# Efficient Network Control for Large and Highly Dense Millimeter Wave Deployments

by

Nina Grosheva

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

Telematic Engineering

Universidad Carlos III de Madrid

Advisor: Joerg Widmer

January 2024

*Efficient Network Control for Large and Highly Dense Millimeter Wave Deployments*

Prepared by:
Nina Grosheva, IMDEA Networks Institute, Universidad Carlos III de Madrid
contact: nina.grosheva@imdea.org

Under the advice of:
Joerg Widmer, IMDEA Networks Institute
Telematic Engineering Department, Universidad Carlos III de Madrid

This work has been supported by:

# Acknowledgements

I would first like to thank my supervisor Dr. Joerg Widmer for his constant and consistent support, guidance and mentorship during my PhD. Thank you for your patience, kindness and invaluable advice. I feel very fortunate to have gotten the opportunity to learn from you over the past years. I would also like to thank all my collaborators who helped me grow and learn and become a better researcher. Dr. Jesus Omar Lacruz, for always having an open door for discussions, advice and a helping hand. And most of all, for being a wonderful friend and an incredible support. Dr. Hany Assasa, for being a kind and patient guide through the dark maze of ns-3 and for all the help and encouragement. Dr. Tanguy Ropitault, for enduring my long emails and ns-3 bugs and always taking the time to provide valuable feedback and help. And Dr. Sofie Pollin who supervised me during my (unfortunately virtual) internship at KU Leuven and taught me a lot about how to think about research problems and come up with interesting and practical solutions.

I would also like to thank all my colleagues at IMDEA Networks who made my time there happy and joyful. The PhD is an extremely challenging experience, and it would've been impossible to accomplish without the companionship and community I found at IMDEA. Thank you for the very long lunches and coffees and for going through this journey with me. Special gratitude goes to Dolo, Yago, Rafa, Serly, Sai, Gaetano, Alejandro and Alex. I would also like to thank all my fellow ESRs from the MINTS project for the best Network Wide Events and the chance to be a part of such a great team. Extra special thanks go to my extra special blog team, Neharika and Enver, with whom I shared the best conversations. Next, I would like to thank all my close friends who supported me throughout the many challenges of the past years and who fill my life with laughter and love. To Oli and Tamara, thank you for the endless Skype calls and always too short trips home. Your friendship keeps me grounded through all the changes and challenges that life brings. To my favourite Germans, thank you for the best holidays and trips that filled me up with positive energy and gave me the strength to keep going. I'm so happy I got to keep you even though I ran away to sunny Spain. And to my newest, Madrid family, thank you for being there to celebrate the good days and make the dark days easier. You made Madrid feel like home. Extra special thanks go to Paloma, Dolo,

and Mustapha who in different ways all supported me through the most difficult times.

And finally, I would mostly like to thank my family, who have always been the greatest support for all my accomplishments. Thank you to my brother Milan for sharing this experience with me and making me never feel alone. I am very proud to be known as your sister and to follow in your footsteps as the second Dr. in our family. And thank you to my parents for their unconditional support and love through all my life. My dad, who has always believed that I can accomplish anything I want and helped me believe it too. And my mom, who has always been an inspiration and a role model, but none more so than in the last few years. Thank you for being here. I love you.

# Published Content

This thesis is based on the following published papers:

**[1]** Hany Assasa **Nina Grosheva**, Tanguy Ropitault, Steve Blandino, Nada Golmie, Joerg Widmer. Implementation and Evaluation of a WLAN IEEE 802.11ay Model in Network Simulator ns-3. Published in *Proceedings of the Workshop on ns-3 (WNS3 '21)*. Association for Computing Machinery, New York, NY, USA, 9-16. https://doi.org/10.1145/3460797.3460799

– This work is fully included and its content is reported in Chapter 2 and Chapter 5.

– The author's role in this work is focused on the design, implementation and evaluation of the mmWave module from the paper and the writing of the paper.

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

**[2]** **Nina Grosheva**, Hany Assasa, Tanguy Ropitault, Pablo Jimenez Mateo, Joerg Widmer, Nada Golmie. A Comprehensive Analysis and Performance Enhancements for the IEEE 802.11ay Group Beamforming Protocol. Published in *the 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (IEEE WoWMoM 2022))*. Belfast, United Kingdom, 2022, pp. 194-203.https://doi:10.1109/WoWMoM54355.2022.00038

– This work is fully included and its content is reported in Chapter 3.

– The author's role in this work is focused on the design, implementation and evaluation of the beamforming protocol from the paper and the writing of the paper.

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

**[3]** **Nina Grosheva**, Sai Pavan Deram, Jesus O. Lacruz, Joerg Widmer. SIGNiPHY: Reconciling random access with directional reception for efficient mmWave WLANs. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services (MobiSys '23)*. Association for Computing Machinery, New York, NY, USA, 150-162. https://doi.org/10.1145/3581791.3596860

– This work is partially included and its content is reported in Chapter 4.

– The author's role in this work is focused on the protocol design, simulation implementation and evaluation of the proposed system and the writing of the paper.

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

**[4] Nina Grosheva**, Rizqi Hersyandika, Joerg Widmer, Sofie Pollin. In-Band Multi-Connectivity with Local Beamtraining for Improving mmWave Network Resilience. In *Proceedings of The 26th International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM'23).* Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3616388.3617529

– This work is fully included and its content is reported in Chapter 6.

– The author's role in this work is focused on the design, implementation and evaluation of the proposed system and the writing of the paper.

