incorporates various stability and feature improvements.

The interface supports the AP*, station, and monitor mode.*

In full control over the operating system and having access the QCA9500 chip via the *wil6210* driver, we establish IEEE 802.11ad connections as any other Wi-Fi link. The embedded Linux supports operating the interface in AP, station, and monitor mode. To achieve control of the internal IEEE 802.11ad operations requires firmware modifications which are described in the following.

4.2.2 *Firmware Extensions*

The firmware does not allow to control the beam training by default.

The Linux system uses the *wil6210* driver to control the IEEE 802.11ad Wi-Fi interface and access the QCA9500 chip. It implements a Wireless Module Interface (WMI) and sends specific WMI commands to the chip which, in turn, responds with corresponding WMI events. All commands and events carry a specific identifier. The driver uses the WMI to control the interface and request connection statistics. For example, it sends specific WMI commands to change the operation mode, set the channel frequency, and scan for APs in range. Moreover, the *wil6210* driver creates auxiliary files in the Linux debug file system that amongst others provide raw memory access and obtain detailed statistics on established connections. These debug interfaces also allow to read the selected sector but still cannot control the beam training.

We adopt the Nexmon firmware patching framework.

To control the beam training and access other internal features, we extend the operation of the Wi-Fi chip. Full control over all IEEE 802.11ad operations and the frame processing is encapsulated in the proprietary firmware running on the QCA9500 baseband chip. This

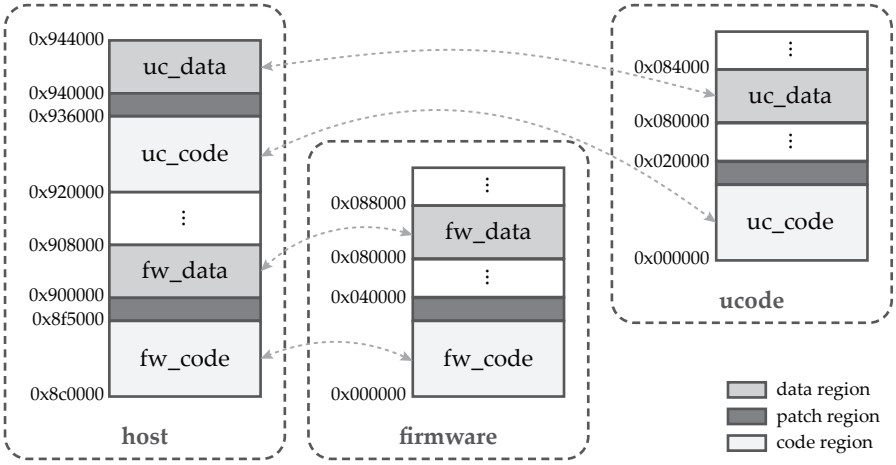

Figure 4.3: Memory layout of the QCA9500 IEEE 802.11ad Wi-Fi chip with two ARC600 processors (ucode and firmware) that have separate write-protected code and writable data memories at low addresses.

chip consists of two ARC600 processor cores for real-time (ucode processor) and other operations (firmware processor), as well as a shared memory. Using the Nexmon firmware patching framework [SWH17], we modify the firmware running on both processors and integrate custom features. The Nexmon framework enables writing binary patches in C instead of an assembly language which eases the patch development. By providing new attributes and pragmas as a GNU Compiler Collection (GCC) plugin, it defines where functions and variables should be placed in the patched firmware. We place patches in the unused memory of the firmware and redirect specific branch instructions to invoke our patches that jump back to the original code after execution. We extend the Nexmon framework for interoperability with the ARC600 instruction set and processor peculiarities. Using ARC600 processors complicates the patch creation because they contain separate read-only code and writable data regions (see Figure 4.3), while Nexmon assumes that the whole patched memory is writable. Fortunately, the code memory is also accessible at different addresses. With an additional address offset the code regions are writable and can be used to place patches.

The chip contains two processors with independent memory regions.

Before writing patches, we first have to understand how the original firmware works to identify the parts handling the information we want to extract and modify. Without any source code and function name strings in this firmware, this was a very tedious process. We started our analysis on firmware version "3.3.3.7759" and later ported our findings to the firmware versions available in the public Linux-firmware repository. After obtaining a rudimentary understanding of the firmware operations, we hooked the initialization methods of both processors and implemented a *printf* function that writes arbitrary strings and values into a ring buffer. In the debug file system of

Understanding the firmware operations is required to develop custom extensions.

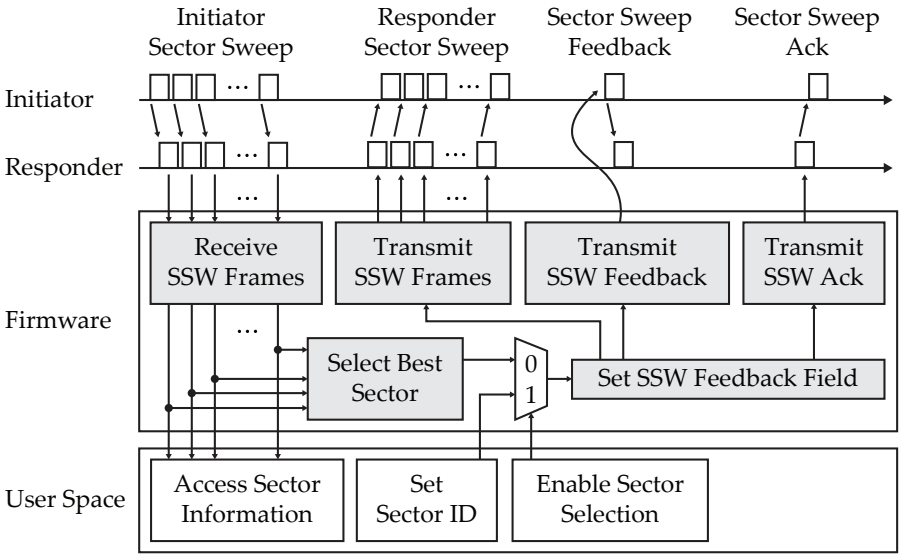

Figure 4.4: Responder sector level sweep with firmware extensions to access the received signal strength and select custom sectors from user space. White boxes indicate the extended functions, gray boxes the original behavior.

the *wil6210* driver, we create a new entity that reads the ring buffer and serves as console output for each of both processor cores. These "*uc-console*" and "*fw-console*" features print custom log information during runtime—a very simplistic yet powerful debugging feature for closed-source binary analysis. Gradually gaining knowledge on the internal operations, we finally identified the code segments that are responsible for handling the beam training. In the upcoming section, we describe the firmware patches necessary to influence the handling of sector sweep frames and extract the signal strength.

We implement a printf function for basic debugging and runtime analysis.

4.2.3 *Signal Strength Extraction*

Our patches access the sector sweep information and select custom sectors for communication.

The default firmware does not allow to control the beam training. It neither provides access to sector sweep information nor does it allow to select a particular sector. Our patches aim at enabling both features from the user space. Sector sweeps are handled in the ucode firmware which is hard to analyze due to the lack of information such as strings that indicate what functions do. To find the code segments handling the sector level sweep, we match the specific patterns of IEEE 802.11ad frames in the firmware memory. Figure 4.4 illustrates the original sector level sweep handling in gray and our extensions to it in white. The original firmware determines the Received Signal Strength Indicator (RSSI) and SNR for all received sector sweep frames and selects the sector with the highest signal strength. Our modified firmware extracts both, the RSSI and SNR, measurements for each sweep into a ring buffer that can be read from the user space using our driver

modifications. Since firmware version "4.1.0.55" the SNR is processed to select a sector. In the initial firmware version ("3.3.3.7759"), sectors cannot be selected for transmission directly. Instead, we manipulate the selected sector ID in the feedback frames that are exchanged at the end of a sweep, as illustrated in Figure 4.4. To this end, we add a switch that allows to either use the sector ID selected by the original algorithm or overwrite it with a custom value. The latter is controlled from the user space by sending a custom WMI command. In either case, the sector ID is set in the sector sweep frames. It is copied into the sector sweep feedback field of the acknowledgment frames sent by the responder and the feedback frames sent by the initiator. As this feature allows us to assess and modify the feedback in all frames, we can control the sector selection at both the initiator and responder and obtain the signal strength for all sectors individually.

Switching between the original behavior and custom sector selections is possible.

4.3 ADVANCED ANTENNA CONTROL

The QCA9500 Wi-Fi chip in the Talon AD7200 performs analog beamforming using a phased antenna array with 32 elements that are individually controllable in phase and amplitude. All these antenna elements are connected via a weighting network of amplifiers and phase shifters to a single RF-chain. This type of beamforming is widely used in off-the-shelf IEEE 802.11ad hardware since more advanced architectures such as digital or hybrid beamforming are not cost-efficient. However, the control over the antenna is fully encapsulated in the firmware of the QCA9500 chip, and only limited access is exposed to the host operating system and device driver. Using the signal strength extraction as described above and tampering with the antenna configuration, we recover the structure of the antenna weighting network. In knowledge of this structure, we adjust the steering parameters and generate arbitrary beam patterns without requiring any hardware modifications. In the following, we describe the architecture of the utilized antenna module, how we obtained control over the beam steering, and generate custom beam patterns.

The antenna allows using custom beam patterns.

4.3.1 *Phased Antenna Array Module*

To adjust the antenna steering parameters, we need to obtain a thorough understanding of the antenna and its configuration capabilities. The QCA9500 chip consists of two modules. First, a baseband IC takes care of the signal and frame processing with the two ARC600 processor cores. Second, an external antenna module with an RFIC drives the antenna elements and controls the radiation characteristics. Both modules, the baseband IC and the antenna, are connected with a coaxial cable for bi-directional transfer of modulated data, control, and clock signals, as well as the power supply for the antenna module. This

Antennas and an RFIC are mounted on an external module.

NAME	DESCRIPTION
psh_hi	**phase shift** values for antenna chains 15 to 0
psh_lo	**phase shift** values for antenna chains 31 to 16
etypeo	**edge amplifier** bit 0 for all antenna chains 31 to 0
etype1	**edge amplifier** bit 1 for all antenna chains 31 to 0
etype2	**edge amplifier** bit 2 for all antenna chains 31 to 0
dtype_swch_off	**distribution amplifier** values (3 bits each) and **X16 switch** value (8 bits)

Table 4.1: Antenna steering configuration parameters.

modular design allows to flexibly place the antenna module at proper locations inside a device chassis to minimize radiation impairments.

Individual antennas are controllable by amplifiers, phase shifters, and an antenna switch.

In transmit mode, the antenna chip mixes up the modulated data signal from IF to the desired RF channel in the 60 GHz band. On the receiver, the RF signal is mixed down to IF again. All antenna elements in the array are driven by an antenna weighting network which is adjusted by the external control signal. This antenna weighting network consists of an antenna switch, eight distribution amplifiers, 32 edge amplifiers, and 32 phase shifters. All of these are controllable from within the firmware running on the baseband chip. For antenna steering, the six 32-bit parameters *psh_hi*, *psh_lo*, *etypeo*, *etype1*, *etype2*, and *dtype_swch_off*, as listed in Table 4.1, are available. A discrete configuration of the antenna with these parameters refers to a so-called sector. The current firmware, in version "5.2.0.18", supports up to 64 different transmit sectors, out of which 35 are defined and used in beam training. Additionally, it defines a single receive sector. Definitions of these sectors are stored in codebook in the firmware.

The firmware stores a codebook of different antenna configurations.

The codebook with individual sector configurations can be either changed statically in the memory of the firmware image or dynamically adjusted during runtime. For the latter, the driver exposes specific netlink vendor commands. By changing the antenna parameters in the sector configurations, we can change the radiation patterns and gains. However, to generate arbitrary beam patterns, we need to understand the internal structure of the antenna as well as the impact of configurable parameters. In the following, we reconstruct the antenna weighting structure.

4.3.2 *Antenna Layout Reconstruction*

We experimentally reconstruct the antenna layout.

At the time of writing, no public documentation for the antenna module or the baseband chip was available. We first had to analyze the internal structure of the antenna elements experimentally to control the antenna weights properly. By iterating over all pairwise combinations of distribution and edge amplifiers and setting all other configura-

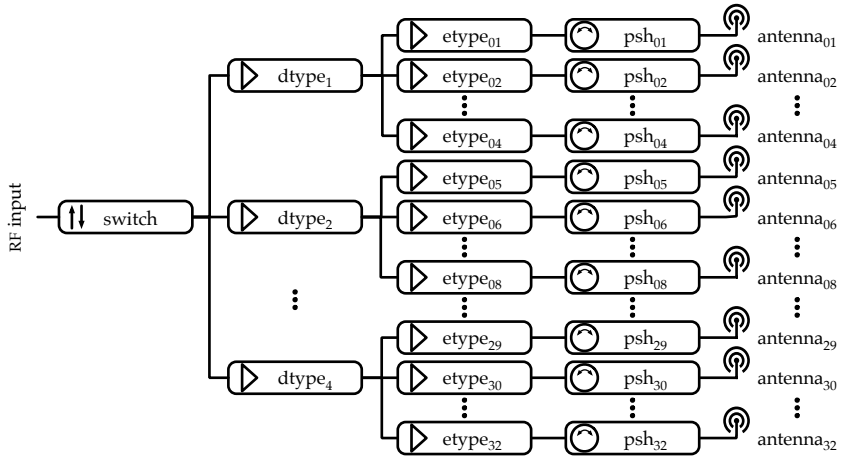

Figure 4.5: Experimentally reconstructed weighting network in the antenna module, which consists of an antenna switch, eight distribution amplifiers (dtype), 32 edge amplifiers (etype) and 32 phase shifters (psh) to drive 32 antenna elements.

tions to zero, we found that each distribution amplifier drives four edge amplifiers. An antenna is only active if both the corresponding distribution and edge amplifier are set to non-zero. Similarly, we verified which phase shifter bits belong to which antenna chain by changing single values and monitoring the received signal strength at an unmodified device. The resulting antenna weighting structure as revealed in our experiments is shown in Figure 4.5. Each phase shifter in the weighting network is driven by two consecutive bits from either *psh_hi* or *psh_lo*. The edge amplifiers use a single bit from each of *etypeo*, *etype1*, and *etype2*, while the distribution amplifiers consume three consecutive bits from *dtype_swch_off*. The most significant bits in *dtype_swch_off* are used to drive the antenna switch.

We identified the physical locations of antenna elements.

By disassembling the antenna from the device and shielding all except one element, we reconstructed the three-dimensional layout of the array as well as the physical element positions as shown in Figure 4.6. In the array 12 patch antennas are located on the front surface of the module in a 2×6 matrix shape. On the back side, there are six patch antennas and the RFIC that blocks the rest of the surface. The remaining 14 are dipole antennas and oriented towards the sides of the Printed Circuit Board (PCB). The three-dimensional layout and its asymmetric assembly, lead to irregular beam patterns, as measured in Chapter 10.

Exploiting the antenna structure enables advanced beam steering.

This know-how of the antenna layout and the capability to control the weighting network is crucial for creating custom beam patterns on the devices. It is also a valuable resource for other researchers using this platform and optimizing the beam steering accuracy. In the following, we describe the process of designing custom antenna patterns and implement them in the firmware.

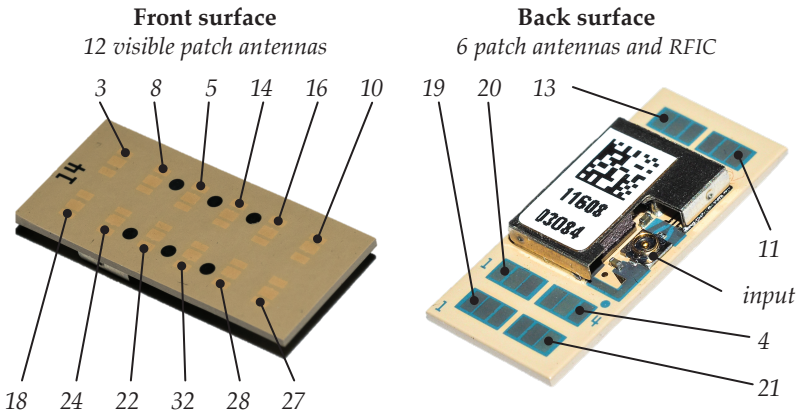

Figure 4.6: Disassembled phased antenna array of the Talon AD7200 with identified antenna elements. Fourteen additional dipole antennas are invisible from the surface.

4.3.3 *Customized Beam Patterns*

Antenna configurations can be adjusted at runtime.

The knowledge about the structure of the antenna obtained in the previous sections allows us to generate arbitrary antenna beam patterns. For each of the antenna elements, a phase shift value (2 bit) and an edge amplifier gain (3 bit) can be set. Moreover, we can configure the distribution amplifiers and the antenna switch. As the configuration for the antenna switch is typically constant, we omit it in the following description. A developed parser extracts the antenna settings and converts the gain and phase parameters for all antenna elements to the configuration registers in the codebook. We either write the codebook directly into the brd-file that hosts the default antenna configuration or update them during runtime by using the WMI commands. Take note that the latter is only available since firmware version "4.1.0.55". Version "3.3.3.7759" does not allow to change the sector definition at runtime and requires the configuration in the static brd-file.

Codebooks with up to 64 sectors are supported.

The number of selectable sectors depends on the firmware. Full support of 64 sectors is only available in firmware version "5.2.0.18". This version allows selecting specific sectors for beaconing and the sector sweep by issuing specific WMI commands. By doing so, we set the sector definitions in the codebook and enable the sectors to be used during the sector sweep dynamically at runtime. With previous firmware versions, we stick to the sectors that are enabled by default and only change their definitions. Using this approach of configuring custom sectors in the codebook, we generate optimized directional beam patterns in Chapter 7.

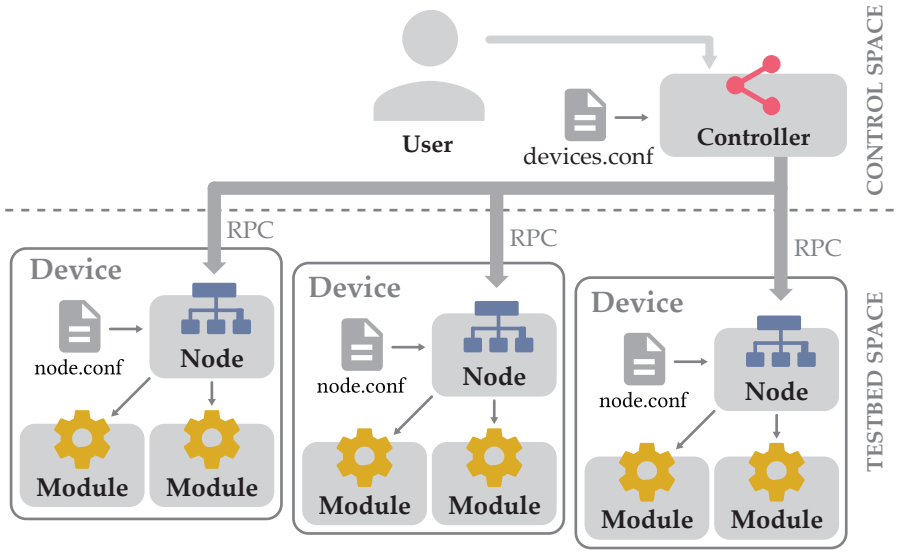

Figure 4.7: Architecture of our testbed experiment automation system.

4.4 EXPERIMENT AUTOMATION

Our testbed automation system controls multiple devices in parallel.

To perform distributed experiments with multiple devices in specific topologies, we present our lightweight framework for interaction with the Talon AD7200 and similar OpenWrt or Linux-based devices. Our framework facilitates experiment-driven research and reduces the effort that is required to set up and configure testbeds. It allows controlling a large number of devices from a single point of control with low organizational overhead. It remotely accesses the features in the OpenWrt system and our extended Wi-Fi firmware as mentioned above. In the following, we describe the architecture of our framework comprising our system model, components, as well as the user interaction.

We consider nodes that implement modules and a single controller.

SYSTEM MODEL. Our system model consists of multiple distributed testbed devices and a single controller to run and coordinate the experiment. To work with the testbed, we distribute the software components and proper configuration files to the devices. Features that are required in the experiments are encapsulated in modules at the devices and are accessible remotely. Devices run a node service to expose modules as specified in the configuration. The controller interacts with all devices and is instantiated with a testbed-specific device configuration. This configuration lists available nodes in the network and describes how they can be reached. The controller consolidates access to remote features allowing to perform various distributed experiments. Figure 4.7 illustrates this system model comprising the controller, nodes, and modules.

Our tool and scripts are implemented in Python.

The software is fully implemented in Python to supported by many operating systems and hardware platforms. Due to its flexibility and the large set of available libraries, Python is the ideal candidate to handle research experiments. To keep our system simple and easy to deploy, we connect to the devices directly without using a dedicated management server. Connections are established via simple TCP/IP sockets to remain independent of specific network topologies; the testbed can spread over multiple subnetworks. Secure Shell (SSH) access to the devices is only required for deployment. Python's interactive mode enables dynamic handling of testbed experiments with "real-time" interaction, which is particularly beneficial for live demonstrations. We encapsulate the configuration of testbed topology in order to separate it from the experiment description and simplify reproducibility.

COMPONENTS. The architecture of our framework consists of multiple components that are distributed among the testbed network. It specifies and distinguishes between *(1) modules* that expose specific features available on the devices, *(2) nodes* that represent the devices in the testbed, and *(3)* a *controller* that manages the testbed, connects to the nodes, and interfaces the modules. In the following, we describe these components in detail.

Modules implement features that are available on testbed devices.

Modules encapsulate the features that are available on the distributed devices. All modules are Python objects, can have a state, store internal properties, and can be initialized with parameters. Modules can implement arbitrary Python code, include other libraries, and invoke systems calls. We provide a module to handle the configuration of the IEEE 802.11ad interface. They control the beam training and obtain verbose debugging information. Specific modules set the IP and MAC address of Wi-Fi interfaces and read properties such as the number of received packets and error rates. These features are achieved by integrating additional Python libraries (e. g. PyRIC [PyRIC]) or invoking respective system commands. Moreover, we provide modules that wrap common system commands such as *ping* and *iperf*. Modules are instantiated and executed locally on the devices and, thus, independent from the actual testbed topology. With its modular concept, our framework supports the integration of new features easily: custom modules can be simply added to the framework.

Devices run a node service that exposes their modules.

Nodes represent devices in the testbed and handle multiple modules that are exposed to the network. A node is a service that binds to a specific TCP/IP port and exposes the available modules via Remote Procedure Calls (RPCs). RPCs invoke methods or procedures on other hosts as if they were local ones. For each remote method, the caller creates a local stub. This stub binds to the remote method, forwards all arguments to invoke the method on the host, and transfers the result back again. In particular, the stub implements the same list of

arguments, serializes them, and sends them over the network. At the remote location, the requested method is executed, and the return value serialized and transmitted back. To achieve this, we implement the Pyro library [Pyro]. As all arguments and results need to traverse the network, we keep our exposed methods as simple as possible and omit complex data structures. The node itself provides only limited features and acts as a gateway to access the modules. Available modules are configured in a *node configuration* file that specifies their properties. They are identified in the network by Unique Resource Identifiers (URIs) that consist of the host address and port as well as the module's name. The node itself exposes a list of available modules.

The controller coordinates the network experiments.

The controller orchestrates the testbed and coordinates the network experiments. It establishes the connection to all nodes and provides direct access to the modules. A user instantiates the controller in the local user space. Available nodes are listed in a *device configuration* that specifies the host address and port at which the node service can be reached. Default parameters can be provided globally to be valid for all nodes in the testbed. The controller connects to the nodes and queries all available modules. For each node, it creates a local stub that binds and forwards all method calls to the remote location such that the modules are accessible as Python objects. As a result, our framework abstracts the distribution of devices and consolidates all features locally.

Users only interact with the controller and provide a configuration.

INTERACTION. To interact with the testbed, the user installs the distribution package of our framework on the devices and instantiates the controller with a specific device configuration. The controller then queries the nodes and exposes the features in the modules to be directly accessible. This makes distributed network experimentation very convenient and straightforward as features can be accessed with only a few lines of code. It also allows to dynamically interact with network devices by using an interactive Python shell.

4.5 DISCUSSION AND SUMMARY

No other testbed with standard compliant beam training exists.

Despite first commercial consumer-grade products being available on the market, researchers struggle with inappropriate mm-wave testbed systems. Such systems only allow for signal-level measurements or use antenna setups with different behaviors than commercial off-the-shelf devices. None of them fully complies with IEEE 802.11ad to allow realistic network evaluations.

Firmware modifications customize the beam training.

In this chapter, we propose our practical testbed experimentation platform that bases on off-the-shelf devices. Our system provides full access to the Linux-based operating system and the wireless interface driver of the Talon AD7200. This device is one of the few that fully support IEEE 802.11ad and comes with an antenna array

of 32 individual elements. It establishes links in the 60 GHz band with up to 4620 Mbps. Unfortunately, the full control over the hardware's capabilities is encapsulated in a closed-source firmware and not exposed to the user. To open the platform for testbed experiments and make the full potential accessible to researchers, we replace the router's vendor-supplied firmware with a customized OpenWrt variant. Since the hardware architecture of the Talon AD7200 was not supported in OpenWrt, we had to adapt the hardware specifications and memory layout manually in order to boot properly. Moreover, we integrated the latest *wil6210* driver from the Linux Kernel to interface the IEEE 802.11ad chip. Using binary firmware patches, we modify the sector sweep algorithm and allow implementing new beam training protocols. By reconstructing the antenna weighting network, we obtained control over the phase and gain of each antenna element. Adjusting the codebook of transmit and receive sectors with these parameters enables the implementation of custom beam-shapes. Our platform supports instant switching between sectors of the codebook at runtime. To provide flexibility of deployments and simplicity of use, we propose an experimentation automation framework with a minimalistic yet powerful design written in Python.

Custom beam patterns can be implemented on the antenna array.

Research with practical mm-wave systems is in its beginnings, and we hope that our Talon AD7200 research platform encourages other researchers to explore the boundaries of off-the-shelf hardware and conduct own experiments. With only about $250 for the whole hardware, our platform provides a low-cost alternative to expensive mm-wave evaluation systems.

Our platform is a low-cost alternative to expensive evaluation systems.

The testbed experimentation platform proposed in this chapter, is applied in Chapter 5, 7, 6, 9, 10, and 11. Additionally, [Bie+18; Sah+18; Ass+18] make use of our tools as well. To support the community and allow others to benefit from our work, we release the relevant source code on our public project page[1] (see Section A.1).

1 The Talon Tools project page is available at: https://www.seemoo.de/talon-tools/

Part III

PERFORMANCE

This part of the thesis covers performance enhancements for mm-wave communication systems. Chapter 5 presents our a compressive sector selection mechanism that reduces the beam training overhead. In Chapter 6, we mitigate lateral interference in antenna side lobes by adaptively switching the receive beam. Finally, Chapter 7 describes our adaptive beam optimization scheme that increases the directionality of transmit beams based on the current channel conditions.

5 COMPRESSIVE SECTOR SELECTION

Compressive sensing exploits spatial similarities in different beam patterns.

Due to the high directionality in mm-wave communications, transceivers perform beam steering to focus their energy in the intended direction. Phased antenna arrays allow adjusting the gains and phases of single elements to change the beam pattern. However, searching all beam combinations results in a huge search space. Thus, practical IEEE 802.11ad systems only use a set of predefined beam patterns, so-called sectors. The sector level sweep then selects the optimal sector to connect to another device. As described in Section 2.2.3, it performs an extensive search, which is very time-consuming as all sectors require probing. To improve the performance, compressive sensing techniques [Ras+17; MRM16; RVM12b; RVM12a] do not evaluate all available antenna settings. Due to spatial similarities in different antenna sectors, these algorithms only probe a subset of sectors and determine the best selection by correlating the probes.

Our compressive sector selection integrates existing compressive sensing approaches into the sector level sweep.

We propose a new compressive sector selection protocol that extends existing compressive path tracking solutions. It accounts for the imperfections of low-cost hardware elements and performs a three-dimensional sector selection that is required for practical phased antenna arrays. Instead of using random beams and abstract beam patterns based on geometrical antenna layouts, we use the already well-performing beam patterns defined in the default IEEE 802.11ad chip's firmware. With our testbed experimentation framework from Chapter 4, we access the signal strength measurements of received frames and overwrite the sector selection in the feedback. Thus, we do not limit ourselves to theoretical approaches or idealistic prototyping platforms, but take into account the imperfections and non-ideal behavior of commercial off-the-shelf devices. With precise three dimensional measurements of the default beam pattern, we optimally tune the sector selection to the hardware. Integrating our compressive sector selection in the sector level sweep protocol, we directly implement it for practical evaluation. Our protocol minimizes the number of probed sectors that are required to find the best antenna steering.

In the following, we first describe the protocol design of our compressive sector selection in Section 5.1. Then, we present the results of our practical evaluation in realistic environments in Section 5.2 and, finally, discuss and summarize our findings in Section 5.3.

5.1 PROTOCOL DESIGN

In our approach, we adopt an existing compressive path tracking algorithm for optimized sector selection on commodity devices. In the following, we state existing solutions, design the compressive sector selection protocol, and describe the integration with the sector level sweep on IEEE 802.11ad devices.

5.1.1 *Existing Solutions*

The sector level sweep is inappropriate for a high number of sectors.

To determine the best sectors for the communication between two devices, the IEEE 802.11ad standard [IEE14; Nit+14] employs the sector level sweep algorithm. It works with predefined beam patterns, so-called sectors as described in Section 2.2.3. In this algorithm, two stations mutually exchange probing frames to determine the received signal strength per sector. They choose the one with maximum signal strength p_n for communication. The selected sector $\hat{n}$ out of N probed sectors is feedback to the transmitter with

$$\hat{n} = \underset{n}{\operatorname{argmax}}\ p_n. \tag{5.1}$$

As this approach scales linearly with the number of sectors in a training set, it becomes inappropriate for a high number of sectors.

Compressive path tracking reduces the search space and achieves a logarithmic complexity.

To lower the training complexity, compressive path tracking [Ras+17; MRM16; RVM12b; RVM12a] reduces the search space to a subset of sectors and thereby provides higher efficiency with lower time overhead and a search complexity of $\log N$. Instead of probing all sectors, compressive path tracking as proposed by Rasekh et al. in [Ras+17] pseudo-randomly varies phase shifts of the antenna elements to produce random beam patterns. The receiver chooses the most dominant path to the transmitter by correlating the received signal strength with the expected beam patterns. Additional phase information also enables multi-path estimation with multiple reflections [MRM16]. Such compressive algorithms achieve high efficiency and accuracy especially for antenna arrays with hundreds of elements.

Channel sensing is typically not supported by commodity devices.

Nevertheless, such solutions cannot directly be integrated into commodity devices which contain low-cost array antennas with a limited number of antenna elements. Due to the placement in a device chassis, surrounding materials may impair radiation. Manufacturers need to calibrate their antennas and already provide a preselected set of beam patterns with good radiation characteristics. Our preliminary experiments to apply random phase shifts substantially reduced the link quality between our devices under test, thus severely limiting the communication range. Given the minimum signal strength requirements that are necessary for proper signal reception, choosing random phase shifts results in less accurate measurements on our low-cost hardware. Additionally, predicting paths between devices only in a

two-dimensional environment is insufficient. The placement of communicating nodes in practical setups is seldom limited to such simple constellations. Hence, for this work, we extend existing compressive path tracking algorithms to cope with realistic three-dimensional setups.

5.1.2 *Compressive Selection*

We use non-coherent path tracking with predefined beam patterns.

To cope with the limitations of commodity IEEE 802.11ad devices, we adopt the non-coherent path-tracking approach provided in [Ras+17]. Instead of relying on random beam patterns, we use the beam patterns defined by sectors that perform well with our device's antenna. Moreover, we extend the search space to three-dimensional patterns to accommodate realistic environments. Our approach works in two steps. First, it takes a random subset of M out of N sectors supplied by the antenna to find the path a signal takes towards a receiver. Then it determines the best of the original N sectors that optimizes the signal strength in the estimated direction. Probing all M sectors results in M received signal strength values p_m with $m \in [1, M]$. Then, we find the sector with the highest expected signal strength according to expected three-dimensional beam patterns $\vec{x}_m(\phi, \theta)$, for azimuth angles ϕ and the elevation angles θ. We correlate a normalized vector of all received signal strength values $\vec{p}$ with a normalized vector of the corresponding expected beam pattern $\vec{x}(\phi, \theta)$ with

Correlations of measurements and expected beam patterns reveal the angular direction.

$$W(\phi, \theta) = \left\langle \frac{\vec{p}}{\|\vec{p}\|}, \frac{\vec{x}(\phi, \theta)}{\|\vec{x}(\phi, \theta)\|} \right\rangle^2, \tag{5.2}$$

where $\langle \vec{u}, \vec{v} \rangle$ indicates the inner product of two vectors $\vec{v}$ and $\vec{u}$ and $\|\vec{u}\|$ is the norm. To estimate the path taken by a signal to arrive at the receiver, we determine the angle of arrival by maximizing the correlation $W(\phi, \theta)$ with

$$\hat{\phi}, \hat{\theta} = \underset{\phi, \theta}{\operatorname{argmax}}\, W(\phi, \theta). \tag{5.3}$$

The sector with the highest expected gain in the measured direction is selected.

In a discrete grid of ϕ and θ, the angles $\hat{\phi}$ and $\hat{\theta}$ with maximum correlation are found numerically. Finding the best out of all N sectors, we use the estimated angles $\hat{\phi}$ and $\hat{\theta}$ to select a sector with ID $\hat{n}$ that provides the strongest gain according to the measured beam patterns $x_n(\phi, \theta)$ with $n \in [1, N]$:

$$\hat{n} = \underset{n}{\operatorname{argmax}}\, x_n(\hat{\phi}, \hat{\theta}). \tag{5.4}$$

Thus, the number of available sectors N can be significantly larger than the number of probing sectors M.

5.1.3 *Sector Level Sweep Integration*

Our firmware modifications allow integrating the compressive selection into the sector level sweep.

Integrating our compressive sector selection into the sector level sweep of commodity IEEE 802.11ad devices requires our modifications to the Wi-Fi chip's firmware as described in Section 4.2.3. First, we need to gain access to the signal strength measurements of received sector level sweep frames, determine the optimal sector, and then insert the selected sector ID in the sector level sweep feedback field. Additionally, precise knowledge of the beam patterns of all available sectors is required. Since these are device-dependent and unique for the Talon AD7200 platform, existing measurement campaigns characterizing mm-wave propagation cannot be used. Instead, we perform extensive measurements to determine the antenna gains in all three dimensions of all of the available beam patterns that are predefined in the device.

Measurements from the firmware are imprecise and exhibit several outliers.

Using our practical testbed platform, we evaluate the signal strength in every sector level sweep during the regular operation of the device. Unfortunately, our measurements occasionally tend to be imprecise, fluctuate, and exhibit several outliers. Especially when observing low gains, the signal strength deviates and sometimes the firmware does not report any measurement at all. This behavior is also observable with the unmodified firmware. It impairs the original sector level sweep algorithm and causes fluctuations in the sector selections.

Combining RSSI *and* SNR *measurements reduces deviations.*

Averaging over multiple measurements is, however, not feasible as sector selection algorithms must react quickly to changes in the environment and adjust the selected sector accordingly. To cope with the effects of measurement fluctuations in our path tracking protocol, we decided to use not only Signal-to-Noise Ratio (SNR) measurements of the sector level sweep frames but also the Received Signal Strength Indicator (RSSI). Both parameters are provided by the firmware and extractable with our modified device driver. Internally, they seem to be differently acquired as fluctuations are not observable in both values at the same time. Nevertheless, the average SNR and RSSI are correlated. Thus, we extend the calculation of the estimation of the angle of arrival presented in Equation 5.3 by considering correlations with both parameter values:

$$\hat{\phi}, \hat{\theta} = \underset{\phi,\theta}{\operatorname{argmax}}\, W_{\text{SNR}}(\phi, \theta) \cdot W_{\text{RSSI}}(\phi, \theta). \tag{5.5}$$

Compressive selection inherently compensates outliers in the measurements.

This approach effectively tolerates outliers and increases the robustness against measurement deviations in either value. By probing only a subset of antenna sectors, we have a huge advantage over the original sector level sweep algorithm: we naturally compensate missing or wrong measurements effectively. Assuming that the firmware misses reporting the signal strength of the best sectors, the sector level sweep algorithm would choose a less optimal one. With compressive path tracking, we still have good chances to find the optimal sector despite a decreased set of probing values.

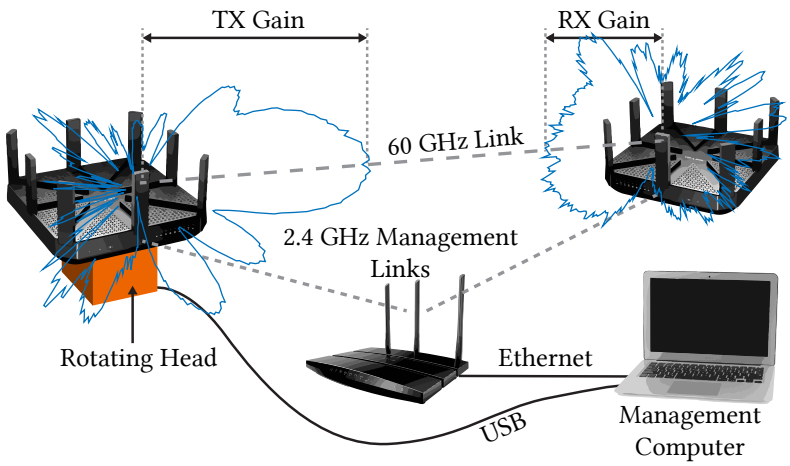

Figure 5.1: Schematic setup of our experiments with two Talon routers and a rotation head controlled by a computer over an additional management network.

5.2 PRACTICAL EVALUATION

In the following, we present our evaluation results to demonstrate that our compressive sector selection algorithm is feasible for off-the-shelf IEEE 802.11ad devices and decreases the training time in comparison to the prevalent sector level sweep. After describing our setup, we analyze the angular estimation accuracy, the impacts of sector selections and evaluate the overhead reduction and throughput.

We compare compressive sector selection with the default sector level sweep.

5.2.1 *Testbed Setup*

For evaluation purposes, we use a setup with two Talon routers and take measurements in a lab environment and a conference room. As shown in Figure 5.1, one device is mounted on a custom rotation head equipped with a step-motor with microstepping support to obtain a high rotation precision in the azimuth plane. In the lab environment, we place the second device at a distance of three meters facing the rotating one; in the conference room, both devices are six meters apart. The conference room contains a couple of potential reflectors such as white-boards so that we can expect reflections and multi-path effects. Experiments in both the lab environment and the conference room constitute typical indoor line-of-sight scenarios for IEEE 802.11ad communication. Both devices establish a connection and transmit pings for 20 s to trigger sector level sweeps and keep the connection alive. Using our firmware patches introduced in Section 4.2.3, we extract the received signal strength regarding SNR and RSSI. A 2.4 GHz network connects all devices to control the testbed remotely and to start the experiments. The step-motor in the rotation head is controlled over USB and achieves ranges of $\pm 60°$. In the lab environment, we manually

Our evaluation comprises typical indoor scenarios with one device mounted on a rotation head.

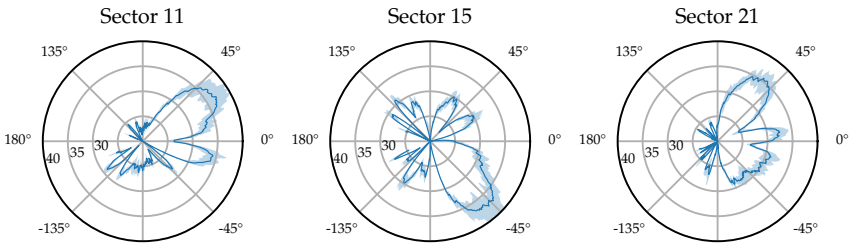

Figure 5.2: Exemplary beam patterns of three sectors predefined on the Talon AD7200 routers.

The rotation head is steered in different azimuth and elevation directions.

tilt the rotation head in steps of 2° from 0° to 30° and use an azimuth resolution of 2.25°. In the conference room, we do not change the elevation angle, but increase the resolution of azimuth angles to 1.3°. A combination of Python and MATLAB scripts automate the experiments and collect the measurement data. The two devices perform sector level sweeps frequently and record the SNR and RSSI for each sector. With off-line analyses in MATLAB, we evaluate the performance of our compressive sector selection algorithm and the original sector level sweep. To assess the desired behavior of compressive estimation with a limited number of probing sectors, we only consider a variable number of random measurements in each sweep. Based on this outcome, we set the corresponding sector in the sector level sweep feedback to force devices to use our selection for data transmission.

We additionally measure the radiation patterns in an anechoic environment.

Knowing the radiation patterns of the sectors is crucial to make an optimal selection using a path-tracking algorithm. The shapes of antenna patterns highly depend on the antenna's geometries and objects in the surrounding. Further, the packaging and placement of the antenna inside a device influence the radiation characteristics. Hence, to obtain device specific radiation patterns, we performed measurements with two Talon AD7200 device in an anechoic chamber to omit disturbing reflections and multi-path effects. Figure 5.2 exemplary shows four out of 36 different beam patterns of the predefined sectors on the Talon AD7200. Measurements for the remaining sectors are provided in Chapter 10. In the following, we describe our evaluation results regarding the performance of our compressive sector selection.

5.2.2 *Path Estimation Error*

Compressive sector selection depends on accurate path estimations.

As our compressive sector selection approach relies on the path an mm-wave signal takes to arrive at its destination, we first evaluate the accuracy of the angle-of-arrival estimation by calculating the estimation error in azimuth and elevation direction. This error is the difference between the physical orientation and the estimated one. From the setup of our experiments, we know the physical orientation of the device mounted on top of our rotation head. To determine the

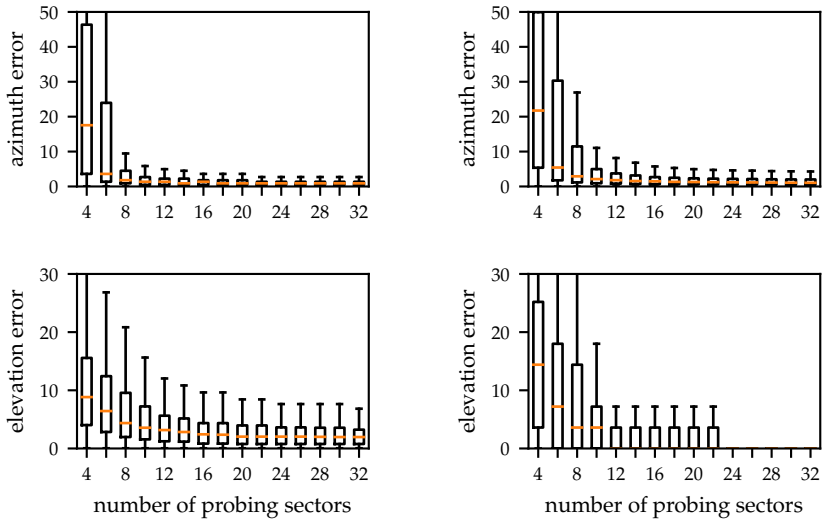

(a) Lab environment (b) Conference room

Figure 5.3: Angular estimation errors with compressive sector selection in the lab environment and conference room.

estimated path, we perform our protocol as presented in Section 5.1.2 by computing the correlation according to Equation 5.5 and estimating the direction with Equation 5.3. Both errors in azimuth and elevation direction are handled independently from each other since we measured them with different resolution and accuracy.

With only a few probing sectors, we achieve a high estimation accuracy.

Our results for both evaluation scenarios are shown in Figure 5.3, where boxes indicate the 50 %, whiskers the 99 % confidence bounds, and the dash the median of all our measurements. A suitable approximation of the azimuth direction is achieved with only a few probing sectors. For example, with ten probing sectors in the lab environment, the error stays with 99 % confidence below 5.8° and expresses a maximum median of 1.3°. With 20 probing sectors, the upper bound decreases to 3.6°, and the median becomes as precise as 0.9°. In the conference room, the errors we observe are slightly higher due to worsened environmental conditions. A higher distance between the devices and stronger multi-paths make the correlation of measurements and estimated patterns less accurate. However, with ten probing sectors, the azimuth estimation error still achieves a median of 2.1° that decreases to 1.3° by using 20 probing sectors. In the elevation axis, the errors in both environments are below 15° with ten probing sectors and below 8.4° with 20 probing sectors. Given that we obtained the beam patterns in elevation direction with half the resolution than in azimuth direction, higher errors such as those are expected. Moreover, elevation errors in the conference room are lower than in the lab environment. We expect some of the errors to be caused by manually tilting the rotation head. Despite using a digital

Manually tilting the rotation head causes inaccuracies.

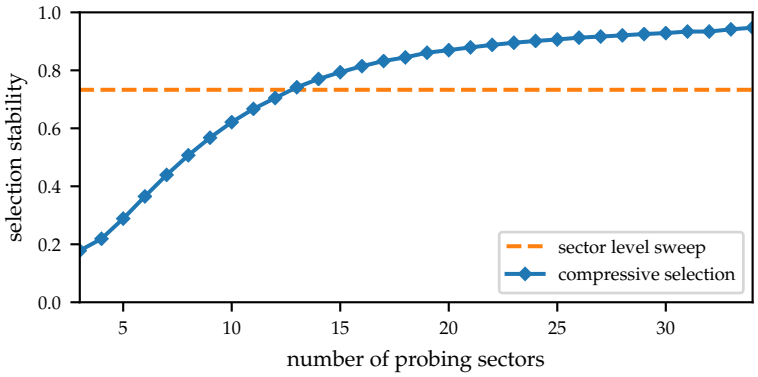

Figure 5.4: The selection stability illustrates the time spent in the most prominent sector. Plotted results are averaged over all evaluated directions in the conference room.

mechanic's level, we could not achieve a sub-degree precision. Since we did not measure different elevation direction in the conference room, the elevation error rate gets low for a sufficient number of probing sectors. In summary, a suitable approximation of the signal path becomes possible with at least 12 probing sectors.

5.2.3 *Sector Selection Accuracy*

The SNR-loss and selection stability indicate the accuracy of sector selections.

