Robust and Reliable Millimeter Wave Wireless Networks

by

Hany Assasa

A dissertation submitted by in partial fulfillment of the requirements for the degree of Doctor of Philosophy in

Telematic Engineering

Universidad Carlos III de Madrid

Advisor/Tutor:

Joerg Widmer

May, 2019

ACKNOWLEDGEMENTS

First, I would like to start by thanking my supervisor Dr. Joerg Widmer for his invaluable guidance and profound support during my Ph.D. period. I sincerely appreciate his patience and academic mentoring. I would also like to express my greatest gratitude for my internship supervisors: Dr. Andres Garcia-Saavedra and Dr. Xavier Costa-Perez from NEC Laboratories Europe, and Dr. Wim Van Thillo from imec and PharrowTech. They provided me with an excellent opportunity to expand my knowledge and obtain experience outside of my working domain. Special thanks go to Dr. Adrian Loch with whom I have worked during the first three years of my Ph.D., and we published several papers together. I want to thank Dr. Tanguy Ropitault from NIST with whom I have been collaborating during the last two years of my Ph.D.

Second, I want to thank all my friends and colleagues at IMDEA Networks Institute, Amr Hussein, Noelia Perez, Mohamed Lamine, Dr. Foivos Michelinakis, Dr. Roderick Fanou, Maurizio Rea, Roberto Calvo, Dario Bega, Dr. Jesus Omar Lacruz, Dr. Danilo De Donno, Dr. Christian Vitale, Dr. Evgenia Christoforou, Dr. Nicola Bui, Dr. Elli Zavou, Dr. Qing Wang, Dr. Gek Hong Sim, Pablo Jimenez, Joan Palacios, Elizaveta Dubrovinskaya, Alvaro Feal, Ander Galisteo, Julien Gamba, Pelayo Vallina, Ricardo Padrino, Javier Hervas, Rafael Garcia, Hector Cordobes, Dr. Claudio Fiandrino, Angel Acosta, Carlos Donato, Cristina Marquez, Oluwasegun Ojo, Adriana Moreno, Alejandro Blanco, Dolores Garcia Marti, Constantine Ayimba, Edgar Arribas, Guillermo Bielsa, Diego Juara, Norbert Ludant, Pavel Chuprikov, Dr. Jose Felix Kukielka, Dr. Vincenzo Mancuso, Dr. Paolo Casari, Dr. Antonio Fernandez Anta, Dr. Sergey Gorinsky, and Dr. Marco Ajmone Marsan. I had a great with them and I hope we can meet again in the near future.

During my Ph.D. internships, I had the pleasure to meet and work closely with these amazing people: Steve Blandino, Takahito Yoshizawa, Dr. Khaled Khalaf, Andy Dewilde, Lanfranco Zanzi, Umar Farooq, Dr. Vincenzo Sciancalepore, Xavier Salvat, and Dr. Kheireddine Aziz.

Finally, I want to express my sincerest thanks and gratitude to my beloved parents and sister in Syria. They were always providing their unconditional love, guidance, and support, even though they were not physically in Spain.

PUBLISHED AND SUBMITTED CONTENT

The research work in this thesis has been published and accepted in several peer-reviewed conferences and workshops. We list in details all the publications that are included in this thesis and their corresponding chapters:

1. **Hany Assasa** and Joerg Widmer, "Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3," in Proceedings of the Workshop on ns-3 (ACM WNS3 2016). pp. 57-64. June 2016, Seattle, WA, USA. [Online]. Available at: http://doi.acm.org/10.1145/2915371.2915377

    - I led the development, implementation, and evaluation process in this work.
    - The item is wholly included in the thesis. The main contribution of this work is located in Chapter 3. Chapter 4 contains the evaluation section of this work.
    - The material from this source included in this thesis is not singled out with typographic means and references.

2. **Hany Assasa** and Joerg Widmer, "Extending the IEEE 802.11ad Model: Scheduled Access, Spatial Reuse, Clustering, and Relaying," in Proceedings of the Workshop on ns-3 (ACM WNS3 2017). pp. 39-46, June 2017, Porto, Portugal. [Online]. Available at: http://doi.acm.org/10.1145/3067665.3067667

    - The work in this publication is an extension to the previous work. Again, I led the development, implementation, and evaluation process in this work.
    - The item is wholly included in the thesis. The main contribution of this work is located in Chapter 3. Chapter 4 contains the evaluation section of this work.
    - The material from this source included in this thesis is not singled out with typographic means and references.

3. **Hany Assasa**, Joerg Widmer, Tanguy Ropitault, and Nada Golmie, "Enhancing the ns-3 IEEE 802.11ad Model Fidelity: Beam Codebooks, Multi-antenna Beamforming Train-

ing, and Quasi-deterministic mmWave Channel," in Proceedings of the Workshop on ns-3 (ACM WNS3 2019). June 2019, Florence, Italy.

- This work is an extension to the previous two works. In collaboration with NIST, I led the development, implementation, and evaluation process in this work.
- The item is wholly included in the thesis. The main contribution of this work is located in Chapter 3. Chapter 4 contains the evaluation section of this work.
- The material from this source included in this thesis is not singled out with typographic means and references.

4. **Hany Assasa**, Joerg Widmer, Tanguy Ropitault, Anuraag Bodi, and Nada Golmie, "High Fidelity Simulation of IEEE 802.11ad in ns-3 Using a Quasi-deterministic Channel Model," in ACM Workshop on Next-Generation Wireless with ns-3 (ACM WNGW 2019). June 2019, Florence, Italy.

   - In this work, I defined several simulations scenarios and carried out the evaluation process using network simulator ns-3. The co-authors from NIST helped in generating the corresponding channel model files for this work.
   - The item is wholly included in the thesis in Chapter 4.
   - The material from this source included in this thesis is not singled out with typographic means and references.

5. **Hany Assasa**, Joerg Widmer, Jian Wang, Tanguy Ropitault, and Nada Golmie, "An Implementation Proposal for IEEE 802.11ay SU/MU-MIMO Communication in ns-3," in ACM Workshop on Next-Generation Wireless with ns-3 (ACM WNGW 2019). June 2019, Florence, Italy.

   - In this work, I proposed a framework for simulating SU/MU-MIMO communications using ns-3. The co-authors from NIST gave their feedback regarding the proposal.
   - The item is wholly included in the thesis in Chapter 3.
   - The material from this source included in this thesis is not singled out with typographic means and references.

6. **Hany Assasa**, Adrian Loch, and Joerg Widmer, "Packet mass transit: Improving frame aggregation in 60 GHz networks," in WoWMoM 2016 - 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks, 2016. [Online]. Available at: https://doi.org/10.1109/WoWMoM.2016.7523522

   - In this work, I designed and evaluated two policies that improve the performance of IEEE 802.11ad WLAN networks utilizing the CSMA/CA protocol.

- The item is wholly included in the thesis in Chapter 5.
- The material from this source included in this thesis is not singled out with typographic means and references.

7. **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 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoWMoM 2018,2018. [Online]. Available: https://doi.org/10.1109/WoWMoM.2018.8449795

   - I performed all the reported measurements in this work to analyze and study the performance of practical 802.11ad COTS devices in dense deployment scenarios.
   - The item is wholly included in the thesis in Chapter 6.
   - The material from this source included in this thesis is not singled out with typographic means and references.

8. **Hany Assasa**, Guillermo Bielsa, Swetank Kumar Saha, Pablo Jimenez Mateo, Adrian Loch, Dimitrios Koutsonikolas, and Joerg Widmer, "Performance Analysis of Medium Access Control and Spatial Reuse for IEEE 802.11ad Deployments," Under submission to Elsevier Pervasive and Mobile Computing (PMC) Journal, 2019.

   - This work is an extension to the previous work item, and it is still under submission to Elsevier Pervasive and Mobile Computing (PMC) Journal. I performed a new set of measurements to study new aspects of COTS devices and report their performance.
   - The item is wholly included in the thesis in Chapter 6.
   - The material from this source included in this thesis is not singled out with typographic means and references.

9. 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 2017). pp. 224-237. [Online]. Available at: http://doi.acm.org/10.1145/3143361.3143395

   - In this work, I have contributed to the simulation part. I used the results collected from the testbed as an input to network simulator ns-3. Inside ns-3, I compare the performance of the developed beam tracking algorithm with respect to the legacy method utilized in the IEEE 802.11ad standard.
   - The item is partly included in the thesis in Appendix A.

- The material from this source included in this thesis is not singled out with typographic means and references.

OTHER RESEARCH MERITS

The following list includes published research items that I have co-authored during the course of my Ph.D. However, these items are not included in this thesis.

1. Joan Palacios, Paolo Casari, **Hany Assasa**, and Joerg Widmer, "LEAP: Location Estimation and Predictive Handover with Consumer-Grade mmWave Devices." In: The 38th IEEE International Conference on Computer Communications (IEEE INFOCOM 2019), 29 Apr - 02 May 2019, Paris, France.

2. 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 2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON 2018), pp. 1-9. [Online]. Available: https://doi.org/10.1109/SAHCN.2018.8397123

3. Claudio Fiandrino, **Hany Assasa**, Paolo Casari, and Joerg Widmer. "Scaling Millimeter-Wave Networks to Dense Deployments and Dynamic Environments". In: Proceedings of the IEEE, vol. 107, no. 4, 2019.

4. Swetank Kumar Saha, Roshan Shyamsunder, Naveen Muralidhar Prakash, **Hany Assasa**, Adrian Loch, Dimitrios Koutsonikolas, and Joerg Widmer, "Poster: Can MPTCP Improve Performance for Dual-Band 60 GHz/5 GHz Clients? " In ACM International Conference on Mobile Computing and Networking (MobiCom 2017). [Online] Available at: https://doi.org/10.1109/JPROC.2019.2897155

Abstract

Millimeter wave (mmWave) technology is one of the main pillars of the next generation Wireless Local Area Networks (WLANs) and 5G mobile networks. The main reason lies in the quantum leap of capacity it provides with respect to wireless networks operating in the sub-6-GHz band. Nevertheless, efficient and reliable communication in the mmWave band demands novel techniques to tackle all the barriers associated with wireless propagation in this band. For example, material penetration and diffraction in the mmWave band are much weaker than the ones experienced in the sub-6-GHz band. As a result, wireless propagation in the mmWave band has a quasi-optical behavior in which the Line-of-sight (LoS) component contributes to the majority of the received signal power. For these reasons, mmWave devices rely on either horn antennas or electronically steerable phased antenna arrays to establish directional beams and focus the energy towards a specific direction in space. Although directional communication compensates for the high power attenuation, it creates a new problem that is uncharted for wireless networks operating in the sub-6-GHz band. It makes mmWave wireless links susceptible to blockage, human mobility, and device rotation. Solving these issues requires developing algorithms to quickly find alternative paths for communication upon link interruption.

The IEEE 802.11ad protocol is the first WLAN standard that supports wireless networking in the unlicensed 60 GHz band. IEEE 802.11ad tackles the aforementioned problems by introducing new mechanisms at the medium access control (MAC) and Physical (PHY) layers such as beamforming training and beam tracking, hybrid channel access scheme, relay operation mode, multi-band operation, etc. The performance of 60 GHz networks is the result of the interaction of all layers of the protocol stack with the wireless medium. To understand this interaction, it is fundamental to consider 60 GHz networks as a whole. Nevertheless, real-world experimentation with mmWave communication is not always feasible due to the significant amount of resources required and its associated costs. For these reasons, we develop in this thesis a high fidelity system-level model to simulate the IEEE 802.11ad standard. This allows us to study large-scale wireless networks operating in the 60 GHz band, taking into account all of the essential features supported by the standard. We investigate the networking aspects of various mmWave wireless network deployments and provide solutions to boost their performance. Additionally, since most of the mmWave devices are anticipated to adopt the carrier sense multiple access with collision avoidance (CSMA/CA) as the primary channel access scheme, we propose two frame-aggregation

policies that significantly improve network throughput and reduce end-to-end delay.

To complement our insights from the simulations, we analyze in-depth the performance of Commercial off-the-shelf (COTS) devices that implement the full 802.11ad protocol stack. Fortunately, a growing number of 60 GHz devices supporting the IEEE 802.11ad standard have become recently available. However, the standard does not specify implementation-dependent characteristics that have a significant impact on device performance. For example, this includes the periodicity of beamforming training, which directly affects ongoing communication and performance under mobility. Also, the placement of the 60 GHz antenna in the device plays a fundamental role regarding self-shadowing, which in turn affects how the device shall be deployed to maximize coverage and improve spatial reuse. Understanding such implementation-dependent issues is crucial to correctly draw the aforementioned system-level insights. In this thesis, we characterize and compare commercial 60 GHz devices which are widely used for research purposes. Particularly, we look at different networking aspects, including spatial sharing, beam patterns synthesis, frame aggregation, and the interactions between CSMA/CA and Transmission Control Protocol (TCP) protocols within dense network settings. In summary, we find significant discrepancies in terms of behavior and performance between these devices and what theory suggests. Additionally, achieving a gigabit of throughput and maintaining low-latency require adopting cross-layer solutions operating between the MAC layer and the transport layer.

The firmware running on these devices provide some experimental capabilities based on the operation of the IEEE 802.11ad protocol. These capabilities include reporting the Channel State Information (CSI) per antenna element for the connected antenna array. This CSI information contains valuable information about the spatial environment. Additionally, this information is sensitive to any minor changes in the environment. This inspires us to build a collaborative scheme consisting of spatially distributed access points (APs) that sense the environment and pinpoint the location of an obstacle without involving the end users.

In this thesis, we first delve into the operation of the IEEE 802.11ad protocol and analyze its performance both in simulations and practice. Second, we propose a set of solutions and recommendations to boost its efficiency. Finally, we extend its scope beyond typical wireless communication and utilize it to build a passive localization system.

Table of Contents

List of Tables

List of Figures

List of Acronyms

A-BFT Association Beamforming Training

AC Access Category

ACK Acknowledgment

ADC Analog-to-Digital Converter

AGC Automatic Gain Control

AID Association Identifier

ANIPI Average Noise plus Interference Power Indicator

AoA Angle of Arrival

AoD Angle of Aeparture

AP access point

AR augmented reality

ATI Announcement Transmission Interval

AWGN Additive white Gaussian noise

AWV antenna weight vector

BA block acknowledgment

BDP bandwidth-delay product

BER bit error rate

BF Beamforming Training

BHI Beacon Header Interval

BI Beacon Interval

BRP Beam Refinement Protocol

BSS Basic Service Set

BT Beam Tracking

BTI Beacon Transmission Interval

CBAP Contention-based Access Period

CCA Clear Channel Assessment

CCSS single centralized coordination service set

CEF channel estimation field

CFO carrier frequency offset

CIR Channel Impulse Response

COTS Commercial off-the-shelf

CRC Cyclic redundancy check

CSI Channel State Information

CSMA/CA carrier sense multiple access with collision avoidance

CSV comma-separated values

CTS Clear-to-Send

CWND Congestion Window

DAC Digital-to-Analog Converter

DCE Direct Execution Mode

DCF Distributed Coordination Function

DMG Directional multi-gigabit

DTI Data Transmission Interval

EDCA Enhanced Distributed Coordination Function (DCF) channel access

EDCF Enhanced Distributed Coordination Function

EDMG enhanced directional multi gigabit

EIRP Equivalent Isotropically Radiated Power

FBCK Feedback

FCS frame check sequence

FD-AF full-duplex amplify-and-forward

FOV field of view

FST fast session transfer

FSTS fast session transfer session

FTP File Transfer Protocol

GP Grant Period

HD-DF half-duplex decode-and-forward

IFS Interframe Space

I-RxSS initiator Receive sector sweep (RXSS)

I-TxSS initiator TXSS

LLT link loss timeout

LoS Line-of-sight

LP Low Power

LUT lookup table

MAC medium access control

MCS modulation and coding scheme

MIMO multiple-input and multiple-output

MLME Mac layer management entity

MMSE minimum-mean-squared-error

mmWave millimeter wave

MPC multipath component

MPDU MAC protocol data unit

MSDU MAC service data unit

MTU maximum transmission unit

MU-MIMO Multi-User multiple-input and multiple-output (MIMO)

NLOS Non-line-of-sight

OFDM Orthogonal Frequency Division Multiplexing

OOMM On-Off Markov Model

PAA Phased Antenna Array

PBSS personal basic service set

PCB printed circuit board

PCP personal basic service set (PBSS) control point

PDF Probability Density Function

PER Packet Error Rate

PHY Physical

PLCP Physical Layer Convergence Protocol

PD Propagation Delay

PN Phase Noise

PP Polling Period

PPDU Physical Protocol Data Unit

PTP Point to Point

Q-D Quasi-deterministic

QoE Quality of Experience

QoS Quality of Service

RDS Relay DMG STA

REDS Relay Endpoint DMG STA

RF Radio Frequency

RLS Relay Link Setup

R-RxSS responder RXSS

RSNI Received Signal to Noise Indicator

RTS Ready-to-Send

RTT Round-Trip Time

R-TxSS responder TXSS

RxBF Receive beamforming

RXSS Receive sector sweep

SACK Selective Acknowledgment

SAP Service Access Point

S-AP Synchronization AP

SC Single Carrier

SIFS short interframe space

SINR signal-to-interference-plus-noise ratio

SLS Sector Level Sweep

SNR Signal-to-Noise Ratio

SP service period

SPCA Service Period Channel Access

S-PCP Synchronization PBSS control point (PCP)

SPR Service Period Request

SS Sector Sweep

SSE Sum of Squared Errors

SSID Service Set Identifier

SSW Sector Sweep

STA Station

STBC space-time block coding

STF short training field

STO Symbol Timing Offset

SU-MIMO Single User MIMO

SVD singular-value decomposition

SQI Signal Quality Indicator

TA Transmitter Address

TCP Transmission Control Protocol

TDMA time-division multiple access

TRN training

TRN-R Receive training

TRN-T Transmit training

TSF timing synchronization function

TxBF Transmit beamforming

TXOP Transmit Opportunity

TXSS Transmit Sector Sweep

UCA Uniform Circular Array

UDP User Datagram Protocol

UHD ultra high definition

ULA Uniform Linear Array

URA Uniform Rectangular Array

VR Virtual Reality

WLAN Wireless Local Area Network

WMI Wireless Module Interface

ZF Zero-forcing

Chapter 1

Introduction

Wireless communication in the mmWave band is a crucial enabler for the next generation wireless technologies such as 5G Mobile Networks. Thanks to the massive amount of spectrum available, mmWave can be used for a wide range of scenarios where, traditionally, wired solutions have been utilized, such as ultra-high definition video streaming, virtual-reality headsets, or wireless front-hauling and back-hauling in mobile networks. With the mature advancements and developments in electronics components and Radio Frequency (RF) circuits operating at these frequencies, manufacturing relatively low-cost mmWave electronics circuits became possible. This resulted in many commercial standards operating in the 60GHz mmWave band such as the IEEE 802.15.3c, and the IEEE 802.11ad [1]. However, mmWave communication is extremely challenging compared to traditional microwave communication. The main reason lies in the mmWave signal characteristics which include high attenuation, susceptibility to blockage and human mobility, and the necessity of a LoS path for stable communication [3,4]. The high signal attenuation in the mmWave band mandates devices to utilize phased antenna array to establish a directional communication link. However, the management of highly directional mmWave links poses some challenges that are alien to the more traditional sub-6-GHz bands. Namely, beam training or user association methods, triggered upon link degradation due to blockages, may incur substantial channel time wastage, particularly in highly-dynamic scenarios with moving receivers, transmitters, and reflectors. As a result, efficient wireless communication in this band is challenging and requires adequate network planning and efficient design of wireless networking protocols.

In this thesis, we propose various techniques to tackle the underlying challenges associated with wireless networking in the mmWave band. First, we create a high-fidelity open-source model for simulating the IEEE 802.11ad protocol in network-simulator ns-3. Using this model, we study the performance of the MAC layer of the 802.11ad protocol, and we identify several deficiencies that penalize its performance for various deployment setups. For each identified problem, we provide novel solutions that enhance wireless networking performance and improve protocol operations and efficiency. Next, we study and analyze the behavior of practical 802.11ad COTS devices. Particularly, we look at different networking aspects including TCP performance over

CSMA/CA for dense mmWave network, spatial reuse when utilizing practical phased antenna arrays, beam patterns synthesis for different frequency channels, etc. For each networking aspect, we study its impact on the overall network behavior and performance. Then, we propose standard-based solutions to tackle each issue independently. Finally, we go beyond just a typical wireless communication in the mmWave band. We build a mmWave sensing system based on COTS 60 GHz devices. The system is capable of sensing human presence and determines its location passively and transparently to the communicating devices.

1.1. Motivations and Contributions

This thesis is dedicated to investigate and improve the performance of the IEEE 802.11ad protocol in simulations and practice. Additionally, it proposes a novel mechanism to utilize the IEEE 802.11ad protocol beyond communication to understand the surrounding environment and provide reliable communication links. To achieve these objectives, the main motivations and contributions in this thesis are as follows:

- **Networking Simulation Tool**: Experimental evaluation of networking in the mmWave band is extremely costly, and available hardware has minimal capabilities [5]. Despite the availability of commercial devices utilizing the IEEE 802.11ad protocol, these devices provide only limited access to the operations of the lower layers of the protocol stack, which hinders in-depth analysis and development of innovative solutions. In such cases, resorting to network simulation is a very useful alternative which abstracts implementation details while providing a good grade of realism. However, there are no publicly available simulation tools supporting IEEE 802.11ad in the mmWave band. For these reasons, in Chapter 3, we present the implementation of a high-fidelity model for simulating the IEEE 802.11ad protocol using network simulator ns-3. The model allows network researchers to understand the interactions and behavior of mmWave devices. Additionally, it paves the way for the development of innovative wireless networking solutions. Our work is the first to provide an open source model for the IEEE 802.11ad amendment in ns-3.

  At the time of writing, the WiFi Alliance is finalizing the next generation multi-gigabit standard to support wireless networking at 60 GHz, the so-called IEEE 802.11ay [6]. IEEE 802.11ay is envisioned to support extremely high data-rates of up to 300 Gbps, achieved through new complex physical layer techniques including MIMO communication, channel bonding and aggregation, and high order modulation schemes. Simulating the IEEE 802.11ay standard in a network-level simulator requires accurate abstraction models to incorporate the effects of those techniques. However, ns-3 still lacks support for Multi-User MIMO (MU-MIMO) communication. Additionally, it requires generating environment dependent Signal-to-Noise Ratio (SNR)-to-bit error rate (BER) look-up tables to accurately simulate Single User MIMO (SU-MIMO) communication. At the end of Chapter 3, we

propose a hybrid implementation that includes minimum signal processing blocks to accurately simulate IEEE 802.11ay SU/MU-MIMO communication in ns-3 with high accuracy and reduced computational complexity.

In Chapter 4, we demonstrate the capabilities of the IEEE 802.11ad model in ns-3. First, we study the performance of each newly introduced technique in 802.11ad independently and validate its behavior. Then, we study the performance of the IEEE 802.11ad protocol for various deployment settings using realistic phased antenna arrays and quasi-deterministic channel models. More particularly, we look at the impact of LoS blockage and the use of Non-line-of-sight (NLOS) paths on link performance. Besides, we show the benefits of deploying multiple access points per room to guarantee gigabit throughput per user. Finally, we evaluate the performance of the IEEE 802.11ad protocol in a typical high-density scenario.

- **Frame Aggregation in Wireless Networks**: The impact of frame aggregation on WLAN performance increases dramatically with higher data rates. The key problem is that the transmission time of packets decreases while the medium access, preamble, and packet header overhead remain the same. Recent 802.11 standards address this issue using frame aggregation, i.e., grouping multiple data frames in a single transmission to reduce the overhead. This already provides substantial efficiency gains in networks operating in the 2.4 GHz and 5 GHz bands, and for the 60 GHz networks such as 802.11ad, gains are even more pronounced due to the order-of-magnitude higher data rates. In 802.11ad, frame aggregation becomes crucial to achieve the multi-Gbps data rates that are possible in theory, since medium access overhead can be 20 times larger than the time required to transmit a single packet. While frame aggregation is essential, it largely depends on the traffic patterns present in the wireless network, and a node may not always have enough packets in the transmit queue to achieve a sufficiently large aggregated frame size. A particularly harmful case occurs when new packets arrive just after a node started the transmission of a small group of aggregated packets. In Chapter 5, we investigate in which case 802.11ad enabled devices should wait to construct a larger aggregated packet before starting the channel access procedure. We present a simple waiting policy for the uplink case that either waits for a minimum number of packets or for a maximum amount of time, whichever comes first. For the downlink case, we utilize a maximum weight scheduling policy with a maximum waiting time. Our results show that both policies significantly improve medium utilization, thus increasing throughput and reducing end-to-end delay.

- **MAC and Transport Layer Aspects in Practical COTS 60 GHz Devices**: mmWave communication promises high spatial reuse at multi-Gbps data rates in dense wireless networks. Existing work studies such networks using commercial hardware but is limited to individual links. In Chapter 6, we study the performance of dense mmWave deployments featuring up to eight stations. We use a testbed of commercial IEEE 802.11ad mmWave

devices that provide access to lower layer parameters. This enables us to analyze the impact of these parameters and other deployment considerations on network performance. We analyze issues such as the impact of channel contention on the buffer size at the transport layer, the effect of frame aggregation, the directivity of practical phased antenna arrays at different frequency channels, and the efficiency of spatial sharing. Our results show that using large buffer sizes with TCP is harmful due to channel contention despite the multi-gigabit-per-second data rates.

For COTS, frame aggregation is only beneficial up to a certain level due to higher error rates for large frames. Besides, utilizing beamforming codebook with static entries has an impact on the overall network performance based on the operating frequency. Finally, spatial reuse is limited by imperfect beam patterns with many side lobes and is only significant for interferer distances of more than 10 m. In contrast, selecting beam patterns that maximize spatial reuse allows doubling the throughput.

- **mmWave Sensing**: mmWave communication systems, albeit having attained great momentum, pose challenges that are alien to the traditional sub-6 GHz bands due to the channel sparsity induced by highly-directional links. Namely, beam (re)alignment methods, triggered upon detection of link blockages, may incur substantial overhead. In line with related literature, we argue that to establish pervasive connectivity, mmWave systems shall integrate mechanisms to predict and react proactively to potential blockages. Towards this goal, in Chapter 7 we design and build `BeamScanner`, a collaborative mechanism across quasi-stationary mmWave APs to detect and track obstacles in indoor scenarios, such as an office room or a living room. The nexus of our approach is the ability of `BeamScanner` to detect weak reflections, even from human bodies, with readily available commodity hardware. Around this, we design a system controller that operates in two stages. First, `BeamScanner` collects CSI from custom-built narrow beams spurred across a set of distributed APs. Second, we use *k-nearest neighbor* (KNN) classifier to infer the location of objects in the environment based on fine-grained data from signal reflections. We evaluate our approach in a typical 6×11-meter indoor room with four commodity APs with a median estimation error of less than 3 m when detecting human obstacles.

- **Beam Tracking**: mmWave devices must use highly directional antennas to achieve Gbps data rates over reasonable distances due to the high path loss. As a consequence, it is crucial to align the antenna beams between sender and receiver precisely. Even minor movement or rotation of a device can result in beam misalignment and thus a substantial performance degradation. Existing work as well as standards such as IEEE 802.11ad [1] tackle this issue through antenna sector probing. This comes at the expense of a significant overhead, which may significantly reduce the performance of mmWave communication, particularly in mobile scenarios. In Appendix A, we present a mechanism that can track both movement and rotation of 60 GHz mobile devices with zero overhead. To this end,

we transmit part of the preamble of each packet using a multi-lobe beam pattern. Our approach does not require any additional control messages and is backward compatible with 802.11ad. We implement our scheme on a 60 GHz testbed using phased antenna arrays to obtain SNR measurements. We use these measurements to perform simulations in ns-3 to validate our approach in a wide range of scenarios. Using our scheme, we achieve up to $2\times$ throughput gain.

1.2. Thesis Overview

The thesis is divided into 8 Chapters. Within each Chapter, we provide a list of state-of-the-art works that are relevant to the topic of the Chapter. First, Chapter 2 presents the necessary background on wireless communications in the 60 GHz band. Additionally, it reviews the first WLAN amendment that operates in this band, the so-called IEEE 802.11ad. Then, Chapter 3 introduces our implementation details for simulating the protocol stack of the IEEE 802.11ad protocol using network simulator ns-3. Chapter 4 evaluates and validates the performance of our IEEE 802.11ad model for different deployment settings. Then, Chapter 5 describes our algorithm for improving overall network throughput and channel utilization for a wireless device utilizing the CSMA/CA protocol in the 60 GHz band. Next, Chapter 6 provides insights into the networking performance of different commodity 60 GHz wireless devices for practical deployment settings. Chapter 7 goes beyond regular communication in the mmWave band and presents a novel system that utilizes the IEEE 802.11ad protocol to sense the environment and localize humans. Chapter 8 concludes the thesis and summarizes our findings and discusses possible research directions. Finally, Appendix A presents our beam tracking algorithm based on the IEEE 802.11ad protocol and evaluates its performance in network simulator ns-3.

Chapter 2

Background on 60 GHz Networks

The penetration of data-intensive applications in the digital market such as augmented and virtual reality services, and ultra high definition (UHD) video streaming and distribution is driving mobile network operators and telecommunications equipment vendors to boost the performance and capacity of the existing cellular and wireless technologies. These services require significantly higher capacity compared to what is currently offered by wireless technologies operating in the sub-6-GHz band. For this reason, the mmWave band between, 30 GHz and 300 GHz, is drawing the attention of the research and industry communities because of the abundant spectrum available there mainly for the unlicensed use in the 60 GHz band. Wireless communication in this band is highly appealing since it provides incredibly high capacity and thus allows for a several-fold increase in data rates and lower latencies. However, transmission in this band has specific signal propagation characteristics and suffers from high attenuation [3, 4] compared to existing technologies working in lower bands and thus requires significant design changes for both MAC and PHY layers. To compensate for the high power loss, mmWave devices utilize phased antenna array to generate directional beam focusing the transmit power towards a specific direction in space. While transmitters use directional antennas to overcome attenuation, recent work shows that consumer-grade phased antenna arrays for 60 GHz systems have many side lobes [7]. Hence, interference among transmitters is significant in spite of directional communication. Further, the MAC layer must operate efficiently to ensure multi-Gbps rates. Since channel bandwidth is typically in the order of a couple of GHz, any inter-frame spacing and control data transmitted at a low MCS significantly reduce MAC efficiency.

Both the Wireless Gigabit Alliance (WiGig) and the Wi-Fi Alliance took the initiative to leverage the wide spectrum in the mmWave band and provide multi-gigabit per second communication in the 60 GHz unlicensed band. They introduced the WLAN IEEE 802.11ad amendment [1, 8] which provides very high throughput of up to 7 Gbps for short-range communication for local area networks. This allows for a range of new high-rate applications, such as wireless docking stations, wireless storage, and instant file synchronization. Compared to IEEE 802.11ac [9] which is capable of supporting multi-gigabit throughput by employing high MCSs and advanced physical

layer technologies such as multi-user-MIMO, IEEE 802.11ad achieves multi-gigabit throughput by utilizing only the wide channels of 2.16 GHz available at the 60 GHz band. In the following sections, we provide a brief description of the significant design changes for both PHY and MAC layers in IEEE 802.11ad and the intuition behind them. The material in this chapter is taken from [10].

2.1. Physical Layer

Emerging applications using the IEEE 802.11ad multi-gigabit capability have different constraints and requirements in terms of power consumption, data rates, processing capabilities, and antenna design complexity. For these reasons, IEEE 802.11ad introduces four different types of PHY layers to cope with these requirements. Each PHY layer supports a set of specific MCSs.

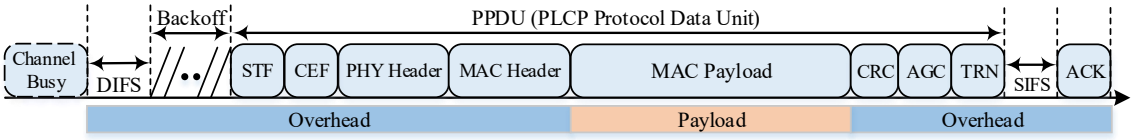

Figure 2.1: IEEE 802.11ad frame structure.

- **Control PHY**: This PHY layer (MCS0) is dedicated to low SNR operation with low throughput communication (27.5 Mbps). It is mainly used during the Beamforming Training (BF) phase.

- **OFDM PHY**: This PHY layer (MCS 13-24) provides the highest data rates of up to 6.76 Gbps. It adopts Orthogonal Frequency Division Multiplexing (OFDM) technology which is very efficient in multi-path environments. However, its implementation is complex, and therefore it targets devices with less stringent power and design constraints, such as docking stations and wireless streaming devices.

- **Single Carrier (SC) PHY**: Power limited and low complexity devices adopt this physical layer which provides a good trade-off between average throughput and energy efficiency compared to the OFDM PHY. Mobile phones and tablet devices will most likely adopt this PHY layer. SC PHY defines MCS 1-12, of which MCS 1-4 are mandatory modes to be implemented in all devices for interoperability.

- **Low Power (LP)-SC PHY**: This PHY layer with MCS 25-31 is similar to the SC PHY layer, but allows for further power reduction by using low-density parity check (LDPC) codes instead of Reed-Solomon codes.

Figure 2.1 depicts the IEEE 802.11ad frame structure. The frame starts with the typical IEEE 802.11 fields such as short training field (STF) and channel estimation field (CEF) which are

used for detecting the type of the PHY layer. These fields are followed by the PHY header which includes information such as payload length in bytes and index of the MCS used in the payload part. This field together with the MAC header and the MAC payload are protected by a Cyclic redundancy check (CRC). Finally, IEEE 802.11ad appends optionally two fields named Automatic Gain Control (AGC) and training (TRN). These new fields are used during the BF phase which we describe in section 2.2.

2.2. Beamforming Training Mechanism

Propagation conditions at 60 GHz band are worse compared to the lower bands due to oxygen absorption [3], high attenuation, weak signal reflectivity, and quasi-optical propagation behavior [11]. For these reasons, IEEE 802.11ad provides a mechanism to establish a directional link through a beamforming training process to compensate for signal quality degradation. In this process, stations focus their energy towards the intended receivers only, which increases antenna gain and may result in reduced interference, allowing for high spatial reuse. The beamforming training process in IEEE 802.11ad is divided into the following two phases:

- **SLS Phase**: In this phase, a DMG STA selects a coarse grain antenna sector for the initial communication. The phase can be used in two ways: 1) as TXSS where a DMG STA tries to select the best transmit antenna sector towards a particular DMG STA by sending Sector Sweep (SSW) frames via each of its antenna sectors or 2) as a RXSS, where a DMG STA trains its receive antenna sector by requesting a peer DMG STA to transmit SSW frames using a fixed antenna pattern while the former is sweeping across its receive antenna sectors.

- **BRP Phase**: IEEE 802.11ad defines multiple optional mechanisms to refine the sectors obtained in the SLS phase. The most important mechanism is the beam refinement mechanism, which is an iterative process where two DMG STAs exchange a special BRP packet ending with either Transmit training (TRN-T) or Receive training (TRN-R) fields. Additionally, the amendment defines a Beam Tracking (BT) option to keep a track of signal quality during an ongoing data transmission by adding the previous TRN fields to the PHY frames.

However, due to the high variability of the 60 GHz wireless channel, the quality of this beamformed link might degrade to a point in which a pair of DMG STAs can no longer communicate with each other. To maintain this beamformed link and ensure its availability and reliability to deliver frames, the IEEE 802.11ad defines a mechanism to monitor and maintain each beamformed link in the network. The mechanism implies running a beam link maintenance timer to monitor the activity of each beamformed link. This timer is reset upon successful exchange of frames between pair of DMG STAs with a beamformed link. The expiration of this timer is an indication of

the unavailability of the link, and it triggers performing beamforming training to find alternative sectors for communication.

2.3. Beacon Interval

All the WLAN standards organize access to the wireless medium in so-called Beacon Intervals (BIs). IEEE 802.11 amendments operating below 6 GHz start their BI by transmitting a single beacon frame in Omni mode to announce network existence. However, due to the directionality requirements for mmWave communication, IEEE 802.11ad transmits multiple directional beacon frames. Additionally, a DMG STA requires an initial access scheme to perform beamforming training and select a directional beam pattern for communication. For these reasons, the BI in IEEE 802.11ad is redesigned to adapt to the aforementioned requirements. In IEEE 802.11ad, each BI is subdivided into different access periods. An access period has different access rules and provides certain functionalities to nearby DMG STAs. Figure 2.2 illustrates a typical BI consisting of Beacon Header Interval (BHI) and Data Transmission Interval (DTI). The BHI compromises the following three sub-intervals:

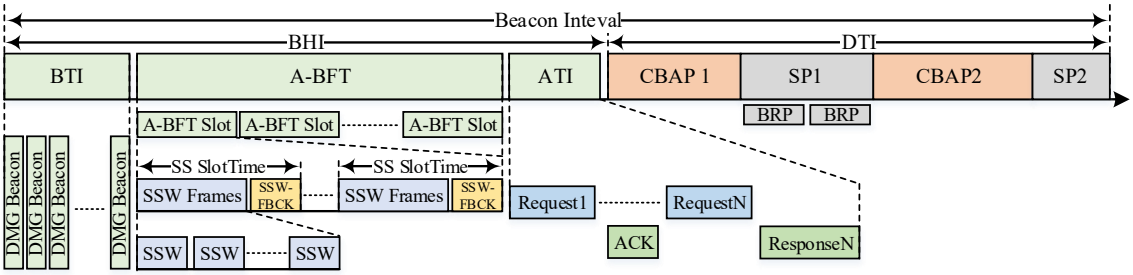

Figure 2.2: IEEE 802.11ad beacon interval with different access periods.

- **Beacon Transmission Interval (BTI)**: In this sub-interval, multiple DMG Beacon frames are transmitted across different sectors by the DMG PCP/AP to announce the network and provide transmit sector training towards nearby stations. DMG Beacons are transmitted using MCS 0 to reach large distances.

- **Association Beamforming Training (A-BFT)**: The A-BFT is used mainly by DMG STAs to train their transmit antenna sectors towards the DMG PCP/AP in a contention based manner.

- **Announcement Transmission Interval (ATI)**: This sub-interval is used mainly for management frame exchange between the PCP/AP and beam-trained STAs. Since communication takes place with beam-trained stations, stations can use high MCSs during the ATI for more efficient communication.

In the DTI period, DMG STAs exchange data frames either in the Contention-based Access Period (CBAP) or the scheduled service period (SP). During the CBAP, DMG STAs contend for the channel access using the IEEE 802.11 Enhanced Distributed Coordination Function (EDCF), whereas in SP, DMG STAs access the channel in a contention-free manner where the channel is reserved for communication between two dedicated DMG STAs.

2.4. Channel Access Schemes

IEEE 802.11ad defines a hybrid medium access [1, 12] scheme which complements the traditional and well-known access technology in the WLAN—the so-called CSMA/CA. It incorporates two new access technologies namely Service Period Channel Access (SPCA) and Dynamic Period Channel Access. Each one of these access techniques serves well for a particular application and scenario.

2.4.1. CSMA/CA Channel Access

CSMA/CA is the legacy distributed MAC mechanism in all the previous WLAN amendments because of its simplicity and robustness. A WLAN network utilizing CSMA/CA does not require a centralized scheduler to manage medium access. All STAs using CSMA/CA access the wireless medium with an equal probability. CSMA/CA supports Quality of Service (QoS) operations through Enhanced DCF channel access (EDCA) which creates virtual MAC queues for each Access Category (AC). Despite these merits, CSMA/CA is not well suited for mmWave wireless networks with their directional transmission and reception. The problem with directionality is that it forces a STA to listen to the medium in a certain spatial direction. This would not fit CSMA/CA operation which in principle requires determining if the channel is free in an omnidirectional manner. This results in stations determining that a medium is free while it is, in fact, busy, thus causing collisions. Since medium access is distributed, an AP does not know from which direction it shall receive so to avoid this problem an AP has to use a quasi-omni pattern for receiving which reduces link budget and the achievable data rate.

2.4.2. Service Period Channel Access

A SP in the context of IEEE 802.11ad wireless networks is designated for reserved time allocation for communication between a pair of stations. A SP is similar to time-division multiple access (TDMA) channel access in cellular networks. Station accessing the channel in a service period communicate with each other directly in a collision freeway. This increases MAC efficiency since no more time is wasted on contending for accessing the channel. Besides, it supports applications with strict QoS requirements. Moreover, it improves energy saving since stations can go into sleep mode whenever they are not scheduled for transmission. Since communication is scheduled, a station is aware of the peer station and thus can steer its antenna beam towards that

station and avoid the use of quasi-omni patterns. However, during the SP allocation phase, the quasi-omni patterns are still needed by the AP. The schedule of the service period allocations is communicated during the network announcement phase.

2.4.3. Dynamic Channel Access

Dynamic channel allocation in IEEE 802.11ad is based on a polling access scheme which is a master-slave protocol. The AP (the master) polls each station (slave) periodically for transmit requests. The polled station specifies the amount of time resources to meet its traffic requirements and QoS constraints. Because of the centralized approach, each station is aware of the direction of the transmission and thus can steer its directional antenna towards the master. This allows the station to avoid using the quasi-omni patterns. Unlike SPCA where modified channel allocations can be announced at the beginning of every BI, the dynamic schedule can be adapted during the course of a BI to react to bursty download traffic. Despite these benefits, the dynamic scheme wastes time resources because of the polling process.

2.5. Fast Session Transfer Technique

Since communication in the 60 GHz band is limited in range and suffers high penetration loss in case of obstacles, IEEE 802.11ad included a fast session transfer (FST) technique or multi-band operation support. With this, an IEEE 802.11 capable device can change its operational band seamlessly from 60 GHz to 2.4/5 GHz. As a result, a device can extend its coverage area and maintain its current sessions. As an example of this technique, a user may stream an Ultra High Definition (UHD) video on his/her device from a wireless docking station over the 60 GHz band when he/she is in the proximity of the docking station. As the user starts to move away from the docking station, signal quality starts to decrease, so the docking station decides to transfer the session to a lower band but continues video streaming using lower video encoding rates.

Currently, the FST technique supports switching between different frequency bands or wireless technologies only, and it utilizes a single wireless interface at a time. Exploiting more than a single wireless interface requires implementing a scheduling algorithm to distribute packets across these interfaces efficiently. Additionally, it mandates implementing a packet re-ordering mechanism to ensure in-order delivery to the higher layers in the protocol stack.

2.6. Relay Operation

The reliance on directional antennas in the mmWave band makes communication links sensitive to blockage and device orientation. For example, in the case of a crowded scenario, the directional communication link between a mobile user and an AP might experience frequent shadowing events within a small interval of time which leads to severe degradation to the link quality. To tackle this issue, IEEE 802.11ad introduces the relay operation mode to improve link resilience

against sudden interruptions. In this mode, two DMG STAs named source Relay Endpoint DMG STA (REDS) and destination REDS can communicate with each other with the assistance of a Relay DMG STA (RDS) which results in coverage area extension, and persistent multi-gigabit throughput. IEEE 802.11ad defines the following two types of relay operation modes:

- **Link Switching Type**: In this type, the source REDS maintains two links to the destination REDS: a direct link and a relay link through RDS. If the direct link is disrupted, the source REDS switches its transmission to the relay link. Communication over the disrupted link can resume once the direct link is recovered. Under this type, the RDS can operate either in full-duplex amplify-and-forward (FD-AF) mode or in half-duplex decode-and-forward (HD-DF) mode. In FD-AF mode, RDS amplifies the received frames and forwards them directly to the destination DMG STA. For this reason, the RDS must include two RF chains for sending and receiving frames simultaneously, which increases the cost and the complexity of these devices. This mode guarantees to maintain the original throughput, but on the other hand, it suffers from self-interference between the transmitting and receiving RF chains. In contrast, for the HD-DF mode the RDS receives multiple frames from the source REDS in one SP and forwards them to the destination REDS in the following SP. This mode is simple and requires a single RF chain; however, it decreases the achieved throughput by half.

- **Link Cooperating Type**: Contrary to the previous mode, in this mode the source REDS utilizes both direct link and relay link simultaneously to improve received signal quality at the destination REDS. Operating in this mode requires the source DMG to be aware of propagation delays over each link.

2.7. Spatial Sharing

Reliable communication in the mmWave band requires the use of directional antennas to compensate for the harsh propagation conditions. The directionality of the communication links allows simultaneous transmissions to take place at the same time without interfering with each other thus increasing the total throughput of the wireless system. To exploit this capability, the IEEE 802.11ad amendment defines a spatial sharing and interference mitigation mechanism to examine a set of established links for parallel communication. The mechanism relies on measuring the interference among communication links separated in time and space. An AP overlaps these links in time and assess the amount of interference introduced and its implication on the quality of the communication. If the AP foresees that the amount of interference induced to each link is negligible and does not affect the performance, then the AP schedule these links in the same time thus achieving spatial sharing and boosting system performance.

2.8. Clustering

Radio propagation characteristics in the 60 GHz band are harsh compared to the microwave band; thus radio coverage in the mmWave is typically confined within a single room. Extending the coverage would require dense deployments of APs to ensure a high probability of radio coverage to all users. However, managing a high number of distributed APs is troublesome. Besides, non-provisioned APs might result in high interference and performance degradation thus hindering the throughput achieved in the 60 GHz band. To mitigate the previous issue and improve spatial sharing between these co-channel APs, the IEEE 802.11ad amendment features clustering of distributed APs. The clustering feature allows co-channel APs to coordinate beaconing to avoid interference and enhance operation in dense environments. This allows a group of APs to schedule their transmissions in non-overlapping periods since each AP can receive scheduling information contained in the DMG Beacons from neighboring APs and thus decides on how to schedule its transmission in a way that avoids interference between these distributed APs.

Chapter 3

Networking Simulation Tool

3.1. Introduction

The IEEE 802.11ad amendment to the 802.11 standard for multi-gigabit communication at 60 GHz was published several years ago, but to date, no precise simulation model for networking in this band is available. In this chapter, we present a model for IEEE 802.11ad implemented in the network simulator ns-3. We model techniques that are essential for IEEE 802.11ad operation such as beamforming training, beamformed link maintenance, relay support, fast session transfer, dynamic and static channel access schemes, decentralized clustering, and spatial sharing. In addition, we provide the design and implementation of various types of beam codebooks. The beam codebooks allow a wireless device to incorporate multiple phased antenna arrays and manipulate their radiation patterns similarly to COTS 801.11ad devices. To improve simulation fidelity, we implement a Q-D wireless channel model in ns-3 that can interface with commercially available ray tracers. The Q-D channel realization allows the user to simulate complex propagation environments and frequency-selective channels. Additionally, the user can study the coexistence between 802.11ad devices and other systems that operate in the 60 GHz band such as IEEE 802.15.3c and IEEE 802.11ay standards. We then evaluate by simulation the performance of IEEE 802.11ad as well as the gains obtained through the aforementioned techniques. At the end of the chapter, we provide a brief introduction to the operations of SU/MU-MIMO techniques in the next-generation 60 GHz standard, the so-called IEEE 802.11ay. Additionally, we propose an implementation for simulating these techniques within network simulators. The material in this chapter is taken from [10, 13–15].