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

# Abstract

Wireless networks have become an integral part of modern society, providing ubiquitous connectivity to a growing number of connected devices. Concepts like Augmented Reality (AR)/Virtual Reality (VR), remote surgery and Industry 4.0 will further increase the number of users and the volume of data being transferred. Satisfying the demands for higher throughput, lower latency and higher reliability necessitates novel technologies and innovative designs. Operation in the high frequency Millimeter-Wave (mmWave) band is foreseen as a crucial part of the design of future wireless networks. The extremely large signal bandwidth offered at mmWave frequencies enables multi-Gbps, low-latency wireless connectivity for a peak performance that far exceeds what can be achieved in the currently used sub-6 GHz bands.

Realising the potential of mmWave technology requires adaptation to the challenging propagation environment at all levels of the protocol stack. Significant work has already been done to enable single-link communication through the use of phased antenna arrays to generate narrow directional beampatterns. Network aspects and interactions in large and dense networks, however, remain largely unexplored. The goal of this thesis is to study the performance of mmWave protocols in dense deployments with many Access Points (APs) and Station (STA) near each other. Such deployments are required for sufficient coverage in real-world implementations, however, they come with unique challenges due to the complex nature of interference in mmWave networks. The thesis studies different proposed architectures for mmWave to gain insight into sources of inefficiency and performance degradation. We then propose solutions to address these challenges and enhance the operation of mmWave networks.

To enable research into dense mmWave networks we implemented the latest mmWave WiFi standard, IEEE 802.11ay, in the network simulator ns-3. We used as a basis for our implementation the existing IEEE 802.11ad model, expanding it to cover advanced features like Multiple-Input and Multiple-Output (MIMO), channel bonding and novel Beamforming Training (BFT) protocols introduced in IEEE 802.11ay. Using our model were able to get in-depth insights regarding the performance of various protocol features of the state-of-the-art mmWave WiFi, including, channel access, BFT, interference management and spatial sharing.

We first focus on BFT scalability in dense deployments, looking at how the accuracy of the training can degrade in high-interference environments, as well as how the growing overhead can limit communication throughput. We propose the use of the novel Group Beamforming protocol introduced in IEEE 802.11ay as it enables simultaneous training of all STA with a Basic Service Set (BSS), rather than relying on a per AP-STA training like legacy BFT from IEEE 802.11ad. We additionally propose performance enhancements for the Group Beamforming protocol that can increase the accuracy to ensure correct beampattern selection even under significant interference. Our analysis demonstrates that the modified Group Beamforming protocol has higher accuracy than legacy BFT and enables higher network throughput due to the reductions in overhead.

We then designed a physical (PHY) layer signalling solution that enhances packet reception in mmWave WiFi devices. We focused on two sources of inefficiency caused by the contention-based random channel access - the use of omnidirectional receiver beampatterns, and the overhearing of unwanted packets. Both of these problems limit the performance in the network and affect spatial re-use. SIGNaling in the PHY Preamble (SIGNiPHY) embeds the user identifier (ID) in the PHY packet preamble, allowing for early user identification. This enables the receiver to use the correct directional beampattern to receive the packet payload, as well as to filter any packets for which it is not the recipient. Thus, SIGNiPHY increases the resilience to interference, enabling packet decoding under challenging conditions and increasing spatial sharing. We evaluated SIGNiPHY in ns-3, as well as an FPGA testbed, revealing significant gains in throughput, latency and fairness.

The next work in the thesis presented our mmWave MIMO implementation with standard compliant MIMO BFT protocols and channel access. We demonstrate how our analog BFT protocol was able to train multiple transmit and receive antennas to find spatially separated, independent streams. Challenges with mobility, the sparsity of the mmWave channel and complex BFT protocols require further research into mmWave MIMO. However, we found promising results regarding the viability of mmWave MIMO even with a fully analog architecture.

We further investigate an alternative architecture for devices with multiple Radio Frequency (RF) chains by introducing multi-connectivity. In multi-connectivity networks, users maintain several simultaneous links with spatially distributed APs. Unlike MIMO networks, multi-connectivity designs aim to not only increase throughput but also enhance resilience and robustness. This makes them extremely suitable for mmWave networks which suffer from frequent outages and service interruptions. Furthermore, mmWave multi-connectivity networks can have reduced implementation complexity by exploiting the spatial separation of the directional links. Therefore, we propose a distributed multi-connectivity design that relies solely on local analog beamforming for interference management. Our architecture was able to enhance resilience and maintain connectivity

at all times even under high interference, as well as exploit the spatial diversity of the multiple links to achieve gains in throughput.

Finally, we study the novel IEEE 802.11bf protocol which aims to standardize sensing operation in WiFi. As a topic of significant interest from both academia and industry, environmental sensing using communication signals opens new possibilities for mmWave networks. We present a first initial system-level study that looks at joint communication and sensing in a mmWave WiFi network. We look at resource allocation and study how the sensing and data traffic interact with each other. We further analyse how the sensing parameters affect the performance and identify network configurations where both sensing and communication can coexist, enabling successful integration of sensing and communication in a single system.

To conclude, in this thesis we present a comprehensive analysis of dense mmWave networks, proposing performance enhancements to enhance scalability and efficiency. We then look at future possibilities for mmWave, by analysing the possibilities of advanced devices with multiple RF chains, as well as novel paradigms that integrate environmental sensing into mmWave WiFi operation.