Based on the path directions that are evaluated in the previous section, our mechanism selects the sector with the highest gain in the path direction from our measurements. To investigate the accuracy of this selection, we consider the SNR-loss and the selection stability. The following states these for both algorithms our compressive sector selection and the sector level sweep.

The time spent with the most prominent sector represents the selection stability.

The selection stability represents the time an algorithm spend in one particular sector. We determine this rate by evaluating the selection of our protocol and the sector level sweep for each captured sweep in the conference room. For each physical path direction, we identify the sector that is selected most and count the occurrences. This number divided by the total number of evaluated sweeps provides the selection stability. In the assumption that each sweep has the same duration, the selection stability directly corresponds to the time spent in one particular sector. Spending more time in one sector indicates higher stability in the static environment.

With at least 13 probing sectors, we achieve a higher stability than the sector sweep.

Figure 5.4 shows the measured stabilities averaged over all evaluated path directions. For the sector level sweep, the results are constant as this algorithm always use the maximum number of probes. It achieves a stability of 73.9 % meaning that its outcome is mostly one constant sector. In 26.1 % of the occasions, the sector level sweep results in inconsistent selections. The sector level sweep tends to switch between sectors due to measurement deviations. Multiple sectors with similar

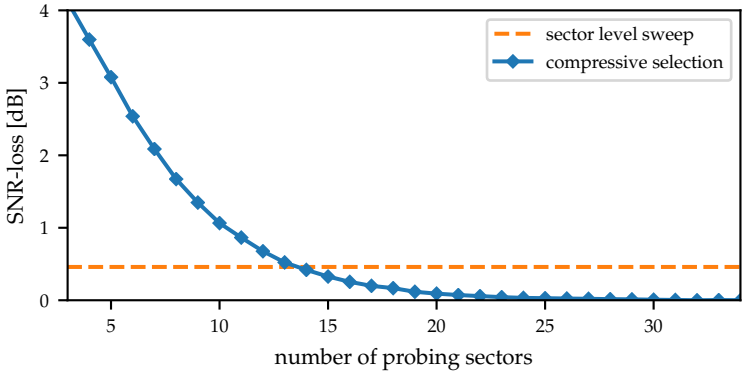

Figure 5.5: Average SNR-loss in compressive sector selection and the sector level sweep in dependency of the number of probing sectors.

gains might be alternating for the highest measurement. With at least 13 probing sectors, compressive sector selection achieves a higher selection stability than the sector level sweep. Applying all probing sectors that are available, we stay on average in 94.7 % of the time in one particular sector. This result demonstrates that compressive sector selection not only achieves the stability of the sector level sweep with fewer probing sectors but also reduces the number of alternating sector selections.

The SNR-loss represents the relative quality of selected sectors.

However, the selection stability does not represent the quality of the selected sectors. To catch up on this, we additionally investigate the loss in SNR achieved by compressive sector selection and the sector level sweep in comparison to the optimal achievable SNR. Doing so, we determine not only the sector selection but also the SNR value in each sweep. The optimal SNR is taken from the sector with the highest measurement as reported in the current and previous measurements. The SNR-loss is determined as difference towards this value. We average the SNR-loss of all evaluated directions and plot our results in dependency of the number of probing sectors in Figure 5.5. In this comparison, the sector level sweep achieves a reasonable SNR on average just 0.5 dB below the optimum. Our protocol already achieves considerable results with only a few probing sectors: six probes are sufficient to find a sector with an average of 2.5 dB below optimum. This gain might be insufficient to achieve high data-rates but already allows for simple data exchange. Greater values than those of the sector level sweep are achieved with at least 14 probing sectors. With about 20 probing sectors the achieved SNR approaches the optimal one quite well. These values exhibit high variations over different directions that are not encountered in the averaged results. Nevertheless, 14 probing sectors are sufficient to make a sector selection that provides the same SNR as the one obtained by the sector level sweep.

14 probing sectors are sufficient to achieve an SNR as high as the default operation.

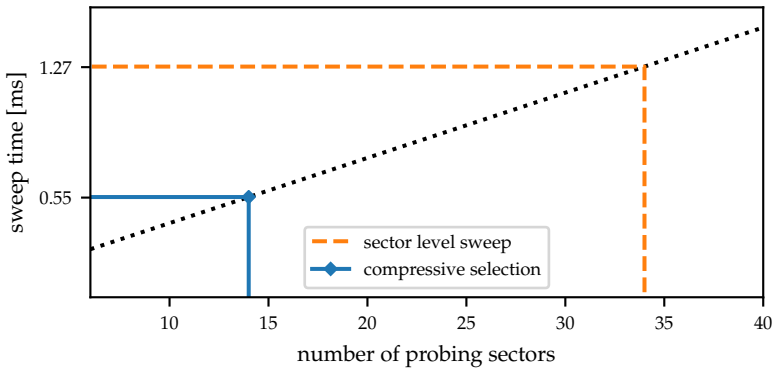

Figure 5.6: Required time to perform a mutual beam training in dependency of the number of probing sectors.

5.2.4 *Overhead Reduction*

Compressive sector selection reduces the training overhead by a factor of 2.3

With at least 14 probing sectors, compressive sector selection achieves a low estimation error as well as an SNR and selection stability in the order of that of the sector level sweep. Based on these results, we evaluate the overhead reduction. The training time directly depends on the number of probing sectors. As experimentally obtained, training transmit sectors in both directions using the sector level sweep takes on average 1.27 ms. From this time 18.0 µs are required for each sweep frame and additional 49.1 µs for initialization and feedback frames. Figure 5.6 illustrates the resulting training time for our compressive selection in dependency of the number of probing sectors. Specifically, with only 14 probing sectors, the time to complete a mutual training of transmit sectors could be performed in 0.55 ms. In comparison to the original sector level sweep that takes 1.27 ms, this effectively speeds up the beam-training by a factor of 2.3.

5.2.5 *Throughput Improvement*

Higher stability slightly increases the TCP *performance.*

To assess the throughput that is achievable with compressive sector selection, we determine the selected sectors for each sweep as done before. With this sector selected in the feedback field, we generate random TCP payload and measure the application-layer data-rate by running `iperf3` [iPerf]. Even in static scenarios, the devices trigger the beam training approximately once every second. Our measured throughput values are taken within a 10 s interval and averaged over all selected sectors to take into account the impacts of suboptimal selections. In the conference room measurements, the rotation head steers to −45°, 0° and 45°. As shown in Figure 5.7, the expected throughput of compressive sector selection is with 1.48 Gbps, 1.51 Gbps, and 1.50 Gbps only slightly better than that of the original sweep. In prac-

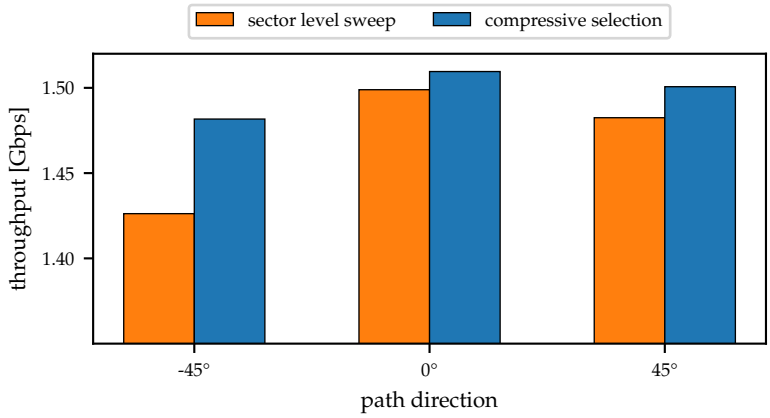

Figure 5.7: Throughput gain from increased stability in compressive sector selection with 14 probing sectors in comparison with the sector level sweep.

tice, these differences might barely be recognizable but yet show the additional performance gain we achieve from higher stability.

Theoretically, a decreased training time leaves more time for data transmission and thereby increases the throughput. This aspect is not covered in our evaluation. The provided throughput comparisons are performed with equal sweep duration. We leave adjusting the duration of the sector level sweep for future work.

With compressive sector selection probing 14 out of 34 beams is sufficient to find the optimal sector.

In this evaluation, we demonstrated the feasibility of compressive sector selection on IEEE 802.11ad off-the-shelf devices. In our protocol, 14 out of 34 probing sectors are sufficient to outperform the sector level sweep. Compressive sector selection estimates the path direction with high accuracy and an error of only a few degrees. The achieved signal strength and throughput are in the same order as those of the sector level sweep. With our approach, we increase the stability of sector selections, thus spend more time transmitting in the optimal sector. By using only 14 probing sectors, compressive sector selection reduces the mutual training time from 1.27 ms to 0.55 ms. However, we support a variable number of sectors, so that our system adjusts to various requirements.

5.3 DISCUSSION AND SUMMARY

Efficient beam training is mandatory to handle mobility.

Our compressive sector selection algorithm is the first working implementation of compressive path tracking that runs on off-the-shelf mm-wave hardware. With only two devices, the possible gain in throughput is low. However, in dense mm-wave deployments, we need to keep in mind that each sector level sweep performed by a pair of devices pollutes the whole mm-wave channel in all directions. This reduces the benefit of using mm-wave hardware to communicate with

many stations in parallel over directional links. In mobile setups, users need to be tracked over time. The shorter the sweeping time, the more often a sweep can be performed without degrading the throughput. Hence, our approach is best suited to increase the performance and frequency of sweeping in emerging mm-wave applications.

Our approach scales well with a high number of supported sectors.

While current mm-wave devices use only a few different sectors, future generations are likely to demand a higher directivity and more fine-grained beam control. Such requirements could be addressed by increasing the number of implemented and predefined sectors. The phased antenna array in the Talon AD7200 with 32 individual antenna elements already suffices to refine the beams. However, increasing the number of sectors adds additional overhead to the training process. With few sectors, the sweep completes in a suitable time, but with a large set of probing sectors, it becomes inefficient. Solutions such as our compressive sector selection overcome this limiting factor and scale well with a high number of sectors. For example, with our approach, we could significantly increase the number of available sectors while keeping the number of probes as low as in the current sweep. As a result, more precise beam patterns could be efficiently selected without adding additional training time overhead.

Adaptively controlling the number of probing sectors could become beneficial for varying scenarios.

Moreover, our approach allows for adaptively controlling the number of sectors that are probed in the sweep. In static scenarios, few probes are sufficient to validate the current antenna settings. Whenever a node starts moving, the number of probes may increase to keep track of the movement. Instead of applying a random selection, predefined probing sectors might provide further benefits. Given an approximate signal direction, sectors with low gain or similar gains in that direction are unlikely to perform efficiently. Choosing the most-efficient probing sectors turned out to be highly context-specific. Still, compressive sector selection is suited for such scenarios by keeping the number of probes as well as the selection of sectors variable.

We increase the efficiency of beam training on practical devices.

We implemented our compressive sector selection protocol that integrates with the sector level sweep on off-the-shelf IEEE 802.11ad devices. Using our testbed experimentation platform (Chapter 4), we integrate precisely measured radiation patterns (Chapter 10). Our implementation estimates the signal direction in practical environments with low error. Thereby, we significantly decrease the number of sectors that need to be probed in order to perform an optimal sector selection. Experimental results show that probing only 14 out of 34 possible sectors is sufficient. The time to search for the optimal sector in a mutual beam training scenario can be reduced from 1.27 ms needed for a complete sector level sweep to only 0.55 ms. Hence, we speed up the beam training on off-the-shelf devices by a factor of 2.3. Since our system design directly integrates with the sector level sweep, it operates orthogonally to other MAC and network layer optimizations.

6 MITIGATING LATERAL INTERFERENCE

Using pencil-shaped beams in mm-wave networks to achieve "pseudo-wire" behavior is a myth. The reason is that real-world antenna arrays suffer from significant side lobes [Nit+15a]. This is particularly notable in the case of consumer-grade hardware such as the Talon AD7200 due to its cost-efficient design. As a result, interference becomes a problem despite the use of directional communication. Faulty nodes or parallel operation of incompatible standards which distort the channel may cause significant interference not just via main lobes but also via side lobes. As already revealed in Section 3.3, the best beam pattern for a receiver may not be the one whose main lobe most accurately points towards the transmitter but the one whose side lobes do not capture any interference. Due to the sparse multi-path environment, interference is one of the main reasons for link instability in quasi-static scenarios. As such impairments are highly direction depended, mitigating interference in mm-wave networks pose a challenging task.

Consumer-grade phased antenna arrays expose significant side lobes.

The straightforward approach to identify the most suitable beam pattern in order to address interference is performing a receive beam sweep (see Section 2.2.3). However, the duration of the sector level sweep in the IEEE 802.11ad standard is in the order of a few milliseconds [EC16; Nit+14]. Since interference may be intermittent and change very often, triggering a sweep each time such a change occurs would strongly impact the performance—nodes would spend a long time sweeping. The key challenge when it comes to avoiding interference is the small timescale at which interference events take place. In contrast, other events which require beam pattern adaptation, such as blockage and device movement, take place at the timescale of seconds. To date, mm-wave networks lack efficient mechanisms to address this order-of-magnitude difference regarding the timescale of interference.

Interference may change frequently.

We present an adaptive beam switching mechanism that adapts receive beam patterns timely and efficiently to changing interference. Our mechanism does not require any location information and only needs limited information regarding the beam pattern shape. Specifically, the intuition behind our mechanism is as follows. Instead of performing a full sector level sweep whenever interference changes, we only probe the beam patterns which are most likely to result in good performance. To this end, we keep track of the probability of interference-free transmission for each beam pattern. Nodes initialize those probabilities based on full sweeps at comparatively large, fixed intervals. Interference changes trigger the aforementioned individual beam pattern probes. The key feature of our mechanism is that, when-

Our beam switching adapts to varying interference in antenna side lobes.

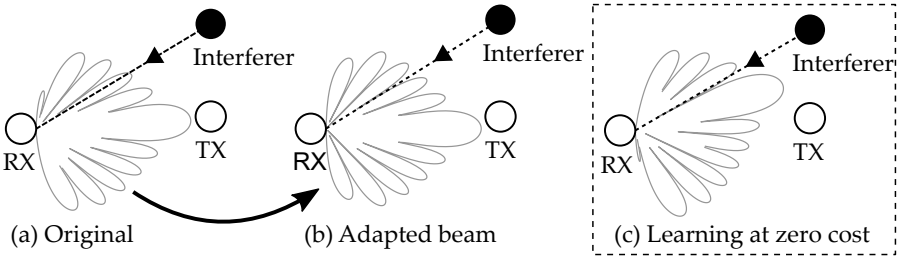

Figure 6.1: Adaptive beam pattern switch example. The beam pattern in case (b) mitigates interference because the interference signal falls into a beam pattern minimum.

ever a node sends such an individual probe, it does not only update the probability for that beam pattern. It also updates the probability of similar beam patterns at zero cost. Basically, we adjust the probabilities of other beam patterns based on the correlation of their lobes with the lobes of the probed beam pattern.

We exploit similarities in beam pattern shapes.

Our algorithm exploits similarities of antenna side lobes based on preliminary simulations with our channel models from Chapter 3. Using our testbed experimentation platform from Chapter 4, we evaluate the practical performance of adaptive beam switching with real-world measurements. Our results show that side lobes play a critical role regarding interference on commercial devices. In the following, we describe the protocol design in Section 6.1. Section 6.2 states the results of our practical evaluation, and, finally, we summarize our findings in Section 6.3.

6.1 PROTOCOL DESIGN

Our mechanism monitors the stability and finds alternative beams.

Our adaptive beam switching mechanism detects and avoids interference by adjusting side lobes in the receive beam patterns. With neither location nor detailed beam pattern information, it monitors the stability of the link and determines alternative beam patterns with minimal search overhead. In the following, we provide a protocol overview and describe the algorithm behind this approach in detail.

6.1.1 *Adaptive Beam Switching*

The general idea of our protocol is to continuously monitor a stability metric that represents changes in the received signal quality. Under low stability, it probes the channel conditions more often. It probabilistically selects beam patterns for probing that are similar to the currently selected one but exhibit different gaps in the beam pattern shape. This probing approach allows finding beam patterns that mitigate interference and still provide a suitable signal gain at each transceiver independently.

Figure 6.1 shows an example of our mechanism. In Figure 6.1(a), the receiver RX is receiving data from transmitter TX. To this end, it uses a receive beam pattern whose main lobe points towards TX. During the communication, a nearby node starts transmitting. Although the node is not aligned with the main lobe, RX still receives interference via one of its side lobes. Our mechanism at RX detects these distortions and switches to a similar beam pattern as shown in Figure 6.1(b). While the main lobe still points roughly towards TX, the interfering signal falls into a minimum of the beam pattern, thus mitigating its impact. Moreover, RX also increases the probability of using the beam pattern in Figure 6.1(c) since its side lobes are similar to the one in Figure 6.1(b). All in all, the design of our mechanism has some significant advantages compared to traditional beam sweeping. With each probe, our mechanism learns about many beam patterns. This approach reduces the overhead dramatically. Our mechanism does not require full beam pattern information but only a beam pattern correlation matrix and operates in a fully distributed manner. No coordination or knowledge from other nodes in the network is needed.

Steering the beam patterns' minima towards the interference direction mitigates the distortions.

6.1.2 *Protocol Specification*

In general, our protocol consists of multiple steps. First, it performs an initialization. Then it monitors the system stability and derives probing probabilities. Based on this, it either probes the channel or continues with data transmission. The flow chart in Figure 6.2 depicts this operation. Details of each processing step are as follows.

INITIALIZATION Each transceiver has a set of N predefined receive beam patterns it can choose from to maximize the signal quality for a respective communication partner. All these beam patterns with $n = 1 \dots N$ exhibit main lobes and side lobes in different directions. In our adaptive beam switching, we keep track of the beam $\hat{n}$ that provides the highest Signal-to-Interference-plus-Noise Ratio (SINR) over time. At time $t = t_0$, the selected beam $\hat{n}$ is initialized by performing a complete sweep:

A complete sweep is performed once for initialization.

$$\hat{n} = \arg\max_{n=1\dots N} \gamma_n. \tag{6.1}$$

The corresponding SINR is stored in $\hat{\gamma}$ and initialized by $\hat{\gamma} = \gamma_{\hat{n}}$ at time t_0. We set up the stability s with the maximum value of $s = 1$. Note that a complete beam sweep is required only during the initialization phase (i.e., at $t = t_0$).

SYSTEM STABILITY. We assume a locally time slotted system, similar to IEEE 802.11ad. At each time slot t, a transceiver either (i) probes the channel or (ii) transmits data with the best-known beam pattern from the previous slot $t - 1$. To make this decision, we use a stability

Transceivers probe alternative beams.

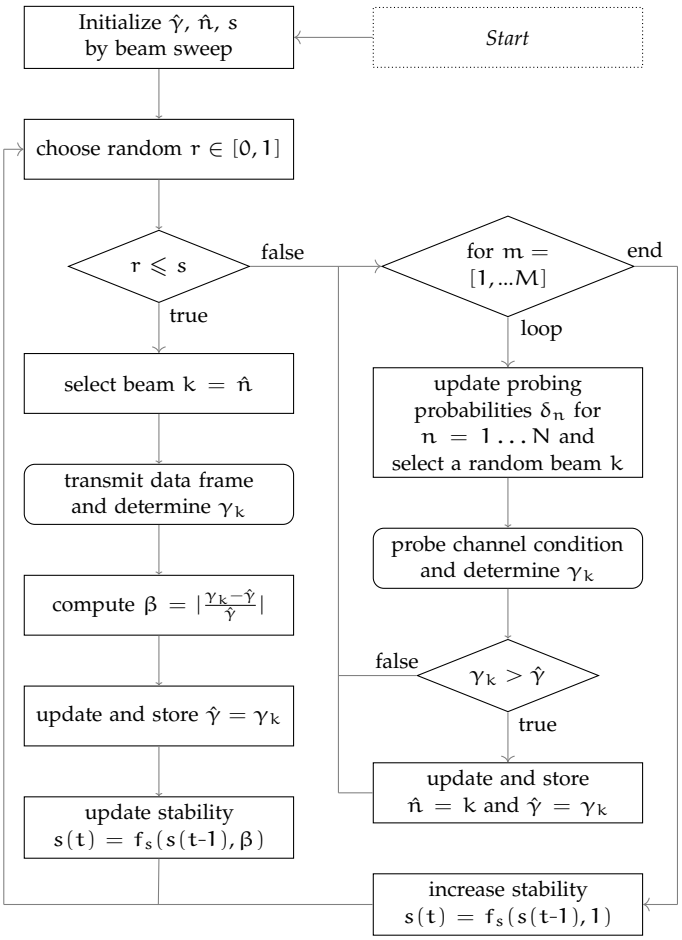

Figure 6.2: Flow chart of our adaptive beam switching algorithm for interference mitigation in antenna side lobes.

parameter $s \in [0,1]$. Given a random value r from a uniform random distribution with $r \in [0,1]$, transceivers continue transmitting data if $r \leqslant s$ or initiates probing otherwise. As channel conditions change over time, the stability s is continuously updated.

Each beam pattern has a probing probability.

PROBING PROBABILITIES. Each beam pattern n is assigned a probing probability δ_n according to the current beam pattern selection $\hat{n}$ and their similarities in beam pattern shape. Let u and v be a pair of beam patterns with $u, v = 1, \ldots, N$ but $u \neq v$. Every beam pattern n features a unique antenna radiation pattern $W_n(\Theta)$ for $\Theta \in [-\pi, \pi]$. We determine the correlation between two beams by their cross-correlation coefficient with zero lag with

$$c_{u,v} = W_u \star W_v[0] = \int_{-\pi}^{\pi} W_u^*(\Theta) \cdot W_v(\Theta) d\Theta, \tag{6.2}$$

where $W_u^*(\Theta)$ is the complex conjugate of $W_u(\Theta)$. Next, the algorithm also determines the correlation of the minima in beam pattern shapes to find beam patterns with similar gaps as

$$\bar{c}_{u,v} = \frac{1}{1+W_u} \star \frac{1}{1+W_v}[0]. \tag{6.3}$$

At runtime, our protocol does not need detailed information on the beam pattern shapes. Only the correlations $c_{u,v}$ and $\bar{c}_{u,v}$ of beam pairs are required. These values can be determined in device calibration once and stored for later use. With this approach, we also encounter for varying hardware inaccuracies in radio circuits among different devices.

Alternative beam patterns should have a high correlation with the current one.

To overcome interference or signal jamming, we aim to search for an alternative beam pattern that (i) maximizes the antenna gain towards the intended direction and (ii) minimizes the impact of interference. Specifically, an alternative beam should have a high correlation with the currently used beam pattern but different zeros to steer the antenna away from the interference direction. Therefore, the probability of probing for beam pattern n given the current beam pattern $\hat{n}$ is

$$\delta_n = c_{\hat{n},n} \cdot (1 - \bar{c}_{\hat{n},n}). \tag{6.4}$$

A subset of available beams is used for channel probing.

CHANNEL PROBING. Channel probing determines if, under the current channel condition, beam switching is beneficial. It is performed in the form of a burst of M control frames. Each frame probes an individual beam pattern. For each $m = 1, \ldots, M$ our algorithm selects a random beam pattern k based on the corresponding δ_n for $n = 1, \ldots, N$. With the chosen beam pattern we measure the SINR γ_k and update the internal state. The measured value γ_k is compared to that of the previously used beam pattern $\gamma_{\hat{n}}$. If $\gamma_k > \gamma_{\hat{n}}$, the algorithm updates the current beam pattern and its corresponding SINR to $\hat{n} = k$ and $\gamma_{\hat{n}} = \gamma_k$, respectively. Doing so, the probed beam pattern gets selected when it provides better quality than the current one.

Probing increases the stability.

Lastly, after probing all M selected beam patterns, the stability value increases. A function $s = f_s(s, \beta)$ updates the current stability based on the previous one and an update parameter β. We implement this update function as moving average with:

$$f_s(s, \beta) = \alpha \cdot \beta + (1 - \alpha) \cdot s \mid \alpha = 0.1. \tag{6.5}$$

This function and especially the adjustment parameter α is replaceable for general applicability of our protocol. In preliminary experiments, we revealed that our implementation with $\alpha = 0.1$ provides good results. More sophisticated update strategies may take into account packet delivery and error rate to adaptively control the stability. However, such strategies are future work. To increase the stability after probing, we set $\beta = 1$ and $s = f_s(s, 1)$.

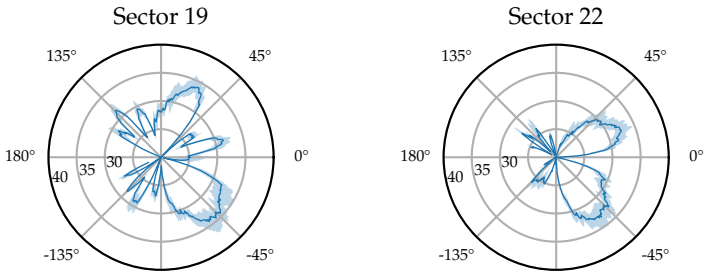

Figure 6.3: Exemplary beam patterns of two sectors with different side lobes. Both exhibit similar gains at −50°, but strongly differ in other directions.

Data is transmitted under high stability.

DATA TRANSMISSION PROCEDURE. If the stability is high so that $r \leqslant s$, the algorithm continues transmitting data with the same beam pattern used before. It selects the beam pattern as $k = \hat{n}$ and transmits a data frame. Still, it updates the stability as the signal quality might change. The relative change in the channel conditions affects the stability. Specifically, the stability update is determined by $\beta = |\frac{\gamma_k - \hat{\gamma}}{\hat{\gamma}}|$, and thus updates to $s = f_s(s, \beta)$. We use the same update mechanism as for probing in the previous paragraph. The instantaneous SINR is stored as $\gamma_{\hat{n}} = \gamma_k$.

Whenever the SINR in the current data transmission drops, the decreasing stability increases the probability of channel probing in the upcoming time slots. Still, the moving average of stability monitoring ensures that our algorithm does not overreact on rare outliers and temporary outage.

6.2 PRACTICAL EVALUATION

In this section, we evaluate the performance of our algorithm in a practical environment. To this end, we utilize our testbed experimentation framework from Chapter 4. The Talon AD7200 routers allow controlling the selected beam pattern and switch between receive beams. The following first describes our testbed setup and then provides the performance results of our interference mitigation approach.

6.2.1 *Testbed Setup*

Parallel operation of incompatible standards causes interference.

Interference may be caused due to parallel operation of incompatible standards. For this reason, we utilize two Talon AD7200 routers that communicate with IEEE 802.11ad and a WiHD transmitter-receiver pair that is based on the older WiGig standard. The latter does not perform clear channel assessment, and thus causes interference on the former. The routers use the predefined beam patterns which are configured in the interface firmware and select them for reception. Figure 6.3

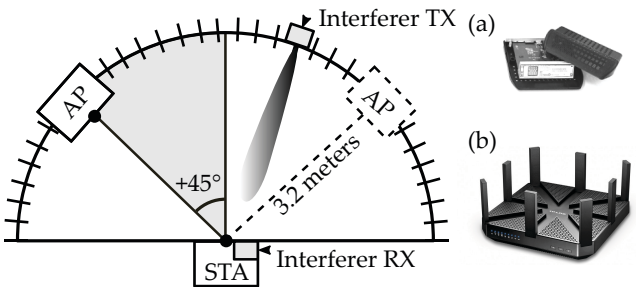

Figure 6.4: Practical experiment setup with an interfering WiHD transceiver (a) and a Talon router (b). The solid line shows the location of the router at +45° and the dashed line at −45°.

exemplary illustrates the shape of two of these beam patterns that allow for interference mitigation. While both patterns exhibit similar gains in the direction of −50°, their lobes completely differ at 60°. All the predefined beam patterns complement each other to provide strong gains in all directions. Still, they feature different beam-widths and main-lobe directions. Switching between those patterns might mitigate the interference.

All devices are deployed on a semi-circle.

Figure 6.4 shows our practical setup. One of the Talon routers is placed in the center of a semicircle and configured in IEEE 802.11ad station mode. The second router is located on the semicircle at 45° or −45° and configured as an Access Point (AP). Additionally, we place the WiHD transmitter on 18 evenly distributed locations on the semicircle, and the WiHD receiver in the center of the semicircle close to the station. As a result, the station receives both the intended transmission of the AP and the interference of the WiHD transmitter from different angles. For our experiment, we set the beam pattern of the AP to point to the center of the semicircle. This ensures that signal quality fluctuations are not due to switching beam patterns. At the receiver, we measure throughput and SNR for each possible receive beam pattern and each location of the interfering WiHD transmitter. All of our experiments are performed in a large and empty sports hall. For each measurement, we generate TCP traffic from the AP to the station using `iperf3` [iPerf].

6.2.2 *Interference Mitigation*

We study the feasibility of interference mitigation.

Our first experiment analyzes whether the intuition sketched in Figure 6.1 holds. That is, we study if switching to a receive beam pattern that has lower gain towards a transmitters helps mitigating lateral interference. Our results are depicted in the heat map in Figure 6.5. It shows the throughput that is achieved for each possible receive beam pattern and each location of the interfering WiHD node. The row marked with an arrow indicates the best pattern when the interfering node is off; the dashed line indicates the best pattern when the inter-

ference is switched on. While the former remains constant, the latter changes for each location of the interfering node. This validates that (a) deviating from the best pattern helps in case of interference, and (b) choosing an alternative pattern is not straightforward. Since practical beam patterns are highly irregular [Nit+15a], the best alternative may be any of the available beam patterns and not just the neighboring sectors. Thus, an efficient probing scheme such as ours is needed.

Alternative beam selections increase the throughput in case of interference.

Figure 6.6 shows the throughput gain that we achieve when selecting the sector with less interference. Our results are averaged over all possible locations of the interfering node. The gain is computed relative to the best sector highlighted in Figure 6.5 and to the default behavior of the router. The latter refers to the case when the router is not forced to use a specific receive sector. Compared to the best sector, our selection achieves an average throughput gain of about 60 %. For individual locations of the interfering node the gains range up to 2- and 8-fold. Since the beam patterns are asymmetric, the results at 45° and −45° differ. At 45° the gain is lower than at −45°. At such alignments, avoiding the WiHD interference becomes more challenging than for others.

6.2.3 *Probing Time*

Our approach required fewer probes than a regular sweep.

Next, we investigate the impact of the number of probes M. If M is small, our protocol only probes a few sectors each time the stability degrades. As a result, it requires more time to select a stable pattern. Beyond a certain threshold of M, the probability of probing a suitable sector increases significantly, and our beam switching quickly converges. Figure 6.7 depicts this effect for a challenging scenario, where the interfering node is located close to the AP. In this case, the station should set $M \geqslant 5$ to achieve a short stabilization time. However, if the value of M is too large, the performance in terms of throughput drops since beam sweeps become inefficient. From our experiments, we observe that $M = 6$ is typically a good trade-off. In comparison to the default operation of the router, this means that the beam sweep duration decreases by 82.4%. Moreover, our approach requires fewer beam sweeps due to its high stability. Thus, the time spent on beam training is much shorter.

6.2.4 *Protocol Operation*

Interference from multiple directions is challenging.

Our last experiment evaluates the protocol operation when reacting to interference. We consider intermittent interference from two WiHD transmitters located at different angles on the semicircle. When one of the interfering nodes starts transmitting, the SNR drops at the station. As a result, the default mechanisms in IEEE 802.11ad trigger a sector level sweep to find a better receive beam. However, this new receive

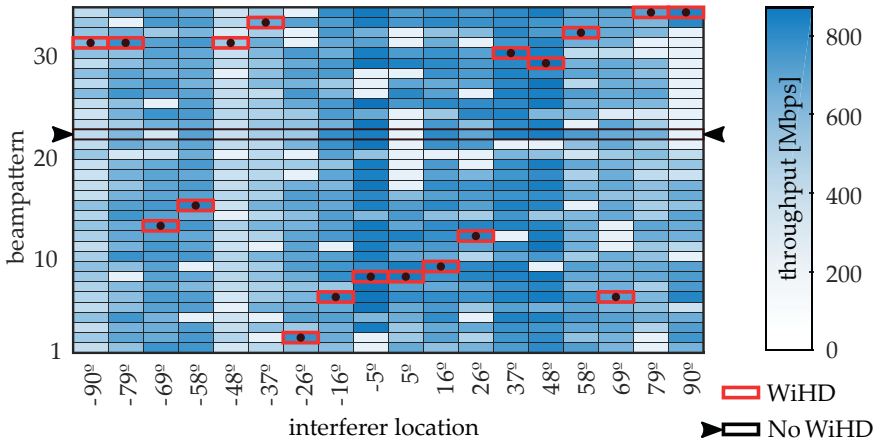

Figure 6.5: TCP throughput for all receive beam patterns and all locations of the interfering WiHD transmitter. The AP is located at −45°.

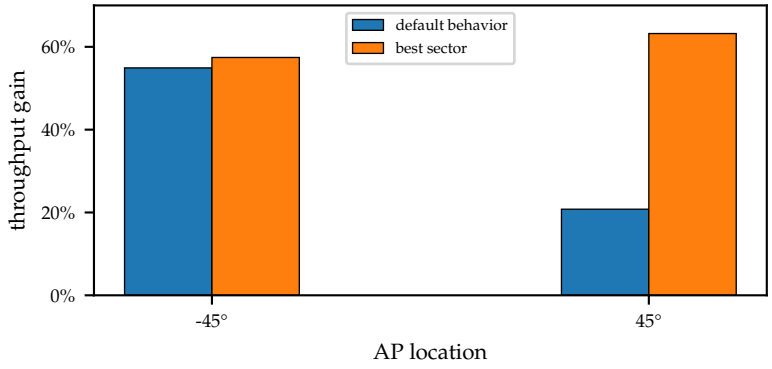

Figure 6.6: Average throughput gain for all locations of the interfering transmitter.

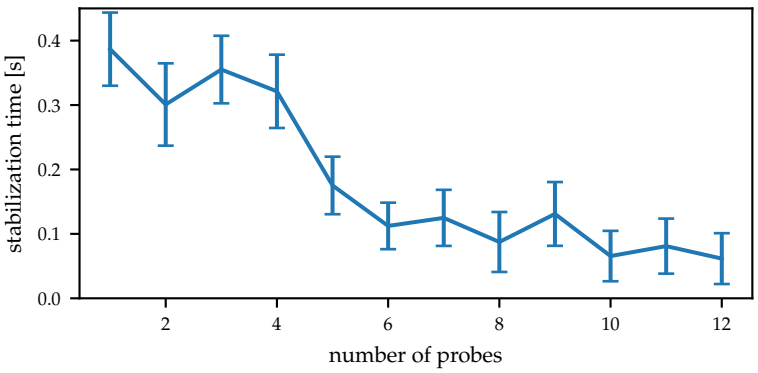

Figure 6.7: Stabilization time of our adaptive beam switching mechanism.

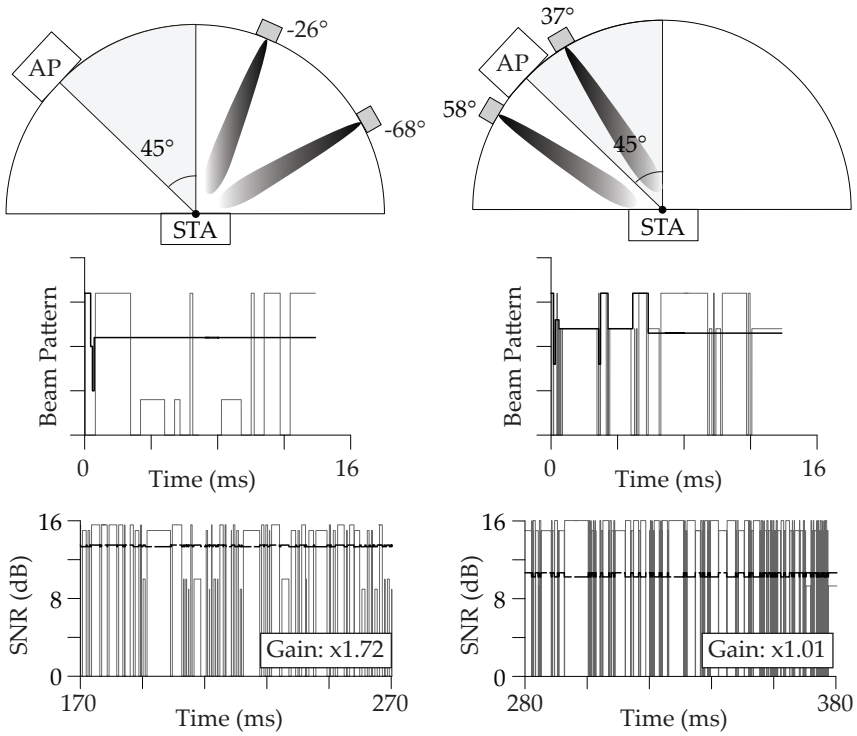

Figure 6.8: Protocol operation with two interfering nodes for (a) a regular case and (b) a challenging case.

beam is unlikely to also minimize the interference from the second WiHD node or a mobile interferer. It causes IEEE 802.11ad to perform sweeps continuously and jumps back and forth between different beam patterns. Such frequent channel probing strongly degrades the performance. In contrast, the system stability s in our protocol prevents such fluctuations. To validate this, we emulate its behavior and that of IEEE 802.11ad on real-world traces collected from our practical testbed.

Our mechanism quickly adapts to changing conditions.

Figure 6.8 depicts two examples of the above scenario. In example (a), the interfering nodes are placed at a reasonable angular distance from the AP. As expected, IEEE 802.11ad continuously triggers beam sweeps and switches among two different patterns. In contrast, our adaptive approach quickly finds a suitable sector and stabilizes within less than a millisecond. Figure 6.8 only shows the selected beam patterns during the first 16 milliseconds of the experiment for clarity. However, the fluctuations of IEEE 802.11ad continue until the end of the experiment. The SNR graph of example (a) shows that our protocol achieves a stable value whereas IEEE 802.11ad experiences frequent SNR drops. As a result, the throughput improves by 72 %. In example (b), the interfering nodes are closer to the AP. This is a particularly challenging scenario since the phased antenna array of the routers is not capable of producing beam patterns which are narrow enough to filter the interference spatially. Figure 6.8 shows

that this increases the stabilization time to about six milliseconds. As shown by the SNR graph, our mechanism chooses a trade-off beam pattern that balances the impact of both interfering nodes. In contrast, IEEE 802.11ad achieves a higher SNR at the expense of costly beam sweeps and high instability. Although there is no throughput gain, we achieve a very high connection stability, which is crucial for upper-layer protocols such as TCP.

Higher layer protocols benefit from high stability.

6.3 DISCUSSION AND SUMMARY

Our adaptive beam switching mechanism mitigates lateral interference in mm-wave wireless networks. It significantly reduces the beam steering overhead in case of intermittent interference from neighboring links. Such interference is critical in practical deployments due to the strong side lobes of phased antenna arrays in consumer-grade mm-wave devices. Whenever interference occurs, our mechanism switches to an alternative beam at the receiver. Thus, the interference falls into a minimum of the beam pattern. The key difference to existing approaches is that ours does not need to probe all possible beam patterns. Only the ones which are most likely to perform better than the current one are needed. We evaluate the performance of our protocol in practical experiments with off-the-shelf devices and achieve average throughput gains of 60 %. In contrast to conventional beam sweeping as in IEEE 802.11ad, the training time reduces by 82.4 %.

Adaptive beam switching reduces the interference among concurrent transmissions.

7 ADAPTIVE BEAM OPTIMIZATION

State-of-the-art beam training uses generic beam patterns.

Designing wireless communication systems that operate at very high frequencies, manufacturers often resort to sub-optimal yet straightforward solutions. This case also applies for beamforming in current off-the-shelf devices. The IEEE 802.11ad standard is limited to perform the sector level sweep with a codebook of generic beam patterns. Such patterns are envisioned to have a uniform narrow shape and cover the entire space around the device. Instead of beamforming towards a specific direction, devices choose the beam pattern out of their codebook which provides the highest gain in that direction. Unfortunately, a large codebook of precise beams causes a significant overhead since all beams are probed sequentially. As a trade-off, practical codebooks typically contain beams that cover a broader area or multiple directions with reduced directionality.

We adapt beam patterns to the current channel conditions.

We design a mechanism that adaptively optimizes the beam for the current channel conditions. Using only non-coherent SNR measurements, we enable full Channel State Information (CSI) extraction on consumer-grade off-the-shelf devices. Our adaptive beam optimization mechanism allows to fully exploit the capabilities of phased antenna arrays in such devices. It probes the channel using carefully engineered beam patterns that allow devices to extract both amplitude and phase information from simple SNR readings. Since SNR measurements are supported on all IEEE 802.11ad devices, our mechanism is widely applicable.

Relative phase measurements reveal the complex channel state information.

By measuring the relative phase among all the elements in the antenna array, we obtain the CSI. Explicitly, we define one antenna element as a reference and measure the relative phase to another element. Both measured elements are switched on while all others are turned off. Then, we transmit a probing frame for different phase values at the antenna, thus cause phase-shifted signals at the receiver. Figure 7.1 illustrates this approach with two-bit phase shifters. In the received signal constellations, the complex gains of the reference (blue) and of the other element (red) sum up. The sums (black) for each of the four probes are affected by slight gain and phase variations. Unfortunately, common off-the-shelf receivers only observe the resulting amplitude regarding the received signal strength; they typically do not expose the complex gains. Since the four probes cover a 2π phase range, their powers must lie on a sinusoidal curve. Our key insight is that the initial phase of that curve is directly related to the relative phase among both antenna elements. By exploiting this property using Fourier analysis to reconstruct the sinusoidal curve,

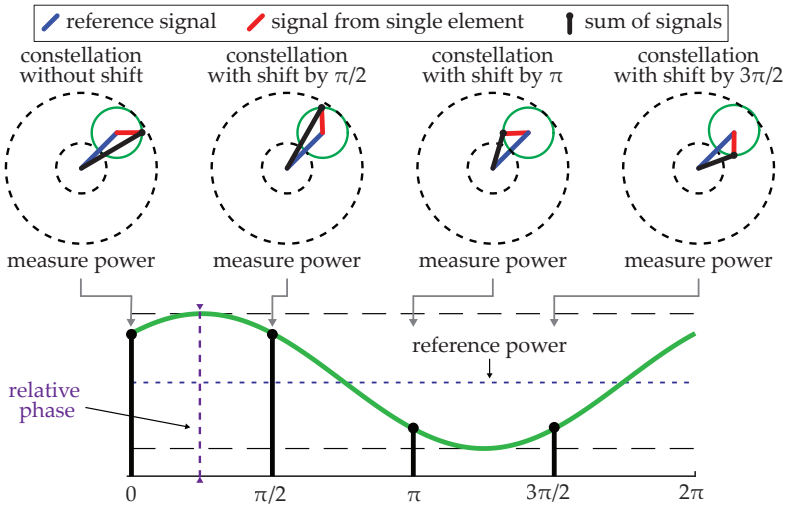

Figure 7.1: Illustration of CSI extraction with phase-shifted SNR measurements. The upper part shows the signal constellations with four different phase shifts, whereas the lower part illustrates how we derive the relative phase using a sinusoidal curve.

we compute the phase shift δ for the received signal. Figure 7.1 shows an example for $\delta = \pi/4$. Repeating the process for each element of the antenna array, we obtain the full CSI and use it to derive beam patterns that maximize the SNR and automatically exploit reflections in the environment.

In the following, we state our protocol design to obtain full CSI using only SNR measurements in Section 7.1. A practical evaluation in an office environment as well as in an anechoic chamber is provided in Section 7.2. Finally, at the end of this chapter, we discuss and summarize our findings in Section 7.3.

7.1 PROTOCOL DESIGN

Our protocol comprises the channel estimation and beam optimization.

Our adaptive beam optimization approach operates between an AP and a station, whereby both benefit from optimized beams. The design of our protocol involves multiple aspects. In the following, we provide on overview on the protocol operations, state our system model, describe the estimation of complex channel gains, and, delve into the generation of optimized beams that fully adapt to the channel.

7.1.1 *Protocol Operation*

Our protocol operation consists of two phases, 1) initialization and 2) continuous adaptation as illustrated in Figure 7.2. In the initialization phase, devices establish a connection using the default IEEE 802.11ad procedure based on generic beam patterns. After that, we exhaustively probe all of the antenna elements of the array. This one-time overhead

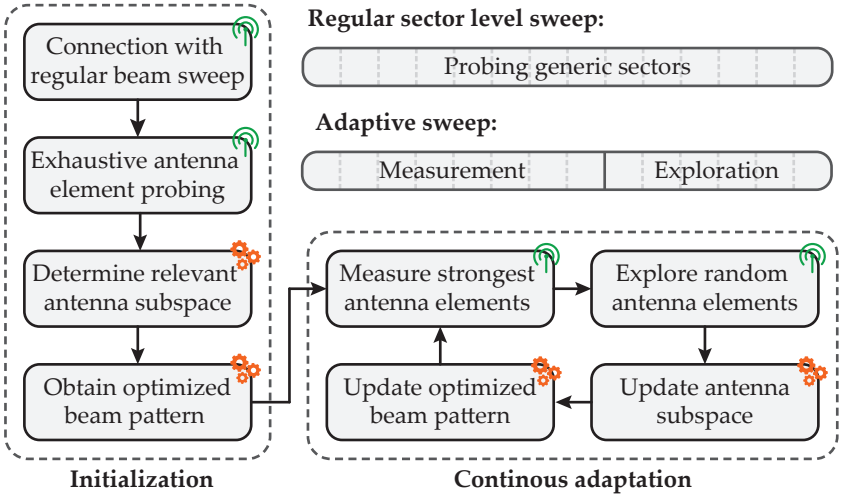

Figure 7.2: Adaptive beam optimization with processing (gears symbol) and communication (antenna symbol) steps.

allows us to obtain the full CSI and to determine which of the elements contribute most to it. Based on this information, we restrict the following channel measurement to those elements. This significant reduction of the probing overhead allows focusing on the subspace of the channel which is relevant for the physical environment. The number of used antenna elements N is determined by

Exhaustive antenna probing causes a one-time overhead.

$$N = \arg\max_n \frac{\left(\sum_{k=1}^{n} a_k\right)^2}{n}, \tag{7.1}$$

where a_k is the amplitude of antenna k, and antenna elements are ordered by strength. To complete the initialization phase, we obtain the beam pattern that maximizes the SNR within the computed subspace and uses it for regular communication. After that, the continuous adaptation phase allows us to re-adjust the beam pattern when the CSI and the environment change during on-going communication. Similarly to IEEE 802.11ad, our mechanisms probes the channel periodically but limits the channel measurement to the strongest antenna elements and thus to the current subspace. However, the subset of relevant antenna elements may change if the environment, for example, changes due to user mobility. To adapt to such changes, we divide the periodic channel probing sweep into a measurement part and an exploration part. The channel measurement is performed over the subspace of strongest antennas, whereas the exploration part probes a subset of randomly chosen antenna elements. In doing so, the latter identifies potential changes in the relevant channel subspace. As shown in Figure 7.2, we choose the overall duration of the sweep such that it is equivalent to that of a regular IEEE 802.11ad sector level sweep. Still, if full CSI is desired, our protocol can easily be reconfigured to probe all antenna elements at the expense of higher overhead. After each periodic sweep, we recompute the beam pattern that maximizes the

Not all antennas require probing during the continuous adaptation.