3.2. Background on Network Simulator ns-3

ns-3 is a discrete event network-level simulator written in C++ with bindings available in Python. It is widely used in academia for teaching networking classes and for research purpose. It includes a comprehensive set of networking and wireless modules to experiment with such as

LTE, WiFi, WiMAX, etc. Additionally, it provides simulation granularity at a packet level using high fidelity TCP/IP stack. Furthermore, ns-3 allows the integration of Linux TCP/IP stack and real applications through the Direct Execution Mode (DCE) module. This allows testing commercial-grade applications with fine-tuned TCP/IP stack over newly developed medium access technologies within ns-3. We decided to select ns-3 as a baseline for building our 802.11ad model for the following reasons:

- ns-3 provides an accurate implementation for the MAC layer of the IEEE 802.11 protocol with a good abstraction to the operations of the PHY layer. Specifically, the MAC layer implementation compromises various techniques including management, data, and control frame serialization and deserialization, frame aggregation, carrier sense multiple access with collision avoidance (CSMA/CA) access scheme, Enhanced DCF channel access (EDCA), Transmit Opportunity (TXOP), automatic transmission request, block acknowledgment, association state-machine, etc. Since the IEEE 802.11ad protocol reuses all of these techniques, utilizing the ns-3 WiFi module would facilitate and speed-up the development process.

- We want to study the performance of the wireless networking protocol in the 60 GHz band within dense network settings. This is not possible with link-level simulators as they are limited to single link performance. Whereas, in ns-3 we can study system-level performance for a high number of devices with heterogeneous capabilities. Additionally, with ns-3 we can understand the impact of mmWave channel variability on the operations of the upper layer of the protocol stack.

3.3. IEEE 802.11ad Model Implementation and Evaluation

In the following section, we provide an overview of the IEEE 802.11 model in ns-3 and how we augmented it to adhere to the IEEE 802.11ad amendment. Our model is publicly available on GitHub [16].

3.3.1. IEEE 802.11 Model in ns-3

The ns-3 WiFi model supports different IEEE 802.11 specifications such as a/b/e/g/n/ac/ax with an accurate implementation of the MAC layer. The model can be divided into the following four layers:

- **MAC High Layer**: It provides some Mac layer management entity (MLME) functionalities depending on the underlying network it supports, such as Infrastructure Basic Service Set (BSS) or Independent BSS.

- **MAC Low Layer**: This layer takes care of Ready-to-Send (RTS)/Clear-to-Send (CTS)/DATA/Normal Acknowledgment (ACK)/BlockACK transmission using the

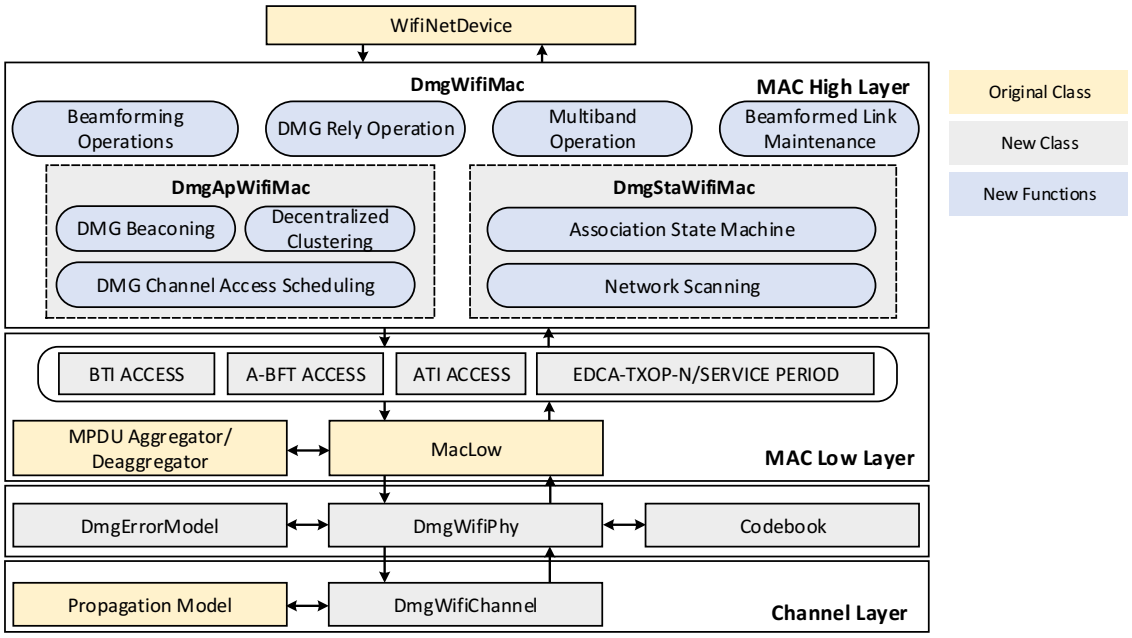

Figure 3.1: IEEE 802.11 architecture in ns-3 with 802.11ad Features.

DCF and EDCA functions. Additionally, it is provides both MAC service data unit (MSDU)/MAC protocol data unit (MPDU) aggregation and deaggregation capabilities.

- **Physical Layer**: It is a simplified model of the real Wifi PHY layer. This layer handles packet transmission and reception over the underlying channel. It calculates interference among different STAs and provides some probabilistic error model for packet reception.

- **Channel Layer**: This layer interconnects different PHY layers of different wireless STAs. Additionally, it simulates and models propagation effects that wireless signals encounter in real environments.

The IEEE 802.11 model in ns-3 is well suited for wireless technologies that use a CSMA/CA scheme with omni-directional transmission and reception. However, IEEE 802.11ad has characteristics that require some major changes to this model. Figure 3.1 shows the existing ns-3 IEEE 802.11 architecture together with the new blocks for the mechanisms introduced in IEEE 802.11ad. An abstract *DmgWifiMac* class provides common capabilities and techniques for DMG operation. These capabilities includes BF, relay operation support, FST. From this class, we derive two classes to represent different BSS types. The first class *DmgStaWifiMac* implements procedures specific to DMG STA such as TXSS in A-BFT, the association state machine, and network scanning. The second class, *DmgApWifiMac*, represents the DMG AP and provides DMG Beaconing, network synchronization, channel resource allocation, and decentralized clustering.

The following subsections provide an in-depth description of the implementation and design

assumptions for each block. Since ns-3 provides packet-level granularity, it was important to provide an accurate implementation of the newly introduced MAC frames and Wifi Information Elements to support various procedures defined in IEEE 802.11ad. Furthermore, representing the actual frame structure facilitates packet flow analysis using any network protocol analyzer that supports the IEEE 802.11ad extension.

3.3.2. DMG PHY Layer

ns-3 provides a simple PHY layer for the operation of IEEE 802.11. In this layer, the reception of the Physical Protocol Data Unit (PPDU) frame is modeled as simulation delay corresponding to the transmission time of this frame plus propagation delay. We follow the same approach in our implementation and we construct a new class named `DmgWifiPhy` which models transmitting and receiving different parts (Preamble, PHY Header, Payload, and TRn Field) of a DMG PPDU frame. Within this class, we provide all the mathematical formulas required for the calculation of PPDU frame transmission time using either control, SC, or OFDM PHY technologies. In addition, the formulas take into account if the transmitted frame has TRN field appended during either BRP or BT phases. All of the wireless the devices are connected to a common channel named `DmgWifiChannel`. This class models propagation effects on each part of an DMG PPDU independently. By default, when a wireless device transmits over this channel, all of the devices connected to the same channel will receive this frame. However, based on the received signal power, the receiver will decide if it can successfully decode the frame. Additionally, it will decide whether the medium is marked as busy or free based on the carrier sensing threshold.

3.3.3. AGC and TRN Fields Modeling

The IEEE 802.11ad standard defines a mandatory phase named BRP to refine the selected sector obtained from a previous SLS phase. The SLS phase is a time consuming process as a DMG STA has to probe with an independent SSW frame for each sector. Each SSW frame consists of 26 Bytes and takes around 14.9 μs when transmitted using the Control PHY. As a result, the duration to complete an SLS increases tremendously for massive antenna arrays, whereas in the BRP phase, evaluating a custom beam pattern takes around 713 ns. This duration is shorter since in the BRP phase, a DMG STA switches between multiple beam patterns during the course of single frame transmission. For this reason, DMG STAs typically probe few sectors with wide beams during an SLS phase to establish a directional communication link at a relatively low data rate. Later, in the BRP phase, each DMG STA refines its selected sector and generates a much narrower beam that yields a higher gain and thus higher throughput. As a result, combining both SLS and BRP phases significantly reduce the total duration to establish a directional link. However, utilizing the BRP protocol requires a phased antenna array with fast switching capabilities in the order of sub-nanoseconds. In contrast, in an SLS phase, switching speed is more relaxed as the standard gives a gap of 1 μs to switch a beam pattern before either DMG Beacon or SSW frame

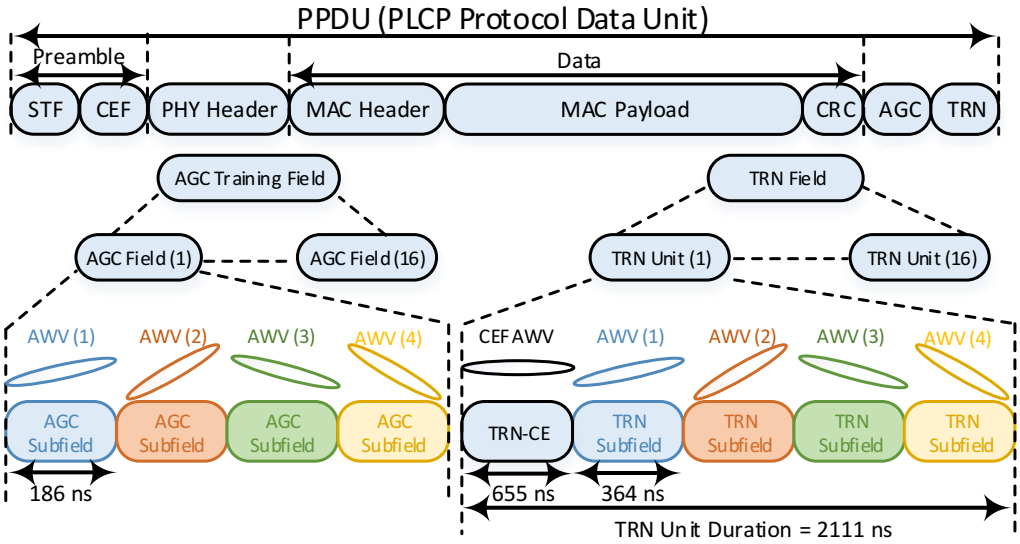

Figure 3.2: DMG Frame Structure with AGC and TRN Subfields for BRP-TX Packet.

transmission.

Figure 3.2 illustrates the general structure of a DMG packet with both AGC Field and TRN Field appended at the tail of a BRP-TX packet. A DMG packet can contain up to 64 TRN subfields. If the TRN Field is used for transmit training, then each TRN subfield is used to evaluate a custom antenna weight vector (AWV). Whereas for receive training, each TRN subfield is transmitted using the same AWV and the receiver switches its receive pattern at the beginning of each TRN subfield. The TRN Field can contain up to 16 TRN Units where each TRN Unit consists of a CEF that has a similar structure as the preamble followed by 4 TRN subfields. The preamble (STF and CEF), the PHY Header, and the Data Part are transmitted/received using the same AWV. The TRN Field is preceded by AGC Fields which are used to allow the receiver to reconfigure its AGC. As a result, each AGC subfield is transmitted/received using the same AWV used by its corresponding TRN subfield. By applying different AWVs, a station can steer its antenna array into different directions in space.

According to the IEEE 802.11ad standard, DMG STAs may utilize the SLS phase for Transmit beamforming (TxBF) and the BRP phase to perform Receive beamforming (RxBF). In our implementation, we provide the user with the flexibility to decide the role of SLS and BRP phases. The TRN Field can also be utilized for Beam Tracking (BT) during data transmission between two communicating devices. In this case, the TRN field is used for tracking the movement/rotation of the peer device and compensating for link quality degradation by choosing a different AWV which yields higher SNR value. Triggering the beam tracking during data transmission is implementation dependent and we leave the choice to the user through a custom TraceSource. We implement the corresponding state machines for transmitting and receiving AGC and TRN subfields as shown on Figure 3.3. The *Codebook* class takes care of iterating over the custom AWVs during AGC and TRN subfields transmission and reception.

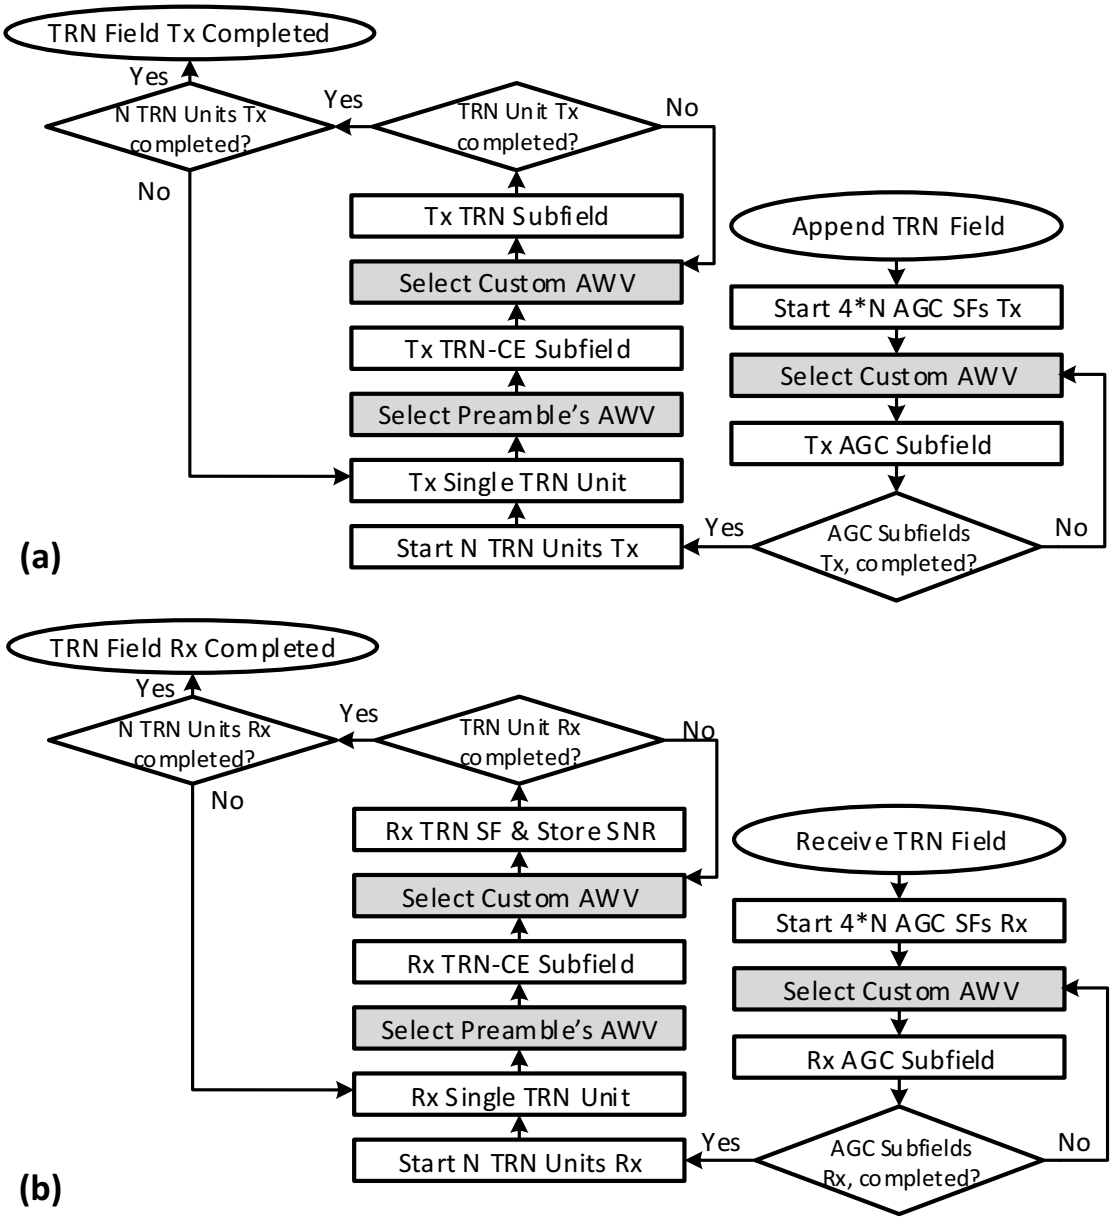

Figure 3.3: State Machines for AGC and TRN Subfields (a) Transmission and (b) Reception.

3.3.4. Quasi-deterministic (Q-D) Channel Model

The previous `DmgWifiChannel` class relies on a simple LoS propagation channel for communication, which neither captures the sparse structure of 60 GHz channels nor allows the user to customize scenario specific parameters. As a result, the performance of the simulated network is independent of the simulation environment and relies only on the geometrical distance between the wireless devices and the number of contending nodes. Measurement campaign analysis in the mmWave band suggests using the Q-D methodology to model the channel impulse response of a mmWave channel [17]. Using this approach, mmWave channel can be characterized using a set of Quasi-deterministic (Q-D) strong Rays (D-rays) that contain most of the power and a set of relatively weak Random Rays (R-rays) with predefined distribution: the *Q-D Channel* model. D-rays are obtained utilizing ray-tracing techniques while R-Rays are modeled in a statistical way.

We utilize a Q-D channel realization software [16] in MATLAB® developed by NIST to generate the spatial channel matrix between every node in the scenario. This software generates first D-Rays thanks to ray-tracing and then, based on material properties, produces the diffuse components associated with the D-Ray due to the rough surface scattering. For each pair of devices within the network, the Q-D channel realization software exports the following group of parameters in a Q-D channel trace file:

- The number of Multipath Components (MPCs). It includes the direct path, the specular, and the diffuse reflections.
- The pathloss in dB for each MPC.
- The phase shift due to propagation, reflection and Doppler Effect (in case of mobility) for each MPC, given in radians.
- The delay in ns for each MPC.
- The Angle of Arrival (AoA) in the azimuth and elevation planes for each MPC in degrees.
- The Angle of Aeparture (AoD) in the azimuth and elevation planes for each MPC in degrees.

It is worth mentioning that our ns-3 Q-D channel implementation is compatible with any Q-D channel realization software as long as the exported file respects the defined format.

Running ns-3 simulations using the Q-D channel model is a three-step process. First, the scenario is defined by specifying the environment geometry, the node locations, and the material properties. Then, the Q-D channel is realized, i.e., D-Rays and R-Rays are generated. Finally, the Q-D channel realization is used by the ns-3 Q-D channel implementation to compute the received power between communicating nodes.

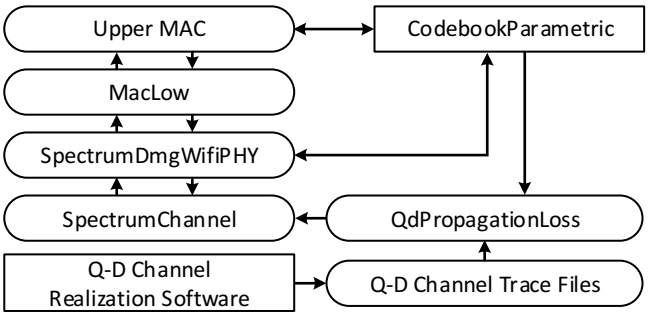

Figure 3.4: Q-D Channel Integration with ns-3 802.11ad Model.

Each Q-D channel trace file allows for a pair of nodes to compute the received power using the Q-D channel. The received power y, when the transmit power is x, is given by Eq. 3.1.

$$y = h_{BF}x \tag{3.1}$$

where h_{BF} is the CIR of the beamformed channel and is given by Eq. 3.2 as defined in [18].

$$h_{BF} = \sum_{i=0}^{N-1} \left(A^{(i)} \left((U_{ch}^{i})^{H} U \left(V_{ch}^{i} \right)^{H} V \right) e^{-j2\pi f t_i} \right) \tag{3.2}$$

where N is the number of MPCs and $A^{(i)}$ is the complex amplitude of the i^{th} MPC. The terms U and V refer to the receiver and transmitter antenna weight vectors respectively. U_{ch}^{i} and V_{ch}^{i} are beamsteering vectors for AoA and AoD respectively, of the i^{th} MPC, as given by $S(k, \theta, \phi)$ in Eq. 3.7. The parameter f is the operational frequency and t_i stands for delay of the i^{th} path. Complex amplitude of i^{th} MPC is given by Eq. 3.3:

$$A^{(i)} = 10^{-PL_i/20} e^{j * phase_i} \tag{3.3}$$

where PL_i and $phase_i$ denote the pathloss and phase shift of the i^{th} path, respectively. Using Eq. 3.6, the CIR can be reduced to Eq. 3.4

$$h_{BF} = \sum_{i=0}^{N-1} \left(A^{(i)} \left(Y_{Rx}^{(i)} Y_{Tx}^{(i)} \right) e^{-j2\pi f t_i} \right) \tag{3.4}$$

where $Y_{Rx}^{(i)}$ and $Y_{Tx}^{(i)}$ is the radiation pattern of the receiver and transmitter array at the i^{th} MPC respectively. It is worth mentioning that the Q-D channel only works with the *ParametricCodebook* as it requires the phased antenna array response.

Figure 3.4 shows a simplified architecture of the integration of the Q-D channel model inside ns-3. Our implementation makes use of the ns-3 Spectrum Model framework. `SpectrumDmgWifiPhy` class (derived from `DmgWifiPhy`) allows a power spectral representation of the 802.11ad signal. We divide the whole 2.16 GHz channel into a number of equally spaced sub-bands. The size of each sub-band corresponds to the sub-carrier spacing for an OFDM PHY which is 5.156 MHz. The power allocation depends on the transmission mode: For Control and Single Carrier mode, we divide the power equally across all the sub-bands while in OFDM mode, the three DC carriers (-1, 0, and +1) are suppressed (zero power), and the power is distributed over the remaining 352 sub-bands.

The newly added `QdPropagationLoss` class (derived from `SpectrumPropagationLossModel`) parses all the Q-D trace files in order to obtain the spatial matrix between every node. For each node, the `QdPropagationLoss` spectrum propagation loss model is added to the Spectrum Channel. It computes the received power per sub-band as described in Eq. 3.1. The power per sub-band is finally turned into a scalar value to represent the total energy apparent to the receiver by applying RF filtering. Our implementation uses the default ns-3 `WifiSpectrumValueHelper` RF filter that applies a flat frequency response over the band. The received power is thus obtained using Eq. 3.5:

$$P_{Rx} = \sum_{i=0}^{N_{subband}} \Big(Rx_i * bw_i \Big) \tag{3.5}$$

where $N_{subband}$ stands for the number of sub-band, Rx_i is the power received for the i^{th} sub-band after applying the filtering, and bw_i is the i^{th} sub-band width.

3.3.5. Beam Codebooks

We introduce a class named `Codebook` which provides high flexibility to define and manipulate all the beamforming aspects of DMG STA. The `Codebook` class encapsulates all the functionalities related to beamforming training (e.g., initiating beamforming, tracking active antenna and sector, switching among sectors) and decouples it from from the *DmgWifiMac* class. Additionally, the class provides the user with the capability to attach multiple antenna arrays per DMG STA to properly simulate IEEE 802.11ad characteristics (e.g., multiple antenna arrays improve the resiliency against blockage typically encountered in the 60 GHz band [19]). This not only improves code usability but also lays the foundation for a future IEEE 802.11ay implementation which relies on multi-antenna communication.

The implementation of the `Codebook` class is inspired by the design of the board file used by the firmware of the QCA9500 [20] chipset from Qualcomm. Figure 3.5 depicts the general architecture of the *Codebook* class. A codebook is designed for a number of RF chains where each RF chain is connected to one or more phased antenna arrays. Only a single antenna array is active per RF chain. This emulates the *Hybrid Beamforming with Partially Connected RF Architecture* [21]. While an 802.11ad device is limited to a single RF chain where up to four antenna arrays can be connected, our implementation allows to have more than one RF chain to facilitate 802.11ay integration with our codebook architecture. Indeed, 802.11ay devices can contain up to 8 digital RF chains. Each phased antenna array in turns consists of a number of predefined sectors where each sector is composed of a set of custom antenna weight vectors (AWVs). The user can decide whether to pre-define those AWVs in advance or in an adaptive manner during the course of the simulation based on the channel measurements. The `DmgWifiHelper` class attaches a `Codebook` instance to each `WifiNetDevice` during the initialization phase. The `Codebook` class contains a list of virtual `PhasedAntennaArrayConfiguration` structures which represent the phased antenna arrays. For each phased antenna array structure, the

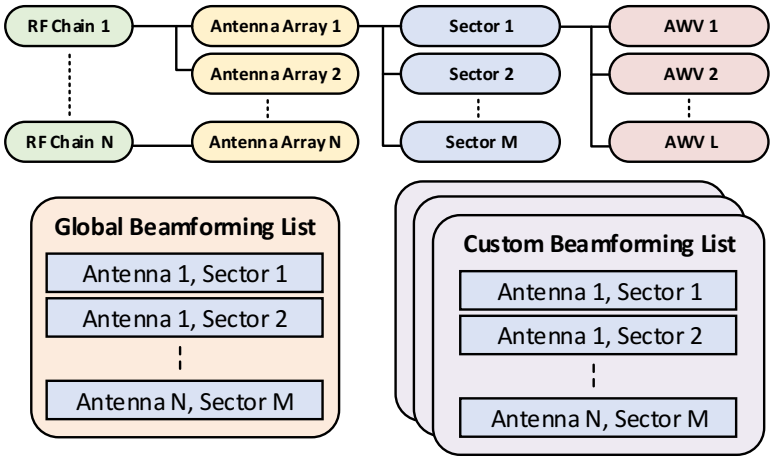

Figure 3.5: Codebook Architecture with Beamforming Lists.

user must define the following set of attributes:

- A list of sectors using the virtual `SectorConfig` structure. The standard limits the number of sectors per antenna array to 64 sectors. The maximum total number of sectors across all phased antenna arrays is 128.

- The orientation of the phased antenna array in both the azimuth and the elevation planes.

Each sector takes the following set of parameters:

- Sector type indicates whether the sector is used for transmission, reception, or both.

- Sector usage to specify whether the sector is used during BHI, SLS, or both.

- A list of virtual `AWV_Config` structures representing custom AWVs. These AWVs will be probed during both BRP and BT phases. The number of AWVs should be multiple of 4 according to 802.11ad [22].

Based on the *Sector usage* and *Sector type* parameters, the codebook generates multiple generic beamforming training lists that facilitate the beamforming operations during both BHI and SLS phases with all the DMG STAs. Besides, the user can customize beamforming training lists and generate a unique list per peer device. Both `SectorConfig` and `AWV_Config` structures are derived from a parent structure named `PatternConfig` since both structures represent the shape of the beam pattern generated by the phased antenna array. Our newly introduced codebook architecture can use three different types of beamforming codebook, based on both the level of fidelity and the complexity required by the simulations. Each codebook has its unique definition for each one of the previous structures.

3.3.5.1. Analytical Codebook

The analytical codebook is the most-basic codebook implementation where the user can either use it to divide the azimuth plane into a number of equally spaced sectors, or manually define the radiation pattern of each sector using a predefined formula. Our implementation uses the Gaussian directional pattern formula to define the shape of our beam patterns [23].

3.3.5.2. Numerical Codebook

With the numerical codebook, a user can import custom radiation patterns obtained through either commercial simulators like HFSS, antennas data-sheets, or antenna beam pattern measurements. We provide a MATLAB® script that imports the radiation patterns [2] of the TP-Link TALON Router AD7200 and generates a codebook file that can be used by the numerical codebook. The TALON router uses a single beam pattern for reception (quasi-omni pattern), 32 sectors for transmitting DMG Beacons, and 35 sectors for transmit beamforming training in the SLS phase. Both the `Analytical Codebook` and the `Numerical Codebook` classes interface with the `DmgWifiPhy` class only.

3.3.5.3. Parametric Codebook

The parametric codebook is a state-of-the-art beamforming codebook that mimics the behavior of 802.11ad COTS devices. This is the most-realistic and complicated codebook implementation in our model. A user can synthesize beam shapes with high granularity and flexibility as it relies on phased antenna array theory. In theory, the radiation pattern Y of an antenna array is computed using Eq. 3.6:

$$Y(k,\theta,\phi) = R(\theta,\phi)AF(k,\theta,\phi) \tag{3.6}$$

where $R(\theta,\phi)$ is the radiation pattern of a single element in the array considering the array has identical antenna elements, $AF(k,\theta,\phi)$ stands for the array factor, i.e., the antenna response, and k is the wave number. Both radiation pattern and array factor depend on θ and ϕ, i.e., the azimuth and elevation angle respectively. The wave number k is equal to $2\pi/\lambda$ where λ is the wavelength. The array factor is obtained using Eq. 3.7.

$$AF(k,\theta,\phi) = W^T S(k,\theta,\phi) \tag{3.7}$$

where W^T is the transpose of the antenna weights vector, the antenna weight vector representing the phase and amplitude excitation between antenna elements, and $S(k,\theta,\phi)$ is the steering vector which represents phase delay among antenna elements for each incoming electromagnetic plane wave impinging on the antenna array. The steering vector $S(k,\theta,\phi)$ depends on the geometry of the phased antenna array, the placement of the antenna elements, and the carrier frequency. In theory, the steering vector should have a constant amplitude across all the elements. In practice, however, the amplitude of the steering vector has non-uniform values due to the mutual coupling between antenna elements and non-idealities in the antenna design.

To make our implementation agnostic to the type of phased antenna array, the `Parametric Codebook` requires the user to provide the following group of parameters:

- The steering vector of the antenna array in the form of a 3D complex matrix of size $M * N * L$, where L is the number of antenna elements. M and N represent the number of azimuth and elevation angles. In our implementation, we limit the angle resolution to one degree. In addition, the steering vector is considered constant across the whole frequency range of the channel.

- The radiation pattern of a single antenna element as a 2D complex matrix of size $M * N$.

- A list of complex AWVs defining the excitation of the antenna elements. The size of each vector is L.

To facilitate the generation of those parameters and the usage of this Codebook, we developed a MATLAB® application named *IEEE 802.11ad Codebook Generator* [16]. With this application, the user can choose either among a set of MATLAB® phased antenna array models such as Uniform Linear Array (ULA), Uniform Rectangular Array (URA), and Uniform Circular Array (UCA) or upload a custom steering vector obtained from antenna measurements. In addition, the user can visualize the 2D/3D directivity of the antenna array and define the parameters of each sector and custom AWV independently. When using MATLAB® phased antenna array models, the user can select between different sets of antenna elements such as omni-element, microstrip, etc. In contrast, in the case of custom phased antenna arrays, the steering vector already comprises the effects of the antenna element pattern. The ns-3 `Parametric Codebook` class parses the file generated by the application and automatically calculates and stores the directivity of each sector and AWV according to Eq. 3.6. The decoupling of the antenna response and the weights of the antenna elements allows the user to adapt the beam pattern within the course of the simulation based on the state of the wireless channel. For example, the user can suppress interference and have a null in the direction of a particular station while increasing the power towards the intended receiver. Simulations utilizing the `Parametric Codebook` have high computational overhead due to the matrix operations which in turn increases the runtime compared to the other two codebooks.

In addition, the user can emulate the procedure of CSI measurements supported in the latest firmware (5.2.0.15) [24] of the Qualcomm chipset. The firmware utilizes the BRP protocol to report the complex Channel Impulse Response (CIR) of the strongest path for each antenna element in the antenna array. Since ns-3 does not model symbol-level signals and only provides packet-level simulations, the same procedure would provide the user with the SNR measurements per antenna element.

Using either the *Numerical Codebook* or the *Parametric Codebook* provides high fidelity simulations and results. Indeed, the mutual coupling among antenna elements and non-idealities

in the electronics are not reflected in the analytical models (e.g., the mutual coupling among antenna elements has a significant impact on the resulting beam pattern).

In practice, when a phased antenna array is deployed in a field, the radiation pattern of the phased antenna array differs from the one reported in the simulations or the antenna data-sheets. The reflections coming from a real environment might interact with the signal coming from the direct path in a destructive way, thus attenuating the signal. To simulate this behavior, we integrate a *Q-D Channel Model* in our implementation that interfaces with the *Parametric Codebook* as described in Section 3.3.4.

3.3.6. Multi-antenna Beamforming Training Support

Commercial-grade APs operating in the mmWave band are envisioned to incorporate multiple antennas to provide ubiquitous coverage for clients in the network. The multi-antenna support is not limited to APs only. Qualcomm proposed the inclusion of two antenna arrays in mobile handsets to provide resiliency and increase spatial diversity against self-blockage [19]. Furthermore, IEEE 802.11ay mandates all the devices to have multi-antenna to support MIMO communication. We extend the beamforming training procedure in the current implementation to support multi-antenna beamforming training during BTI, A-BFT and DTI channel access periods for both transmit and receive beamforming training.

For beamforming training in the BTI and A-BFT access periods, the IEEE 802.11ad standard states that only a single phased antenna array is used for probing SSW frames to obtain the best sector for communication within the same BI. However, since this sector might not be the best one among all the phased antenna arrays, a device can send a request to the AP to allocate a SP during DTI to continue beamforming training using the remaining phased antennas arrays. Once a DMG STA has contended successfully in an A-BFT access period and obtained the best sector for transmission, it does not participate further in any future A-BFT access periods. This reduces contention in A-BFT and allows more stations to associate successfully with the AP. A device can later switch its current active phased antenna array at the beginning of the next BI to listen for incoming DMG Beacons. This ensures that if one antenna array is blocked and cannot detect those DMG Beacons, another antenna array with different field of view (FOV) can take over.

For beamforming training in the DTI period, the two DMG STAs involved in the beamforming training need to be aware of the beamforming capabilities of each other. This includes first reporting both the number of supported transmit and receive antennas, and the number of transmit and receive sectors in the *DMG Capabilities Element* to the AP. Then, the beamforming initiator can retrieve this information from the AP and calculates the exact duration of an SP allocation to complete the beamforming training with the beamforming responder. In our implementation, we consider all the possible beamforming training scenarios based on the number of antennas of both the initiator and the responder and the type of the beamforming training, either TXSS or RXSS. For brevity, Figure 3.6 shows the general case when both initiator and responder have multiple antenna arrays.

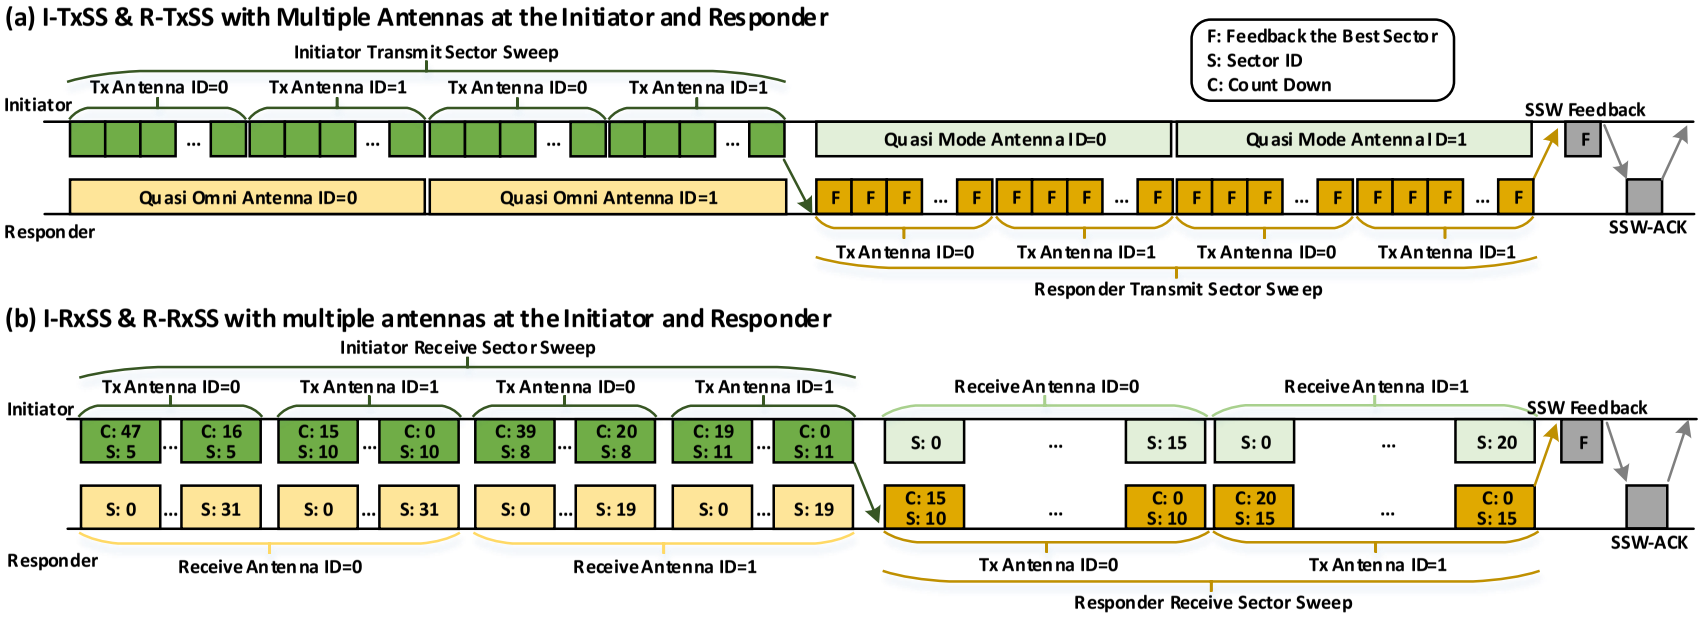

Figure 3.6: Multi-antenna Beamforming Training Examples: (a) I-TxSS & R-TxSS with Multiple Antennas at the Initiator and Responder (b) I-RxSS & R-RxSS with Multiple Antennas at the Initiator and Responder.

In Figure 3.6(a), both the initiator and the responder perform TXSS during an SLS phase. During the I-TxSS phase, the responder activates one of its antennas and sets it to receive in quasi-omni mode. At the same time, the initiator starts probing SSW frames across all its antenna arrays until it has covered all its sectors. Then, the responder switches to the next antenna array and the initiator sweeps again across all its antenna arrays. The operation continues until the responder has used all its antenna arrays. In the responder TXSS (R-TxSS), both the initiator and the responder exchange their roles and perform the same operations as in the I-TxSS phase. In the end, both the initiator and the responder obtain the best transmit sector for communication across all their antenna arrays. In Figure 3.6(b), both the initiator and the responder perform RXSS during an SLS phase. It is worth mentioning that in order to perform the RXSS phase, a DMG STA should have successfully realized the TXSS phase. In the initiator RXSS (I-RxSS) phase, we train the receive antenna of the responder. For each receive antenna at the responder side, the initiator iterates over its transmit antennas. For each transmit antenna, the receiver switches among its receive sectors while the initiator transmits SSW frame using the best transmit sector for this particular antenna. In the responder RXSS (R-RxSS) phase, we perform the same operations as before but to train the receive antenna of the initiator.

3.3.7. DMG Error Model

We build a robust and realistic error model inside ns-3 named `DmgErrorModel`. The model compromises several lookup tables (LUTs) for mapping SNR to BER for each MCS defined in the IEEE 802.11ad standard. Starting from MATLAB® R2017b, the *WLAN Toolbox* provides the capability to perform an accurate end-to-end link-level simulations for an 802.11ad PHY layer. We generate these LUTs by performing the following steps:

1. First, we generate an 802.11ad compliant frame for one of the following physical layers: Control, SC, and OFDM.

2. Then, we transmit the constructed frame over an Additive white Gaussian noise (AWGN) channel using specific SNR value. We assume perfect frame synchronization and perfect channel estimation. Additionally, we omit any physical layer impairments such as phase noise, carrier frequency offset, etc.

3. Finally, we compare the content of the received frame with respect to the content of the original transmitted frame. If there is a difference in the content, then we increase the number of erroneous packets by one.

To ensure a high confidence in our results, we generate each LUT for a particular MCS until either successfully decoding 20K packets or reaching 2K erroneous packets whichever comes first. We generate these LUTs using a discrete set of SNR values with a step size of 0.2 dB. We store these LUTs using comma-separated values (CSV) file format. The `DmgErrorModel`

parses the CSV file and uses a simple linear interpolation method to predict the Packet Error Rate (PER) values from non-existing SNR values.

3.3.8. DMG Access Periods

The `DmgApWifiMac` class organizes medium access by initiating BI through transmission of DMG Beacons across all its antenna sectors. The remaining time for each access period is announced in the duration field of the MAC header. This allows DMG STAs to synchronize their clocks with the DMG AP clock. During BTI, the DMG AP ensures the medium is free before it starts DMG beacon transmission. For this reason, the value of the duration field is calculated once the DMG AP is granted access to the channel. A DMG STA that receives at least one DMG beacon from the DMG AP schedules an event to start the A-BFT access period at the end of the current BTI. The DMG APs divides the A-BFT into slots, where the duration of each slot is calculated based on the number of SSW frames to be transmitted. DMG STAs choose one of these slots randomly using a uniform distribution. If two DMG STAs select the same slot, they will collide and do not receive an SSW-Feedback (FBCK) frame within a pre-determined period of time. These two DMG STAs then have to select a new slot while ensuring that they do not exceed the duration of the A-BFT. This period is followed by the ATI access period, where the DMG AP initiates management frame transmission. Currently, we use this period to perform the BRP setup phase and exchange BRP transactions. Any packet that arrives during the previous access periods is queued for transmission until the beginning of the DTI.

3.3.9. DMG Channel Access Schemes

3.3.9.1. CSMA/CA Channel Access

We modify DCF function to support operation using directional transmission and reception in IEEE 802.11ad. During a transmission, a STA should steer its antenna beam towards the intended receiver. Whereas if the station is in idle state, it configures its receive antenna to quasi-omni pattern to receive any frames transmitted by nearby STAs covered by this antenna. A special case is when an DMG STA expects frames exchange only with the PCP/AP, then it directs its receiving antenna towards the PCP/AP. In addition, we add the support for RTS/DMG CTS control frames exchange to protected against hidden terminals. The difference between legacy CTS and DMG CTS is that the latter incorporates Transmitter Address (TA) field.

3.3.9.2. Service Period Channel Access

To allocate a SP, a non-PCP/non-AP DMG STA sends an `ADDTS Request` to the PCP/AP. The `ADDTS Request` carries `DMG TSPEC` element which identifies the source Association Identifier (AID) and the destination AID of the allocation. An AID is a unique identifier that distinguishes a STA within BSS. It is assigned to the STA upon successful association with the

PCP/AP. In addition, the `DMG TSPEC` characterizes the SP allocation in terms of duration and format. An allocation can compromise multiple SPs. The format describes the type of the traffic as either isochronous or asynchronous. An isochronous traffic type satisfies applications with periodic payload such as wireless display applications. In contrast, the asynchronous traffic type is convenient for applications with non-periodic traffic, e.g, rapid file download.

The PCP/AP either accepts or rejects the request based on the employed admission control policy and the available time resources in the DTI access period. The PCP/AP sends `ADDTS Response` to the STAs identified as source and destination. If the request is accepted, the PCP/AP announces the schedule of the allocation in next the DMG Beacon or Announcement Frame. In addition, the PCP/AP decides the position of the SPs that makes up the allocation within the BI. Communication in a service period is allowed in uni-directional direction from source DMG STA towards destination DMG STA. The source DMG STA holds the channel and initiates the transmission of data frames towards the destination DMG STA which stays in receive mode.

In our implementation, we leave the decision of allocating the requested resources in the `ADDTS Request` to the user. This is done by adding a new trace source *ADDTSReceived* in the `DmgApWifiMac` class. Once the PCP/AP receives an `ADDTS Request` it invokes this trace source. This trace source provides two parameters: the MAC address of the requesting STA and the allocation characteristics. The user accumulates all the requests for SP allocations and provides an implementation for admission and resource scheduling algorithms.

3.3.9.3. Dynamic Allocation of Service Period

The dynamic allocation mechanism in IEEE 802.11ad takes place during DTI access period. It compromises two periods: the Polling Period (PP) and the Grant Period (GP). During the PP, the PCP/AP polls each station that has declared its willingness to participate in the dynamic allocation. A station declares its wish to participate by setting the `PP Available` in the `STA Availability` element to true. The previous block of poll frames is answered by a series of Service Period Request (SPR) frames. In the SPR frame, each station declares its resource requirements to satisfy its traffic constraints. Then, the PCP/AP allocates resources based on these requests and announces the allocation in a separate GP period. STAs are allowed to transmit frames exclusively during the allocated time in the GP period.

Similar to Section 3.3.9.2, the admission of the requested resources is left to the user. This provides the user with flexibility to design its custom resource scheduler. To accomplish this task, we provide one call back *RegisterSPRequestFunction* in the `DmgStaWifiMac` class and one trace source (*PPCompleted*) in the `DmgApWifiMac`. The *RegisterSPRequestFunction* callback is invoked each time the PCP/AP polls a station. A user registers a function to this callback and declares the resources required in a dynamic way. The PCP/AP triggers the trace source (*PPCompleted*) upon the completion of PP period. At this point, the PCP/AP has received all the resource requests and the user can decide the allocation of the resources in the following GP period. Fig-

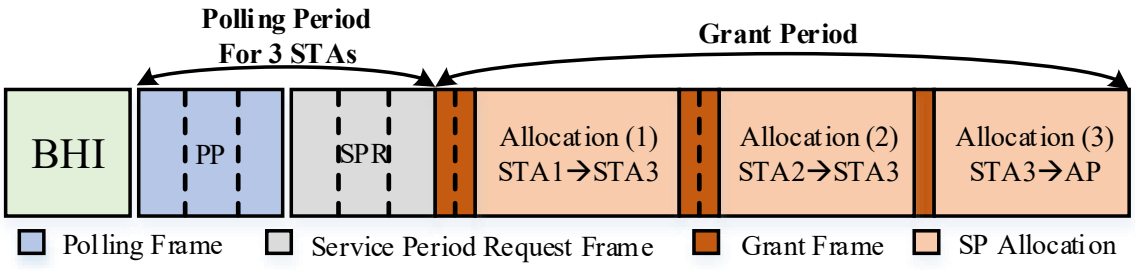

Figure 3.7: Example of dynamic channel allocation for 3 STAs.

ure 3.7 illustrates an example of a PCP/AP polling 3 STAs for transmission opportunity.

3.3.10. Beacon Generation in Infrastructure BSS

To minimize the interface caused by the transmission of DMG Beacons to nearby PCPs/APs, a PCP/AP changes the sequence of directions through which a DMG Beacon is transmitted at the beginning of each beacon interval. The PCP/AP chooses the sequence of directions pseudo-randomly from a set of directions that fully cover all the spatial directions. Figure 3.8 depicts an example of DMG Beacon transmission by PCP/AP for a set of N directions.