# Table of Contents

# List of Tables

# List of Figures

# List of Acronyms

**A-BFT** Association Beamforming Training

**AGC** Automatic Gain Control

**AID** Association Identifier

**A-MPDU** Aggregate MAC Protocol Data Unit

**ATI** Announcement Transmission Interval

**CoMP** Coordinated Multi-Point

**AP** Access Point

**AR** Augmented Reality

**AWV** Antenna Weight Vector

**BER** Bit Error Rate

**BFT** Beamforming Training

**BI** Beacon Interval

**BP** Beam Pattern

**BPSK** Binary Phase-Shift Keying

**BRP** Beam Refinement Protocol

**BRP TXSS** BRP Transmit Sector Sweep (TXSS)

**BSS** Basic Service Set

**BTI** Beacon Transmission Interval

**CBAP** Contention-Based Access Period

**CDF** Cumulative Distribution Function

**CEF** Channel Estimation Field

**CFO** Carrier Frequency Offset

**CIR** channel Impulse Response

**COTS** Commercial Off-the-Shelf

**CSI** Channel State Information

**CSMA/CA** Carrier Sense Multiple Access with Collision Avoidance

**CTS** Clear-to-Send

**DDIR** Double Directional channel Impulse Response (CIR)

**DMG** Directional Multi-Gigabit

**DOA** Direction of Arrival

**DOD** Direction of Departure

**DL** downlink

**DSC** Dynamic Sensitivity Control

**DTI** Data Transmission Interval

**EDMG** Enhanced Directional Multi-Gigabit

**FST** Fast Session Transfer

**FWA** Fixed Wireless Access

**GI** Guard Interval

**ID** identifier

**IoT** Internet Of Things

**L-CEF** Legacy-Channel Estimation Field (CEF)

**LOS** Line-Of-Sight

**L-STF** Legacy-Short Training Field (STF)

**LUT** lookup table

**MAC** Medium Access Control

**MCS** Modulation and Coding Scheme

**MIMO** Multiple-Input and Multiple-Output

**mmWave** Millimeter-Wave

**MU** Multi User

**MPC** Multipath Component

**MU-MIMO** Multi User MIMO

**NAV** Network Allocation Vector

**NIST** National Institute of Standards and Technology

**NLOS** Non-Line of Sight

**OFDM** Orthogonal Frequency Division Multiplexing

**PAA** Phased Antenna Array

**PER** Packet Error Rate

**PHY** physical

**PLCP** Physical Layer Convergence Protocol

**PPDU** PLCP Protocol Data Unit

**PSR** Packet Success Rate

**Q-D** Quasi-Deterministic

**RAA** Rate Adaptation Algorithm

**RAT** Radio Access Technology

**RF** Radio Frequency

**RSS** Responder Sector Sweep

**RTS** Ready-to-Send

**RX** receive

**SC** Single Carrier

**SIGNiPHY** SIGNaling in the PHY Preamble

**SINR** Signal-to-Interference-plus-Noise Ratio

**SISO** Single-Input and Single-Output

**SLS** Sector Level Sweep

**SNR** Signal-to-Noise Ratio

**SP** Service Period

**SSW** Sector Sweep

**STA** Station

**STF** Short Training Field

**SU** Single User

**SU-MIMO** Single User MIMO

**TCP** Transmission Control Protocol

**TDMA** Time-Division Multiple Access

**TGbf** Task Group IEEE 802.11bf

**TRN** Training

**TRN-R** Receive Training

**TRN-T** Transmit Training

**TRN-R/T** Receive/Transmit training

**TX** transmit

**TXSS** Transmit Sector Sweep

**UDP** User Datagram Protocol

**UL** uplink

**URA** Uniform Rectangular Array

**VR** Virtual Reality

**WLAN** Wireless Local Area Network

# 1 Introduction

The past decade has seen an exponential rise in wireless traffic as both the number of connected devices and the throughput demands have increased multifold. Emerging applications like Augmented Reality (AR)/Virtual Reality (VR), remote surgery, vehicular connectivity, wireless backhauling, Fixed Wireless Access (FWA) and factory automation will even further increase the demands from wireless networks [5]. The need for higher data rates, lower latency, and higher reliability demands innovative designs and an evolution of current wireless networks, as existing technologies are becoming over-congested and unable to satisfy the ever-growing demands for wireless communication [6]. One key issue is the scarcity of available bandwidth in the sub-6 GHz bands where both cellular and WiFi systems, as well variety of other technologies like bluetooth and Zigbee are currently deployed. Advanced signal processing techniques like Multiple-Input and Multiple-Output (MIMO) as well as network densification have been used to increase capacity and throughput, however, there are fundamental limits to the amount of data that can be transmitted in the bandwidth currently used. Therefore, a shift towards the higher frequency bands has been proposed as a key enabling technology for 5G and future 6G systems [7]. While communication up to THz frequencies is being considered, the Millimeter-Wave (mmWave) band (between 30 GHz and 300 GHz) is the most promising and well-researched high-frequency technology for timely real-world deployment.

mmWave systems can offer very high bandwidth for communication due to the large amounts of spectrum available in this frequency band, leading to the ability to support multi-Gbps communication even with a Single-Input and Single-Output (SISO) link [8]. In addition, they can satisfy the low latency requirements of emerging real-time applications like AR/VR and factory automation. Furthermore, the high signal bandwidth also enables mmWave to support features like high accuracy localization and environment sensing as it enables much higher resolution than current sub-6 GHz systems [9]. The main obstacle to the broad application of mmWave systems lies in coping with the challenging propagation environment for high-frequency signals, as they suffer from high path loss, oxygen absorption and sensitivity to blockage. This leads to low communication ranges

and frequent link breakages, making it difficult to establish stable links and provide uninterrupted service to users [10, 11]. Significant research effort has already gone into designing solutions to enable communication at mmWave frequencies, resulting in the development of Commercial Off-the-Shelf (COTS) devices and the integration of mmWave technology in the latest cellular and WiFi standards.