Our training duration is equal to that of the default operation.

SNR. The complexity of the processing steps shown in Figure 7.2 is negligible due to the very high computational efficiency of our method. It enables our protocol to operate both on powerful APs as well as resource-constrained stations.

7.1.2 *System Model*

Off-the-shelf devices use analog beamforming with a single RF*-chain.*

In our system model, we assume that both transceivers utilize a phased antenna array that is capable of analog beamforming with an arbitrary number of antenna elements. All these antenna elements are connected to a single RF-chain and controllable via a weighting network of amplifiers and phase shifters. This architecture is common for off-the-shelf devices. More advanced beamforming techniques such as digital or hybrid beamforming are expensive and not yet available in consumer-grade devices.

When a transmitter sends a signal x towards a receiver, the signal gets distorted during the transmission. The received signal y is expressed as

$$y = c'Hpx + c'N, \tag{7.2}$$

where c and p indicate the complex antenna gains at the receiver and transmitter, respectively. H denotes the propagation effects of the channel and N is uncorrelated additive white Gaussian noise. The channel H represents geometrical propagation effects such as reflections and blockage but also fading and free-space attenuation. By this definition, the received signal power is computed as $|c'Hp|^2$. For simplicity reasons, we ignore the noise power. Taking into account that most off-the-shelf devices only expose the SNR or RSSI in dB-scale, we obtain the received signal power by $|c'Hp|^2 = 10^{\frac{SNR}{10}}$ using the well-known formula of the SNR in logarithmic units.

7.1.3 *Complex Gain Estimation*

We obtain the complex channel gains from simple SNR *measurements.*

The core contribution of our mechanism is to obtain the complex gain for each antenna element using simple strength readings. To achieve this, we transmit multiple probing signals with different phase shifts and use the received signal strengths to reconstruct the complex gains. In particular, we obtain the relative phase shift of individual antenna elements concerning a reference. Regarding this reference, we distinguish between two different estimation variants for high and low SNR scenarios. The former is the intuitive approach in which we use a single antenna element as the reference and directly measure the relative phase shift of other antenna elements. Unfortunately, the beam patterns of single antenna elements expose a quasi-omnidirectional shape with low gain. It allows channel estimation in all directions when the link quality is sufficient, and devices are close to each other.

However, when devices are far away, or only a non-line-of-sight link is available, the gain of a single antenna element is insufficient to reach the receiver. As a result, the reference cannot be received which makes the channel estimation impossible. For these low SNR scenarios, we propose the second variant that estimates the complex gains by using a directional beam pattern with multiple active antenna elements as the reference instead of the beam of a single antenna. It provides a much stronger reference beam, but only allows the estimation of the complex gains in the specific directions of that particular beam. In case of sudden movement, when the receiver moves out of the reference beam, the channel estimation is likely to fail. Thus, we have a trade-off between both variants that address different application scenarios.

STRONG LINK ESTIMATION. In our first variant, we retrieve the complex gains by enabling a single antenna element as a reference. Given a phased antenna array with multiple elements, where each is identified using an index k, the reference element is denoted with $\bar{k}$. All these elements expose a complex gain p_k on the transmitted signal, such that the overall gain of the array p is composed of $p_k \forall k$. In static environments, the channel H and the complex receiving gain c are assumed to be constant. Thus, p can be measured from the received signal strength.

Strong link estimation uses the best antenna elements as the reference.

Measuring the amplitude of all antenna elements is straightforward. For each antenna element, we generate a beam pattern by switching on a single element at a time. Such a probing pattern for each antenna element k is defined by:

$$p_{k'} = \begin{cases} 0 & k' \neq k \\ 1 & k' = k \end{cases}, \qquad p = \frac{p'}{\|p'\|}, \tag{7.3}$$

where p' refers to the non-normalized version of the gain. Using this approach that requires n probes to measure the amplitude of n antenna elements, we select the element that provides the highest received signal strength as reference $\bar{k}$.

For the phase shift measurements, we need to probe each antenna element, except the reference, multiple times with different phase shifted gains. For the reference element, zero relative phase shift is assumed per definition. Being particularly interested in the relative phase shift, we need to enable both antenna elements k and $\bar{k}$ in a probing pattern. For the element k, we probe four different phase values $p'_k \in \{1, i, -1, -i\}$, while $p'_{\bar{k}} = 1$ and all other elements are set to zero. That is, each probing beam contains the sum of two signals—the signal transmitted by the measured antenna element k with one out of four different phase shifts and the reference signal from element $\bar{k}$. Using a wave decomposition expression and a fast Fourier transform, we then extract the relative phase among the phase shifted measurements as illustrated in Figure 7.1. The detailed formulation

Each antenna is probed multiple times with different phase shifts.

	id	psh (element phase)	etype (element amplitude)
amp. measurements	0	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	**4**, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	1	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, **4**, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	2	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, 0, **4**, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	⋮	⋮	⋮
	29	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, **4**, 0, 0
	30	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, **4**, 0
	31	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, **4**
relative phase measurements	32	**0**, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	**4**, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	33	**1**, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	**4**, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	34	**2**, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	**4**, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	35	**3**, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	**4**, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	36	0, **0**, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, **4**, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	37	0, **1**, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, 0	0, **4**, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, 0
	⋮	⋮	⋮
	150	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, **2**, 0	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, **4**, 0
	151	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, **3**, 0	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, **4**, 0
	152	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, **0**	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, **4**
	153	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, **1**	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, **4**
	154	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, **2**	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, **4**
	155	0, 0, 0, 0, 0, 0, ⋯ 0, 0, 0, 0, 0, **3**	0, 0, 0, **4**, 0, 0, ⋯ 0, 0, 0, 0, 0, **4**

Table 7.1: Sample channel estimation codebook for complex gain retrieval in strong links. The upper part shows 32 sectors for the amplitude measurements, while the lower part 124 sectors for relative phase measurements with the eights element as reference.

Probing beams contain the sum of two signals.

for this mathematical problem is stated in [Pal+18a]. The total number of probes that are required directly relates to the number of elements in the array. For the reference element, a single probe is sufficient to obtain the amplitude; it has a relative phase shift of zero. For all other elements, we need one probe to obtain the amplitude and four probes to estimate the phase shift relative to $\bar{k}$. Thus, for a total number of n antenna elements, our approach requires $n + 4(n-1)$ probes to retrieve the complex gain of the antenna array.

Probing all 32 antenna elements requires 156 beams.

Many of current off-the-shelf devices, such as the Talon AD7200 (Chapter 4) use a phased antenna array with 32 elements. Using the 2-bit phase shifters and 3-bit amplifiers, we set-up a probing codebook as stated in Table 7.1. For the phase shifters, we iterate through the values 0 to 3 and set the amplitude to 4. In preliminary experiments, this amplitude setting provided the best estimation accuracy. In total, this channel estimation results in a codebook with 156 different probing beams. As the devices only support to configure up to 64

different beams at a time, we split the full estimation process into three consecutive sweeps or reduce the number of interested antenna elements. Using a full codebook, we could at least estimate 13 elements simultaneously in a single sweep.

WEAK LINK ESTIMATION. As the channel estimation is only feasible with a strong reference signal and the previous variant becomes inaccurate with a low gain of $\bar{k}$, we propose a second channel estimation variant that operates with weak links. Instead of using a single antenna element as a reference, it selects a directional beam with multiple active antenna elements as the reference. The directionality results in a higher received signal strength of the reference beam thus increase the channel estimation accuracy. One of the generic beam patterns that are configured by default on the devices can be chosen as reference. Due to many active antenna elements, these beams provide a strong gain even when single antenna elements are inactive. They enable the estimation of complex gains on weak links. This becomes useful when the link is blocked or devices communicate over long distances. For this variant, we assume that the phase of the reference beam is zero and determine the relative phase of all other antenna elements.

Weak link estimation uses a directional beam as the reference.

To estimate the channel on weak links, we need to distinguish between two cases when a) the measured antenna element is active in the reference beam and whether it is not. For antenna elements that are inactive in the reference beam, four different phase shifts need probing. In addition to the active elements in the reference beam, the measured element is activated with phase shifts $\{1, i, -1, -i\}$. This allows us to estimate the relative phase as in the strong link estimation variant. To assess the amplitude, we also compute the difference with respect to the reference beam. For antenna elements that are already active in the reference beam, the phase difference cannot be simply extracted. The reference is already affected by this particular element. In this case, only three other phase shifts must be probed and compared to the existing one. It requires a more protracted analytical analysis, but effectively reduces the number of required probes. A detailed explanation of the mathematical model behind his computation is provided in [Pal+18a; Pal+18b].

We differentiate between active and inactive antenna elements in the reference beam.

The channel estimation on weak links requires to generate of probing codebook for a specific reference beam dynamically. Thus, defining fixed codebook with static probing sectors as for the strong link estimation in Table 7.1 is impossible in this scenario. The probing beams highly vary with the antenna elements that are active in the reference beam. Active antenna elements need to measured by the reference beam itself and three additional probes with different phases. For inactive antenna elements, we require four probes to estimate the complex gain. In total, we need $1 + 3k + 4(n - k)$ probes for an antenna array

Channel estimation on weak links requires a dynamic probing codebook.

with n elements of which k are active in the reference beam. Using the default beam patterns of the Talon AD7200 with 8 active elements on average, we need 121 probes to estimate all complex gains of the antenna array.

7.1.4 *Optimized Beam Patterns*

Optimized beam patterns are generated by maximizing the expected signal strength.

Using the previously described channel estimation, we become able to compute beams that directly adapt to the current channel conditions and maximize the SNR at the receiver. As discussed above, the received signal strength is $|c'Hp|^2$. We derive a beam pattern p that maximizes this equation by aligning the phases of all the antenna elements such that their signals do not cancel each other and constructively interfere. As the phase shifters in the array operate in a discrete space, we round the phases to match the hardware's resolution (2-bit for the Talon AD7200). The computational complexity for this beam generation is minimal, as the process requires simple matrix operations only. Thus, we can compute the optimal beam pattern given specific channel conditions highly efficient.

7.2 PRACTICAL EVALUATION

We implement full CSI extraction on commodity hardware.

For performance evaluation in two experiments, we implement our adaptive beam optimization mechanism on commodity off-the-shelf hardware. In particular, we use the Talon AD7200 devices of our testbed experimentation platform from Chapter 4. Setting specific beams requires full access to the beamforming control and the understanding of the antenna layout. To include custom beam patterns for the channel probing and communication, we modify the sector sweep behavior. Our performance measurements are conducted in an open-plan office environment as well as an anechoic chamber. By embedding our implementation in the regular operation on the devices, our results are obtained from actual data transmissions in the testbed and not generated in post-processing. This approach enables us to quantify the performance concerning the SNR, data rate, and throughput using the Transmission Control Protocol (TCP). In the following, we first state our testbed setup. Then, we evaluate the signal strength, throughput, as well as the data rate, and finally investigate the pattern shapes of optimized beams.

7.2.1 *Testbed Setup*

Our prototype implementation consists of an IEEE 802.11ad AP that serves one or multiple stations. We use the Talon AD7200 devices to realize both, the AP and the stations. For ease of experimentation, we control all devices from a laptop using our testbed automation

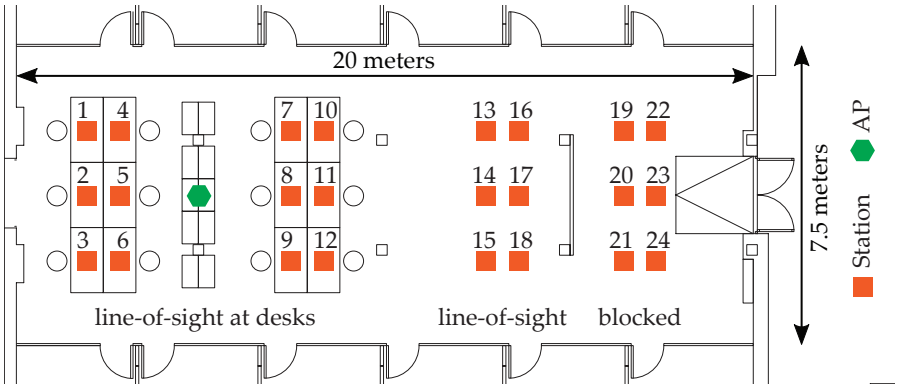

Figure 7.3: Experiment setup in an indoor open-plan office environment.

The evaluation scenario contains one AP and multiple stations.

system over a common 2.4 GHz Wi-Fi network. Using this control and management network, we instruct the individual stations to connect to the AP, generate traffic in the IEEE 802.11ad network, and collect statistics. For simplicity, we run specific components of our algorithm on the laptop, such as the reconfiguration of the codebooks. However, these portions of the code could also run directly on the devices since our approach has very low computational requirements.

We select different codebooks and measure the signal strength and throughput.

The operation of our prototype is as follows. First, all devices, the AP and the stations, load their default codebook with generic beam patterns. Then, the stations establish a connection to the AP using the regular IEEE 802.11ad procedure. After that, we load an initial measurement codebook on the devices and collect the resulting SNR readings to compute a particular transmit beam pattern that maximizes the signal quality for each link. We do this for the AP and the stations to analyze both the up- and downlink. To specifically optimize the beam patterns under low SNR conditions, we employ the weak link estimation scheme described in Section 7.1.3. Finally, all devices are re-configured with the optimized beam patterns and generate traffic in the network. In particular, we use `iperf3` [iPerf] to sequentially generate uplink and downlink traffic for ten seconds on each link. In each experiment, we repeat this measurement twenty times while collecting statistics such as throughput, data rate, and signal quality.

Measurements are taken at 24 distributed stations in an office environment.

We evaluate our approach with performance measurements in an open-plan office environment as shown in Figure 7.3. Additionally, we characterize the generated beam patterns by radiation measurements within an anechoic chamber. The open-plan office area has a size of $20\,\mathrm{m} \times 7.5\,\mathrm{m}$ and is surrounded by twelve individual offices. The AP is deployed below the ceiling next to the desks. Further, we place the stations at 24 different positions within the room. Twelve of these positions correspond to typical workspaces at the desks. The stations are oriented such that they mimic the usual placement of a laptop. Next, six positions are in an area without desks but still in line-of-

sight of the AP. Finally, six additional positions are chosen behind an isolated wall that separates the entry area from the open-plan office space. These positions are not in line-of-sight, which means that reflections are required to reach the AP.

7.2.2 *Signal Gain Maximization*

Optimized beams increase the received signal strength.

As described in Section 7.1.4, we design our mechanism to maximize the signal strength at receivers. In the first part of our evaluation, we conduct experiments to measure the SNR gains that are achievable with our approach in comparison to the default IEEE 802.11ad operation with generic beam patterns. For both we measure the SNR at all positions of our office testbed on the up- and downlink. As shown in our results in Figure 7.4, the generic beams patterns achieve an average SNR of 6.46 dB for the uplink and 5.98 dB for the downlink. Our optimized beams that adapt to the current channel conditions increase the SNR up to 10.81 dB and 10.38 dB for the up- and downlink, respectively. The Cumulative Distribution Function (CDF) reveals that 90 % of the measurements with generic beam pattern expose an SNR lower than 12.55 dB. In contrast, the optimized beams achieve an SNR higher than this in about 45 % of our measurements.

Up- and downlink results are uncorrelated and location dependent.

Table 7.2 lists the SNR gains achieved at individual measurement positions in our testbed environment. We find that the distance between AP and stations cause only a minor effect. However, the variance among different positions is high. In some spots, only marginal gains are observable while in others they are up to 8.00 dB. Some positions can be easily reached via reflectors while others cannot. At location 20, for instance, the station could not establish a link at all. Moreover, the gains on the up- and downlink are uncorrelated and may exploit different reflection paths. In general, we observe an average increase of the SNR by 4.35 dB on the uplink and by 4.40 dB on the downlink when using our optimized beams instead of the generic ones.

Constellations of optimized beams are less noisy.

To validate the high SNR gains, we additionally analyze the signal constellations at the physical layer. As the off-the-shelf devices do not allow to access the raw signal samples, we utilize a platform similar to our channel sounding platform as described in Section 3.1. In particular, we place a Sivers IMA FC2221V/01 V-band down-converter with a horn antenna at the location of the receiver and capture the raw samples using a Keysight DSOS254A oscilloscope. The bandwidth of the oscilloscope is sufficient to capture and decode the full IEEE 802.11ad signal. Doing so, we obtain the In-phase and Quadrature (IQ) constellations of data frames transmitted with both the generic and optimized beam patterns and different Modulation and Coding Schemes (MCSs). Figure 7.5 shows an example for frames with MCS 8, which use a $\pi/2$-QPSK modulation. Our optimized beam patterns show a visible accuracy improvement and lead to constellations that

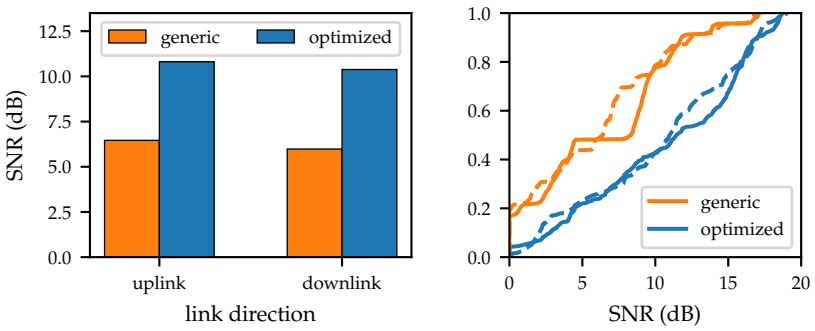

Figure 7.4: Average SNR and CDF of generic and optimized beam pattern for uplink (solid line) and downlink (dashed line) measurements.

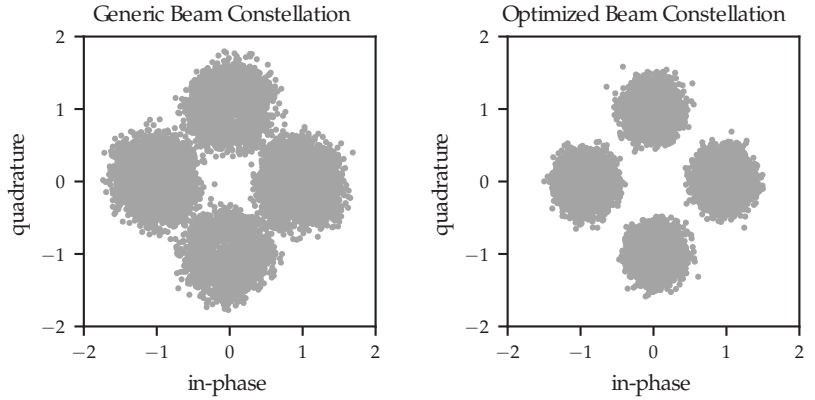

Figure 7.5: Constellation diagrams of MCS 8 encoded frames transmitted with generic (left) and optimized beam patterns (right).

LOC.	UPLINK	DOWNLINK	LOC.	UPLINK	DOWNLINK
1	3.46 dB	2.08 dB	13	5.37 dB	6.01 dB
2	−0.26 dB	2.42 dB	14	2.00 dB	2.18 dB
3	2.27 dB	1.47 dB	15	6.89 dB	5.16 dB
4	6.76 dB	6.99 dB	16	4.28 dB	3.33 dB
5	6.47 dB	1.75 dB	17	2.08 dB	3.43 dB
6	5.11 dB	6.12 dB	18	3.71 dB	3.94 dB
7	5.45 dB	6.43 dB	19	4.45 dB	5.46 dB
8	1.40 dB	1.92 dB	20	n/a	n/a
9	6.80 dB	5.42 dB	21	4.97 dB	4.68 dB
10	7.39 dB	5.13 dB	22	2.85 dB	1.45 dB
11	3.93 dB	8.00 dB	23	4.18 dB	4.20 dB
12	5.66 dB	6.67 dB	24	4.81 dB	7.02 dB

Table 7.2: SNR of our optimized beams relative to the generic ones. Each row corresponds to a specific station location in our evaluation. The station at location 20 could not connect.

are significantly less noisy than those of the generic beam patterns. They effectively reduce the number of symbol errors, which in turn improves the stability of the communication link.

7.2.3 *Pattern Shapes*

Optimized beams exhibit a higher directionality.

To characterize the optimized beam patterns that are computed with our CSI estimation, we perform additional antenna radiation measurements in an anechoic chamber. In particular, we mount an AP in the center of the chamber on a pan-tilt unit to orient it in different azimuth and elevation directions. A second device is configured as station and placed few meters away with a fixed orientation to measure the CSI as described above. First, the AP is oriented to different directions in the azimuth plane with zero elevation. In all these orientations, we compute the optimized beam patterns according to the strong link estimation approach. Second, these beams are configured in a codebook and configured on the devices. While the AP iterates through all possible orientations, both devices record the received signal strengths for each of the configured beams. Doing so, we obtain their radiation patterns without any environmental reflections. Figure 7.6 illustrates some of these beam patterns shapes for azimuth orientations of $-30°$, $-10°$, $10°$ and $30°$ and zero elevation. Besides the optimized beam patterns, the figure also depicts the generic beam patterns that the device would choose according to the default IEEE 802.11ad operation. These results show that our approach increases the achievable SNR by about 5.17 dB and visibly lead to much more narrow beam lobes. Thus, adaptive beam optimization provides a higher directionality.

A uniform coverage enables more flexible device deployments.

While rotating the transmitter, we continuously measure the coverage of the optimized and generic beam patterns. In each orientation, the AP measures the CSI and computes an optimized beam. The station measures the signal strength of the best generic beam as well as of the beam that is optimized for that particular orientation. Thereby, we assess the coverage of the default IEEE 802.11ad operation and our adaptive beam optimization approach as shown in Figure 7.7. With the generic beam patterns, we find a strong gain towards the front of the device while the coverage in the back is significantly lower. During the default operation, devices should face each other to achieve the best link quality. Our approach provides uniform coverage in all directions. It facilitates a more flexible device deployment and allows APs to be mounted in locations where they do not necessarily face the stations.

Chapter 10 provides additional measurements of our optimized beam patterns in three-dimensional environments and details the complex gains of individual elements in the antenna array.

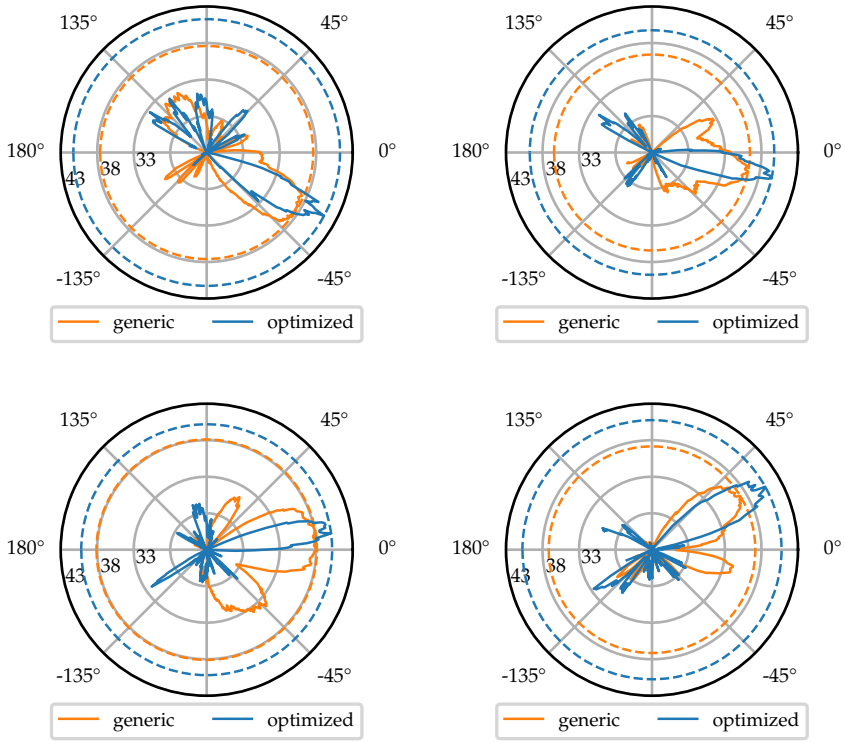

Figure 7.6: Radiation patterns of optimized and generic beams in comparison. Dashed lines show the maximum achievable gain.

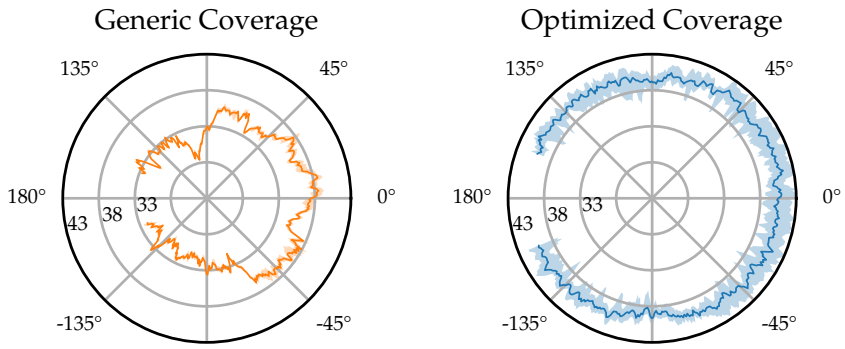

Figure 7.7: Measured coverage of generic and optimized beam patterns.

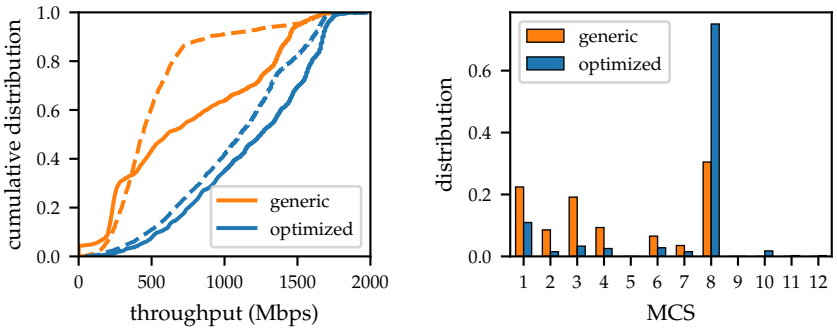

Figure 7.8: Line-of-sight throughput CDF and MCS histogram. Dashed lines indicate the downlink, whereas solid lines the uplink.

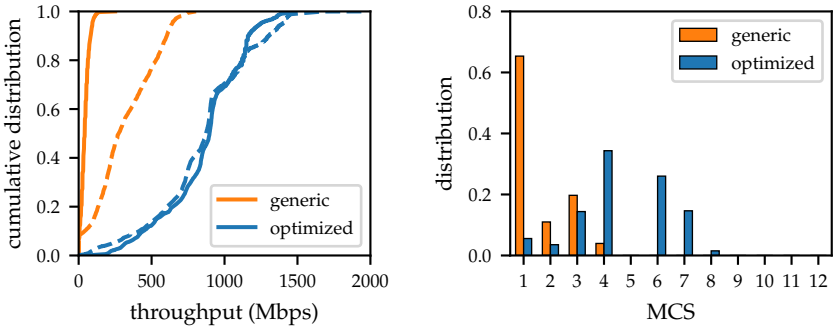

Figure 7.9: Non-line-of-sight throughput CDF and MCS histogram. Dashed lines indicate the downlink, whereas solid lines the uplink.

7.2.4 *Throughput Improvement*

Optimized beams allow devices to select higher MCS

The vast SNR improvement in Section 7.2.2 enables our devices to switch to a higher MCS at most locations. As a result, we obtain significant throughput gains when generating TCP traffic with `iperf3` in the network. Figure 7.8 exemplary depicts the CDF of the achieved TCP throughput at a location in line-of-sight of the AP in our testbed environment. We achieve gains of 58 % for the uplink and 102 % for the downlink, respectively; that is, the downlink throughput effectively doubles. Moreover, the right graph in Figure 7.8 depicts the distribution of the used MCS in all transmitted up- and downlink frames for that location. As expected by the higher SNR the devices tend to switch to higher MCS by using the optimized beams instead of the generic ones. For instance, with our optimized beams, the devices select MCS 8 for more than 70 % of the transmitted frames. The default operation sticks to a lower MCS for about 70 %.

In Figure 7.9, we provide a similar analysis for a blocked location without any line-of-sight link to the AP (see Figure 7.3). Due to the blockage, the AP and the station must communicate via reflections. Our results indicate that the generic beamforming mechanism in

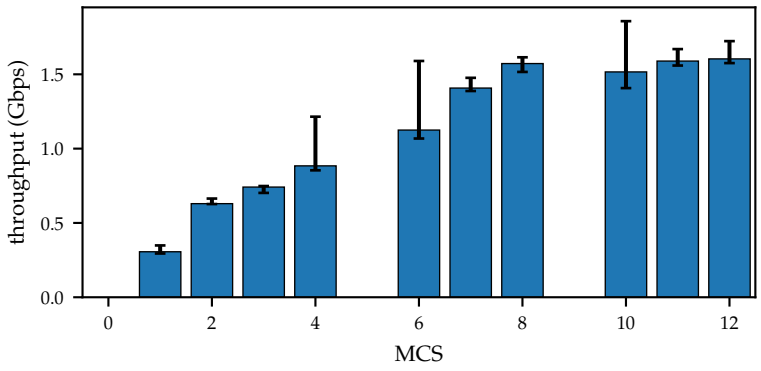

Figure 7.10: Throughput saturation at high MCS. MCS 5 and MCS 9 are not implemented.

IEEE 802.11ad performs poorly in such a scenario. It only achieves an average throughput of 42.7 Mbps on the uplink and 314.54 Mbps on the downlink while mostly operating with MCS 1 to MCS 3. Despite missing a strong link, our optimized beams inherently align towards the reflection path and thereby increase the signal quality such that transmissions with modulation higher than MCS 3 become feasible. On both, the up- and downlink, we achieve a throughput of about 855 Mbps, which is a remarkable gain. The generic beam patterns are inappropriate for complex deployments with a high amount of blockage. Our approach, in contrast, achieves order-of-magnitude gains and facilitates a considerable throughput.

We enable efficient communication over reflected signal paths.

7.2.5 *Expected Data Rate*

Despite high SNR improvements and clear gains in the pattern shapes, the throughput improvements at certain locations are marginal. To explain this behavior, we delve into the internal operations of the device. In many situations, it does not switch to higher modulations even though the signal strength would be sufficient and increases by factor 3 or more. In particular, the rate adaptation often does not exceed MCS 8 even though the router implements up to MCS 12. Figure 7.10 shows that in a typical scenario the throughput saturates with MCS 8 at about 1.5 Gbps. We observe this effect on many locations of the testbed independent of the length of the link and other aspects that impair the signal quality. Only situations in which the link is not already saturated lead to high-performance gains. The rate adaptation disfavors MCS configurations beyond MCS 8 that all base on a Quadrature Amplitude Modulation (QAM) with 16 constellation points (16-QAM). We assume that this is an intentional design decision of the hardware manufacturer that tolerates certain limitations to lower the production costs. The typical traffic on the router is limited by the Gigabit Ethernet interface anyway such that IEEE 802.11ad is

Devices are already saturated with an MCS below maximum.

not the bottleneck and already saturates the system with an MCS below maximum. In this regard, a conservative rate adaption that does not switch to the maximum is sufficient and more robust against decoding errors. As a result of these limitations, the benefits of our beam optimization only become visible when a link using the generic beam patterns of IEEE 802.11ad operates below MCS8. Otherwise, the SNR improvement does not translate into an effective throughput gain.

Next generation systems supporting higher data rates are likely to benefit from our optimization.

Next generation devices that implement the IEEE 802.11ay standard are expected to provide better support for data rates beyond 1 Gbps. Those could highly benefit from optimized beams for throughput enhancements. To illustrate the benefits that would be achieved on hardware that can process higher data rates, we compute the data rate according to the expected MCS for a given SNR. To this end, we build on the signal quality thresholds recommended in the IEEE 802.11ad standard [IEE14] as listed in Table 2.1. Transforming the measured SNRs to sensitivity thresholds, we look up the recommended MCS and obtain the expected data rates. The values are then cross-validated with our measurements.

Expected data rates increase by up to a factor of 2.49.

Table 7.3 lists the expected data rate for the generic and our optimized beams as well as the achieved gains averaged over the up- and downlink. In average, we increase the data rate by 60 % and reach 3.20 Gbps while the default operation with generic beam patterns only facilitates 2.13 Gbps. Strong gains are achieved across most locations of the testbed. At specific stations, the gains reach up to a factor of 2.49. In a few cases, the achieved data rate is still limited. This effect occurs in particularly challenging locations, such as the desks which are very close to the AP. Since the AP hangs from the ceiling, the vertical angle of the link towards the station located at the desks is very steep. Due to the layout of the antenna array (see Section 4.3.2), the vertical steering capability of the devices is limited. While those links are inherently weak, our approach still provides a substantial improvement even in such challenging scenarios. It enables the system to reach a higher MCS in most cases.

7.3 DISCUSSION AND SUMMARY

The training overhead is adjustable.

Our adaptive beam optimization requires additional SNR feedback. Instead of feeding back only the ID of the best antenna beam pattern as with conventional beam training, the SNRs of all the beam patterns that were probed are required. Furthermore, the overhead to acquire full channel state information would be higher than that of the conventional IEEE 802.11ad training. However, measuring the full CSI for all antenna elements is usually not necessary. Our mechanism can flexibly adjust the number of active antennas based on how rapidly the sub-channel changes, thus, allowing us to reduce the overhead and achieve a desired accuracy. Probing a single sector in the sector level

LOCATION	GENERIC	OPTIMIZED	GAIN
1	0.79 Gbps	1.48 Gbps	1.86
2	0.77 Gbps	0.95 Gbps	1.23
3	0.77 Gbps	1.09 Gbps	1.42
4	2.50 Gbps	4.56 Gbps	1.82
5	3.98 Gbps	4.62 Gbps	1.16
6	2.70 Gbps	4.09 Gbps	1.52
7	2.52 Gbps	4.58 Gbps	1.82
8	4.62 Gbps	4.62 Gbps	1.00
9	2.29 Gbps	4.01 Gbps	1.75
10	2.03 Gbps	3.62 Gbps	1.78
11	2.68 Gbps	4.51 Gbps	1.69
12	1.93 Gbps	2.81 Gbps	1.46
13	2.66 Gbps	4.57 Gbps	1.72
14	4.24 Gbps	4.62 Gbps	1.09
15	2.50 Gbps	4.18 Gbps	1.67
16	2.20 Gbps	2.70 Gbps	1.23
17	2.41 Gbps	2.81 Gbps	1.17
18	1.12 Gbps	2.26 Gbps	2.03
19	1.30 Gbps	2.44 Gbps	1.87
20			
21	2.09 Gbps	3.47 Gbps	1.66
22	0.77 Gbps	1.17 Gbps	1.53
23	0.77 Gbps	1.92 Gbps	2.49
24	1.34 Gbps	2.48 Gbps	1.86

Table 7.3: Expected data rate with generic and our optimized beams for each specific station location in our testbed.

sweep takes about 2.8 µs (see Chapter 5). Hence, probing all $N = 32$ antenna elements would need $N + 4(N - 1) = 156$ probes, which results in a total training time of about 2.8 ms. With 13 active antenna elements, we perform the beam training as fast as the IEEE 802.11ad sector level sweep with 64 sectors in about 1.2 ms. As not all antenna elements contribute to good beam patterns, learning a sub-space of the channel with less active antenna elements is usually sufficient and provides a good trade-off between overhead and performance.

The estimation should adapt to the channel dynamics.

This flexible overhead also allows our approach to adapt to mobility which requires to train the beams continuously. By adapting the number of active antennas in the array, we can select a training overhead that matches the dynamics of the devices and the environment. Specifically, we only need to probe a small sub-space of the channel in case the movement is small.

The overhead does not increase with the hardware precision.

Throughout our evaluation, our CSI estimation mechanism probes each possible value of the phase-shifters. The resulting overhead is limited for two-bit phase shifters that are prevalent in off-the-shelf devices but would increase significantly for higher-resolution hardware. However, four samples are sufficient to reconstruct the sinusoidal curve. More probes only compensate for inaccuracies in the SNR readings. To keep the overhead low, we propose to probe four phase shift values even if more combinations are possible.

Our approach enables CSI extraction from non-coherent signal strength measurements.

In summary, obtaining full CSI on mm-wave devices enables beamforming to achieve higher directionality. It allows for more accurate beam steering and exploits multi-path effects caused by reflectors and obstacles in the environment. Unfortunately, this requires complex channel measurements which are not provided by commercial off-the-shelf devices. We propose an adaptive beam optimization scheme that extracts full CSI on such devices from non-coherent signal strength measurements. In particular, we engineer beam patterns with constant phase shifts that allow extracting CSI regarding the phase and magnitude from simple SNR readings. As such readings are available on most devices, our approach is easily portable to any hardware.

Optimized beams fully adapt to the environments.

Moreover, we use the obtained CSI to compute beam patterns that adapt to the environment and maximize the signal strength at receivers. In this way, we avoid destructive interference and inherently choose the best available path between transmitter and receiver. Using our testbed experimentation framework, we implement our scheme on IEEE 802.11ad tri-band routers by accessing the beamforming control of the integrated 32-element phased antenna array. Measurements of the optimized beam-pattern reveal that our scheme significantly increases the directionality. Evaluations in a real-world office environment show that our optimized beams increase the signal strength by 4.38 dB and achieve a 60 % higher data rate.

Part IV

SECURITY

This part of the thesis addresses security aspects of mm-wave communication systems. We investigate the feasibility of eavesdropping on reflections with environmental objects in Chapter 8. Followed by Chapter 9, in which we analyze attacks on the beam training in IEEE 802.11ad and tamper with the beam steering.

8 EAVESDROPPING ON REFLECTIONS

Directionality impedes eavesdropping and other network attacks.

The high carrier frequency of mm-wave communications induces very specific signal propagation effects. Since path loss increases quadratically with the frequency, high antenna directionalities are required to realize links at Wi-Fi scale distances. Indeed, IEEE 802.11ad specifies beamwidths as small as 3 degrees [Nit+14] to compensate the attenuation. Because of the narrow beamwidth, it is often asserted that mm-wave networks are inherently resilient to eavesdropping. Many publications [Fre13; Yan+15] assume that eavesdropping is infeasible if the eavesdropper is forced to locate itself several degrees off the path between the transmitter and the receiver. However, at the mm-wave frequencies, even small-scale objects can cause significant reflections. In this chapter, we practically investigate the vast impacts of inconspicuous objects on mm-wave security. We assess the feasibility of eavesdropping on directional mm-wave links in Section 8.1, perform experimental measurements in Section 8.2, and discuss and summarize our findings in Section 8.3.

8.1 ATTACK METHOD

Mm-wave signals reflect on many environmental objects.

Eavesdropping is a prevalent attack in wireless communications and easy to mount with off-the-shelf equipment. Attacks typically exploit the omnidirectional propagation characteristics to overhear the communication at remote locations. Due to the high directionality, this is challenging in mm-wave communications. Mm-wave signals are easy to reflect on a variety of plain surfaces, walls, and buildings [Rap+14]. Exploiting these reflections may compensate the challenge of directionality from a security perspective and facilitate eavesdropping outside of the intended signal beam's coverage.

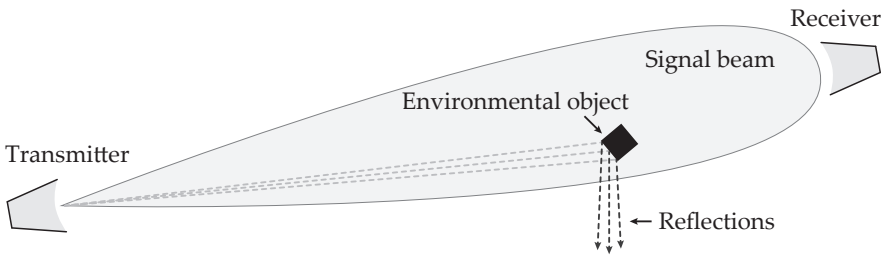

Figure 8.1: Small-scale object exploited by an eavesdropper to create a virtual periscope and reflect the signal out of the intended signal beam.

Reflections may facilitate eavesdropping on directional links.

Eavesdroppers may intercept highly directional transmissions by creating a virtual periscope using a small-scale object as a reflector. In contrast to large-scale reflections that make eavesdropping evident to the receiver, we consider inconspicuous environmental objects are sufficiently small not to impede the intended communication. These objects are placed to redirect only parts of the transmitted signal outside of the designated signal beam as illustrated in Figure 8.1. Varying the geometry and material properties of potential reflectors facilitates an eavesdropper to decode the reflected signal and simultaneously impede the attack from being detected. In particular, we consider three classes of attackers:

We consider different attacker classes: the object manipulator, the nomadic attacker, as well as the opportunistic stationary attacker.

- *Object manipulator.* This attacker tampers with the environment, by placing or moving small-scale physical objects. It carefully manipulates objects in proximity of the signal beam. Thus, these objects cause reflections that are directed towards the eavesdropping antenna.
- *Nomadic attacker.* This attacker is mobile, but cannot manipulate physical objects. Nomadic attackers do not place any additional objects but instead, find a favorable location to exploit existing reflections from the environment.
- *Opportunistic stationary attacker.* This attacker neither moves itself nor any environmental object. It avoids any suspicion and possibly leaves only an eavesdropping 'bug' in the environment. Consequently, this attacker must solely rely on reflections towards its specific location, which are caused by environmental objects within the narrow beam of the intended communication.

In the following, we state our system model, detail the attacker classes, describe the topology and environment, and indicate our performance metrics.

8.1.1 *System Model*

Our system model comprises Alice, Bob, and Eve.

Our system model consists of three communication parties *(1)* a transmitter Alice, *(2)* a receiver Bob, and *(3)* an eavesdropper Eve. Alice transmits a signal towards Bob that she wants to keep secret from Eve by using a narrow beamwidth. We assume that both antennas of Alice and Bob are perfectly aligned and transmit in the optimal direction. Eve aims at revealing the information that Alice sends to Bob without obstructing it. She tries to receive reflections from objects in the signal beam. For convenience, Eve is assumed to use the same hardware as Alice and Bob, acts passively, and is only listening for signals from Alice.

To set up a link budget, we use the Free-Space Path Loss (FSPL) model. It is a good approximation for the propagation loss experienced

at a certain distance from the transmitter in free space without any environmental obstacles. According to [Whi05], the FSPL is defined as

We set up a link budget model to assess the received signal strength in free space.

$$\mathrm{FSPL} = \left(\frac{4\pi d f}{c}\right)^2, \tag{8.1}$$

where f is the carrier frequency of the signal, d the distance to the transmitter, and c the speed of light. We express reflections and blockage of objects inside the signal beam by additional gains and blockage. In particular, the received signal strength with reflections in dB-scale is determined by

$$P_r[\mathrm{dB}] = P_{tx} + G_{tx} - \mathrm{FSPL}_d + G_r + G_{rx}, \tag{8.2}$$

where P_{tx} is the transmitted power, G_{tx} and G_{rx} are the antenna gains at the transmitter and receiver, FSPL_d denotes the FSPL at a distance d, and G_r is the reflection gain. In terms of blockage, the received signal strength is expressed by

$$P_b[\mathrm{dB}] = P_{tx} + G_{tx} - \mathrm{FSPL}_d - L_b + G_{rx}, \tag{8.3}$$

where L_b denotes the blockage loss.

We make the same assumptions for all devices.

Assuming that all devices have identical hardware components and operate at the same beamwidth, we set G_{tx} equal to G_{rx} and consider P_{tx} to be constant. Thus, the received signal strength only depends on the reflectivity, blockage, and distance between the devices. This assumption implies that variations in distance compensate bad reflectivity or blockage and vice versa. Practical antennas for mm-waves achieve directivity gains above 20 dB. Objects of different size, shape, and orientation may feature different characteristics. Although this model does not incorporate any misalignment of the antennas and reflectors, it provides a basic understanding of the scenarios covered in this work and guides us in specifying the attacker models in the following.

8.1.2 *Attacker Classes*

Throughout this work, we distinguish three attacker classes, *(1)* the object manipulator that moves and places objects to cause reflections, *(2)* the nomadic attacker that moves itself and exploits reflections of existing objects, and *(3)* the opportunistic stationary attacker that can neither move nor manipulate the environment. In detail, the attackers are specified as follows.

The object manipulator tampers with objects in the environment.

THE OBJECT MANIPULATOR. The object manipulator considers Eve to be located at a fixed location outside of the coverage of the intended signal beam. From this location, it is impossible to receive the signal directly. However, Eve tampers with the environment and places arbitrary objects to cause reflections towards her direction. She is able

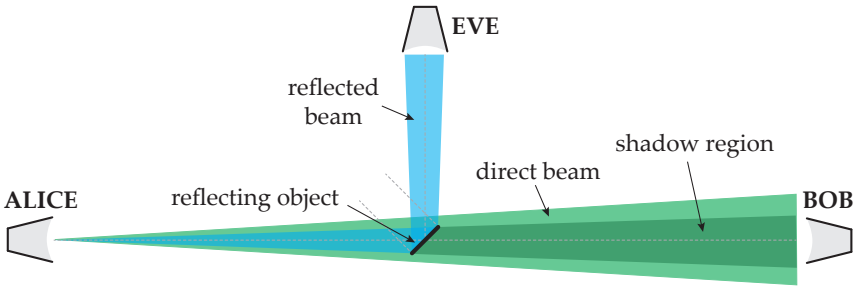

Figure 8.2: System model showing the setup of Alice, Bob, and Eve with reflections of the signal.

to steer her antenna towards this object to optimally receive reflections of the transmitted signal. By doing so, she aims at obtaining a signal quality sufficient for information decoding. At the same time, Eve tries to remain invisible to Alice and Bob by causing only marginal blockage of the direct signal transmission.