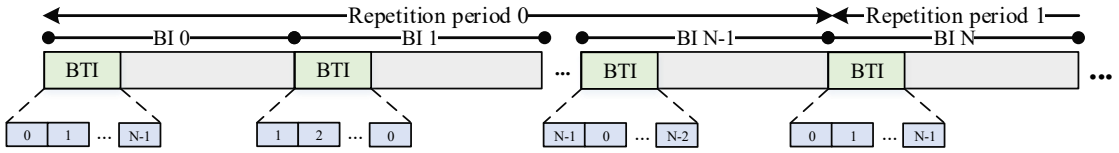

Figure 3.8: DMG Beacon transmission by PCP/AP during the BTI [1].

3.3.11. Synchronization

Network synchronization within a DMG BSS is extremely important to ensure that all the STAs access the wireless channel in their corresponding allocation slots. Upon network initialization, a DMG AP starts a local timer known as timing synchronization function (TSF) which has a microsecond resolution. All of the DMG STAs associated to the same DMG AP should update their local TSF to match the one running in the DMG AP. As a result, all the DMG STAs will be aware of the exact start and end of each access period in the BI. Without correct synchronization, a DMG STA might access the channel in a wrong period which results in high collision and wrong network operations. The DMG AP is responsible for ensuring synchronization within a network through a periodic transmission of DMG Beacons that carry synchronization information. A DMG Beacon contains the following information:

1. Remaining duration in the BTI access period.

2. The length of the A-BFT access period.
3. The duration of the ATI access period.
4. The remaining number (CDOWN) of DMG beacons until the end of the current BTI access period. The last DMG Beacon will have its CDOWN subfield set to zero.
5. The length of the BI in microseconds.
6. Timestamp indicating the value of the TSF timer when we start MPDU transmission (MAC HDR + MAC Payload).

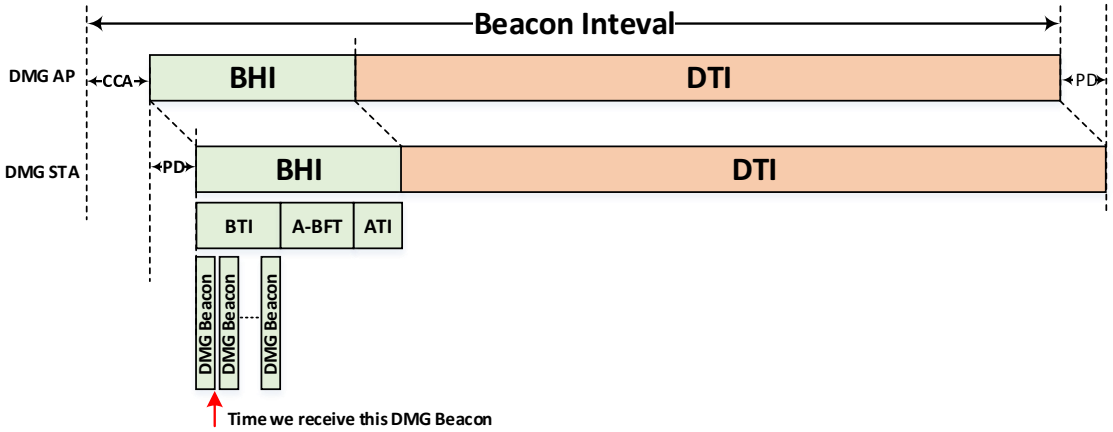

Figure 3.9: Synchronization in DMG BSS.

Figure 3.9 depicts delays within a BI at the DMG AP and DMG STA. A DMG AP performs the Clear Channel Assessment (CCA) procedure to ensure that the wireless medium is idle before starting its DMG Beacon transmission. A DMG AP must ensure that the BHI does not overlap with static allocations in the DTI access period because of the CCA procedure. At the DMG STA, the amount of time that has elapsed since the actual start of the BI at the DMG AP is:

$$t = TxTime(DMGBeacon) + CCA + PropagationDelay(PD) \tag{3.8}$$

Upon the reception of a DMG Beacon, a DMG STA update its local TSF Timer using the time stamp information contained in the received DMG Beacon. However, all the DMG STAs will have their local TSF drifted from the one in the DMG AP by an amount equals to the Propagation Delay (PD). The PD is extremely small and is in the order of nanoseconds. In our implementation, we do not model any of the following delays:

- PHY processing delay which includes frame encoding and decoding.
- MAC processing delay which includes frame processing and response preparation.
- PHY Tx Latency: Radio turn around time to switch from reception and transmission.

3.3.12. DMG Beamforming Operation

We provide a generic implementation for both SLS and BRP phases in the `DmgWifiMac` class. The implementation can be used either as part of the initial BT between DMG AP and DMG STAs, as a scheduled SP between two DMG STAs, or within a CBAP in the DTI between a DMG STA and DMG AP. The `DmgWifiMac` class has two data structures: one for storing and mapping antenna sector configurations for each received frame from a peer DMG STA with its corresponding SNR and another data structure for storing the best transmit and receive antenna sectors towards a particular station. The latter is updated at the end of each BF operation. In the current implementation, all decisions regarding the best transmit and receive antenna sectors are based on SNR measurements. The `MacLow` class uses the second data structure to select the antenna sector based on the receiver address in the MAC header. In the current implementation, we assume DMG STAs perform TXSS in the A-BFT. In addition, the BRP phase by default will be utilized to train antenna receive sectors for all DMG STAs instead of refining the selected antenna transmit sector. All BF frames are transmitted using MCS-0.

3.3.13. Fast Session Transfer

IEEE 802.11ad supports multi-band operation for fast session transfer (FST). FST operation can be either in transparent or non-transparent mode. In transparent mode, all MAC sub-layers in the STA expose a single MAC-Synchronization AP (S-AP) to the upper layers, i.e., a single MAC address. In non-transparent mode, each MAC sub-layer exposes its own MAC-S-AP to the higher layers which adds more complexity. In our implementation, we use the transparent mode where we design a new `NetDevice` named `MultiBandNetDevice`. This new `NetDevice` encapsulates different IEEE 802.11 technologies as depicted in Figure 3.10. For each technology, a user defines a `WifiMac`, `WifiPhy`, `WifiRemoteStation` and `WifiChannel` objects. One technology should be active at any point for any pair of devices. A STA that supports multi-band operation should announce this in its Beacon, Association Request, Association Response, Probe Request, Probe Response, and DMG Beacon a MultiBand Information Element.

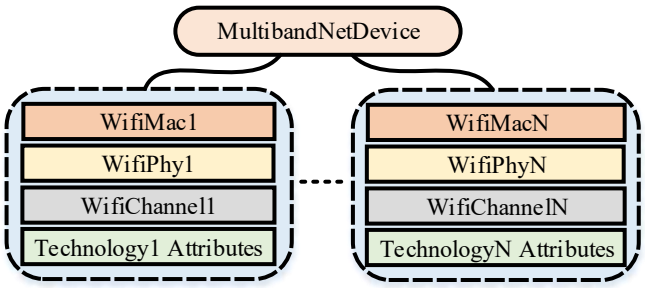

Figure 3.10: MultiBandNetDevice implementation in ns-3.

Figure 3.11 illustrates various states a STA goes through to establish a unique fast ses-

sion transfer session (FSTS) ID with a peer STA. At the beginning, each STA is in the **INITIAL_STATE** where they communicate in the old band/channel. A station that wishes to set-up a FSTS is called FST Initiator and the peer station is FST Responder. To proceed to the **SETUP_COMPLETION_STATE** and obtain a unique FSTS ID, both Initiator and Responder have to exchange FST Setup Request/Response frames successfully. In this new state, STAs keep communicating in the old band/channel. However, depending on the value of the link loss timeout (LLT) field in the Session Transfer Information Element, both STAs shall either transfer their current session to the new band/channel immediately if the value of LLT is equal to zero, or they shall start a Link Loss countdown equal to LLT $*$ 32μs if LLT $>$ 0. The LLT defines the amount of time that has to elapse since the initiating STA received an MPDU frame from the responding STA until the initiating STA should perform FST. Once the value of LLT reaches zero, both Initiator and Responder move to the **TRANSITION_DONE_STATE** and start communicating in the new band/channel. If the two STAs exchange normal MPDU frames or FST ACK Request/Response in the new band/channel successfully, the two STAs move to the **TRANSITION_CONFIRMED_STATE**, otherwise the two STAs move to the **INITIAL_STATE** and resume communication in the old band/channel.

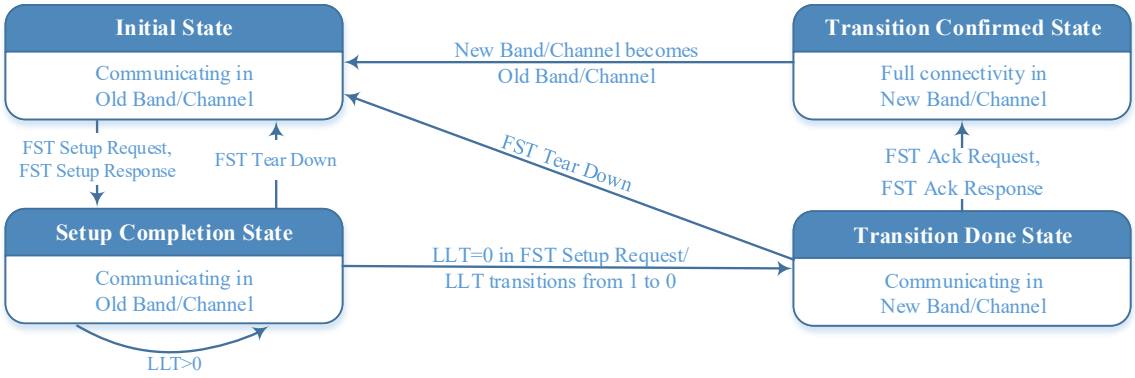

Figure 3.11: FST state machine.

3.3.14. DMG Relay Operation

We implement link switching type relay operating in either HD-DF or FD-AF mode. In addition, we incorporate frame exchange rules during a service period allocation as defined in the amendment for both FD-AF and HD-DF relay modes. At the beginning, a source REDS and a destination REDS sets up a relay link with a RDS following the Relay Link Setup (RLS) signaling procedure. Later, if the PCP/AP receives `ADDTS Request` frame with the `Source AID` and the `Destination AID` fields within the `DMG TSPEC` element are equal to the previous pair of source REDS and destination REDS then frame exchange rules follow the ones depicted in Figure 3.12.

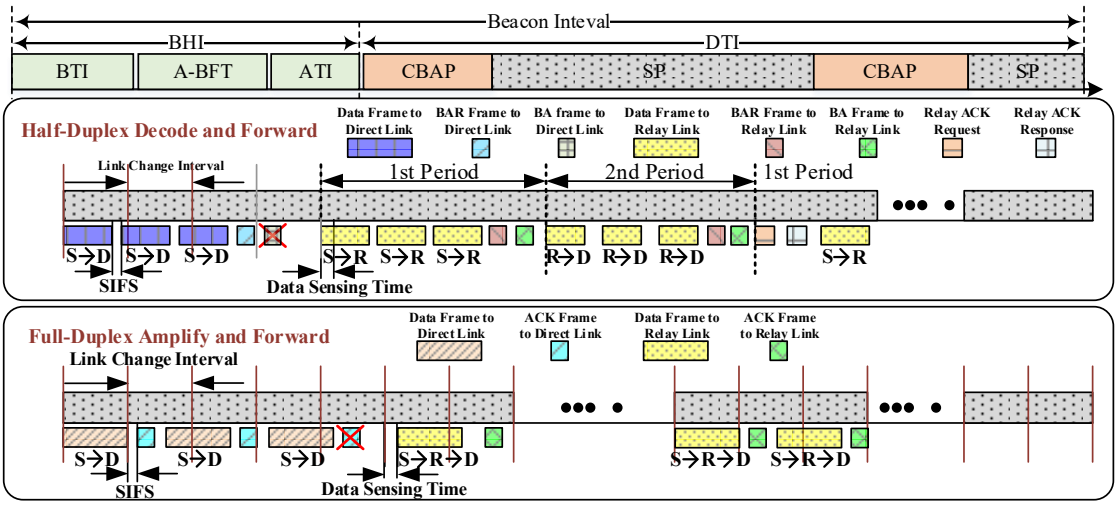

Figure 3.12: Relay Operation Modes in IEEE 802.11ad [1].

3.3.14.1. Common Frame Exchange Rules

Following the completion of the RLS procedure and SP allocation, the contiguous SP shall be divided into *Link Change Intervals*. This interval indicates when the link between source REDS and destination REDS is changed i.e. at each time instant of the *Link Change Interval* can change the current operating link. This implies that source REDS, destination REDS, and RDS within a *Link Change Interval* should all use the same link at the beginning of the *Link Change Interval* period.

If the source REDS transmits frames to the destinations REDS via the direct link but does not receive an ACK/block acknowledgment (BA) from the destination REDS during the *Link Change Interval*, the source REDS should fall back to the relay link at the beginning of the next *Link Change Interval* and forward frames via the RDS to the destination REDS. In the following SPs, the source REDS uses the link in which the frame exchange towards the destination REDS was successful.

3.3.14.2. Additional Rules for FD-AF RDS

To inform the destination REDS about the link switch, the source REDS defers its transmission by *Data Sensing Time*. This gives implicit signaling to the destination REDS to switch to the relay link and steer its antenna beam towards the RDS. However, the source REDS might not have any frame at the beginning of a *Link Change Interval* to transmit. In this case, the destination inspects the value of the `More Data` field in the last frame it received from the source REDS. If the value is equal to 0, then the destination REDS shall keep the direct link.

3.3.14.3. Additional Rules for HD-DF RDS

When the current link is the relay link, the frame exchange is performed in two periods. The SP is divided into alternating periods named *First Period* and *Second Period*. In the *First Period*, the source REDS transmits a frame to the RDS and then the RDS acknowledges the reception after a short interframe space (SIFS). In the *Second Period*, the RDS forwards the received frames from the source REDS to the destination REDS and then the destination REDS acknowledges the reception after SIFS. If the source REDS decides to change to the relay link, it should suspend its frame transmission to the destination REDS and starts its frame transmission towards the RDS in the following *Link Change Interval* which will be the start of the *First Period*. The destinations REDS stops receiving frames in the following of *Link Change Intervals*. This gives an implicit signaling to the destination REDS that link switching happened.

3.3.14.4. Relay signaling procedure

Figure 3.13 summarizes different procedures to establish a relay link with a destination Relay Endpoint DMG STA (REDS) in a DMG BSS. A STA should acquire the DMG capabilities of the DMG STA it wishes to establish a relay operation with before it initiates any relay setup operation. This is done by sending an Information Request frame to the DMG AP after the DMG STA completes its association with the DMG AP.

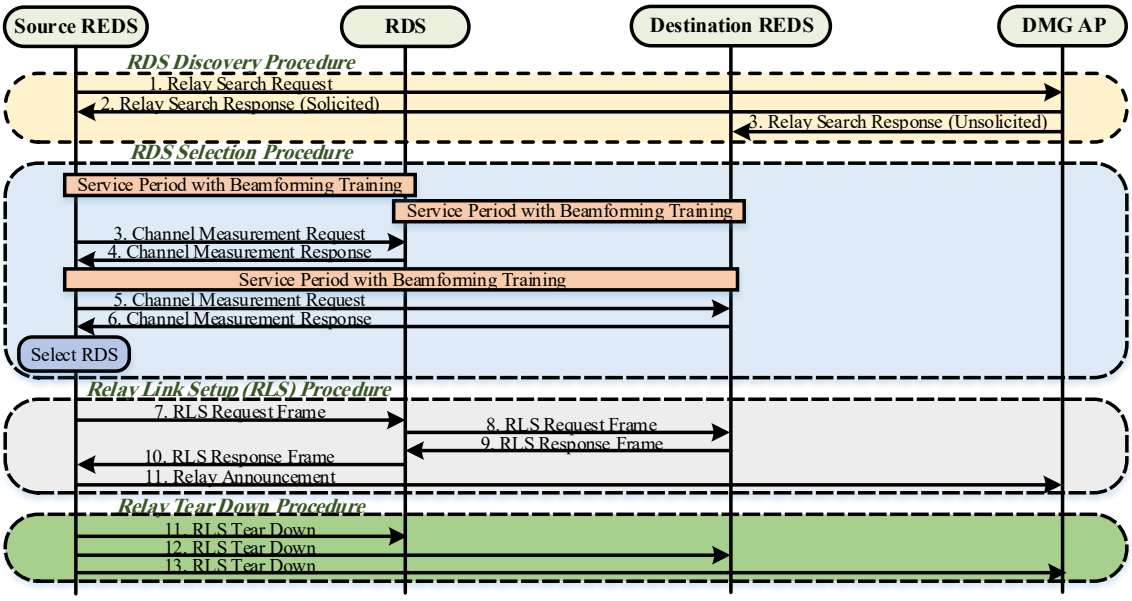

Figure 3.13: Relay Signaling Procedure.

- **RDS Discovery Procedure**: In this phase, a source REDS searches for candidates Relay DMG STAs (RDSs) in the DMG BSS. The DMG AP informs both source REDS and destination REDS about the available REDS in the network with their DMG capabilities.

- **RDS Selection Procedure**: At this point, the DMG AP schedules several SPs for BF training between all the available RDSs together with source REDS and destination REDS consecutively. After that, the source REDS request for channel measurements with the candidates RDSs. Later on, the DMG AP schedules a SP for BF between source REDS and destination REDS. After finishing the BF, the source REDS requests destination REDS to send channel measurements with the available RDSs. As a result, the source REDS will be aware of all channel states in the network. Based on this information, the source REDS selects the best RDS for relaying. In our implementation, we select the RDS which receives frames from both source REDS and destination REDS with the highest SNR.

- **RLS Procedure**: In this phase, the source REDS decides to forward its current transmission through the selected REDS in the previous phase. Thus, it sends an RLS Request frame to the selected RDS. The selected RDS in return forwards this frame to the destination REDS. At this point, the destination REDS replies to the selected RDS with an RLS Response with a status equal to Success if it accepts to communicate through the relay link. The selected REDS forwards this frame to the source REDS with a status equal to Success if it accepts to act as a relay. If both destination REDS and RDS accept to switch the link, the source REDS sends an announcement frame to the DMG AP regarding the newly established relay link in the network.

- **Relay Teardown Procedure**: If the source REDS decides to terminate its relay link through the selected RDS, it shall transmit an RLS Tear Down frame to the selected RDS, destination REDS and DMG AP.

3.3.15. Spatial Sharing Technique

The spatial sharing technique allows SP allocations for different DMG STAs in the same spatial domain with the same DMG BSS to be scheduled concurrently during a DTI access period. To support spatial sharing, a DMG PCP/AP sets the `SPSH and Interference Mitigation` field equal to 1 in the `DMG Capabilities` information element. Establishing spatial sharing requires executing the following two phases:

3.3.15.1. Assessment Phase

In this phase, the PCP/AP initiates radio resource measurement procedure with the intended STAs to assess the possibility to perform spatial sharing. Initially, the PCP/AP sends a `Directional Channel Quality Request` as part of the `Radio Measurement Request` action frame to the intended STAs. These STAs should have done beamforming training before starting any radio measurements. In the context of spatial sharing, we use the term `candidate SP` to refer to an SP to be assessed for spatial sharing with existing SPs. A candidate SP could be either a new SP to be scheduled in the next BI or a scheduled SP with

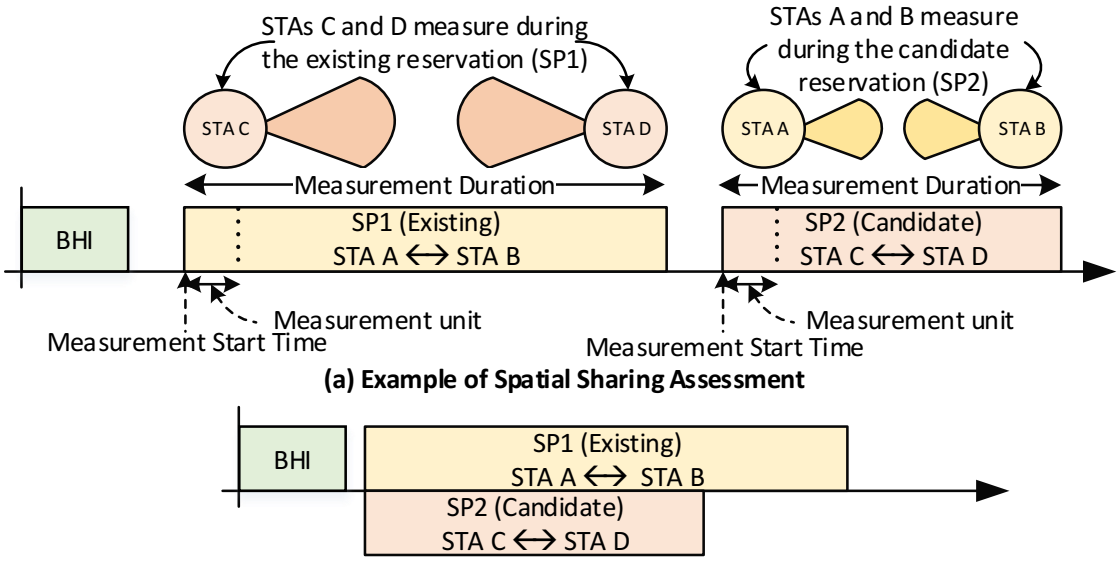

Figure 3.14: Spatial Sharing and Interference Assessment [1].

an allocated channel time in the DTI access period. The PCP/AP sets the `Target STA` field in the `Directional Channel Quality Request` to the peer STA's MAC address involved in the candidate SP and the `Measurement Method` field to indicate the Average Noise plus Interference Power Indicator (ANIPI). If the candidate SP is already allocated, the PCP/AP additionally transmits a `Directional Channel Quality Request` to the STAs involved in the existing SP to assess the possibility for spatial sharing with the candidate SP. In this request, the PCP/AP sets the `Target STA` to the corresponding peer STA involved in the existing SP and the `Measurement Method` field to indicate ANIPI. The recipient of the `Directional Channel Quality Request` carries out channel measurements using same receive antenna configuration as is used when receiving from the target STA. When the recipient STA completes the requested measurements, it reports back to the PCP/AP the result of these measurements using the `Directional Channel Quality Report` as part of the `Radio Measurement Response` action frame.

Figure 3.14 shows an example of spatial sharing assessment between two SPs where SP1 is the existing SP and SP2 is the candidate SP. At the beginning, the PCP/AP transmits a `Directional Channel Quality Request` to STA C and STA D to perform channel measurement over SP1's channel allocation. Then, it transmits a `Directional Channel Quality Request` to STA A and STA B to measure over SP2's channel allocation.

3.3.15.2. Execution Phase

Based on the channel measurements results in the previous phase, the PCP/AP estimates the channel quality across STAs and decides whether to implement spatial sharing. The PCP/AP

overlaps the candidate SP with the existing SP in its BI if the performance is expected to maximize. The determination of the performance improvement is based on the measurements reports received by the PCP/AP and it is implementation dependent. Figure 3.14 (b) depicts the allocation of SPs in the BI due to spatial sharing. To ensure those SPs involved in the time-overlapped schedule do not exhibit additional interference to each other due to sudden changes in the propagation environments, the PCP/AP periodically transmits a `Directional Channel Quality Request` to each STA involved in those SPs. The PCP/AP sets the `Target STA` to the peer STA involved in the same SP and set the `Measurement Method` field to indicate Received Signal to Noise Indicator (RSNI). Then, each STA performs channel measurements as indicated in the request and sends back `Directional Channel Quality Report` to the PCP/AP. Based on the report, the PCP/AP decides whether to sustain spatial sharing between SPs based on the quality of the link links involved in the spatial sharing. However, the decision is beyond the scope of the standard.

We provide the user with a new trace source *ChannelQualityReportReceived* in the `DmgApWifiMac` class. This trace source is triggered every time the PCP/AP receives a `Directional Channel Quality Report` and it provides two parameters: the MAC address of the reporting STA and the measurement report. The measurement report is calculated based on the measurement method specified in the `Directional Channel Quality Request`. Based on these reports, a user can decide whether to re-allocate existing SPs allocations in the BI and achieve spatial sharing or to stop spatial sharing between these SPs if the performance is degraded.

3.3.16. DMG PCP/AP Clustering

IEEE 802.11ad defines two types of clustering: decentralized clustering and centralized clustering. The former one requires no centralized controller between the distributed PCPs/APs, whereas the later requires the existence of a single centralized coordination service set (CCSS) entity. In this work, we focus on the decentralized clustering implementation. Forming a decentralized cluster among a group of spatially distributed PCPs/APs operating on the same channel requires one of these PCPs/APs to act as Synchronization PCP (S-PCP)/S-AP. Each cluster has a unique ID which corresponds to the MAC address of the S-PCP/S-AP. The BI in a cluster is divided into `ClusterMaxMem` Beacon SPs. A member can transmit exclusively during one of these SPs while other members stays in receive mode. The duration of each SP equals the length of the beacon interval of the S-PCP/S-AP divided by `ClusterMaxMem`. The first Beacon SP is reserved for S-PCP/S-AP. Establishing a decentralized cluster requires the following steps:

3.3.16.1. Cluster Formation

Forming a decentralized cluster requires at least one S-PCP/S-AP. A PCP/AP becomes PCP/S-AP by transmitting DMG Beacon that includes `Clustering Control` field with the

`Cluster Member Role` subfield set to the value for an S-PCP/S-AP. Each PCP/AP on the channel that receives a DMG Beacon from an S-PCP/S-AP starts monitoring the channel for DMG Beacon transmission during each Beacon SP for a duration of at least equals to $1.024s$. A Beacon SP is considered empty if no DMG Beacon is received during the monitoring interval. Figure 3.15 illustrates an example of decentralized clustering with three PCPs/APs.

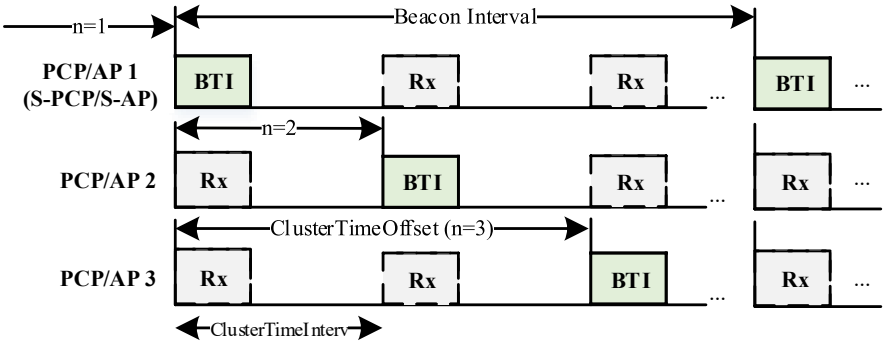

Figure 3.15: Decentralized Clustering for 3 PCPs/APs [1].

3.3.16.2. Cluster Maintenance

The S-PCP/S-AP is responsible for providing synchronization among cluster members. This facilities sharing medium resources fairly among all the members. For this reason, it is important to provide a procedure for cluster maintenance in case of S-PCP/S-AP failure. This is done through S-PCP/S-AP handover. The handover procedure compromises two *Cluster Monitoring Periods*. The first period ends when a member of the decentralized cluster stops receiving DMG Beacons from the S-PCP/S-AP within 8 beacon intervals. In the second period, the PCP/AP monitors the channel for DMG Beacons transmitted by other members and at the same time it keeps transmitting its own DMG Beacons in the selected Beacon SP. Hence data communication in the existing allocations is not affected. At the end of this period, if the member PCP/AP receives DMG Beacon with the `Cluster Member Role` equals to the S-PCP/S-AP, the member PCP/AP forms a cluster as described in Section 3.3.16.1. Otherwise, each member compares its own MAC address against all the MAC addresses of the received DMG Beacons. If a PCP/AP has the lowest MAC address, then it shall become S-PCP/S-AP. However, if the MAC address is not the lowest then the PCP/AP starts a new *Cluster Monitoring Period*.

3.3.17. DMG Beamformed Link Maintenance

To maintain a beamformed link between STAs, these STAs should negotiate the value of the `Link Maintenance Timer` during beamforming training. If one of the STAs sets this value to 0, then beamformed link maintenance is not supported. Figure 3.16 shows the operation of the beam link maintenance timer between STA-A and STA-B. In BI (n), STA-A and STA-B complete beamforming training and establish link maintenance timer. In the following BI (n+1), the STAs

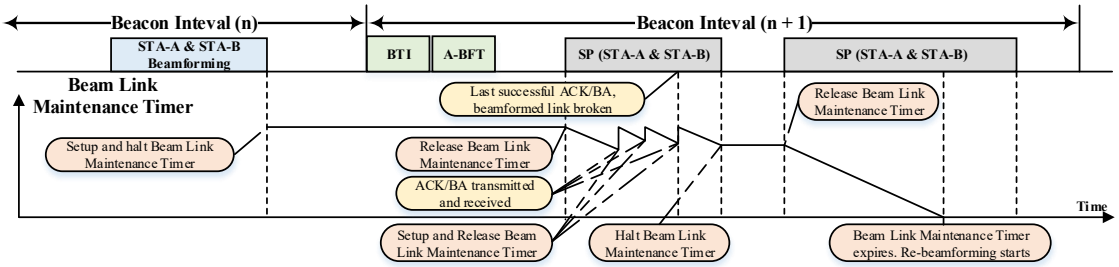

Figure 3.16: Example of Beamformed Link Maintenance [1].

have two SPs for direct communication. The beam link maintenance timer is set-up, halted, and released according to the rules described in Section 2.2. During the last frame exchange of the first SP, the beamformed link gets blocked, so the frame exchange fails, and as a result, the timer continues counting down. In the following SP, the link is still interrupted so that STA-A and STA-B cannot reliably exchange frames, and thus the timer keeps counting down until it expires. Upon the expiration of the timer, the STAs redo re-beamforming.

3.4. SU/MU-MIMO Communication for IEEE 802.11ay in ns-3

The main goals of the next generation wireless and mobile networks are ubiquitous connectivity, a high number of connected devices, ultra high-speed links, and extremely low latency. Communication in the mmWave band fulfills all the requirements as mentioned earlier and paves the way for a new set of applications that are not possible with the existing wireless technologies such as 4G/LTE and 802.11n/ac standards. The IEEE 802.11ad was the first WLAN standard that utilizes the 60 GHz band and provides multi-gigabit throughput over the air using both SC and OFDM PHY technologies. The use cases for IEEE 802.11ad range from ultra high definition video streaming and HDMI cable replacement to fast file synchronization. Despite the very high throughput provided by IEEE 802.11ad, it is still insufficient for some applications such as wireless back-hauling and front-hauling solutions. IEEE 802.11ay is the next generation mmWave standard that employs a variety of techniques to dramatically increase PHY capacity and throughput from 7 Gbps to up to 300 Gbps. This is mainly achieved by introducing complex PHY techniques, including MIMO communication, channel bonding, and aggregation [6].

The ns-3 WiFi module does not provide sufficient support for MIMO communication. The only support of MIMO in ns-3 is a basic optimistic implementation of a single user (SU)-MIMO, where the throughput is increased by a fixed ratio depending on the SU-MIMO configuration. This approach was used for simplicity and scalability but has rather low fidelity. Accurately representing SU/multi-user (MU)-MIMO communication requires intensive signal processing operations first to estimate and extract the coefficients of the wireless channel and then, to perform digital precoding to eliminate or reduce inter-stream interference. Simulating the complete signal processing chain would make network-level simulations last longer and more complex to analyze.

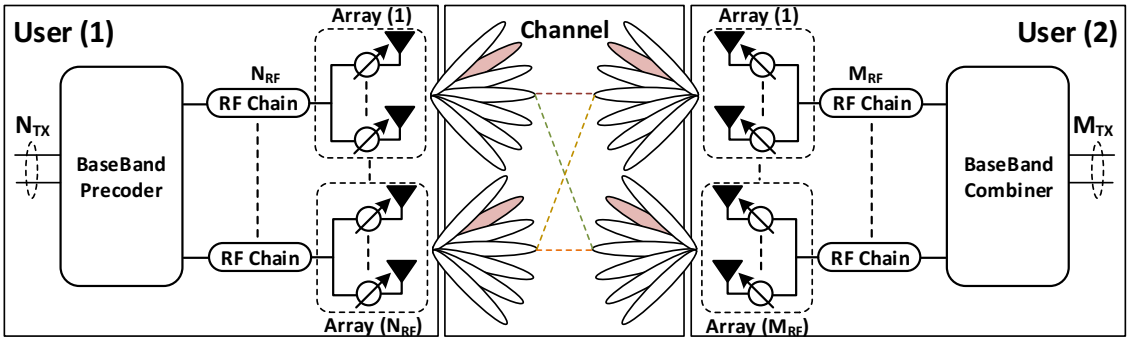

Figure 3.17: SU-MIMO Partially Connected Beamforming Architecture.

For these reasons, network-level simulators rely on PHY layer abstraction techniques, which must be carefully chosen to guarantee results close to real-PHY behavior.

In this section, we survey different options for implementing SU/MU-MIMO communication for the IEEE 802.11ay standard in ns-3, and we compare their advantages and disadvantages. Finally, we provide a proposal for the design of a system that incorporates minimum signal processing steps to simulate SU/MU-MIMO communication with high fidelity and low computational complexity.

3.4.1. Millimeter Wave MIMO Systems

Wireless communication in the mmWave band requires the use of steerable Phased Antenna Arrays (PAAs) to generate analog beams and focus the energy towards the intended receiver. Achieving full digital beamforming requires a dedicated RF chain per each antenna element in the PAA. This approach is not practical due to both high cost and high power consumption of the Digital-to-Analog Converters (DACs) and Analog-to-Digital Converters (ADCs) operating for high bandwidth signals. For these reasons, different hybrid beamforming architectures that combine both analog and digital beamforming have been proposed in the literature [25]. In these approaches, a single RF chain is connected to a subset of the antenna elements. Figure 3.17 illustrates the partially connected beamforming architecture. In this architecture, each RF chain is connected to a single PAA or a sub-array of a larger PAA. Also multiple PAAs can be connected to a single RF chain, but then only one of those PAAs is active at a time. The baseband precoder maps N_{TX} spatial streams to N_{RF} chains. Hybrid approaches are practical, consume less power, and achieve similar performance to fully digital beamforming architectures [26].

However, the joint design of analog and digital beamforming is a rather challenging task. For this reason, mmWave wireless systems tend to perform beamforming training in two-stages. In the first stage, the two communicating devices perform analog beamforming training using coarse analog beams of the PAA. However, these coarse analog beams introduce inter-stream interference among the spatial streams at the receiver. As a result, in the second stage, the two communicating devices perform digital precoding or digital beamforming in which the receiver

estimates the equivalent baseband channel and feeds it back to the transmitter. The equivalent baseband channel contains the CSI between each pair of transmit and receive PAA. Then, the transmitter builds a spatial mapping matrix Q to eliminate inter-stream interference. The IEEE 802.11ay standard defines a set of procedures to perform beamforming training for both SU-MIMO and MU-MIMO communications. These procedures allow devices to determine the best transmit and receive antenna configuration for simultaneous transmission and receptions of multiple spatial streams. In the following sections, we provide a brief background on SU/MU-MIMO communication in the IEEE 802.11ay standard and how to simulate it in ns-3.

3.4.2. SU-MIMO Communication

3.4.2.1. Background

In this communication mode, a single data stream is scrambled, encoded, and then mapped to multiple spatial streams N_{SS}. For each spatial stream, the transmitter can use an independent MCS. If space-time block coding (STBC) is enabled, then each single spatial stream will be further mapped to two space-time streams. In that case, a transmitted stream over the air does not correspond to a single MAC frame but rather to the symbols of a fraction of that frame. Then, for each spatial stream, we generate an independent preamble and header fields. Finally, the space-time streams N_{STS} are mapped to multiple transmit chains N_{TX} based on a spatial mapping matrix Q.

To simulate SU-MIMO in ns-3, we need to implement the previous signal processing chain at the transmitter and its corresponding part at the receiver, which increases computational complexity and thus simulation runtime. The authors in [27] integrated the OFDM PHY layer of the IEEE 802.11a/p standards in ns-3 and found that simulation runtime increased by a factor of 300 to 14000. The increase depends on many factors including the amount of generated traffic, packet sizes, and the number of simulated stations. We expect this to be even worse when simulating the IEEE 802.11ad/ay PHY layer. Indeed, this standard can generate gigabits of throughput per second which results in tens of thousands of packets per second. On the other hand, simulating the whole signal processing chain improves simulation accuracy.

The current implementation of SU-MIMO for the IEEE 802.11n/ac standards in ns-3 abstracts all the signal processing steps to generate and transmit multiple parallel spatial streams N_{SS}. Specifically, transmitting multiple N_{SS} affects only the number of OFDM symbols in the payload part of the transmitted frame. This approach requires generating an offline LUT for mapping SNR to PER. The SNR in this case is an effective SNR value that incorporates all the spatial streams SNRs. However, this approach requires generating a large set of LUTs based on the number of transmit and receive PAAs, and the type of channel between each pair of transmit and receive antennas. For this reason, such an approach does not scale well for dense scenarios.

The scalability of the LUTs method is even worse when simulating SU-MIMO in the mmWave band. This is because the mmWave channel is very sparse and few multipath compo-

nents (MPCs) exist between pairs of transmit and receive PAAs. As a result, transmitting multiple spatial streams is not always feasible as it depends on channel diversity to multiplex these spatial streams. Additionally, since PAAs within the same device are located close to each other, analog beamforming is not sufficient to separate the concurrent streams and avoid inter-stream interference at the receiver side. For these reasons, we need to generate an extensive set of LUTs to consider both channel condition for each channel instance and the digital precoding technique utilized to orthogonalize the spatial streams at the receiver.

3.4.2.2. Implementation Proposal

Figure 3.18 shows the building blocks for our proposal to simulate SU/MU-MIMO for the IEEE 802.11ay standard in ns-3. We explain in detail each block:

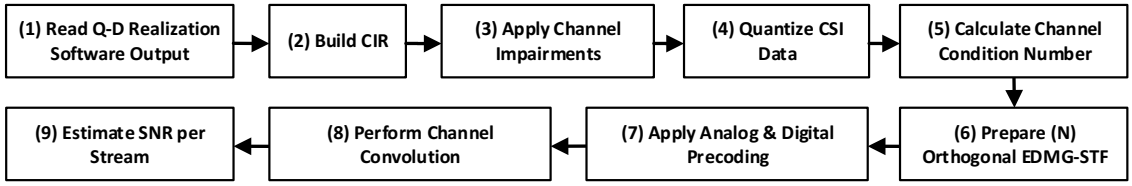

Figure 3.18: MIMO Simulation Blocks for IEEE 802.11ay in ns-3.

1. We generate Q-D channel traces using our Q-D realization software [16].

2. The generated trace files include all the MPCs in the simulated environment, allowing to construct the CIR between each pair of transmit and receive PAAs. As a result, we get the full channel matrix $H_{BB}(t, \tau)$. This step eliminates the channel estimation on a per-packet basis.

3. Then, we apply the effects of channel impairments on the CIR to emulate the channel estimation error. These effects include carrier frequency offset (CFO), Phase Noise (PN), non-linearity in the amplifiers, and Symbol Timing Offset (STO). Executing this step is optional.

4. Here, we quantize the calculated CSI to emulate quantization errors in real systems where devices use a limited number of bits.

5. Based on the quantized CSI information, we compute the channel condition number to estimate the number of N orthogonal spatial streams the fading channel can support. The channel condition number depends on the Eigenvalues of the channel matrix and is not impacted by analog and digital precoding.

6. At this point, we start the data transmission phase by constructing N orthogonal enhanced directional multi gigabit (EDMG)-STF fields as defined in the IEEE 802.11ay standard.

7. We perform both analog and digital precoding to these EDMG-STFs. The analog precoding works as a spatial filter to the CIR whereas the digital precoding reduces inter-stream interference. We perform digital precoding using one of the existing techniques, e.g., Zero-forcing (ZF) or minimum-mean-squared-error (MMSE).

8. Then, we perform channel convolution on the previous EDMG-STF fields.

9. We estimate the signal-to-interference-plus-noise ratio (SINR) for each stream from the convoluted EDMG-STF fields. The EDMG-STF field contains long repetitions of identical Golay sequences, which allows estimating the noise of the signal. Additionally, the number of transmitted spatial streams is known, and these streams are orthogonal due to the properties of the utilized Golay sequences, and we can thus estimate the power of all the streams at all the PAA. Finally, we map the SINR value for each stream to PER value. If the receiver has multiple PAAs, we map the SINRs values of all its streams to an effective PER value.

While our proposal introduces extra processing overhead and increases simulation runtime, it avoids generating a high number of LUTs to consider all the possible configurations for each simulated scenario. However, it is still necessary to create a few LUTs that map the SNRs of multiple streams at each receiver to an effective PER value. The heavy computation and processing overhead come from the operations related to constructing the full channel matrix $H_{BB}(t, \tau)$ between all the transmit and receive PAAs. This step has to be executed every time a single channel instance between a single pair of transmit and receive PAAs changes, for example, due to mobility. In summary, our implementation proposal is a compromise between performing symbol-level simulations and entirely abstracting PHY layer operation.

3.4.3. MU-MIMO Communication

In MU-MIMO communication, an AP transmits multiple spatial streams to different users at the same time. The number of users can be more than the number of transmit chains at the AP. For these reasons, researchers have been proposing different user scheduling algorithms for MU-MIMO systems. These algorithms select which group of users should participate in the current MU-MIMO transmission to increase overall system capacity [28].

At the time of writing, the WiFi module in ns-3 lacks support for MU-MIMO. To simulate IEEE 802.11ay MU-MIMO communication in ns-3, we propose two implementations based on network complexity. In the first implementation, we assume that each user in the network has a single RF chain connected to a single PAA to transmit and receive a spatial stream. Besides, the AP has multiple RF chains but it transmits a single spatial stream per user. As a result, each spatial stream corresponds to a data packet from the MAC layer for a specific user. Additionally, we consider that analog beamforming can generate orthogonal beams to avoid inter-stream interference at each user and thus we can omit the digital beamforming part (open-loop solution). This

implementation gives an upper-bound on the system's performance and capacity.

For the second implementation, the AP can transmit multiple spatial streams to each user, and each user can use two or more RF chains to transmit and receive. Here, we reuse the same implementation proposal for simulating SU-MIMO as we need to eliminate inter-stream interference among streams belonging to the same user. In addition, we need to reduce interference between users as the analog beams are not orthogonal. This implementation requires simulating the effect of digital beamforming. The overhead of this implementation increases with the number of users participating in the MU-MIMO communication.

3.5. Conclusions and Future Work

We provide an architecture for modeling WLAN IEEE 802.11ad with its various techniques in ns-3. We describe the implementation of new features that improve the reliability and fidelity of the ns3 IEEE 802.11ad model. Our work is a great step towards reducing the discrepancies between the results coming from practical testbeds and the results obtained through system-level simulations. Besides, we developed a set of powerful software tools that facilitate the usage of the newly implemented features and provide the user with unprecedented capabilities to mimic real networks.

Our next implementation step is to extend the ns-3 IEEE 802.11ad model to support the latest draft of the IEEE 802.11ay amendment. We plan to implement our proposal for simulating SU/MU-MIMO in ns-3. The next step is to evaluate the performance of our scheme against a link-level simulator that performs symbol level modeling.

Chapter 4

IEEE 802.11ad Performance Evaluation in ns-3

4.1. Introduction

Networking aspects of the IEEE 802.11ad protocol have been studied in great detail using COTS devices. However, these devices provide limited information about the operations and the performance of the MAC and PHY layers. In addition, analyzing the interactions between those layers and the upper layers of the protocol stack using COTS devices is challenging due to the lack of full control over COTS devices. In practice, network operators tend to study the performance of mmWave networks using high fidelity network-level simulators before actual deployment. The ns-3 IEEE 802.11ad model in Chapter 3 provides researchers with a framework to study and analyze the performance of the IEEE 802.11ad protocol in complex settings and with high fidelity. Moreover, using a simulation approach allows users to analyze the behavior of each component of the system, from the lowest to the highest layer, and to understand how the different components interact with each other. The material in this chapter is taken from [10, 13, 14, 29].

4.2. IEEE 802.11ad Techniques Evaluation and Validation

In this section we provide some evaluation results for our IEEE 802.11ad model. In all our simulations, we use the `DmgWifiPhy` class and its corresponding `DmgWifiChannel` class to evaluate the performance of the implemented IEEE 802.11ad techniques, unless otherwise stated. Additionally, we assume all DMG STAs and PCPs/APs have a single phased antenna array covering all directions using `CodeBookAnalytical` class. This codebook generates 8 predefined virtual sectors where each sector covers 45°. We use the parameters listed in Table 4.1 in our simulations:

Table 4.1: Simulations parameters using `DmgWifiPhy` class

Parameter Name	Parameter Value
Application Traffic Pattern	Continuous Data Stream
Payload Size	1000 Bytes
Transport Protocol	UDP
MAC Queue Size	1000 Packets
Rx Noise Figure	10 dB
Propagation Loss Model	Friis loss model

4.2.1. Evaluating Achievable Throughput

In this experiment, we demonstrate the obtained throughput for different MCSs for both SC and OFDM PHY layers. The setup compromises two nodes: one DMG AP and one DMG STA. These nodes are spaced 2 m apart from each other. The DMG STA generates a flow of User Datagram Protocol (UDP) messages towards the DMG AP. The announced A-BFT by the DMG AP consists of 8 Sector Sweep (SS) slots where each slot contains 8 SSW frames.

Figure 4.1 depicts the obtained throughput for two different sets of MCSs. The highest throughput achieved for SC is almost 4 Gbps, and for OFDM is 5.2 Gbps. However, the achieved throughput for OFDM is 1.5 Gbps less than the theoretical maximum IEEE 802.11ad throughput of 6.72 Gbps. This is mainly due to the overhead imposed by the CBAP access mechanism. Besides that, the data rate reported in the standard assumes a continues stream of OFDM symbols without any PHY and MAC overhead and any Interframe Space (IFS).

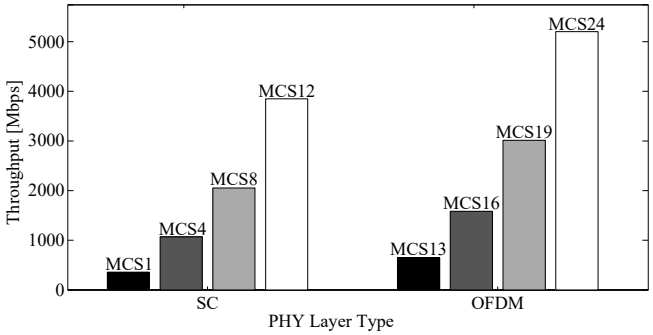

Figure 4.1: Throughput for different MCSs

4.2.2. Comparing Channel Access Schemes

In this experiment, we compare between two channel access schemes the CBAP and the Service Period Channel Access (SPCA) in terms of achievable application goodput. The scenario consists of a single PCP/AP and one DMG STA and they are separated by 1m. An `OnOffApplication` is installed on the DMG STA and it generates traffic towards the

`PacketSinkApplication` running on the PCP/AP. Figure 4.2 shows the obtained goodput with respect to the frame aggregation level. We plot the results using both MCS12 with a SC PHY layer and MCS-24 with an OFDM PHY layer. From the figure, we can notice the difference in the achievable goodput even for the case of a single packet transmission. The goodput for both access schemes increases with the aggregation level until we saturate the channel. The high efficiency of the SPCA comes from the fact that the a DMG STA using SPCA can transmit immediately without having to contend for access to the wireless medium.