The key modification is the switch to directional communication by means of phased antenna arrays. Unlike current sub-6 GHz systems where antennas are usually omnidirectional and radiate energy evenly in all directions, mmWave networks require beamforming to focus the signal in the direction of the peer device. This generates narrow directional beams that allow for higher transmission and reception antenna gains to combat the high propagation loss, thus enabling communication at longer ranges. However, the handling of highly directional mmWave links introduces unique challenges to mmWave networks. First, before communication can take place, it is necessary to perform a Beamforming Training (BFT) procedure to find the optimal beam for communication with every communication peer. Additionally, it is necessary to continuously monitor and update the chosen beam to prevent misalignment caused by mobility. Finally, environmental changes and mobility can introduce blockages and fast link degradation, requiring the use of speedy recovery mechanisms to ensure uninterrupted service. Furthermore, directionality requires modifications to many procedures across the protocol stack designed with the assumption of omnidirectional reach of signals. Initial access [12], association [13] and handover [14] procedures all require efficient and low-overhead solutions to reduce the complexity and increase the speed. Interference management and channel access also differ significantly from traditional networks due to the highly dynamic nature of the interference which is localized in certain directions and varies not only with the location and orientation of devices but also with the transmit and receive beampatterns. On the one hand, spatial sharing allows multiple directional beams sufficiently separated in space to operate simultaneously as they do not interfere with one another, leading to high aggregate network throughput [15]. On the other hand, random channel access protocols like Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), which represent a foundational aspect of WiFi, rely on accurate sensing of the occupation of the wireless medium and are negatively affected by the directional, low-range nature of mmWave signals [16–18]. Practical research has shown that adaptation to mmWave frequencies will require modification of everything from the aggregation policies [19], to rate adaption mechanisms [16] and Transmission Control Protocol (TCP) variants [17, 20]. mmWave networks with their unique propagation characteristics, directional communication and extremely high data rates fundamentally differ from traditional sub-6 GHz networks in many ways. Thus, they require innovative solutions to realize their full potential and provide high-performance wireless connectivity for the next generation of applications.

## 1.1 Standardization and COTS development

mmWave technology has been included in the latest cellular and WiFi standards, leading to the development of the first COTS devices with mmWave support and the first real-world deployments. This has significantly progressed research by studying more complex, realistic scenarios and complete systems with a full protocol stack.

The IEEE 802.11ad standard [21], introduced in 2012, is the first Wireless Local Area Network (WLAN) standard to provide Medium Access Control (MAC) and physical (PHY) layer specifications for wireless networking in the unlicensed 60 GHz band, referred to as Directional Multi-Gigabit (DMG) operation. It represents a technical achievement and a large step forward for mmWave development as it standardized for the first time procedures for directional communication including initial access, BFT and medium access. New procedures such as relaying and fast session transfer to the sub-6 GHz band address challenges due to the reduced range and unreliability of mmWave links. Finally, a very robust control PHY layer is designed for transmitting control messages at very low Signal-to-Interference-plus-Noise Ratio (SINR) before BFT has been completed. IEEE 802.11ad allows for data rates of up to 6.72 Gbps with the highest Modulation and Coding Scheme (MCS), using the highest performing Orthogonal Frequency Division Multiplexing (OFDM) modulation. This is enabled by the extremely wide bandwidth of 2.16 GHz for each WiFi channel, with 4 channels available from 57 GHz to 64 GHz. The approval of IEEE 802.11ad led to the development of a diverse range of commercial products that support mmWave WiFi, including routers [22–24], docking stations [25] and even mobile phones [26]. This allowed for practical evaluation of mmWave WiFi on COTS devices, leading to great insight regarding mmWave performance, both in terms of real-world propagation and protocol challenges that arise from the specific nature of mmWave links.

While incredibly valuable, IEEE 802.11ad failed to fully exploit the vast capabilities of the 60 GHz band and offer support to the highest-performing applications [7]. This led to the development of its successor, IEEE 802.11ay [27], which defined Enhanced Directional Multi-Gigabit (EDMG) operation in the 60 GHz band. IEEE 802.11ay was designed to be fully compatible with IEEE 802.11ad but to offer much higher performance, up to 100 Gbps. Most of the performance enhancements are accomplished through the inclusion of advanced physical layer solutions like MIMO and channel aggregation and bonding, allowing for multifold increases in throughput. Both Single User (SU) and Multi User MIMO (MU-MIMO) with a maximum of 8 spatial streams are supported by IEEE 802.11ay, as well as aggregation or bonding of up to 4 channels. In addition, IEEE 802.11ay significantly expands the BFT procedures supported to improve the scalability and support different deployments and heterogeneous end-devices. This includes the inclusion of new protocols for MIMO BFT where multiple Radio Frequency (RF)

chains need to be simultaneously trained, the redesign of existing IEEE 802.11ad BFT mechanisms and the addition of new SISO BFT mechanisms that address different use cases. IEEE 802.11ay was only recently approved in 2021 with the first devices emerging recently [28]. The inclusion of many new features offers exciting opportunities for research into the next generation of mmWave networks, particularly as concepts like mmWave MIMO have not yet been thoroughly evaluated.

Finally, 3GPP has included operation in the high-frequency bands (>24 GHz) in the specifications for 5G New Radio (NR) [29] to provide support for ultra-high throughput and capacity performance. Key aspects of the standardization process included the design of antenna beamforming and beam management procedures, as well as dealing with the interoperability between the mmWave Radio Access Network (RAN) and other parts of the 5G network [30]. Deployment of 5G mmWave NR has recently started and the first real-world studies are analysing aspects like coverage, beam management and throughput performance [31]. While many of the research challenges considered in this thesis also apply to 5G NR, we focus on evaluating mmWave WiFi due to the lower protocol complexity, easier implementation and greater accessibility and flexibility.