The nomadic attacker is mobile a seeks for good eavesdropping locations.

THE NOMADIC ATTACKER. In contrast to the previous class, the nomadic attacker does not assume Eve to change the environment but tries to exploit existing environmental reflections. Eve freely chooses a location outside the beam coverage and steers her antenna towards any reflector in the environment. Unable to affect the blockage at all, she only aims at maximizing her received signal quality by seeking for the optimal eavesdropping location and orientation. Relying on existing reflections is more challenging than manipulating the objects. However, detecting the nomadic attack is more difficult because nothing in the environment changes from the normal operation.

The opportunistic stationary attacker acts completely passive.

THE OPPORTUNISTIC STATIONARY ATTACKER. In the opportunistic stationary attacker model, Eve can neither manipulate the environment nor move herself to an optimal position. This implies that she must entirely rely on environmental objects in the hope that a signal may eventually reflect towards her location. Similar as in the nomadic attacker model, Eve does not cause any blockage. Her only option is to steer her antenna from a fixed location. This class constitutes the weakest attacker in our consideration but makes Eve nearly impossible to detect. Neither the environment nor Eve's location causes any suspicious change.

8.1.3 *Topology and Environment*

In our testbed environment, reflecting objects are located on the center line of the narrow beam between Alice and Bob. This is the optimal case that causes the highest reflection and blockage. As illustrated in Figure 8.2, Bob is located in the shadow region behind the blocking and

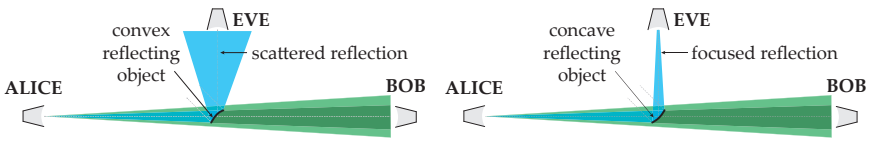

(a) convex reflecting object (b) concave reflecting object

Figure 8.3: Variation of the system model with different surface shapes of the reflecting object.

reflecting object. Eve resides outside the direct coverage of the signal beam and only receives the reflections that bounce off the object. We consider reflecting objects with different characteristics of reflectivity and blockage. Both of these characteristics vary with the material and size, as well as with the structure and shape of the object. With planar reflector surfaces the transmitted and reflected beams expose the same beamwidth. Convex surfaces, in contrast, disperse the reflections in different directions as illustrated in Figure 8.3a. On the contrary, with concave reflector surfaces, the signal focuses towards a certain focal point, as shown Figure 8.3b. At this point the signal bundles its power and can be received with high strength. Throughout our investigation, we only consider first-order reflections and omit additional reflectors.

Objects may have convex or concave surfaces.

8.1.4 *Performance Metrics*

To evaluate the performance of eavesdropping, we measure the signal strength and Bit Error Rate (BER) during data transmission. The signal strength at Bob and Eve reveal the effective reflectivity r and blockage b of an object as

We measure the signal strength and bit errors.

$$r = \max\left(s_{\text{Eve}}\right) \;/\; \max\left(s_{\text{opt}}\right) \tag{8.4}$$

$$b = 1 - \left(\max\left(s_{\text{Bob}}\right) \;/\; \max\left(s_{\text{opt}}\right)\right). \tag{8.5}$$

In this equation s_{Bob} and s_{Eve} are the received signal strengths at Bob and Eve in linear scale. The optimal received signal strength from a direct transmission without reflectors is denoted by s_{opt}. Further, the secrecy capacity [BR06] expresses the performance of eavesdropping as

$$c_s = \log_{10}\left(1 + s_{\text{Bob}}\right) - \log_{10}\left(1 + s_{\text{Eve}}\right). \tag{8.6}$$

The normalized secrecy capacity represents the advantage in signal quality of Bob over Eve by residing in the beam. Its maximum value of 1 implies that Eve is unable to decode anything from the signal. In contrast, the lowest value of 0 indicates that the signal strength at Eve is at least as high as at Bob which facilitates perfect eavesdropping. Additionally, the effective blockage represents the attenuation of the

The secrecy capacity describes the eavesdropping resistance.

We determine the effective blockage and relative reflectivity.

signal by placing the reflector in the beam. The reflectivity is the relative reflected signal strength compared to the signal strength at Bob in unblocked transmission. Blockage and reflectivity are both considered in the secrecy capacity. Eavesdropping achieves the best performance by lowering the secrecy capacity. Within our attacker model, we aim at low secrecy capacities while simultaneously reducing the effective blockage to remain undetected.

8.2 PRACTICAL INVESTIGATION

Experiments are performed with our channel sounding platform.

To develop an understanding of the aforementioned attacker classes, we evaluate the eavesdropping performance using our channel sounding platform from Section 3.1 in practical experiments. Using the Software-Defined Radio (SDR) based setup, we transmit custom signals and evaluate the received and decoded data. Our evaluation scenario comprises a transmitter and a receiver with directional antennas as well a variety of small-scale objects between them. For each object, we measure the reflections towards an eavesdropper and the blockage of the targeted transmission. We vary the shape, size, and material of the object to represent a wide variety of ordinary small-scale objects spanning from coffee cups to cell phones. Such objects could be placed in a particular location by the *object manipulator* or represent those that are common in the environment without object manipulation. The eavesdropper is placed in an exhaustive set of locations in order to represent the *nomadic attacker* and also to evaluate the spatial footprint of an *opportunistic stationary attacker*. In particular, we conduct five different testbed experiments that show:

(1) the feasibility of eavesdropping on reflections,

(2) reflector location optimization,

(3) freedom-of-space from scattering,

(4) focused reflections for improved signal strength, and

(5) reflections on common communication devices.

The transmitter is optimally steered towards the receiver.

In all these experiments, the transmitter is mounted on a rotation head to steer the signal in different directions. This ensures that the signal beam is in at least one orientation optimally aligned to the receiver. The receiving antenna is placed at various locations as seen in Figure 8.4. Depending on the experiment, we evaluate different distances of Bob in direct line of the transmission and change the offset angles of Eve with a constant distance to the reflector. Eve's initial orientation is perpendicular to the transmission direction. In every experiment, we conduct 100 iterations and state the 95 % confidence intervals for measurements of the BER and signal strength. The signal strength is expressed in dB and normalized to the noise floor.

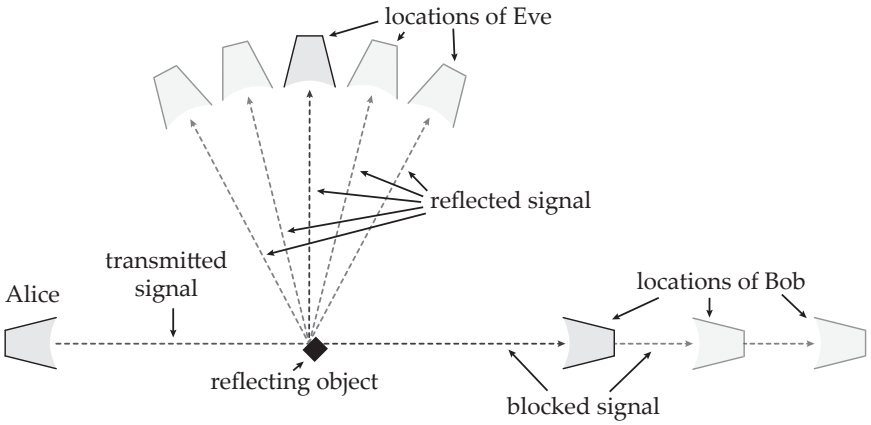

Figure 8.4: Experimental setup showing the communication parties and location variations analyzed throughout our evaluation.

8.2.1 *Baseline and Setup*

Measurements without any reflector are used as a baseline.

Before starting the individual experiments, we analyze the baseline and measure the performance of Bob and Eve without any objects in the beam. These measurements produce s_{opt}, which is needed for the effective reflectivity and blockage calculations via Equation 8.4 and Equation 8.5 in the subsequent experiments. The antennas for Alice and Bob are distributed at a fixed distance of 2 m away from each other. Eve resides 1 m away from the direct link and is oriented perpendicular to the transmission direction. In optimal transmission without any objects in the beam, Bob's signal strength peaks at perfect alignment of 0° with 23.9 dB. Since we are using an antenna with a beamwidth of 7°, the signal strength drops by about 3 dB at an offset angle of 3.5°. Due to the narrow beamwidth, no signal is measured at Eve which is unsurprising as she resides outside the beam.

8.2.2 *Feasibility of Eavesdropping*

We analyze the blockage and reflectivity of various objects of different sizes.

In our first experiment, we evaluate the impact of an object manipulator by placing arbitrary objects in the signal beam to cause reflections towards Eve's antenna. Although it is well-known that mm-waves reflect off metallic reflectors and many other materials, we investigate the critical point between reflector size and material and the effective blockage and reflectivity. An optimal reflector maximizes reflectivity for eavesdropping but simultaneously minimizes blockage to avoid being detected. In particular, we use a metal block and small metal sheets of different sizes. Besides, reflections of a block of wood and acrylic glass are evaluated as well. Given the narrow beamwidth, all objects (except the acrylic glass) are significantly smaller than the beam, which has a width of 12 cm at 1 m distance.

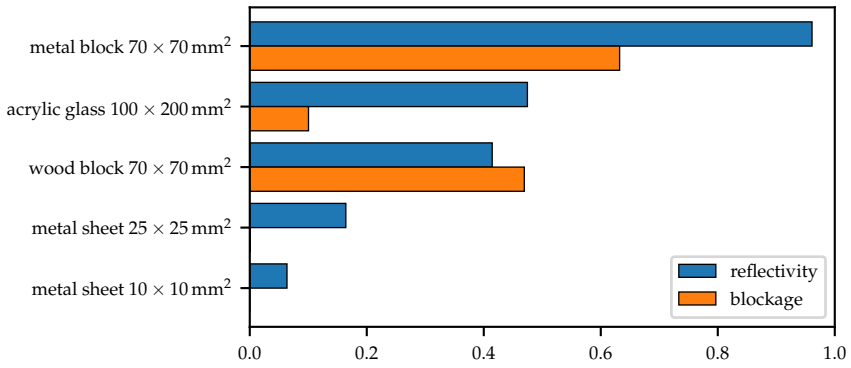

Figure 8.5: Reflectivity and blockage of different objects in the beam.

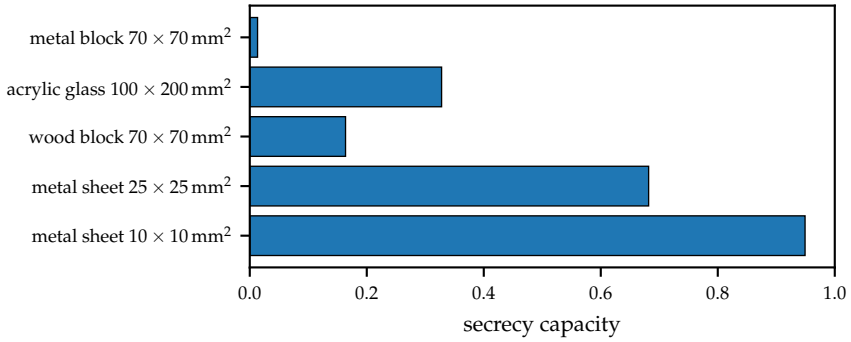

Figure 8.6: Achievable secrecy capacity with different objects in the beam.

Most objects with high reflectivity block the intended signal.

The larger metal sheet and the wood block cause high attenuation, while all other objects only marginally block the signal. Apparently, metal causes strong reflections, but their strength highly depends on the reflector's size. Only the smallest $10 \times 10\,\mathrm{mm}^2$ metal sheet results in a low received signal strength at Eve. Even the wood block and acrylic glass, both with plain surfaces, cause considerable reflections. The effective reflectivity r (see Equation 8.4) and blockage b (see Equation 8.5) of each object is shown in Figure 8.5. With an effective reflectivity of 96 %, as for the metal block, the signal quality at Eve is nearly as good as at Bob. Unfortunately, this high reflectivity comes with high blockage of 63 % and might be easily detectable. While the smaller metal sheets block less than 1 % of the signal, they still provide a notable effective reflectivity of up to 16 %. With a reflectivity of 47 % and a blockage of 10 %, the acrylic glass is a good trade-off between both characteristics.

All reflecting objects decrease the secrecy rate.

All analyzed objects in our experiment lower the secrecy capacity as shown in Figure 8.6. Featuring a small size, the $25 \times 25\,\mathrm{mm}^2$ metal sheet decreases the secrecy capacity by 32 %. The metal block of size $70 \times 70\,\mathrm{mm}^2$ diminishes the secrecy capacity to 1 %, which means that Eve's reception becomes nearly as good as Bob's.

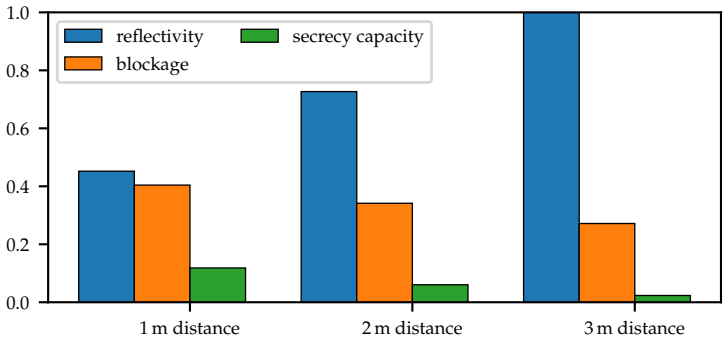

Figure 8.7: Effective reflectivity, blockage, and secrecy capacity for Alice and Eve at fixed positions and Bob with varying distance. An increasing distance of Bob leads to higher reflectivity and lower blockage and secrecy capacity.

8.2.3 *Reflector Location Optimization*

By moving the receiver, we vary the relative eavesdropping distance.

While the previous section reveals the feasibility of eavesdropping on reflections, the question on where to place the reflector is still unanswered. In our second experiment, we aim to determine the optimal reflecting location that diminishes the secrecy capacity of Alice's transmission. The object manipulator optimizes its performance by varying the relative object location within the beam. To cause the reflections, we use a medium size metal sheet of size $70 \times 70\,\mathrm{mm}^2$. While the reflector and the eavesdropper are located at fixed locations separated by 1 m, the receiving antenna for Bob is placed at distances of 1 m, 2 m, and 3 m away from the reflecting object. This setup results in a communication distance of 2 m, 3 m, and 4 m between Alice and Bob, respectively. As in the previous experiment, the eavesdropping distance is kept constant at 2 m. By varying only the distance of Bob and not that of Eve, we ensure to maintain the same reflections in all evaluation steps. However, varying this distance affects the optimal received signal strength s_{opt} at Bob in direct transmission without blockage. For larger distances between Alice and Bob, the relative eavesdropping distance decreases and perform differently.

Small objects are affected by diffraction.

Only considering shadowing effects, we would expect Bob's signal strength s_{Bob} to decrease similar to s_{opt} and the effective blockage to be constant. In the line-of-sight setting, Bob always resides in the shadow region. However, as seen in Figure 8.7, the blockage decreases with Bob's distance. This effect must be caused by diffraction, which still occurs in mm-waves around small obstacles [Jac+12; Kle+12]. Diffraction, by implication, assists the eavesdropper in remaining invisible by lowering the effective blockage. Furthermore, the effective reflectivity increases with Bob's distance. The reflected signal does not change, but, since s_{opt} decreases, the relation between Eve's and Bob's signal strength becomes greater.

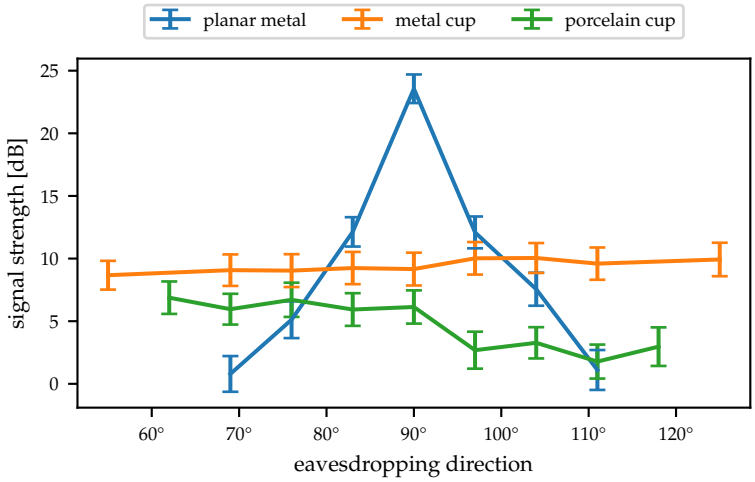

Figure 8.8: Effective reflectivity for different eavesdropper locations.

Reflectors are ideally placed close to Alice.

Since the effective reflectivity increases more than the blockage decreases, the secrecy capacity decreases with Bob's distance. In this particular experiment, we observe the secrecy capacity to decrease by 81 % when moving Bob's antenna from 1 m to 3 m away. The optimal position for placing the reflecting object is, apparently, close to Alice. By placing an object there, Eve not only increases the reflected signal strength but also can be less afraid of blocking too much of the beam.

8.2.4 *Freedom-of-Space from Scattering*

Scattering compensates small misalignments of reflectors.

Finding the optimal reflector orientation is a challenging task. Small misalignments can have significant impacts on the signal quality at Eve. This experiment analyzes the freedom-of-space that a nomadic attacker has in choosing its location for a fixed reflector. In addition to that, we investigate whether scattering helps the opportunistic stationary attacker. The more locations from which Eve can successfully eavesdrop, the more likely an opportunistic stationary attacker can be successful despite not moving itself nor altering the environment. We set up this experiment as the first one (Section 8.2.2), but instead of placing Eve only perpendicular to the beam directions, we move her on a circle around the reflecting object. Figure 8.8 shows the reflected signal strength with different reflectors over varying eavesdropping angles. With the metal block of $70 \times 70\,\mathrm{mm}^2$ placed as the reflector, the signal strength strongly decreases when moving away from the optimal position of 90°. At an offset of 7°, the signal strength already drops by around 10 dB. Round objects like a porcelain cup and a metal shielded cup reflect much weaker than the planar metal sheet. However, they feature a nearly constant signal strength over a wide range of eavesdropping positions; they scatter the signal to multiple directions.

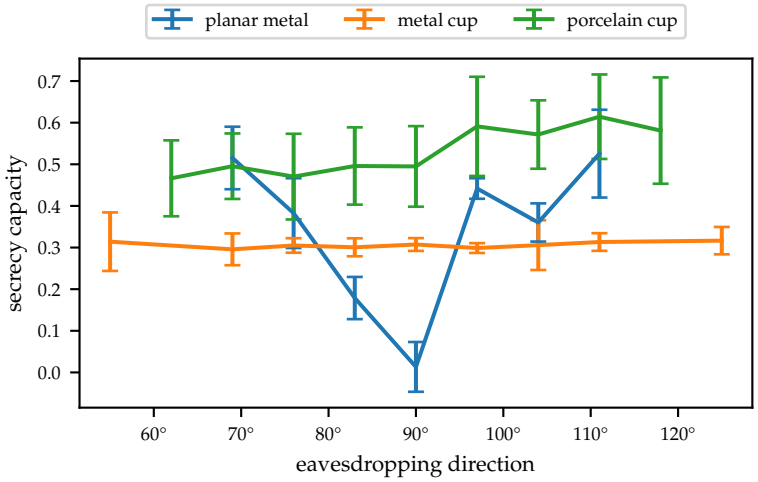

Figure 8.9: Secrecy capacity in dependency of the eavesdropper's location.

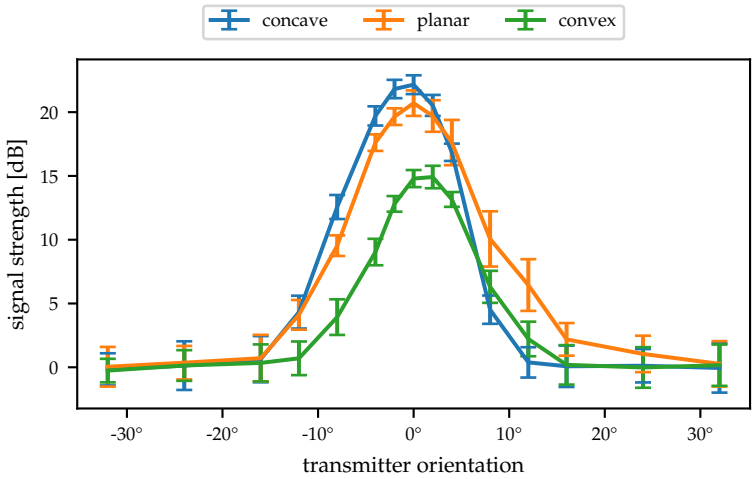

Figure 8.10: Signal strength of the eavesdropped signal with bent reflectors.

Attackers may benefit from scattering on round objects.

The secrecy capacities shown in Figure 8.9 support our finding that attackers residing not in the optimal reflection direction might improve their performance with round reflectors. In this particular scenario, we observe that the round metal reflector provides better reflections than the planar one at an offset angle of approximately 10°. When only the coarse beam direction is known, an attacker benefits from these scattering effects. They become particularly important for the opportunistic stationary attacker that is limited in mobility.

8.2.5 *Reflection Focusing*

To analyze if reflectors can focus the signal towards the attacker, we use a planar, concave, and convex metal sheet of size $100 \times 200\,\mathrm{mm}^2$ as reflecting object. In contrast to the previous experiments, the reflectors

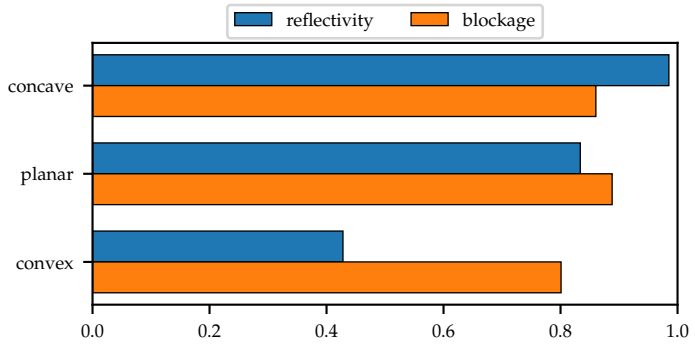

Figure 8.11: Blockage and reflectivity of objects with different bendings.

Concave surfaces focus the reflections.

are larger, because they are easier to bend to the correct curvature. Figure 8.10 shows the signal strength at Eve with these objects over the transmitter's azimuth angle. With the insights from the previous experiment, it is obvious that the convex reflector provides a low signal strength. However, slight deformations of only a few millimeters are sufficient to optimally focus the beam for our evaluation setup where Alice and Eve are both 1 m away from the reflector. Compared to the planar reflector, we achieve an increase in the signal strength of 1.5 dB. Since the size of the object does not vary, the blockage remains constant. Still, the effective reflectivity varies as depicted in Figure 8.11. Concave reflectors bundle the reflected signal towards a focal point at which very high signal strengths are achievable. Object manipulators, as well as nomadic attackers, can exploit this to obtain better eavesdropping performance. However, they only benefit from focusing the signal beam when residing precisely at the focal point. For that reason, this approach is unfavorable for the opportunistic stationary attacker. Even small misalignments in position and orientation may lead to massive losses in the signal strength.

Finding the focal point is challenging.

8.2.6 *Reflections on Commodity Devices*

Commodity devices cause reflections on their surfaces.

Regular communication devices—which will be equipped with mm-wave hardware soon—cause reflections towards potential eavesdroppers. These devices are typically made of materials with high reflectivity. Both nomadic attackers and opportunistic stationary attackers can take advantage of such reflections despite not being able to place their own reflectors. To analyze how strong these reflections are in practice, we place several communication devices in our evaluation setup. In particular, we use an iPhone 6, a laptop, and a Mac mini in different orientations. Similar to the previous experiments, they are deployed at a distance of 1 m away from the transmitter and receive the reflected signal at 1 m distance perpendicular to the transmission direction. As shown in Figure 8.12, the signal that is reflected at a

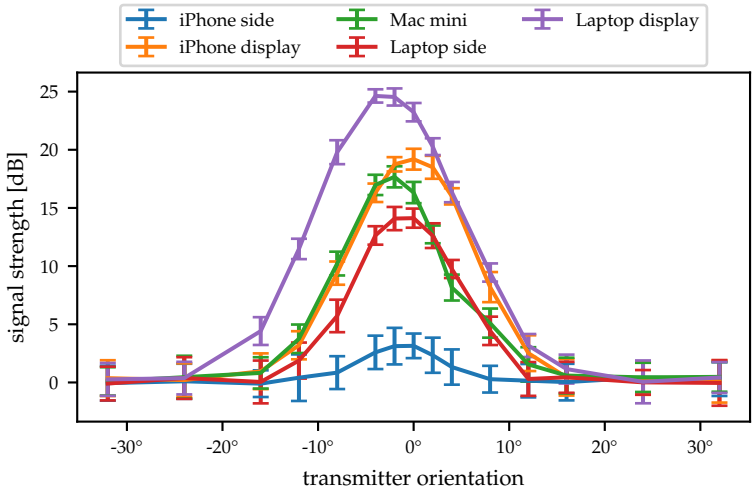

Figure 8.12: Reflected signal strength with common communication devices.

laptop display achieves the highest strength of 24 dB. This is as high as we observed in direct transmission without reflectors. The side of the laptop, the iPhone display, and the Mac Mini also achieve reflected signal strengths between 14 and 19 dB. Only the side of an iPhone is too small and curved to cause significant reflections towards the eavesdropping antenna; the signal strength remains below 4 dB. These results imply that reflections can be caused not only by specifically placed reflectors but also by inconspicuous consumer devices. In a typical communication scenario, the signals may reflect on the surface of the receiver and enable a nomadic attacker as well as an opportunistic stationary attacker to eavesdrop without changing the environment.

Eavesdropping on environmental reflections is feasible.

Our practical evaluation with testbed experiments shows that eavesdropping on reflected mm-wave transmissions is possible and enables attackers to reside outside the designated signal beam. Further, it shows that varying the reflector position affects the eavesdropping performance and that small-scale diffraction helps to be undetectable. By bending objects, we reveal that concave surfaces scatter the signal and increase the freedom-of-space at the cost of signal strength. In contrast, convex surfaces focus the beam and increase the signal strength at certain positions. Finally, we demonstrate that not only additional reflectors but also the intended recipient's devices can cause significant reflections of mm-waves. A summary of all evaluated objects along with their measurement results is provided in Table 8.1.

8.3 DISCUSSION AND SUMMARY

Although mm-wave communication systems are often addressed as intrinsically secure against eavesdropping from afar, our practical

Analyzed Object	Size	Section	Effective Reflectivity	Effective Blockage	Secrecy Capacity	Impacts
metal block	70x70 mm^2	8.2.2, 8.2.4	0.96	0.63	0.01	very good reflections but also high blockage
wood block	70x70 mm^2	8.2.2	0.41	0.46	0.16	medium reflections and high blockage
acrylic glass	100x200 mm^2	8.2.2	0.47	0.10	0.33	good tradeoff between reflections and blockage
metal sheet	25x25 mm^2	8.2.2	0.16	0.00	0.68	no blockage but still some reflections
metal sheet	10x10 mm^2	8.2.2	0.06	0.00	0.95	too small for significant reflections
metal sheet	70x70 mm^2	8.2.3	0.45	0.40	0.12	medium reflections and significant blockage
metal cup	cup size	8.2.4	0.20	0.62	0.30	scattering to all directions
porcelain cup	cup size	8.2.4	0.14	0.48	0.46	poor reflections but still scatters
metal sheet	100x200 mm^2	8.2.5	0.83	0.80	0.00	good reflections but very high blockage
convex metal sheet	100x200 mm^2	8.2.5	0.42	0.89	0.00	good scattering but very high blockage
concave metal sheet	100x200 mm^2	8.2.5	0.98	0.86	0.00	focuses the reflections but very high blockage
iPhone display	n/a	8.2.6	0.58	—	—	good reflections
iPhone side	n/a	8.2.6	0.09	—	—	poor reflections due to non-optimal surface
laptop display	n/a	8.2.6	1.00	—	—	perfect reflections
laptop side	n/a	8.2.6	0.32	—	—	low reflections on non-solid surface
Mac mini	n/a	8.2.6	0.49	—	—	acceptable reflections

Table 8.1: Summary of all evaluated objects with reflection and blockage characteristics. Reflections are considered towards an eavesdropping antenna that is located perpendicular to the intended transmission direction.

work demonstrates that this is not the case. We introduce three distinct attacker models and evaluate their performance in practical experiments. Our findings prove that highly directional mm-wave transmissions are not intrinsically secure against attackers outside the beam. Despite evaluating only low distances in our laboratory setup, we show that environmental objects cause strong reflections outside of the expected coverage.

Manipulating objects in the environment can direct reflections towards an attacker.

Attackers of the object manipulator class can tamper with the environment by placing objects in the signal beam to cause reflections towards a fixed eavesdropping antenna. Using this method, they can achieve good reflections with low blockage. Our experiments show that it is possible to place objects in such a way that reflections facilitate eavesdropping: objects as small as $25 \times 25\,\mathrm{mm}^2$ decrease the secrecy capacity by 32 % without any effective blockage. For a fixed object size, the blockage remains nearly independent from the object's shape and orientation, but the effective reflectivity towards a certain eavesdropping position varies. Sophisticated object structures with concave surfaces can focus the signal beam towards a certain eavesdropper. For example, bending a metal reflector leads to a received signal strength at the eavesdropper as high as that at the intended receiver. An attacker with physical access to the environment can achieve good eavesdropping performance.

Nomadic attackers exploit existing reflections in the environment.

Nomadic attackers fall into a significantly weaker attacker class and cannot actively manipulate the environment. Yet, they can also achieve a good eavesdropping performance by exploiting reflections that occur in the environment and at the surface of the communication devices itself. The nomadic eavesdropper seeks to find a location at which strong reflections are receivable with a minimal position change to avoid detection. We demonstrate that device-incident reflections of common communication devices, such as a laptop or a mobile phone, are sufficient to enable eavesdropping: essentially the recipient becomes the traitor to itself. In such a scenario, the attacker only has to find a good location, point its antenna toward the receiver, and then eavesdrops on the reflected signals that bounce off of the device.

Completely passive attackers need to compensate poor signals.

The opportunistic stationary attacker is less powerful; it must eavesdrop from a given location without being able to manipulate objects in the environment. Our results show that round objects disperse the signal into multiple directions, thus facilitating attacks from opportunistic attackers. In most cases, however, we found the signal to be too weak for effective eavesdropping. If the attacker is off by only a few degrees from the optimal angle of reflection, the signal strength strongly impairs. Consequently, this attack likely needs to compensate for poor signal quality. Besides using expensive antenna apertures, this attack might be launched by multiple cooperative eavesdroppers in a distributed attack with several eavesdropping elements planted throughout the environment.

9 BEAM STEALING

Beam stealing tampers with the beam selection.

Compensating for high attenuation, mm-wave communications apply beam steering. This provides strong directionality and allows to transmit signals only in a particularly preferred direction. As revealed experimentally in the last chapter, this directionality leaves a false sense of security in case of environmental reflections. However, also the underlying training protocols such as the sector level sweep in IEEE 802.11ad [Nit+14; IEE14] are unprotected against malicious behavior. Finding the optimal antenna steering typically takes place before any secure channel is established. The sector level sweep sends probing frames over all available beams. Receivers report the one that is received with the highest signal strength such that its beam can be selected for further transmissions. Since this process is neither authenticated nor encrypted, devices cannot distinguish whether the feedback comes from the expected peer or an attacker. An Attacker might tamper with the beam training process for his benefit. He could inject forged feedback from distant locations to steer the beam towards his direction as illustrated in Figure 9.1. While this slightly impairs the signal quality at the legitimate devices, it boosts that at the attacker. Thus, the attacker increases his chances for eavesdropping and may act as a Man-in-the-Middle (MITM) to relay and intercept packets between the legitimate devices. Launching such an attack could have a severe impact, as it cannot be prevented with conventional security mechanisms that operate on higher layers.

In this chapter, we practically demonstrate the feasibility and impact of an mm-wave beam stealing attack and propose an authentication scheme to defend against this threat. Section 9.1 specifies an at-

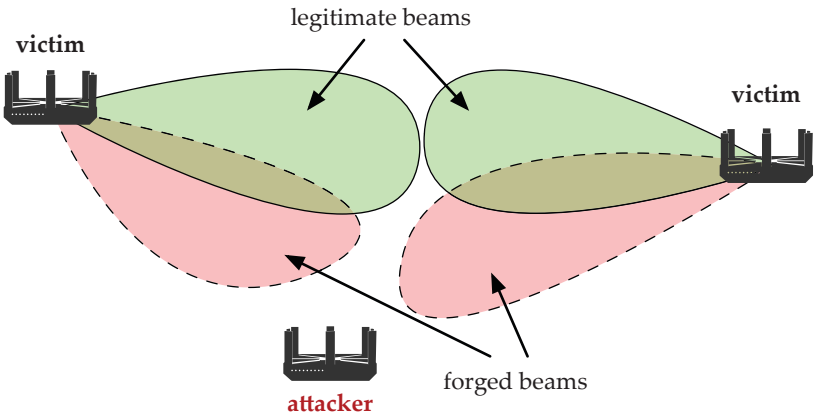

Figure 9.1: Injecting forged feedback into the sector sweep operation turns devices to select sectors that better serve the attacker.

tacker model that targets the sector level sweep of IEEE 802.11ad and feedbacks forged sector information on behalf of other devices. In Section 9.2, we implement a proof-of-concept of such an attack on off-the-shelf devices and experimentally evaluate our findings. Additionally, we propose an authentication scheme, to protect against forged sector sweep feedback in Section 9.3. Finally, Section 9.4 discusses and summarizes our findings.

9.1 ATTACK METHOD

The sector level sweep in IEEE 802.11ad networks is prone to attacks that inject forged feedback to tamper with the sector selection on remote devices. The following describes the vulnerabilities in the sector level sweep, the attack method and possible scenarios such as active eavesdropping and the MITM.

9.1.1 *Vulnerabilities in the Sector Level Sweep*

Forged feedback selects arbitrary sectors.

Beam training mechanisms, such as the sector level sweep algorithm in IEEE 802.11ad, probe all sectors in a predefined codebook. In doing so, they find the sector that provides the highest signal strength at a receiver. This protocol is performed mutually between two devices and works as described in Section 2.2.3. The initiating device transmits sector sweep frames sequentially on all available sectors with an identifier encoded in the payload. During the transmission, the second device selects a quasi-omnidirectional sector for receiving and determines the received signal strength for all the frames it overhears. It then chooses the sector for which it receives the highest signal strength and reports back the respective identifier. Finally, the initiating device selects this sector for further transmissions. To cope with changing channel conditions, all devices periodically repeat the sector sweep. Since the transmitted frames are unprotected IEEE 802.11 control frames, attackers can inject forged payload and report arbitrary sector to be selected by other devices.

9.1.2 *Forged Sector Sweep Feedback*

Attackers redirect the beams in arbitrary directions.

Victims of a beam stealing attack are two legitimate devices, an Access Point (AP) and a station, that are communicating via IEEE 802.11ad and perform the sector level sweep to determine aligned transmit sectors. The attacker's goal is to change the sectors both devices select and force them to redirect their beams. Steering the victims' beams towards his location, allows him to receive frames that otherwise could not be decoded. He can freely take any position in the surrounding but is unlikely to reside directly between the devices on the line-of-sight. Throughout our investigation, we focus on typical indoor

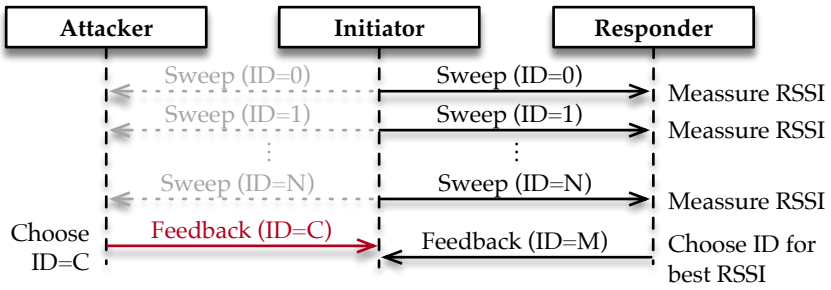

Figure 9.2: Sequential overview of the sector level sweep with forged feedback.

environments. As mm-waves do not penetrate walls or obstacles, the attacker stays in distances of a few meters. Besides, we assume that he has full control over the capabilities of its IEEE 802.11ad interface, can listen to all frames received at its antenna, and inject custom ones. He has neither access to the legitimate device's hardware nor the software and does not cooperate with others.

Forged sector sweep feedback turns the beams into arbitrary directions.

To tamper with the sector selection, the attacker forges the feedback of the sector level sweep. After an initiator sends out sector sweep frames in all sectors, it typically awaits the feedback from the responder. The attacker listens to the sector sweep frames itself, makes an own selection and sends back his feedback to the initiator on behalf of the legitimate responder as shown in Figure 9.2. The attacker just needs to ensure that his forged feedback gets accepted, for example by achieving a faster reply time or jamming the legitimate frames. Whenever the initiating device completes the sweep, it chooses those sectors the attacker has forged. This allows an attacker to perform active eavesdropping or launch a MITM as described in the following.

9.1.3 *Active Eavesdropping*

Directional transmissions make eavesdropping more challenging.

Eavesdropping on wireless communication networks is typically considered to be a passive attack in which an attacker captures and processes all the frames it could receive. While this might have a high success rate in conventional wireless systems with omnidirectional radiation characteristics, it becomes challenging and location dependent in mm-wave networks. Whether an eavesdropper could receive frames depends not only on its distance to the transmitter but also on the direction, and side lobe gains of the selected transmit beam. For good eavesdropping results, an attacker should ideally be 1) located on the line-of-sight between the communicating devices, 2) in directions with significant side lobes gains, or 3) redirect the signal by exploiting environmental reflections (Chapter 8). However, achieving this is seldom possible. To relax the location requirements and facilitate eavesdropping from any location, we forge the sector

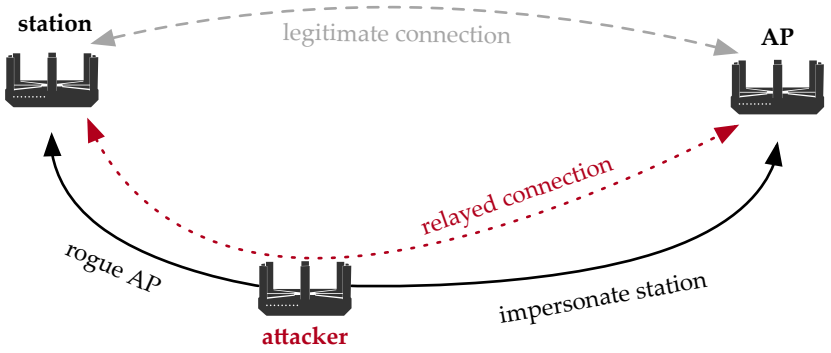

Figure 9.3: Man-in-the-Middle attack scenario in which the attacker impersonates the station and launches a rogue AP.

sweep feedback. As this approach requires injecting custom frames, we refer to it as active eavesdropping.

Forged sector sweep feedback facilitates eavesdropping with distant locations.

Eavesdropping on sector sweep frames is rather simple as they are transmitted sequentially in all sectors and modulated with low Modulation and Coding Schemes (MCSs). At least those that are transmitted in the direction of the attacker should be receivable. Data frames are more challenging to decode. They are modulated with higher MCS and transmitted only in the chosen sector, which is unlikely to provide sufficient gain towards a distant eavesdropper. To listen on mutual communication, an active eavesdropper forges the sector sweep feedback of legitimate devices. He does this such that he can decode the transmitted frames but still allows for a high data-rate between the legitimate devices.

9.1.4 *Man-in-the-Middle*

A MITM *relays packets between two devices.*

To launch a MITM attack on directional mm-wave networks and relay packets between legitimate devices, an attacker needs to impersonate the station and connect to the AP. The station must be persuaded to connect to the attacker instead of legitimate AP. A typical setup of such a scenario is illustrated in Figure 9.3. The attacker impersonates the station and establishes a connection to the AP. Hence, the AP believes to be connected to the station, when in fact it is connected to the attacker. Spoofing the station's MAC and IP address and using a higher signal power allows the attacker to trick the AP. Forging its beam selections provides the required signal power towards his location. While the legitimate link gets degraded, a reliable connection to the attacker becomes possible. The rogue AP pretends to be the legitimate one. This scenario, in which the attacker aims at catching the connection from the legitimate station, is also known as 'evil twin' attack. Both APs compete for the station's connection. In contrast to station impersonation, the rogue AP cannot initiate the association itself. It periodically announces its presence in beacon frames and

waits for the stations to connect. A MITM attack combines the station impersonation and rogue AP, thus receives frames from the legitimate station and forwards them towards the AP and vice versa. Simultaneously, traversing packets can be forged, dropped, and analyzed. Since both victims believe in communicating with their counterparts, the MITM is transparent and hard to detect. Next, we provide details on our proof-of-concept implementation on off-the-shelf devices.

It impersonates the station and launches a rogue AP.

9.2 PRACTICAL INVESTIGATION

This section provides the results of launching the beam stealing attack in practical IEEE 802.11ad networks. We provide our proof-of-concept implementation with off-the-shelf devices and experimentally evaluate our the performance of forging the sector sweep feedback. After describing our setup, we investigate attack scenarios of active eavesdropping, station impersonation, and rogue access points. The combination of them constitutes a MITM attack that relays all traffic between legitimate devices through a malicious link. Additionally, we discuss possible countermeasures and examine detection schemes.

We investigate the impact on off-the-shelf devices.

9.2.1 *Experiment Setup*

Our proof-of-concept implementation of the MITM attack on directional mm-wave networks is built on off-the-shelf IEEE 802.11ad devices. Talon AD7200 routers from our testbed experimentation platform (see Chapter 4) constitute the attacker as well as the legitimate devices. The OpenWrt operating system with the *wil6210* driver running on the devices allows us setting up the connection between the devices. Unfortunately, the closed-source firmware for the mm-wave module does not allow to inject arbitrary frames without being connected to an AP. To overcome this limitation, we indirectly inject our forged sector sweep feedback via the legitimate devices. With our binary firmware patches, which allow integrating custom beam training protocols, devices report custom sectors in the sector sweep feedback. This forces the beam of the transmitter to steer towards an arbitrary direction on behalf of the attacker.

Using our testbed platform, we select arbitrary sectors.

Sending custom association and disassociation requests helps the attacker to capture the connection between the legitimate devices. This is achieved by sending custom management frames with the debug feature of the *wil6210* driver. As the format of IEEE 802.11 association and disassociation frames is well documented, we captured default parameters from the legitimate operation and adjust the source, destination, and Service Set Identification (SSID) fields, respectively. Doing so allows disassociating the station and associating the attacker.

Re-associations help the attacker to capture the connection.

For the evaluation of the feasibility of our attack, we set up two Talon AD7200 devices for the legitimate station and AP and up to

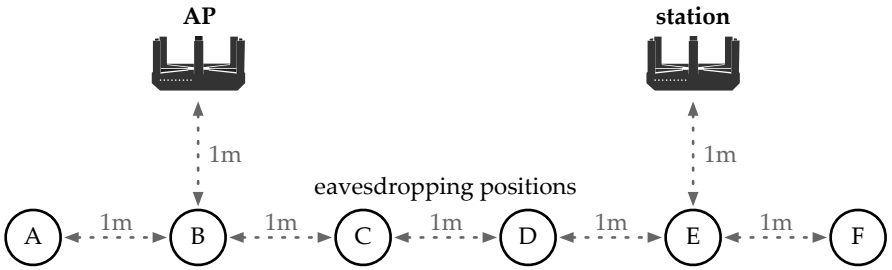

Figure 9.4: Evaluation scenario for active and passive eavesdropping with an attacker placed at six different locations.

We set up all devices in a single room.

two additional devices for the attacker. They are interconnected via a 2.4 GHz Wi-Fi network for management and control. The two legitimate devices communicate via the 60 GHz link, while the attacker aims to intercept that connection. All respective network interfaces are configured with static addresses. To control the attack and automate our experiments, we further utilize our testbed automation system. Custom scripts start a rogue AP and fetch the properties from legitimate devices for impersonation. We use sockets to generate custom traffic using the User Datagram Protocol (UDP) and analyze the frames traversing the attacker in relaying or eavesdropping with the Python packet manipulation tool scapy[1]. Throughout our experiments, we consider a typical indoor scenario. As walls and obstacles absorb mm-wave signals, we assume the attacker to be located in the same room and only a few meters away from the legitimate devices. The attacker's performance is evaluated in four steps. First, we experimentally study the outcome of injecting forged sector sweep feedback in an eavesdropping scenario. Second, we investigate the feasibility of station impersonation by setting up a fake station. Third, we analyze the success of deploying a rogue AP. Finally, we combine the previous scenarios in a MITM attack and evaluate how well the attacker can reroute data packets. Our results are described in the following sections.

9.2.2 *Active Eavesdropping*

The eavesdropper takes six different locations.