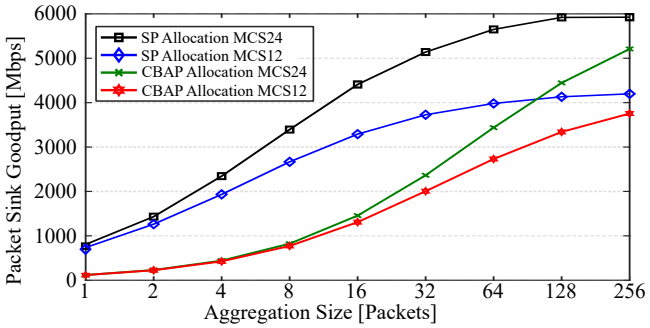

Figure 4.2: Comparing channel access schemes.

4.2.3. Evaluating Fast Session Transfer

In this experiment, we demonstrate the capability of transferring an on-going data session smoothly from the 60 GHz band to the 2.4 GHz band. The simulation setup is similar to the one in Section 4.2.1 with the addition that the nodes can communicate in the 2.4 GHz band using IEEE 802.11n. We set the value of LLT to 1000 which corresponds to a link loss countdown value of 32 ms. After the nodes establish the directional link, they setup a unique FSTS between each other.

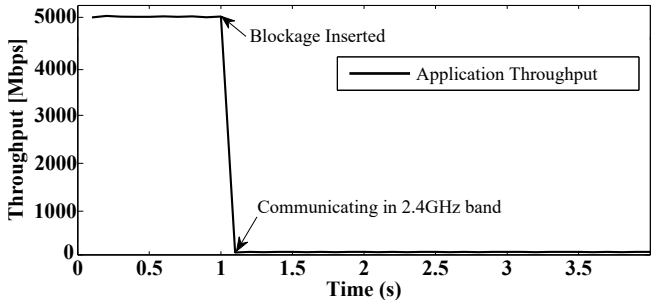

Figure 4.3: FST setup results

A the beginning of the simulation, the nodes communicate with each other normally, and the achieved throughput is around 5 Gbps as shown in Figure 4.3. After one second, we introduce a blockage in the link of -45 dBm. This blockage breaks the link, so the nodes start a link loss

countdown. When the timer expires, the two nodes switch to the 2.4 GHz band and continue their session. We notice the degradation in the achieved throughput around 60 Mbps due to the limited capacity available in the lower frequency band.

4.2.4. Evaluating Half-Duplex Relay Operation

Here, we study the impact of using HD-DF relay operation on the application goodput. Theoretically using a relay node operating in half duplex mode should reduce the throughput of a wireless link by half. Figure 4.4 depicts the simulation setup which consists of one PCP/AP with 3 DMG STAs. Two DMG STAs act as REDS and one DMG STA supports RDS. The source REDS runs `OnOffApplication` and generates data flow towards the `PacketSinkApplication` installed on the destination REDS.

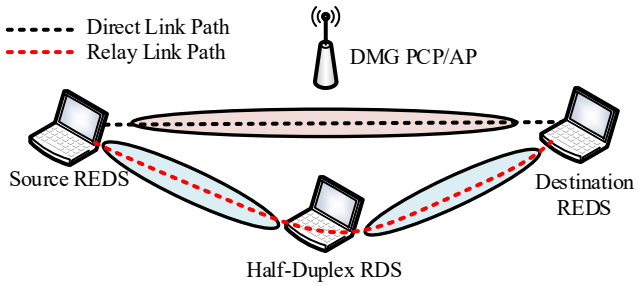

Figure 4.4: Relay network topology

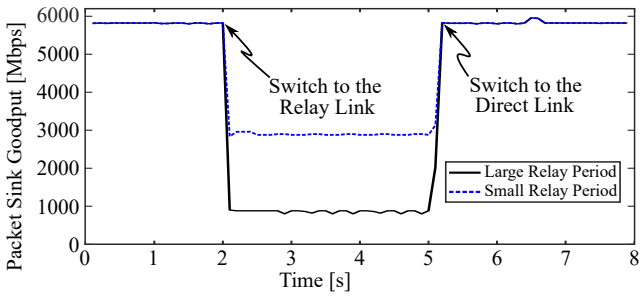

Figure 4.5: Relay setup results

At the beginning, the source REDS executes the RLS procedure to establish a half-duplex relay link towards the destination REDS through the RDS to protect any prospective SP allocations. Later on, the source REDS requests the PCP/AP to allocate a SP for data communication with the destination REDS. The PCP/AP accepts the allocation request and announces the schedule of this SP allocation in the following BIs. The RDS receives this schedule and switches into relay mode at the start time of this SP allocation for possible link switching. During the course of data communication, we explicitly signal all the STA to switch to the relay link.

Figure 4.5 illustrates the achievable application goodput before and after link switching for two different values of *Relay Period*, where $RelayPeriod = FirstPeriod + SecondPeriod$. The $FirstPeriod$ indicates the allocated time for a unidirectional data transmission from the source REDS to the RDS and the $SecondPeriod$ refers to the amount allocated time for a unidirectional transmission the RDS to the destination REDS.

The long relay period corresponds to $8ms$, whereas the small relay period corresponds to $2ms$. We can notice the impact of the time the RDS remains communicating with each REDS. The shorter the period is the lower the impact of switching to the relay link on the application goodput. This is related to the size of the MAC queue. For the large relay period case, the queue keeps buffering frames for a long time, and at some point, it gets full and starts dropping incoming frames. In contrast, for the small relay period case, the queue of the relay node does not reach its limit and can drain frames much faster.

4.2.5. Evaluating Spatial Sharing Technique

In this simulation, we demonstrate the importance of spatial sharing in the context of mmWave wireless networks. Figure 4.6 shows our simulation setup which compromises 4 STAs and a single PCP/AP. Once all STAs have successfully associated with the PCP/AP, the PCP/AP allocates two SPs for beamforming training between STA A and STA B and between STA C and STA D. Upon completion of the beamforming training, the PCP/AP schedules two recurring SPs allocations for data communication as in Figure 3.14. Table 4.2 summarizes allocation parameters for each SP.

Table 4.2: Service periods allocations parameters

SP	Start Time[ms]	Length[ms]	Flow Direction
SP1	0	20	STA A $\rightarrow$ STA B
SP2	20	12	STA C $\rightarrow$ STA D

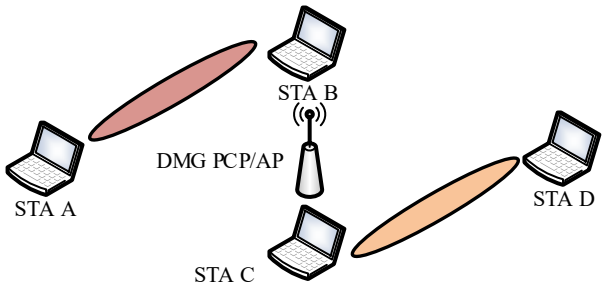

Figure 4.6: Spatial sharing evaluation topology.

During the course of data communication, the PCP/AP decides to evaluate the two SPs for spatial sharing, so it follows the procedure described in Section 3.3.15. After receiving *Channel Quality Reports* from all the STAs, the PCP/AP evaluates ANIPI values in those reports and

decides that it is safe to proceed with the spatial sharing. As a result, the PCP/AP reallocates SP2 to start at the same time of SP1. In addition, the PCP/AP exploits the extra time resource gained in the DTI access period due to the spatial sharing and decides to extend the lengths of the two SP to be 32 ms. As a result of this operation, the total network goodput increases. Figure 4.7 shows the total network goodput before and after performing spatial sharing.

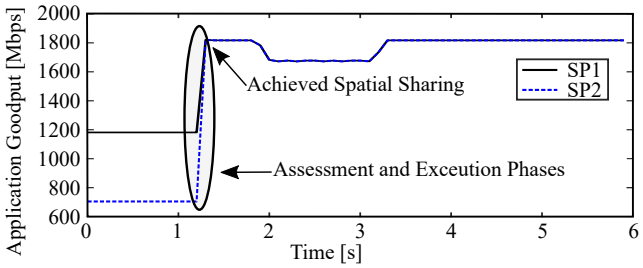

Figure 4.7: Spatial sharing results

4.2.6. Beamforming Training in SLS vs BRP

In this simulation, we compare the duration of the beamforming training when solely using the SLS mechanism with respect to the case when utilizing the BRP protocol. Figure 4.8 depicts the total beamforming duration for both SLS and BRP mechanisms for variable number of sectors/custom AWVs. All the frames are transmitted using MCS-0.

The BRP protocol is simulated by transmitting a BRP packet with TRN-TX Subfields appended to it. The BRP protocol assesses the custom AWVs in a group of 4 elements. Probing 64 SSW frames using SLS takes around 1 ms compared to 67 μs when using the BRP protocol.

To validate that our implementation changes the power during the course of a frame transmission when TRN Subfields are appended to that frame, we extended the current ASCII tracing system in the 802.11ad module. The new ASCII tracing system tracks PHY layer activities for all the DMG STAs operating on the same channel. The tracing system outputs the following set of information for each activity on the channel:

- The starting time of the activity in ns.
- The type of channel activity (transmission or reception).
- The ID of the transmitting node.
- The ID of the receiving node.
- The duration of the activity in ns.
- The type of the Physical Layer Convergence Protocol (PLCP) part that we are transmitting/receiving either Preamble, Header, Data, AGC-SF, TRN-CE, or TRN-SF.

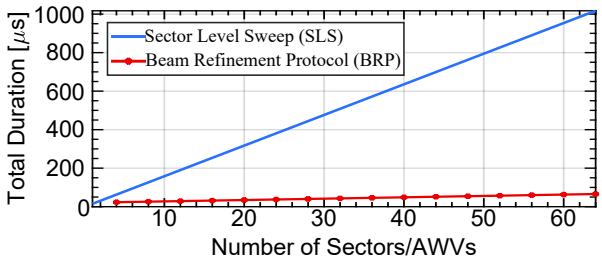

Figure 4.8: Beamforming Duration for BRP and SLS

- The Equivalent Isotropically Radiated Power (EIRP) for the corresponding PLCP part for the transmission activity in dBi or the total received power for the reception activity (EIRP + Receive Antenna Gain) in dBm.

The previous information is stored in a text file using the CSV format. Using this tracing capability, we run a simulation where a DMG STA transmits an 802.11ad frame to a DMG AP with the training field length set to 1 which corresponds to 4 TRN Subfields. Both the DMG STA and the DMG AP utilize the analytical codebook for synthesizing their beam patterns. The EIRP is evaluated at the receiver. Figure 4.9(a) shows EIRP variations over the transmitted frame. We can see that based on the AWV used, the EIRP is varying over different parts of the transmitted frame. The *Frame* part in the legend corresponds to the preamble, the PHY header, and the data part as they are transmitted using a specific sector. Besides, the TRN-CE is transmitted using the previous specific sector. The custom AWVs provides a 6 dBi gain compared to the sector, and they steer toward different directions in space. Each TRN Subfield together with its corresponding AGC Subfield are transmitted using a custom AWV as shown in Figure 4.9(b).

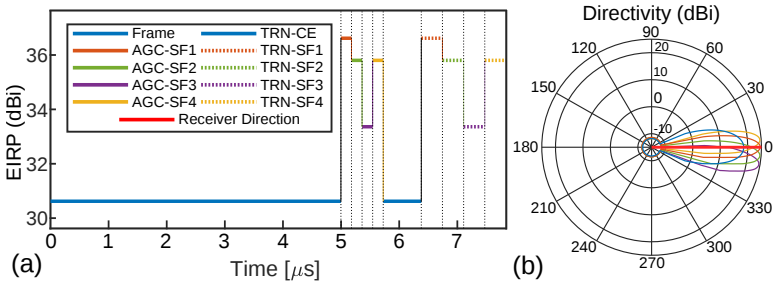

Figure 4.9: (a) 802.11ad Frame with AGC and TRN Subfields Highlighted (b) Directivity of the Sector and the Custom AWVs

4.2.7. Quasi-deterministic (Q-D)-Channel Model

The Q-D channel Model implementation allows to compute the received power by relying not only on the LoS propagation but also on the specular reflections coming from the simulation environment. To demonstrate this ability, we developed a proof-of-concept 3D visualizer: the Q-D Visualizer. The Q-D Visualizer, developed with Mayavi Python's library, allows to visualize:

- The geometry of the environment, the nodes' positions and the MPCs corresponding to the direct path and specular reflections (obtained from the Q-D channel realization software).

- The 3D antenna response pattern corresponding to each sector/AWV (provided by the *IEEE 802.11ad Codebook Generator*).

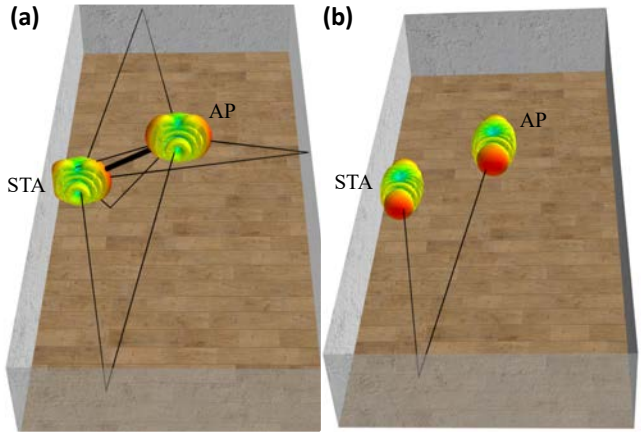

Figure 4.10: Q-D Visualizer Software. (a) AP I-TxSS using all MPCs (b) AP I-TxSS using one MPC

The visualizer uses the ns-3 output traces and can be used to visualize 802.11ad beamforming phases, i.e., SLS and BRP. Figure 4.10(a) presents a screenshot of the Q-D Visualizer. The simulated environment is a rectangular room. There are 6 MPCs between the AP and the STA (1 direct path represented by the thick black line and 5 specular reflections depicted with thin black lines). The visualizer displays the outcome of the AP and the STA TXSS SLS phase simulated in ns-3 using the Q-D channel model (the TXSS phase determines the best transmit sector that a node must use to establish a directional communication link with another node). We can see that the chosen sector 3D response patterns are aligned with the direct path to yield a maximum gain. Figure 4.10(b) presents the same scenario except that we configure the Q-D channel realization software to generate the channel only with a specular reflection (the one bouncing in the front wall of the room). We can see that the TXSS chosen antenna patterns are steered along with the specular reflection, showing that the Q-D channel in ns-3 makes use of the spatial path diversity to compute the received power.

4.3. Simulation Scenarios

In this section, we evaluate three different scenarios that demonstrate the versatility of our ns-3 IEEE 802.11ad model. All the scenarios use a 10 x 19 x 3 m room. Device parameters are summarized in Table 4.3. We configure each DMG STA to perform BF each 10 BIs during the DTI access period. The BF is limited to TXSS only. Both DMG STA and DMG AP use 2x8

elements URA PAA. This choice has been made as usually, for indoor scenarios, the steering of the PAA is done in the azimuth plane rather than in the elevation plane. All the DMG APs are mounted on the ceiling at a height of 3 m and all DMG STAs are placed at a height of 1.2 m.

Table 4.3: Simulations parameters using Q-D channel model

Parameter Name	Parameter Value
Application Type	OnOffApplication
Data Rate	3.6 Gbps
Flow Direction	Uplink
Payload Size	1448 Bytes
Transport Protocol	UDP
MAC Queue Size	1000 Packets
Aggregation Type	A-MSDU and A-MPDU
MAC Protocol	CSMA/CA
PHY Layer Type	SC MCS-12
Transmit Power	10 mW
Rx Noise Figure	10 dB
Operating Frequency	60.48 GHz
Number of Transmit Sectors	37 Sectors
A-BFT Sectors	13 Sectors
PAA Quantization Bits	8 Bits

4.3.1. Blockage Scenario

Electromagnetic-waves in the mmWave band propagate in a quasi-optical way. As a result, the received signal is dominated by the LoS path and first order reflections coming from reflecting materials. Diffraction in this band is almost negligible, and thus electromagnetic-waves in this band cannot bend over corners. This makes communication in this band sensitive to blockage and human mobility. Besides, due to the high frequency, the Free Space Path Loss (FSPL) is 20 dB higher compared to wireless technologies operating in the microwave band. For these reasons, communication in this band is more suitable for indoor scenarios, and coverage is confined to a small area. In this scenario, we study the impact of a blockage on a mmWave communication link for a single AP deployed in the corner of an L-shaped room.

Figure 4.11 depicts the user mobility stages within the L-shaped room. The thick black line represents the LoS component and the other lines represent the MPCs coming from specular reflections from side walls, floor, and ceiling. Figure 4.12 shows both the variations of the application layer throughput for MCS9 and MCS12 and the SNR changes over the course of the simulation. The vertical dashed lines represent the time the DMG STA completes its TXSS BF. We observe the effect of the Beamforming Training (BF) as the SNR value increases after its completion.

At the first stage (a), the DMG STA has a clear LoS path towards the DMG AP which results in a high SNR value and thus the DMG STA can use MCS-12 for communication. At stage (b),

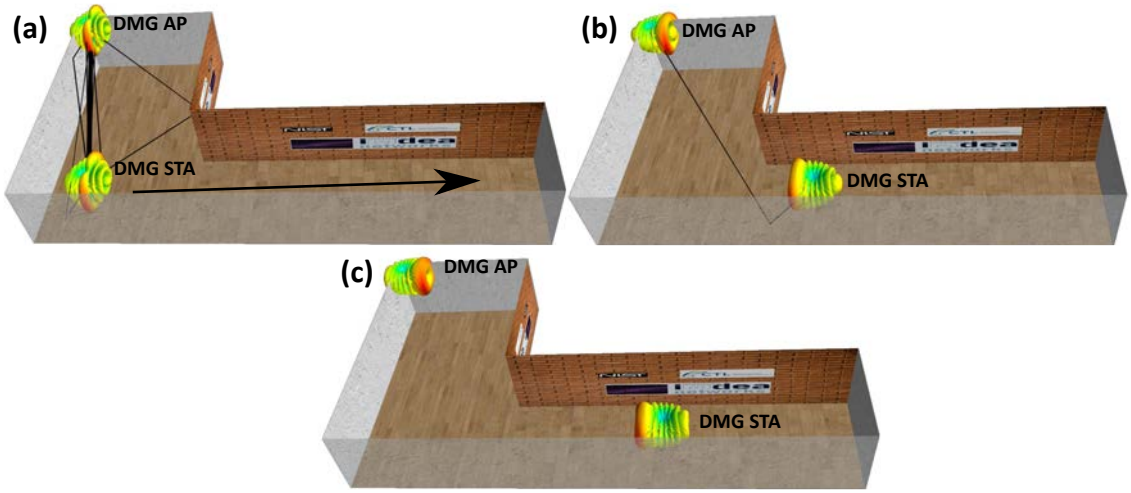

Figure 4.11: L-Shaped Room Scenario: (a) LOS Communication (b) NLOS Communication (c) No Communication.

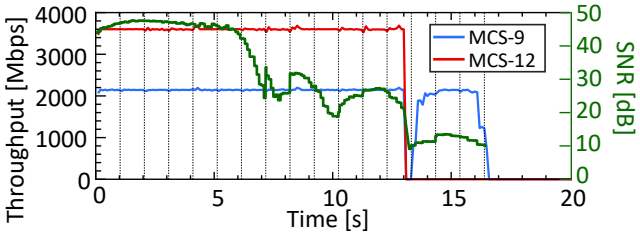

Figure 4.12: L-Shaped Room Throughput and SNR Variations.

the DMG STA has moved to the other part of the room, where the LoS component is blocked by the corner of the room, and communication with the DMG AP is NLOS through a first order reflection coming from the wall facing the AP. As a consequence, the received SNR decreases considerably, which in turn increases the PER and requires a lower MCS, and thus affects application throughput. This stage starts at time 13 s. The SNR becomes too low to sustain a communication link using MCS-12 and thus the links break down. However, the SNR value is suitable for communication using MCS-9 which is the value then selected by the rate adaptation mechanism. In the last stage (c) which starts at time 16.2, the DMG STA cannot sustain the NLOS link with the DMG AP anymore since the reflections become too weak, and the communication link breaks down for any MCS value.

Several solutions can be envisioned to overcome the previous problem. For example, a relay node could be installed at the opposite side from the DMG AP. The relay node could maintain a LoS path towards the DMG AP and at the same time, it could reach the user at the other side of the room. This is a standard compliant solution as the IEEE 802.11ad standard supports two types of relay operation modes namely FD-AF and HD-DF. Another solution would be to deploy two

APs inside the room. However, this solution would require the support of a handover mechanism between those DMG APs through a centralized controller.

4.3.2. Medium Sharing and Spatial Reuse

In this scenario, we look at how mmWave devices share the wireless medium when utilizing the CSMA/CA protocol. In addition, we show the ability of achieving full spatial reuse when deploying multi-AP. Our scenario is composed of two DMG BSSs. Each DMG BSS comprises one DMG AP and one DMG STA associated with it. Figure 4.13 shows our scenario for two different locations of DMG STA (2).

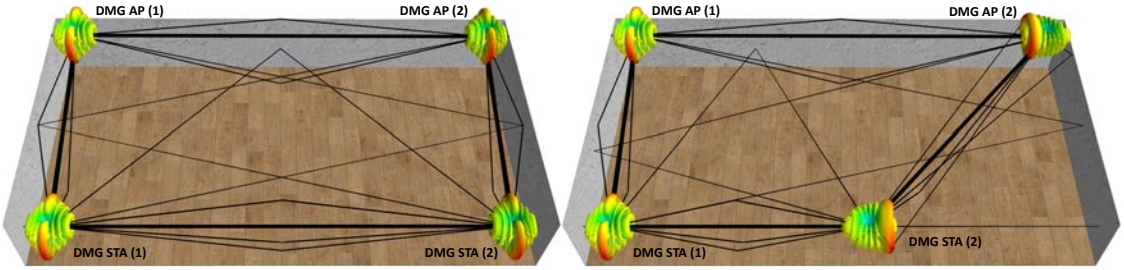

Figure 4.13: Spatial Sharing Scenario

In this scenario, a single DMG AP would be sufficient to provide full coverage for all the DMG STAs within the room. However, all the devices would share the wireless medium since they utilize the CSMA/CA protocol. To maintain multi-gigabit throughput for all wireless devices, we increase the number of DMG APs. This would be a typical case for a cubical office environment, where multiple DMG APs are deployed on the ceiling. Each DMG AP serves a single cubicle office to ensure gigabit of throughput for the user. Due to the use of directional beam patterns, each DMG AP will have its collision domain and all of the devices within the same collision domain will be sharing the wireless medium in a fair way.

Figure 4.14 shows the evolution of the application layer throughput with respect to time. At time 0 s, the DMG STAs achieve full spatial sharing as the directionality of the chosen patterns minimizes the interference between the two links. Each DMG STA is thus able to reach its maximum throughput (3.6 Gbps). Then, DMG STA (2) starts moving on a straight line towards DMG STA (1). Until time 3.8 s, the two DMG STAs are still capable of achieving full spatial sharing. After this point the two DMG STAs start overhearing each other's transmissions, since the PAAs have strong side-lobes which are capable of hearing nearby transmissions. Thus, as CSMA/CA is utilized to access the channel for data transmission, the two DMG STAs start sharing the wireless medium. At time 7.1 s, DMG STA (2) starts to move away from DMG STA (1) on the same line and stops in the middle of the room, yielding a throughput increase.

In general, utilizing a PAAs with a high number of antenna elements is not sufficient to improve spatial sharing in a network. To achieve full spatial sharing, all the devices must esti-

mate the wireless channels between each other. This implies decomposing the channel using the singular-value decomposition (SVD) technique to understand how many MPCs exist between all the nearby devices. Based on this information, those devices can optimize their beam patterns to generate directional beams aligned with the strong path towards the intended receiver, while avoiding signal energy on the paths leading to the neighboring devices.

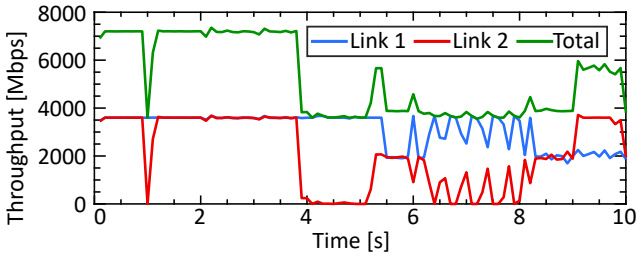

Figure 4.14: Spatial Sharing Scenario Throughput Evolution during Mobility for 2x8 PAA

4.3.3. Dense Deployment

Finally, we look into the performance of the IEEE 802.11ad protocol in a dense deployment setting. In this scenario, up-to ten clients are associated with a single DMG AP. This is a typical scenario for a meeting room within an office environment, where multiple 802.11ad capable devices in close proximity are connected to the same DMG AP. Figure 4.15(a) shows the distribution of the DMG STAs around the DMG AP. All of the DMG STAs are beamforming towards the DMG AP. Figure 4.15(b) shows the results of the TXSS BF in A-BFT for DMG STA (6).

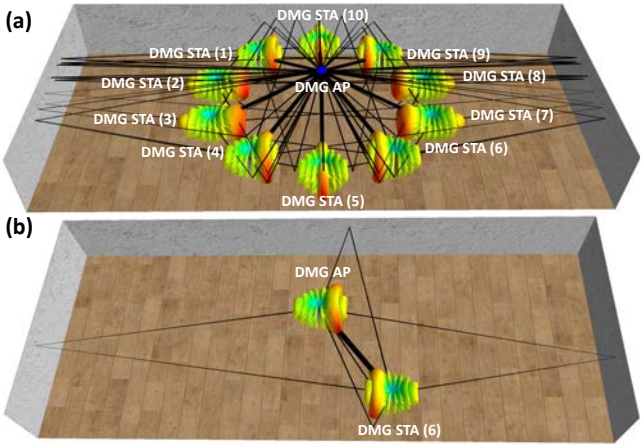

Figure 4.15: Dense Deployment Scenario (a) All the DMG STAs Beamforming Towards the DMG PCP/AP (b) SLS TxSS Phase output between DMG STA (6) and the DMG PCP/AP.

Figure 4.16 depicts throughput values for the 10 wireless links when using TCP protocol. We observe that all the links equally share the wireless medium as they have approximately the same

median for the throughput. This is a complex wireless networking scenario, and there is a broad range of parameters which impact the overall system throughput. These parameters include TCP layer transmit buffer size, queue size at the MAC layer, frame capture effect, resource allocation algorithms, traffic prioritization, etc.

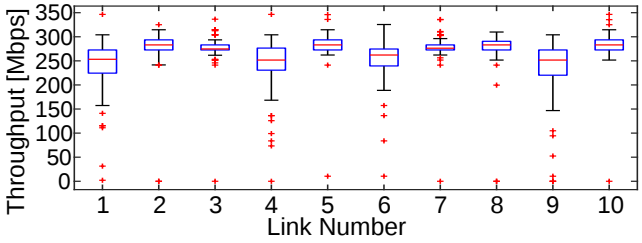

Figure 4.16: Links Throughput in the Dense Deployment.

4.4. Conclusions and Future Work

In this work, we evaluated typical deployment scenarios for 802.11ad networks. For each scenario, we generated a realistic propagation environment using our Q-D channel realization software. Based on the obtained results, we conclude the following: (1) To ensure persistent connectivity and maintain high Quality of Experience (QoE) for the end user, network operators must deploy either multiple DMG APs or a single DMG AP with multiple relay nodes. (2) Achieving full spatial sharing requires phased antenna arrays with good steering capabilities accompanied by intelligent beamforming techniques. (3) Improving the performance of a dense WLAN network requires implementing complex cross-layer solutions that optimize various parameters at each layer of the protocol stack.

For future work, we plan to conduct a detailed and comprehensive study to analyze the performance of the IEEE 802.11ad protocol in more diverse scenarios with devices of heterogeneous capabilities. In addition, we plan to model head rotation and self-blockage for Virtual Reality (VR) and augmented reality (AR) applications. Finally, since the WLAN module in ns-3 does not support any mechanism to perform handover between APs, we plan to add the support for both IEEE 802.11r for fast BSS transition and the IEEE 802.11k amendments for fast handover.

Chapter 5

Improving Frame Aggregation in 60 GHz Networks

5.1. Introduction

Introducing artificial delay can improve the performance of wireless networks. This somewhat counter-intuitive idea becomes particularly relevant for wireless systems that achieve multi-Gbps data rates, such as 802.11ad [30]. The underlying reason is frame aggregation, which plays a fundamental role in recent 802.11 standards. Its impact is amplified in the case of IEEE 802.11ad due to the large bandwidth available in the 60 GHz band—802.11ad channels are 2.16 GHz wide—and the resulting very high data rates. Any time spent for MAC backoff, inter-frame spacing, or retransmissions is highly detrimental to performance. To give an intuition, transmitting a single packet of size 1500 bytes at a moderate 802.11ad rate requires around $3\,\mu s$. In contrast, the MAC overhead for channel access is $60\,\mu s$, that is, $20\times$ larger. Thus, transmitting large frames that include as many data packets as possible is crucial; 802.11ad must use frame aggregation. Moreover, this has such a large impact on 802.11ad performance that it is crucial to exploit as many aggregation opportunities as possible. Figure 5.1 gives an example. The first timeline shows the arrival of packets at the transmit queue of a wireless 802.11ad station. Packets p_1 and p_2 arrive while the medium is busy, and p_3 shortly after it becomes free. Case I in Figure 5.1 shows the behavior of existing 802.11 devices, that is, aggregate all the data which is available when the medium is free and transmit. However, this behavior misses the opportunity of aggregating p_3 and incurs significant overhead to transmit it later. In contrast, in this chapter, we design a mechanism that introduces an artificial delay t_{wait} before transmitting. In Case II, this allows the station to aggregate all packets in a single transmission, and thus significantly reduce medium utilization. Also, the average delay is reduced. While p_1 and p_2 are slightly delayed in Case II of Figure 5.1, p_3 arrives at the receiver much earlier than in Case I, and the average delay is smaller.

Waiting for packets is particularly beneficial for bursty packet arrivals such as in Figure 5.1, which is often the case for typical Internet traffic. In infrastructure-based networks, this is highly

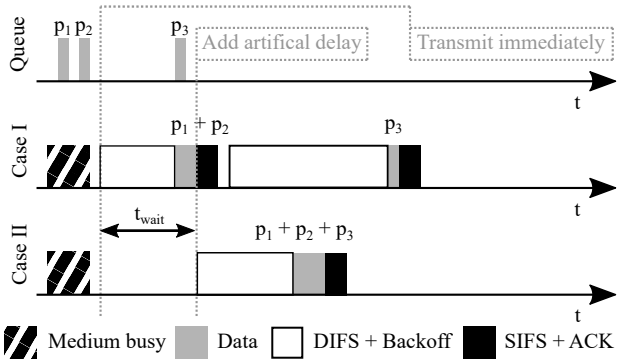

Figure 5.1: Aggregation Opportunity Example.

relevant for the uplink, where waiting time may be used by the AP or other STAs which might have accumulated more packets. Moreover, introducing a waiting time improves performance in case of high contention since nodes access the medium less often, the more packets they aggregate. As a result, fewer collisions occur. Introducing a waiting time for the AP is more problematic since traffic is usually asymmetric—forcing the AP to idle may lead to undesirable queue build-up and delay.

While the underlying idea is simple, the above scheme raises many vital questions. When should a station wait? And, most importantly, for how long it should wait? Without clear guidelines on how to answer both questions, a waiting scheme may reduce performance rather than improve it.

In this chapter, we design uplink and downlink scheduling policies that take aggregation opportunities into account [31]. Our design is local; that is, it incurs no communication overhead among nodes. It requires a STA to either wait for a certain minimum number of packets, or for a maximum amount of time, whichever comes first. This is in contrast to most existing queuing work (c.f. Section 5.2), and our work is the first to use such a policy for MAC-level frame aggregation. For the downlink case, it is beneficial to transmit to the STA with the highest number of packets in the AP queue, i.e., the one that allows for the highest level of aggregation, rather than introducing a waiting time. To this end, we use a maximum weight scheduling policy together with a waiting time limit to ensure fairness among stations with different traffic patterns. Throughput improves for both uplink and downlink since the medium is used more efficiently. We make the following contributions:

- We design and implement downlink and uplink scheduling policies that are optimized for aggregation.

- We study the benefits of maximum weight scheduling and introducing artificial delays for a range of the above thresholds and traffic patterns.

- We show experimentally that wrong parameterizations of our waiting policy are not harmful in most cases.

5.2. Related Work

Frame aggregation is a key feature to reduce the aforementioned MAC inefficiency. Starting with the 802.11n standard, wireless networks may use two types of aggregation, namely, A-MSDU and A-MPDU [32]. A-MSDU aggregates few frames has a single frame check sequence (FCS). That is, if one packet is lost during transmission, the entire aggregated frame needs to be retransmitted. In contrast, A-MPDU can aggregate more frames since it has individual FCSs for each packet, but at the same time, this incurs higher overhead. Existing studies show that aggregation can achieve up to 95% channel utilization [33], and that hybrid of A-MSDU and A-MPDU performs best [34]. Related work discusses improvements to these 802.11 aggregation schemes, such as including additional headers to allow A-MSDU to retransmit individual packets [35], and compressing the per-packet subheaders in an A-MSDU frame [36]. Moreover, other approaches allow aggregating packets addressed to different stations. This is feasible if stations use the same physical layer rate [37], or if the transmitter transmits the packets for each station on disjoint 802.11 subcarriers [38]. In [39], the authors propose a joint spatial multiplexing and packet aggregation scheme for MU-MIMO 802.11ac. However, they neither take into account the two types of aggregation supported by the standard nor consider the uplink case which is relevant for aggregation. Finally, the authors in [40] use fuzzy control to determine the optimum aggregation buffer delay before accessing the channel. Unlike our work where we consider two variables for the waiting policy, namely the maximum waiting time and the maximum number of packets, the authors only evaluate the impact of buffer delay on the end-to-end latency.

Frame aggregation can be modeled as a batch service queue since the transmitter provides service to all packets in a frame together. Early models for such queues consider that the service time is independent of the batch size [41]. However, this does not hold in our case since the transmission time increases with the number of aggregated packets. More sophisticated models [42, 43] take this into account and study additional features such as limiting the maximum/minimum batch size N [44], and allowing for server vacations [45], which model waiting times. The design that we sketch in Section 5.1 considers both a minimum batch size and a maximum waiting time. Hence, it is a batch queue with N-policy and interrupted vacation [46]. We apply this queuing policy to wireless networking. Earlier work on 802.11 networks only considers aggregation based on a non-interruptible waiting time [47]. Further, [47] does not allow the transmitter to aggregate packets while waiting for the channel to become available. In contrast, our design allows for this, which is a more realistic assumption.

Adjusting the *maximum* batch size can be beneficial, too. For instance, the optimal length of A-MSDU depends on the PER since these frames have a single FCS [48]. Further, 802.11ac only performs channel equalization at the beginning of A-MPDU/A-MSDU frames. Hence, frames should not be longer than the coherence time of the medium [49]. While these schemes deal with the maximum aggregation size, we focus on how to achieve that size given bursty traffic. Hence, these approaches are orthogonal to our work.

5.3. Scheduling Policy

For our design, we consider an infrastructure-based 60 GHz network with one AP and N STAs. We first present the details of the aggregation-aware uplink and downlink scheduling policies we introduced in Section 5.1 and then outline a method to estimate suitable policy parameters.

5.3.1. Uplink case

The key idea is to allow STAs to wait for a limited amount of time to receive more packets of the current burst and thus increase aggregation. A basic waiting policy would be to wait for a fixed amount of time whenever a STA is ready to transmit. While simple, this approach may wait unnecessarily. For instance, if the STA had to wait for the medium to become available, the chances are that its queue already holds a sufficient number of packets. Since the STA aggregates all of these packets, the per-packet overhead for that medium access may already be acceptable. In this case, waiting for a fixed amount of time may allow aggregating more packets, but the additional benefit is limited since the per-packet overhead decreases as $1/k$, where k is the number of packets in the queue of the STA. In contrast, if k is small when the medium becomes available, any additional packet that the STA receives during the waiting time significantly reduces the per-packet overhead. That is, the number of packets k in the queue of a STA is critical to decide whether to wait for more packets. Hence, we design our waiting policy based on two thresholds, namely:

1. The **packet threshold** P_s is the minimum number of packets that a STA must have, to transmit *before* the elapsed waiting time reaches the time threshold.

2. The **time threshold** T_s is the maximum waiting time duration that a STA must not exceed even if it has *fewer* packets than the packet threshold requires.

Hence, if $k \geq P_s$ the STA transmits. It also transmits if it waited for a time of T_s even if $k < P_s$. This prevents an uncontrolled increase of the artificial delay that we introduce through our waiting policy.

The transmit queue of the STA contains packets addressed to the AP and the STA must check the above thresholds whenever a new packet arrives at the queue. Specifically, when the *first* packet arrives at an empty transmit queue, the STA sets a timer to expire after a time of T_s. Whenever a packet arrives at the queue which results in $k = P_s$, the STA initiates medium access according to 802.11 (i.e., start the backoff procedure in case of CSMA/CA), and the timer is canceled. Once the STA is granted access to the medium, it aggregates *all* packets that arrived up to this point in time (including packets that may have arrived during backoff), if the maximum possible frame length of the device and the standard allow this. Otherwise, the maximum allowed number of packets are aggregated and sent, and the timer is reset to expire after $T_s - (t - t_a)$, where t is the current time and t_a is the arrival time of the oldest packet in the queue.

To avoid excessive out-of-order packet delivery, the waiting policy does not apply to retransmissions. These are handled separately from regular data packets, and medium access follows the usual 802.11 procedure. Note that also for immediate retransmission, any other packets that are in the queue at that time are aggregated with it.

STAs locally decide when and for how long to wait. Thus, our waiting policy does not incur any control overhead—any benefit that results from waiting is at *zero cost* (other than the delay itself). The specific values of P_s and T_s have a significant impact on the performance of our policy. Hence, we study them in detail in Section 5.5.

5.3.2. Downlink case

The downlink case is different from the uplink case since the transmit queue of the AP usually contains packets for multiple STAs, and forcing the AP to wait can be detrimental to performance since downlink traffic usually exceeds uplink traffic. Instead, the AP uses a maximum weight scheduling policy to exploit aggregation opportunities. When the AP gains a transmission opportunity, it transmits to the STA with the highest number of packets in the queue. This way, AP adds some *implicit* waiting time to the packets for other STAs in its queue, providing the opportunity for more packets to these STAs to arrive.

To prevent starvation for STAs with low traffic, also the AP uses a maximum waiting time T_{ap}. Whenever the AP gains access to the channel, it only transmits to the STA with the highest number of packets if $t - t_a < T_{ap}$, where t_a is again the arrival time of the oldest packet in the queue. Otherwise, it transmits to the STA that is the destination of the head-of-queue (i.e., oldest) packet.

5.3.3. Setting the Parameter

Parameters P_s, T_s and T_{ap} determine the performance of our waiting policy. The AP's maximum waiting time threshold T_{ap} essentially limits unfairness among downlink flows. It is primarily of importance in case low rate flows with strict delay constraints compete with high rate flows. In this case, T_{ap} should be set to the desired maximum wireless delay.

Section 4.2 shows that the optimal value of P_s and T_s depends on the traffic pattern and the number of nodes in the network. Designing in detail a mechanism to estimate the number of nodes in the network a priori is out of our scope since we focus on analyzing the waiting policy itself. However, we provide a simple *a posteriori* method, and evaluate in Section 5.5 whether it is suitable for our waiting policy.

Essentially, a node can just follow an adaptive trial-and-error approach for P_s and T_s. That is, while transmitting data, it tries P_s and T_s values and observes their impact on performance. If performance improves, i.e., an increase in P_s and T_s did allow to include more packets in a frame, the node continues increasing the parameters. Otherwise, it returns to a previous value known to provide gains. This requires no coordination among the nodes of the network since

each node can probe P_s and T_s values independently. Nodes try new values in two cases—first, periodically to determine if a different parameterization is more beneficial, and second, whenever performance decreases without having changed the parameters. This allows nodes to adapt dynamically to changes in the network. Additionally, if network conditions are stable, nodes can deduce throughput and delay trends after probing some P_s and T_s combinations. As a basic approach, nodes can attempt to fit curves on the values they observe, and use the curves to estimate performance for new parameter combinations. For obvious reasons, P_s should never be set higher than the maximum aggregation level allowed by the standard.

In Section 5.5 we analyze whether the behavior of our waiting policy is suitable for such a trial-and-error approach. Moreover, we show that the above fitting approach performs well, too, yielding good Sum of Squared Errors (SSE) values when comparing our estimation to the actual performance.

5.4. Scenario

In the following, we describe the network scenario and the traffic pattern that we consider in our evaluation.

5.4.1. Network

We consider an indoor 60 GHz network with one AP and N STAs following the 802.11ad standard. All STAs are located in one room and have a line-of-sight link to the AP. Even though the STAs may use directional beamforming patterns to transmit data to the AP, recent work [2, 7] shows that consumer-grade phased antenna arrays for 60 GHz devices exhibit significant side lobes. Hence, we consider that all STAs interfere but also overhear each other if they transmit simultaneously. This reduces the number of collisions due to deafness, and is thus a worst case scenario for our waiting policy, which particularly benefits from collisions in the uplink. That is, in a scenario with deafness the performance of our waiting policy would be even better than our results in Section 5.5. Further, we use CSMA/CA at the MAC layer, which is likely to become the main protocol for medium access in 802.11ad hardware. At the time of writing, existing consumer-grade 60 GHz devices only implement CSMA/CA at the MAC layer [7]. We consider an error-free channel where all packet drops are due to collisions only. In addition, both AP and STAs use the same transmission rate for communication.

5.4.2. On-Off Markov Model Traffic Pattern

The statistical characteristics of packet arrivals have a strong impact on aggregation opportunities. The burstier the traffic, the higher are the benefits of the aggregation policy described in this work. For the evaluation, we use a basic On-Off Markov Model (OOMM) traffic generator to simulate bursty application layer traffic.

The OOMM models the traffic as a two-state Markov model. In the "on" state, the traffic generator sends packets with a fixed length at a constant data rate, whereas in the "off" state it does not transmit at all. The traffic generator performs random experiments and based on the outcome switches back and forth between both states. The duration of the "on" and "off" periods follows truncated exponential distributions. Figure 5.2 shows an OOMM traffic output example. The model allows us to adjust the burstiness of the traffic by tuning the probabilities of switching between states (and thus the distributions of state duration) as well as the packet generation rate. Table 5.1 gives an overview on the statistical parameters of OOMM.

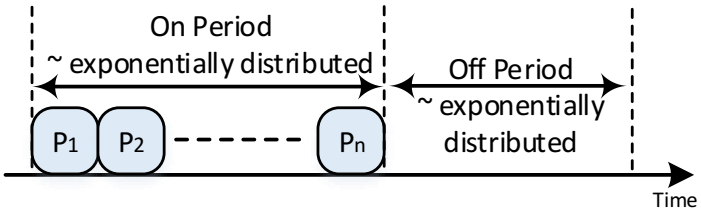

Figure 5.2: On-Off Markov Model example

Table 5.1: On-Off Markov Model Traffic Pattern Parameters

Parameter	Distribution	Value
"On" duration	Truncated Exponential	Constant Data Rate = 1 Gbps
		Mean = 600 μs, Bound = 6 ms
"Off" duration	Truncated Exponential	Mean = 2 ms, Bound = 20 ms

5.5. Evaluation

5.5.1. Simulation Setup

Practical evaluation of wireless networking in the 60 GHz band is challenging due to the lack of flexible consumer-grade and experimental hardware. This prevents us from evaluating our waiting policy in practice since it interacts with the MAC layer and the PHY. Hence, we must resort to simulations to assess our scheme in 802.11ad networks. Specifically, we use our ns-3 IEEE 802.11ad model.

5.5.2. Experiment Design

We classify our results by scenario, that is, uplink/downlink using UDP as a transport protocol. Table 5.1 summarizes further parameter values of our traffic patterns. For our results, we consider the following metrics:

- **Medium busy time**, which we compute as the ratio of the total time spent transmitting—including preamble, header, and payload of all frames—-to the total simulation time.

- **Total throughput** received at the MAC S-AP of the access point (uplink case) or the sum of this throughput received at all STAs (downlink case).

- **Packet Delivery Delay**, which we measure from the moment the application generates a packet until it is received successfully by the intended receiver.

We note that the results of each experiment are the average of 20 to 60 seconds of simulation time. Moreover, all our results in this section show the performance gain of our waiting policy *on top* of regular frame aggregation gains.

5.5.3. Results

5.5.3.1. Uplink Scenario

We start with the uplink scenario in which multiple STAs have some data to upload to the AP. For each experiment, we study a range of P_s and T_s values. In most cases, we set $P_s \in [1..200]$ in steps of 20 packets, and $T_s \in [0..2]$ ms in steps of 100 μs. Further, we consider $N \in [5, 10, 15]$ STAs. Figure 5.3 depicts the medium usage for different packet thresholds P_s and time thresholds T_s for the case of 15 STAs. The black dot at $P_s = 1$ and $T_s = 0$ is the baseline, since with these parameters a STA does not introduce any artificial delay and transmits as soon as it has at least one packet. In Figure 5.3, the medium busy time initially rises for all packet thresholds until reaching a certain time threshold, beyond which it decreases significantly. The underlying reason is that each STA can aggregate more packets, thus reducing MAC overhead and specifically time spent for backoff. Figure 5.4 confirms the corresponding throughput increase. The medium usage decreases again beyond the aforementioned time threshold occurs as soon as we aggregate enough to deliver all packets. Beyond this threshold, further aggregation increases the channel idle time. Again, this matches Figure 5.4, which shows that the throughput increase stabilizes after that time threshold.

Further, Figure 5.4 shows that the higher the packet threshold, the higher the throughput gain that we can achieve for a certain time threshold. The reason is that for low values of P_s we often wait for less than T_s since we receive enough packets to satisfy the packet threshold before reaching the time threshold. Hence, we aggregate less. Conversely, if we set T_s to a large value, we wait longer on average and thus aggregate more.