## 1.2 Research challenges

mmWave technology is poised at the moment to enter widespread deployment and application with sufficient standardization support and availability of commercial devices that support it. However, open questions still remain regarding the stability, robustness and resilience of mmWave networks. This is particularly true when considering network aspects like coverage, interference, channel access and spatial sharing. While single-link deployments have been extensively studied, interactions in large and dense networks still pose a lot of challenges. However, real-world implementations will require dense deployments with many Access Points (APs) and users in close proximity to ensure sufficient coverage. Performance evaluations of current multi-AP deployments [16–18, 32–34] show that the throughput and reliability drop significantly as the network density grows. Understanding the causes of these performance trends and improving the efficiency and robustness of large mmWave deployments is crucial to the further development of mmWave technology. In the following, we summarize some of the key research challenges for large and dense mmWave networks.

- **Research tools:** The study of dense mmWave networks can be quite challenging, as few tools exist that provide reliable and comprehensive performance evaluation. Commercial devices do not implement all features of the mmWave WiFi protocols [35] and in particular, do not support advanced capabilities like MIMO, fast beamtraining and channel aggregation which are particularly relevant for research. In addition, they offer limited interfaces to the lower protocol layers

and only minor modifications to the operation are possible, making it difficult to evaluate novel solutions. FPGA-based testbeds [36] allow for flexibility and insight into the performance, however, their high cost makes them unsuitable for dense network evaluation. Furthermore, they often lack fidelity in the implementation of higher protocol layers and therefore, cannot provide insight into network aspects. Therefore, network simulation tools can be very valuable for large-scale, detailed evaluation and fast implementation of new features. Unlike signal processing simulation tools, network simulators have implementations of the full protocol stack and allow exploring aspects like channel access, packet collisions and application performance. The IEEE 802.11ad module [37–39] of the network simulator ns-3 represented the first open-source model for mmWave WiFi in ns-3. The model features a detailed implementation of the novel MAC features of IEEE 802.11ad including BFT, initial access, relaying and fast session service. Furthermore, it has a high-fidelity PHY layer as it includes support for interfacing with commercially available ray tracers, as well as realistic antenna models with directional radiation patterns. This made it an extremely valuable tool to the research community and allowed for interesting insight into mmWave networks. In this thesis, we present the extension of the model to support the latest advancements in mmWave technology introduced in the IEEE 802.11ay standard. In particular, support for MIMO required significant updates to ns-3 operation and represents the first implementation of mmWave MIMO operation from a protocol perspective.

- **Scalability and efficiency of beam training/tracking algorithms:** Beampattern selection is a crucial factor in mmWave performance as high antenna gains are fundamental to overcoming the large path loss in the high-frequency band [40, 41]. The BFT process used to find the optimal beampattern for communication is particularly challenging in dense deployments. The training overhead grows with the number of devices in the network, leading to reduced network efficiency and data rates. In addition, the high-interference environment can corrupt the training and lead to erroneous beampattern selection which will in turn reduce communication efficiency. Finally, the dynamic nature of dense networks might require quick beampattern adaptation which requires fast, low-overhead training. Identifying BFT algorithms that scale well in dense networks and maintain a high accuracy can lead to significant performance enhancements.

- **Resilience and robustness:** The increased attenuation, sensitivity to blockage and need for directionality lead to problems with low coverage ranges, unstable links and reduced reliability in mmWave networks [10, 11]. Ensuring uninterrupted and robust service remains a key challenge in current mmWave systems. In addition to outages caused by mobility or blockages, dense networks

additionally suffer from high interference which can also disrupt data transmission and reception. Moreover, link recovery algorithms need to be scalable to avoid high overhead and latency. The handover and beam management processes also suffer from increased complexity in dense networks and could be disturbed by high interference. All of these challenges make dense and large mmWave networks particularly vulnerable to link breakages and loss of service, requiring innovative solutions to design robust systems.

- **Interference management and increased spatial re-use:** The use of narrow directional beams in mmWave enables high spatial re-use by allowing multiple links located in close proximity to be simultaneously active as the interference they create is limited to a certain spatial direction. This is crucial to enabling high data rates in large networks which require dense AP deployment for coverage. However, creating and identifying highly directional links which have full spatial separation so as to not interfere with each other is not always achieved in current networks, leading to low spatial re-use and low network throughput [16, 17]. In addition, dense networks suffer from high interference and packet collisions, further reducing the efficiency. In particular, the dynamic interference environment caused by directional transmissions can lower the reliability of the carrier sensing during channel access. This leads to packet collisions, errors in BFT and higher unreliability of data transmissions. Interference management and spatial re-use are closely linked in mmWave networks and pose many unique challenges. For example, unreliable carrier sensing leads to packet collisions, however, efforts to improve channel access can unintentionally lead to reduced spatial re-use. Interactions between devices in the network are highly complex as the interference varies with both the beampattern selection (which itself can be influenced by interference during the BFT) and the location and orientation of the devices. Exploring the dynamics in dense networks and understanding the performance trade-offs is key to designing efficient mmWave networks.