In our first experiment, we distribute the AP and the station at a distance of 3 m, which is a typical use-case in IEEE 802.11ad networks [Mal+10a]. To assess the active and passive eavesdropping performance, we place a third device as eavesdropper at six different positions parallel to the communication direction as illustrated in Figure 9.4. Given the irregularity of the beam-pattern, exploring all possible locations and directions do not provide meaningful results. Thus, we focus on this simple proof-of-concept scenario with an eavesdropper at different angles and distances to the transmitting

1 http://secdev.org/projects/scapy/

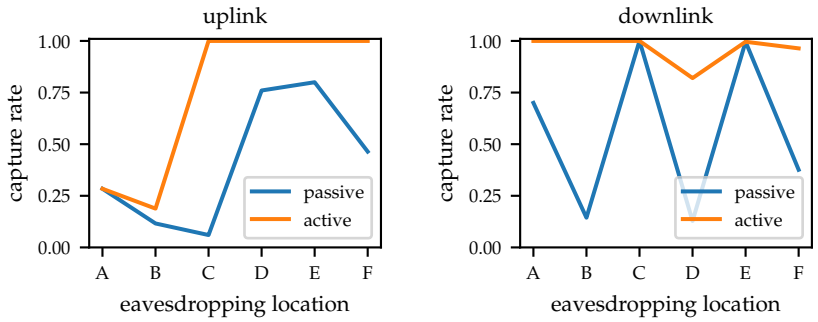

Figure 9.5: Capture rates of active and passive eavesdropping at six different locations.

devices. We configure the eavesdropper in monitor mode and transmit 50 UDP packets with a length of 8192 Byte from the station to the AP. The AP responds to each incoming packet with the same payload. Overhearing the same packets, the eavesdropper records and validates received requests and responses for active and passive operation. In passive eavesdropping, the sectors are selected using the default operation. For the active eavesdropper, we forge the sector selection such that the beam directs towards the eavesdropper's location. All of our experiments are repeated five times.

The active eavesdropper is less location dependent.

Our results reveal that the passive eavesdropping success is location dependent and varies enormously. As shown in Figure 9.5, we observe good capture rates at certain positions where the attacker decodes most frames correctly. The eavesdropping location might already fall into one of the sidelobes of the transmitting antenna beam. In other locations, the packet rate goes below 20 %, which can be explained with the irregular antenna pattern shapes. With active eavesdropping, by steering the signal directly towards the attacker, we compensate these gaps and become less dependent on the eavesdropper's location. In all evaluated locations, we increase the capture rate of requests and responses. The achievable capture rates vary with the asymmetric beam-patterns deployed on those devices. Active eavesdropping increases the packet rate by 38 % on average.

Redirecting the beam slightly decreases the throughput.

Steering the beam away from the intended direction impairs the link quality between the devices. To address this impact, we measure the achievable throughput on this link under passive and active eavesdropping as well. Our results in Figure 9.6 exhibit a constant throughput with the default beam alignment in passive eavesdropping. The active attack that forged the sector selection lowers the signal strength at the intended receiver, thus impairs the performance. Depending on the eavesdropper's location the impact varies. While active eavesdropping decreases the average performance between the legitimate devices by 15 %, it still achieves a remarkable throughput of about 1.4 Gbps.

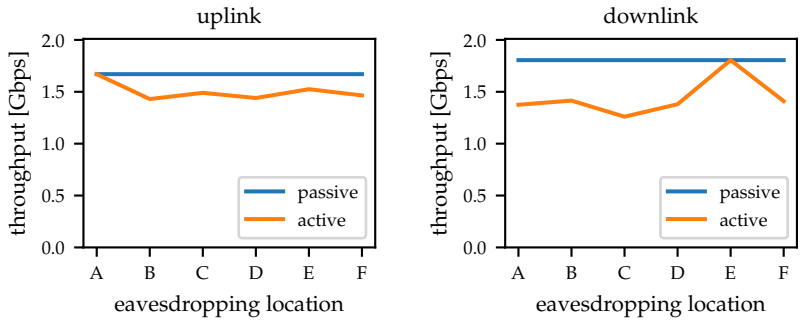

Figure 9.6: Achievable throughput with active and passive eavesdropping at six different locations.

9.2.3 *Station Impersonation*

A fake station impersonates the legitimate ones.

For the impersonation attack, a fake station imitates the configuration of the legitimate one. To achieve this, it duplicates the MAC and IP address. In the following experiment, we consider two different scenarios, one with a higher distance and one with a lower distance of the fake station to the AP. All three devices, the attacker, the station, and the AP, are placed on a line, whereas the AP is located in the middle. This setup constitutes the most challenging case, as the attacker needs to force the AP to steer its signal in the complete opposite direction and cannot reuse the existing beam. First, we deploy the station at a distance of 1 m and the attacker 3 m away from the AP. In the second scenario, we swap the distances and set up the station at 3 m while moving the attacker closer. By swapping the distances, we enforce that either the station or the attacker becomes favored and receives signals with higher strength than that of the other one. Our setup for both scenarios is depicted in Figure 9.7. During the experiment that is repeated 20 times, the station pings the AP. After starting the attack, we observe the connection for 10 seconds. Throughout this interval, the AP logs the received signal strength and the chosen sector. The stations continuously check the connectivity and records all control frames to reveal which one is associated.

We record the chosen sectors during transmissions.

From the frames that we recorded throughout the experiment, we determine the sectors that the AP choose in the sector level sweep and whether they serve the station or attacker. For all observed sector sweeps, we count how often the sector towards the fake station is chosen and express this as success-rate of the sector sweep competition. A sector level sweep is triggered after each association and repeated continuously throughout the communication. It can be initiated by all devices and is performed mutually over multiple rounds. However, for less stable connections the sweeps are performed more often. By manually sending association requests, we also trigger a new sweep, which offers another chance for the fake station to set its forged

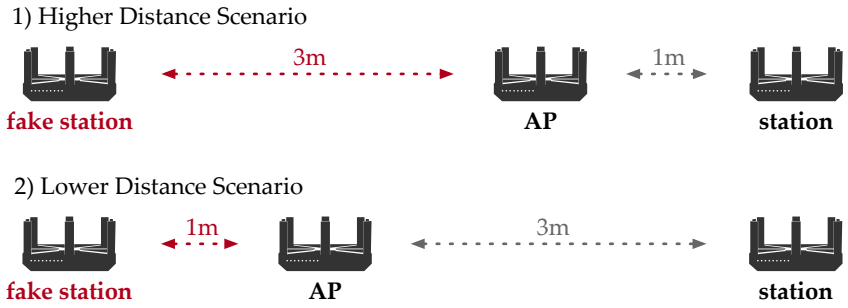

Figure 9.7: Placement of devices for the station impersonation attack. We evaluate two scenarios with the attacker in higher and lower distance to the AP than the legitimate station.

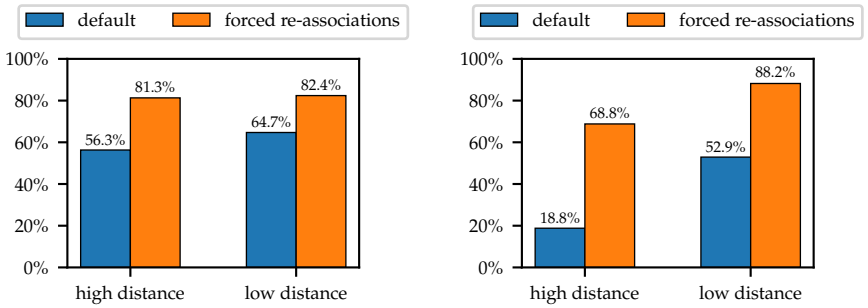

(a) success in sector selection

(b) network connectivity

Figure 9.8: Success in sector selection and network connectivity rate of the fake station for the default behavior and by forcing re-associations.

sector. If a station invokes this process, its outcome depends on the last round of the sector sweep which is initiated by the AP. While both stations compete to respond, only one gets accepted. In most cases both stations start the association process simultaneously, the AP coordinates the channel access such that their sweeps are executed alternately. The station which initiates the last sweep is accepted and used until the next sweep starts. Hence, the signal strength has only minor influence, but sending association attempts, whenever indicated, increases the chances of setting the own sector. The success rate without re-associations, as shown in Figure 9.8a, achieves 56.3 % for the higher and 64.7 % for the lower distance. Forcing re-associations from the fake station increases the success rate on average by about 18.0 % to 81.3 % and 82.4 %, respectively.

The attacker must acquire a connection to the AP.

The performance of station impersonation depends on the success of the fake station to acquire an association with the AP. Once the attacker attempts to associate, the AP reacts with a disassociation as it believes to be already associated. Then, both stations, the legitimate and the fake station, compete with each other to associate again. Whether the attack succeeds depends on the relative signal strength between both

stations. With the fake station at the higher distance, we observe an average Received Signal Strength Indicator (RSSI) at the fake station that is 12.1 dB lower than that at the legitimate station. When the fake station is closer to the AP, its RSSI increases by 13.4 dB while that at the legitimate station decreases by 12.5 dB. We assume a stable connection if devices receive at least ten ping replies continuously. Our first scenario, where the signal strength of the fake station is low, we accomplish a successful connection in only 18.8 % of our experiments without invoking re-associations. In the second scenario, we observe that higher signal strengths are prioritized, such that the success rises to 52.9 %, as shown in Figure 9.8b. For winning the association competition, the attacker manually injects association frames whenever necessary. Forcing the devices to run another round of the association process increases the chances to obtain a stable connection. We achieve a connectivity of 68.8 % at the higher distance and 88.2 % at the lower distance. This results in an average gain of 42 %. Hence, an attacker must ensure a continuous association and a high signal strength simultaneously to impersonate a station.

9.2.4 *Rogue Access Points*

The rogue AP *announces its presence and waits for stations to connect.*

The rogue AP cannot invoke the association itself but relies on the station to initiate the connection. Therefore, the duration of the Beacon Interval (BI) at which it advertises its presence by sending beacons takes a crucial role. The attacker sends out beacons alternately to the legitimate AP. Stations in range associate with the AP from which they received the last beacon. When they are already associated, they disassociate first and initiate a new association. This allows the rogue AP to respond and capture the connection. Unfortunately, the alternating beacons lead to many re-associations. Our experiments consider this in relation to the BI in three different scenarios. First, we keep the BI of the legitimate and rogue AP constant at the default setting of 100 ms. Second, we investigate a higher BI for the rogue AP of 200 ms. Last, the rogue AP sends beacons more frequent than the legitimate AP by keeping its BI at 100 ms, whereas that of the legitimate AP increases to 200 ms. As expected, we observe many re-associations in the former two scenarios. Taking a closer look at the packet dumps reveals that the station sends out many disassociation frames. This is an indicator that it gets confused by too many association attempts. If we set different BIs for both APs, the number of disassociation frames decreases. It seems that whenever a beacon is received before it is expected, the station re-initiates the association. Hence, both APs compete, and the one that lastly transmitted a beacon obtains a change to catch the connection until the next BI. In the third scenario, the attackers BI is lower than that of the legitimate AP. Thus, intervals exist that are not interrupted by new beacons from the legitimate AP.

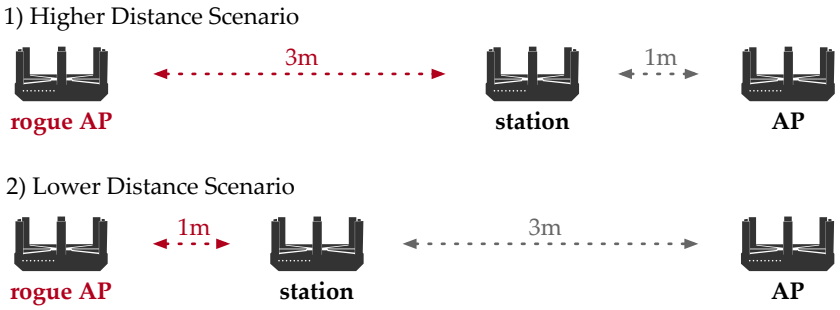

Figure 9.9: Placement of devices for the rogue AP attack. We evaluate two scenarios with the attacker in higher and lower distance to the station than the legitimate AP.

Only in this scenario, we achieve a stable connection for a complete interval. Hence, to successfully launch a rogue AP, the attacker should ensure a BI lower than that of the legitimate AP.

We decrease the BI to achieve a higher connectivity.

To evaluate the success of the rogue AP, we use the decreased BI and set up the experiments similar to the station impersonation ones. We distribute the rogue AP and the legitimate devices as shown in Figure 9.9 on a straight line with the station in the center. First, we keep the rogue AP at a higher distance and place it 3 m away from the legitimate station while the legitimate AP has a distant of 1 m. Second, we switch the distances and move the rogue AP closer to a distance of 1 m and move the legitimate AP to 3 m distance. These two setups allow us to investigate the feasibility of launching a rogue AP with different relative signal strengths between the attacker and the legitimate devices.

We collect all frames to determine the attack performance.

During our experiments, both APs continuously ping the station. Once the rogue AP is associated, all devices continuously record the network state and the selected transmit sectors. Due to association issues, the network connection might fail even though the sectors are selected properly. To address this, we consider both properties individually. While the rogue AP pings the legitimate station, they capture frames on all interfaces. Similar to the station impersonation experiments, we assume a stable network connectivity, if the station receives ten ping replies successively. Moreover, we check the captured frames to reveal the sector that are chosen by the station and which AP it favors. All measurements are repeated 20 times. Our results for the network connectivity and sector success rate are shown in Figure 9.10. With higher distance, the rogue AP achieves a network connectivity of 44 % and a sector success rate of 56 %, meaning that in 44 % of probes, we successfully ping the station for at least 10 seconds. During that time, we choose the optimal sector for the attacker in 56 % of the time. The results for the network connectivity in the lower distance are significantly higher and rise by 24 % up to 68 %. The sector success rate is rarely affected and achieves 60 %. Hence, the rogue AP depends less on high signal strengths than the station impersonation.

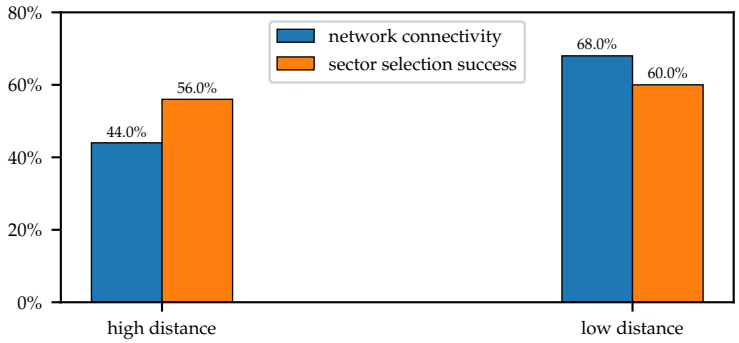

Figure 9.10: Network connectivity and success rate in selecting the sector of the rogue AP for the attacker at s higher and lower distance than the legitimate AP (observed over all experiments).

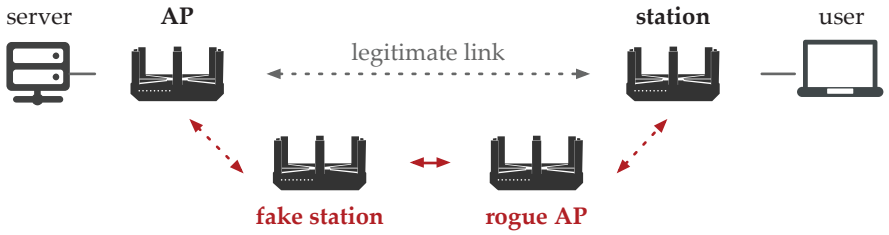

Figure 9.11: Experiment setup for the MITM attack with legitimate devices, the fake station, and the rogue AP. The attacker aims at intercepting the connection between the user and the server.

9.2.5 *Man-in-the-Middle*

The MITM *combines all previous attacks.*

Based on the previous results, we evaluate a combined MITM attack. For this, we tamper with the transmit sectors at legitimate devices and apply the station impersonation, as well as the rogue AP attack simultaneously. In the experiment, our devices are placed as shown in Figure 9.11. We set up a legitimate station and an AP at a distance of 3 m. Both establish a wireless link and, in addition to it, connect a user to a remote server. The MITM consists of a fake station and a rogue AP. The former is located at a closer distance to the legitimate AP, namely 0.5 m, while the latter is 0.5 m further from the legitimate station. Both are interconnected via Ethernet to relay packets. As in the rogue AP experiment, the attacker uses the decreased BI. The fake station makes use of re-associations to capture the connection from the legitimate one. For both malicious links, between the attacker's and legitimate devices, we forge the sector sweep feedback such that the selected beams favor the attacker. To prevent the fake station from connecting to the rogue AP, we locate them back-to-back and block the space between them. We evaluate the MITM attack by sending UDP packets and intercepting an HTTP connection.

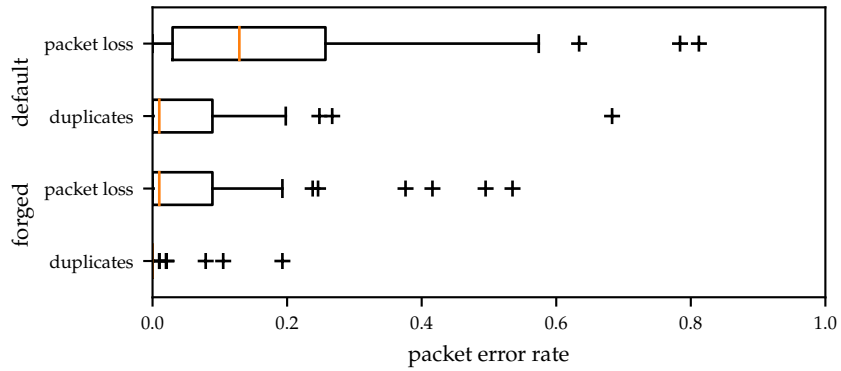

Figure 9.12: Packet error rates on the malicious link in the MITM evaluation scenario with the default behavior and by forging the sector.

We transmit packets and determine the path they take.

In the first part of our experiment, we send 100 UDP packets in 50 ms intervals from the user towards the server. Ordinarily, these packets would traverse the legitimate station and AP only. By starting the rogue AP and fake station, we aim at rerouting those packets through the man-in-the-middle such that the attacker obtains access. We choose the payload of the transmitted packets to carry a unique sequential number for identification. The server is configured to reply to all received packets and acknowledge the reception by sending back the same payload. This allows us to identify the packets that are traversing the rogue AP and the fake station. Throughout this experiment, we evaluate the packet loss and duplicate rates with and without forging the sector sweep feedback. We first launch the rogue AP and station impersonation as mentioned in the previous evaluation steps. Second, we additionally force the legitimate devices to select a sector with high gain into the direction of the rogue AP or fake station, respectively. By doing so, we investigate the impact of forged sector feedback on the performance of the MITM attack.

Forged feedback reduces the packet errors.

Our evaluation results show that forging the sector sweep feedback decreases the packet error rates as shown in Figure 9.12. With the default sector selection, as reported from all devices, we observe a median of the packet loss rate of 0.13 and a duplicate rate of 0.01. This means that many packets still take the legitimate link from the station to the AP without traversing the attacker. These results agree with our findings in the rogue AP and station impersonation: the legitimate devices alternately use the malicious and legitimate link. However, we can minimize these effects by forging the feedback and steer the legitimate devices towards the attacker. This decreases the median of the packet loss rate by 92 % to only 0.01. In most of the time, we do not observe any packet duplicate at all. Additionally, we calculate the rate of valid replies received by the legitimate station. As shown in Figure 9.13, the performance improves with forged sector sweep

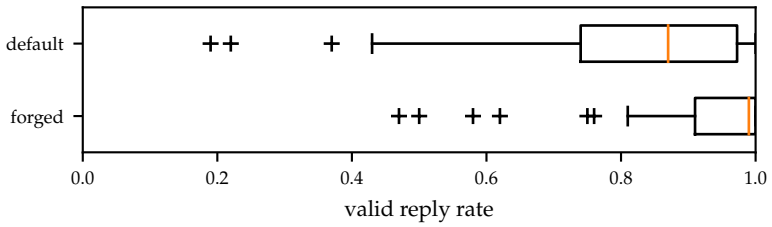

Figure 9.13: Rate of valid replies received at the legitimate station in the default behavior and by forging the sector.

feedback. The rate of correctly received reply packets increases from 87 % to 99 %.

Our attack facilitates an HTTP *interception.*

To demonstrate the practical impact of such a MITM attack, we utilize an HTTP interception example. We set up an HTTP proxy in transparent mode at the attacker to replace all image contents in a traversing HTML file. The required tools are provided by Mitmproxy[2] that easily enables to intercept, modify, and relay HTTP traffic flows. The result is obvious. While browsing the Internet, we see replaced images all over the websites. Hence, we successfully demonstrate that a MITM attack can be practically launched against a directional IEEE 802.11ad network.

9.2.6 *Detection Schemes*

We investigate four detection metrics.

Detection is the first step to protect against attacks such as the MITM presented above. In this section, we discuss four possible detection metrics to indicate whether an attack is happening or not. In particular, we consider the 1) the switching between sectors, 2) changes in the received signal strength, 3) beacon interval length, and 4) beacon counters as potential attack indicators.

Selected sectors and signal strength variations indicate an attack.

Since active eavesdroppers and the MITM attacker send additional IEEE 802.11ad frames, they are visible to the legitimate devices and detectable. Unfortunately, detecting if additional frames are received from an unintended transmitter is challenging. The best metrics that are available on off-the-shelf devices are to monitor the selected sectors and the resulting signal strength. If the attacker launches an attack, victims are forced to select another transmit sector instead of the legitimate one. Occasionally, feedback might be received from the legitimate device and the attacker, such that the sector selections start to fluctuate. Frequent switches between sectors with different patterns, might indicate an attack that injects forged feedback information. This also affects the signal strength, as choosing a sector that serves the eavesdropper most likely decreases the signal strength at the legitimate device. A simple detection scheme might trigger an alarm whenever

2 https://mitmproxy.org

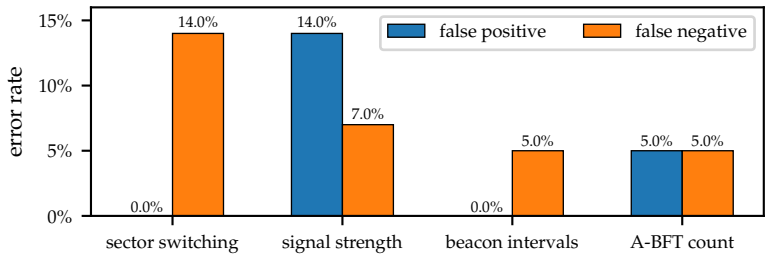

Figure 9.14: Error rates of four detection metrics.

the difference between consecutive signal strength measurements is higher than a specific threshold. However, natural effects such as blockage and mobility also cause switching beam patterns and high variations of the signal strength over time. For evaluation of these aspects, we record samples of the signal strength and the selected transmit sector during a mobile and attacking scenario. Recorded traces with 100 frames in 5 repetitions indicate that an attack can be differentiated from mobility. Monitoring the sector switches with a detection window of 20 sweeps, we obtain no false positives and a false negative error rate of 14 %, as shown in Figure 9.14. Analyzing the signal strength provides similar results: the detection achieves a false positive and false negative error rate 14 % and 7 %, respectively. It appears that observing the sector switches is slightly more accurate than considering the signal strength, which incorporates naturally high noise. Nevertheless, both detection schemes based on switching sectors and the signal strength exhibit high error rates.

The rogue AP *sends additional beacons.*

Detecting the rogue AP, in particular, we analyze the time difference of beacon intervals, compare it with the announced values, and count if more beacons are received than expected. Additional beacons from an attacker shorten the measured interval as new beacons appear earlier than expected. If we detect beacons, which deviate the schedule of the previous beacons, it is quite obvious that a rogue AP exists. Further, the IEEE 802.11ad beacon frame format contains an A-BFT field that counts the number of beacons that have been already transmitted. Considering this value reveals if too many beacon phases by different devices take place. For evaluation of these two detection metrics, the beacon interval and the A-BFT count, we record 20 samples in a regular and in an attack scenario with the rogue AP. We repeat our measurements 5 times. The detection of beacon irregularities is processed in monitor mode by capturing all control frames. In our implementation, we exemplary use a threshold for the time difference of 2.5 % and trigger an alarm whenever it is exceeded. For the A-BFT count, we compare the values of two consecutive beacon intervals and classify an attack when the current value is not higher than the previous one. The results of these two detection schemes are also shown in Figure 9.14. Their performances are better than those of

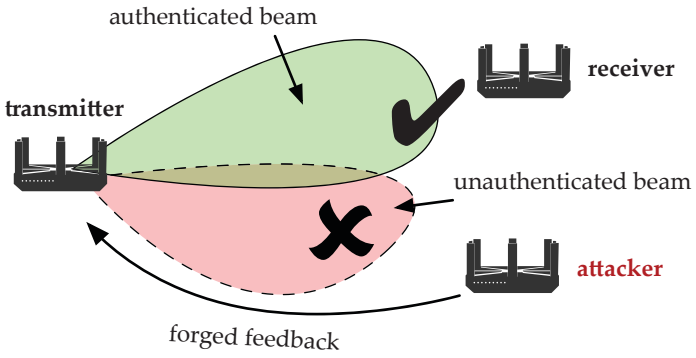

Figure 9.15: Schematic illustration of authenticated sectors that protect against forged feedback in beam stealing attacks.

observing the transmit sector and signal strength. Taking the beacon interval into account reveals no false positives and a false error rate of 5 %. Analyzing the A-BFT reveals a false positive and negative error rate of 5 %, each.

Detection requires that at least some frame are received from the legitimate device.

All described schemes only detect the forged sector sweep feedback if at least a few frames or beacons from the legitimate devices are processed. A highly sophisticated attacker might jam all these frames such that the victims get completely isolated. By distorting the legitimate frames, jamming attackers may also reduce the number of re-associations that we observed in our experiments. In such a situation, when a victim only receives frames from the attacker, the proposed detection schemes are likely to fail. The upcoming section presents an authentication scheme as a possible countermeasure against beam stealing attacks.

9.3 AUTHENTICATION SCHEME

Our authentication scheme prevents beam stealing.

Due to the lack of any authentication mechanism in the sector sweep, IEEE 802.11ad devices are vulnerable to low-layer attacks such as beam stealing or denial-of-service. We propose a sweep authentication scheme, which ensures that devices only accept valid feedback from intended devices, as illustrated in Figure 9.15. It extends the existing sector sweep with a mutual feedback authentication by amending the frame format. The following first describes our scheme in detail and then provides results from our evaluation.

9.3.1 *Authenticated Sector Sweep*

Devices authenticate the sector sweep feedback.

The operation of our sweep authentication is split into a session establishment phase and an authenticated sweep phase. Devices have an asymmetric pair of keys, of which the public key is assumed to be verifiable by others. In the session establishment phase, two devices exchange their public keys. Both then agree on a session secret

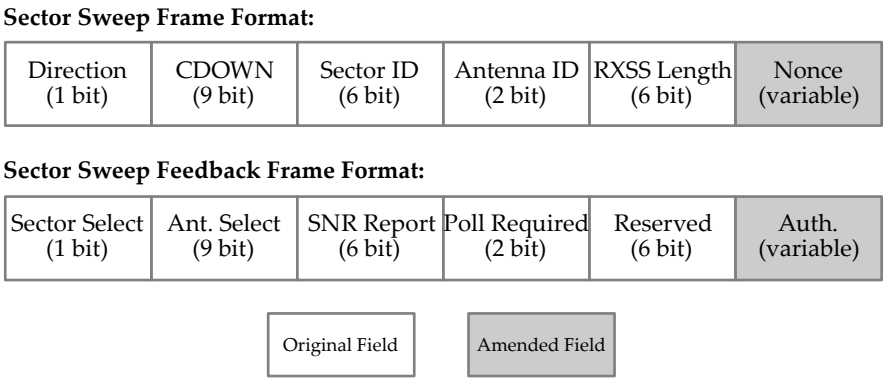

Figure 9.16: Amended sector sweep frame format.

that is derived from their keys. As secret keys are kept private, this session secret is unique for every link. During the authenticated sweep phase, devices perform the usual sector level sweep with specific additions. They generate a random cryptographic nonce for each sweep and append it to the sweep frames. Receivers select the best sector and determine an authenticator for the selected sector based on the received nonce as well as the session secret. This authenticator is appended to the sector sweep feedback. Devices then take the reported sector from the feedback and verify its authenticity using the respective nonce and session secret. The nonces ensure that selecting sectors is only possible for correctly received sweeps and protect against replaying. Due to the session secret, only responses from legitimate peers are accepted. Since attackers cannot forge the authenticator they are unable to alter the sweep feedback.

Running our protocol, we assume that the attacker can listen to all IEEE 802.11ad traffic and generate and inject arbitrary frames but does not interrupt the initial association. He can transmit more powerfully than any other device, but neither jam nor remove frames from the air. Devices are assumed, to be honest with each other and verify their identities. Common cryptographic protocols such as SHA-256 are assumed secure and cannot be broken in feasible time.

9.3.2 *Extended Frame Format*

We extend the frame structure.

To authenticate the sector sweep feedback, we extend the structure of sector sweep frames. Such frames are regular IEEE 802.11 control frames and contain a sector sweep, or a sector sweep feedback field, as shown in Figure 9.16. The sector sweep field specifies the direction of the sweep, a countdown of remaining sectors in the sweep, the sector and antenna ID, and the length of an optional receive sweep. The sector sweep feedback field contains the ID of the selected sector and antenna, the measured SNR for this sector and a poll request indicator. As common mm-wave devices use a single antenna and

perform a transmit sector sweep only, we omit the antenna selection and receiver sector sweep for simplicity in our design. However, this could be easily integrated in future work. In our sweep authentication, we extend the sector sweep field by an additional nonce and add an authenticator to the sector sweep feedback field. Both values have variable length, such that our design remains flexible and support extensible security and performance requirements. In the following, we provide details on our authentication scheme and describe how the nonce and authenticator fields are used.

9.3.3 *Protocol Design*

Our protocol is performed during a sector level sweep.

The authentication procedure of the sector sweep feedback is illustrated in Figure 9.17 and operates as follows. Given are an Initiator I and a Responder R that both perform our authentication protocol. Both devices have a key pair $\{pk_I, sk_I\}$, and $\{pk_R, sk_R\}$ respectively with which they can verify their identity. In the session establishment phase, they broadcast their public keys pk_I and pk_R as well as a random seed ϵ_I and ϵ_R. Using common cryptographic key exchange mechanisms (such as Diffie-Hellman), both determine a session secret s using their secret keys and the exchanged public keys. To this end, they apply a function $dh(pk, sk)$ that provides $s = dh(pk_I, sk_R) = dh(pk_R, sk_I)$. Moreover, s is randomized using the seeds ϵ_I and ϵ_R. Due to the cryptographic properties of asymmetric key exchange, only I and R can determine the session secret s and can use it for mutual authentication. During the initiator sweep with M sectors, I generates a unique cryptographic nonce ν_I. R receives at least some of the sweep frames, determines the best sector $\hat{m}$, and computes an authenticator α_R of length l_α as

$$\alpha_R = auth_{l_\alpha}(\hat{m}, \nu_I, s) = trunc(h(\hat{m}, \nu_I, s), l_\alpha), \tag{9.1}$$

where $auth_l(m, \nu, s)$ computes an authenticator of size l as a truncated hash function of the concatenation of m, ν, and s. With $trunc(x, l)$ we denote a function that truncates an arbitrary input x to a length of l bit. According to the NIST guidelines [Dan12], the output of a hash function can be truncated by extracting the leftmost bits. In the responder sweep with N sectors, R generates a nonce ν_R. I computes its authenticator as $\alpha_I = auth_{l_\alpha}(\hat{n}, \nu_R, s)$ for the selected sector $\hat{n}$. After both devices complete their sweeps, I sends $\hat{n}$ and α_I in a feedback frame and R acknowledges with $\hat{m}$ and α_R. Finally, both devices verify the received authenticators. I verifies α_R with $\hat{m}$ and s by computing the truncated hash $auth_{l_\alpha}(\hat{m}, \nu_I, s)$ itself. Similarly, R verifies α_I with $\hat{n}$ and s. If both verifications succeed, the devices set their sectors accordingly and continue the communication.

As the space for nonces is limited, and they might repeat, we renew the session secret to protect against replaying outdated authenticators.

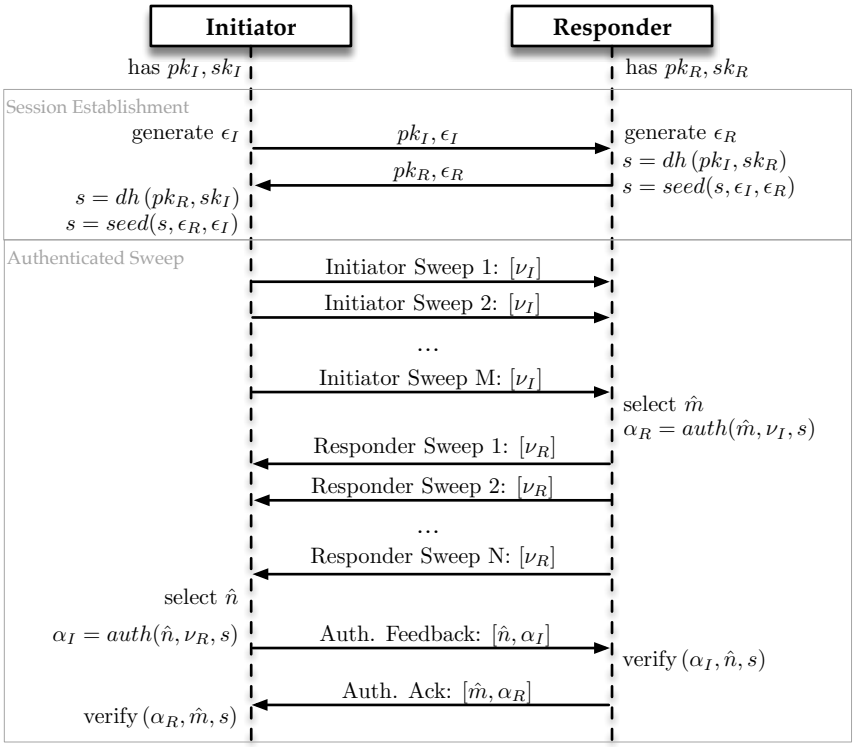

Figure 9.17: Simplified sequential operation of the sector sweep with authentication.

Session secrets are renewed to prevent repeating nonces.

To achieve this, we periodically start over with the session establishment to generate new session secrets. Devices should implement a countdown that decreases with every sweep, such that it can trigger the renewing process before the limited space of nonces expires.

Sweep authentication is independent of specific protocols.

Sweep authentication only extends the operation of the sector level sweep without changing any internal operation. Therefore, our design is globally applicable and independent of specific cryptographic algorithms. Particular algorithms, as well as the length of keys, nonces, and authenticators, can be chosen on demand. In the following section, we present an evaluation of our proof-of-concept implementation.

9.3.4 *Performance Overhead*

We evaluate the performance overhead of our sweep authentication in a combination of simulations and practical measurements. This section first describes our evaluation setup and then presents the results of our simulations and measurements. Finally, we end this section with a summary of our findings.

EVALUATION SETUP. To assess the performance of transmitting sweep authentication frames, we provide a proof-of-concept implementation of our amended sector sweep frame format in the Network

We combine simulations with practical measurements.

Simulator 3 (ns-3) [NS3]. In particular, we use the IEEE 802.11ad implementation from Assasa and Widmer [AW16] and integrate our extended frame format. The existing implementation already provides the essential components of the IEEE 802.11ad protocol, such as the sector sweep and the respective MAC layer frame formats. For our evaluation, we reserve additional space at the end of the sector sweep and sector sweep feedback frames to incorporate the nonces and authenticators. Our simulation environment determines the required transmission time of our authentication scheme and the regular sector level sweep. Unfortunately, the ns-3 simulator does not allow to consider the packet processing delay that is caused by cryptographic operations. To compensate for this drawback, we additionally evaluate the cryptographic computation overhead on commodity mm-wave devices. In particular, we use Talon AD7200 routers (see Chapter 4) and implement common cryptographic algorithms from the OpenSSL [SSL] library. Benchmarking the performance of these algorithms on the resource-constrained router hardware, allows us to estimate the average processing delay. Combining simulation results and practical measurements, we determine the effective performance overhead of sweep authentication in comparison to the original sweep. Our results are described in the following.

The size of nonces has a high impact on the performance.

SIMULATION RESULTS. We first evaluate the impact of different nonce sizes in our authentication. To this end, we run simulations of the sector sweep in ns-3 with 60 sectors at both devices. In the configuration, we set the space that is reserved for the authenticators to 8 Byte and use the SHA-256 algorithm to compute the cryptographic hashes. Our results for nonces with up to 5 Byte are shown on the left in Figure 9.18. The time overhead that is caused by our sweep authentication increases linearly with the size of the transmitted nonces. As nonces are transmitted for every sector in the sweep, they have a high impact on the total completion time and cause the most significant overhead. Nonces of 1 Byte causes a time overhead of about 2 %. With a size of 5 Byte, the overhead increases to about 9 %. To achieve a feasible trade-off between provided protection and time overhead, we continue evaluating the authenticator size with nonces of 4 Byte.

Authenticators are only transmitted once per device.

In contrast to the size of nonces, the size of authenticators in the sector sweep feedback has only a minor impact on the overall performance. For our second set of simulations, we configure 60 sectors, use the SHA-256 algorithm, and evaluate authenticator sizes of up to 8 Byte with 4 Byte nonces. Our results are shown on the right in Figure 9.18. Increasing the authenticator size from 1 Byte to 8 Byte only causes a difference of about 0.1 %. Independent of the number the sweeps, the authenticator is only transmitted twice (once by each device) in the sector sweep feedback or the acknowledgment. Thus, one can be less strict with the size of the authenticators. In summary,

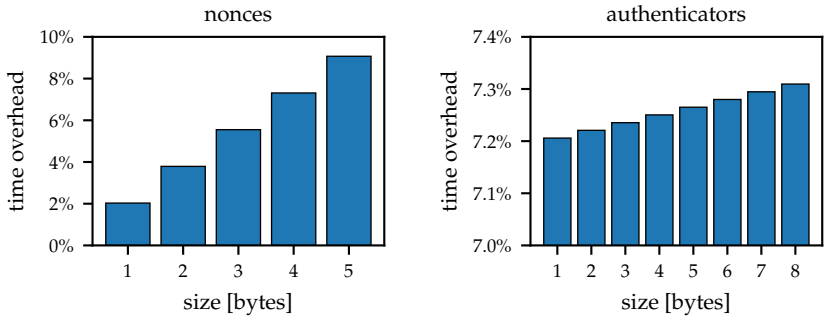

Figure 9.18: Transmission time overhead with different nonce (left) and authenticator (right) sizes compared to the original sector sweep.

we achieve a performance overhead of 7.31 % with 4 Byte nonces and 8 Byte authenticators.

The number of active sectors has no impact on the relative performance overhead.

In our third simulation, we investigate the impact of the number of used sectors on the performance overhead. The more sectors being used, the more sweep frames need to be transmitted. Despite adding nonces to each of these frames, the overhead is following a linear trend. The number of transmitted authenticators is independent of the number of sectors. As a result, the relative transmission overhead remains constant and the number of sectors itself does not affect the relative performance of our sweep authentication.

We benchmark different hash algorithms.

PRACTICAL MEASUREMENTS. In practical measurements on off-the-shelf devices, we investigate the performance of different cryptographic hash implementations. In particular, we consider the most commonly used algorithms SHA-1, SHA-256, SHA-512, as well as HMAC(MD5) with block sizes of 64 Byte. On the router, we measure the time that is needed to complete the required cryptographic operations and integrate the measurements as additional delay in the simulation. We hash random values over an interval of 3 seconds and determine their average duration. Our measurements are shown in Figure 9.19. The computation of an SHA-512 hash takes about 2.52 µs on average. Processing with SHA-256 is about 0.99 µs faster and takes only 1.53 µs. As in the previous simulations, we consider the performance overhead in a sweep with 60 sectors in which we set the sizes of nonces to 4 Byte and add 8 Byte authenticators. Despite variations in the processing time, the impact of different hash algorithms on the overhead is only about 0.1 %.

The generation of the session secret cause a one-time overhead.

Computing a session secret causes additional time overhead. Investigating this, we measure the average processing time for performing an Elliptic-curve Diffie-Hellman (ECDH) key exchange. In experiments, the devices generate a 160 bit and 192 bit session secret in about 1.3 ms and 1.8 ms, respectively. Since shared secrets are only generated to start a session, the overhead of a few milliseconds is sustainable.

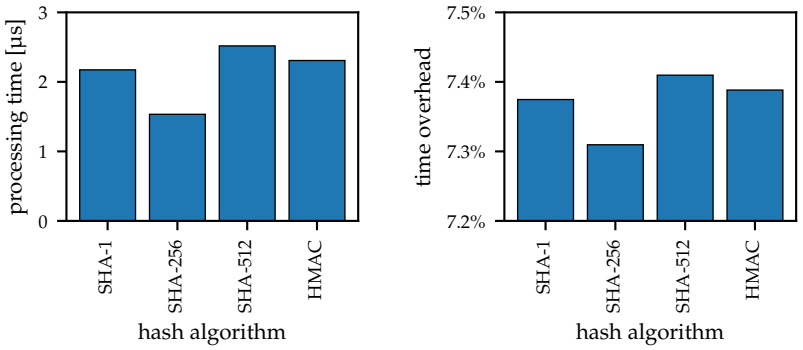

Figure 9.19: Processing time and overhead with different hash algorithms compared to the original sector sweep.

The session secret is renewed every few minutes.

As the available space to generate random nonces is limited, the devices periodically renew the session secret to mitigate collisions. As described by Zenner in [Zen09], the relation of maximum collision probability p_{max}, bitlength l, and the number of nonces θ can be expressed by $\frac{\theta^2-\theta}{2*2^l} \leqslant p_{max}$. With nonces of 4 Byte ($l = 32$) and assuming a maximum collision probability of $p_{max} = 2^{-9}$, we can use $2^{12} = 4096$ collision-free nonces. Following the maximum link maintenance time of 128 ms from the IEEE 802.11ad specification [IEE14], the devices should renew the secret every $2^{12} \times 128\,ms = 8.7\,min$. Hence, the generation session secrets cause additional overhead only one every few minutes.

Sweep authentication protects against beam stealing with feasible overhead.

RESULTS. Our evaluation shows that adding authenticators and nonces to the sector level sweep causes an additional time overhead due to the transmission of larger frames. With SHA-256 and embedding 4 Byte nonces and 8 Byte authenticators, the overhead compared to the original sweep is only 7.31 %. In the simulation, we find that the most significant performance impact is caused by nonces that must be transmitted for every sector. In contrast, the authenticators are only transmitted once by every device. We measured the processing time that is required to compute the cryptographic operations on the Talon AD7200 tri-band routers. Common algorithms affect the performance overhead by only 0.1 % on average. In conclusion, our sweep authentication achieves a feasible overhead protecting the feedback in sector level sweeps.

9.4 DISCUSSION AND SUMMARY

Lower layer amendments increase the attack surface.

The IEEE 802.11ad standard adopts directional communication to overcome the challenge of the increased signal attenuation in the mm-wave frequency band. In this regard, several new features on the lower layers are introduced such as predefined antenna sectors with directional

transmission features and the sector level sweep protocol for beam training. In this chapter, we identify these extensions as a new attack vector that threatens wireless networks. A distant attacker can manipulate the beam training and force victims to steer their antennas to arbitrary directions. Forging the sweep feedback increases the chances of eavesdropping or launching a MITM attack. This is particularly harmful if no higher layer encryption and authentication mechanisms are applied. But also if higher protocol layers are properly secured beam stealing constitutes a threat. An attacker that gathers encrypted data frames has still access to unprotected signaling and channel control. It might steer the beams to off-site directions and effectively cause a denial-of-service. In contrast to classical jamming, this becomes particular energy efficient as only a few frames have to be transmitted. Individual devices could be effectively isolated from the network. Forcing multiple devices to steer their signals towards one specific location may result in distributed jamming. Furthermore, selfish devices may exploit the beam training of others to minimize interference and maximize their signal quality. The more applications make use of directional mm-wave communications; the more attack scenarios are likely to appear.

Beam stealing enables active eavesdropping and a MITM *attack.*

We demonstrate the feasibility of a beam stealing attack for active eavesdropping and for acting as MITM with off-the-shelf IEEE 802.11ad devices. Our results from practical experiments highlight the significance of such an attack. Beam stealing increases the performance of eavesdropping by 38 % while still achieving a notable throughput of 1.4 Gbps. A MITM attack successfully tampers with an HTTP connection and achieves a median packet error rate of only 1 %. Current devices typically perform transmit-side beamforming only and deploy a fixed quasi-omnidirectional sector for receiving. Tampering the receiving beams could further increase the attacker's possibilities. Devices could not only be prevented from transmitting in particular directions but also from receiving. Steering receive beams in the direction of an attacker could turn jamming highly efficient. Such weaknesses cannot be mitigated in particular implementations as the standard lacks arrangements for protection. With the upcoming IEEE 802.11ay standard [Gha+17], new beam training protocols will be introduced. These protocols induce higher complexity to make beam training more efficient. We believe that they are likely to increase the attack surface as well.

Although our experimental analyses only cover line-of-sight scenarios with short distances, the described beam stealing attack is likely to be successful in non-line-of-sight scenarios with larger distances as well. The impaired signal quality of the attacker that resides further away may be compensated with higher transmission power. Reflections and additional attenuations make injecting forged feedback more challenging but not impossible. Our authentication scheme is inde-

pendent of the attacker's link quality, thus mitigates attacks in both scenarios.

New protection schemes are required to handle attacks on directional networks.

Current detection and protection schemes are insufficient to address attacks on mm-wave beam training. In contrast to data frames, the standard provides no authentication mechanisms to protect sector sweep frames against misuse. The sector level sweep is performed before any link is established, thus before the communication channel could be secured. This makes forging the feedback particular hard to defend. New low-layer protocols that rely on such feedback to control the beam alignment must be secured from bottom-up. For example, authenticating the sector sweep frames could prevent attackers from injecting arbitrary feedback.

Our sweep authentication protects against beam stealing with feasible overhead.

We propose a sweep authentication that protects against beam stealing attacks. Performing a key exchange and authenticating the sector sweep feedback, it ensures that responses to the sweep are only accepted from legitimate devices. In a combination of simulations and measurements, we demonstrate the feasibility of our approach and find that the incurred time overhead is only about 7.3 % higher than that of the original sweep. The protocol behind our approach is flexible and allows to configure a trade-off between protection and performance, thus, adapts to various applications and requirements. Lightweight authentication schemes with pure symmetric cryptography are likely to cause less overhead. In this work, we focus on the transmission overhead that is caused by embedding additional authentication fields in the frame format and assess the general feasibility of authentication in the sector level sweep. A thorough security analysis might be required to address different attack scenarios and mitigate physical layer attacks such as jamming. Distance bounding along with secure device localization based on mm-wave signal characteristics appear promising to also detect and defend against beam stealing. We strongly suggest to come up with proper solutions to address the efficiency and security issues that currently impede directional mm-wave communications.