Figure 5.5 shows the delay for 5 and 10 STAs. The delay includes propagation delay and channel access delay. For both cases, the delay decreases until reaching a minimum. Beyond that, it increases quasi-linearly with the time threshold but changes its slope at a certain point. This slope change occurs for the time threshold beyond which, *on average*, the policy hits the P_s threshold before it reaches the T_s threshold. This is the reason why the slope change takes

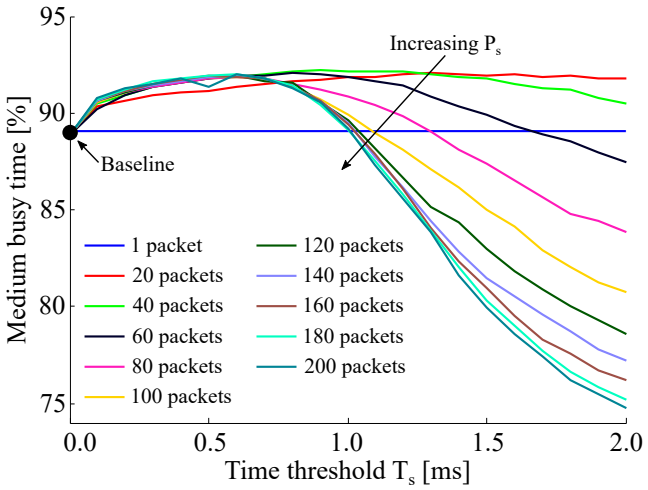

Figure 5.3: UDP uplink scenario: medium usage

place at a lower T_s for lower P_s. We observe similar slope changes in Figure 5.4—the throughput increases less beyond that point since we wait for fewer packets. Note that Figure 5.4 shows the case for 15 nodes, while Figure 5.5 depicts the results for 5 and 10 nodes. While we do not show the throughput figure for 10 nodes due to space constraints, we observe similar effects as those for the case of 15 nodes. These slope changes show the importance of the packet threshold. While P_s limits the throughput increase, it also limits the delay increase, allowing a node to use a larger T_s. For traffic with a highly irregular burst spacing, we expect P_s to have a large impact since it prevents waiting if a node has enough packets.

Further, Figure 5.5 shows that the delay improvement is much larger for 10 STAs than for 5 STAs. This difference is because, the more STAs contend for access to the channel, the higher is the probability of collisions. Such packet losses result in very high MAC overhead. Introducing an artificial delay allows nodes to aggregate more and thus access the channel less frequently, hence reducing the probability of collisions. This has a more significant impact for 10 STAs (and even more so for 15 STAs), since with 5 STAs the collision probability is low. Thus, waiting only provides a slight delay improvement in the latter case.

Figure 5.6 shows an overview of the gains for different number of STAs. For clarity, we only show the results for the P_s and T_s values that maximize the gain in each scenario and for each metric. As expected, the higher the number of STAs the higher the medium usage, delay, and throughput gains that we obtain. We observe that our waiting policy is particularly beneficial in terms of air interface delay, achieving up to 80% improvement. This highlights the relevance of artificial delay in contention scenarios.

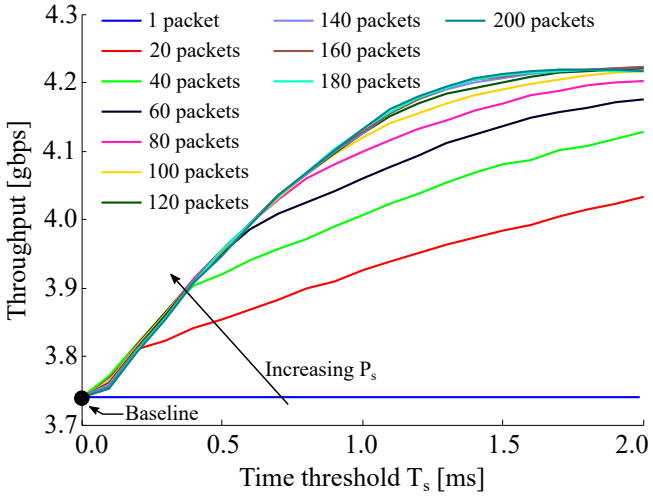

Figure 5.4: UDP uplink scenario: total throughput

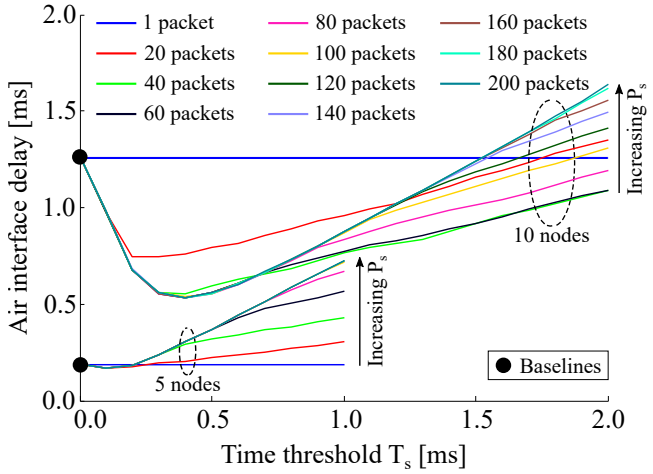

Figure 5.5: UDP uplink scenario: air interface delay

5.5.3.2. Downlink Scenario

The UDP downlink is fundamentally different from the uplink case, since only the AP accesses the channel, and thus no collisions occur. Figure 5.7 depicts the throughput and delay gains in this scenario where all STAs generate the same amount of traffic. In this case, our baseline to compare against is a FIFO scheduler with aggregation, which sends packets in the order of arrival. Figure 5.7 shows that the maximum weight scheduling policy—and the implicit delay for STAs that did not yet accumulate as many packets—allows us to achieve up to 247 Mbps net throughput gain. In this case, the gain is exclusively due to the burstiness of the traffic pattern, since this scenario does not suffer from collisions.

Finally, Figure 5.8 shows the average end-to-end delay with respect to the time threshold T_{ap} for stations with different traffic patterns. We note that the higher the value of T_{ap}, i.e., the

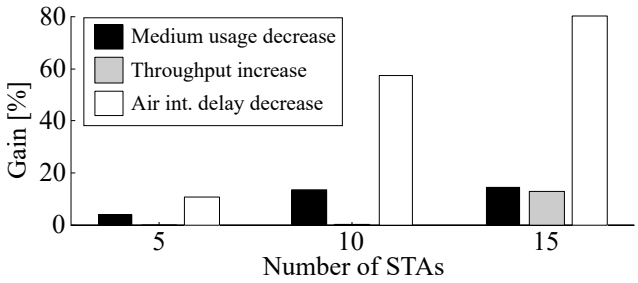

Figure 5.6: UDP uplink results

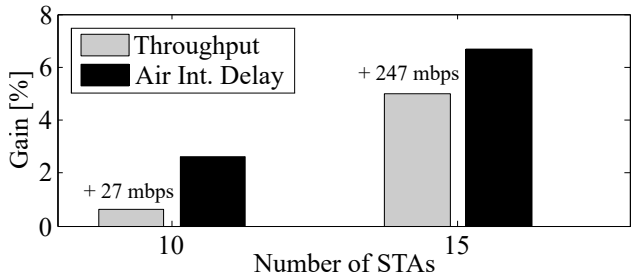

Figure 5.7: UDP Downlink Results

longer the packets may reside in the queue, the more the average end-to-end delay increases. This increases short-term unfairness as packets are delayed until flows which allow for more aggregation are served. Reducing the time threshold, in turn, reduces this unfairness and ensures that a STA with low traffic intensity is served on time.

5.5.3.3. Parameter Estimation

Finally, we discuss whether the parameter estimation method that we suggest in Section 5.3.3 is suitable for our waiting policy for the uplink case. Our results show that small increases in P_s and T_s do not cause abrupt changes in performance. That is, following a simple trial-and-error approach is unlikely to produce significant performance penalties even in case of wrong P_s/T_s parameterizations. Finally, we also evaluate the curve fitting approach sketched in Section 5.3.3 for the case of delay performance. We obtain an SSE value of 0.0282 ms, which suggests that this method is also well suited to estimate P_s and T_s. The corresponding graphs are omitted due to space constraints.

5.6. Discussion

Our results in Section 5.5 show that although it may be counter-intuitive, introducing an artificial delay in wireless networks may significantly increase performance. Most importantly, we provide insights into when and how long a STA should wait, as well as the fundamental tradeoffs

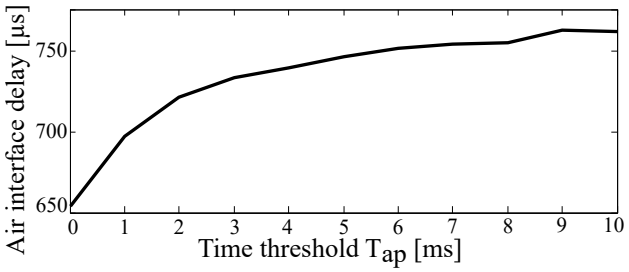

Figure 5.8: UDP Downlink Scenario: Air Interface Delay

of such artificial delay.

Performance. Waiting exhibits a significantly different behavior in the uplink compared to the downlink due to contention (c.f. Section 5.5). However, this effect is not necessarily limited to the uplink, since a network with multiple APs in the same interference domain would have to deal with potentially high contention. Such a scenario is particularly relevant for 60 GHz networks since such networks may require multiple APs per room to ensure coverage. Further, consumer-grade 60 GHz devices are likely to exacerbate contention due to significant side lobes [7].

Parameterization. In most cases, using our waiting policy without prior knowledge of network conditions is safe. Conservative time and packet thresholds typically provide gains, as discussed in Section 5.5.3.3. Based on our results, we provide some rough recommendations on how to set P_s and T_s, as an alternative to estimation methods such as the one we suggest in Section 5.3.3. First, a node should evaluate whether (a) it observes frequent collisions and (b) its transmission queue becomes empty periodically. This provides a basic notion on the network conditions. If neither (a) or (b) occur, the node need not use our waiting policy since the baseline aggregation policy provides the same gains. However, such a permanently backlogged case is unusual. Second, a node should set P_s and T_s more conservatively the less contention it observes, and the longer the intervals are at which its queue becomes empty.

Cost. A key feature of our waiting policy is that it causes *zero communication overhead*. In other words, a node may *test it* to estimate the possible gains with only a minimal risk, i.e., there is little to lose in terms of potential delay, but significant performance improvements to gain.

5.7. Conclusion

We design an aggregation-aware MAC scheduling policy for 802.11ad wireless networks which introduces a maximum weight scheduling policy for the downlink and an artificial delay for the uplink. In case of bursty traffic, this enables nodes to aggregate more packets of the current burst at the MAC layer. This is particularly beneficial for 802.11ad since its channel access overhead per transmission is extremely large. Our policy allows STAs to wait for either a minimum amount of packets to aggregate or a maximum amount of time, whichever comes first. We implement and evaluate this policy in ns-3 for uplink and downlink scenarios. Our policy is

beneficial in two ways. First, since more aggregation reduces the number of channel accesses, it reduces collisions in case of high contention. Second, for bursty traffic, it avoids that a small number of packets at the end of a burst require costly individual medium access. In our experiments, we achieve up to 480 Mbps throughput increase and 80% channel access delay reduction.

Chapter 6

Practical Evaluation of IEEE 802.11ad COTS Devices

6.1. Introduction

mmWave communications narrow the gap between wired and wireless networks. While both domains are fundamentally different, mmWave wireless networks feature highly directional links and large bandwidths which result in low interference and high throughput, similarly to the wired case. Early work in this area even dubs mmWave links as "pseudo-wired" [50]. Such links can alleviate the characteristics of wireless links that hinder the seamless integration with the wired part of the network. For instance, the limited interference results in less random packet losses, which in turn improves the operation of the Transmission Control Protocol (TCP) that would otherwise interpret such losses as congestion [51]. Also, directional links are less prone to fluctuations due to multi-path effects, yielding high and stable throughput as in the wired case.

On the other hand, some of the features of mmWave links are still inherently wireless. For example, packet loss can be high due to the inherent characteristics of the wireless medium triggering link adaptation algorithms (rate adaptation, beamforming); and mmWave access points (APs) still transmit beacons to allow clients to discover them, which may interfere with data transmission. More importantly, practical phased array antennas in Commercial off-the-shelf (COTS) devices typically generate wide beam patterns with strong side lobes instead of perfect "pencil beams" [2, 7]; such patterns can increase interference among neighboring links. Thus, the characteristics of mmWave links correspond neither to the ones of wired links nor to the ones of traditional wireless links. To date, it is not well understood how such links with mixed characteristics behave in practical WLAN deployments.

In this chapter, we fill this gap by studying the performance of 60 GHz WLANs comprised of COTS APs and clients. Earlier work in this area is limited to isolated mmWave links [7, 52–54], whereas we consider a complete network with up to eight stations that fully comply with the IEEE 802.11ad standard [1]. This is the first study of larger IEEE 802.11ad deployments. We extract

detailed information regarding the operation of the network across all layers of the protocol stack directly from our 802.11ad hardware. To this end, we install LEDE [55] on a COTS router operating in the 60 GHz band. This allows us to obtain unprecedented insights into the operation of mmWave networks for the case of 802.11ad. In contrast, earlier work only has limited access to the inner workings of 802.11ad hardware and thus can only conjecture regarding the effects it measures. As a result of the aforementioned mixed characteristics of mmWave links, many of our observations are unexpected and challenge the prevailing intuition regarding communication links. Specifically, our insights are as follows:

1. The very high data rates in mmWave networks exacerbate the inefficiencies of CSMA/CA resulting in strong link fluctuations. Thus, dynamic TCP buffer sizing [56] is crucial and much more relevant than in wireless networks operating in the 2.4 GHz and 5 GHz bands. The deafness due to the directional communication in the mmWave band further exacerbates this issue.

2. The susceptibility of mmWave to blockage and human mobility requires frequent beamforming training and tracking to avoid link interruption. This signaling overhead is particularly significant in large dense deployments with many clients, and as a result, degrades the performance of TCP flows. This requires either good network planning to distribute clients among access point to reduce potential blockages or intelligent tracking methods to exchange fewer messages to establish and maintain a communication link. We highlight this effect in section 6.4.

3. While mmWave frame aggregation is highly beneficial to reduce the impact of medium access overhead, the packet error rate increase due to high aggregation levels is much more pronounced than for traditional wireless and wired networks. The former supports aggregation up to ten milliseconds and the latter in the form of very large “jumbo frames”. In contrast, mmWave links are limited to frame aggregation on the order of microseconds [7].

4. The beacons that 802.11ad APs periodically transmit for each of their beam patterns to enable directional communication can interact with data transmissions. As a result, we observe spikes in the RTT which in turn impact the behavior of TCP and may cause the above unfairness. Other effects, such as increased backbone delays, result in the same issues.

5. COTS mmWave devices use beamforming codebooks with static entries for all the frequency channels to steer the antenna array into different directions in space. In section 6.8, we highlight the effect of those static entries on the directivity of practical phased antenna arrays.

6. Link orientation, the distance to an interferer, and the specific choice of beam patterns play a significant role for spatial reuse and determine whether two links can operate simulta-

neously or not. This is different from traditional wireless networks where interference is primarily related to the distance among nodes.

7. Current devices select beam patterns based on the received signal strength during beamforming training. Given that beam pattern selection plays an essential role in achieving spatial sharing in dense deployments, taking interference into account when selecting the beam pattern improves network performance and allows for enhanced spatial sharing. This allows improving overall throughput by up to a factor of two.

The chapter is structured as follows. In Section 6.2.3 we explain our experimental methodology. In Section 6.3 we survey related work in the field of practical mmWave network measurements. After that, in Sections 6.4 to 6.8 we present our results on carrier sense multiple access with collision avoidance (CSMA/CA) behavior, frame aggregation, delay, spatial sharing, and beam patterns respectively. Then, we discuss the implications of our results in Section 6.9. Finally, Section 6.10 concludes the chapter. The material in this chapter is taken from [57, 58].

6.2. Experimental Methodology

In this section, we list the set of devices that we use and the 60 GHz wireless chipset they utilize and then present our WLAN setup.

6.2.1. Devices

The **TP-Link Talon AD7200 [59]** was the first commercially available 802.11ad router released in June 2016. For 802.11ad, it uses the QCA9008-SBD1 module with the QCA9500 chipset from Qualcomm, supporting single-carrier data rates up to 4.6 Gbps. The 32-element phased antenna array is located on a separate board and connected to the chipset with an MHF4 cable. The router also includes an 802.11n/ac solution from Qualcomm. The routers can achieve up to 2.2 Gbps using MCS-12 over the 60 GHz wireless interface itself, but connectivity to other networks is limited by the speed of the 1 Gbps Ethernet interface.

The **Netgear Nighthawk X10 Smart WiFi Router [60]** was released around October 2016. It uses the same module from Qualcomm for 802.11ad as the one used by Talon. In addition to the 1G Ethernet ports, it also has a 10-Gigabit LAN SFP+ interface. While the 10G port theoretically provides multi-Gigabit speeds, we found that in practice the maximum throughput (with MCS-12) is limited to around 2.3 Gbps.

The **MikoTik wAP 60G Solution [61]** was released around August 2017. Similar to the two previous devices, the MikroTik Router uses the QCA9500 chipset from Qualcomm. The device integrates a Uniform Rectangular Array (URA) of 6x6 elements. Besides, the Equivalent Isotropically Radiated Power (EIRP) is around 35 dBm which allows the device to reach a distance of 200 m. Similar to the TALON Router, the MikroTik Router is limited to 1 Gbps of throughput due to

the use of a 1 Gigabit Ethernet port. The device uses a proprietary operating system (RouterOS) which gives the device all the necessary functions to operate as a router. The RouterOS provides a graphical user interface to configure and control some aspects of the physical layer such as the selected beam pattern. The device is mainly intended for Point-to-Point scenarios only and cannot support multiple clients.

The **Acer Travelmate P446-M [62]** laptop, released in April 2016, has the client-version QCA9008-TBD1 of the module used in the Nighthawk and Talon routers, which includes 802.11ac, 802.11ad and Bluetooth chipsets. The host connects to the module using an M.2 slot, runs Linux OS (Fedora 24, kernel 4.x) and uses the open source wil6210 wireless driver to interface with the chipset. It comes with same 32-element phased antenna array as the routers.

6.2.2. QCA9500 Module Overview

The QCA9500 [20]chipset from Qualcomm Qualcomm supports the IEEE 802.11ad protocol with a maximum data rate of 4.6 Gbps. The module comes with two boards:

- Radio frequency integrated circuit (RFIC) board with a phased antenna array of 32 elements.
- Baseband (BB) board connected to the RFIC board through coaxial cable. The BB has a processor which runs the corresponding firmware.

In the following subsections, we provide an overview of the PHY and MAC capabilities of the module. Additionally, we briefly describe the features of the module's firmware and its corresponding wireless driver in Linux OS.

6.2.2.1. QCA9500 Linux Wireless Driver

Linux OS uses the open source driver *wil6210* to communicate with the firmware running on QCA9500 module. *wil6210* is a loadable Linux Kernel Module (LKM) that uses *cfg80211* framework only. This implies that the driver is a *HardMAC* where the MLME is completely implemented in the hardware. The module uses a closed binary firmware named *wil6210.fw* to realize 802.11ad functionalities. Besides the firmware, the module utilizes a file named *wil6210.brd* which serves as a codebook for defining the beam patterns generated by the PAA. Linux OS automatically loads the *wil6210* driver upon boot-up which in turn downloads the firmware to the Baseband processor. This allows the module to use an up-to-date firmware which includes new capabilities and fixes some performance issues instead of having hard-coded firmware. Without the firmware, the module cannot start up. The driver provides many experimental features that can be used for research purposes. We list some of the *wil6210* driver features at the time of writing:

- Support AP operation mode using either *wpa_supplicant* or *hostapd* applications. Up-to eight stations can attach and communicate with a single AP.

- Managed mode which allows the host device to work in station mode. In this mode, a device scans all the frequency channels to find all the operational APs and reports their attributes to the user-space. These parameters include Service Set Identifier (SSID), EDCA parameters, vendor-specific information, and signal strength. Additionally, this mode allows the device to authenticate and associate with one of the available APs.

- Sniffer mode to capture raw MAC frames. The sniffer attaches a *RadioTap* header to each received frame. The *RadioTap* header includes details about the PLCP such as MCS, signal quality, channel center frequency, etc.

- A set of Wireless Module Interface (WMI) commands and events to communicate with the QCA9500 module and report information back to the *wil6210* driver.

- Debugfs is a debugging facility to examine and change the networking state of a firmware. For example, the user can check the index of the current active transmit and receive sectors for both the local station and the peer station.

Figure 6.1 shows the network model for *wil6210* driver in Linux OS. *iw* is a netlink (nl80211) command line interface that can be used to configure and control the Linux driver of wireless devices.

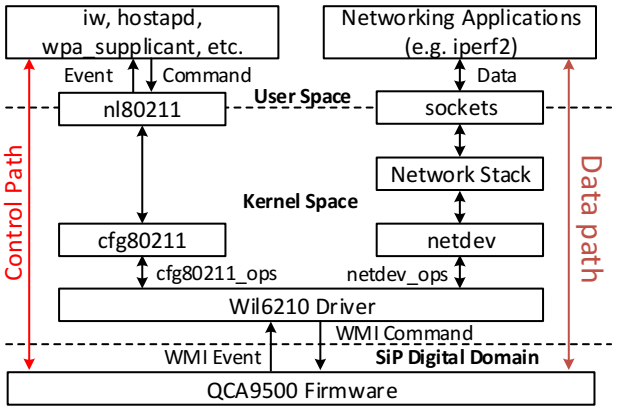

Figure 6.1: Wil6210 Network Model

6.2.2.2. Physical Layer

Although the IEEE 802.11ad amendment supports OFDM PHY specifications for data transmission, the latest release of the WLAN standard [1] considers the use of OFDM as obsolete. This because the PHY Header of the the OFDM does not provide enough information for devices using SC PHY to determine the transmission duration of OFDM frames. This creates a problem when the network is made of heterogeneous devices as those devices cannot operate properly. As a result, 802.11ad chipset manufacturers implement SC PHY in their devices only. QCA9500

supports all the data rates MCS (1-12) for SC PHY layer, producing a theoretical throughput that ranges from 385 Mbps to 4620 Mbps. Management and control frames (such as DMG Beacons, SSW frames, Associations, Probe, etc.) are transmitted using Control PHY which corresponds to MCS-0. This ensures a high probability of successful reception as MCS-0 is ideal for low SNR region. For multicast traffic, the driver enforces the use of MCS-1 as it provides the highest reliability for data packet delivery. The module supports three channels 1,2, and 3 with center frequencies of 58.32, 60,48, and 62.64 GHz.

6.2.2.3. Channel Access

The wireless medium is divided into a sequence of BIs where each BI in turn is divided into both BTI and DTI access periods. Both A-BFT and ATI access periods are not present. The whole DTI is reserved for CBAP which uses EDCA framework. The EDCA defines a set of access categories (ACs) to provide QoS provisioning within the WLAN. Table lists EDCA parameters as announced by the QCA9500 module when operating in an AP mode. At the time of writing, both SPCA and dynamic allocation of service periods are not implemented in the firmware. The firmware uses TXOP to reserve channel for a period of time after channel contention.

6.2.2.4. Beamforming Training

Due to the high overhead and complexity to establish optimal directional beams in the intended direction, existing technologies such as the IEEE 802.11ad standard rely on a set of predefined beam patterns referred to as codebook. The current chipset uses a phased antenna array of 32 elements. The elements are distributed on the surfaces of the printed circuit board (PCB) in an irregular way. Authors in [2] measured the 2D and 3D radiation patterns for all the sectors in the codebook. During the BTI access period, the firmware uses 32 sectors for advertising the network and performing initial beamforming training. Since the A-BFT is not present, beamforming training is performed in the DTI access period where both devices utilize 36 sector to perform TXSS BF. The initiator of the beamforming training sends SSW frames to the responder. A station can perform beamforming training only with the AP and vice versa. The selection of the best sector for communication is based on the highest SNR value. Authors in [2], performed reverse engineering to the firmware of the QCA9500 module to print the SNR for each received SSW frames after the completion of an SLS phase.

6.2.2.5. Frame Aggregation

Frame aggregation is one of the most important features that has been introduced since the 802.11n amendment. It improves MAC layer performance by aggregating multiple small packets into a larger packet which reduces the overhead introduced by the contention procedure. The current firmware of the QCA9500 module is limited to A-MPDU aggregation only. For each A-MPDU aggregated packet, the MAC layer adds a sub-header for each aggregated frame in addition

to the MAC layer header. In this way, if the subframe is erroneous, the receiver can selectively request the sender to re-transmit the erroneous subframe. Due to memory constraints, the current hardware permits to aggregate either up-to 32 packets or 65535 bytes whichever comes first. This puts a limit on the achievable throughput as opposed to the one defined by the standard.

6.2.2.6. Rate Adaptation

Rate adaptation is an integral part of all the wireless and mobile systems. Due to the dynamic nature of the wireless channel, it is almost impossible to set the optimal transmission rate for a wireless link. A rate adaptation algorithm monitors the quality of the communication link and frequently adapts transmission rate based on various metrics such as SNR, number of successfully transmitted packets, the number of failed packets during transmission window, etc. At the time of writing, no rate adaptation algorithm has been proposed in the literature for the IEEE 802.11ad standard. The current firmware implements a proprietary rate adaptation algorithm and exposes a set of parameters to manipulate it.

6.2.2.7. Codebook Design

The user can change the excitation of the antenna elements and generate custom beam patterns. The codebook can accommodate up to 128 beam patterns. The user can define both the number and the order of the beams to be used during both BTI and SLS. Further, the firmware provides a granularity to set this per station. By default, the firmware uses 32 transmit sectors during BTI and 35 SLS phase. Only a single reception pattern is used which has a quasi-omni shape [2].

6.2.3. Measurement Methodology

We use the devices listed in Section 6.2.1 for our WLAN setup. Besides, we differentiate between four sets of experiments based on the number of clients. The first and last sets of experiments (Sections 6.4 and 6.7) use multiple Talon routers, with some in AP mode and the rest in client mode, to have a WLAN-like environment. The second set of experiments (Section 6.5) uses a router again in AP mode and a laptop in client mode. The third set of experiments (Section 6.8) uses a router in AP mode together with custom hardware to measure the beam patterns. The fourth set of experiments (Section 6.6) consists of a router running in AP mode connected over a 1G Ethernet link to a high-end desktop and a laptop operating in client mode and associated to that AP.

We flash LEDE [63] on the Talon routers – a Linux operating system based on OpenWrt. This allows us to manage these routers and collect measurement performance in a centralized way. For these routers, we either use the Ethernet interface or the 2.4/5 GHz WLAN interface for management and control. Note that the driver, firmware, chipset, and the phased antenna array in the Talon router, in this case, is the same as in the laptop.

The Wilocity driver (wil6210) available for the Linux operating system allows setting the device in either AP mode, monitor mode, or client mode. Besides, it provides some useful information regarding the wireless link such as the current modulation and coding scheme (MCS), MAC layer throughput, Signal Quality Indicator (SQI), Beamforming Training Status (OK/-Failed/Retrying), Tx/RX sector IDs for itself and the peer device, and the number of received packets per MCS. We log all these parameters every 150 ms. The Wilocity firmware implements both a proprietary beamforming algorithm for selecting the best transmit sector and a rate adaptation mechanism for MCS selection. As a result, it does not allow fixing either the MCS or transmit/receive sector IDs. However, this changed with the firmware version «4.1.0.55», which helps to set the beam pattern index to be used, even though the device still maintains the beamforming algorithm when this index is not set manually. Experiments from Section VI have been done with this newer firmware «4.1.0.55», while the rest of the experiments were performed with the «3.3.3.7759» version.

We use iPerf2, a network benchmarking tool, for generating TCP/UDP traffic. Unless otherwise stated, the reporting period of iPerf is 100 ms. When using TCP with iPerf, iPerf reports the throughput, the Round-Trip Time (RTT), and the Congestion Window (CWND).

6.3. Related Work

60 GHz Practical Work. Related work studies the performance of practical 60 GHz networks using both custom hardware platforms such as software-defined radios [64–66], and commercial off-the-shelf devices [7, 52, 67–69]. The former are typically limited in terms of bandwidth and focus on the physical layer only, thus not allowing for an analysis of upper-layer issues such as in this work. The latter use pre-standard 60 GHz devices that do not allow for networking and thus typically study individual links only. While such papers analyze the performance of TCP on mmWave links and emulate the deployment of multiple APs [7, 53, 70], the lack of multiple stations and a backbone network strongly limits their insights compared to our analysis.

mmWave MAC Layer. Medium access control plays a fundamental role on the alleged "pseudo-wired" behavior of mmWave links. Earlier work studies its performance both analytically [33] and in simulation [71]. While 802.11ad allows for multiple access based on time-division [1, 72], existing hardware implements CSMA/CA only. The key problem of CSMA/CA in directional mmWave networks is deafness. Related work suggests virtual carrier sensing [73, 74] to address this issue. However, practical work on 802.11ad devices shows that spatial sharing is limited due to the highly irregular beampatterns of commercial off-the-shelf devices [7]. In contrast to the above earlier work, we are the first to study the medium sharing behavior of 802.11ad in practice.

Frame Aggregation. Due to the large bandwidth available in the 60 GHz band, frame aggregation has a particularly large impact on the throughput performance of 802.11ad as explained in Chapter 5. While WiFi networks operating at lower frequencies must generate frame sizes in the order of milliseconds to achieve significant gains [32, 49, 75], 802.11ad aggregation in the order

of microseconds results in even larger gains [7]. Still, the challenges of frame aggregation at low frequency [37, 76] are often valid also for higher frequencies. For instance, related work shows that high levels of aggregation increase packet error rates in 802.11n/ac [35, 36, 49]. We extend this analysis for the case of 802.11ad.

Beam Pattern Synthesis and Selection. The authors in [2, 7] report the beam patterns for the current COTS 802.11ad enabled devices. They found that these devices generate very wide beam patterns instead of directional ones with a clear steering angle. To overcome this issue, the authors in [77] develop a technique to dynamically adapt the beam patterns in these devices to optimize the link quality. Other state-of-the-art beamforming techniques are surveyed in [78].

6.4. Behavior of CSMA/CA

Current studies on WiGig and 802.11ad in the literature focus on characterizing and benchmarking the performance of a single communication link for different placements and in different scenarios [7]. While this provides valuable insights, it is very important to move one step further and study the performance of COTS 802.11ad devices in large, dense deployments, i.e., more complex WLAN setups. In addition, these studies ignore the interaction between the protocol aspects of the IEEE 802.11ad standard and the operations of the transport protocols such as TCP. In this section, we try to answer the following questions: i) how do 802.11ad devices share the wireless medium? ii) What is the impact of transport layer buffering on the perceived throughput, latency, and fairness between competing TCP flows? iii) What is the optimal size for TCP buffer? iv) How does the directionality in mmWave networks impact the performance of transport layer protocols? To answer these questions, we carry out extensive measurement campaigns where we deploy eight stations and a single AP within a room of size 6x3x5 meters as shown in Figure 6.2. All the stations are placed at the same height, and each one starts a single TCP flow at the same time towards the AP, or receives a single TCP flow from the AP. Each experiment lasts for 60 seconds, and we repeat each experiment between 5 to 10 times. We differentiate between two scenarios based on the direction of the TCP flows, namely uplink or downlink.

Figure 6.2: Dense Deployment Setup Layout.

6.4.1. Downlink Scenario

In the downlink scenario, the AP generates a single bulk TCP flow towards each station. The channel access comes mainly from the AP, since it has to multiplex transmissions among several stations. Thus, the impact of the stations contending to transmit TCP ACKs back to the AP is small compared to the channel contention required to transmit the high amount of TCP data segments to those stations. This isolates to some degree the interactions between TCP behavior and CSMA/CA protocol aspects. Moreover, this allows us to understand the impact of Linux TCP receive buffer size on the aggregated throughput and the RTT.

Figure 6.3 shows the aggregated throughput for all the flows with respect to the number of stations and for various TCP receive buffer sizes. By default, LEDE uses 4 MBytes for the size of both TCP send and receive buffers. We benchmark the performance using additional values for the TCP receive buffer size: a small value (1 Mbytes), the default value (4 MBytes), and a large value (32 MBytes). For a fixed number of stations, the throughput increases proportionally with the size of the TCP buffer. For 6 and 8 stations, the throughput starts to stabilize for each of the buffer size values. However, the most important observation is that the aggregate throughput drops with the number of stations. This is due to the inefficiencies of CSMA/CA. Additionally, due to the nature of the mmWave, stations cannot distinguish if a packet loss is due to collision, deafness, or shadowing. For that reason, a station frequently tries to access the wireless medium to perform beamforming training with the corresponding peer station. This is a further reason why the aggregated throughput degrades with increasing number of stations. This additional training overhead that does not exist for wireless networks operating in the microwave band. There are numerous experimental and analytical studies [79–82] that evaluate TCP performance in dense WLAN settings in the microwave band. All of these studies show that TCP downlink throughput scales quite well with increasing number of clients if stations do not generate any uplink traffic. For 802.11ad networks, stations have to access the channel occasionally to perform beamforming training even if they do not have any data for transmission.

Figure 6.4 depicts the average RTT for all of the flows. The RTT rises as we increase both the number of stations and the TCP buffer size. This is due to the fact that the AP has to multiplex between different flows and TCP flows tend to fill all available buffer space, which results in queue build up at the wireless interface of the AP. In addition, the contention of the returning TCP ACKs and the frequent beamforming training play a role in the rise of the RTT. As a result, efficient methods for enabling fast beam alignment are key aspects in the design of mmWave protocols since the overhead of beamforming training has a clear impact on the performance of TCP in the downlink.

6.4.2. Uplink Scenario

In contrast to the downlink scenario where a single station frequently contends to access the channel, the uplink scenario multiple stations try to access the channel to transmit their data to

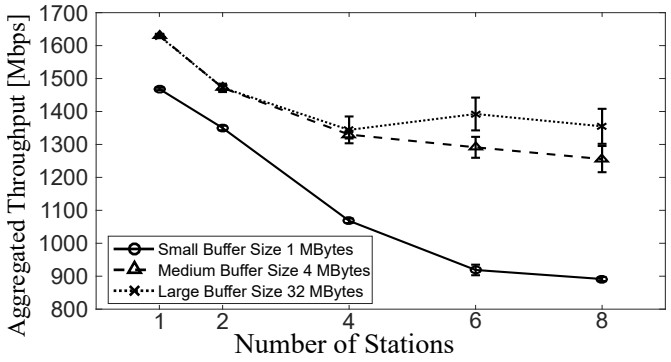

Figure 6.3: Aggregated Throughput Comparison for Downlink Scenario.

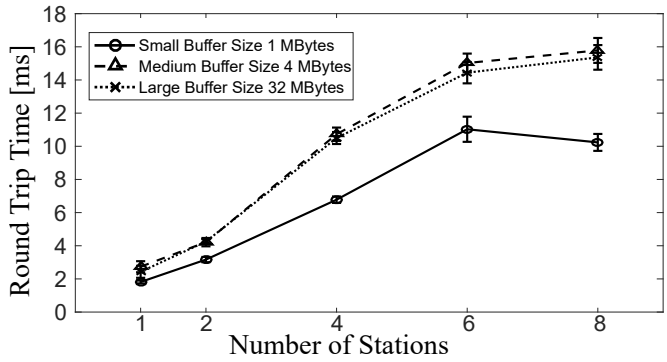

Figure 6.4: RTT Comparison for Downlink Scenario.

the AP. As a result, CSMA/CA has a direct impact on the performance of the transport protocol and the amount of buffering in the transport layer. Similar to the downlink scenario, we study the impact of the size of TCP receive and send buffers together with the number of contending stations on the achievable throughput and RTT. Additionally, we study Jain's fairness index to have a better understanding of the fairness between the TCP flows. Here, Jain's fairness index ranges between 0.125 and 1; 0.125 is the lowest value corresponding to completely unfair channel allocation whereas 1 means that the allocation scheme is fair and all the flows share the bandwidth equally.

Figure 6.5 shows that using a large TCP buffer for high contention scenarios improves the total aggregated throughput. A large buffer masks the channel changes which are more frequent in the 60 GHz band but this comes at the expense of increasing the RTT as can be seen in Figure 6.6. The RTT is almost 5 to 10 times fold compared to the downlink scenario even for a small number of stations and with small TCP buffer size.

We plot Jain's fairness index in Figure 6.7 with a confidence interval of 95% for different fairness window sizes. The fairness window refers to a moving mean that varies from 100 ms to 3 s over which fairness is measured and thus allows to compare short-term and long-term

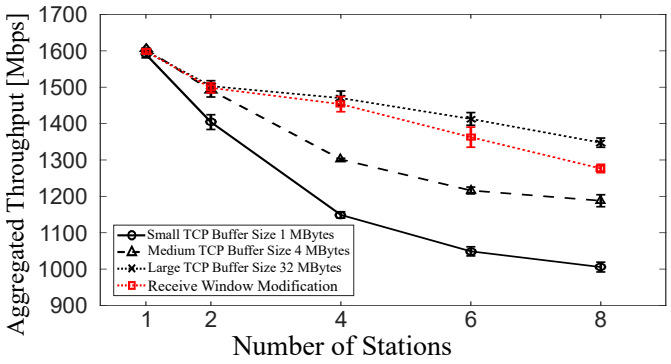

Figure 6.5: Aggregated Throughput Comparison for Uplink Scenario.

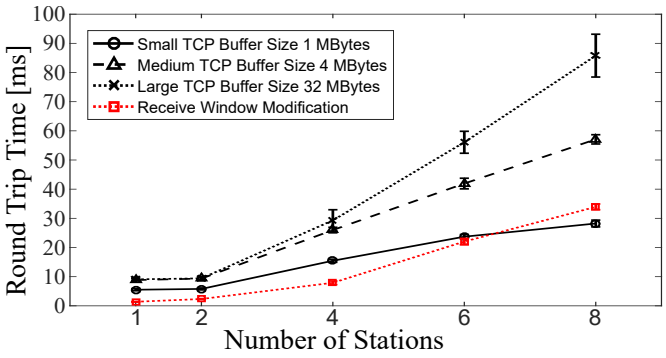

Figure 6.6: RTT Comparison for Uplink Scenario.

fairness. If Jain's index is very high for small window size, this means that all the clients are using the channel fairly and none of them is starving. On the other hand, if Jain's index becomes high only for very large window size, this implies while medium sharing is fair in the long term, over shorter time frames there is unfairness, and some of the flows may even be fully stalled.

In our experiments, for a small number of stations, Jain's index is high even for small fairness window sizes. Fairness among flows decreases when we increase both the number of active stations and the size of TCP buffer even for a fairness window of 3 s. This is true especially for the case of large buffer size. To understand the reason behind this low fairness, we look at the instantaneous throughput variation for eight stations in Figure 6.8(a). We observe that TCP is very aggressive and the instantaneous throughput can reach up to 1 Gbps. Besides, some stations have periods of zero throughput; those stations were idle for a couple of hundreds of milliseconds since the wireless medium is highly utilized. Because of that, such stations start to buffer packets until they obtain a transmission opportunity, which increases the round trip time and decreases fairness between the flows.

In summary, there is no TCP buffer size configuration that achieves acceptable throughput, RTT and fairness at the same time. To avoid excessive buffering and improve fairness, we need

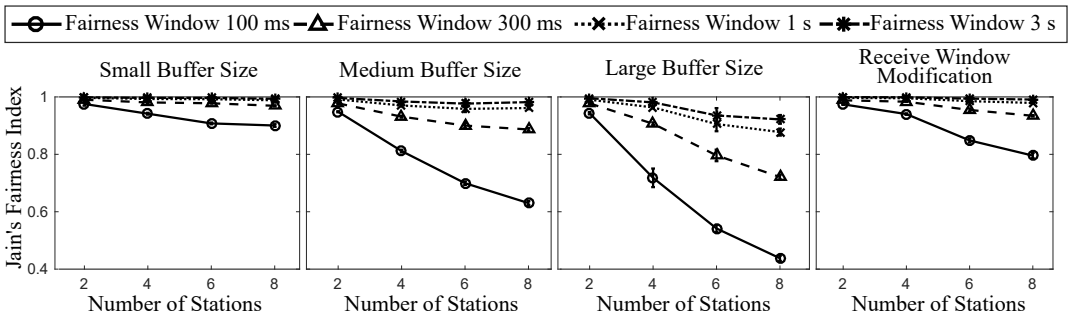

Figure 6.7: Fairness Comparison for Uplink Scenario.

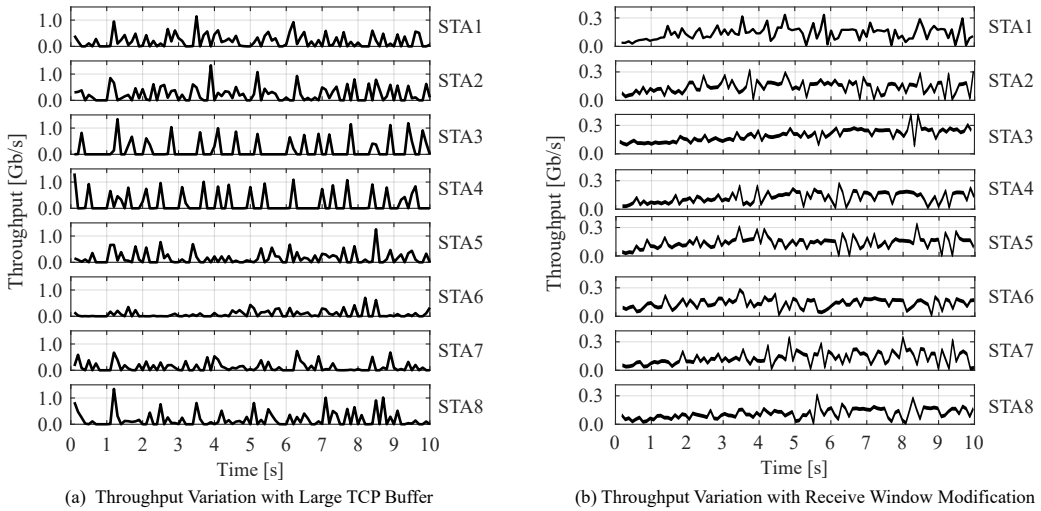

(a) Throughput Variation with Large TCP Buffer

(b) Throughput Variation with Receive Window Modification

Figure 6.8: IPERF Throughput Variation for 8 Stations with Large Buffer Size.

to adapt the rate at the sender side to be below the available capacity. For this purpose, we rely on the idea of receive size window modification to adapt the rate and thus solve the problem of high RTT while achieving high throughput. The receiver window sizing has been proposed to mitigate the buffer bloat problem in cellular networks [83]. In our setup, the optimum rate depends on the number of the contending stations since the CSMA/CA part contributes to the RTT as high number of contending stations results in an increased channel access delay per station.

We take the value of the lowest RTT from Figure 6.6, and we use it as the base RTT. For the bottleneck capacity or the desirable throughput, we take the value of the highest aggregated throughput in Figure 6.5 for large buffer size and divide it by the number of stations. We use these values to set the receive window size at the AP and we plot again the RTT, aggregated throughput, and Jain's fairness index. Figures 6.6, 6.5, and 6.7 show the performance of the receive window modification. The performance improves dramatically. For a different number of stations, we can achieve almost a throughput similar to the case of large buffer size but with much lower latency. Figure 6.8(b) shows the instantaneous throughput after modifying the receive window size. TCP

becomes less aggressive, and the throughput becomes significantly smoother over the period of during the communication. As a result, adapting the TCP receive window size prevents TCP from pushing burst of packets into the MAC layer and causing timeouts.

The only reasonable solution for a station to obtain the maximum throughput in the uplink case even if it is alone is to set the TCP send buffer large enough. However, due to the inefficiencies of the CSMA/CA protocol, a large buffer causes unfairness between flows in the case of high contention. We cannot adapt the transmit buffer at the station since it requires the station to estimate the number of other stations in the wireless network. Since mmWave networks use directional antennas, the estimation becomes hard due to deafness. However, it is possible at the AP since it is aware of all the stations that are connected and sees the flows of all the stations; thus, it can determine the right per-flow rate. This means that the AP can perform the receive window overwriting to throttle flows that are too aggressive and cause medium access inefficiencies with CSMA/CA. Another solution is for the AP to use a scheduled channel access scheme (TDMA) and assign slots for each station. In this way, a station can determine the exact flow rate during its time slot. However, the current firmware running on the 802.11ad chipset is limited to CSMA/CA operation.

6.5. Frame Aggregation

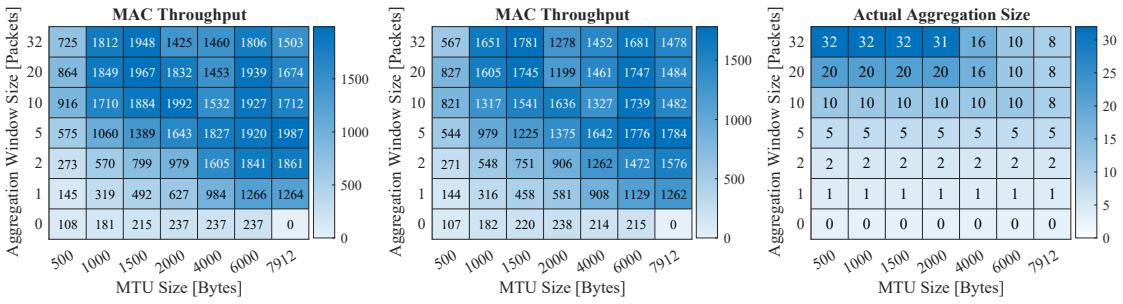

(a) MAC Throughput for Short Distance. (b) MAC Throughput for Long Distance. (c) Actual Aggregation Window Size.

Figure 6.9: A-MPDU Aggregation Performance in wil6210.

802.11ad, like its predecessors 802.11n/ac, allows for both A-MSDU and A-MPDU aggregation to improve MAC efficiency. While the A-MSDU aggregation is limited to 7935 bytes, a station using A-MPDU aggregation can aggregate either up to 64 packets or up to 262143 bytes as long as the duration of the frame does not exceed 2 ms. Although both A-MSDU and A-MPDU aggregation can be combined to further increase MAC efficiency, the chipset firmware used by our devices only supports A-MPDU aggregation with either a window size of at most 32 packets or up to 65,536 bytes.

In this section, we analyze the impact of changing the aggregation window size on the achievable MAC throughput. In addition, we emulate the behavior of A-MSDU aggregation by varying the maximum MTU size (the wil6210 driver supports MTU size of up to 7912 bytes). We use

UDP traffic in the uplink direction, with a packet size always kept smaller than the maximum MTU size to avoid IP fragmentation. We enable the `rx_large_buf` parameter in the wil6210 driver to allocate a large buffer for high MTU values.

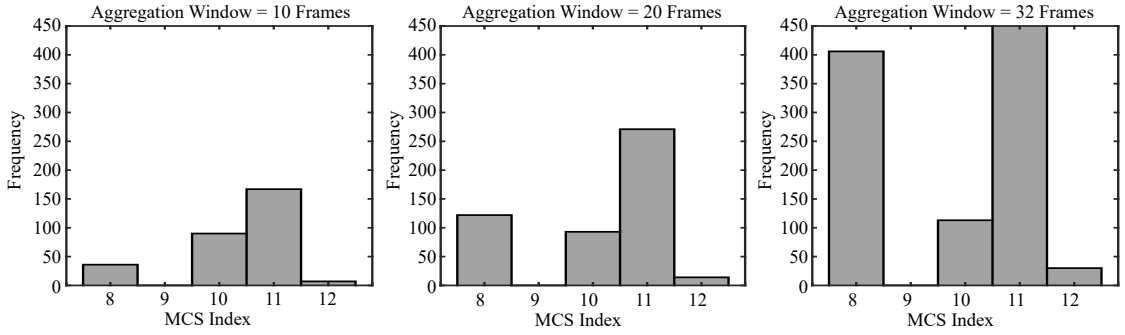

Figure 6.10: MCSs Distribution for Small Distance with MTU = 2000 Bytes.