- **Multi-RF deployments:** The latest mmWave WiFi standard supports MIMO operation, requiring advanced devices with multiple RF chains. This enables communication with multiple spatial streams for a multifold increase in traffic. MIMO operation at mmWave frequencies, however, is very challenging. The wide bandwidth exacerbates linear and non-linear impairments in the RF devices [42]. The high number of antenna elements prohibits a fully digital precoding architecture, instead relying on analog or hybrid beamforming [42]. BFT protocols for MIMO require the training of multiple RF chains at the same time, resulting in a large number of possible beam combinations which makes an exhaustive search of all combinations impossible. The resulting beamtraining process is lengthy and

complex and can provide sub-optimal results, requiring practical evaluations of the performance. In addition, due to the sparsity of the mmWave channel successfully establishing multiple independent streams can challenging. An alternative design proposed for devices with multiple RF chains is a multi-connectivity [43–46] or cell-free deployment [47, 48], where users maintain multiple communication links with *different* APs simultaneously. In addition to the spectral efficiency improvements, this design can address key issues of resilience and robustness in mmWave networks [49, 50]. Most previous work, however, proposes a coherent design which requires tight synchronization between APs in the network which is extremely challenging to implement. As advancements in the design and fabrication of mmWave electronics are paving the way towards high-performance, robust, low-power and low-cost RF integrated circuits, multi-RF devices represent the next step in the evolution of mmWave technology. Exploring and evaluating different proposed architectures will be necessary to realize the full potential of multi-RF mmWave.

- **Joint communication and sensing at mmWave frequencies:** Wireless sensing is one of the most prominent research areas currently, leveraging wireless communication signals to sense everything from motion detection [51–53] to health monitoring [54, 55] to 3D vision applications [56, 57]. Joint communication and sensing WiFi systems benefit from existing infrastructure and ubiquitous deployment. Significant research [51–55, 58] has demonstrated the potential of sensing using WiFi signals, prompting the development of IEEE 802.11bf [59], the latest WiFi standard which aims to standardize sensing procedures and facilitate the real-world deployment of joint communication and sensing WiFi systems. IEEE 802.11bf addresses both sensing at sub-7 GHz as well as in the 60 GHz band. Sensing at mmWave frequencies offers higher range and angular resolution due to the high signal bandwidth and can therefore enable high-precision sensing. However, the unreliable nature of WiFi signals, particularly in the mmWave band, can affect the sensing accuracy. Similarly, allocating resources to sensing traffic can negatively affect data communication. Thus it is very important to study resource allocation in joint communication and sensing systems, as well as determine the practical performance of these systems in larger networks.

## 1.3 Contributions of the thesis

The overall goal of this thesis is to explore dense and large mmWave networks from many aspects, looking at proposed deployment architectures (standard WiFi, MIMO, multi-connectivity), and analysing different performance trends (throughput, reliability, BFT accuracy) with regard to the density of the network and the level of interference. Based on the insights gained during the analysis we then design solutions to current

mmWave networks. In the following, we provide a summary of the contributions of this thesis:

- **An IEEE 802.11ay model in ns-3:** Our implementation expands the existing IEEE 802.11ad model to capture the specifics of IEEE 802.11ay such as the new frame format, MCSs, information elements and management frames. The novel BFT protocols of IEEE 802.11ay are supported by implementing the new Training (TRN) field format and new elements such as the Short Sector Sweep (SSW) packet. We support channel bonding with up to 4 bonded channels. Finally, a major extension introduces support for MIMO operation including the ray-tracing channel model and the full standard-compliant MIMO BFT protocols. The implementation has been updated to also support non-standard compliant features such as SIGNaling in the PHY Preamble (SIGNiPHY) - our PHY layer signalling solution, and multi-connectivity networks. The latest update covers the implementation of sensing procedures from the novel IEEE 802.11bf protocol. Our model is open-source and publicly available for the use of the research community.

- **BFT for dense networks:** We perform a detailed analysis of BFT algorithms in dense networks by looking at the accuracy of the training and the training overhead effect on data communication. Our work is the first analysis of the novel Group Beamforming protocol introduced in IEEE 802.11ay which is particularly suited to dense networks as it allows simultaneous training of multiple stations and has very low training overhead. We compare the performance of Group Beamforming against the legacy 802.11ad BFT - the Sector Level Sweep (SLS) protocol in a variety of simulation scenarios. Our results show that interference during the training can lead to errors in the beampattern selection which results in significant drops in the achievable Signal-to-Noise Ratio (SNR) (up to 35 dB in our scenarios). Although the errors are mostly transient and devices are able to recover in the next training, they can lead to service outages and loss of connection which is shown to be quite harmful. Moreover, we found that the legacy SLS protocol has difficulty scaling in dense networks due to the large overhead, leading to significant channel time spent on BFT procedures. In contrast, Group Beamforming proved to be a very promising alternative, allowing for gains in achievable throughput of up to 500 Gbps. The main drawback of Group Beamforming is that in very dense scenarios the AP training is quite vulnerable to interference leading to significant BFT errors. We therefore designed two enhanced versions of Group Beamforming that significantly improve its performance. Our modified version of Group Beamforming outperforms both the legacy 802.11ad BFT and the standard Group Beamforming and provides the best network performance in all scenarios we tested, both in terms of network throughput and BFT accuracy.