Part V

MEASUREMENTS

In this part, we provide complementary measurements with our testbed experimentation platform. We obtain the radiation patterns of the default sectors, individual antenna elements, and custom directional beams in Chapter 10. In Chapter 11, we measure the performance of different beams in a practical large-scale deployment and outline the advantages of directionality.

10 ANTENNA RADIATION PATTERNS

Measurements of beam patterns allow to asses their performance in different directions.

Knowing the beam patterns of the sectors that are used in mm-wave devices is crucial to get an in-depth understanding of propagation effects and systems performance. Their shape highly depends on the antenna's geometries and the placement inside the device. Due to the complex layout of conventional antenna arrays with dozens of elements in off-the-shelf devices, their radiation patterns are hard to predict by calculations. For example, the antenna array in the Talon AD7200 exposes 12 patch antennas on the front surface, 6 on the back, and 14 additional dipole antennas on the side. Uniform rectangular arrays as often used in theory are not the norm. Precise measurements in all spherical directions are required to assess the performance of particular beam patterns. Using our practical testbed experimentation platform, we precisely measure different radiation patterns of the antenna array in the Talon AD7200.

Section 10.1 describes our setup that is used to measure the default beam patterns in Section 10.2. In Section 10.3, we reveal the patterns of individual elements in the antenna array, which lead to array factor in Section 10.4. Section 10.5 provides measurements of custom beams with different characteristics such as strong directionality and multiple lobes. Finally, Section 10.6 summarizes our findings.

10.1 MEASUREMENT SETUP

We perform measurements in an anechoic chamber and precisely steer one device mounted on a pan-tilt unit.

Our measurement setup, with which we measure the radiation patterns, consists of two Talon AD7200 devices in an anechoic chamber. As shown in Figure 10.1, we mount one device on a pan-tilt unit that uses two precise step-motors to orient it in different directions. In particular, we use a FLIR PTU-E46-70 that achieves an accuracy of up to 0.003°. The Talon AD7200 is mounted with a custom adapter plate directly above the center of the rotation axis. The second device is placed three meters away on a tripod facing the first one. The walls, floor, and ceiling of the room are covered with radio-wave absorbing foam to omit disturbing reflections and multi-path effects. For each measurement, we establish a connection between both devices and perform frequent sector level sweeps. During this, both devices record the Signal-to-Noise Ratio (SNR) for received frames in all sectors. The codebook with the sector definitions varies among our measurements. We start with the predefined codebook that comes with the firmware to obtain the default beam patterns. In the subsequent measurements with custom beams, we adjust the codebook accordingly. Specifi-

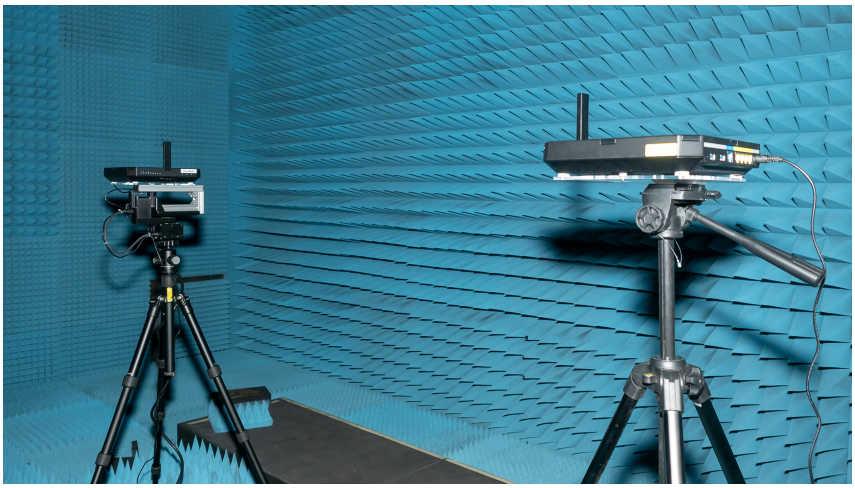

Figure 10.1: Measurement setup with two Talon AD7200 devices in an anechoic chamber. The device on the left is mounted on a pan-tilt unit that steers it to different spherical directions.

cally, we use the features of our testbed experimentation platform described in Chapter 4. Our automation system remotely triggers the measurements on the devices over a 2.4 GHz management network. The pan-tilt unit, remotely controllable via Ethernet, is handled by our automation as well.

10.2 DEFAULT BEAM PATTERNS

Devices use transmit sectors that are defined in a codebook.

Access Points (APs) periodically announce their existence by successively transmitting beacon frames in different sectors. Stations that overhear at least one of these beacons initiate the sector level sweep to find the optimal transmit beam. While the sectors used for beaconing and the sector level sweep can differ, they are all configured in a single codebook. The recent firmware for the mm-wave module in the Talon AD7200 supports up to 64 different transmit sectors of which 36 are defined by default (see Chapter 4). Beam training is only applied for transmit sectors; the devices use a single quasi-omnidirectional sector for reception. APs trigger the beaconing every 102.4 ms [IEE14]. When devices communicate, sector level sweeps are initiated at least once per second or whenever necessary. As experimentally revealed, performing a mutual beam training with the default sectors takes on average 1.27 ms. Probing a single sector needs 18.0 µs; additional 49.1 µs are required for hand-shakes and feedback. In the following, we measure the SNR of the predefined sectors used in the sector level sweep in two- and three-dimensional setups.

2D PATTERN MEASUREMENTS. In our first experiment, we measure the predefined transmit sectors of the Talon AD7200. The pan-tilt unit steers the device in steps of 0.75° from −159° to 159° in pan-

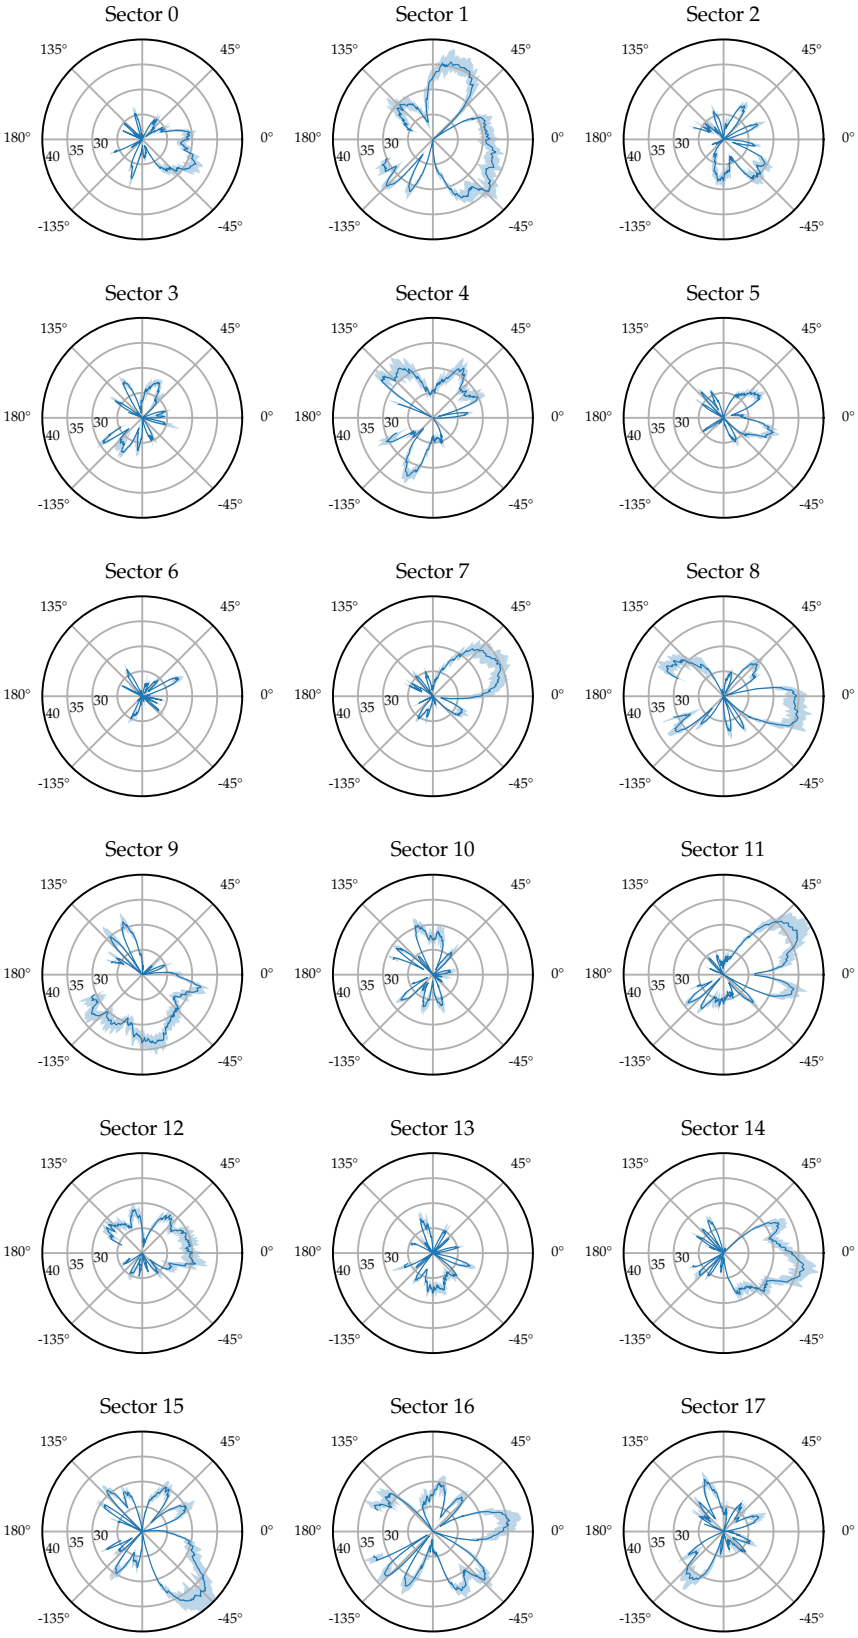

Figure 10.2: Measured beam patterns for the default sectors 0 - 17.

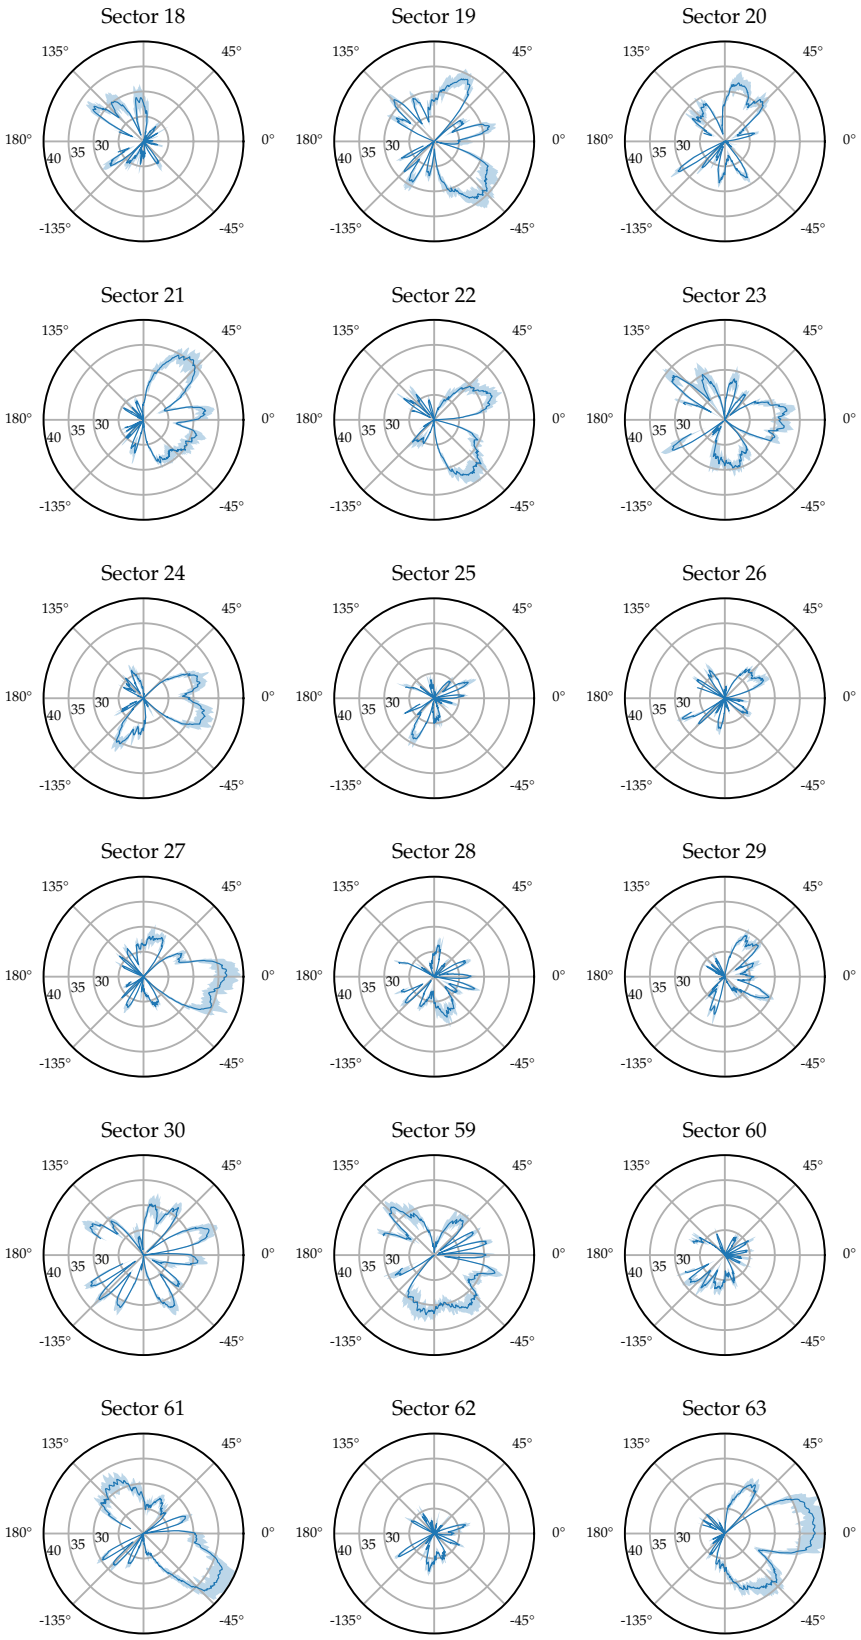

Figure 10.3: Measured beam patterns for the default sectors 18 - 63.

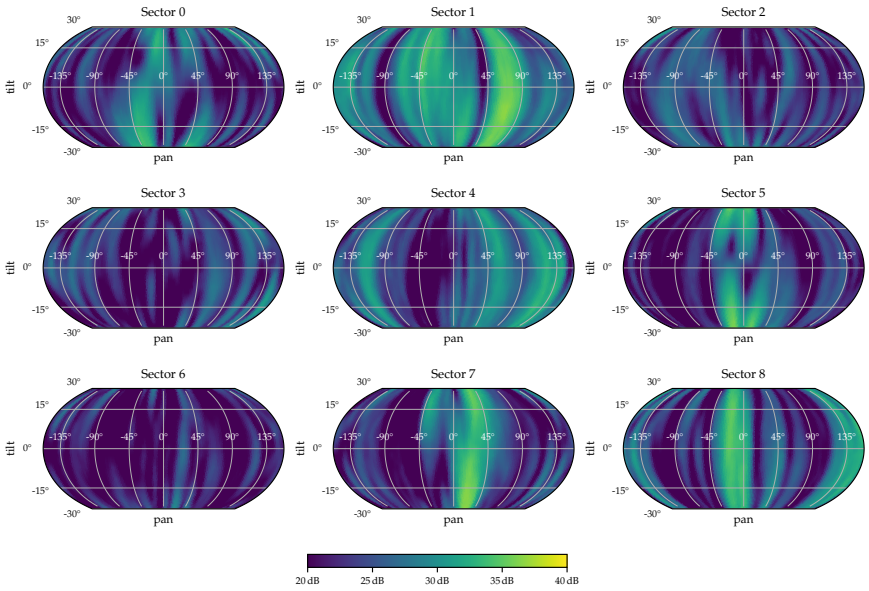

Figure 10.4: Measured three-dimensional beam patterns for the default sectors 0 - 8.

direction, while keeping a constant tilt of 0°. In each step, the second device records the SNR while the first one sends sector sweep frames. Averaging over at least 20 measurements, we obtain the radiation patterns. Figure 10.2 and Figure 10.3 illustrate our results for all 36 default sectors. Shadows in the plots represent the 95 % confidence intervals. The measured beam patterns show different characteristics. We find that some transmit beam patterns, such as those of sectors 7, 8, 14, 15, 21, 27, 61, and 63, expose a clear main-lobe. They provide a strong gain in one particular direction. Other beams, such as those of sectors 1, 4, 16, 19, 23, and 30, come with multiple, equal powered lobes. Sectors 9 and 59 cover a wider range. Due to these strong variations and the irregular beam shapes, we do not provide specific beamwidths or steering angles. Towards the back of the antenna—for angles higher than 120° or lower than −120°—most beam patterns get distorted. This is not surprising since the antenna array is partially shielded in this direction. Few beam patterns, such as those of sectors 3, 5, 25, 26, 28, 29, 60, and 62, exhibit low gains in all directions. It is likely that they have their maximum outside the evaluated horizontal plane. To account for these directions as well, we extend our measurements to the third dimension with different tilts of the device.

We first measure the antenna radiation patterns in the planar space.

Most beam patterns expose irregular shapes.

3D PATTERN MEASUREMENTS. To map the antenna radiation patterns in spherical space, we repeat our measurements and additionally tilt the rotation head from −30° to 30°. We use the same setup as before, take measurements at different pan angles from −159° to 159°, but decrease the accuracy to 2.25°. Still, this results in more than 3800 iterations. Figure 10.4 and Figure 10.5 illustrate the measured beam

We extend our measurements to the spherical space.

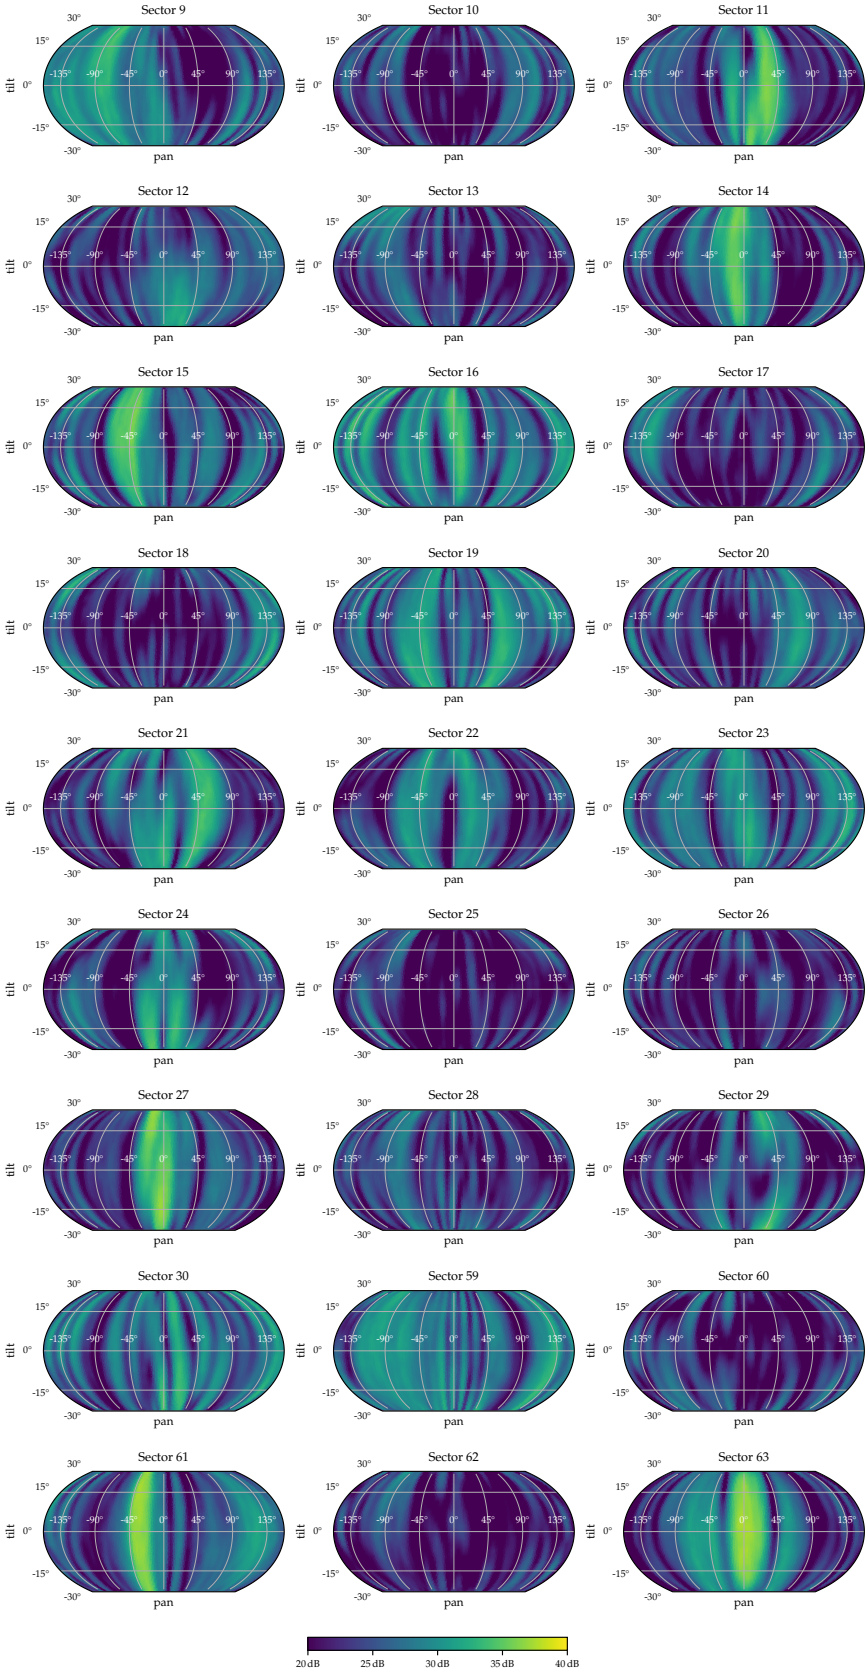

Figure 10.5: Measured three-dimensional beam patterns for the default sectors 9 - 63.

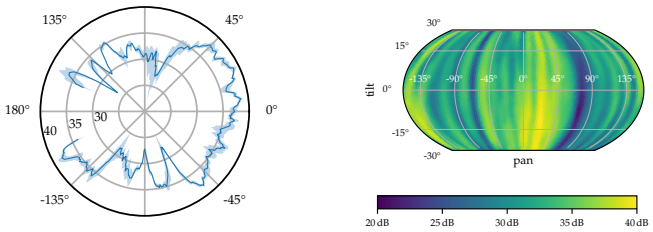

Figure 10.6: Radiation pattern of the default receive sector.

patterns projected to a spherical surface. The x- and y-axes represent different pan and tilt angles, respectively. Our results indicate that most beam patterns vary at different elevations. For example, sector 24 shows a decent lobe to the front in high elevation which turns into a gap at tilt angles below zero. Comparing the spherical beam patterns to those of the planar measurements in Figure 10.2 and Figure 10.3, we find that some sectors increase their gain when tilted. Especially for sectors 5, 12, 13, 18, 24, 28, 29, and 62 the highest signal strength increases by about 3 dB to 6 dB. Sectors 6, 25, and 26 still not expose a clear main-lobe within the measured space. Nevertheless, we obtained sufficient information to assess the default transmit beam pattern performance in the most relevant communication directions.

The receive sector provides a quasi-omnidirectional coverage with small gaps.

RECEIVE BEAM PATTERN. The single receive sector is intended to provide a quasi-omnidirectional coverage. To obtain its beam pattern, we switch the roles of the two devices in our measurements. While the fixed device continuously transmits sector sweep frames, the one on the pan-tilt unit evaluates the SNR. Minimizing the noise-floor in our results, we only consider transmit sector 63, which provides the highest gain. Since the transmit power and direction remains constant, the measured SNR corresponds to the receive gain of the antenna. Figure 10.6 shows the beam pattern in two- and three-dimensional measurements. In contrast to all transmit beam patterns, it has fewer variations and features a constant gain in most directions. However, at specific angles, it clearly exposes lower gains which might impede the signal reception. These irregularities can be particularly harmful as the default operation cannot switch to different sectors.

10.3 INDIVIDUAL ANTENNA ELEMENTS

Patch antennas provide a strong gain to either the front or the back.

Radiation patterns are a combination of the signals from multiple antenna elements in the array. To analyze their individual beam patterns, we repeat the previous measurement with a custom codebook of sectors that have only a single antenna element active. The gains of all other elements and all phase values are set to zero. As before, we pan the device from $-159°$ to $159°$ and tilt it from $-30°$ to $30°$. The

accuracy is kept at 0.75° and 2.25°, respectively. Figure 10.7a shows the beam patterns for three exemplary elements on the front surface of the array. They all expose a strong gain between ±45°. Depending on the location of the element in the array, one side might be slightly favored. For example, the pattern of antenna element 10 exposes a lower gain for orientations above 10°. Elements on the back surface clearly show a lobe in the opposite direction, shown in Figure 10.7b.

The RFIC partially blocks the signals.

On this side, the Radio Frequency Integrated Circuit (RFIC) partially blocks the radiation. The four antenna elements, 4, 19, 20, and 21, on the left side of the RFIC, expose strong gains at about −160°. For the two antenna elements, 11 and 13, on the other side, the beam patterns orient towards 160°, respectively. The dipole antennas on the sides of the array expose different characteristics, as shown in Figure 10.7c. Elements 2 and 26 are located on one side while elements 9 and 17 are on the other. All the remaining radiation patterns show a diffuse shape. These antennas are likely to be used to fine-tune the beams in various directions. In the upcoming section, we extend the measurements of all these elements and additionally take into account the relative phase among them.

10.4 ANTENNA ARRAY FACTOR

The antenna array factor describes the complex gains of all individual elements.

The antenna array factor is a characteristic that describes the complex gain of an array in different directions. It allows to compute the radiation pattern for a specific antenna configuration and also enables to derive particular steering characteristics. This possibility facilitates to generate beams with high directionality or multiple lobes.

To obtain the array factor, we measure the Channel State Information (CSI) as described Chapter 7 for each of the different device orientations in our setup. Since the devices are rather close to each other, the strong link estimation (see Section 7.1.3) is sufficient. In contrast to the previous measurements of individual antennas, we also record the relative phase between the elements and estimate their complex gains. Since any multi-path propagation effects do not distort our measurements, these gains correspond to the antennas' physical radiation characteristics. Combining them, we construct the array factor. For an orientation defined by the azimuth and elevation angles (ϕ, θ), we define the array factor $AF_{(\phi,\theta)}$ as a vector of the elements' complex gains. In particular, we derive

$$AF_{(\phi,\theta)} = [c_0(\phi,\theta), c_1(\phi,\theta), \ldots, c_{31}(\phi,\theta)]^T, \tag{10.1}$$

where $c_i(\phi,\theta)$ represents the complex gain at the i's antenna element in measured in (ϕ,θ) direction. These values are recorded for azimuth angles ranging from −159° to 159° and elevation angles from −30° to 30°. Thus, our measurements obtain the physical antenna radiation characteristics in the most relevant directions.

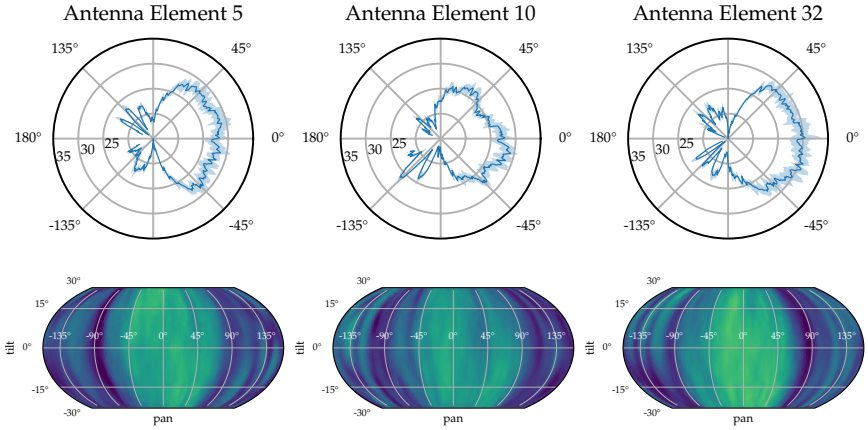

(a) Antenna elements on the front surface

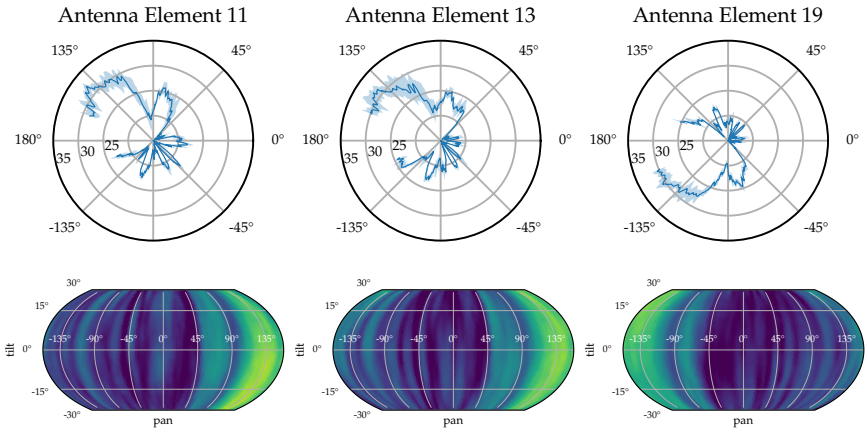

(b) Antenna elements on the back surface

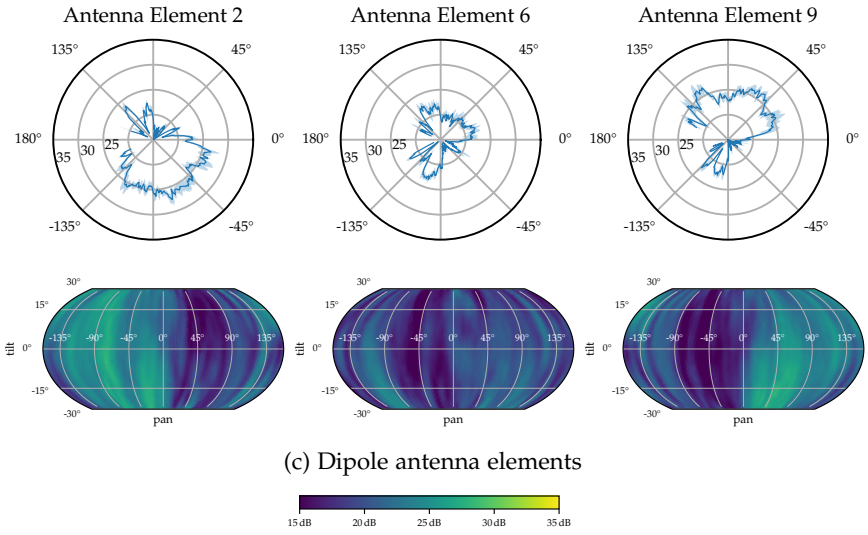

(c) Dipole antenna elements

Figure 10.7: Excerpt of the measured beam patterns of individual antenna elements.

The array factor allows estimating the radiation of particular beam configurations.

In the knowledge of the array factor, we estimate the beam pattern shapes given particular antenna configurations. Radiation patterns are a superposition of the complex gains of individual elements with specific weights applied. The antenna array factor and the weights from the configuration allow estimating the radiation pattern without explicitly measuring it. Given the weights for all antennas elements in $W = [w_0, w_1, \ldots, w_{31}]$, we compute the radiation pattern from the scalar product with the array factor $\mathrm{AF}_{(\phi,\theta)}$ as $p(\phi,\theta) = \left|W \cdot \mathrm{AF}_{(\phi,\theta)}\right|$. This computation reveals the relative strength a beam in arbitrary directions.

Beam patterns with specific geometrical characteristics can be derived from the array factor.

Moreover, the array factor enables to derive beams with specific geometrical characteristics. For example, knowing the $\mathrm{AF}_{(\phi_0,\theta_0)}$ for a direction (ϕ_0, θ_0) allows to obtain the weights W_0 that optimize the beam in this particular direction as $W_0 = \mathrm{conj}\left(\mathrm{AF}_{(\phi,\theta)}\right)$ (see Chapter 7). Multi-lobe beam patterns can be generated by either averaging over multiple of such beams or solving a matrix equation as

$$W_{\mathrm{ml}} = \begin{bmatrix} c_0(\phi_0,\theta_0) & c_1(\phi_0,\theta_0) & \cdots & c_{31}(\phi_0,\theta_0) \\ c_0(\phi_1,\theta_1) & c_1(\phi_1,\theta_1) & \ldots & c_{31}(\phi_1,\theta_1) \\ \vdots & \vdots & \ddots & \vdots \\ c_0(\phi_n,\theta_n) & c_1(\phi_n,\theta_n) & \ldots & c_{31}(\phi_n,\theta_n) \end{bmatrix}^{\mathrm{P}} \cdot [1, 1, \ldots, 1], \tag{10.2}$$

where M^{P} is the pseudo-inverse of a matrix M. Similar to this approach, we design beam patterns that minimize the radiation in specific directions. With this tool at hand, beamforming on devices becomes more accurate and set up a beam that exactly steers towards the intended receivers.

We compute directional, multi-lobe, and side lobe suppression beam patterns.

Exemplary beam patterns with normalized power levels are shown in Figure 10.8. Figure 10.8a illustrates directional beam patterns that point into one specific direction. They are ideal for high throughput transmissions but still expose few side lobes. Figure 10.8b shows multi-lobe beams with two and three steering directions. Keeping the steering directions close, as shown in the pattern on the right, we obtain a wide lobe. Moreover, isolating specific directions reduces side lobes. Figure 10.8c shows beam patterns with the same steering directions than in Figure 10.8a but additional isolation directions in the side lobe directions. The isolations effectively suppress the side lobes while only slightly decreasing the gain in the steering direction. Depending on the application scenario, different beam steering possibilities exist. Within this work, we only provide an excerpt of possible beam patterns that could be derived from the array factor. In the following, we detail the characteristics of directional beam patterns.

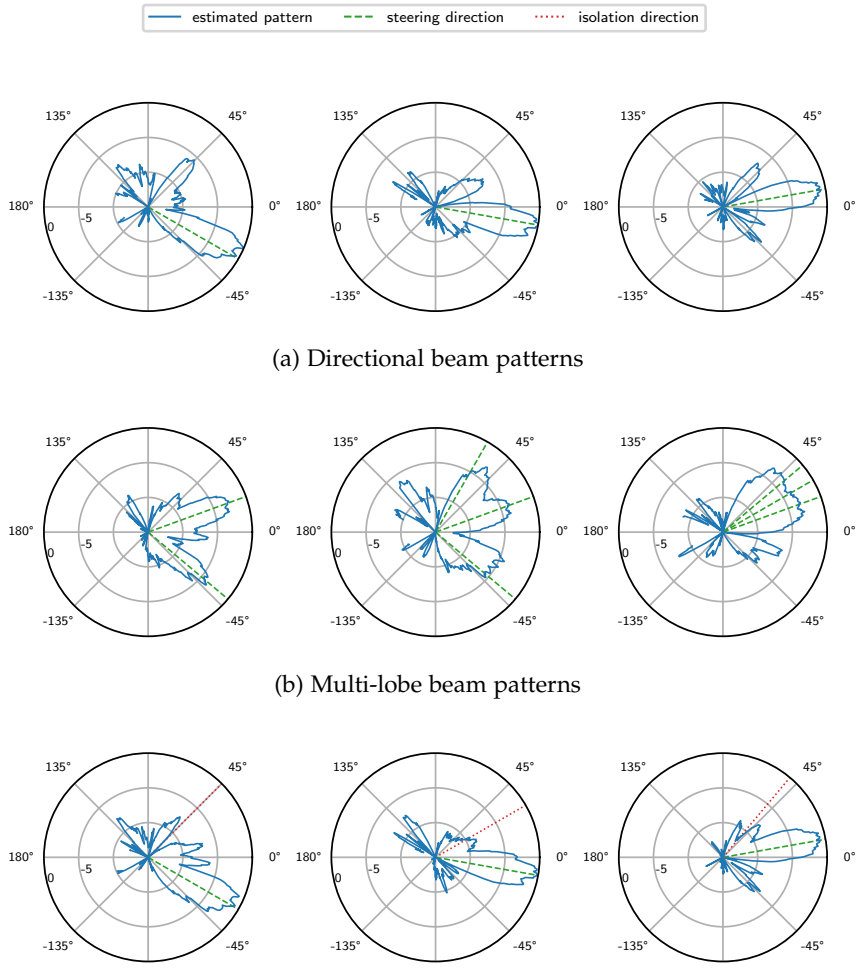

(a) Directional beam patterns

(b) Multi-lobe beam patterns

(c) Side-lobe suppression beam patterns

Figure 10.8: Estimated shape of custom directional, multi-lobe and side lobe suppression patterns. All patterns are generated with the array factor and use 16 active antenna elements.

10.5 CUSTOM DIRECTIONAL BEAMS

We measure custom directional beam patterns.

This final measurement reveals the beam patterns of custom directional sectors. Following our adaptive beam optimization approach from Chapter 7, we first steer the device to nine different directions of −60°, −30°, −20°, −10°, 0°, 10°, 20°, 30° and 60° and estimate the CSI. Based on this, we compute an optimal beam pattern for each direction that maximizes the signal strength. Since our environment does not expose any reflectors, the resulting beams are likely to expose a strong main-lobe in the specified direction. Setting these beams in a codebook and repeating our measurements, we obtain the two- and three-dimensional radiation patterns shown in Figure 10.9. Indeed, our measurements expose a strong directionality. In contrast to the predefined beam patterns, they exhibit a stronger main-lobe and lower

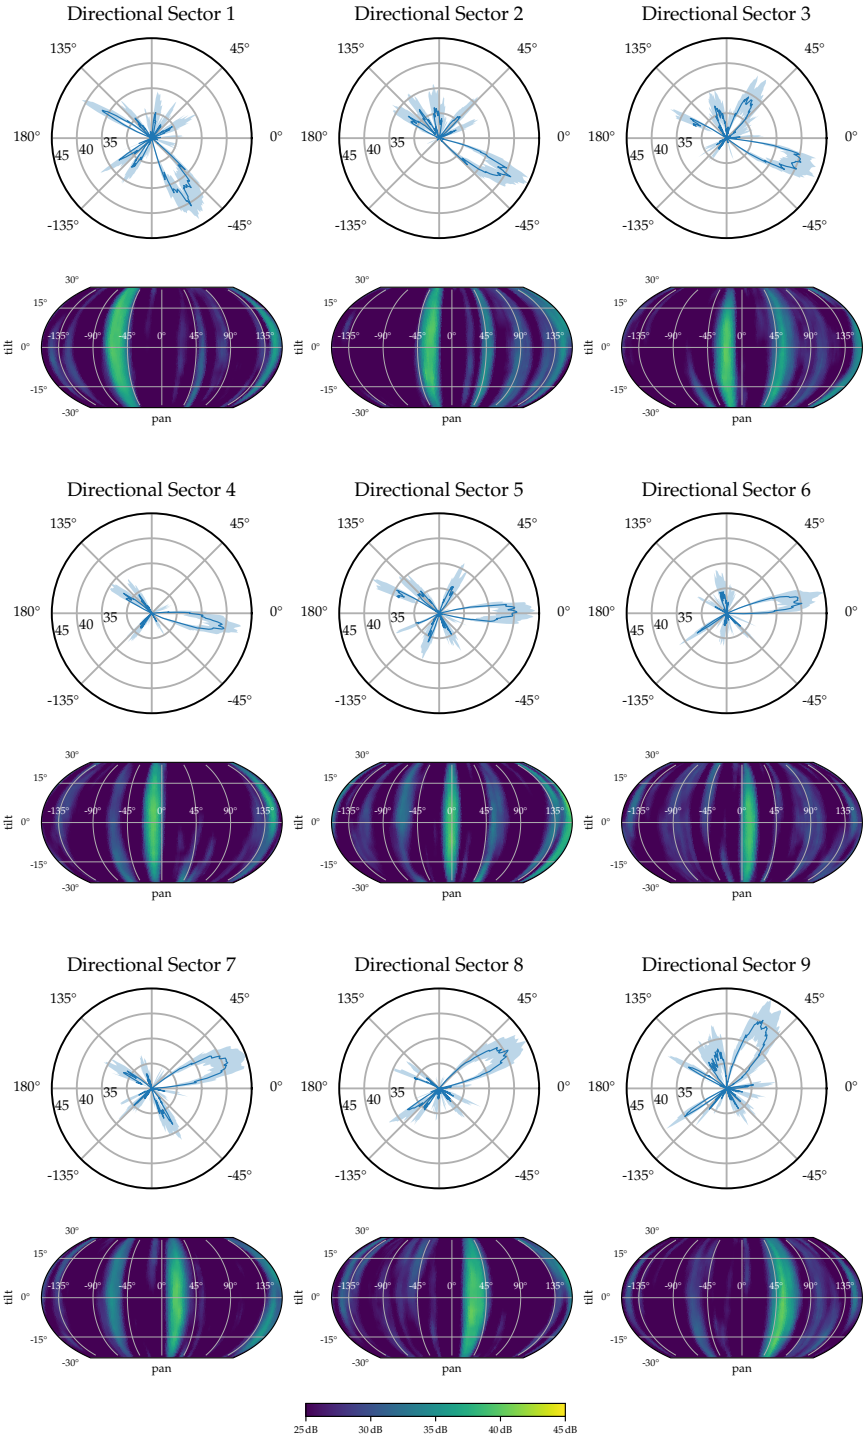

Figure 10.9: Custom directional antenna patterns.

side lobes. Specifically, the gain is about 5.17 dB higher than in the default beams. Under different tilting, the beams show a smooth gradient with low variations. These results confirm that the hardware components in the phased antenna array are capable of generating more precise beam patterns than those defined by the default sectors.

Our beams steer precisely towards the receiver.

10.6 DISCUSSION AND SUMMARY

Our measurements, as presented above, capture the radiation characteristics of the Talon AD7200 with different beams. The default beam patterns feature an irregular shape and expose strong side lobes. Identifying a clear main-lobe is hardly possible. With our control over the antenna weighting network, we isolate individual antenna elements and measure their radiation characteristics as well. Patch antenna elements on the surface of chip expose typical radiation to the front or back. Dipole antennas have lower gains and can be used to fine-tune the beam pattern. Obtaining the array factor with complex antenna gains, we derive custom beam patterns with strong gains in one or multiple directions. Additional measurements approve that our beam patterns provide much higher directionality than the default ones.

We provide precise radiation patterns of default and customized beams.

Despite using an antenna array that is used in many devices, our results are very specific for Talon AD7200. Other devices may mount the antenna at different positions in the device, thus causing different radiation effects. The Netgear Nighthawk R9000, for example, mounts the antenna directly in front of the mainboard and blocks signals towards to back. In such a scenario our measurements are not applicable and likely deviate from the physical characteristics. Nevertheless, we derived our beam patterns from measurements with multiple pairs of devices. Thereby, we confirm that all Talon AD7200 expose similar characteristics. Our beam patterns have general applicability and can be reused for other setups with Talon AD7200 devices.

Other devices may expose different characteristics.

The array factor measurements are highly valuable and enable to fully exploit the beamforming on the Talon AD7200. They provide the foundation to generate codebooks of beam patterns with specific geometric features. For example, the default beam patterns are unlikely to perform well when the device is mounted under the ceiling or in any corner of a room. Beams that are oriented towards the blocking walls cause an unnecessary training overhead. Custom beams derived from the array factor can steer the signal directly to different positions in the room. As a result, the coverage maps to the environment and the beam resolution increases. We highly encourage users to consider custom codebooks when deploying devices in productive environments.

Custom beam patterns can be adjusted to the environment.

The measurements in this chapter partially support our contributions in Chapter 5 and 7. Additionally, we believe that future projects could make use of the generated data as well. To support the commu-

nity and allow others to benefit from our work, we release the relevant data on our public project page[1] (see Section A.6).

1 The measured antenna radiation patterns are available at: https://github.com/seemoo-lab/talon-sector-patterns

11 PRACTICAL DEPLOYMENT

Practical deployments expose sub-optimal characteristics.

The measurements in the previous chapter demonstrate the performance of our directional beams under ideal conditions. To obtain the radiation patterns, transmitter and receiver were close to each other, and no multi-path effects distort the observations. Practical deployments, in contrast, have different characteristics; signals are seldom received in high accuracy from only a single direction. Deployed devices are mostly not oriented face-to-face and stay further away from each other. Additionally, multiple communications may occur at the same time, thus causing interference. In order to evaluate such practical performance aspects of directional beams, we take extensive measurements with different beam configurations in a large-scale deployment with multiple devices.

Section 11.1 describes our large-scale measurement setup. In Section 11.2, we investigate the performance gain of directional beams over the default configuration. Section 11.3 reveals the spatial isolation that is achieved between concurrent transmissions. Section 11.3 details the interference among these transmissions with directional and side lobe suppression beams. Multi-lobe beam patterns that reach multiple devices are evaluated in Section 11.5. Finally, we discuss and summarize our findings in Section 11.6.

11.1 MEASUREMENT SETUP

Our measurement setup consists of 28 devices deployed among three floors.

The atrium of our University and State Library (ULB) on campus features an open-space among multiple floors. This environment is ideal to evaluate parallel mm-wave communication links over longer distances. On each floor, a balustrade separates a corridor with workspaces and book-shelves from the open atrium. It provides an excellent position to deploy our devices since people in the environment can not block signals from there. For our measurements, we occupy the upper three floors and set up 28 Talon AD7200 devices from our testbed experimentation platform as shown in Figure 11.1. While 12 devices are placed in each of the upper two floors, four are located on the lower one. Since our measurements are performed during regular opening hours of the library, we mark all devices with an information sheet and avoid unnecessary disturbances of other visitors as much as possible. To remotely control all devices and measurements, we configure an additional management network in the 2.4 GHz band.