We perform these measurements in two different settings. In the first setting, we place the laptop close to the AP (Small Distance Setting), whereas the second setting corresponds to a 9 m separation between the laptop and the AP (Long Distance Setting). Figures 6.9 (a) and (b) depict the heat map of the achievable MAC throughput for different MTU values and aggregation window sizes. Note that since the current firmware limits the number of bytes that can be aggregated, in practice the actual window size is smaller than the specified window size. Figure 6.9 (c) depicts the actual window size for all combinations of specified window size and MTU size based on wil6210 driver.

One would expect that increasing MTU and aggregation window sizes would always lead to a throughput increase. However, we observe that some combinations of MTU sizes and aggregation window sizes, in fact, degrade the performance. This is because a large aggregation window size or frame size increases the probability of having erroneous sub-frames inside an A-MPDU frame. Erroneous frames trigger the rate adaptation mechanism to lower its current MCS to cope with instantaneous channel changes. Figure 6.10 shows the distribution of MCSs for an MTU size of 2000 bytes for different aggregation window sizes in the small distance setting. We can see that for an aggregation window size of 20 and 32 frames; we have a high occurrence of both MCSs-8 and 11 compared to the case of aggregation window size of 10 frames. This indicates that the rate adaptation algorithm tries to use MCS 11 for its transmission but, due to the high probability of error with a large aggregated packet size, the rate adaptation mechanisms lowers its current MCS and utilizes MCS-8 more frequently which results in lower throughput.

Interestingly, for both short and long distances, Figures 6.9 (a) and (b) show two areas of maximum throughput: (i) large MTU sizes (6000 bytes or more) and small to medium aggregation window size (5-20) and (ii) the default MTU size (1500 bytes) and large aggregation window sizes (20-32). In contrast, the combination of maximum MTU and aggregation window size yields lower throughput.

Note that for an MTU size of 7912 bytes with A-MPDU aggregation disabled, we were unable to establish a connection to the AP. Further, setting the aggregation window size to 0 disables A-MPDU frame aggregation in the firmware. Interestingly, if we set the aggregation window size to a single packet, we achieve higher throughput compared to the case of disabled aggregation. The reason is that when we set the aggregation window size to a value equal to or larger than one packet, the firmware enables the Transmission Opportunity (TxOP) feature that allows a station to hold the wireless medium for a certain amount of time depending on the traffic class.

6.6. Delay

The analysis in the previous sections focused largely on achievable throughput. However, for many applications like web-browsing, RTT matters more than throughput in determining the user-perceived quality. Here, using the first of the setups described in Section 6.2.3 and ICMP ping packets, we measure the RTT between the laptop and the host machine attached directly to the AP. We experimented with both routers in various indoor locations. In the interest of space, we only present the results with the Nighthawk at one location, but other router-location combinations gave similar results.

Figure 6.11(a) shows the mean RTT for 100 ICMP REQUEST-REPLY pairs, with a 1 s interval between two consecutive REQUESTs, as measured at different distances between the AP and the client. The mean RTT at all distances is longer than 40 ms. Additionally, all distances show significant variance in the RTT, with RTTs as large as 100 ms and as small as 3 ms. These values, undoubtedly, are too high for a WLAN setup involving two high-speed links (802.11ad and 10G Ethernet). The RTTs to the same host, when reached through the 802.11ac interface, were always lower than 2 ms, indicating that the issue is unique to the 802.11ad path. We repeated our measurements with a single-hop ping to the router itself instead of the desktop host behind the router, removing the possibility of a delayed REPLY from the end-host. Our observations were similar confirming that delays are caused in the 802.11ad wireless path.

To understand the underlying cause, we go back to the original setup, but we now capture the packets at both the wired and wireless interfaces. Careful analysis of the packet traces shows that ICMP REPLY packets are *held-up* at the AP before being transmitted over the 802.11ad interface for almost the entire RTT measured by the ping program. Further, when we reverse the ping's direction (from the desktop connected to the AP to the laptop), the ICMP REQUEST packets suffer the same hold-up as experienced by the ICMP REPLY packets before. We further find that the amount of additional delay experienced by the ICMP packets is always equal to the time between the arrival of the ICMP REQUEST/REPLY at the AP and the *next 802.11ad's Directional Multi-Gigabit (DMG) Beaconing start*. The 802.11ad DMG Beaconing for our APs occurs every 102.4 ms and takes about 0.6 ms to complete. We make two observations: (i) Packets are held at the AP regardless of their arrival time, hence, the added delay can be as long as 103 ms (if a packet arrives right after the completion of a DMG Beaconing and is held up until the next DMG

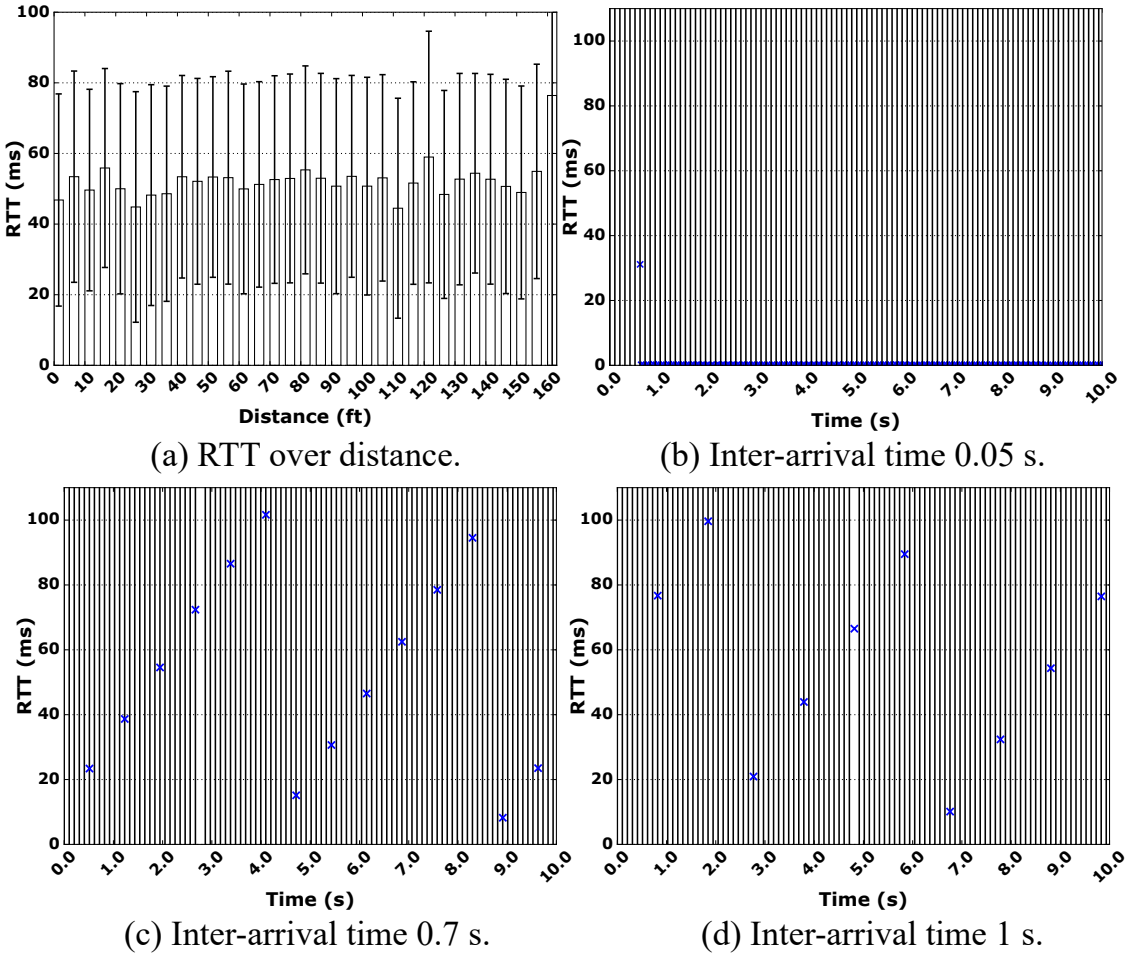

(a) RTT over distance. (b) Inter-arrival time 0.05 s.

(c) Inter-arrival time 0.7 s. (d) Inter-arrival time 1 s.

Figure 6.11: Variation of RTT with distance (a) and packet inter-arrival time (b, c, d).

Beaconing). (ii) Among all packets arriving at the router at a given amount of time before the next DMG Beaconing, only a fraction of them experiences the RTT inflation, while others are transmitted immediately. These two observations together explain the large standard deviations in Figure 6.11(a). However, they also suggest that there is another logic explaining the packet hold up and the resulting RTT inflation.

We suspect a firmware mechanism that prefers delaying packets at the AP in anticipation of higher aggregation opportunities. To explore this, we varied the interval between consecutive pings from 0.01 s to 1 s and measured the RTTs. Figures 6.11(b), 6.11(c) and 6.11(d) plot the RTT of consecutive pings sent with an inter-arrival times of 0.05, 0.7 and 1.0 s, respectively. The vertical lines mark the beginning of a DMG Beaconing. For the 0.05 s case, surprisingly there is no RTT inflation (except for the very first packet). We observe the same behavior for all inter-arrival times shorter than 0.1024 s. On the other hand, pings with inter-arrival times longer than 0.1024 s suffer inflated RTT which seems to follow a distinct pattern. For example, see Figure

6.11(c) for the 0.7 s case; the first packet has an RTT of about 20 ms, the second one about 40ms and so on. After the increasing value of RTT reaches about 100 ms, the next packet's RTT goes back to 20 ms. Note that 100 ms is very close to the DMG Beaconing period of 102.4 ms. We observe a similar pattern for the 1.0 s case in Figure 6.11(d). This zig-zag pattern of RTT inflation further explains why we observe large variations in RTT in our original experiments. Unfortunately, we were not able to explain the logic behind this pattern, which seems to be the result of a complex interplay among the packet inter-arrival time, the Beaconing period, and AP-internal counters/timers governing packet transmission.

Note that this RTT inflation does not affect the throughput of backlogged transfers, where the packet inter-arrival times are much shorter than 0.1024 s. However, this observation has a significant implication for non-backlogged traffic. We verified that this RTT inflation is experienced by TCP packets as well, suggesting that many applications can be affected. For example, page load times in web access would be artificially inflated because of this behavior, as the first RTT (for the SYN-SYNACK exchange in this case) always experiences inflation (Figure 6.11(b)).

6.7. Spatial Sharing

The order-of-magnitude shorter wavelengths in mmWave bands compared to the 2.4/5 GHz bands makes it possible to pack a very large number of antenna arrays into small a form factor, enabling highly directional transmissions. Directional communication allows for spatial reuse, i.e., the ability to establish multiple concurrent transmission links at the same time without interfering with each other. However, generating highly directional beams with minimum side lobes requires complex phased antenna arrays with a large number of antenna elements. In contrast, all current COTS 802.11ad enabled devices use a simple phased antenna array architecture with only 32 antenna elements and low-resolution phase shifters. As a result, the constructed beams are not very directional and have many side lobes. It is thus essential to analyze at which distance can we establish concurrent links using COTS devices without degrading single link performance.

To this end, we conduct different sets of experiments to cover multiple deployment scenarios. In each experiment, we have two pairs of wireless devices (referred to as "left" and "right" link) operating in the same frequency channel. Each pair consists of a single client and a single AP and each AP broadcasts a unique SSID. The AP and its corresponding client are separated with a distance of 2 m creating a strong link. Each client starts a TCP session towards its respective AP (uplink scenario). Each link is capable of achieving between 1.4 Gbps and 1.7 Gbps of throughput during standalone operation depending on the environment, the selected MCS and channel state.

To fix the index of the transmit/receive beam pattern in the TALON router, we update the firmware running on Wilocity chipset to version «4.1.0.55». Additionally, we upgrade the LEDE OS running on the TALON router to include the corresponding vendor commands that enable the user to fix both the transmit and the receive sectors for a particular communication link. The user can select the index of those sectors from a list of predefined beam patterns in a codebook file.

6.7.1. Link Separation

We study the impact of operating two links in the same frequency channel on the throughput for various values of separation distances between the two links. We start our measurement by shifting the right link in steps of 1 m towards the left link. We repeat this distance shifting three times where at each position we run three iPerf sessions, giving a total of 9 iPerf measurements per position, each lasting for 30 seconds. In this case, each link is capable of achieving around 1.4 Gbps of throughput during standalone operation.

Figure 6.12 depicts the aggregated throughput for the two links with respect to separation distance with a confidence interval of 95%. As expected, a higher separation distance results in a higher aggregated throughput for the different placements. Conversely, the smaller the separation distance, the higher the probability that the two pairs can sense each other and, as a result, start sharing the wireless medium. The high variation of the throughput is due to many factors: the proprietary rate adaptation algorithm, the contribution of reflections inside the measurement area, and the beam selection algorithm.

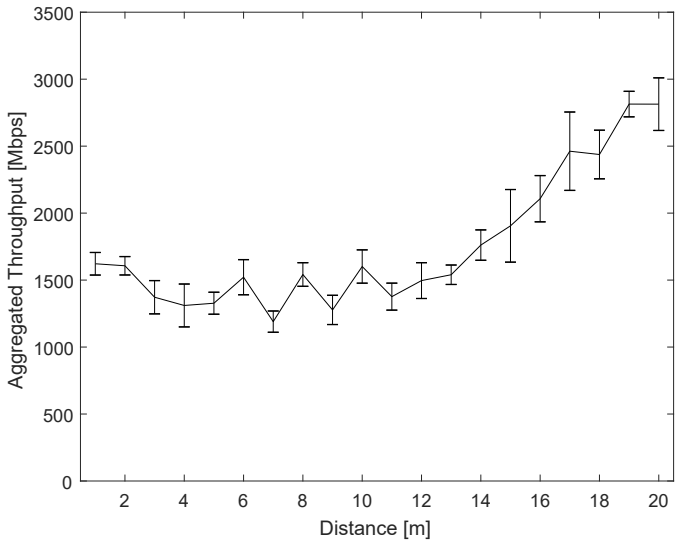

Figure 6.12: Total Aggregated Throughput with respect to Separation Distance.

When the separation distance is larger than 18 m, each client can achieve around 1.4 Gbps, the same as in standalone operation, which implies that they achieve full spatial reuse. For shorter separation distances, the two links start interfering with each other via their side lobes: Between 13 m and 19 m, the aggregated throughput decreases by 1.3 Gbps and remains around 1.5 Gbps. For shorter separation distances, the aggregated throughput varies between 1.6 Gbps and 1.2 Gbps.

Overall, we observe that COTS devices often do not allow for high spatial reuse due to the use of wide beam patterns with significant side lobes. Figure 6.13 depicts two example antenna patterns that the Talon router uses, taken from [2]. As can be seen, the Talon router uses a

Table 6.1: Throughput for Different Link Orientations

Orientation	Left Link (Mbps)	Right Link (Mbps)	Total (Mbps)
↑ ↑	1362	1027	2389
↑ ↓	954	753	1707
↑ →	694	884	1578
← →	785	696	1481
↑ ←	654	761	1415
→ →	588	787	1375
→ ←	521	666	1187

single beam pattern for receiving which has a quasi-omni shape in the azimuth plane. Besides, it frequently uses sector (20) which has a strong lobe when directly facing the peer station but also has many side lobes. Thus, the wide beam of the receiving sector and the strong side lobes in the different transmit sectors used by the Talon router explain the need for a large separation distance between parallel links to achieve a high degree of spatial reuse.

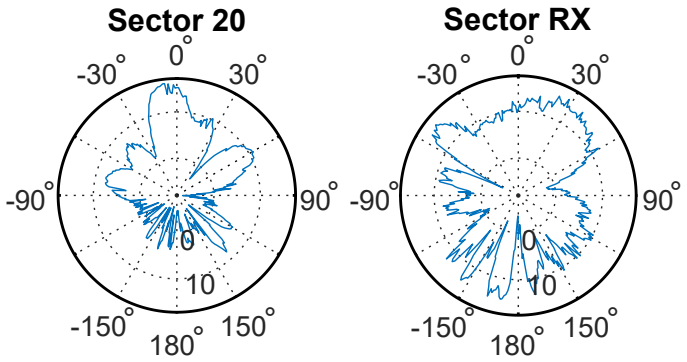

Figure 6.13: Two antenna radiation patterns in azimuth plane [2].

6.7.2. Link Orientation

In this subsection, we study the impact of link orientation on the aggregated throughput when sharing the wireless medium. The two links are separated with a distance of 15 m, and the separation between each AP and its corresponding client is 2 m. In this measurement campaign, each AP-client pair is placed either in a horizontal or vertical orientation. Table 6.1 shows the results of this set of experiments in descending order based on the total aggregated throughput. The first column contains the orientation of each link, while the remaining three columns show both the individual and aggregated throughputs.

The parallel links topology with communication flows in the same direction is the one with the highest aggregated throughput around 2.4 Gbps. The high aggregated throughput is because the clients do not sense each other as often as in the other scenarios. In addition, since clients do not face each, this makes a balanced communication link with low interference and therefore lower

probability of packet collisions. However, it can be seen that the throughput is not identical for both links, with the left link having higher throughput compared to the right link. This is due to the asymmetry in the shape of the utilized beam patterns. On the other hand, the topology of two horizontal links where the two clients face each other and the two APs are between them performs the worst in terms of aggregated throughput. This is because the two clients can overhear each other's transmissions frequently and thus defer their transmissions. Besides, each client causes high interference at the AP of the peer link.

Between these two cases lie all the remaining scenarios with different links orientations. Based on the relative orientation of the client in each link with respect to the neighboring link, the probability of deferring transmissions either increases or decreases. If the client is the one interfering with the neighboring link, it will interfere more often as it transits large frames containing multiple aggregated TCP Segments, while the AP interferes less as it only sends short TCP ACK frames. Many other factors play a significant role in determining the performance of a particular link orientation. These factors include the logic behind the rate adaptation algorithm in case of high interference and high packet collisions. Besides, the specifics of frame capture for a given chipset have a significant impact on the performance.

We conclude that the orientation of links has a substantial impact on spatial reuse, much more so than with sub-6GHz wireless networks. This is a unique feature of mmWave networks.

6.7.3. Impact of Beam Pattern Selection

The IEEE 802.11ad beam pattern training mechanism is purely based on the SNR achieved by the different beam patterns of the two devices forming the link, without taking into account any other devices with which they share the wireless medium [1]. This beam pattern selection mechanism is optimal when there is no spectrum sharing among different networks, but can lead to substantial interference when various networks share the wireless medium.

In this set of measurements, we manually select the transmit beam patterns and compare the overall network performance with respect to the case when the default beam pattern algorithm in COTS selects the transmit beam. The intuition behind this analysis is that some of the beam patterns have similar power in the desired direction but lower side lobes in the direction of devices of nearby networks. As a consequence, selecting those beam patterns reduces interference to/from neighboring networks and thus improves spatial reuse.

To study this phenomenon, we carry the following study: we set up two parallel links and iterate among all the possible beam pattern configurations in their codebooks. We do this for separation distances of 2, 5 and 15 meters, once for different the transmit beam patterns and using the default omnidirectional receive beam pattern, and once where we iterate over both transmit and receive beam patterns (which for simplicity we select to be the same, assuming that we have reciprocal channel).

This measurement is realized symmetrically, meaning that each of the diagonal items to use the same beam pattern from an interference and power point of view as they share the equal

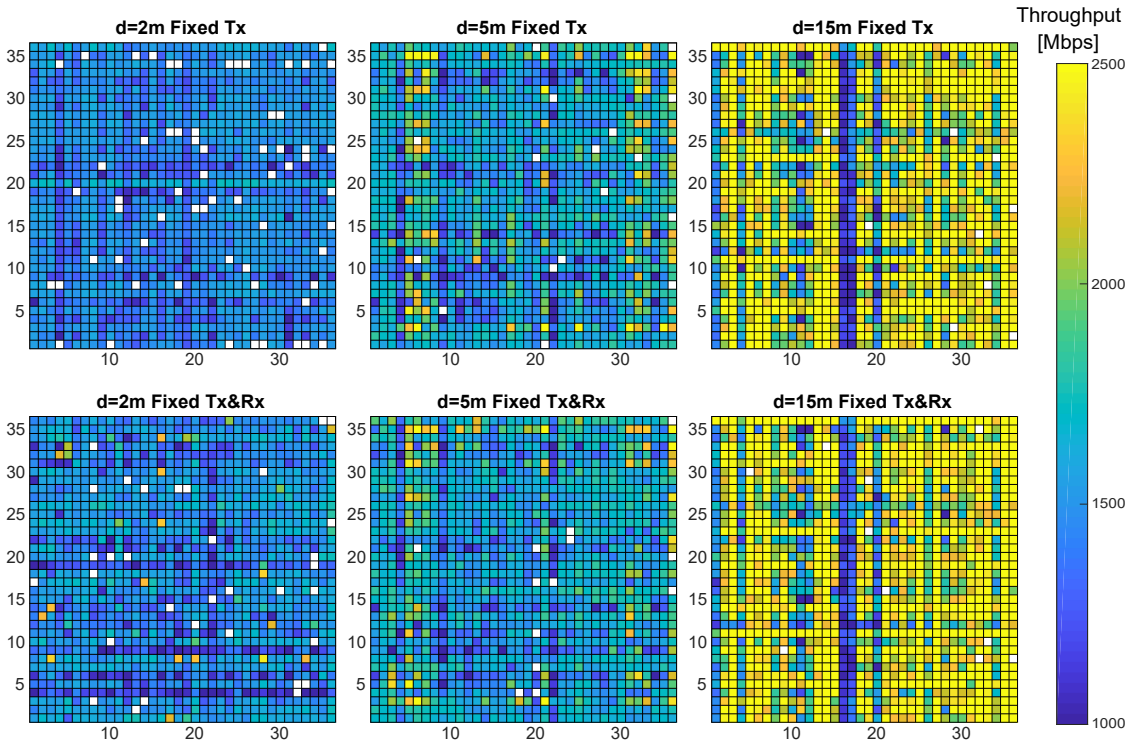

Figure 6.14: Aggregated throughput for each of the beam pattern combinations: the top row has different transmit patterns and omni-directional receive pattern, the bottom row modifies both transmit and receive beam pattern.

relative angles, simplifying the beam pattern selection measurements to reduce the number of combinations and thus the overall measurement time. Doing so, we iterate among all the pairwise possible beam pattern configurations from the codebook, using the same beam pattern in the left link's AP and the right link's client, at the same time that we fix the same beam pattern in right link's AP and the left link's client, resulting in a total of 1296 beam pattern combinations.

Figure 6.14 shows the resulting aggregated throughput for each of the beam pattern combinations while Table 6.2 compares the default throughput with respect to the maximum achieved rate among the different measurement configurations. In the figure, the columns from left to right show the distances 2 m, 5 m, and 15 m. A throughput lower than 1000 Mbps is represented with the same dark blue, and a throughput larger than 2500 Mbps is represented with the same yellow. White represents the cases where a link cannot be established.

Table 6.2: Beam pattern comparison

Distance	Default	Fixed Tx	Fixed Tx&Rx
2 m	1670 Mbps	1746 Mbps	2448 Mbps
5 m	1327 Mbps	2753 Mbps	2797 Mbps
15 m	1905 Mbps	3332 Mbps	3453 Mbps

At first glance, it can be seen that the default beam pattern selection mechanism can be improved, as in every scenario it is possible to achieve a more significant throughput when selecting beam patterns different from the default configuration. When distances are short, the gains are not very large as stations usually sense each other's transmission, given the enormous interference power at other stations independently of the beam pattern they use. However, this changes when link distances increase as we can see that selecting a different beam pattern can avoid interference and boost the throughput, even doubling it for the case of 5 m distance.

Fixing the receive beam pattern can further improve spatial sharing. In this case, not only the interference emitted towards other nodes decreases but also the interference observed through the station's receive beam pattern is reduced. This is extremely important in the case where the links are close to each other.

It can be seen that the throughput is very uniform among the beam pattern combinations for a distance if 2 m, as given the amount of interference causes the stations to carrier sense each other. Nevertheless, some gains can be achieved when fixing the transmission beam pattern. Besides, changing the omnidirectional beam pattern into a directional one gives the possibility of increasing the throughput around by 700 Mbps. At 5 m, some beam patterns provide significant gains compared to the default. For example, we can see higher throughputs in the columns corresponding to beam patterns 5, 6 and 7. These beam patterns transmit less power at the angle where the neighboring stations are located. In contrast, in columns 4, 9 or 22 the bitrate decreases due to the high power transmitted towards the neighbors. Fixing the receiving beam pattern does not significantly affect the performance of the link in this case, and the gains are similar to the ones achieved when fixing only the transmit pattern. For both cases, the maximum throughput can be doubled, providing gains of 1400 Mbps. For distance equal to 15 m, we can see many cases with substantial gains, and some combinations where the link quality is degraded. Specifically, it can be seen that beam patterns 16, 17 and 20 are suboptimal, having 1400 Mbps less aggregated throughput. Beam pattern with index equals 20 is the one used by the default configuration, showing that even though the received power in the single link case is significant, the interference created is also large. For this link distance, fixing the receiving beam pattern can increase the link throughput by another 100 Mbps.

It is also important to note that, given the asymmetric beam patterns, the behavior of the beam pattern selection also is not symmetric, so a pairwise switch of the beam patterns leads to different rates. Even though we have bi-directional communication in the link, the acknowledgments require less time than the payload, making the interference-free operation more critical in the transmitter node than in the receiving one.

6.8. Beam Radiation Patterns

mmWave devices rely on electronically steerable phased antenna arrays to establish directional links for communication; thus it is of high importance to characterize the directivity of

those antenna arrays for proper network operation. The directivity of an antenna array, in general, relies on the operating frequency and the physical spacing between its elements. Since the spacing between the antenna elements cannot be modified after antenna array production, the antenna array's directivity depends solely on the operating frequency. For wireless devices operating in the microwave band with megahertz of spacing between neighboring channels, the changes in the directivity of an antenna array are almost negligible compared to mmWave devices with a gigahertz of frequency spacing. The IEEE 802.11ad standard channelizes the 60 GHz band into four orthogonal channels each of which has a bandwidth of 2.16 GHz. A large amount of frequency spacing between the channels drastically changes the radiation patterns drastically for those antenna arrays. Additionally, current COTS devices [59–61] rely on a predefined codebook whose entries represent the weights applied to the complex signal do not change based on the operating frequency channel. As a result, taking into account the operating frequency is crucial for the optimal configuration of the device since a sector which is optimal for a specific frequency channel might not be the best choice for communication in different frequency channel.

In the following set of experiments, we measure the radiation patterns of the phased antenna arrays in the TALON Router and the MikroTik RouterBoard. We conduct our measurements using a rotating table that holds the device in study. The device is set to work as AP and thus periodically transmits directional Beacons across each of its predefined sectors during each beacon interval (BI). Besides, the length of the BI is set to 102.4 ms. For the reception, a Vubiq 60 GHz receiver [84] is used, and it is equipped with a highly directional horn antenna of 7°-beamwidth. The Vubiq, in turn, is connected to an oscilloscope that captures and stores the raw signal of all the detected Beacons to be post-processed in MATLAB. mmWave RF absorbers are used to ensure that all the reflections coming from the surrounding environment are suppressed. The device is set to rotate in the azimuth plane from -160° to 158° with a step of 2°. For each rotation angle, we capture 20 BIs and average the power of all the received Beacons within those BIs. We repeat those measurements for each channel supported by the previous devices. At the time of the writing, all the COTS devices operating in the 60 GHz band support three frequency channels only (Ch.1: 58.32, Ch.2: 60.48, and Ch.3: 62.64 GHz).

Figure 6.15 shows the comparison of the radiation patterns for both the TALON Router and the MikroTik Router over three frequency channels. Since those devices use different transmit power for each frequency channel, we normalize the power of the captured Beacons for a fair comparison. As expected, the radiation pattern changes for both routers with respect to the selected channel. For a specific frequency channel, the radiation patterns have directional lobes at certain angles, but those directional lobes either become side-lobes or shift to be directional at a different angle.

Figure 6.16 shows the difference in antenna array directivity when the best sector for communication in Ch. 1 is utilized in the other frequency channels. For each steering angle in the lowest channel (Ch. 1), we obtain the sector with the highest directivity. Then, we measure the changes in directivity when we use the same sector in the two other frequency channels.

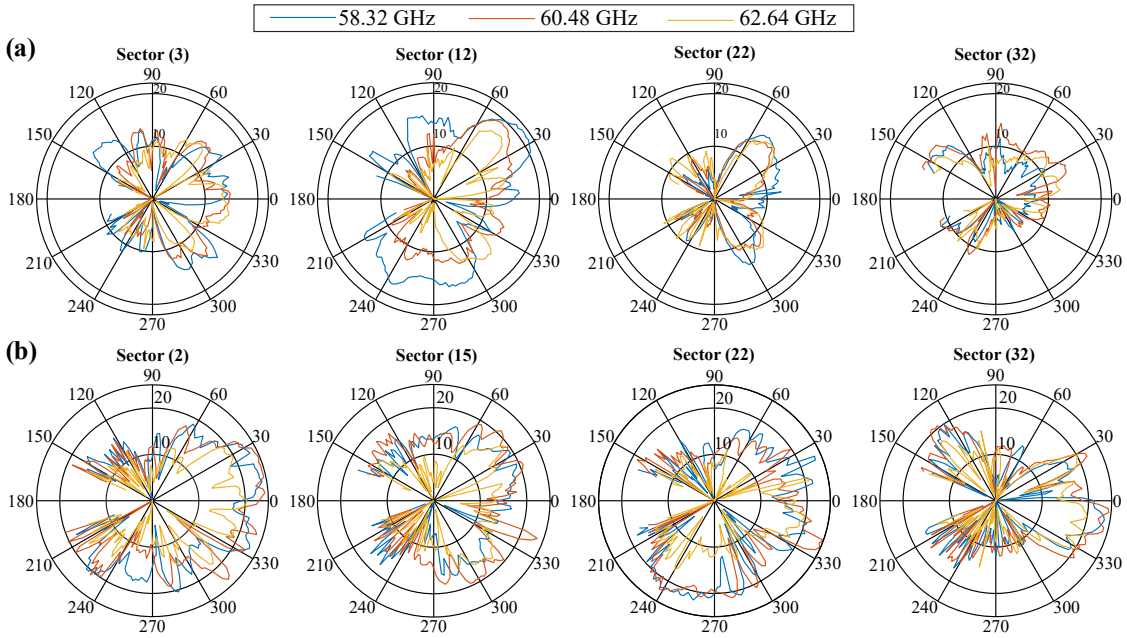

Figure 6.15: Beam Patterns Comparison for Different Frequency Channels (a) Beam Patterns for the TALON Router (b) Beam Patterns for MikroTik Router.

As observed in the Figure, when the separation between the channels increases the difference in directivity increases. It can reach up to 15 dB which might break the communication when using high MCS. This demonstrates the importance of characterizing the directivity of practical antenna arrays before network deployment and operation. To overcome the previous problem, COTS mmWave devices should embed adaptive codebooks that modify its entries based on the operating frequency. This would ensure that mmWave devices have either identical beam patterns or beam patterns with minimum changes across the frequency channels.

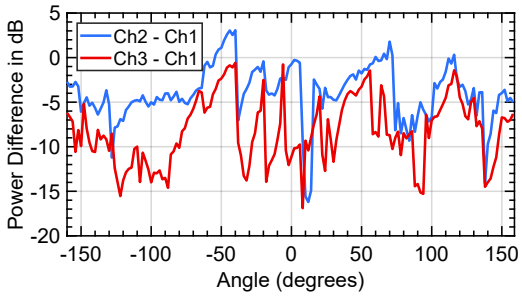

Figure 6.16: Antenna Array Directivity Difference with Respect to the Lowest Channel

6.9. Discussion

After analyzing the performance of the COTS 802.11ad devices, we get some insights into the existing problems in those devices, so we discuss some of the possible solutions.

Phased Antenna Array. Both the Talon and the NightHawk routers use a phased antenna array of 32 antenna elements that implement analog beamforming. In [2], the authors measured the 2D radiation patterns of the sectors defined in the codebook of the Talon router in an anechoic chamber. All the sectors have irregular and non-directional patterns. However, in theory, a linear phased antenna array of 32 elements can generate very narrow beams with a beamwidth of 7 degrees assuming the elements are spaced by a half lambda. However, the antenna array used in those devices arrange the elements in an asymmetric way where some of the elements are located on the surface of the antenna module while others are located on the back and the rest are located on the sides of the board. This arrangement results in irregular beam patterns. The wide irregular beams are suitable for robust communication during user mobility within an indoor environment. This is because the device does not have to continuously sweep across all its sectors to avoid link interruption thus resulting in a smooth experience to the end user. Besides, since those devices are limited for indoor usage with few numbers of devices, it makes sense to use wide beams for mobility and robustness. Moreover, those devices work well enough with those beam patterns as they can achieve Gbps throughput across the room except for the blockage case. However, as the network becomes dense and we try to maximize the number of clients, interference increases significantly. Thus, achieving spatial reuse becomes hard. Another problem that exacerbates in the mmWave band is that antenna arrays require precise calibration during the manufacturing process so that they can operate as expected. Since this is a complex process that mandates many measurements to characterize the antenna array and its elements, many vendors may eliminate this step as it would increase the cost of end devices. As a result, this may make end devices of the same manufacturer have different and unpredicted performance [85]. Thus one can anticipate that the low-end commercial devices in the future will exhibit similar performance to those routers. On the other hand, future devices will most likely integrate multiple phased antenna arrays to overcome the blockage case [19].

Beamforming. As mentioned earlier, the current 802.11ad off-the-shelf devices implement analog beamforming in which a single RF chain is connected to all the antenna elements. To steer the antenna into a specific direction in space, we modify the excitation, i.e., the phase and the amplitude of each element. The current firmware provides this capability. It allows the end user to configure the beam patterns and modify the weighting network that controls the antenna elements, thus generating relatively narrow beam patterns as opposed to the default beams defined in the codebook of those devices [2]. This allows us to reduce interference which is vital for dense deployment and spatial reuse. On the other hand, this is a problem for beamforming training and user tracking with narrow beams. Another possibility to generate narrow beams would be by using different beamforming architectures such as a hybrid beamforming architecture [86, 87].

The hybrid beamforming provides similar performance to digital beamforming but with lower complexity and less power consumption.

Operating Frequency. The IEEE 802.11ad standard is envisioned to be used for ultra high-speed point-to-point and mesh networks replacing wired connections when cable installation is prohibitively expensive. In such kind of networks, proper network planning before network deployment is mandatory since the directivity of practical phased antenna arrays changes significantly with respect to the operating frequency compared to theoretical results. Thus, mmWave networks that have little interference operating at a specific frequency channel might interfere at another frequency due to the changes in antenna directivity. Therefore, COTS mmWave devices should adopt multi-frequency codebook designs that update the weights of the antenna array elements based on the operating frequency to compensate for possible changes in the directivity.

COTS Firmware. The current firmware running on Qualcomm 11ad chipsets exposes some Netlink commands to the user space to control the chipset and query information regarding network operation. However, it neither provides a capability to modify the rate control mechanism nor change the logic of sector selection after the beamforming training phase. This is important as the algorithmic part of those mechanisms change based on the deployment scenario. For example, deploying an 802.11ad network in a busy place such as a mall with dense and high user mobility requires fast beam tracking and rate adaptation algorithms to ensure a smooth experience and avoid frequent link blockage. Whereas, deploying an 802.11ad network in an office environment where human mobility is not a concern, allows the usage of low-complex algorithms.

Channel Access Schemes. We showed in Section 6.4 the impact due to the interaction between between CSMA/CA and TCP for different TCP buffer sizes on network throughput and delay when using CSMA/CA with a variable number of stations. The 802.11ad standard introduces hybrid medium access which compromises both contention-based access (CSMA/CA) and time-based access (TDMA). Both the idle time and the contention in CSMA/CA have a high impact on TCP performance and fairness. In addition to that, we have deafness issues which did not exist in omnidirectional communications. On the other hand, a time-based channel access scheme improves fairness and network efficiency and guarantees QoS but requires a complex resource scheduling algorithm. The current firmware only supports CSMA/CA channel access and guarantees QoS through the Enhanced DCF channel access (EDCA).

AP Deployment. In our study, we focus on transport and MAC layer aspects for a network with a single AP and multiple devices located at the same height within proximity of each other. Although this setup allows us to isolate the effect of deafness, AP orientation, and placement within the room, it does not reflect a realistic deployment. A realistic deployment would incorporate multiple APs distributed across a large area where each AP is mounted on the ceiling and pointing towards a certain direction. The authors in [2,63] measured the 3D beam patterns for the Talon routers, and they found that some of the beams were directional in the 3D space at a certain azimuth and elevation angles. The authors in [88] perform an extensive measurement campaign to understand how to deploy COTS mmWave AP to achieve the best coverage and throughput.

These two studies encourage further research to understand how to place many APs efficiently to achieve high spatial reuse, optimize AP and client association, and reduce interference among clients.

Beam Pattern Selection. The current beam pattern selection mechanism in COTS devices work efficiently when only a single network is utilizing the wireless medium. However, when a second network operates on the same frequency channel, the overall performance decreases dramatically due to the non-ideal selection of beam patterns that cause high interference to the neighboring networks. In our measurements, we showed the possibility of improving the overall performance when taking into account the existence of nearby networks during the selection of the transmit beam patterns. Furthermore using directional reception pattern instead of the default quasi-omni pattern also helps in improving network performance and decreases interference coming from nearby networks. Using a directional reception pattern has a clear impact when the interfering link is close to the user.

6.10. Conclusions

We study protocol aspects of the MAC layer and TCP in practical mmWave networks. In contrast to earlier work, we consider a complete IEEE 802.11ad network with one AP and up to eight stations. We perform our analysis on a practical testbed using commercial off-the-shelf hardware. Our results show that dynamic TCP buffer sizing is crucial in the uplink because the multi-gigabit-per-second speeds at the physical layer and the directional antennas exacerbate the inefficiencies of CSMA/CA. We also investigate the impact and limitations of frame aggregation. Regarding the delay, we find that the regular beacon transmission process of IEEE 802.11ad APs can inflate the round trip time, thus degrading performance. Further, we show the impact of link orientation and beam pattern selection on spatial sharing and show that interference aware beam pattern selection can double overall throughput compared to default IEEE 802.11ad operation. Finally, we showed the importance of supporting adaptive codebook mechanism in COTS devices to minimize beam pattern changes across frequency channels.

Chapter 7

Millimeter wave Sensing

7.1. Introduction

Achieving pervasive connectivity in the mmWave with receivers, transmitters and reflectors moving, has become the topic of substantial research over the last few years [89–92]. The standard approach is to react, upon blockage detection, and search for a strong enough NLOS reflection over the (rather large) space of beam patterns (beam realignment or training), or for another AP (handover). To alleviate the overhead incurred in such tedious approach (see e.g., [89]), several authors have studied preemptive alternatives [89–92], discussed in §7.2. Interestingly, all this literature exploits the concept of *context awareness*, such as out-of-band information [90], pose information [91], predictable traffic patterns [92] or device localization [93, 94]. Evidently, the ability to locate devices precisely enables interesting optimizations, e.g., by executing an immediate handover upon blockage detection without the need for invoking a beam training or AP search procedure. This is indeed the case of [93, 94], which leverage AoA spectrum information as a basis for device localization.

In contrast, in this chapter, we address the dual problem that consists of *obstacle localization*. This approach has a few advantages, namely, (i) it does not require user device collaboration and (ii) *environment awareness* can be used to predict (and then avoid) link obstructions, as opposed to reacting once a blockage is detected. In this sense, our approach is similar in nature to `RadMAC` [95], which proposes to augment mmWave APs with radar equipment to carry out the object detection legwork. Conversely, our approach does not require additional hardware nor specialized mmWave equipment or modifications to standard MAC/PHY protocols. On the contrary, our approach relies exclusively on readily available commodity APs with no involvement of end devices.

Similarly to prior literature [93, 96], we consider a typical indoor deployment (e.g., an office or a living room) comprised of multiple mmWave APs to provide coverage. In this context, we design and build `BeamScanner`, a system that (i) exploits collaboration across spatially distributed APs, (ii) does not involve end users and (iii) is implementable onto off-the-shelf

standard-compliant hardware, to *scan and map obstacles in the environment* with very high accuracy. The design of `BeamScanner` pivots around the following key features.

First, we implement an user-accessible interface to collect fine-grained Channel State Information (CSI) from commodity chipsets. Second, we construct a method to generate custom-built beams that allows us to scan the environment with controllable (and narrow) beam patterns. This helps to mitigate the additive noise of multiple and diverse signal paths of highly imperfect beam patterns available in consumer-grade equipment. Third, we treat the collected data with a custom designed filter to remove spurious data from imperfect measurements. As an example, commodity devices report channel information *after* AGC and thus we need to compensate for the lack of an absolute reference value. Fourth, we design a simple multi-AP collaborative beam scanning protocol that can be seamlessly integrated into the upcoming IEEE 802.11ay standard and, with mild overhead, into the legacy IEEE 802.11ad. The combination of these four schemes enables us to scan accurately indoor scenarios, crawling fine-grained information *even from weak reflections coming from human obstacles*. Finally, we then process this data using a simple classifier, previously trained with labeled fingerprints, to estimate the presence of objects and their location[1]. In this way, even small non-linear effects from weak reflections contribute to a more accurate localization without compromising complexity. In summary, our contributions are the following:

- We illustrate how weak reflectors such as human obstacles cause a footprint that can be inferred with commodity mmWave receivers.
- We develop a user-friendly interface between the firmware of a commodity AP and user-space to trigger and collect low-level channel state information.
- We design a method to detect the footprint of obstacles in indoor environments. Our method is based on the coordination of distributed APs.
- We evaluate the obstacle localization accuracy with a human obstacle in a typical indoor room with four commodity APs.

The remainder of this chapter is structured as follows. Section 7.2 provides a summary of relevant literature. Section 7.3 lays the foundations of our approach, presenting some motivating toy experiments. In Section 7.4, we introduce the design of `BeamScanner` and then present an experimental evaluation in §7.5. Section refsec:conclusions provides concluding remarks.

7.2. Related Work

7.2.1. mmWave Beam Steering Methods

The cumbersome process of a beam (re-)steering and/or AP hand-over, once a link blockage is detected, incurs in high overhead. This has motivated substantial work on preemptive meth-

[1] We note that triangulation methods are not possible in this case because discerning which reflections come from a specific obstacle is inherently hard.

ods that exploit context information to proactively guide the beam search procedure [89–92]. BeamSpy [89] exploits Channel Impulse Response (CIR) information of the beam in use to build a full *path skeleton* model that allows quality prediction across all possible beams. However, this approach is functional in quasi-stationary scenarios only, and it requires knowledge of the radiation patterns of the AP. The latter is a rather important issue because the imperfect beam patterns of commodity APs [93] makes per-device calibration a rather limiting requirement. MUST [90] uses available out-of-band information from co-located WiFi chipsets to (roughly) estimate good 60-GHz beams and, in case of LOS blockage, quickly switch over to WiFi communication. However, out-of-band channel estimation is rather inaccurate, and dual-connectivity is battery-involved. Alternatively, Pia [91] exploits pose information (3D position and polar/azimuth orientation) of clients connected to a multi-AP setting to predict the best AP at each given time. Finally, MAP [92] is a cross-layer solution that exploits particularities in traffic burst patterns of video streaming applications to opportunistically select an appropriate AP. However, this approach is hard (if possible at all) to generalize to other types of applications.

In contrast to all this work, BeamScanner does not require the end device collaboration nor does it rely on inaccurate side information.

7.2.2. mmWave-based Localization Methods

The quasi-optical nature of mmWave systems make them inherently suited for device localization [97–100]. Assuming idealized hardware, [97] devises a mechanism to estimate not only the location of a device but also its orientation. A multiple-input-single-output (MISO) Orthoghonal Frequency-Division Multiplexing (OFDM) mmWave system for localization is used in [98]. Consider real, but non-commodity, hardware, triangulation [99] and ranging with multilateration [100] are also studied. Finally, angle of arrival (AoA) spectrum information available in commodity equipment is considered in more recent work [93, 94].

However, as we mentioned in §7.1, our work addresses the dual problem that consists on *obstacle localization*. The main advantage of these type of systems is that they do not require collaboration for the user devices and provides a system controller with context information that can be used for, e.g., proactive beam steering methods. This is indeed the case of RadMAC [95], a radar-enhanced mmWave transceiver that exploits object-tracking information to predict link blockages during which traffic is re-steered throughout alternative non-LOS links (or different APs, potentially). In contrast, our approach does not require specialized equipment like RadMAC does and we exploit collaboration across APs installed in an indoor location to construct an overall view of the environment.

7.2.3. Sub-6 GHz Object Detection Mechanisms

In addition to device localization, which has been addressed (via triangulation or multilateration) in substantial literature, e.g., [101, 102], object detection and tracking has also been con-

sidered in lower-frequency technologies. However, the granularity of the detection information is overall rather rough using this type of mechanisms. The work of [103] was the first example of the utilization of WiFi signals for human presence detection. Later improvements led a few works to detect humans even behind walls, e.g., [104, 105]. However, the low sparsity of such type of technologies imposes severe limits to accurate detection *and* positioning of obstacles such as humans.

7.3. Preliminaries

7.3.1. Notation

We use conventional notation. We use bold $\mathbf{a}$ to denote a vector, upper-case bold $\mathbf{A}$ to denote matrices, and caligraphic $\mathcal{A}$ to denote a set. We let $\mathbb{R}$ and $\mathbb{Z}$ denote the set of real and integer numbers. We use $\mathbb{R}_+$, $\mathbb{R}^n$, and $\mathbb{R}^{n\times m}$ to represent the sets of non-negative real numbers, n-dimensional real vectors, and $m \times n$ real matrices, respectively. $\|\mathbf{a}\|$ is the ℓ^2 norm of $\mathbf{a}$, $\|\mathbf{A}\|_F$, $|\mathbf{A}|$, $\mathbf{A}^T$, $\mathbf{A}^H$ and $\mathbf{A}^{-1}$ are, respectively, the Frobenius norm, determinant, transpose, Hermitian, and inverse of $\mathbf{A}$.

7.3.2. Weak reflections provide valuable information

There exists certain consensus in the literature to filter out high-order reflections to understand the environment by exploiting mmWave channel information [89, 93, 96]. We argue in this work, however, that even weak reflections can provide rich contextual information.

Let us illustrate this with a toy experiment. We first set up a commodity TP-Link TALON AD7200 IEEE 802.11ad-compliant AP periodically transmitting standard DMG beacons [106].[2] This is depicted in the bottom-left corner of Figure. 7.1. We also place a highly-directional horn antenna at 0° azimuth and 2 meters apart and connect it to a down-converter to feed base-band information to a conventional Keysight oscilloscope. As shown by the figure, we force[3] the AP to use a beam pattern at an azimuth of 45° (Sector 6 shown in [2]). Our objective is to use the oscilloscope to measure the impact of potential NLOS reflections of the beacons emitted by the AP onto different obstacles. The whole setup is illustrated in Figure. 7.1.

We first measure the power of the signal received by the horn antenna with no LoS obstacles. This is shown in black color in Figure. 7.2 when capturing a single beacon, and gives us a baseline to compare with. As expected, the signal power is very low in this case as the signal received by the horn antenna is mostly coming from spurious lobes due to the (imperfect) beam pattern of our AP. Then, we place different obstacles in LoS, at 45° azimuth of both transmitter and horn

[2]The TP-Link TALON AD7200 uses the QCA9008-SBD1 module with the QCA9500 chipset from Qualcomm, embeds a 32-element phased-antenna array and supports single-carrier data rates up to 4.6 Gbps.