- **Enhanced efficiency and spatial reuse for mmWave WiFi:** Our study of dense mmWave networks revealed that one of the main causes of inefficiency is the contention-based channel access with CSMA/CA. Due to the lack of a transmission schedule devices have to use omni-directional reception beampatterns, as they could receive packets from *any* user in the networks. This affects both the single-link performance due to the lower antenna gains and reduces spatial reuse since devices receive interference from all directions. Furthermore, it significantly contributes to the problem of unwanted packet overhearing. Packet overhearing is another consequence of CSMA/CA where due to the lack of a transmission schedule devices attempt to decode all packets that they overhear, including many packets not meant for them. While in sub-6 GHz networks, this is mainly an energy-efficiency problem, in mmWave networks it causes much more harm as it limits spatial sharing, keeping devices busy decoding unwanted packets so that they are unable to receive their own data. To solve both of these issues, we designed SIGNiPHY - a PHY layer signalling solution that allows the embedding of the user identifier (ID) in the packet preamble, at the very start of the packet. Identifying the user early in the reception process enables fast switching to a directional receive beampattern to receive the packet payload with higher gain and reduce the received interference. Additionally, it facilitates fast packet filtering of unwanted packets, leaving the device free to receive any packets for which it is the recipient. SIGNiPHY does not incur any overhead and ensures backwards compatibility and interoperability with legacy devices. It is additionally designed to be low-complexity and require minimal modifications for easy implementation. We implemented SIGNiPHY in our ns-3 model and found that it enhances the throughput between 13% and 230% compared to different baseline schemes while ensuring higher fairness and boosting the performance of the lowest-performing users. Moreover, we also evaluated SIGNiPHY real-time operation in our FPGA-based testbed and showed that SIGNiPHY significantly enhances the resilience to interference, enabling the decoding of all packets at SINR as low as 0 dB while coping with up to 3 different interferers.

- **IEEE 802.11ay mmWave MIMO evaluation:** We provide the first protocol evaluation of the standardized MIMO operation for IEEE 802.11ay networks. Our ns-3 implementation includes the full analog BFT procedure for both Single User MIMO (SU-MIMO) and MU-MIMO as well as the SU-MIMO channel access procedure and data transmission. We demonstrate the complexity and length involved in simultaneously training multiple transmit and receive beampatterns and how trade-offs can occur between the scalability and duration of the process and the spatial separation achieved. We also highlight how because an exhaustive search of all combinations of beampatterns is not possible, the selection of good candidate beampatterns to test is crucial to the performance. We validate the performance

in static and mobile scenarios and look at how the BFT is able to cope with user mobility. Finally, we demonstrate how even with analog beamforming it is possible to successfully establish independent spatial streams for a multifold increase in throughput.

- **Enhanced resilience through in-band multi-connectivity:** While MIMO communication can offer significantly higher throughput rates, the BFT process is complex and lengthy and the mmWave channel sparsity limits the number of spatial streams that can be established. In addition, MIMO does not address resilience and robustness which are arguably much bigger problems in mmWave where the large bandwidth allows for very high data rates even with a SISO connection. Therefore we considered alternative solutions to better utilize the enhanced capabilities of multi-RF devices. Multi-connectivity approaches enable users to have multiple simultaneous connections to different APs in the network which can significantly improve the resilience and robustness to link breakage. Furthermore, unlike designs at sub-6 GHz which have to rely on phase-coherent processing for interference management, mmWave networks offer an opportunity for simpler implementation by relying on spatial separation of signals through analog beamforming. This prevents the need to have tight synchronization and global knowledge of Channel State Information (CSI) in the network, making mmWave multi-connectivity networks easier and simpler to implement and deploy. We validate this approach by designing a multi-connectivity mmWave network with minimal network synchronization, relying solely on analog beamforming for spatial separation. Our evaluation results demonstrate that in-band multi-connectivity with 4 asynchronous and independent links can provide uninterrupted service even in dense, high-traffic scenarios, compared to up to 20% of service loss in a standard single-connectivity deployment. Distributing the traffic across multiple APs also has throughput gains of up to 30%, showing that multi-connectivity mmWave networks can provide a high-throughput, reliable and stable service for next-generation applications.

- **mmWave sensing with IEEE 802.11bf:** mmWave sensing is one of the most promising applications, receiving significant interest from both academia and industry. IEEE 802.11bf offers an opportunity for widespread deployment of joint communication and sensing WiFi systems by standardizing interfaces and sensing procedures to enable interoperability and coordination between multiple devices. We present the first system-level performance evaluation of mmWave sensing with IEEE 802.11bf which can capture effects like reduced sensing accuracy due to packet collisions and disruptions to the sensing timing due to the uncoordinated channel access. In addition, we observe the effect that different sensing configurations have

on data communication and identify scenarios which ensure accurate sensing without compromising data communication.

## 1.4 Outline of the thesis

The remainder of this thesis is divided into 8 chapters. Chapter 2 presents the implementation details of extending the IEEE 802.11ad ns-3 model to support the IEEE 802.11ay protocol and a performance validation of the new features. Using this model we then provide a comprehensive analysis of SISO dense networks and design enhancements to address the main performance limitations we encountered. Chapter 3 addresses BFT in dense networks as a crucial factor dictating the performance of mmWave networks, while Chapter 4 presents a PHY layer signalling solution that aims to improve spatial reuse and efficiency. In the next part of the thesis, we evaluate multi-RF networks, with two proposed architectures. Chapter 5 evaluates the classical MIMO design as standardized in IEEE 802.11ay, while Chapter 6 considers an alternative design with a multi-connectivity approach based on local analog beamforming. Lastly, Chapter 7 presents a joint communication and sensing deployment based on the novel IEEE 802.11bf protocol which looks at resource allocation and trade-offs between sensing accuracy and communication performance. Chapter 8 provides a summary of our findings and concludes the thesis.