Figure 11.1: Large-scale experiment setup with multiple Talon AD7200 devices in an atrium.

We limit the measurements to specific scenarios.

Running experiments with 28 individual devices means that 756 possible links exist. Each device can pair with 27 others. Measuring all possible combinations with multiple links active at a time would turn out to be very time-consuming. To keep the complexity of our measurements low and finish within a day, we had to limit our scenario to specific scenarios of link combinations. They are split into three groups. First, we consider a group of scenarios with up to four active links on the same floor. In the second, the set of scenarios expands to links spanning two floors. Finally, our third group considers up to 8 active links on all three floors. In the latter two, devices can also establish links on different floors. Our scenarios comprise parallel links and also those that cross each other. In total, they cover different characteristics. Figure 11.2 gives an impression on the placement of devices and the links between them.

After obtaining the CSI, we measure the performance with different beams.

Our measurements for each scenario are performed in multiple steps. First, we estimate the CSI as described in Chapter 7 between all devices on the up- and downlink. This CSI allows devices to compute optimized directional, multi-lobe, and side lobe suppression beam patterns. They derive an optimized directional beam pattern that maximizes the signal strength at their link partners. The multi-lobe beam pattern exposes lobes to all the other devices and the side lobe suppression pattern aims at minimizing the interference at other receiver but the intended one. For each of those beams, the intended receiver and all other devices in the scenario record the SNR. While having all links in the scenario we measure the achievable throughput with default and directional beams.

For all links in these scenarios, we obtain the CSI and SNR of particular beam patterns and also invest the effects of these patterns on other

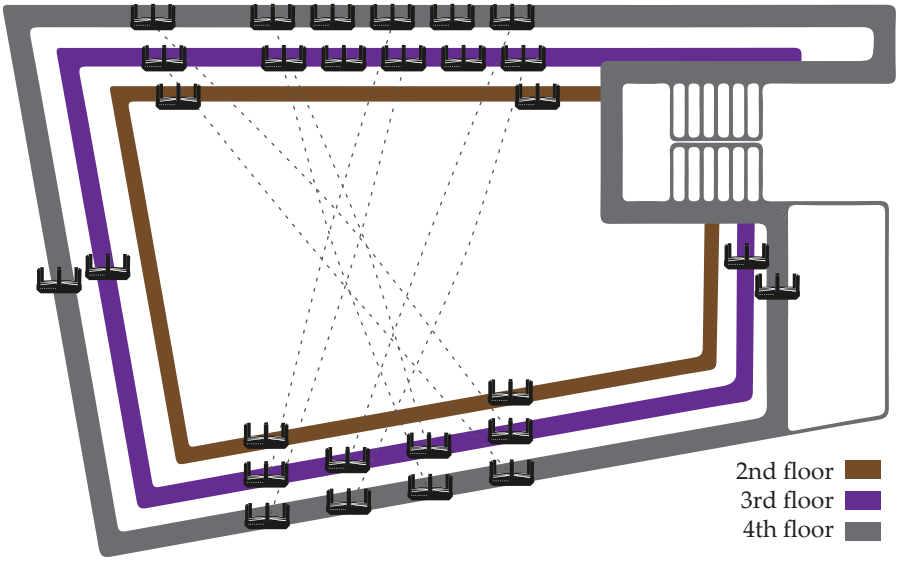

Figure 11.2: Deployment of 28 devices spanning three floors in an atrium for practical measurements. Connections represent an exemplary scenario with eight concurrent links.

devices in the same scenarios. Moreover, we measure the throughput with directional beams as well as with the default ones. All throughput measurements are performed with `iperf3` [iPerf] and repeated three times to achieve sufficient confidence. In total, we obtained the characteristics of 402 links in 61 different scenarios. The following example scenario illustrates the measurements we perform.

An example describes our measurements with three active links.

EXAMPLE SCENARIO. A scenario with three active links may consist of the combinations A->B, C->D, and E->F among the six devices A to F. First, to be able to generate custom beam patterns, we estimate the CSI between all devices. Then, these devices compute a directional beam that follows the link directions. A selects a beam that maximizes the signal strength at B. Similarly, C and E optimize the beam pattern regarding D and F, respectively. To derive a multi-lobe beam, all transmitters A, C, and E, aim to steer the signal towards the receivers B, D, and E simultaneously. They create a beam with a lobe in each transmission direction employing their CSI. For spatial isolation of the links, all transmitters suppress the side lobes towards all but the intended receiver. C creates a beam pattern that maximizes the signal strength at D and simultaneously suppresses the reception at B and F. A and E perform similarly and create a beam towards B and F with low interference on the other links, respectively. Configuring these beams on the transmitters, we record the SNR at each receiver. Additionally, we measure the throughput with traffic on all links simultaneous. Scenarios with more or less active links are evaluated correspondingly.

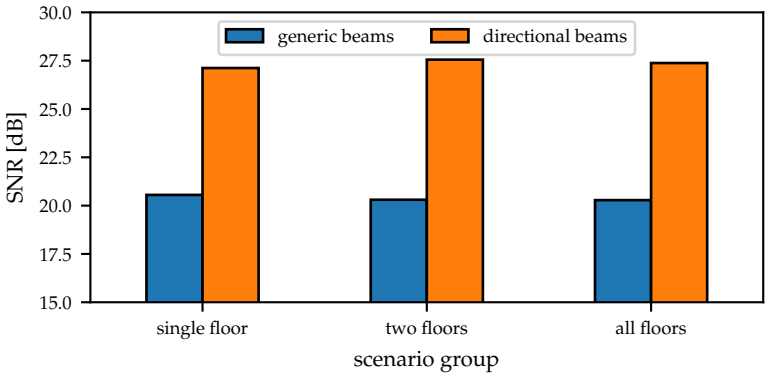

Figure 11.3: Average signal-to-noise ratio in different groups of links.

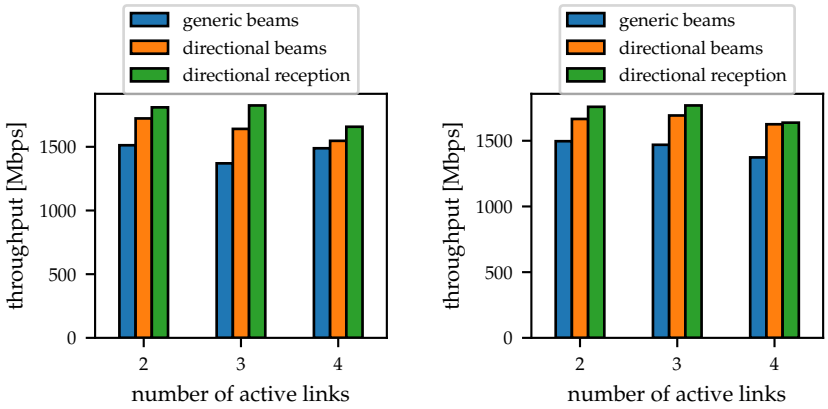

Figure 11.4: Total throughput with up to four parallel transmissions.

11.2 DIRECTIONAL GAINS

Devices select an optimized directional beam pattern.

The first part of our practical experiment considers the directional beam patterns. Similar to our measurements in the anechoic chamber (see Chapter 10), we compare the achievable signal strength with directional beams to those of the default beam patterns. Figure 11.3 illustrates the results averaged over our groups of link scenarios. When considering only links on a single floor, our directional beams provide an SNR of 27.1 dB at the receiver. With links spanning over two or three floors, the average SNRs with 27.6 dB and 27.4 dB does not indicate any significant difference. In contrast to the default beam patterns that achieve an SNR of 20.6 dB, 20.3 dB, and 20.3 dB respectively, the average performance with directional beams is about 6.94 dB higher. This difference is even slightly higher than the results in our previous measurements and indicates the strong practical application.

11.3 SPATIAL ISOLATION

Spatial-reuse requires highly directional beams without interference.

Spatial-reuse is one of the most considerable advantages of directional mm-wave communication. When links become very narrow and obtain wire-like characteristics, they do not distort each other and enable simultaneous interference-free transmissions. How close our directional beams are to this theoretical consideration, we evaluate with parallel throughput measurements on multiple active links. Iterating through the scenarios of different link combinations, our measurements with generic and directional beams reveal the accumulated throughputs shown in Figure 11.4. In scenarios covering a single floor with two active links, the total throughput by using the generic beams is about 1512.3 Mbps. With two and four active links, the throughput does not significantly differ and reaches values of 1370.2 Mbps and 1488.6 Mbps, respectively. Directional beams slightly increase these values by about 12.5 % and achieve a total accumulated throughput on the active links of 1722.9 Mbps, 1640.8 Mbps, and 1547.0 Mbps respectively. Using directional beams not only on the transmitter but also on the receiver, we increase the throughput by about 307.0 Mbps. Similar results are observable with links covering two or three floors, as shown for the former in the graph on the right in Figure 11.4. In all scenarios, directional beams provide a constant performance gain.

Our optimized beam cause interference.

Surprisingly, the measurements do not indicate any spatial-reuse. With perfect link isolation, the accumulated throughput would scale linearly with the number of active links. Instead, the evaluated links share the available resources as they would do in omnidirectional communication systems. The spatial separation between those links is likely too low to operate entirely independent. Before transmitting a frame, the devices sense the channel and only start the transmission when they do not detect any other transmission. While this reduces the collisions of transmitted frames, it leads to unnecessary high congestion back-offs in directional networks. Unfortunately, we could not adjust the sensitivity of the carrier sensing on the Talon AD7200 devices. The only option we have is to decrease the interference on other links using side lobe suppression in transmission beams.

11.4 INTERFERENCE MITIGATION

Side-lobe suppression reduces the interference.

Interference among parallel communication in spatial proximity has vast impacts on the performance of each link. When two devices that are close to each other start transmitting at the same time, they halve their performances. Transmitted frames are time-multiplexed and transmitted on after the other. Each device only obtains access to the channel for half of the time. While this is a common issue in lower frequency bands with omnidirectional communications, directional communications should overcome this limitation when links

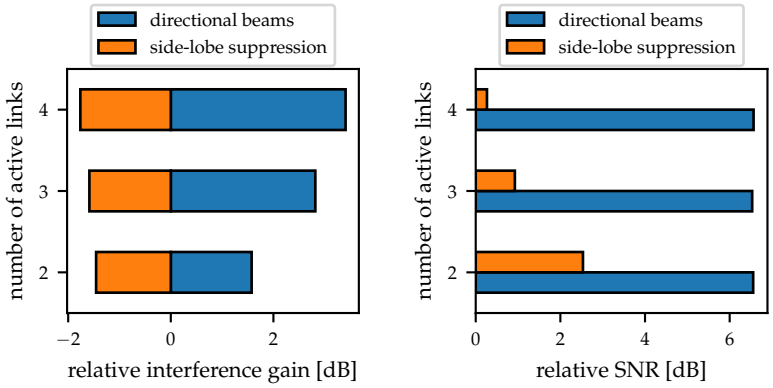

Figure 11.5: Interference and signal gain with directional and side lobe suppression beams relative to the default beam patterns.

are unaffected from each other. Unfortunately, our directional beams not only increase the gain into the intended directions but also introduce additional interference to other receivers. Despite enhancing the performance on individual links, our directional beams limit the overall network performance. Using side lobe suppression in our beam patterns, we decrease the signal strength in the directions of the other devices, thus limit the interference. Figure 11.5 shows the interference of directional and side lobe suppression beams relative to the generic ones for links on a single floor. Directional beams increase the interference of about 1.58 dB, 2.82 dB, and 3.41 dB for 2, 3 and 4 active links, respectively. Thus, it is unsurprising that links distort each other. The side lobe suppression reduce the interference. With two active links, the interference is 1.46 dB lower than with the default beams. Similar differences occur for two and three active links with −1.58 dB and −1.76 dB, respectively.

Spatial isolation in dense deployments is challenging.

Another issue of side lobe suppression is the lower SNR the resulting beam provides in the intended direction. While this difference was small under ideal circumstances in the anechoic chamber, the differences between side lobe suppression and direction beam patterns can take multiple dB. As shown in Figure 11.5 (right), our directional beams perform about 6.55 dB better than the generic beams. With sidelobe-suppression, the gain with two active links is only 2.53 dB and further reduces to 0.92 dB and 0.27 dB for three and four active links. Hence, interference mitigation with side lobe suppressions in such a dense scenario is insufficient to achieve spatial isolation.

11.5 MULTI-LOBE BEAM PATTERNS

Multi-lobe patterns serve multiple devices at the same time and enable multicast transmissions. In our experiments, we obtain the SNR that is achievable with such multicast beams. The lowest SNR of each

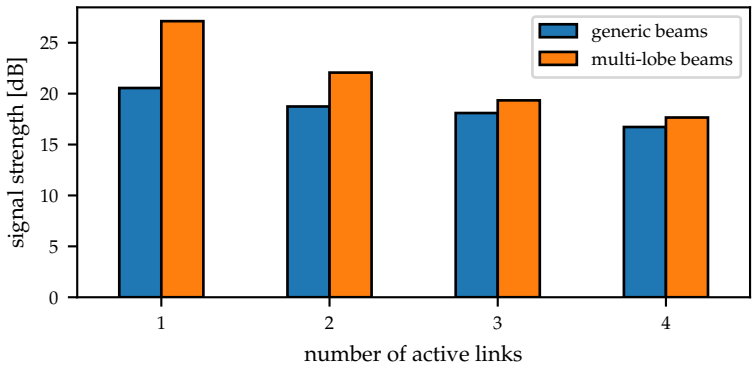

Figure 11.6: Effective multicast signal strength of beam patterns with multiple lobes.

receiver described the performance of the beam as it effectively limits the modulation and data-rate that can be used to reach all receivers. We derive this multicast signal strength for all scenarios within our experiment as shown in Figure 11.6. Generic beam patterns already achieve satisfying results. Most of them expose a broad coverage with good chances to reach multiple receivers. Still with an increasing number of receivers, the provided signal-strength decreases. It reaches about 20.55 dB when serving a single receiver that decreases to 18.74 dB, 18.10 dB, and 16.72 dB for two, three and four receivers. Our optimized multi-lobes beam patterns specifically point into the direction and thus achieve higher precision. The lowest signal strength at all receivers stays higher as with the default beam patterns. Serving one receiver only, the multi-lobe patterns expose a single lobe and correspond to the directional beam patterns. They provide an SNR of 27.12 dB. Beams featuring two lobes obtain an SNR of 22.07 dB, thus performs 3.33 dB better than the generic beam. With three and four lobes the differences are smaller but still increase the signal gains by 1.24 dB and 0.93 dB respectively. Generating beams with few lobes, our mechanism achieves better multi-cast performances than the generic coverage. These results implicate that directional multicasting is possible with specific aligned multi-lobe beam patterns.

Multi-lobe beam patterns enable efficient multicasting.

11.6 DISCUSSION AND SUMMARY

Measurements in a practical deployment demonstrate the advantages of beamforming in mm-wave communication networks. Generic beam patterns predefined on the devices provide decent coverage and facilitate different application scenarios. Still, they leave a high potential for improvements. Our directional beams achieve SNRs that are about 6.94 dB higher. Serving multiple receivers, specific multi-lobe patterns increase the effective SNR for multicasting to two users simultaneously by 3.33 dB. When multiple links are active, our side lobe suppres-

Beam patterns should adapt to the environmental context.

sion beams decrease the interference by 1.46 dB, thus enhancing the spatial separation among concurrent transmissions. Applying directional beams on the transmitter side, we increase the throughput by 12.5 %. Directional beams on both transmitter and receivers antennas exceed the throughput of generic beams by 307.0 Mbps. Despite these remarkable advantages, we do not achieve any spatial-reuse. The deployment of devices is too dense for links to become independent. These results indicate the high importance of cross-layer considerations. Beamforming alone achieves remarkable performance gains. For optimal network operations, it should jointly cooperate with scheduling and carrier sensing algorithms to increase the total performance of practical wireless mm-wave networks.

Custom beam patterns perform better than generic ones.

The measurements in this chapter partially support our contribution in Chapter 7. Additionally, we believe that future projects could make use of the generated data as well. To support the community and allow others to benefit from our work, we release the relevant data on our public project page[1] (see Section A.7). The available data-set contains the SNR, CSI, and throughput measurements at the links covered in our scenarios.

1 The measured link statistics are available at: https://github.com/seemoo-lab/talon-library-measurements

Part VI

DISCUSSION AND CONCLUSIONS

In this part, we discuss our findings, possible applications, and future work in Chapter 12. Afterwards, Chapter 13 concludes this thesis.

12 DISCUSSION AND OUTLOOK

We faced the issue of lacking evaluation platforms.

Having great ideas in mind, we started our research on performance and security aspects in mm-wave communication systems. Unfortunately, our work soon faced the limitations of mm-wave prototyping platforms. Existing solutions, by then, did not provide the desired capabilities. Due to the lack of alternatives, we delved ourselves into operating off-the-shelf devices. In contrast to early prototyping systems, they already provide decent Gigabit-links with integrated antenna arrays. As soon as the first IEEE 802.11ad compatible router, the Talon AD7200, became available on the market, we vowed to get the maximum benefits out of it. This intention resulted in our testbed experimentation platform, which significantly amplified our research. Our platform is flexible and still standard compliant. Obtaining access to the beam training parameters on the devices enables our performance improvements presented in Chapter 5, Chapter 7, and Chapter 6 as well as the security considerations in Chapter 9. Complementing this, our setup for signal propagation analysis additionally allows for investigating reflection and environmental aspects in Chapter 8. Combined, both platforms provide an unprecedented framework that simplifies practical mm-wave research. In the following, we state application scenarios for our framework and outline future work that is worth to address in follow-up research projects.

12.1 APPLICATIONS

Our practical testbed framework has many applications.

Our testbed experimentation framework provides access to the beam training operation and the capability to implement custom beam pattern shapes. Thereby, it allows implementing even more advanced beam training techniques. For example, compressive sector selection and adaptive beam optimization might be combined to reduce the training overhead further while increasing the beam accuracy. Joint training of both, transmit and receive sectors could additionally amplify the directional gain. Carefully selecting the active antennas in the array might decrease the power consumption of transmissions. With the Channel State Information (CSI), a high amount of environment information could be collected. In combination with precise directional probing beams, sensing applications such as gesture recognition and in-band radar technologies are conceivable. Practical measurements in large-scale environments with numerous obstacles should be performed to identify optimal Access Point (AP) and base station deployments. Managed mm-wave Wi-Fi networks tend to use ultra-

dense topologies to serve stations with high data rates. Due to the risk of blockage, planning and maintaining such networks is challenging; they may expose dead spots without coverage. At these locations, a seamless handover to other communication technologies that are less affected by sudden impairments should be considered. Dealing with substantial channel variations also raises the question of optimal adaptation. In contrast to omnidirectional systems, which typically adjust the coding rate and transmission power only, antenna steering and connection handover increase the degree-of-freedom in adaptation algorithms. Devices should not only rely on the signal strength when selecting an AP and communication technology.

With their directional propagation characteristics, mm-wave transmissions are suitable to achieve a high localization precision. Leveraging the angles-of-arrival of links, we exemplary describe two different localization systems that use a single or multiple APs in the following. Following this, we discuss the challenges and opportunities of practical mm-wave applications.

Reflection paths can be used to localize devices.

PSEUDO-LATERATION. User localization in indoor environments with a single AP is still an unsolved challenge. Most existing schemes require multiple infrastructure nodes or prior knowledge of the environment. We propose a pseudo-lateration scheme that exploits environmental reflectors as pseudo anchors to derive a user's location. By combining these directions of reflected signal paths with time-of-flight measurements, we obtain a typical triangulation problem. Our protocol operates in three steps. First, the AP performs a sector level sweep with narrow beams. For each beam, the stations report the received signal strength. Second, they refine their receive beams for the identified signal paths and determine their angular offsets. Processing these measurements, the AP determines the angular directions of the line-of-sight and all reflected signal paths. Finally, both nodes cooperatively measure the time-of-flight to compute the paths' distances. Combining the distance and the angular direction, a user localization becomes possible. We implement the angular estimation of this localization scheme on our channel sounding platform (see Section 3.1). Traces from an indoor environment are further processed in simulations. Our publication [Che+17] provides extended results and a detailed description of our protocol. In summary, pseudo-lateration, with narrow beams and users in proximity, achieves a centimeter scale accuracy.

The default beam patterns allow device triangulation.

LOCALIZATION WITH MULTIPLE APS. The localization scheme in the previous section comes with certain requirements that are challenging for deployments with commodity devices. Time-of-flight measurements are either imprecise or unavailable at all. As seen in Chapter 10, typical beam patterns are highly affected by irregulari-

ties and lack a narrow shape. Taking such technical limitations into account, we propose a second indoor localization scheme that operates with off-the-shelf devices. In particular, we consider a typical IEEE 802.11ad network with multiple APs. Each AP periodically performs the sector level sweep and collects the measured Signal-to-Noise Ratios (SNRs) from all stations in range. Due to a low SNR accuracy, transmission angles cannot be directly obtained. Instead, our localization scheme employs particle filters along with linear programming and uses a Fourier analysis to derive a station's location from measurements of multiple APs. Using our testbed experimentation platform, we implement and evaluate our algorithm in practical office scenarios. Our publication [Bie+18] provides detailed protocol descriptions and extended results. In summary, our system operates in real-time and achieves a sub-meter accuracy in 70 % of the cases. Given the limited hardware capabilities, this constitutes a remarkable result.

CHALLENGES AND OPPORTUNITIES. Emerging mm-wave application scenarios often provide location-based services that require to precisely localize a device. Despite mm-wave communication systems with narrow beams address these demands, practical results are not as accurate as theory expects. Overcoming these limitations, we demonstrate that device localization is feasible. While the first approach is impractical with off-the-shelf devices, it achieves high accuracy gains due to receive beamforming. The second approach only exploits the generic beam patterns that are predefined on the devices. They use a fixed receive beam pattern that provides a quasi-omnidirectional coverage. Extending this localization scheme with custom directional beam patterns and receive beamforming is likely to increase the accuracy. In other words, practical device localization leaves space for various improvements.

Directional beams facilitate localization system.

With our open research platform, we provide an ecosystem for firmware modifications on off-the-shelf mm-wave devices that is not limited to the applications above. Besides using our tools in our research projects described in this work, they already lead to various findings from other researchers. For example, Zhang, Garude, and Pathak [ZGP18] use our environment simulation to evaluate their proactive blockage mitigation technique that uses joint transmissions from multiple APs. Loch et al. [Loc+17] use specific multi-lobe beam patterns to track the movement and rotation of mobile devices with zero overhead. By integrating our simulation environment, they validate their approach in a wide range of scenarios. Using our testbed platform, Assasa et al. [Ass+18] investigate medium access and transport protocol aspects in dense mm-wave networks. In doing so, they reveal practical issues caused by channel contention and frame aggregation on higher layer network performance. A measurement

Our framework is easily extensible and not limited to the presented application.

campaign by Saha et al. [Sah+18] investigates the beam training and rate control characteristics of various off-the-shelf devices. They use our platform to adjust lower layer parameters. While these examples demonstrate a few possibilities of our framework, we expect numerous extensions in future work.

12.2 FUTURE WORK

We give ideas for further feature extensions.

Our research has not yet fully exploited the capabilities of our framework. Especially the features of our testbed experimentation platform could be further extended. With extensive analyses of the firmware running in the IEEE 802.11ad chip, we aim to access the Automatic Gain Control (AGC) and carrier sensing settings. Controlling these two parameters in a realistic setup provides more accurate insights on directional propagation effects and further allows to investigate the limitations of spatial reuse. By enabling frame injection and raw In-phase and Quadrature (IQ) sample processing, off-the-shelf devices could be turned into powerful Software-Defined Radios (SDRs). Facilitating a peer-to-peer mode and simultaneous AP and station mode operations could enable directional device-to-device communications. Analyzing and understanding the proprietary one-wire protocol between the baseband and antenna module of the IEEE 802.11ad Wi-Fi chip could enable to interchange the antennas with other evaluation systems. In particular, this could allow using a cheap off-the-shelf antenna in SDR based prototyping systems.

Academic researcher and engineers from industry should join their forces to shape the future of mm-wave communications.

On the long-term, academic researchers and engineers from industry should work hand-in-hand. Unfortunately, in the domain of mm-wave communications, both areas have drifted apart. The research directions above only address problems that industry has already solved. While industry came up with market ready consumer products—which is a great effort and highly honorable—many research groups still stick to their ancient prototyping systems. Although our contribution narrows the gap between these parties, it is far from being a solution to the stated problem. Every time manufactures release a new product the process of analysis and integration of our tools starts over again. Literary spoken, we fix the symptoms and provide a short-term solution. Industry and academia must combine their forces to succeed in the long-term. Academia comes with great ideas but struggles with practical implementations. Industry solution mostly focuses on standardized specifications. However, what remains left when standardization lacks new ideas? We believe that innovation is an interdisciplinary process that cannot be driven by either party. Manufacturers should find a way to open their products for research without losing their intellectual property. Evaluation platforms should be publicly available and not restricted to those with non-disclosure

agreements. Researchers, in contrast, should focus on the practicality and applicability of their proposed solutions. In the end, this way is beneficial for all and lead to novel mm-wave applications.

13 CONCLUSIONS

In this work, we propose performance and security enhancements for practical mm-wave communication systems. Current mm-wave systems operating in the unlicensed 60 GHz band underutilize the advantages of directional communication. Despite the substantial available spectrum, IEEE 802.11ad does not achieve throughput gains that are significantly higher than those in the latest sub-6 GHz standards. Off-the-shelf devices with integrated phased arrays of dozens of antenna elements exhibit irregular beam pattern with sub-optimal directionality. Due to several strong side lobes, they make spatial-reuse almost impossible. Additionally, state-of-the-art beam training, such as the sector level sweep in IEEE 802.11ad, is imperfect for three reasons. First, all available sectors are frequently probed even though some are unlikely to perform well. Second, the utilized fixed codebook with generic beam patterns does not adapt to the environment. Third, the protocol does not provide any protection against forged feedback. Unfortunately, due to the lack of proper mm-wave evaluation systems, addressing these issues in practical scenarios becomes a challenge.

State-of-the-art beam training is imperfect. It causes a high overhead, uses generic beams, and is vulnerable to forged feedback.

Instead of relying on existing prototyping systems with limited features, our practical studies base on off-the-shelf hardware with custom software modifications. We propose a holistic evaluation framework that consists of two components. First, our channel sounding platform and ray-tracing based simulations allow analyzing the signal propagation in various environments with arbitrary reflectors and blockage. Second, our testbed experimentation platform facilitates evaluations and measurements in practical deployments. Using commodity tri-band wireless routers allows an IEEE 802.11ad compliant operation. We open these devices for research experiments by revealing access to the internal operations of the proprietary Wi-Fi firmware. With modified software components, our platform provides full access to the operating system, device driver, and network interface. Integrating binary patches in the firmware, we uncover specific beam training parameters. Such modifications allow implementing custom beam training protocols with full control over all elements in the phased antenna array. This platform taps the full potential of the resource-restricted hardware components and provides a cost-efficient and scalable evaluation framework for practical standard-compliant experiments.

Our framework enables practical experimentation with off-the-shelf devices.

Firmware modifications allow for beam training customizations.

Using the testbed experimentation platform above, we implement and validate performance improvements for the beam training in IEEE 802.11ad devices. Our compressive sector selection integrates existing path tracking approaches into the sector level sweep. By

We increase the directionality of beams, lower the training overhead, and avoid interference.

only probing a random subset of available sectors, it speeds up the training by a factor of 2.3. Still, it finds the best sector from a given codebook with predefined beam patterns. Since antenna side lobes may cause distortions to other devices, we propose to steer the receiving beam away from any interference direction. Our adaptive beam switching mechanism continuously tracks the interference and signal strength in all the main and side lobes. It identifies a trade-off that maximizes the Signal-to-Interference-plus-Noise Ratio (SINR) and switches the beam whenever necessary. This approach facilitates the parallel operation of incompatible protocols in the same frequency range. By this means, we increase the throughput by about 60 % in practical testbed experiments. Dynamically adapting the beam patterns to the wireless channel conditions maximizes the signal strength at receivers. Our beam optimization derives beams that have a much higher directionality than the default ones. With specific probing beams including constant phase shifts, we obtain the full Channel State Information (CSI). The protocol behind our approach uses this CSI and automatically exploits reflections and destructive interference in the environment. Practical evaluations in different environments show that this approach increases the signal strength by 5 dB on average. In summary, our contributions reduce the training overhead and distortion caused by interference while enhancing the directionality.

Attackers can eavesdrop on reflections and remotely tamper with the beam training.

Directional beams have minimal coverage that is limited by any blockage in the environment. Thus, intercepting an mm-wave link is typically assumed to be more challenging than in omnidirectional communications. Our experimental studies show that, despite a strong directionality, attack vectors exist. In practical experiments, we reveal that environmental objects in the beam may reflect the signal towards the outside of the designed coverage area and facilitate eavesdropping from remote locations. Inconspicuous small-scale objects are sufficient to cause significant reflections. Consumer devices with metal bodies are perfect reflectors and enable an eavesdropper to obtain a signal strength in order of that of the intended receiver. Even the beam training itself can become an attack vector. Attackers can forge the feedback in the sector level sweep and report arbitrary sectors to choose. This beam stealing attack forces victims to steer their beams to random directions. As a result, Denial-of-Service (DoS) and Man-in-the-Middle (MITM) attacks become possible as practically demonstrated in our testbed. To protect against this threat, we discuss countermeasures and propose a simple authentication scheme. With about 7.3 % overhead, our authenticated sector level sweep enforces that responses are accepted from legitimate devices only. These contributions emphasize the threats and security challenges that still exist in directional communication networks. Proper solutions should be considered to enhance the security in future standards.

We outline current threats in mm-wave communication and discuss possible countermeasures.

Emerging mm-wave applications scenarios benefit from our contributions.

Emerging application scenarios, such as information shower and cable replacements, heavily rely on the mm-wave specific propagation effects with high directionality. Unfortunately, current systems and devices do not yet turn these advantages into practice. We believe that our contribution narrows this gap between theory and practice. Both, our adaptive beam optimization and switching mechanisms, enhance the directionality and minimize the interference among multiple devices. They increase the *spatial isolation* and provide a significant enhancement over the standard operation. Compressive sector selection integrates the imperfections of low-cost hardware components into an effective and *efficient beam training* with reduced probing overhead. Protecting the sector level sweep with a simple yet effective authentication scheme, we increase the *attack resistance* of devices concerning beam stealing. Our framework for testbed experimentation and simulation facilitates extensive *practical evaluations*. It allows investigating various aspects of mm-wave communications from a practical perspective. Publicly releasing our framework and measurements to the community, we ease the *reproducibility* of our results. The provided artifacts enable others to adapt our tools to their purposes and to build on our findings. We strongly believe that our testbed platform is a great contribution to the community, simplifies practical mm-wave research experiments, and leads to novel experimental-driven results.

Our tools and measurements allow others to build on our findings.

Part VII

APPENDIX

A

SOFTWARE AND DATA RELEASES

During our research, we engineered various tools that enabled us to implement our systems. To allow others to build on our achievements we release selected software components and measurements. This simplifies reproducing our results and develop custom extensions. In the following, we describe our software and data releases.

A.1 TALON TOOLS

The Talon Tools project consolidates a set of software tools for practical research with commodity IEEE 802.11ad devices. It bases on TP-Link's Talon AD7200, which is the first wireless router that supports the IEEE 802.11ad standard and was released in 2016. Using this platform allows investigating various aspects of 60 GHz mm-wave communications in realistic on-site experiments. With our framework, we support various kinds of measurements and evaluations performed with multiple routers in arbitrary environments.

The Talon Tools repository consolidates multiple contributions that are required to start testbed experimentation with the Talon AD7200 devices. It links to other of our repositories and integrates:

- our OpenWrt system image,
- the Nexmon firmware patching framework for ARC,
- the TPy testbed experiment automation tool, and
- measured antenna radiation patterns.

The Talon Tools project page is available at:
https://www.seemoo.de/talon-tools/

A.2 OPENWRT SYSTEM IMAGE

The released OpenWrt variant contains the source code to compile a Linux Embedded Development Environment (LEDE) system image for the TP-Link Talon AD7200. With the *wil6210* driver and firmware already integrated, it supports to configure the 60 GHz interface in Access Point (AP), managed, and monitor mode. Therewith it allows to easily establish mm-wave links between multiple devices. Through the Linux based operating system, this allows to use the Talon AD7200 routers for arbitrary application scenarios.

The OpenWrt system image for the Talon AD7200 is available at:
https://github.com/seemoo-lab/lede-ad7200

A.3 NEXMON FIRMWARE PATCHING FRAMEWORK FOR ARC

The Nexmon C-based firmware patching framework for Wi-Fi chips enables to write firmware patches and integrate custom features. Our repository hosts the adapted variant for ARC based Wi-Fi chips such as the QCA9500. It allows, for example, extending the beam training operation.

The Nexmon for ARC project is available at:
https://github.com/seemoo-lab/nexmon-arc

A.4 TPY: TESTBED EXPERIMENT AUTOMATION

The TPy project contains scripts to remotely control networked devices, such as the Talon AD7200, to run custom experiments. It essentially consists of two components, a generic node service and a central controller. While the former is designed to run on the network devices, the latter constitutes a central network orchestrator. The node is modularized and supports generic testbeds devices with different features (e.g., IEEE 802.11ad links) and software (e.g., measurement tools such as `iperf` or `ping` and special operating systems such as OpenWrt). The code requires Python3, and a working SSH connection all the devices only.

The source code of TPy is available at:
https://www.seemoo.de/tpy

A.5 ENVIRONMENT SIMULATION WITH MMTRACE

Our mmTrace simulation is a deterministic image-based ray tracing framework for mm-wave propagation developed in MATLAB. It supports the design of mm-wave specific protocols and, in contrast to common statistical models, deals with multiple transceivers. The strengths of mmTrace constitute signal variations at different receivers and interference of multiple transmitters, which are crucial in certain situations. It generates channel impulse responses and determines signal characteristics in arbitrary scenarios.

The source code of the mmTrace simulator is available at:
https://github.com/seemoo-lab/mmTrace

A.6 ANTENNA RADIATION PATTERNS

We measured the beam patterns of the predefined sectors of a Talon AD7200 router inside an anechoic chamber. The measured patterns are provided for validation and integration in other projects. Please note that these results might be device dependent and differ on other hardware components or firmware versions.

The measured antenna radiation patterns are available at:
https://github.com/seemoo-lab/talon-sector-patterns

A.7 PRACTICAL DEPLOYMENT TRACES

Measurements in a practical deployment of the atrium of our library reveal the performance of custom beam patterns. Our traces contain the Channel State Information (CSI) of individual links as well as the received signal strength and throughput measurements with different beam configurations in specific scenarios.

The measured link statistics are available at:
https://github.com/seemoo-lab/talon-library-measurements

AUTHOR'S PUBLICATIONS

CONFERENCE PAPERS

[1] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices." In: *24th Annual International Conference on Mobile Computing and Networking (MobiCom '18)*. New Delhi, India: ACM, Oct. 2018, pp. 241–255. ISBN: 9781450359030. DOI: 10.1145/3241539.3241576 [Pal+18a]. Part of this thesis.

[2] Daniel Steinmetzer, Yimin Yuan, and Matthias Hollick. "Beam-Stealing: Intercepting the Sector Sweep to Launch Man-in-the-Middle Attacks on Wireless IEEE 802.11ad Networks." In: *11th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '18)*. Stockholm, Sweden: ACM, June 2018, pp. 12–22. ISBN: 9781450357319. DOI: 10.1145/3212480.3212499 [SYH18]. Part of this thesis.

[3] Swetank Kumar Saha, Hany Assasa, Adrian Loch, Naveen Muralidhar Prakash, Roshan Shyamsunder, Shivang Aggarwal, Daniel Steinmetzer, Dimitrios Koutsonikolas, Joerg Widmer, and Matthias Hollick. "Fast and Infuriating : Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware." In: *15th International Conference on Sensing, Communication, and Networking (SECON '18)*. Hong Kong: IEEE, June 2018. ISBN: 9781538642818. DOI: 10.1109/SAHCN.2018.8397123 [Sah+18].

[4] Guillermo Bielsa, Joan Palacios, Adrian Loch, Daniel Steinmetzer, Paolo Casari, and Joerg Widmer. "Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points." In: *International Conference on Computer Communications (INFOCOM '18)*. Honolulu, HI, USA: IEEE, Apr. 2018, pp. 2384–2392. ISBN: 978-1-5386-4128-6. DOI: 10.1109/INFOCOM.2018.8486232 [Bie+18]. Summarized in this thesis.

[5] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices." In: *13th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '17)*. Incheon, Republic of Korea: ACM, Dec. 2017, pp. 414–425. ISBN: 9781450354226. DOI: 10.1145/3143361.3143384 [Ste+17b]. Part of this thesis.

[6] Matthias Schulz, Francesco Gringoli, Daniel Steinmetzer, Michael Koch, and Matthias Hollick. "Massive Reactive Smartphone-based Jamming using Arbitrary Waveforms and Adaptive Power Control." In: *10th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '17)*. Boston, MA, USA: ACM, July 2017, pp. 111–121. ISBN: 9781450350846. DOI: 10.1145/3098243.3098253 [Sch+17]. Best Paper Award.

[7] Joe Chen, Daniel Steinmetzer, Jiska Classen, Edward W. Knightly, and Matthias Hollick. "Pseudo Lateration: Millimeter-wave Localization using a Single RF Chain." In: *Wireless Communications and Networking Conference (WCNC '17)*. San Francisco, CA, USA: IEEE, Mar. 2017. ISBN: 9781509041831. DOI: 10.1109/WCNC.2017.7925882 [Che+17]. Summarized in this thesis.

[8] Daniel Steinmetzer, Joe Chen, Jiska Classen, Edward W. Knightly, and Matthias Hollick. "Eavesdropping with Periscopes: Experimental Security Analysis of Highly Directional Millimeter Waves." In: *Conference on Communications and Network Security (CNS '15)*. Florence, Italy: IEEE, Sept. 2015, pp. 335–343. ISBN: 9781467378765. DOI: 10.1109/CNS.2015.7346844 [Ste+15]. Part of this thesis.

[9] Daniel Steinmetzer, Matthias Schulz, and Matthias Hollick. "Lock-picking Physical Layer Key Exchange: Weak Adversary Models Invite the Thief." In: *8th ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec '15)*. New York City, NY, USA: ACM, June 2015. ISBN: 978-1-4503-3623-9. DOI: 10.1145/2766498.2766514 [SSH15]. Best Paper Award.

WORKSHOP PAPERS

[10] Daniel Steinmetzer, Milan Stute, and Matthias Hollick. "TPy: A Lightweight Framework for Agile Distributed Network Experiments." In: *12th International Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH '18)*. New Delhi, India: ACM, Nov. 2018, pp. 38–45. ISBN: 978-1-4503-5930-6. DOI: 10.1145/3267204.3267214 [SSH18]. Part of this thesis.

[11] Daniel Steinmetzer, Saad Ahmad, Nikolaos A. Anagnostopoulos, Matthias Hollick, and Stefan Katzenbeisser. "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks." In: *2nd Workshop on Millimeter Wave Networks and Sensing Systems (mmNets '18)*. New Delhi, India: ACM, Oct. 2018, pp. 3–8. ISBN: 978-1-4503-5928-3. DOI: 10.1145/3264492.3264494 [Ste+18]. Part of this thesis.

[12] Daniel Steinmetzer, Adrian Loch, Amanda García-García, Joerg Widmer, and Matthias Hollick. "Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks." In: *1st Workshop on Millimeter-Wave Networks and Sensing Systems (mmNets '17)*. Snowbird, UT, USA: ACM, Dec. 2017, pp. 29–34. ISBN: 978-1-4503-5143-0. DOI: 10.1145/3130242.3130244 [Ste+17a]. Part of this thesis.

[13] Jiska Classen, Daniel Steinmetzer, and Matthias Hollick. "Opportunities and Pitfalls in Securing Visible Light Communication on the Physical Layer." In: *3rd Workshop on Visible Light Communication Systems (VLCS '16)*. New York City, NY, USA: ACM, Oct. 2016, pp. 19–24. ISBN: 9781450342537. DOI: 10.1145/2981548.2981551 [CSH16].

[14] Daniel Steinmetzer, Jiska Classen, and Matthias Hollick. "mmTrace: Modeling Millimeter-wave Indoor Propagation with Image-based Ray-tracing." In: *1st Millimeter-wave Networking Workshop (mmNet '16)*. San Francisco, CA, USA: IEEE, Apr. 2016, pp. 429–434. ISBN: 9781467399555. DOI: 10.1109/INFCOMW.2016.7562115 [SCH16b]. Part of this thesis.

[15] Jiska Classen, Joe Chen, Daniel Steinmetzer, Matthias Hollick, and Edward W. Knightly. "The Spy Next Door: Eavesdropping on High Throughput Visible Light Communications." In: *2nd International Workshop on Visible Light Communications Systems, VLCS '15*. Paris, France: ACM, Sept. 2015, pp. 9–14. ISBN: 9781450337021. DOI: 10.1145/2801073.2801075 [Cla+15].

JOURNAL ARTICLES

[16] Mikhail Fomichev, Flor Álvarez, Daniel Steinmetzer, Paul Gardner-Stephen, and Matthias Hollick. "Survey and Systematization of Secure Device Pairing." In: *IEEE Communications Surveys and Tutorials* 20.1 (2017), pp. 517–550. ISSN: 1553877X. DOI: 10.1109/COMST.2017.2748278 [Fom+17].

POSTERS AND DEMONSTRATIONS

[17] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. "Demo: Channel Estimation and Custom Beamforming on the 60 GHz TP-Link Talon AD7200 Router." In: *12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH '18)*. New Delhi, India: ACM, Nov. 2018, pp. 10–11. ISBN: 9781450359306. DOI: 10.1145/3267204.3268070 [Pal+18c]. Best Demo Award. Part of this thesis.

[18] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. "A Practical IEEE 802.11ad Research Platform: The Hidden Potential of Off-the-Shelf Devices." In: *3rd NSF Millimeter-Wave Research Coordination Network (RCN) Workshop*. Tucson, AZ, USA: NSF, Jan. 2018 [SWH18]. Part of this thesis.

[19] Tobias Schultes, Markus Grau, Daniel Steinmetzer, and Matthias Hollick. "Demo: Far Away and Yet Nearby - a Framework for Practical Distance Fraud on Proximity Services for Mobile Devices." In: *9th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '16)*. Darmstadt, Germany: ACM, July 2016, pp. 205–207. ISBN: 9781450342704. DOI: 10.1145/2939918.2942416 [Sch+16].

[20] Daniel Steinmetzer, Jiska Classen, and Matthias Hollick. "Exploring Millimeter-Wave Network Scenarios with Ray-tracing based Simulations in mmTrace." In: *International Conference on Computer Communications Workshops (INFOCOM WKSHPS '16)*. San Francisco, CA, USA: IEEE, Apr. 2016. ISBN: 978-1-4673-9955-5. DOI: 10.1109/INFCOMW.2016.7562269 [SCH16a]. Summarized in this thesis.

TECHNICAL REPORTS (NON PEER-REVIEWED)

[21] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. *Addendum to Adaptive Codebook Optimization for Beam Training on Off-The-Shelf IEEE 802.11ad Devices*. Tech. rep. TR-IMDEA-Networks-2018-1. IMDEA Networks, July 2018 [Pal+18b].

CURRICULUM VITÆ

Personal Information

Name	Daniel Steinmetzer
Date of Birth	February 2, 1987
Place of Birth	Northeim, Germany
Nationality	German

Education

2014 – 2019 **Doctor of Engineering**
Computer Science, Technische Universiät Darmstadt, Darmstadt, Germany

2012 – 2014 **Master of Science**
Information Systems Technology, Technische Universität Darmstadt, Darmstadt, Germany

2008 – 2012 **Bachelor of Science**
Information Systems Technology, Technische Universität Darmstadt, Darmstadt, Germany

2006 – 2008 **IT Specialist in System Integration**
Certificate of Chamber of Industry and Commerce, Clausthal University of Technology, Clausthal-Zellerfeld, Germany

1999 – 2006 **General Higher Education Entrance Qualification**
Ernst-Moritz-Arndt Gymnasium, Herzberg am Harz, Germany

Work Experience

2014 – 2019 **Research Associate**
Secure Mobile Networking Lab, Technische Universität Darmstadt, Darmstadt, Germany.

2013 – 2014 **Student Research Assistant**
Secure Mobile Networking Lab, Technische Universität Darmstadt, Darmstadt, Germany.

2011 – 2012 **Student Research Assistant**
System Security Lab, Technische Universität Darmstadt, Darmstadt, Germany.

2009 – 2011 **Student Assistant for IT Administration**
Department of Management and Logistics, Technische Universität Darmstadt, Darmstadt, Germany.

2008 – 2008 **IT Specialist in System Integration**
Clausthal University of Technology, Clausthal-Zellerfeld, Germany.

Awards

Publication **Best Demo Award at ACM WiNTECH 2018**
Paper: "Channel Estimation and Custom Beamforming on the 60 GHz TP-Link Talon AD7200 Router" [Pal+18c].

Publication **Best Paper Award at ACM WiSec 2017**
Paper: "Massive Reactive Smartphone-Based Jamming using Arbitrary Waveforms and Adaptive Power Control" [Sch+17].

Publication **Best Paper Award at ACM WiSec 2015**
Paper: "Lockpicking Physical Layer Key Exchange: Weak Adversary Models Invite the Thief" [SSH15].

Thesis **Advancement Award in IT Security 2014**
For the third-ranked master thesis by the Competence Center for Applied Security Technology, CAST e.V.

Studies **Deutschlandstipendium 2012/13 and 2013/14**
Outstanding achievement award with financial and non-material support for talented and dedicated students by the German Federal Government and private sponsors.

Thesis **Advancement Award in IT Security 2012**
For the third-ranked bachelor thesis by the Competence Center for Applied Security Technology, CAST e.V.

Supervised Master and Bachelor Theses

B.Sc. Thesis **Fabian Franke** "Learning the Beams: Applying Evolutionary Algorithms for Optimized IEEE 802.11ad Beam Training".

M.Sc. Thesis **Saad Ahmad** "Using Physical Unclonable Functions (PUFs) for Data-Link Layer Authenticity Verification to Mitigate Attacks on IEEE 802.11ad Beam Training".