[3]The legacy firmware provides the capability to fix any beam pattern and to choose the order and number of beams during beamforming training and beaconing. For this work, we have implemented a custom set of driver modifications to make use of the firmware API.

antenna, as depicted in Figure. 7.1, and repeat the experiment. Sequentially, we place a laptop with some metal casing (in blue) and a human body (in red). Not surprisingly, the reflections onto the laptop have substantial strength. However, it is interesting to observe that, albeit weaker than the laptop's, human obstacles (red line) also cause distinguishable reflections. This motivates us to seek the footprint of (e.g., human) obstacles in the environment across higher-order (weak) reflections, as opposed to filtering these out.

Figure 7.1: Toy Example Setup.

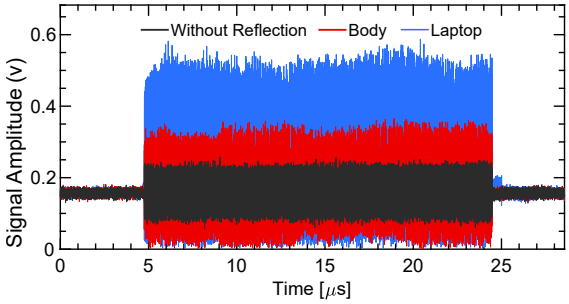

Figure 7.2: Received Signal Power for different Cases.

7.3.3. Analysis of the perturbations caused by obstacles

Understanding amplitude and phase changes due to (weak) reflectors in the environment is inherently hard. Consider a scenario with M homogeneous APs $\{s_1, s_2, \ldots, s_M\}$ providing indoor coverage with identical number of antenna elements equal to N and available beam patterns K ($N = 32$ in our testbed comprised of $M = 4$ TP-Link TALON AD7200 APs). Let us focus on the link that can been created between either pair of APs. Due to its short wavelength, the wireless channel between any transmitter s_i and any receiver s_j, $i \neq j$, has a quasi-optical behavior and so, similarly to [107], we model its channel using a geometric channel model with $l \in \mathcal{L}$ paths:

$$\mathbf{H}^{(i,j)} = \sum_{l \in \mathcal{L}} \alpha_l \mathbf{a}_{\text{RX}}^{(j)}(\theta_l) \mathbf{a}_{\text{TX}}^{(i)}(\phi_l)^H \tag{7.1}$$

where α_l is the complex gain on the l^{th} sub-path, $\mathbf{a}_{\text{RX}}^{(j)}(\cdot)$ and $\mathbf{a}_{\text{TX}}^{(i)}(\cdot)$ is the planar array steering vector of the receiver and the transmitter device, respectively, and θ_l and ϕ_l are the l^{th} path's azimuth angles of arrival of the receiver and angles of departure of the transmitter.

When fixing a beam pattern for transmission, like we did in §7.3.2, we get a reception channel described by

$$\mathbf{h}_{\text{RX}}^{(i,j)} = \sum_{l \in \mathcal{L}} \alpha_l \mathbf{a}_{\text{RX}}^{(j)}(\theta_l) \mathbf{a}_{\text{TX}}^{(i)}(\phi_l)^H \mathbf{p} \tag{7.2}$$

where $\mathbf{p}$ is a vector containing the weights applied to the transmitter's steering vector $\mathbf{a}_{\text{TX}}^{(i)}$. The above can be conveniently rewritten as

$$\mathbf{h}_{\text{RX}}^{(i,j)} = \sum_{l \in \mathcal{L}} \beta_l \mathbf{a}_{\text{RX}}^{(j)}(\theta_l) \tag{7.3}$$

by defining $\beta_l = \alpha_l \mathbf{a}_{\text{TX}}^{(i)}(\phi_l)^H \mathbf{p}$.

Now let us consider the same scenario with an object that has a partially reflective surface. In such a case, we can make the following two observations: (i) the object will not modify the already existing paths unless they are blocked by the obstacle; and (ii) when it does block paths, the object will likely create a new first order reflection. Both effects can be seen as the creation of a new path (and additional higher-order reflection paths that are irrelevant to the channel contribution). In case of blockage, it can be seen as a new path that cancels destructively the path (or paths) blocked by the obstacle. As a result, we describe the channel with the perturbation (object) as

$$\hat{\mathbf{h}}_{\text{RX}}^{(i,j)} = \sum_{l \in \hat{\mathcal{L}}} \alpha_l \mathbf{a}_{\text{RX}}^{(j)}(\theta_l) \mathbf{a}_{\text{TX}}^{(i)}(\phi_l)^H \mathbf{p},$$

with $\mathcal{L} \subset \hat{\mathcal{L}}$ (to include the newly created path(s)). Like before, this can be rewritten as

$$\hat{\mathbf{h}}_{\text{RX}}^{(i,j)} = \sum_{l \in \hat{\mathcal{L}}} \beta_l \mathbf{a}_{\text{RX}}^{(j)}(\theta_l)$$

Since $\mathbf{p}$ is designed to be directional, we expect the effect of the complex coefficients β_l to be amplified if $\mathbf{p}$ is steered towards the path direction ϕ_l.

Given a specific beam pattern $\mathbf{p}_n$, which leads to channel $\mathbf{h}_{\text{RX},n}^{(i,j)}$, our objective is to understand how the channel changes when an obstacle appears, that is,

$$\bar{\mathbf{h}}_{\text{RX},n}^{(i,j)} = \hat{\mathbf{h}}_{\text{RX},n}^{(i,j)} - \mathbf{h}_{\text{RX},n}^{(i,j)} = \sum_{l \in \mathcal{L} \cap \hat{\mathcal{L}}} \beta_{l,n} \mathbf{a}_{\text{RX}}^{(j)}(\theta_l). \tag{7.4}$$

This can be seen as *the virtual channel* that comprises only the effects of the perturbation.

7.4. System Design

As we mentioned in Section 7.1, we consider a typical indoor deployment with multiple mmWave APs distributed in a room that ensure coverage. Assuming idealized sector shapes, the IEEE 802.11ad standard specifies clients to periodically transmit sector sweep messages of each available to cover all possible directions. However, the shape of the underlying beam patterns

generated by commodity hardware is highly imperfect [2, 5] and thus is hard to map sectors and angular directions. This has also been noted in prior work [89, 93], which declare such imperfections as noise and resort to filter them out to create a sparse representation of the channel. As remarked earlier, not only do we not filter these higher-order effects but seek them instead to build an accurate map of the environment.

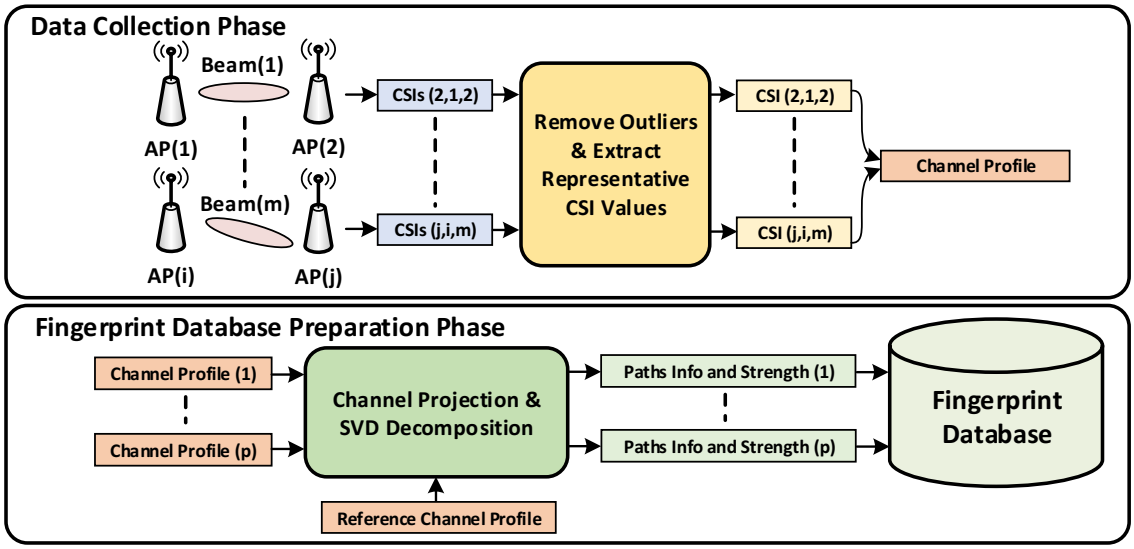

Figure 7.3: `BeamScanner` System Design.

The overall system is summarized in Figure. 7.3. `BeamScanner` is essentially divided into two processes. On the one hand, a data collection phase is periodically triggered. In this stage, the goal is to build a complete profile of the room by aggregating channel samples $\bar{\mathbf{h}}_{\mathrm{RX}}$ with different beam patterns i and APs j (illustrated at the left-most part of the data collection phase in the figure). In order to obtain an accurate channel snapshot, the CSI measurement of each s_i-s_j mmWave link and beam configuration is repeated attractively. This yields a sequence $\langle g_k \rangle$ of channel measures that are processed by a filter purposely designed to remove spurious data and obtain a single channel sample g. The filtered channel samples are then aggregated to construct $\mathcal{S}$.

On the other hand, each channel sample is compared to a reference profile (a room with no obstacles) to obtain $\bar{\mathbf{h}}_{\mathrm{RX},n}$, i.e., a virtual channel representation containing only the perturbations of the obstacle (if any). As we explain later, this entails some challenge as commodity hardware reports channel samples *after* AGC (Adaptive Gain Control) and so measurements are dynamically normalized. The set of all virtual channel samples builds a *fingerpring* matrix that we pass through a simple classifier to infer the presence of obstacles and their location. In the sequel, we detail the design of each of the components shown in Figure. 7.3.

7.4.1. Collection of CSI measurements

The first task of `BeamScanner` is to collect channel samples to construct $\mathcal{S}$. In our case, we use low-level channel state information (CSI), which is the key to infer the footprint of weak reflectors. To this aim, we first note that the latest version of Qualcomm's QCA9500 (commodity) chipset can report the strongest channel impulse response (CIR) across antenna elements in the array. In short, the firmware reports the amplitude and phase across each antenna element after each measurement procedure. To the best of our knowledge, such low-level information provided by commodity devices has not been exploited in the literature so far. Evidently, this data is a much richer source of information regarding the environment, giving `BeamScanner` an advantage. To trigger and access CSI samples, we implement the corresponding vendor commands in the driver (`wil6210`) and develop an interface with the `iw` user-space tool.

Now, to understand how CSI samples are measured, we employ the same experimental testbed introduced in §7.3.2. We now reverse the role of the TALON device (used before as an AP) to act as a legacy client instead and associate it to a second device, co-located with the horn antenna, which acts as the AP. In this case, we select a beam pattern with a strong lobe at 0° azimuth, that is, a strong LoS beam, and feed the raw I/Q samples captured by the oscilloscope into an IEEE 802.11ad decoder to sniff the transmitted frames, including all PHY-related headers.[4] In this case, we instrument the AP to trigger CIR reports periodically and study the set of frames exchanged between AP and client. Figure. 7.4 (top) depicts the signal strength of the frames exchange. Once decoded, it becomes clear that the firmware is using the Beam Refinement Protocol (BRP) to measure CSI [106], a protocol that is originally specified by the IEEE 802.11ad standard to refine the transmit/receive beam patterns *after* a Sector Level Sweep (SLS) phase. However, how the receiver changes the weights of the antenna elements to select an appropriate receive pattern and ultimately how to estimate CSI is proprietary and not publicly available.

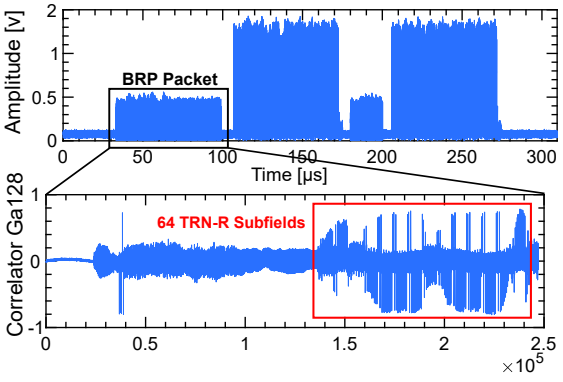

Figure 7.4: CSI Measurements Signaling.

In this case, the standard BRP protocol is triggered by the firmware when a CIR estimate is

[4] We have used the software 81199A Wideband Waveform Center.

requested. The first and third frames of the exchange shown in Figure. 7.4 (top) corresponds to BRP requests from the initiator (the client) to train the receive beam of the responder (the AP). The first frame has the flag `RX-train-response` set and appends a training sequence within `TRN-R` subfields. We confirm this by correlating the frame with a 128-bit Golay complementary sequence using our decoding software, which results in a positive test as illustrated in Figure. 7.4 (bottom). It is not the case for the second BRP request (third frame). As a response to each request, the responder (the AP) sends back to the initiator (the client) BRP requests (second and fourth frames in the figure) including in both training sequences. We thus hypothesize that the client's firmware uses the training sequence of one these two messages to make a CIR estimation of the channel between the AP and the client, and that the remaining frames simply complete the standard BRP protocol and are redundant.

As a result of the aforementioned process, the firmware provides a CIR measurement report with a data structure shown in Table 7.1. Through experimentation, we learned that the phase values are represented using 10 bits, ranging between 0 and 1023, to encode phase values between 0 and 2π radians; whereas the amplitude values range between 0 and 178. After appropriate scaling, this essentially corresponds to an estimate of $\mathbf{h}_{\text{RX},n}^{(i,j)}$, for a given beam pattern n between AP s_i and s_j.[5]

Type	Length	Description
16-byte unsigned integer	2	Phase and amplitude of the strongest path
2-byte unsigned integer	32	Amplitude at each antenna element
2-byte unsigned integer	32	Phase at each antenna element

Table 7.1: CIR measurement report

7.4.2. Hardware measurements

The aforementioned mechanism provides channel samples encoded as a $2N$-dimensional vector with phase and amplitude information across all the antenna elements. As an example, we illustrate in Figure. 7.5 a collection of 60 measurements for one antenna element (index 1) across two experiment sets where transmitter and receiver are 1 and 3 meters apart, respectively. The remaining parameters such as fixed beam pattern configuration is the same in both scenarios. From the figure, we can observe that, perhaps surprisingly, both scenarios report practically the same CSI estimates, despite the different link lengths.

It becomes evident, in light of this result, that the hardware reports a normalized version of $\mathbf{h}_{\text{RX,n}}$ (or $\hat{\mathbf{h}}_{\text{RX,n}}$), probably because such estimates are reported *after* dynamic gain control (AGC).

[5]Hereafter, we remove the superscript (i, j) whenever it is evident from context or not needed for explanations to reduce notation clutter.

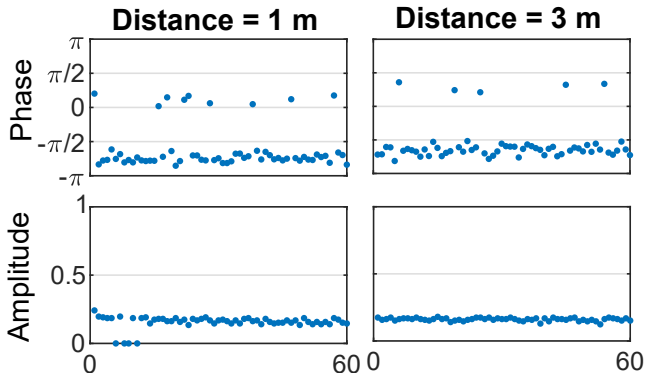

Figure 7.5: CSI Measurements vs Distance.

This implies that the data available does not have an absolute reference point, namely

$$\gamma \frac{\mathbf{h}_{\text{RX,n}}}{\|\mathbf{h}_{\text{RX,n}}\|}$$

where $\gamma \in \mathbb{C} : |\gamma| = 1$ is a random phase shift indicating the lack of absolute reference point. This means that extracting a virtual channel of perturbations $\bar{\mathbf{h}}_{\text{RX},n}$ as explained in §7.3.3 becomes a challenge.

However, in order to compute a close approximation of $\bar{\mathbf{h}}_{\text{RX},n}$, we can rely on the fact that the paths are not modified but created/destroyed and so there are two possibilities to estimate $\bar{\mathbf{h}}_{\text{RX},n}$: either the new path is created/destroyed with direction of departure in the direction in which $\mathbf{p}$ is steering, in which case the contribution to the receiving channel is going to be significant; or not, in which case it is not going to be significantly modified.

Note that, in the first case, since two different paths cannot follow the same trajectory, a path in the direction of departure of $\mathbf{p}$ *and* the direction of the arrival in the new path cannot exist. In light of this, we assume that the receiving channel is composed of paths that are distinguishable from the one created upon the presence of an obstacle, and so, the receiving channel with and without the object $\mathbf{h}_{\text{RX,n}}$ is orthogonal to the virtual channel $\bar{\mathbf{h}}_{\text{RX},n}$ created by the perturbation. Due to this orthogonality, we can also compute the virtual channel as an orthogonal subtraction

$$\bar{\mathbf{h}}_{\text{RX},n} = \hat{\mathbf{h}}_{\text{RX,n}} - \frac{1}{\|\mathbf{h}_{\text{RX,n}}\|^2} \mathbf{h}_{\text{RX,n}} \mathbf{h}_{\text{RX,n}}^H \hat{\mathbf{h}}_{\text{RX,n}} \tag{7.5}$$

which is an expression that is not affected by the power normalization.

7.4.3. Filtering tones from imperfect CSI data

In addition to the above, Figure. 7.5 also presents some spurious samples in both amplitude and (especially) phase estimates. We confirm that this phenomenon occurs across all antennae, as depicted in Figure. 7.6 for three different antenna elements. Hence, it is important to devise a mechanism that accounts for imperfections in our measurements.

To this aim, we let $\langle \mathbf{g}_k : k \in \mathbb{Z}_+ \rangle$ denote a sequence of collected CSI samples. Our task

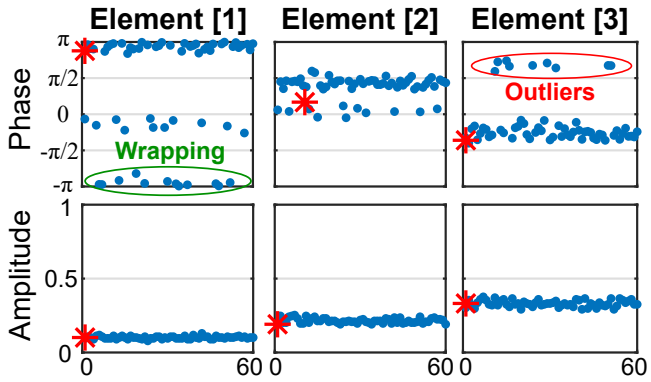

Figure 7.6: Phase and Amplitude Samples for 3 Antenna Elements.

is to design a function $f : \langle \mathbf{g}_k \rangle \rightarrow \mathbf{g}$ that filters out spurious information to provide a close approximation to the *real* $\mathbf{g}$. Our approach lies upon the observation that the proportion of spurious information $\tilde{\mathbf{x}}$ is very small, compared to the values of interest $\mathbf{x}$, when k is sufficiently large, i.e., $|\mathbf{x}| > |\tilde{\mathbf{x}}|$. Hence, if we denote the projection matrix of the sampling sequence as

$$\mathbf{\Pi} = \sum_k \frac{1}{\|\mathbf{g}_k\|^2} \mathbf{g}_k \mathbf{g}_k^H,$$

we will be able to decompose $\mathbf{\Pi}$ into 2 submatrices, $\mathbf{P}$ and $\mathbf{E}$, with the contribution of the values close to the real CSI and an error on such estimation, respectively, using SVD decomposition. Considering that $|\mathbf{x}| > |\tilde{\mathbf{x}}|$ and a sufficiently large k, then we know that for any $\mathbf{v} \in \mathcal{C}^N$

$$\mathbf{v}^H \mathbf{E} \mathbf{v} = \sum_k \frac{1}{\|\mathbf{g}_k\|^2} \mathbf{v}^H \mathbf{g}_k \mathbf{g}_k^H \mathbf{v} < \sum_k 1 = |\tilde{\mathbf{x}}|$$

and for real CSI, $\mathbf{g}$, we have that

$$\mathbf{P} = \sum_{k:\mathbf{g}_k \sim \mathbf{g}} \frac{1}{\|\mathbf{g}_k\|^2} \mathbf{g}_k \mathbf{g}_k^H \sim \sum_{k:\mathbf{g}_k \sim \mathbf{g}} \frac{1}{\|\mathbf{g}\|^2} \mathbf{g}\mathbf{g}^H = \frac{|\mathbf{x}|}{\|\mathbf{g}\|^2} \mathbf{g}\mathbf{g}^H.$$

So, since the contribution of $\mathbf{P}$ to the eigenvalues of $\mathbf{\Pi}$ is much larger than the one of $\mathbf{E}$, then we know that $\frac{1}{\|\mathbf{g}\|}\mathbf{g}$ is a good approximation of the maximum eigenvector λ of $\mathbf{\Pi}$. Then, by using the Schwartz inequality, $\mathbf{g}_k : |\mathbf{g}_k^H \lambda| > \cos(\epsilon)$, where ϵ is a tolerance error in radians, we can use λ to filter samples that contribute to the information of the channel state and remove spurious data. After this, $\mathbf{g}$ is estimated to have the mean amplitude and phase of all $\mathbf{g}_k e^{-\arg(\mathbf{g}_k^H \lambda)i}$ with $\mathbf{g}_k$ selected by the filter. The exponential term is a phase correction to have a common reference value (λ). Figure. 7.6 highlights with a red star symbol the resulting amplitude and phase information from the estimated $\mathbf{g}$ in these experiments.

7.4.4. Creation of custom beams

A key design requirement of `BeamScanner` is the need to collect CSI measurements when using highly-directional (narrow) beams. This minimizes the number of paths between transmitter and receiver and helps the filter designed above to expose only information that is useful to detect

obstacles in the environment. However, the shape of available beam patterns $\{\mathbf{p}\}$ in commercial devices do not satisfy this requirement, and hence we need to find new ones.

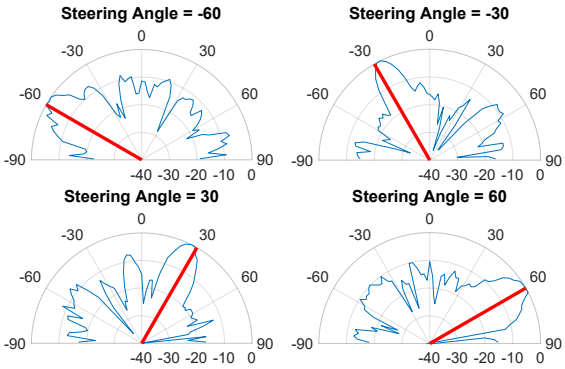

Figure 7.7: Custom Radiation Patterns.

The key challenge to designing new beam patterns is to estimate the steering vector of the antenna array of the device of interest across the whole azimuthal dimension ($\mathbf{a}_{\mathrm{RX}}$). The steering vector is a complex matrix of size $K \times N$ where K is the number of azimuth angles, and N is the number of antenna elements. If this were known, to generate a directional pattern towards a specific direction, we would need to set the antenna weights $\mathbf{p}$ to the conjugate of the steering vector in the desired azimuth angle. We obtain the steering vector of the phased antenna array of the TALON router from [63]. As a result, we can generate custom radiation patterns with very narrow beams in the desired direction. This can be observed in Figure. 7.7 for four beam patterns with a strong narrow lobe at $\{-60, -30, 30, 60\}$ degrees of azimuth angle, respectively, and low power in other azimuth directions.

7.4.5. Protocol integration

Building upon IEEE 802.11ad, the upcoming IEEE 802.11ay contains several amendments that enhance the standard beamforming training procedure, in addition to supporting new transmission modes and other features [6]. One of these techniques is the asymmetric beamforming training which allows simultaneous training of the transmit sectors on the AP side and the receive sectors at the client side. This is mainly achieved by appending TRN-R subfields at the end of each DMG beacon. These TRN-R subfields allow the extraction of the CSI information per antenna element at the receiver side. Additionally, the IEEE 802.11ad standard allows the clustering of neighboring APs so they do not interfere with each other. This allows all the APs in the same proximity to measure CSI towards each other within the same BI. At the end of the BI, the centralized controller would have collected all CSI between all the APs and constructed a channel profile as in Section 7.4. Figure 7.8 shows the integration of `BeamScanner` in 802.11ay.

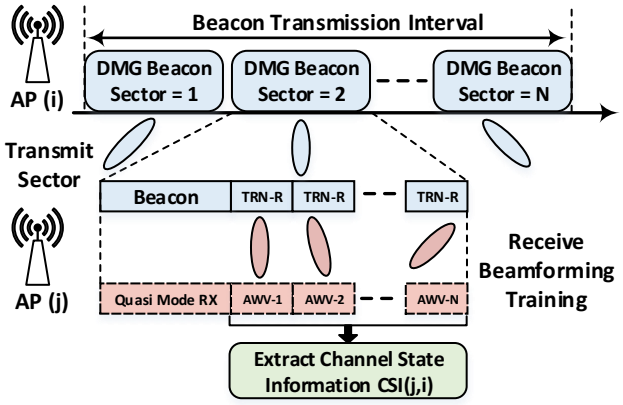

Figure 7.8: `BeamScanner` Protocol Integration in IEEE 802.11ay.

7.4.6. Classification

Our method periodically scans the environment, using narrow mmWave links among distributed APs in the room, for the fingerprint of *alien* obstacles, compared to baseline empty room. The main challenge is that such fingerprint may consist of high-order reflections and/or some other non-linear effects that are hard to infer. In this way, we resort to *k-nearest neighbor* (KNN), a machine learning (ML) based method to do the classification legwork. KNN is a simple non-parametric classification method that allows us to contrast channel state fingerprints to baseline scenarios built in an offline phase.

7.5. Experimental Evaluation

Note that we focus on the presence of human obstacles (indeed, the more compelling and challenging scenario). In this way, we *train* the classifier with a baseline scenario (with no human obstacles), and with the presence of a human obstacle at different predefined locations.

We perform our measurements in a large room of a size 6x11 m. In this room, we deploy one AP at each corner to provide wide coverage. Each AP uses a custom codebook that has 19 directional beam patterns with steering angles that changes from -90 to 90 with a step of 10°. We divide the room into 50 blocks of 60x60 cm. For each block, we place a human obstacle, and we collect CSI measurements from all the APs using the procedure described in Section 7.4. Collecting CSI measurements for each location spans 90 mins. Eventually, we end up with a database of complex CSI values of size 12x19x50x32.

The goal is to identify the presence of a human obstacle and try to estimate its location using our fingerprint database. We generate the entries for our fingerprint database using the following steps:

- For each location of the obstacle, we compare how did the CSI measurements for a specific beam change with respect to the baseline (empty room). We do not care about

which antenna element changed and how did it change but rather how much was the change for each antenna element.

- Based on a predefined threshold that we set as a hyper-parameter, we identify whether the CSI measurements have changed or not. In the current work, we check the median for the phase and amplitude values and compare it to the predefined threshold. If the changes exceed the threshold, we mark the CSI values as 1 otherwise 0. As a result, the CSI measurements for a certain beam will be reduced to a single boolean value.

- As a result, we end up with a vector of 1's and 0's for each location in the room. The number of vectors depends on the granularity of the measurements.

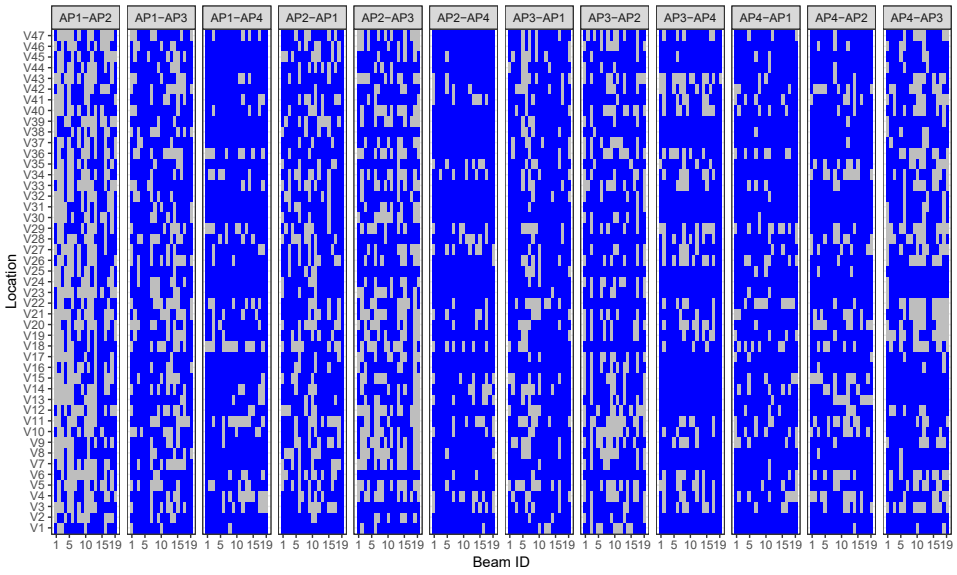

Figure 7.9: Footprint of the obstacle for all trained locations. Blue color means a positive footprint in either phase or amplitude compared to baseline (empty room).

Figure 7.9 shows the footprint for our fingerprint database using the previous procedure, where the blue areas represent a positive CSI footprint. For each new set of CSI measurements, we repeat the previous steps again. Besides, we perform the following:

- We look at similarities between the newly generated vector and all the vectors stored in our database.

- The vector that receives the highest score in terms of similarity is considered as the location of our object.

We generate a new set of samples that correspond to a human walking in a particular trajectory that does not contain any of the points in our database. We plot the empirical CDF for the distance

error for different values of phase threshold as shown in Figure. 7.10. From the plots, we can see that setting the threshold to a value of 0.04π performs the best and we can localize the obstacle with a median error of less than 3 m. This is considered good as our database is small and contains a limited number of samples.

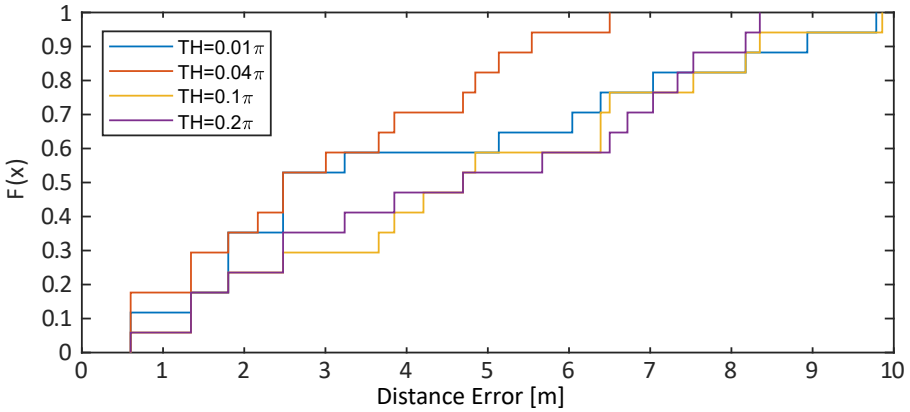

Figure 7.10: Empirical CDF.

The small localization errors in Figure 7.10 correspond to the case when the object is close to one pair of the deployed APs. This is because the power of the received signal through reflection off the object is strong enough to surpass the LoS signal. Additionally, the object might interrupt with high probability the primary communication path, and as a result, the communicating pair selects an alternative path which affects the measured CSI. However, when the object is located in the center of the room, reflections from the objects become very weak; thus, they do not affect the measured CSI values.

7.6. Discussion

The fusion between communication and sensing in the mmWave band paves the way for new applications that are not possible with the current bandwidth-limited wireless technologies. Electronically steerable phased antenna array operating in the mmWave band integrates a high number of antenna elements in a small footprint. Typically these elements are spaced by a half lambda which allows each element to observe a different version of the channel. This makes it attractive for sensing applications that require rich spatial information from the environment.

We enforced the TALON routers to use custom directional beam patterns with specific steering direction. However, these beam patterns exhibited strong side lobes and back lobes. Besides, for some of them, the difference between the main lobe and the side lobe is around 10 dB only. This difference is not sufficient to affect the measured the CSI when the object is far from the APs. The strong side lobes in the beam pattern are due to the intrinsic characteristics of the utilized antenna array including the irregular arrangement of its element, the mutual coupling between these

elements, and the variability of its directivity across the operating channel.

Building analytical models to analyze how the CSI changes for each scenario is difficult and analytically intractable. For these reasons, machine learning and deep learning approaches are convenient as they are capable of discovering intrinsic and hidden patterns in the data without manual intervention. However, finding the right algorithm and the correct parameters is an open research problem. Additionally, these methods require collecting a high number of measurements and data labeling to perform well. Although we have collected a high number of measurements at different reference points within our measurements setup, still our database is not sufficiently large for advanced machine learning algorithms.

As a proof-of-concept, we have used a simple machine learning algorithm to find the closest match of the current channel profile in our fingerprint database. Although our classification method is capable of identifying the presence of a human within a large room, it generates very high localization error for some case. This is because the QCA9500 chipset reports only the strongest path (single tap) in the channel for each antenna element of the phased antenna array. As this chipset is designed for communication only, relying on the strongest path is enough. Additionally, based on the orientation and the location of the human within the room, the newly created paths might not be strong enough to change the strongest tap in the channel. However, these new paths will affect other channel taps. Additionally, the QCA9500 chipset does not support OFDM modulation, which prevents obtaining CSI information per subcarrier similar to the commercial wireless devices operating in the microwave band [108]. To improve the system's performance and reduce localization error, we plan to do the following:

- Perform reverse engineering to the firmware running on the QCA9500 chipset to extract CSI information for all the channel taps.

- Extract CSI information in the lower frequency and correlate this information with the CSI information obtained in the 60 GHz band. This would provide more spatial information about the environment and helps reducing localization error.

7.7. Conclusions and Future Work

In this work, we have implemented a novel system to proactively localize human obstacles within an environment using COTS 60 GHz devices. Our system runs on commodity 60 GHz devices and does not require any hardware modifications as it is based on the operations of the IEEE 802.11ad protocol. We plan to extend our work and leverage the localization information to predict human mobility pattern. In this way, an AP can predict whether a walking human is going to interrupt or block its communication link and thus proactively switch its current active transmit and receive sectors without the need for time-consuming beamforming training. Additionally, we plan to study the accuracy of our system and its ability to detect multiple objects within the same environment when we decrease/increase the number of deployed APs.

Chapter 8

Conclusions

In a nutshell, this thesis provides a set of algorithms and solutions to enable robust and reliable wireless communication in the 60 GHz band. In particular, our work focuses on improving the performance of the IEEE 802.11ad protocol both in simulation and practice.

First of all, due to the lack of system-level simulators that incorporate a wireless technology operating in the mmWave band and the fact that COTS devices operating in the 60 GHz band prevents accessing lower layer information, we took the initiative and provided the design and implementation of the first open source wireless model for simulating the IEEE 802.11ad protocol in network simulator ns-3. The model facilitates analyzing and debugging wireless networking performance for highly dense wireless networks incorporating a large number of devices and access points, which is extremely complicated and costly using COTS devices. Additionally, it allows studying the impact of mmWave channel dynamics on the operations of the upper layer of the protocol stack. Next, using this model, we demonstrated the MAC layer capabilities of IEEE 802.11ad protocol. Furthermore, we studied the performance of the protocol for various deployment settings and provided: (i) Recommendations on how to deploy mmWave WLAN networks to ensure persistent gigabit of throughput per user. (ii) A novel aggregation-aware algorithm that alleviates high channel contention for dense mmWave networks relying on the CSMA/CA protocol for data transmission.

Then, we delved into the performance of practical and standard compliant COTS 60 GHz devices. These devices exhibit unique characteristics with respect to how 60 GHz devices should perform in theory. We looked at various wireless networking aspects and provided solutions at different layers of the protocol stack to enhance the efficiency and performance of these devices when operating within a wireless network. In particular, we analyzed the performance of the TCP protocol over CSMA/CA in highly dense WLAN network settings, and we found the following:

- TCP in the uplink does not perform well when the number of stations increases. This is because of the mismatch between the actual channel capacity and the configured TCP buffer size. Solving this issue would require deploying dynamic buffer sizing strategies at the transport layer to adapt the buffer size to the actual bandwidth-delay product (BDP),

which improves overall network throughput and reduces latency. This problem already exists for wireless networks operating below 6 GHz. However, it is exacerbated in the mmWave band because of the multi-gigabit-per-second speeds at the physical layer and the deafness introduced by directional antennas.

- In the downlink, TCP shows some inefficiencies due to the collisions between TCP packets and the returning TCP ACKs. Since these devices cannot distinguish if a packet loss is due to collision, deafness, or shadowing, stations frequently try to access the wireless medium to perform beamforming training with the corresponding AP. The beamforming training consumes channel resources and affects TCP throughput.

Interference is a major limiting factor in current wireless networks operating in the microwave band. Directional communication through electronically steerable phased antenna arrays alleviates the interference issue and allows multiple wireless devices to communicate at the same time with minimum interference. Since communication in the mmWave band requires directional links to compensate for the high path loss in this band, this makes mmWave a viable choice for highly dense wireless networks. However, we found that practical phased antenna arrays do exhibit non-ideal behavior and tend to generate imperfect beam patterns with many strong side-lobes in contrast to what theory suggests. This has an impact on network deployment and mandates careful network planning. The wide beam patterns of these antenna arrays solve the problem of device mobility and avoid the need for frequent beam scanning. However, this comes at the expense of reducing the degree of spatial reuse. Possible solutions to compensate for the consequences of non-directionality in practical phased antenna array on spatial sharing include incorporating adaptive carrier sensing threshold and exploiting the frame capture effect to enhance the network's throughput.

Finally, the mmWave band provides substantial bandwidth of several hundreds of megahertz compared to wireless networks operating in the microwave band. This makes it attractive for ranging, localization, and radar-based applications as those applications require high time resolution. Motivated by the above, we utilized the IEEE 802.11ad protocol to build a multi-AP collaborative beam scanning protocol that is capable of localizing human obstacles within an environment passively and transparently without end-user intervention. We utilized a simple machine learning algorithm for this purpose, which resulted in good accuracy. We are planning to improve the performance of our system to detect multiple objects and pinpoint their locations with a marginal error. This can be achieved through the use of advanced signal processing techniques combined with state-of-the-art machine learning algorithms like deep learning. Our system can be seamlessly integrated into the upcoming IEEE 802.11ay standard and, with mild overhead, into IEEE 802.11ad enabled COTS devices.

8.1. Future Work

We present here possible research directions based on our insights gained from this thesis. Precisely, we identify open issues and ideas for future work per chapter.

The WiFi Alliance is currently finalizing the IEEE 802.11ay amendment and commercial products that adopt this protocol will be available soon in the market. In Chapter 3, we provided a proposal for simulating SU/MU-MIMO communications in ns-3 using a minimum number of signal processing blocks. The next step would be to implement this proposal and compare its performance to a full-fledged link-level simulator that incorporates the complete signal processing chain at the physical layer. Upon completing this step, we can study various networking aspects for the IEEE 802.11ay protocol for large network settings. These aspects include (i) Users scheduling algorithms for MU-MIMO communication. The challenge here would be how to group the users while ensuring minimum interference between the spatial streams, channel access fairness, and improved network throughput. (ii) Beam patterns synthesis for SU/MU-MIMO communication. (iii) Hybrid beamforming architectures with different digital precoding algorithms. (iv) Rate-adaptation algorithms for multi-stream communication.

There has been a tremendous amount of research work on cross-layer solutions for wireless networks operating in the microwave band. However, this is not the case for mmWave networks. Cross-layer solutions require accurate and reliable implementation and modeling for the operations at each layer of the protocol stack. Additionally, these solutions cannot be implemented on COTS devices as these devices do not allow modifying their operations as we have seen in Chapter 6. As a result, researchers resort to high fidelity network/system-level simulators to try out their ideas. Our simulation model in ns-3 satisfies all the aforementioned requirements and paves the way to develop and test innovative solutions that improve the performance of 60 GHz wireless networks. In particular, we can study the interactions between rate adaptation algorithm, frame aggregation, beamforming training mechanism, channel access schemes, buffer management strategies, TCP congestion control, and traffic flows and patterns. Optimizing the operations of all of these mechanisms together and their impact on the overall network performance is a complex problem, and closed-form solutions are infeasible. However, heuristic solutions could be tested and validated using network-level simulations.

Another possible research direction is to look at wireless networking scenarios for extreme dense mmWave networks. These networks will incorporate a high number of access points and mobile devices with heterogeneous capabilities and service requirements constraints. In these scenarios, interference is the performance limiting factor. As a result, coordination mechanisms between access points must be developed to eliminate and minimize the interference and provide high QoE to these mobile devices. These coordination mechanisms include deploying a central controller that has a global view on the network and perform the following tasks: (i) Distributing the users between access points for load balancing. (ii) Provide handover services to avoid link interruption. (iii) Schedule data transmissions in a way that guarantees to serve a high number of

users while still minimizing interference.

mmWave technology accommodates extreme bandwidth, which makes it attractive for radar, sensing, and ranging applications. These applications require high time resolution, which is not possible using wireless technologies operating in the microwave band. Additionally, these applications are essential ingredients for advanced driver assistance systems for the automotive industry. An autonomous driving system must combine various information from many sensors to provide a seamless driving experience and ensure passengers safety. Some of this information is obtained from either nearby vehicles or static infrastructure that has a global view on the status of the road. Exchanging this amount of information within short periods requires a new wireless paradigm that is beyond the capabilities of the state-of-the-art wireless technologies. For these reasons, the automotive industry is heavily investing in the mmWave technology as it provides extremely low-latency for communication together with ultra-high data rates. However, exploiting the full capabilities of the mmWave technology for connected vehicles requires answering the following research questions: (i) How frequent should a vehicle perform beam training? Moreover, how to use out-of-band information to configure the antenna array? (ii) Should mmWave exclusively be used to transfer sensitive information, or should it be combined with another wireless technology? (iii) What type of phased antenna arrays should be deployed in the vehicle?

Appendices

Appendix A

Device Tracking in 60 GHz Wireless Networks

A.1. Introduction

Most mmWave systems exchange beam training control messages such as Sector Level Sweep (SLS) frames to address device tracking. This is the usual approach both in related work (e.g., [109–111]) as well as in mmWave networking standards such as IEEE 802.11ad [1]. Existing COTS devices operating according to this standard [52] typically use simple SNR and time-based mechanisms to decide when to do beam training. A basic solution is to monitor the SNR and trigger a beam-sweep whenever a big drop in SNR occurs or after a certain amount of time has passed. During a beam-sweep, devices sequentially send a SLS frame on each antenna sector to find a better steering [1]. In the worst case, a beam sweep on a COTS device can take seconds [52, 109]. Since data rates in the mmWave band are in the range of gigabits per second, such interruptions are extremely harmful. While the duration of the beam-sweep itself can be reduced down to a few milliseconds [112], we show that short but frequent interruptions still have a substantial effect on throughput. Recent work focuses on further reducing the duration of beam-sweeps. This includes sending beacon and SLS frames via multi-lobe beam patterns [109, 110], preceding each data frame with a very short control frame that probes neighboring sectors [113], using past steering information to only probe the most promising sectors [114], or resorting to beacons on lower-frequency bands for angle-of-arrival estimation [111]. The standard also defines basic beam tracking [1]. Still, all of these approaches require the exchange of dedicated control messages, which may result in excessive overhead if device movement and/or rotation is significant, or if fragile reflected paths are used due to blockages such as human persons. Moreover, due to the cost of frequent periodic retraining, most approaches only react to SNR drops *after* they occurred, causing transient link degradation and interruptions.

In this work, we implement a device tracking mechanism for IEEE 802.11ad that incurs zero overhead [115]. Moreover, it does not rely on SNR measurements to detect beam misalignment

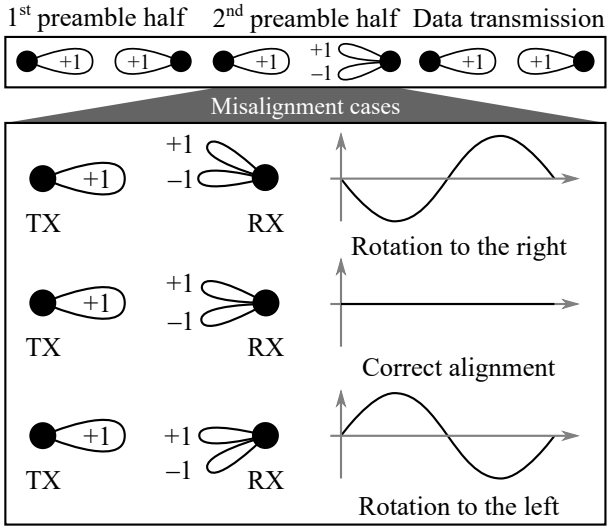

Figure A.1: Toy example of our device tracking mechanism.

but continuously adjusts the beam-steering as the user moves. The approach can track both LoS and NLOS paths, thus being suitable for scenarios with blockage. We only perform one initial beam-sweep to find the path. After that, we can track it as long as the path changes gradually. Abrupt changes are due to sudden blockages. Finding a new path during such a blockage is an orthogonal problem to tracking a given LoS/NLOS path, and thus out of the scope of this work. However, we discuss the impact of this problem in our evaluation in Section A.3. Further, our scheme does not require any feedback but operates on each node independently. We do not modify the 802.11ad protocol. Thus, a device implementing our mechanism is fully backwards-compatible with regular 802.11ad nodes. More importantly, *a device still benefits from our approach when communicating with a regular node*.

The key idea of our mechanism is to use a beam pattern with two strong adjacent lobes (Figure A.1) instead of just one main lobe when receiving part of the preamble of each packet. The rest of the preamble, as well as the headers and the data payload of the packet, are received with a regular beam pattern, as shown at the top of Figure A.1. Recent work [116, 117] shows that such fast switching among beam patterns is feasible in practice. The segment of the preamble that the node receives with the two-lobe beam pattern allows inferring whether the node has rotated or moved. To this end, we design our two-lobe beam pattern such that the signal received via one of the lobes has a 180° phase shift compared to the other lobe, indicated by +1 and -1 in the figure. This is feasible with a conventional phased array, and a single transmit chain since the phase shift is due to the beam pattern itself. By comparing the phase of the first half of the preamble to that of the second half, we can deduce whether the node received the second half primarily via the positive or the negative lobe. This reveals whether a rotation has taken place. As shown in Figure A.1, a rotation to the left results in both preamble parts being in phase, whereas a rotation to the right results in a 180° phase shift. Further, correct alignment results in very low

received signal strength, which we can easily detect. In this case, the receiver only receives half of the preamble, but we show that this does not compromise signal acquisition quality. Depending on the outcome, the receiver steers the beam to the right, to the left, or not at all. The same principle is valid not only for data packets but also for acknowledgments, allowing both sides of the link to track one another. The undesired side-lobes of practical beam patterns do not hinder the operation of our scheme since they do not point to the transmitter. We validate this in our practical experiments. Since our mechanism operates on the regular 802.11ad preamble, it does not incur any additional overhead. Our design becomes even more critical for future antennas with a higher number of elements and narrower beam width. Such antennas incur an even higher steering overhead with conventional approaches, while they allow designing more fine-grained multi-lobe patterns and thus more efficient preamble based beam steering.

For our design, we leverage several physical layer effects inherent to mmWave communication:

1. **Beam Pattern Design.** Conventional phased arrays allow to design multi-lobe beam patterns with a pre-defined phase shift among the lobes. We exploit this characteristic to transmit a signal with 0° and 180° phase shift simultaneously.

2. **Indoor Propagation.** Due to the high absorption of walls, 802.11ad APs typically cover a single room only. We show that this results in a sufficiently high SNR to allow signal acquisition with only the first part of the preamble *when the antenna beams are roughly aligned*. In the case of strong misalignment, we resort to the full preamble.

3. **Correlation Properties.** The excellent correlation properties of the Golay sequences in the 802.11ad preamble [1] allow to reliably estimate the phase shift between the first and the second preamble half even at very low SNRs. This enables our system to work even if the lobes of the two-lobe beam are relatively far apart.

We build on these effects to implement zero-overhead mmWave device tracking. Our contributions are as follows:

- We design a mechanism that enables accurate rotation and movement tracking on a per-packet basis, without incurring any control overhead.

- We present a method to infer link changes by detecting phase inversion compared to a reference phase in the 802.11ad packet preamble.

- Our mechanism is fully backward compatible and works when communicating with legacy 802.11ad devices. It only requires changes to the antenna pattern with which packet preambles are received.

- We evaluate the steering accuracy of our approach in practice using 60 GHz electronically steerable antenna arrays. For most cases, our error is below 5°.

- We perform an extensive simulation campaign using an accurate ray-tracing-based channel model along with ns-3 and achieve up to $2.39\times$ better throughput (c.f. Section A.3).

A.2. Related Work

Device Tracking. Earlier work uses 60 GHz communication for object tracking following a radar-based approach [118, 119]. However, this requires nodes to simultaneously transmit and receive to, for example, track hand gestures or a pen. In contrast, we aim at tracking an independent transmitter and operate our mechanism at the receiver side only. Other approaches use external radars in the 60 GHz band to track the movement of objects [120, 121]. Our work stands apart from such schemes since our approach does not require any additional devices and is backward compatible with 802.11ad.

AoA Estimation. Similarly to related work, our mechanism tracks changes in the AoA [122–124]. The key difference to our approach is that we only compute relative changes, which simplifies our design significantly and allows us to operate with a single receive chain. Moreover, our mechanism does not need to resort to lower frequency bands for AoA estimation such as [111, 125].

Efficient Steering. Instead of using sophisticated AoA estimation techniques, a large body of work focuses on sector-level accuracy, improving on the techniques standardized in 802.11ad [1]. To this end, such approaches often resort to beaconing on different sectors to determine the approximate angle to the receiver [126, 127]. This includes transmitting beacons on multi-lobe beam patterns, using beam patterns with random phase shifts for compressive sensing based tracking [109, 110, 128], or collecting historical information to reduce the duration of the beam-sweep [114, 129]. Still, in contrast to our work, such approaches still incur training overhead. Other approaches use hybrid beam-forming to receive beacons from multiple directions simultaneously [87]. This requires a sophisticated antenna design as well as multiple receiver chains, while we operate with a regular analog beam-forming antenna.

Hardware Validation. A key challenge in 60 GHz research is validating that mechanisms work in practice. Related work builds hardware for this purpose [65, 130–132] but typically does not achieve the full bandwidth of 802.11ad and uses antenna arrays with a limited number of elements. As a result, researchers often must resort to mimicking the behavior of 60 GHz radios using horn antennas. While this does provide fundamental insights into the performance of a mechanism, it cannot fully capture how it would operate on actual 60 GHz COTS hardware. In contrast, we validate our approach on a full-bandwidth system with a highly flexible phased antenna array. This allows us to get deep insights into our mechanism which otherwise would not be feasible.

A.3. Simulative Evaluation

We implement our device tracking mechanism in MATLAB using standard-compliant 802.11ad signals as a basis. Further, we use the resulting SNR and MCS traces as input to the ns-3 simulator [133]. This allows us to understand the impact of rotations and movement not only at the physical layer but throughout the protocol stack for a wide range of scenarios.

A.3.0.1. Simulation Setup

Figure A.2 depicts the architecture of our simulator. It consists of a symbol-level part and a packet level part. The former is based on MATLAB and the latter on ns-3. As a result, we obtain an accurate model of the physical layer while at the same time understanding the impact on the transport layer.

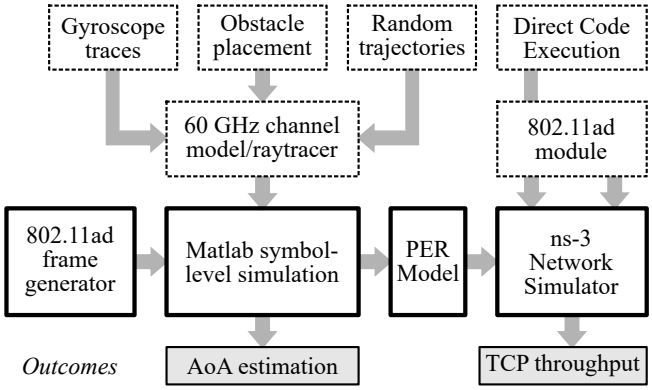

Figure A.2: Architecture of our full protocol stack simulator. Bold boxes indicate the simulation tool-chain, dashed boxes are inputs, and gray boxes are outputs.

A.3.0.1.1. Packet-Level Simulation To evaluate the impact of the physical layer dynamics on upper-layer protocols, we input the SNR and MCS from the symbol-level simulations to ns-3. To translate these values to the packet-level domain, we use the 802.11ad PER model released by the authors of [4]. We use an ideal rate control algorithm, which always predicts the correct MCS. This is the worst case algorithm for our device tracking mechanism since it mitigates the impact of the frequent SNR drops of the baseline scheme by choosing the optimal MCS and thus avoiding packet errors. Further, we use our 802.11ad module for ns-3 in Chapter 3 which recreates the full frame format of the protocol. We run ns-3 in DCE mode [134]. That is, instead of using the ns-3 implementation of upper-layer protocols such as the TCP, we use the actual implementations of the protocols from the Linux kernel. Our experiments run TCP Cubic [135] with enabled Selective Acknowledgment (SACK). In our infrastructure-based scenarios (c.f. Section A.3.0.2), we also use ns-3 to model a backbone network and a remote server. The delay to reach the server is 40 milliseconds. In all of our scenarios, we use the `BulkSendApplication` of ns-3 at the application layer, which behaves similarly to a File Transfer Protocol (FTP) server. As a result,

we obtain the TCP throughput for each of the trajectories that we generate in the symbol-level part of the simulation.

A.3.0.2. Scenarios

To generate our scenarios, we consider three basic parameters: the size of the room, type of the network (AP or point-to-point), and type of obstacles. As a baseline mechanism, we recreate the behavior of 802.11ad. That is, whenever the SNR decreases due to misalignment or blockage, the baseline performs a beam-sweep to correct for the misalignment or find a new path.

A.3.0.2.1. Room Size The room size has a direct impact on the strength of the reflections. We consider medium rooms with a size of 7×7 meters such as open-plan areas in offices, and large rooms with a size of 15×15 meters such as a conference hall.

A.3.0.2.2. Network Type We consider both infrastructure-based and point-to-point scenarios. In the point-to-point case, two mobile devices communicate within the same room. Both follow arbitrary trajectories. In the infrastructure-based case, we place an AP in a randomly chosen corner of the room, and let the mobile device move on an arbitrary trajectory. This mimics a scenario where a user is accessing a remote server on the Internet via 802.11ad. The key difference between the two network types is that the delay is much higher in the infrastructure-based case. While TCP recovers quickly in the point-to-point case due to the low delay, the impact of SNR fluctuations is much higher in the remote server case.

A.3.0.2.3. Blockage Type We generate scenarios both with and without blockages at random locations. Their size is 1×1 meters, and the default material is metal. We also consider self-blockage scenarios. In this case, we model a person holding a phone as a mobile obstacle that follows the node movement. The material is such that it absorbs most of the energy, similar to a human. The size of the obstacle is 0.3×0.875 meters, based on the average human size [136].

A.3.0.3. Results

For readability and conciseness, in the following we refer to the infrastructure-based case as the AP scenario and to the point-to-point case as the PTP scenario.

A.3.0.3.1. Outage Gain In the following, we analyze the impact of our tracking mechanism outages on TCP throughput. Figure A.3 depicts the Probability Density Function (PDF) of the throughput for both the PTP and the AP scenarios. In the PTP case, our device tracking approach achieves 3 Gbps almost all of the time. In contrast, 802.11ad fluctuates among throughput values ranging from 1.5 Gbps to 3 Gbps. As expected, our approach is significantly more stable than 802.11ad. For this case, we achieve a throughput gain of up to $1.34\times$. In the AP case, the comparatively high delay to the remote server has a clear impact on the distribution of the throughput.

Our approach is again more stable than 802.11ad, but fluctuations are higher for both mechanisms since TCP needs more time to recover in case of link quality variations. For 802.11ad, this means that the throughput occasionally falls below 1 Gbps, whereas our approach is above 2 Gbps most of the time. Overall, this results in a throughput gain of up to $2.39\times$.

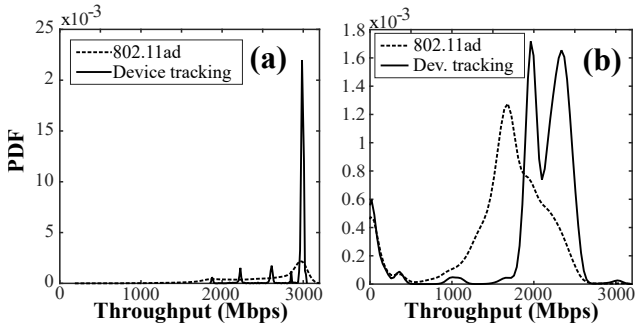

Figure A.3: Throughput PDF of (a) PTP and (b) AP scenarios for a large room with no obstacles. 100 repetitions.

Further, in Figure A.4(b) we depict the number of outages in the case with obstacles. As expected, this number is higher than in the case without obstacles. We reduce the number of outages by roughly 40% in a scenario with two obstacles, and by 30% in the self-blockage case.

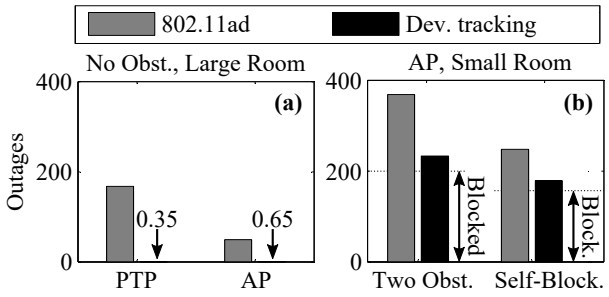

Figure A.4: Number of outages due to steering misalignment and/or blockage. The arrows in (a) show the values of the bars that are too small to be legible. The arrows in (b) indicate the fraction of outage which occurs because of no working (reflected) path being available due to the blockage.

We achieve this gain because our system is only affected by blockage outages, whereas 802.11ad additionally suffers from misalignment outages. Still, the gain is much smaller than for the case without obstacles. The reason is twofold. First, often no reflected path is available during a blockage. Thus, the outage cannot be avoided. Second, once blockage occurs, we often keep tracking the reflection even after the LoS is available again. Due to the robustness of Golay sequences, NLOS tracking continues even if the reflection is too weak for communication. That is, our tracking mechanism *is too accurate*. To address this, our mechanism can be paired with a LoS recovery approach.

A.3.1. Comparison to Existing Approaches

Related work on beam steering suggests several techniques to deal with beam misalignment. We survey them in Section A.2. Table A.1 gives an overview of how such approaches compare to our multi-lobe device tracking. In particular, we show the overhead in terms of time that each approach incurs when a misalignment occurs. In some cases, this time includes both tracking existing paths *and* finding new paths in case of blockage. While these problems are orthogonal, elaborate techniques that find new paths may also track existing ones as a by-product. Further, we use our simulator to assess the stability of the connection in terms of throughput. Existing approaches are typically reactive, that is, they only probe alternative sectors *after* a link disruption. As a result, throughput fluctuations are large. In contrast, our approach continuously tracks nodes at zero cost, thus avoiding misalignment outages altogether. While the device tracking mechanism suggested in the 802.11ad standard may achieve similar performance, it requires continuous probing. Specifically, nodes can request a beam refinement during on-going communication using training fields appended to each packet [1]. The standard defines the corresponding protocol but does not specify when such a beam refinement shall take place since those parameters are implementation-specific. While this hinders a detailed evaluation of 802.11ad tracking, we assess its performance based on the standard and related work [1, 137]. In 802.11ad, nodes must explicitly request a beam refinement to track each other. In response to such a request, nodes append training fields to their frames. In each training field, nodes transmit on a different sector, allowing their communication partner to asses the performance when using that sector. Finally, nodes exchange feedback on the assessed sectors. Each training field incurs an overhead of $5.52\mu s$. To achieve a performance similar to our approach, nodes would have to probe constantly one sector to the left and one sector to the right of the current sector, incurring an overall overhead per frame of $11.04\mu s$. Since practical frame sizes in 802.11ad often only go up to $25\mu s$ [7], an overhead of $11.04\mu s$ has a strong impact. Moreover, the 802.11ad tracking mechanism transmits training fields on adjacent beam patterns, which increases interference on neighboring links. In contrast, our mechanism operates at the receiver-side only and does not require feedback. Thus, we conclude that our approach is significantly more practical.

A.4. Discussion

Cross-Layer Effects. Further, we observe that taking into account not only the physical layer but also the link, network, and transport layers are crucial to fully understand the impact on the performance of the characteristics of 60 GHz communication. Beyond the influence of TCP dynamics discussed in Section 4.2, rate adaptation at the link layer also plays a crucial role. We evaluated our system for the worst-case, that is, for a perfect rate adaptation algorithm that always chooses the optimal MCS and thus mitigates the impact of the high SNR fluctuations of

[1] Values based on improvement claimed in the corresponding references.
[2] Simulation result for a large room with no obstacles and a point-to-point link.

Table A.1: Average overhead and throughput stability of existing approaches for a 16-element antenna array

Mechanism	Overhead[1]	σ Throughput [2]
802.11ad Default [1, 112]	640 μs	$\sigma > 1$ Gbps
Agile-Link [109]	385 μs	
BBS [111]	121 μs	
Reduced Search [129]	91 μs	
Adaptive Search [114]	70 μs	
MOCA [113]	20 μs	
802.11ad Tracking [1]	11.04 μs	$\sigma = 0.32$ Gbps
Our approach	0 μs	

802.11ad. To avoid rate adaptation errors resulting in high packet losses, maintaining a stable SNR is crucial. Our approach tackles this problem extremely well.

Multi-User Scenarios. A natural question arising from our design is how it performs in case of multiple users. Similarly to the 802.11ad protocol itself, our mechanism discards AoA estimations obtained from frames which are not addressed to the current node. While a receiver could receive higher interference due to the wider shape of the two-lobe beam pattern, this does not affect the data part of the packet, which is still received via the one-lobe beam pattern. To entirely avoid this type of interference, nodes could perform clear channel assessment using the two-lobe beam pattern.

A.5. Conclusions and Future Work

Tracking mobile devices in 60 GHz networks is challenging but crucial to prevent harmful beam pattern misalignments. Existing approaches to maintain alignment typically resort to periodic probing of adjacent beam patterns to keep track of the rotation and/or movement of a device, and thus incur significant overhead. In contrast, we present a device tracking mechanism that accurately tracks both rotation and movement of a device without incurring any overhead. The key idea is to receive part of the preamble of each packet via a custom two-lobe beam pattern. One of the lobes introduces a phase shift of $180°$. This allows the receiver to identify via which lobe it received the preamble, and thus estimate its relative rotation towards the transmitter. We implement our mechanism in practice using a full-bandwidth IEEE 802.11ad setup with an electronically steerable phased antenna array. We also perform a simulation campaign in a broad range of scenarios. We show that our mechanism can reduce the steering error to below $5°$ and achieves up to $2\times$ throughput gain. The mechanism becomes even more beneficial for future phased antenna arrays with a higher number of antenna elements and narrower beamwidths.

The next steps of our device tracking approach include the integration with techniques to recover the LoS path after a blockage, the design of more accurate beam patterns when steering

towards the edge of a phased array in practical deployments, and the improved filtering of AoA estimates based on typical human movement to better identify outliers due to measurement errors. Further, we consider extending our scheme to 3D.

References

[1] Standards_Committee, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," *IEEE Std 802.11g-2003 (Amendment to IEEE Std 802.11, 1999 Edn. (Reaff 2003) as amended by IEEE Stds 802.11a-1999, 802.11b-1999, 802.11b-1999/Cor 1-2001, and 802.11d-2001)*, no. June, pp. i–67, 2003.

[2] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.1 lad Devices," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '17. Incheon, Republic of Korea: ACM, 2017, pp. 414–425. [Online]. Available: http://doi.acm.org/10.1145/3143361.3143384

[3] U. The, "Exploiting the 60 GHz Band for Local Wireless Multimedia Access : Prospects and Future Directions," *IEEE Communications Magazine*, vol. 40, no. January, pp. 140–147, 2002.

[4] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall, "Augmenting data center networks with multi-gigabit wireless links," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, 2011, p. 38.

[5] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and bane of 60 GHz networks," in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '15. New York, NY, USA: ACM, 2016, pp. 1–13. [Online]. Available: http://doi.acm.org/10.1145/2716281.2836102

[6] Y. Ghasempour, C. R. C. M. Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: 60 GHz Communication for 100 Gb / s Wi-Fi," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 186–192, 2017. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/8088544/

[7] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and Bane of 60 GHz Networks : Practical Insights into Beamforming , Interference , and Frame Level Operation," in *CoNEXT 2015*, 2015.

[8] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and I. N. P. Aper, "IEEE 802.11ad: Directional 60 GHz Communication," *Communications Magazine, IEEE*, no. December, pp. 132–141, 2014.

[9] IEEE, "IEEE Std 802.11ac. Enhancements for Very High Throughput for Operation in Bands below 6 GHz," *IEEE 802.11 Working Group*, 2013.

[10] H. Assasa and J. Widmer, "Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3," in *Proceedings of the Workshop on Ns-3*, ser. WNS3 '16. New York, NY, USA: ACM, 2016, pp. 57–64. [Online]. Available: http://doi.acm.org/10.1145/2915371.2915377

[11] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and temporal characteristics of 60-GHZ indoor channels," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 3, pp. 620–630, 2002.

[12] X. Tie, K. Ramachandran, and R. Mahindra, "On 60 GHz Wireless Link Performance," in *Proc. of PAM 2012*, 2012, pp. 147–157.

[13] H. Assasa and J. Widmer, "Extending the IEEE 802.11ad Model: Scheduled Access, Spatial Reuse, Clustering, and Relaying," in *Proceedings of the Workshop on Ns-3*, ser. WNS3 '17. New York, NY, USA: ACM, 2017, pp. 39–46. [Online]. Available: http://doi.acm.org/10.1145/3067665.3067667

[14] H. Assasa, J. Widmer, T. Ropitault, and N. Golmie, "Enhancing the ns-3 IEEE 802.11ad Model Fidelity: Beam Codebooks, Multi-antenna Beamforming Training, and Quasi-deterministic mmWave Channel," in *Proceedings of the Workshop on Ns-3*, ser. WNS3 '19. New York, NY, USA: ACM, 2019.

[15] H. Assasa, J. Widmer, J. Wang, T. Ropitault, and N. Golmie, "An Implementation Proposal for IEEE 802.11ay SU/MU-MIMO Communication in ns-3," in *ACM Workshop on Next-Generation Wireless with ns-3 (WNGW 2019)*. ACM, 2019.

[16] H. Assasa and T. Ropitault. (2019) A Collection of Open-source Tools to Simulate IEEE 802.11ad/ay WLAN Networks in ns-3. [Online]. Available: https://github.com/wigig-tools

[17] A. Maltsev, A. Pudeyev, I. Karls, I. Bolotin, G. Morozov, R. J. Weiler, M. Peter, W. Keusgen, M. Danchenko, and A. Kuznetsov, "Quasi-Deterministic Approach to MmWave Channel Modelling in the FP7 MiWEBA Project," in *Wwrf'33*, Austin, USA, 2014.

[18] I. P. G. for Wireless Local Area Networks (LANs), "Channel Models for IEEE 802.11ay," 2017.

[19] G. Brown, "The Promise of 5G mmWave - How Do We Make It Mobile ? Today ' s Presenters," 2016. [Online]. Available: https://www.qualcomm.com/media/documents/files/the-promise-of-5g-mmwave-how-do-we-make-it-mobile.pdf

[20] Qualcomm, "QCA9500 SoC," 2019. [Online]. Available: https://www.qualcomm.com/products/qca9500

[21] D. Zhang, Y. Wang, X. Li, and W. Xiang, "Hybridly Connected Structure for Hybrid Beamforming in mmWave Massive MIMO Systems," *IEEE Transactions on Communications*, vol. 66, no. 2, pp. 662–674, 2018.

[22] I. 802.11, "IEEE Standard for Information technology-Telecommunications and information exchange between systems Local and metropolitan area networks- Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY)Specifications," *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)*, vol. 11, pp. 1–2793, 2012.

[23] T. Seki, K. Iigusa, H. Sawada, Y. Fujita, N. Orihashi, K. Science, and O. Siga, "Reference antenna model with side lobe for TG3c evaluation," pp. 1–13, 2006.

[24] Qualcomm, "Linux Wil6210 Driver," 2017. [Online]. Available: https://github.com/torvalds/linux/tree/master/drivers/net/wireless/ath/wil6210

[25] R. Lenner, G. J. Schilero, M. L. Padilla, and A. S. Teirstein, "A Survey on Hybrid Beamforming Techniques in 5G: Architecture and System Model Perspectives A Survey on Hybrid Beamforming Techniques in 5G: Architecture and System Model Perspectives," *Sarcoidosis Vasculitis and Diffuse Lung Diseases*, vol. 19, no. 2, pp. 143–147, 2002.

[26] F. Sohrabi and W. Yu, "Hybrid Digital and Analog Beamforming Design for Large-Scale Antenna Arrays," *IEEE Journal on Selected Topics in Signal Processing*, vol. 10, no. 3, pp. 501–513, 2016.

[27] J. Mittag, S. Papanastasiou, H. Hartenstein, and E. G. Ström, "Enabling accurate cross-layer PHY/MAC/NET simulation studies of vehicular communication networks," *Proceedings of the IEEE*, vol. 99, no. 7, pp. 1311–1326, 2011.

[28] R. Pal, K. V. Srinivas, and A. K. Chaitanya, "A Beam Selection Algorithm for Millimeter-Wave Multi-User MIMO Systems," *IEEE Communications Letters*, vol. 22, no. 4, pp. 852–855, 2018.

[29] H. Assasa, J. Widmer, T. Ropitault, A. Bodi, and N. Golmie, "High Fidelity Simulation of IEEE 802.11ad in ns-3 Using a Quasi-deterministic Channel Model," in *ACM Workshop on Next-Generation Wireless with ns-3 (WNGW 2019)*. ACM, 2019.

[30] M. Mezzavilla, S. Dutta, M. Zhang, M. R. Akdeniz, and S. Rangan, "5G mmWave Module for ns-3 Network Simulator," in *Proceedings of the 18th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '15. New York, NY, USA: ACM, 2015, pp. 283–290. [Online]. Available: http://arxiv.org/abs/1506.08801

[31] H. Assasa, A. Loch, and J. Widmer, "Packet mass transit: Improving frame aggregation in 60 GHz networks," in *WoWMoM 2016 - 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks*, 2016. [Online]. Available: https://doi.org/10.1109/WoWMoM.2016.7523522

[32] D. Skordoulis, Q. Ni, H. H. Chen, A. P. Stephens, C. Liu, and A. Jamalipour, "IEEE 802.11N MAC frame aggregation mechanisms for next-generation high-throughput WLANs," *IEEE Wireless Communications*, vol. 15, no. 1, pp. 40–47, 2008.

[33] K. Chandra, R. V. Prasad, and I. Niemegeers, "Performance Analysis of IEEE 802.11ad MAC Protocol," *IEEE Communications Letters*, vol. 21, no. 7, pp. 1513–1516, jul 2017.

[34] E. H. Ong, J. Kneckt, O. Alanen, Z. Chang, T. Huovinen, and T. Nihtila, "IEEE 802.11ac: Enhancements for very high throughput WLANs," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC*, 2011, pp. 849–853.

[35] A. Saif and M. Othman, "A reliable A-MSDU frame aggregation scheme in 802.11n wireless networks," *Procedia Computer Science*, vol. 21, pp. 191–198, 2013.

[36] A. Saif, M. Othman, S. Subramaniam, and N. A. W. A. Hamid, "An enhanced A-MSDU frame aggregation scheme for 802.11n wireless networks," *Wireless Personal Communications*, vol. 66, no. 4, pp. 683–706, oct 2012.

[37] A. Majeed and N. B. Abu-Ghazaleh, "Packet aggregation in multi-rate wireless LANs," in *Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks workshops*, vol. 1, 2012, pp. 452–460.

[38] J. Gross and O. Pu, "Multi-user OFDMA Frame Aggregation for," in *Ifip International Federation For Information Processing*, ser. Lecture Notes in Computer Science, 2009, vol. 5550, pp. 220–233.

[39] I. M.-m. Wlans, B. Bellalta, J. Barcelo, D. Staehle, A. Vinel, and M. Oliver, "On the Performance of Packet Aggregation," *IEEE Communications Letters*, vol. 16, no. 10, pp. 1–7, 2012.

[40] P. Dely and A. J. Kassler, "FUZPAG: A Fuzzy-Controlled Packet Aggregation Architecture for Wireless Mesh Networks," in *Performance Evaluation*, 2009, pp. 1–15.

[41] N. T. J. Bailey, "On Queueing Processes with Bulk Service," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 16, no. 1, pp. 80–87, 2018.

[42] J. H. Hong and K. Sohraby, "On modeling, analysis, and optimization of packet aggregation systems," *IEEE Transactions on Communications*, vol. 58, no. 2, pp. 660–668, 2010.

[43] M. F. Neuts, "The Busy Period of a Queue with Batch Service," *Operations Research*, vol. 13, no. 5, pp. 815–819, 2008.

[44] ——, "A General Class of Bulk Queues with Poisson Input," *The Annals of Mathematical Statistics*, vol. 38, no. 3, pp. 759–770, 2007.

[45] M. M. S. Hyo-Seong Lee, "Control Policies for the M^x/G/1 Queueing System," *Management Science*, vol. 35, no. 6, 1989.

[46] V. Makis, "Optimal Control of a Batch Service Queueing System With Bounded Waiting Time." *Kybernetika*, vol. 21, no. 4, pp. 262–271, 1985.

[47] A. Razi, A. Abediy, and A. Ephremides, "Delay minimization with channel-adaptive packetization policy for random data traffic," in *2014 48th Annual Conference on Information Sciences and Systems, CISS 2014*, 2014.

[48] Y. Lin, "WSN01-1: frame aggregation and optimal frame size adaptation for IEEE 802.11 n WLANs," in *Conference, 2006. GLOBECOM'06. IEEE*, 2006. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4151555

[49] S. Byeon, K. Yoon, O. Lee, and S. Choi, "MoFA : Mobility-aware Frame Aggregation in Wi-Fi," in *CoNext*, 2014.

[50] R. Mudumbai, S. Singh, and U. Madhow, "Medium access control for 60 GHz outdoor mesh networks with highly directional links," in *Proceedings - IEEE INFOCOM*, 2009, pp. 2871–2875.

[51] Y. Tian, K. Xu, and N. Ansari, "TCP in wireless environments: Problems and solutions," *IEEE Communications Magazine*, vol. 43, no. 3, 2005.

[52] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60GHz outdoor picocells," in *Proc. ACM Mobicom*, 2014, pp. 5–16.

[53] S. K. Saha, A. Garg, and D. Koutsonikolas, "A first look at TCP performance in indoor IEEE 802.11ad WLANs," in *Proceedings - IEEE INFOCOM*, vol. 2015-August, 2015, pp. 63–64.

[54] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, "A feasibility study of 60 GHz indoor WLANs," in *2016 25th International Conference on Computer Communications and Networks, ICCCN 2016*, 2016.

[55] "LEDE Project." [Online]. Available: https://lede-project.org/

[56] T. Li, D. Leith, and D. Malone, "Buffer sizing for 802.11-based networks," *IEEE/ACM Transactions on Networking*, vol. 19, no. 1, pp. 156–169, 2011.

[57] H. Assasa, S. K. Saha, A. Loch, D. Koutsonikolas, and J. Widmer, "Medium Access and Transport Protocol Aspects in Practical 802.11 ad Networks," in *19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoWMoM 2018*, 2018. [Online]. Available: https://doi.org/10.1109/WoWMoM.2018.8449795

[58] H. Assasa, G. Bielsa, S. K. Saha, P. Jimnez Mateo, A. Loch, D. Koutsonikolas, and J. Widmer, "Performance Analysis of Medium Access Control and Spatial Reuse for IEEE 802.11ad Deployments," *Pervasive and Mobile Computing*, 2019.

[59] TP-Link, "Talon AD7200 Multi-Band Wi-Fi Router," 2016. [Online]. Available: http://www.tp-link.com/us/products/details/cat-5506_AD7200.html

[60] Netgear, "Nighthawk X4 Smart WiFi Router," p. 1, 2014. [Online]. Available: http://netgear.co.uk/home/products/networking/wifi-routers/R7500.aspx

[61] "MikroTik wAP 60G AP." [Online]. Available: https://mikrotik.com/product/wap_60g_ap

[62] "Acer TravelMate P446-M." [Online]. Available: https://www.acer.com/ac/en/US/content/professional-series/travelmatep4

[63] D. Steinmetzer, D. Wegemer, and M. Hollick. (2017) Talon Tools: The Framework for Practical IEEE 802.11ad Research. [Online]. Available: https://seemoo.de/talon-tools/

[64] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking through Flexible Beams : A Link-Level Profiling," in *Acm Sigmetrics*, 2015, pp. 71–84.

[65] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan, "OpenMili: a 60 GHz software radio with a programmable phased-array antenna," in *Proceedings of the 22Nd Annual International Conference on Mobile Computing and Networking*, 2016, pp. 485–486.

[66] X. Tie, K. Ramachandran, and R. Mahindra, "On 60 GHz wireless link performance in indoor environments," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 7192 LNCS, 2012, pp. 147–157.

[67] G. Bielsa, A. Loch, I. Tejado, T. Nitsche, and J. Widmer, "60 GHz Networking: Mobility, Beamforming, and Frame Level Operation From Theory to Practice," *IEEE Transactions on Mobile Computing*, p. 1, 2018.

[68] S. K. Saha, T. Siddiqui, D. Koutsonikolas, A. Loch, J. Widmer, and R. Sridhar, "A detailed look into power consumption of commodity 60 GHz devices," in *18th IEEE International Symposium on A World of Wireless, Mobile and Multimedia Networks, WoWMoM 2017 - Conference*, 2017.

[69] Y. Zhu, X. Zhou, Z. Zhang, L. Zhou, A. Vahdat, B. Y. Zhao, and H. Zheng, "Cutting the cord: a robust wireless facilities network for data centers," in *Proceedings of the MOBICOM '14*, 2014, pp. 581–592.

[70] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, "Multi-Gigabit Indoor WLANs : Looking Beyond," in *2016 IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.

[71] X. Zhu, A. Doufexi, and T. Kocak, "Throughput and coverage performance for IEEE 802.11ad millimeter-wave WPANs," in *IEEE Vehicular Technology Conference*, may 2011, pp. 1–5.

[72] J. Qiao, X. Shen, J. W. Mark, Z. Shi, and N. Mohammadizadeh, "MAC-layer integration of multiple radio bands in indoor millimeter wave networks," in *IEEE Wireless Communications and Networking Conference, WCNC*, 2013, pp. 889–894.

[73] M. X. Gong, D. Akhmetov, R. Want, and S. Mao, "Multi-user operation in mmwave wireless networks," in *IEEE International Conference on Communications*, 2011.

[74] M. X. Gong, R. Stacey, D. Akhmetov, and S. Mao, "A directional CSMA/CA protocol for mmWave wireless PANs," in *IEEE Wireless Communications and Networking Conference, WCNC*, 2010.

[75] B. Ginzburg and A. Kesselman, "Performance analysis of A-MPDU and A-MSDU aggregation in IEEE 802.11n," in *2007 IEEE Sarnoff Symposium, SARNOFF*, 2007.

[76] B. Bellalta, J. Barcelo, D. Staehle, A. Vinel, and M. Oliver, "On the performance of packet aggregation in ieee 802.11ac mu-mimo wlans," *IEEE Communications Letters*, vol. 16, no. 10, pp. 1588–1591, October 2012.

[77] J. Palacios, D. Steinmetzer, A. Loch, M. Hollick, and J. Widmer, "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New York, NY, USA: ACM, 2018, pp. 241–255. [Online]. Available: http://doi.acm.org/10.1145/3241539.3241576

[78] S. Kutty and D. Sen, "Beamforming for Millimeter Wave Communications: An Inclusive Survey," *IEEE Communications Surveys and Tutorials*, vol. 18, no. 2, pp. 949–973, 2016.

[79] M. Maity, B. Raman, and M. Vutukuru, "TCP download performance in dense WiFi scenarios: Analysis and solution," in *IEEE Transactions on Mobile Computing*, vol. 16, no. 1, 2017, pp. 213–227.

[80] G. Kuriakose, S. Harsha, A. Kumar, and V. Sharma, "Analytical models for capacity estimation of IEEE 802.11 WLANs using DCF for internet applications," in *Wireless Networks*, vol. 15, no. 2, 2009, pp. 259–277.

[81] M. A. Ergin, K. Ramachandran, and M. Gruteser, "An experimental study of inter-cell interference effects on system performance in unplanned wireless LAN deployments," in *Computer Networks*, vol. 52, no. 14. Elsevier North-Holland, Inc., 2008, pp. 2728–2744.

[82] R. Bruno, M. Conti, and E. Gregori, "Modeling TCP Throughput Over Wireless LANs," in *Proc. 17th IMACS World Congress Scientific Computation, Applied Mathematics and Simulation*, 2005, pp. 11–15.

[83] M. Zhang, M. Mezzavilla, J. Zhu, S. Rangan, and S. S. Panwar, "The bufferbloat problem over intermittent multi-gbps mmwave links," *CoRR*, vol. abs/1611.02117, 2016. [Online]. Available: http://arxiv.org/abs/1611.02117

[84] VubIQ Inc. (2013) V60WGD03 60 GHz Waveguide Development System. [Online]. Available: http://www.pasternack.com/60-ghz-development-systems-category.aspx

[85] T. Eichler, "Challenges and Techniques for Characterizing Massive MIMO Antenna Systems for 5G," pp. 0–14, 2017.

[86] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, "Hybrid precoding for millimeter wave cellular systems with partial channel knowledge," *IEEE Journal of Selected Topics in Signal Processing*, pp. 1–5, 2013.

[87] J. Palacios, D. De Donno, D. Giustiniano, and J. Widmer, "Speeding up mmWave beam training through low-complexity hybrid transceivers," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC*, 2016.

[88] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick, "Fast and infuriating: Performance and pitfalls of 60 GHz WLANs based on consumer-grade hardware," in *2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking, SECON 2018*, 2018, pp. 1–9.

[89] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "Beamspy: Enabling robust 60 ghz links under blockage," in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, ser. NSDI'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 193–206. [Online]. Available: http://dl.acm.org/citation.cfm?id=2930611.2930625

[90] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Wireless Networks," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '17. New York, NY, USA: ACM, 2017, pp. 28–41. [Online]. Available: http://doi.acm.org/10.1145/3117811.3117817

[91] T. Wei and X. Zhang, "Pose Information Assisted 60 GHz Networks," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking - MobiCom '17*, ser. MobiCom '17. New York, NY, USA: ACM, 2017, pp. 42–55. [Online]. Available: http://dl.acm.org/citation.cfm?doid=3117811.3117832

[92] F. Zhou, M. Y. Naderi, K. Sankhe, and K. Chowdhury, "Making the Right Connections: Multi-AP Association and Flow Control in 60GHz Band," in *Proceedings - IEEE INFOCOM*, vol. 2018-April, 2018, pp. 1214–1222.

[93] G. Bielsal, J. Palacios, A. Loch, D. Steinmetzer, P. Casari, and J. Widmer, "Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points," in *Proceedings - IEEE INFOCOM*, vol. 2018-April, 2018, pp. 2384–2392.

[94] J. Palacios, G. Bielsa, P. Casaril, and J. Widmer, "Communication-Driven Localization and Mapping for Millimeter Wave Networks," in *Proceedings - IEEE INFOCOM*, vol. 2018-April, 2018, pp. 2402–2410.

[95] L. Simić, J. Arnold, M. Petrova, and P. Mähänen, "RadMAC," in *Proceedings of the 3rd Workshop on Hot Topics in Wireless*, ser. HotWireless '16. New York, NY, USA: ACM, 2016, pp. 61–65. [Online]. Available: http://doi.acm.org/10.1145/2980115.2980134

[96] T. Wei, A. Zhou, and X. Zhang, "Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors," in *14th {USENIX} Symposium on Networked Systems Design and Implementation, {NSDI} 2017, Boston, MA, USA, March 27-29, 2017*, ser. NSDI'17. Berkeley, CA, USA: USENIX Association, 2017, pp. 213–226.

[97] A. Shahmansoori, G. E. Garcia, G. Destino, G. Seco-Granados, and H. Wymeersch, "5G position and orientation estimation through millimeter wave MIMO," in *2015 IEEE Globecom Workshops, GC Wkshps 2015 - Proceedings*, dec 2015, pp. 1–6.

[98] A. Jafari, J. Sarrazin, D. Lautru, A. Benlarbi-delaï, L. Petrillo, and P. D. Doncker, "NLOS influence on 60 GHz indoor localization based on a new TDOA extraction approach," in *European Microwave Conference*, oct 2013, pp. 330–333.

[99] A. Olivier, G. Bielsa, I. Tejado, M. Zorzi, J. Widmer, and P. Casari, "Lightweight indoor localization for 60-GHz millimeter wave systems," in *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking, SECON 2016*, jun 2016, pp. 1–9.

[100] J. Chen, D. Steinmetzer, J. Classen, E. Knightly, and M. Hollick, "Pseudo lateration: Millimeter-wave localization using a single RF chain," in *IEEE Wireless Communications and Networking Conference, WCNC*, mar 2017, pp. 1–6.

[101] B. Cook, G. Buckberry, I. Scowcroft, J. Mitchell, and T. Allen, "Indoor Location Using Trilateration Characteristics," in *Proceedings of the London Communications Symposium*, no. 1, 2005, pp. 2–5. [Online]. Available: http://discovery.ucl.ac.uk/136687/

[102] D. Vasisht, S. Kumar, and D. Katabi, "Decimeter-Level Localization with a Single WiFi Access Point," in *Nsdi*, vol. 16, 2016, pp. 165–178. [Online]. Available: https://www.usenix.org/system/files/conference/nsdi16/nsdi16-paper-vasisht.pdf

[103] M. Youssef, M. Mah, and A. Agrawala, "Challenges in Device-free passive localization for wireless environments," in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking - MobiCom '07*, no. January. ACM, 2007, p. 222. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1287853.1287880

[104] K. Chetty, K. Woodbridge, G. E. Smith, H. Guo, \textbfAl-Ashwal, W.A., and M. Ash, "Through-Wall Sensing of Personnel at Standoff Distances using Passive Bistatic {WiFi} Radar," *Proc. American Electromagnetics Conference*, vol. 50, no. 4, pp. 1218–1226, 2010.

[105] S. D. Domenico, M. D. Sanctis, E. Cianca, and M. Ruggieri, "WiFi-based through-the-wall presence detection of stationary and moving humans analyzing the doppler spectrum," *IEEE Aerospace and Electronic Systems Magazine*, vol. 33, no. 5-6, pp. 14–19, may 2018.

[106] IEEE, "IEEE Std 802.11ad-2012. Enhancements for Very High Throughput in the 60 GHz Band," *IEEE 802.11 Working Group*, 2012.

[107] D. De Donno, J. P. Beltrán, D. Giustiniano, and J. Widmer, "Hybrid analog-digital beam training for mmWave systems with low-resolution RF phase shifters," in *2016 IEEE International Conference on Communications Workshops, ICC 2016*, may 2016, pp. 700–705.

[108] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Tool Release: Gathering 802.11n Traces with Channel State Information," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 1, 2011, p. 53.

[109] Y. Wang and Z. Shi, "Millimeter-wave mobile communications," in *5G Mobile Communications*, ser. HotNets '16, 2016, pp. 117–134. [Online]. Available: http://dl.acm.org/citation.cfm?doid=3005745.3005766

[110] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmWave Path Tracking," in *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications*, 2017, pp. 13–18.

[111] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: Mm-Wave beam steering without in-band measurement," in *Proceedings - IEEE INFOCOM*, vol. 26, 2015, pp. 2416–2424.

[112] A. Eitan and C. Cordeiro, *Short SSW Format for 11ay (IEEE 802.11-16/0416-01-00)*, 2016.

[113] M. K. Haider and E. W. Knightly, "Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs," in *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2016, pp. 61–70.

[114] A. Patra, L. Simić, and M. Petrova, "Experimental evaluation of a novel fast beamsteering algorithm for link re-establishment in mm-wave indoor WLANs," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC*, sep 2016, pp. 1–7.

[115] A. Loch, H. Assasa, J. Palacios, J. Widmer, H. Suys, and B. 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*, ser. CoNEXT '17. New York, NY, USA: ACM, 2017, pp. 224–237. [Online]. Available: http://doi.acm.org/10.1145/3143361.3143395

[116] R. Bonjour, M. Singleton, S. A. Gebrewold, Y. Salamin, F. C. Abrecht, B. Baeuerle, A. Josten, P. Leuchtmann, C. Hafner, and J. Leuthold, "Ultra-fast millimeter wave beam steering," *IEEE Journal of Quantum Electronics*, vol. 52, no. 1, 2016.

[117] R. Bonjour, M. Burla, F. C. Abrecht, S. Welschen, C. Hoessbacher, W. Heni, S. A. Gebrewold, B. Baeuerle, A. Josten, Y. Salamin, C. Haffner, P. V. Johnston, D. L. Elder, P. Leuchtmann, D. Hillerkuss, Y. Fedoryshyn, L. R. Dalton, C. Hafner, and J. Leuthold, "Plasmonic phased array feeder enabling ultra-fast beam steering at millimeter waves," in *Optics Express*, vol. 24, no. 22, 2016, p. 25608.

[118] J. Lien, N. Gillian, M. E. Karagozler, P. Amihood, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, "Soli: Ubiquitous Gesture Sensing with Millimeter Wave Radar," *ACM Trans. Graph*, vol. 35, no. July, p. 142, jul 2016. [Online]. Available: http://doi.acm.org/10.1145/2897824.2925953.

[119] T. Wei and X. Zhang, "mTrack: High-Precision Passive Tracking Using Milimeter Wave Radios," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (ACM MobiCom '15)*, 2015, pp. 117–129. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2789168.2790113

[120] L. Simi, J. Arnold, M. Petrova, and P. Mähönen, "RadMAC : Radar-Enabled Link Obstruction Avoidance for Agile mm-Wave Beamsteering," in *HotWireless'16*, 2016, pp. 61–65.

[121] J. Arnold, L. Simić, M. Petrova, and P. Mähönen, "Radar-enhanced Mm-wave Agile Beam-steering: Demo," in *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*, 2016.

[122] J. Xiong and K. Jamieson, "ArrayTrack : A Fine-Grained Indoor Location System," in *10th USENIX Symposium on Networked Systems Design and Implementation*, no. 279976, 2013, pp. 71–84.

[123] S. Bellofiore, J. Foutz, C. Balanis, and A. Spanias, *Smart antennas for wireless communications*. McGraw-Hill Professional, 2002.

[124] H. Jiang, D. Fang, J. Xiong, J. Wang, and X. Chen, "D-Watch: Embracing "Bad" Multipaths for Device-Free Localization With COTS RFID Devices," in *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, 2017, pp. 3559–3572.

[125] A. Patra, L. Simić, and M. Petrova, "Design and experimental evaluation of a 2.4 ghz-aoa-enhanced beamsteering algorithm for ieee 802.11ad mm-wave wlans," in *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2017.

[126] M.-g. W. P. A. N. Systems, K. Hosoya, N. Prasad, K. Ramachandran, N. Orihashi, and S. Kishimoto, "Multiple Sector ID Capture (MIDC): A Novel Beamforming Technique for 60-GHz Band," *IEEE Transactions on Antennas and Propagation*, vol. 63, no. 1, pp. 81–96, 2015.

[127] T. Baykas, M. Rahman, R. Funada, F. Kojima, H. Harada, and S. Kato, "Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems," *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 8, pp. 1390–1399, 2009. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5262295

[128] Z. Marzi, D. Ramasamy, and U. Madhow, "Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells," *IEEE Journal on Selected Topics in Signal Processing*, vol. 10, no. 3, pp. 514–527, 2016.

[129] A. Patra, L. Simić, and P. Mähönen, "Smart mm-Wave Beam Steering Algorithm for Fast Link Re-Establishment under Node Mobility in 60 GHz Indoor WLANs," in *Proceedings of the 13th ACM International Symposium on Mobility Management and Wireless Access*, ser. MobiWac '15, 2015, pp. 53–62.

[130] J. Arnold, L. Simic, M. Petrova, P. Mähönen, L. Simić, M. Petrova, and P. Mähönen, "Demo: Spectrum-agile mm-wave packet radio implementation on USRPs," in *2015 Workshop on Software Radio Implementation Forum, SRIF 2015*, 2015, pp. 5–8. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2801676.2801681

[131] O. Abari, H. Hassanieh, M. Rodreguiz, and D. Katabi, "A millimeter wave software defined radio platform with phased arrays," in *Proceedings of the 22Nd Annual International Conference on Mobile Computing and Networking*, 2016, pp. 419–420.

[132] P. Zetterberg and R. Fardi, "Open source SDR frontend and measurements for 60-GHz wireless experimentation," *IEEE Access*, vol. 3, pp. 445–456, 2015.

[133] G. F. Riley and T. R. Henderson, *The ns-3 network simulator*. Springer Berlin Heidelberg, 2016, pp. 15–34.

[134] H. Tazaki, "Direct Code Execution : Revisiting Library OS Architecture for Reproducible Network Experiments Our target : experimentation reproducibility Ideally one should be able to easily," in *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT'13)*, 2013.

[135] H. Sangtae, R. Injong, and X. Lisong, "CUBIC: a new TCP-friendly high-speed TCP variant," *SIGOPS Oper. Syst. Rev.*, vol. 42, no. 5, pp. 64–74, 2008.

[136] E. Neufert, *Architects' Data*. Lockwood, 1970.

[137] Y. M. Tsang, A. S. Poon, and S. Addepalli, "Coding the beams: Improving beamforming training in mmWave communication system," in *GLOBECOM - IEEE Global Telecommunications Conference*, 2011.