B.Sc. Thesis **Damir Mehmedovic** "Wi-Fi based Key Exchange on Android Smartphones".

M.Sc. Thesis **Yimin Yuan** "Investigating Practical Man-in-the-Middle Network Attacks on IEEE 802.11ad".

M.Sc. Thesis **Abdul Hannan** "Neighbor Discovery and Maintenance under Mobility in mmWave-based Mesh Networks".

B.Sc. Thesis **Steffen Kreis** "Unified Multi-modal Secure Device Pairing for Infrastructure and Ad-hoc Networks".

M.Sc. Thesis **Johannes Löffler** "Efficient Neighbour Discovery with Adjustable Antenna Beams in Highly Directional 60 GHz Mesh Networks".

Miscellaneous

Organization Breakout Session Discussion Leader at the 2nd NSF mmW RCN Workshop in Tucson, Arizona, USA.

Organization Publication and Registration Chair of the 9th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2016), Darmstadt, Germany.

Teaching Organization, planing and supervision of the practical student course "Secure Mobile Networking Lab Exercise and Project".

Teaching Organization, planing and supervision of the student seminar course "Seminar and Advanced Seminar on Networking, Security, Mobility, and Wireless Communications".

Teaching Preparation and teaching of selected topics on key agreement schemes as parts of the annual lecture series "Physical Layer Security in Wireless Systems (PhySec)".

Teaching Preparation and teaching of selected topics on mm-wave communication systems as part of the annual lecture series "Secure Mobile Systems (SeMoSy)".

Darmstadt, 28. Januar 2019

BIBLIOGRAPHY

[Ahm+09] Javad Ahmadi-Shokouh, Sima Noghanian, Ekram Hossain, Majid Ostadrahimi, and James Dietrich. "Reflection Coefficient Measurement for House Flooring Materials at 57-64 GHz." In: *Global Communications Conference (GLOBECOM '09)*. Honolulu, HI, USA: IEEE, Mar. 2009. ISBN: 978-1-4244-4148-8. DOI: 10.1109/GLOCOM.2009.5425497.

[AH17] Anum Ali and Robert W. Heath. "Compressed beam-selection in millimeterwave systems with out-of-band partial support information." In: *International Conference on Acoustics, Speech and Signal Processing (ICASSP '17)*. New Orleans, LA, USA: IEEE, Mar. 2017, pp. 3499–3503. ISBN: 978-1-5090-4117-6. DOI: 10.1109/ICASSP.2017.7952807.

[Alk+14a] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W. Heath. "Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems." In: *IEEE Journal of Selected Topics in Signal Processing* 8.5 (Oct. 2014), pp. 831–846. ISSN: 1932-4553. DOI: 10.1109/JSTSP.2014.2334278.

[Alk+14b] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W. Heath. "Single-sided Adaptive Estimation of Multi-path Millimeter Wave Channels." In: *15th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC '14)*. Toronto, ON, Canada: IEEE, June 2014, pp. 125–129. ISBN: 978-1-4799-4903-8. DOI: 10.1109/SPAWC.2014.6941330.

[AH16] Ahmed Alkhateeb and Robert W. Heath. "Frequency Selective Hybrid Precoding for Limited Feedback Millimeter Wave Systems." In: *IEEE Transactions on Communications* 64.5 (May 2016), pp. 1801–1818. ISSN: 0090-6778. DOI: 10.1109/TCOMM.2016.2549517.

[An+09] Xueli An, Chin-Sean Sum, R. Venkatesha Prasad, Junyi Wang, Zhou Lan, Jing Wang, Ramin Hekmat, Hiroshi Harada, and Ignas Niemegeers. "Beam Switching Support to Resolve Link-blockage Problem in 60 GHz WPANs." In: *20th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC '09)*. Tokyo, Japan: IEEE, Sept. 2009, pp. 390–394. ISBN: 978-1-4244-5122-7. DOI: 10.1109/PIMRC.2009.5449837.

[Ans+15] Junaid Ansari, Nikos Perpinias, Alexander Nahring, Petri Mahonen, and Marina Petrova. "Empirical Characterization of mm-Wave Communication Links in Realistic Indoor Scenarios." In: *Wireless Communications and Networking Conference (WCNC '15)*. IEEE, Mar. 2015, pp. 1799–1804. ISBN: 978-1-4799-8406-0. DOI: 10.1109/WCNC.2015.7127741.

[Ara+14] Daniel C. Araújo, André L. F. de Almeida, Johan Axnäs, and João C. M. Mota. "Channel Estimation for Millimeter-Wave Very-Large MIMO Systems." In: *22nd European Signal Processing Conference (EUSIPCO '14)* (Sept. 2014).

[Ass+18] Hany Assasa, Swetank Kumar Saha, Adrian Loch, Dimitrios Koutsonikolas, and Joerg Widmer. "Medium Access and Transport Protocol Aspects in Practical 802.11 ad Networks." In: *19th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM '18)*. Chania, Greece: IEEE, June 2018, pp. 1–11. ISBN: 978-1-5386-4725-7. DOI: 10.1109/WoWMoM.2018.8449795.

[AW16] Hany Assasa and Joerg Widmer. "Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3." In: *Workshop on ns-3 (WNS3 '16)*. Seattle, WA, USA: ACM, June 2016, pp. 57–64. ISBN: 9781450342162. DOI: 10.1145/2915371.2915377.

[Aya+12a] Omar El Ayach, Robert W. Heath, Shadi Abu-Surra, Sridhar Rajagopal, and Zhouyue Pi. "Low Complexity Precoding for Large Millimeter Wave MIMO Systems." In: *International Conference on Communications (ICC '12)*. Ottawa, ON, Canada: IEEE, June 2012, pp. 3724–3729. ISBN: 978-1-4577-2053-6. DOI: 10.1109/ICC.2012.6363634.

[Aya+12b] Omar El Ayach, Robert W Heath, Shadi Abu-Surra, Sridhar Rajagopal, and Zhouyue Pi. "The Capacity Optimality of Beam Steering in Large Millimeter Wave MIMO Systems." In: *13th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC '12)*. Cesme, Turkey: IEEE, June 2012, pp. 100–104. ISBN: 978-1-4673-0970-7. DOI: 10.1109/SPAWC.2012.6292865.

[Baj+10] Waheed U. Bajwa, Jarvis Haupt, Akbar M. Sayeed, and Robert Nowak. "Compressed Channel Sensing: A New Approach to Estimating Sparse Multipath Channels." In: *Proceedings of the IEEE* 98.6 (June 2010), pp. 1058–1076. ISSN: 0018-9219. DOI: 10.1109/JPROC.2010.2042415.

[BR06] Joao Barros and Miguel R. D. Rodrigues. "Secrecy Capacity of Wireless Channels." In: *International Symposium on Information Theory* (July 2006), pp. 356–360. DOI: 10.1109/ISIT.2006.261613.

[Bas+09] Michael Bass, Casimer DeCusatis, Jay M. Enoch, Vasudevan Lakshminarayanan, Guifang Li, Carolyn MacDonald, Virendra N. Mahajan, and Eric Van Stryland. *Handbook of Optics, Volume 1: Geometrical and Physical Optics, Polarized Light, Components and Instruments*. 3rd Editio. McGraw-Hill Education, 2009. ISBN: 978-0-07-162925-6.

[Ber+14] Daniel S. Berger, Francesco Gringoli, Nicolò Facchi, Ivan Martinovic, and Jens Schmitt. "Gaining Insight on Friendly Jamming in a Real-world IEEE 802.11 Network." In: *7th Conference on Security and Privacy in Wireless & Mobile Networks (WiSec '14)*. Oxford, United Kingdom: ACM, July 2014, pp. 105–116. ISBN: 9781450329729. DOI: 10.1145/2627393.2627403.

[Ber+16] Daniel S. Berger, Francesco Gringoli, Nicolo Facchi, Ivan Martinovic, and Jens B Schmitt. "Friendly Jamming on Access Points: Analysis and Real-World Measurements." In: *IEEE Transactions on Wireless Communications* 15.9 (Sept. 2016), pp. 6189–6202. ISSN: 1536-1276. DOI: 10.1109/TWC.2016.2581165.

[Bie+18] Guillermo Bielsa, Joan Palacios, Adrian Loch, Daniel Steinmetzer, Paolo Casari, and Joerg Widmer. "Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points." In: *International Conference on Computer Communications (INFOCOM '18)*. Honolulu, HI, USA: IEEE, Apr. 2018, pp. 2384–2392. ISBN: 978-1-5386-4128-6. DOI: 10.1109/INFOCOM.2018.8486232.

[Che+17] Joe Chen, Daniel Steinmetzer, Jiska Classen, Edward W. Knightly, and Matthias Hollick. "Pseudo Lateration: Millimeter-wave Localization using a Single RF Chain." In: *Wireless Communications and Networking Conference (WCNC '17)*. San Francisco, CA, USA: IEEE, Mar. 2017. ISBN: 9781509041831. DOI: 10.1109/WCNC.2017.7925882.

[Cho15] Jinho Choi. "Beam Selection in mm-Wave Multiuser MIMO Systems Using Compressive Sensing." In: *IEEE Transactions on Communications* 63.8 (Aug. 2015), pp. 2936–2947. ISSN: 0090-6778. DOI: 10.1109/TCOMM.2015.2449860.

[Cis18] Cisco. *Cisco Visual Networking Index: Forecast and Trends, 2017–2022*. Tech. rep. CISCO, 2018, pp. 1–38.

[Cla+15] Jiska Classen, Joe Chen, Daniel Steinmetzer, Matthias Hollick, and Edward W. Knightly. "The Spy Next Door: Eavesdropping on High Throughput Visible Light Communications." In: *2nd International Workshop on Visible Light Communications Systems, VLCS '15*. Paris, France:

ACM, Sept. 2015, pp. 9–14. ISBN: 9781450337021. DOI: 10.1145/2801073.2801075.

[CSH16] Jiska Classen, Daniel Steinmetzer, and Matthias Hollick. "Opportunities and Pitfalls in Securing Visible Light Communication on the Physical Layer." In: *3rd Workshop on Visible Light Communication Systems (VLCS '16)*. New York City, NY, USA: ACM, Oct. 2016, pp. 19–24. ISBN: 9781450342537. DOI: 10.1145/2981548.2981551.

[Dai+13] Hong-Ning Dai, Qiu Wang, Dong Li, and Raymond Chi-Wing Wong. "On Eavesdropping Attacks in Wireless Sensor Networks with Directional Antennas." In: *International Journal of Distributed Sensor Networks* 9.8 (Aug. 2013), p. 760834. ISSN: 1550-1477. DOI: 10.1155/2013/760834.

[Dan12] Quynh Dang. *Recommendation for Applications Using Approved Hash Algorithms*. Gaithersburg, MD, 2012. DOI: 10.6028/NIST.SP.800-107r1.

[Dha15] Aditya Dhananjay. "Iris: Mitigating Phase Noise in Millimeter Wave OFDM Systems." PhD thesis. New York University, 2015.

[Dua+15] Qiyou Duan, Taejoon Kim, Huang Huang, Kunpeng Liu, and Guangjian Wang. "AoD and AoA Tracking with Directional Sounding Beam Design for Millimeter Wave MIMO Systems." In: *26th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC '15)*. Hong Kong, China: IEEE, Aug. 2015, pp. 2271–2276. ISBN: 978-1-4673-6782-0. DOI: 10.1109/PIMRC.2015.7343676.

[EC16] Alecsander Eitan and Carlos Cordeiro. *Short SSW Format for 11ay (IEEE 802.11-16/0416-01-00)*. 2016.

[Elt+16] Mohammed E. Eltayeb, Junil Choi, Tareq Y. Al-Naffouri, and Robert W. Heath. "On the Security of Millimeter Wave Vehicular Communication Systems Using Random Antenna Subsets." In: *84th Vehicular Technology Conference (VTC-Fall '16)* (Sept. 2016), pp. 1–5. DOI: 10.1109/VTCFall.2016.7881128.

[Fom+17] Mikhail Fomichev, Flor Álvarez, Daniel Steinmetzer, Paul Gardner-Stephen, and Matthias Hollick. "Survey and Systematization of Secure Device Pairing." In: *IEEE Communications Surveys and Tutorials* 20.1 (2017), pp. 517–550. ISSN: 1553877X. DOI: 10.1109/COMST.2017.2748278.

[FY10] Michael A. Forman and Derek Young. "The Generation of Shared Cryptographic Keys Through Half Duplex Channel Impulse Response Estimation at 60 GHz." In: *International Conference on Electromagnetics in Advanced Applications*. Sydney, NSW, Australia: IEEE, Sept. 2010, pp. 627–630. ISBN: 978-1-4244-7366-3. DOI: 10.1109/ICEAA.2010.5652964.

[For+95] Steven J. Fortune, David M. Gay, Brian W. Kernighan, Orlando Landron, Reinaldo A. Valenzuela, and Margaret H. Wright. "WISE Design of Indoor Wireless Systems: Practical Computation and Optimization." In: *IEEE Computational Science and Engineering* 2.1 (1995), pp. 58–68. ISSN: 10709924. DOI: 10.1109/99.372944.

[Fre13] Louis E. Frenzel. "Millimeter Waves Will Expand the Wireless Future." In: *Electronic Design* 04/2013 (2013), pp. 30–36.

[GDW16] Zhen Gao, Linglong Dai, and Zhaocheng Wang. "Channel Estimation for mmWave Massive MIMO Based Access and Backhaul in Ultra-dense Network." In: *International Conference on Communications (ICC '16)*. Kuala Lumpur, Malaysia: IEEE, May 2016, pp. 1–6. ISBN: 978-1-4799-6664-6. DOI: 10.1109/ICC.2016.7511578.

[Gao+16] Zhen Gao, Chen Hu, Linglong Dai, and Zhaocheng Wang. "Channel Estimation for Millimeter-Wave Massive MIMO With Hybrid Precoding Over Frequency-Selective Fading Channels." In: *IEEE Communications Letters* 20.6 (June 2016), pp. 1259–1262. ISSN: 1089-7798. DOI: 10.1109/LCOMM.2016.2555299.

[Gen+10] Zulkuf Genc, Umar H. Rizvi, Ertan Onur, and Ignas Niemegeers. "Robust 60 GHz Indoor Connectivity: Is It Possible with Reflections?" In: *71st Vehicular Technology Conference (VTC '10)*. Taipei, Taiwan: IEEE, May 2010, pp. 1–5. ISBN: 978-1-4244-2518-1. DOI: 10.1109/VETECS.2010.5493722.

[Gha+17] Yasaman Ghasempour, Claudio R. C. M. da Silva, Carlos Cordeiro, and Edward W. Knightly. "IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi." In: *IEEE Communications Magazine* 55.12 (Dec. 2017), pp. 186–192. ISSN: 0163-6804. DOI: 10.1109/MCOM.2017.1700393.

[HK16] Muhammad Kumail Haider and Edward W. Knightly. "Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs." In: *17th International Symposium on Mobile Ad Hoc Networking and Computing*

(MobiHoc '16). Paderborn, Germany: ACM, 2016, pp. 61–70. ISBN: 9781450341844. DOI: 10.1145/2942358.2942380.

[HK18] Muhammad Kumail Haider and Edward W. Knightly. "iTrack: Tracking Indicator LEDs on APs to Bootstrap mmWave Beam Acquisition and Steering." In: *19th International Workshop on Mobile Computing Systems & Applications (HotMobile '18)* (Feb. 2018), pp. 107–112. DOI: 10.1145/3177102.3177105.

[Hal+11] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. "Tool Release: Gathering 802.11n Traces with Channel State Information." In: *ACM SIGCOMM Computer Communication Review* 41.1 (Jan. 2011), p. 53. ISSN: 01464833. DOI: 10.1145/1925861.1925870.

[Hur+13] Sooyoung Hur, Taejoon Kim, David J. Love, James V. Krogmeier, Timothy A. Thomas, and Amitava Ghosh. "Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks." In: *IEEE Transactions on Communications* 61.10 (Oct. 2013), pp. 4391–4403. ISSN: 0090-6778. DOI: 10.1109/TCOMM.2013.090513.120848.

[IEE04] IEEE 802.11 Working Group. *IEEE P802.11 Wireless LANs TGn Channel Models*. Tech. rep. IEEE 802.11-03/940r4. 2004.

[IEE14] IEEE Standards Association. *IEEE Std 802.11ad-2012: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. ISO/IEC/IEEE 8802-11:2012/Amd.3:2014(E)*. 2014.

[iPerf] *iPerf - The ultimate speed test tool for TCP, UDP and SCTP*. URL: https://iperf.fr.

[Jac+12] Martin Jacob, Sebastian Priebe, Robert Dickhoff, Thomas Kleine-Ostmann, Thorsten Schrader, and Thomas Kurner. "Diffraction in mm and Sub-mm Wave Indoor Propagation Channels." In: *IEEE Transactions on Microwave Theory and Techniques* 60.3 (Mar. 2012), pp. 833–844. ISSN: 0018-9480. DOI: 10.1109/TMTT.2011.2178859.

[Pyro] Irmen de Jong. *Pyro - Python Remote Objects - 4.60*. URL: https://pythonhosted.org/Pyro4/.

[Jun+09] Junyi Wang, Zhou Lan, Chang-woo Pyo, Tuncer Baykas, Chin-sean Sum, M.A. Rahman, Jing Gao, Ryuhei Funada, Fumihide Kojima, Hiroshi Harada, and Shuzo Kato. "Beam Codebook Based Beamforming Protocol for Multi-Gbps Millimeter-wave WPAN Systems." In: *IEEE Journal on Selected Areas in Communications* 27.8 (Oct. 2009),

pp. 1390–1399. ISSN: 0733-8716. DOI: 10.1109/JSAC.2009.091009.

[KHK17] Meejoung Kim, Eenjun Hwang, and Jeong-Nyeo Kim. "Analysis of Eavesdropping Attack in mmWave-based WPANs with Directional Antennas." In: *Wireless Networks* 23.2 (Feb. 2017), pp. 355–369. ISSN: 1022-0038. DOI: 10.1007/s11276-015-1160-4.

[Kle+12] T. Kleine-Ostmann, M. Jacob, S. Priebe, R. Dickhoff, T. Schrader, and T. Kurner. "Diffraction Measurements at 60 GHz and 300 GHz for Modeling of Future THz Communication Systems." In: *37th International Conference on Infrared, Millimeter, and Terahertz Waves* (Sept. 2012), pp. 1–2. DOI: 10.1109/IRMMW-THz.2012.6380411.

[KGH15] Preeti Kumari, Nuria Gonzalez-Prelcic, and Robert W. Heath. "Investigating the IEEE 802.11ad Standard for Millimeter Wave Automotive Radar." In: *82nd Vehicular Technology Conference (VTC'15-Fall)*. IEEE, Sept. 2015, pp. 1–5. ISBN: 978-1-4799-8091-8. DOI: 10.1109/VTCFall.2015.7390996.

[KS16] Shajahan Kutty and Debarati Sen. "Beamforming for Millimeter Wave Communications: An Inclusive Survey." In: *IEEE Communications Surveys & Tutorials* 18.2 (2016), pp. 949–973. ISSN: 1553-877X. DOI: 10.1109/COMST.2015.2504600.

[LLH94] B. Langen, G. Lober, and W. Herzig. "Reflection and Transmission Behaviour of Building Materials at 60 GHz." In: *5th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Wireless Networks - Catching the Mobile Future.* (1994), pp. 505–509. DOI: 10.1109/WNCMF.1994.529141.

[LFW] *Linux Firmware: Repository of firmware blobs for use with the Linux kernel*. URL: https://git.kernel.org/pub/scm/linux/kernel/git/firmware/linux-firmware.git.

[Loc+17] Adrian Loch, Hany Assasa, Joan Palacios, Joerg Widmer, Hans Suys, and Björn Debaillie. "Zero Overhead Device Tracking in 60 GHz Wireless Networks using Multi-Lobe Beam Patterns." In: *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies - CoNEXT '17*. Incheon, Republic of Korea: ACM, Dec. 2017, pp. 224–237. ISBN: 9781450354226. DOI: 10.1145/3143361.3143395.

[LBW16] Adrian Loch, Guillermo Bielsa, and Joerg Widmer. "Practical Lower Layer 60 GHz Measurements using Commercial Off-the-shelf Hardware." In: *10th International*

Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization (WiNTECH '16). New York City, New York, USA: ACM, Oct. 2016, pp. 9–16. ISBN: 9781450342520. DOI: 10.1145/2980159.2980165.

[Lu+14] Jonathan Lu, Daniel Steinbach, Patrick Cabrol, Philip Pietraski, and Ravikumar V. Pragada. "Propagation Characterization of an Office Building in the 60 GHz Band." In: *8th European Conference on Antennas and Propagation (EuCAP '14)* (Apr. 2014), pp. 809–813. DOI: 10.1109/EuCAP.2014.6901885.

[Mal10] Alexander Maltsev. *Channel Models for 60 GHz WLAN Systems*. Tech. rep. doc.: IEEE 802.11-09/0334r8. 2010.

[Mal+10a] Alexander Maltsev, Vinko Erceg, Eldad Perahia, Chris Hansen, Roman Maslennikov, Artyom Lomayev, Alexey Sevastyanov, Alexey Khoryaev, Gregory Morozov, Martin Jacob, Sebastian Priebe, Thomas Kürner, Shuzo Kato, Hirokazu Sawada, Katsuyoshi Sato, and Hiroshi Harada. *Channel Models for 60 GHz WLAN Systems (IEEE 802.11-09/0334r8)*. Tech. rep. IEEE P802.11 Wireless LANs, 2010, pp. 1–152.

[Mal+09] Alexander Maltsev, Roman Maslennikov, Alexey Sevastyanov, Alexey Khoryaev, and Artyom Lomayev. "Experimental Investigations of 60 GHz WLAN Systems in Office Environment." In: *IEEE Journal on Selected Areas in Communications* 27.8 (Oct. 2009), pp. 1488–1499. ISSN: 0733-8716. DOI: 10.1109/JSAC.2009.091018.

[Mal+10b] Alexander Maltsev, Roman Maslennikov, Alexey Sevastyanov, Artyom Lomayev, Alexey Khoryaev, Alexander Davydov, and Vladimir Ssorin. "Characteristics of Indoor Millimeter-wave Channel at 60 GHz in Application to Perspective WLAN System." In: *Fourth European Conference on Antennas and Propagation*. Barcelona, Spain: IEEE, Apr. 2010. ISBN: 978-84-7653-472-4.

[Mal+10c] Alexander Maltsev, Eldad Perahia, Roman Maslennikov, Alexey Sevastyanov, Artyom Lomayev, and Alexey Khoryaev. "Impact of Polarization Characteristics on 60-GHz Indoor Radio Communication Systems." In: *IEEE Antennas and Wireless Propagation Letters* 9 (2010), pp. 413–416. ISSN: 1536-1225. DOI: 10.1109/LAWP.2010.2048410.

[MMI96] Takeshi Manabe, Yuko Miura, and Toshio Ihara. "Effects of Antenna Directivity and Polarization on Indoor Multipath Propagation Characteristics at 60 GHz." In: *IEEE Journal on Selected Areas in Communications* 14.3 (Apr. 1996), pp. 441–448. ISSN: 07338716. DOI: 10.1109/49.490229.

[MRM16] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. "Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells." In: *IEEE Journal of Selected Topics in Signal Processing* 10.3 (Apr. 2016), pp. 514–527. ISSN: 1932-4553. DOI: 10.1109/JSTSP.2016.2520899.

[ML10] Roman Maslennikov and Artyom Lomayev. *Implementation of 60 GHz WLAN Channel Model*. Tech. rep. doc.: IEEE 802.11-10/0854r3. 2010.

[Mat] MathWorks. *Propagation Channel Models: Channel models for 802.11*. URL: http://de.mathworks.com/help/wlan/ref/wlantgnchannel-system-object.html.

[MG09] Minyoung Park and Praveen Gopalakrishnan. "Analysis on Spatial Reuse and Interference in 60-GHz Wireless Networks." In: *IEEE Journal on Selected Areas in Communications* 27.8 (Oct. 2009), pp. 1443–1452. ISSN: 0733-8716. DOI: 10.1109/JSAC.2009.091014.

[Nee+07] Behnam Neekzad, Kamran Sayrafian-Pour, Julio Perez, and John S. Baras. "Comparison of Ray Tracing Simulations and Millimeter Wave Channel Sounding Measurements." In: *18th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC '07)*. Athens, Greece: IEEE, Dec. 2007, pp. 1–5. ISBN: 978-1-4244-1143-6. DOI: 10.1109/PIMRC.2007.4394537.

[Nit+15a] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. "Boon and Bane of 60 GHz Networks." In: *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT '15*. Heidelberg, Germany: ACM, 2015, pp. 1–13. ISBN: 9781450334129. DOI: 10.1145/2716281.2836102.

[Nit+14] Thomas Nitsche, Carlos Cordeiro, Adriana B. Flores, Edward W. Knightly, Eldad Perahia, and Joerg Widmer. "IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-Second Wi-Fi." In: *IEEE Communications Magazine* 52.December (2014), pp. 132–141. DOI: 10.1109/MCOM.2014.6979964.

[Nit+15b] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. "Steering with Eyes Closed: mm-Wave Beam Steering Without In-band Measurement." In: *Conference on Computer Communications (INFOCOM '15)*. Kowloon, Hong Kong: IEEE, Apr. 2015, pp. 2416–2424. ISBN: 978-1-4799-8381-0. DOI: 10.1109/INFOCOM.2015.7218630.

[Niu+15] Yong Niu, Yong Li, Depeng Jin, Li Su, and Athanasios V. Vasilakos. "A Survey of Millimeter Wave Communications (mmWave) for 5G: Opportunities and Challenges." In: *Wireless Networks* 21.8 (Nov. 2015), pp. 2657–2676. ISSN: 1022-0038. DOI: 10.1007/s11276-015-0942-z.

[NZL17] Song Noh, Michael D. Zoltowski, and David J. Love. "Multi-Resolution Codebook and Adaptive Beamforming Sequence Design for Millimeter Wave Beam Alignment." In: *IEEE Transactions on Wireless Communications* 16.9 (2017), pp. 5689–5701. ISSN: 15361276. DOI: 10.1109/TWC.2017.2713357.

[NS3] *ns-3 Network Simulator*. URL: https://www.nsnam.org/.

[SSL] *OpenSSL: Cryptography and SSL/TLS Toolkit*. URL: https://www.openssl.org/.

[OpenWrt] *OpenWrt Project*. URL: https://openwrt.org/.

[Pal+18a] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices." In: *24th Annual International Conference on Mobile Computing and Networking (MobiCom '18)*. New Delhi, India: ACM, Oct. 2018, pp. 241–255. ISBN: 9781450359030. DOI: 10.1145/3241539.3241576.

[Pal+18b] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. *Addendum to Adaptive Codebook Optimization for Beam Training on Off-The-Shelf IEEE 802.11ad Devices*. Tech. rep. TR-IMDEA-Networks-2018-1. IMDEA Networks, July 2018.

[Pal+18c] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. "Demo: Channel Estimation and Custom Beamforming on the 60 GHz TP-Link Talon AD7200 Router." In: *12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH '18)*. New Delhi, India: ACM, Nov. 2018, pp. 10–11. ISBN: 9781450359306. DOI: 10.1145/3267204.3268070.

[Pas] Pasternack Enterprises. *Pasternack 60 GHz Transmit/Receive (Tx/Rx) Development System (PEM003-KIT)*.

[PKF07] Michael Peter, Wilhelm Keusgen, and Robert Felbecker. "Measurement and Ray-tracing Simulation of the 60 GHz Indoor Broadband Channel: Model Accuracy and Parameterization." In: *2nd European Conference on Antennas and Propagation (EuCAP 2007)*. Institution of Engineering and Technology, 2007, pp. 432–432. ISBN: 978 0 86341 842 6. DOI: 10.1049/ic.2007.1555.

[PyRIC] *PyRIC: Python Radio Interface Controller*. URL: https://iperf.fr.

[RVM12a] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. "Compressive Adaptation of Large Steerable Arrays." In: *Information Theory and Applications Workshop*. San Diego, CA, USA: IEEE, Feb. 2012, pp. 234–239. ISBN: 978-1-4673-1472-5. DOI: 10.1109/ITA.2012.6181796.

[RVM12b] Dinesh Ramasamy, Sriram Venkateswaran, and Upamanyu Madhow. "Compressive Tracking with 1000-element Arrays: A Framework for Multi-Gbps mm Wave Cellular Downlinks." In: *50th Annual Allerton Conference on Communication, Control, and Computing (Allerton '12)*. Monticello, IL, USA: IEEE, Oct. 2012, pp. 690–697. ISBN: 978-1-4673-4539-2. DOI: 10.1109/Allerton.2012.6483285.

[Rap+14] Theodore S. Rappaport, Robert W. Heath, Robert C. Daniels, and James N. Murdock. *Millimeter Wave Wireless Communications*. Prentice Hall, Sept. 2014. ISBN: 978-0132172288.

[RMG11] Theodore S. Rappaport, James N. Murdock, and Felix Gutierrez. "State of the Art in 60-GHz Integrated Circuits and Systems for Wireless Communications." In: *Proceedings of the IEEE* 99.8 (Aug. 2011), pp. 1390–1436. ISSN: 0018-9219. DOI: 10.1109/JPROC.2011.2143650.

[RST91] Theodore S. Rappaport, Scott Y. Seidel, and Koichiro Takamizawa. "Statistical Channel Impulse Response Models for Factory and Open Plan Building Radio Communicate System Design." In: *IEEE Transactions on Communications* 39.5 (May 1991), pp. 794–807. ISSN: 00906778. DOI: 10.1109/26.87142.

[Ras+17] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. "Noncoherent mmWave Path Tracking." In: *18th International Workshop on Mobile Computing Systems and Applications (HotMobile '17)*. Sonoma, CA, USA: ACM, Feb. 2017, pp. 13–18. ISBN: 9781450349079. DOI: 10.1145/3032970.3032974.

[RGH15] Cristian Rusu, Nuria Gonzalez-Prelcic, and Robert W. Heath. "An Attack on Antenna Subset Modulation for Millimeter Wave Communication." In: *International Conference on Acoustics, Speech and Signal Processing (ICASSP '15)*. Brisbane, QLD, Australia: IEEE, Apr. 2015, pp. 2914–2918. ISBN: 978-1-4673-6997-8. DOI: 10.1109/ICASSP.2015.7178504.

[Sah+18] Swetank Kumar Saha, Hany Assasa, Adrian Loch, Naveen Muralidhar Prakash, Roshan Shyamsunder, Shivang Aggarwal, Daniel Steinmetzer, Dimitrios Koutsonikolas, Joerg Widmer, and Matthias Hollick. "Fast and Infuriating : Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware." In: *15th International Conference on Sensing, Communication, and Networking (SECON '18)*. Hong Kong: IEEE, June 2018. ISBN: 9781538642818. DOI: 10.1109/SAHCN.2018.8397123.

[Sah+17] Swetank Kumar Saha, Daniel Uvaydov, Josep Miquel Jornet, Edward W. Knightly, Dimitrios Koutsonikolas, Dimitris Pados, Zhi Sun, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, and Owen Torres. "X60." In: *11th Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH '17)*. Snowbird, Utah, USA: ACM, Oct. 2017, pp. 75–82. ISBN: 9781450351478. DOI: 10.1145/3131473.3131479.

[SV87] Adel A. M. Saleh and Reinaldo Valenzuela. "A Statistical Model for Indoor Multipath Propagation." In: *IEEE Journal on Selected Areas in Communications* 5.2 (Feb. 1987), pp. 128–137. ISSN: 0733-8716. DOI: 10.1109/JSAC.1987.1146527.

[Sat+97] Katsuyoshi Sato, Takeshi Manabe, Toshio Ihara, Hiroshi Saito, Shigeru Ito, Tetsu Tanaka, Kazuyoshi Sugai, Norichika Ohmi, Yasushi Murakami, Masanori Shibayama, Yoshihiko Konishi, and Tsuneto Kimura. "Measurements of Reflection and Transmission Characteristics of Interior Structures of Office Building in the 60-GHz Band." In: *IEEE Transactions on Antennas and Propagation* 45.12 (1997), pp. 1783–1792. ISSN: 0018926X. DOI: 10.1109/8.650196.

[Sch+16] Tobias Schultes, Markus Grau, Daniel Steinmetzer, and Matthias Hollick. "Demo: Far Away and Yet Nearby - a Framework for Practical Distance Fraud on Proximity Services for Mobile Devices." In: *9th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '16)*. Darmstadt, Germany: ACM, July 2016, pp. 205–207. ISBN: 9781450342704. DOI: 10.1145/2939918.2942416.

[Sch+17] Matthias Schulz, Francesco Gringoli, Daniel Steinmetzer, Michael Koch, and Matthias Hollick. "Massive Reactive Smartphone-based Jamming using Arbitrary Waveforms and Adaptive Power Control." In: *10th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec*

'17). Boston, MA, USA: ACM, July 2017, pp. 111–121. ISBN: 9781450350846. DOI: 10.1145/3098243.3098253.

[SWH17] Matthias Schulz, Daniel Wegemer, and Matthias Hollick. *Nexmon: The C-based Firmware Patching Framework*. 2017. URL: https://nexmon.org.

[SF15] Hossein Shokri-Ghadikolaei and Carlo Fischione. "Millimeter Wave Ad Hoc Networks: Noise-Limited or Interference-Limited?" In: *2015 IEEE Globecom Workshops (GC Wkshps)* (Dec. 2015), pp. 1–7. DOI: 10.1109/GLOCOMW.2015.7414085.

[SR14] Jaspreet Singh and Sudhir Ramakrishna. "On the Feasibility of Beamforming in Millimeter Wave Communication Systems with Multiple Antenna Arrays." In: *IEEE Global Communications Conference (GLOBECOM '14)*. Austin, TX, USA: IEEE, Dec. 2014, pp. 3802–3808. ISBN: 978-1-4799-3512-3. DOI: 10.1109/GLOCOM.2014.7037400.

[SR15] Jaspreet Singh and Sudhir Ramakrishna. "On the Feasibility of Codebook-Based Beamforming in Millimeter Wave Systems With Multiple Antenna Arrays." In: *IEEE Transactions on Wireless Communications* 14.5 (May 2015), pp. 2670–2683. ISSN: 1536-1276. DOI: 10.1109/TWC.2015.2390637.

[Sin+09] Sumit Singh, Federico Ziliotto, Upamanyu Madhow, Elizabeth M. Belding, and Mark Rodwell. "Blockage and Directivity in 60 GHz Wireless Personal Area Networks: from Cross-layer Model to Multihop MAC Design." In: *IEEE Journal on Selected Areas in Communications* 27.8 (Oct. 2009), pp. 1400–1413. ISSN: 0733-8716. DOI: 10.1109/JSAC.2009.091010.

[SW92] Peter F. M. Smulders and Anthony G. Wagemans. "Wideband Indoor Radio Propagation Measurements at 58 GHz." In: *Electronics Letters* 28.13 (1992), p. 1270. ISSN: 00135194. DOI: 10.1049/el:19920804.

[Ste+18] Daniel Steinmetzer, Saad Ahmad, Nikolaos A. Anagnostopoulos, Matthias Hollick, and Stefan Katzenbeisser. "Authenticating the Sector Sweep to Protect Against Beam-Stealing Attacks in IEEE 802.11ad Networks." In: *2nd Workshop on Millimeter Wave Networks and Sensing Systems (mmNets '18)*. New Delhi, India: ACM, Oct. 2018, pp. 3–8. ISBN: 978-1-4503-5928-3. DOI: 10.1145/3264492.3264494.

[Ste+15] Daniel Steinmetzer, Joe Chen, Jiska Classen, Edward W. Knightly, and Matthias Hollick. "Eavesdropping with Periscopes: Experimental Security Analysis of Highly

Directional Millimeter Waves." In: *Conference on Communications and Network Security (CNS '15)*. Florence, Italy: IEEE, Sept. 2015, pp. 335–343. ISBN: 9781467378765. DOI: 10.1109/CNS.2015.7346844.

[SCH16a] Daniel Steinmetzer, Jiska Classen, and Matthias Hollick. "Exploring Millimeter-Wave Network Scenarios with Ray-tracing based Simulations in mmTrace." In: *International Conference on Computer Communications Workshops (INFOCOM WKSHPS '16)*. San Francisco, CA, USA: IEEE, Apr. 2016. ISBN: 978-1-4673-9955-5. DOI: 10.1109/INFCOMW.2016.7562269.

[SCH16b] Daniel Steinmetzer, Jiska Classen, and Matthias Hollick. "mmTrace: Modeling Millimeter-wave Indoor Propagation with Image-based Ray-tracing." In: *1st Millimeter-wave Networking Workshop (mmNet '16)*. San Francisco, CA, USA: IEEE, Apr. 2016, pp. 429–434. ISBN: 9781467399555. DOI: 10.1109/INFCOMW.2016.7562115.

[Ste+17a] Daniel Steinmetzer, Adrian Loch, Amanda García-García, Joerg Widmer, and Matthias Hollick. "Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks." In: *1st Workshop on Millimeter-Wave Networks and Sensing Systems (mmNets '17)*. Snowbird, UT, USA: ACM, Dec. 2017, pp. 29–34. ISBN: 978-1-4503-5143-0. DOI: 10.1145/3130242.3130244.

[SSH15] Daniel Steinmetzer, Matthias Schulz, and Matthias Hollick. "Lockpicking Physical Layer Key Exchange: Weak Adversary Models Invite the Thief." In: *8th ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec '15)*. New York City, NY, USA: ACM, June 2015. ISBN: 978-1-4503-3623-9. DOI: 10.1145/2766498.2766514.

[SSH18] Daniel Steinmetzer, Milan Stute, and Matthias Hollick. "TPy: A Lightweight Framework for Agile Distributed Network Experiments." In: *12th International Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization (WiNTECH '18)*. New Delhi, India: ACM, Nov. 2018, pp. 38–45. ISBN: 978-1-4503-5930-6. DOI: 10.1145/3267204.3267214.

[SWH18] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. "A Practical IEEE 802.11ad Research Platform: The Hidden Potential of Off-the-Shelf Devices." In: *3rd NSF Millimeter-Wave Research Coordination Network (RCN) Workshop*. Tucson, AZ, USA: NSF, Jan. 2018.

[Ste+17b] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices." In: *13th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '17)*. Incheon, Republic of Korea: ACM, Dec. 2017, pp. 414–425. ISBN: 9781450354226. DOI: 10.1145/3143361.3143384.

[SYH18] Daniel Steinmetzer, Yimin Yuan, and Matthias Hollick. "Beam-Stealing: Intercepting the Sector Sweep to Launch Man-in-the-Middle Attacks on Wireless IEEE 802.11ad Networks." In: *11th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '18)*. Stockholm, Sweden: ACM, June 2018, pp. 12–22. ISBN: 9781450357319. DOI: 10.1145/3212480.3212499.

[Sur+17] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. "WiFi-Assisted 60 GHz Wireless Networks." In: *23rd Annual International Conference on Mobile Computing and Networking (MobiCom '17)*. Snowbird, Utah, USA: ACM, Oct. 2017, pp. 28–41. ISBN: 9781450349161. DOI: 10.1145/3117811.3117817.

[Sur+15] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. "60 GHz Indoor Networking through Flexible Beams." In: *International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '15)* (2015), pp. 71–84. DOI: 10.1145/2745844.2745858.

[Sur+16] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. "BeamSpy: Enabling Robust 60 GHz Links Under Blockage." In: *13th Usenix Conference on Networked Systems Design and Implementation (NSDI '16)*. Santa Clara, CA, USA: USENIX, Mar. 2016, pp. 193–206. ISBN: 978-1-931971-29-4.

[TCK15] Timothy A. Thomas, Mark Cudak, and Tom Kovarik. "Blind Phase Noise Mitigation for a 72 GHz Millimeter Wave System." In: *International Conference on Communications (ICC '15)*. IEEE, June 2015, pp. 1352–1357. ISBN: 978-1-4673-6432-4. DOI: 10.1109/ICC.2015.7248511.

[TPA11] Y. Ming Tsang, Ada S. Y. Poon, and Sateesh Addepalli. "Coding the Beams: Improving Beamforming Training in mmWave Communication System." In: *Global Telecommunications Conference (GLOBECOM '11)*. Kathmandu, Nepal: IEEE, Dec. 2011, pp. 1–6. ISBN: 978-1-4244-9268-8. DOI: 10.1109/GLOCOM.2011.6134486.

[VLH13] Nachiappan Valliappan, Angel Lozano, and Robert W. Heath. "Antenna Subset Modulation for Secure Millimeter-Wave Wireless Communication." In: *IEEE Transactions on Communications* 61.8 (Aug. 2013), pp. 3231–3245. ISSN: 0090-6778. DOI: 10.1109/TCOMM.2013.061013.120459.

[WARP] *WARP Project*. URL: https://warpproject.org.

[WARP+] *WARP Project: 802.11 Reference Design for WARP v3*. URL: https://warpproject.org/trac/wiki/802.11.

[WZ17] Teng Wei and Xinyu Zhang. "Pose Information Assisted 60 GHz Networks." In: *23rd Annual International Conference on Mobile Computing and Networking (MobiCom '17)*. Snowbird, Utah, USA: ACM, Oct. 2017, pp. 42–55. ISBN: 9781450349161. DOI: 10.1145/3117811.3117832.

[WZZ17] Teng Wei, Anfu Zhou, and Xinyu Zhang. "Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors." In: *USENIX Symposium on Networked Systems Design and Implementation (NSDI '17)*. Boston, MA, USA: USENIX, Mar. 2017. ISBN: 978-1-931971-37-9.

[Whi05] Jerry C. Whitaker. *The Electronics Handbook*. 2nd Editio. CRC Press, Apr. 2005. ISBN: 978-0849318894.

[Won+14] Wonbin Hong, Kwang-Hyun Baek, Youngju Lee, Yoongeon Kim, and Seung-Tae Ko. "Study and Prototyping of Practically Large-scale mmWave Antenna Systems for 5G Cellular Devices." In: *IEEE Communications Magazine* 52.9 (Sept. 2014), pp. 63–69. ISSN: 0163-6804. DOI: 10.1109/MCOM.2014.6894454.

[XKR02] Hao Xu, Vikas Kukshya, and Theodore S. Rappaport. "Spatial and Temporal Characteristics of 60-GHz Indoor Channels." In: *IEEE Journal on Selected Areas in Communications* 20.3 (Apr. 2002), pp. 620–630. ISSN: 07338716. DOI: 10.1109/49.995521.

[YDX15] Guang Yang, Jinfeng Du, and Ming Xiao. "Maximum Throughput Path Selection with Random Blockage for Indoor 60 GHz Relay Networks." In: *IEEE Transactions on Communications* 63.10 (2015), pp. 3511–3524. ISSN: 00906778. DOI: 10.1109/TCOMM.2015.2463284.

[Yan+15] Nan Yang, Lifeng Wang, Giovanni Geraci, Maged Elkashlan, Jinhong Yuan, and Marco Di Renzo. "Safeguarding 5G Wireless Communication Networks Using Physical Layer Security." In: *IEEE Communications Magazine* 53.4 (Apr. 2015), pp. 20–27. ISSN: 0163-6804. DOI: 10.1109/MCOM.2015.7081071.

[Zeno9] Erik Zenner. "Nonce Generators and the Nonce Reset Problem." In: *12th International Conference on Information Security (ISC '09)*. Pisa, Italy: Springer, 2009, pp. 411–426. DOI: 10.1007/978-3-642-04474-8_33.

[ZGP18] Ding Zhang, Mihir Garude, and Parth H. Pathak. "MmChoir: Exploiting Joint Transmissions for Reliable 60GHz mmWave WLANs." In: *International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '18)*. 2018, pp. 251–260. ISBN: 9781450357708. DOI: 10.1145/3209582.3209608.

[Zha+16] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. "OpenMili: A 60 GHz Software Radio with a Programmable Phased-array Antenna." In: *22nd Annual International Conference on Mobile Computing and Networking (MobiCom '16)*. New York City, New York, USA: ACM, Oct. 2016, pp. 485–486. ISBN: 9781450342261. DOI: 10.1145/2973750.2985614.

[ZZM17] Anfu Zhou, Xinyu Zhang, and Huadong Ma. "Beam-Forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering." In: *Conference on Computer Communications (INFOCOMM '17)*. Atlanta, GA, USA: IEEE, May 2017, pp. 1–9. ISBN: 978-1-5090-5336-0. DOI: 10.1109/INFOCOM.2017.8057188.

[Zhu+16] Yongxu Zhu, Lifeng Wang, Kai-Kit Wong, and Robert W. Heath. "Physical Layer Security in Large-Scale Millimeter Wave Ad Hoc Networks." In: *Global Communications Conference (GLOBECOM '16)*. Washington, DC, USA: IEEE, Dec. 2016, pp. 1–6. ISBN: 978-1-5090-1328-9. DOI: 10.1109/GLOCOM.2016.7842143.

[Zhu+17] Yongxu Zhu, Lifeng Wang, Kai-Kit Wong, and Robert W. Heath. "Secure Communications in Millimeter Wave Ad Hoc Networks." In: *IEEE Transactions on Wireless Communications* 16.5 (May 2017), pp. 3205–3217. ISSN: 1536-1276. DOI: 10.1109/TWC.2017.2676087.

ERKLÄRUNG ZUR DISSERTATIONSSCHRIFT

gemäß § 9 der Allgemeinen Bestimmungen der Promotionsordnung
der Technischen Universität Darmstadt
vom 12. Januar 1990 (ABI. 1990, S. 658)
in der Fassung der 8. Novelle
vom 1. März 2018

Hiermit versichere ich, Daniel Steinmetzer, die vorliegende Dissertationsschrift ohne Hilfe Dritter und nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die Quellen entnommen wurden, sind als solche kenntlich gemacht worden. Eigenzitate aus vorausgehenden wissenschaftlichen Veröffentlichungen werden in Anlehnung an die Hinweise des Promotionsausschusses FB Informatik zum Thema „Eigenzitate in wissenschaftlichen Arbeiten“ (EZ-2014/10) in Kapitel *„Previously Published Material“* auf Seite XXIII ff. gelistet. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen. In der abgegebenen Dissertationsschrift stimmen die schriftliche und die elektronische Fassung überein.

Darmstadt, 17. Dezember 2018

Daniel Steinmetzer