# 2 An IEEE 802.11ay Model in Network Simulator 3 (ns-3)

## 2.1 Introduction

The approval of the IEEE 802.11ad [21] standard in 2012 pushed Millimeter-Wave (mmWave) research and enabled the first system-level evaluations of multi-Gbps wireless operation in the unlicensed 60 GHz band. However, the performance gains were considered minor in comparison to the sub-6 GHz IEEE 802.11ac [7]. Moreover, emerging wireless applications such as mmWave distribution networks, uncompressed content streaming for Augmented Reality (AR)/Virtual Reality (VR) technologies, and dense network deployments required even higher data rates, lower latency and higher robustness. This motivated the WiFi alliance to form the Task Group ay in 2015 to define the next-generation mmWave standard, named IEEE 802.11ay [27]. The following design factors were taken into account during the standardization phase: i) the standard must support a throughput of at least 20 Gbit/s, ii) it must maintain backward compatibility with IEEE 802.11ad, and iii) it must extend the set of possible use cases and scenarios by introducing novel solutions at the Medium Access Control (MAC) and physical (PHY) layers. Most of these requirements are achieved thanks to the incorporation of advanced physical layer solutions that are predominant in wireless systems operating at sub-6 GHz. These solutions include Multiple-Input and Multiple-Output (MIMO), channel bonding and aggregation, fast Beamforming Training (BFT), and multi-user transmission. The goal was to boost the performance and reliability by several orders of magnitude, to improve the network scalability and to support diverse deployments and end devices.

At the time of writing, there is only a single IEEE 802.11ay Commercial Off-the-Shelf (COTS) device [28] available, although there are no details provided on which features it supports and whether advanced solutions like MIMO and channel bonding are incorporated. Experimental mmWave testbeds which are IEEE 802.11ay compatible exist, however, they are not able to support all new features and it is difficult to evaluate network aspects with this approach. Additionally, the implementation of new solutions is expensive and time-consuming and there are limitations of the network size that can be studied. The

lack of available tools hinders research progress and innovation into the protocol aspects of IEEE 802.11ay. In this chapter, we fill this gap by introducing our IEEE 802.11ay implementation in the popular network simulator ns-3. Network simulators can be very valuable for fast and cost-effective evaluation of new protocol features. Unlike link-level simulators, they can provide insight into the performance with dense network settings and have system-level evaluation. Ns-3 is an open-source, packet-level simulator which is widely used in academia and research and supports a comprehensive set of technologies including WiFi, LTE, WiMAX, etc. An existing IEEE 802.11ad model [60] implements the key MAC and PHY features for directional mmWave WiFi, such as BFT, channel access and packet structure. In addition, the model was integrated with the National Institute of Standards and Technology (NIST) Quasi-Deterministic (Q-D) Channel Realization Software [61] to allow for high-fidelity modelling of mmWave propagation. This provided an excellent basis for our IEEE 802.11ay implementation, particularly as it helped us ensure backward compatibility between the two standards.

In Chapter 5 we present the details of the MIMO implementation, while in this Chapter we focus on the baseline IEEE 802.11ay implementation and evaluate the performance with the new model. Section 2.2 gives a brief overview of the IEEE 802.11ad model which we extend, Section 2.3 presents our implementation and Section 2.4 shows the validation and evaluation of the proposed model. Finally, Section 2.5 concludes this Chapter.

The main contributions of this chapter are as follows:

- We upgrade the ns-3 IEEE 802.11ad model [37–39] to support IEEE 802.11ay. This includes the 802.11ay frame structure, Modulation and Coding Schemes (MCSs), channelization, and error model.
- We implement channel bonding of up to 4 contiguous channels.
- We support the higher aggregation enabled by IEEE 802.11ay.
- We implement the new Training (TRN) field structure and the new Short Sector Sweep (SSW) frame to support the new BFT protocols.
- We evaluate the performance of the new model with a focus on network scalability and look at the performance gains with regard to IEEE 802.11ad.
- Finally, we make our implementation publicly available to the research community.

## 2.2 Background on the ns-3 IEEE 802.11ad model

Directional Multi-Gigabit (DMG) operation in the 60 GHz band required significant PHY and MAC modifications. In this section, we briefly introduce the ns-3 IEEE 802.11ad model on which we base our implementation.

### 2.2.1 MAC layer operation

Many new features are required to support directional communication. In addition to the need for BFT mechanisms, procedures such as initial access were redesigned as they assumed omnidirectional reach of mmWave signals. Three new classes - `DmgWifiMac`, `DmgApWifiMac` and `DmgStaWifiMac` were added to capture the specifics of DMG operation in the MAC high layer. In addition, modifications to the MAC low layer were necessary to organize medium channel access and to ensure directional transmission and reception. The MAC features supported by the ns-3 IEEE 802.11ad model are summarized below:

• **DMG Access Periods**: Channel access in WiFi networks is organized in Beacon Intervals (BIs) which typically last 100 ms and start with a beacon transmission from the Access Point (AP) to advertise the existence of the network to new users. IEEE 802.11ad redesigns the BI to adapt to directional operation, as seen in Figure 2.1. Instead of a single beacon transmission, the Beacon Transmission Interval (BTI) was added to support multiple DMG beacons transmitted with different directional beampatterns to reach users all around the AP. The Association Beamforming Training (A-BFT) then allows unassociated Stations (STAs) to perform initial BFT to be able to communicate with the AP. The Announcement Transmission Interval (ATI) can be used to communicate management information directionally to beamtrained STAs for more efficient communication. The rest of the BI is used for data transmissions, either using the standard Contention-Based Access Period (CBAP) or with a new scheduled access in Service Periods (SPs). The timing of the BI scheduling in ns-3 is organized by the `DmgApWifiMac` class and new functions were added to establish the rules for transmitting in each access period.

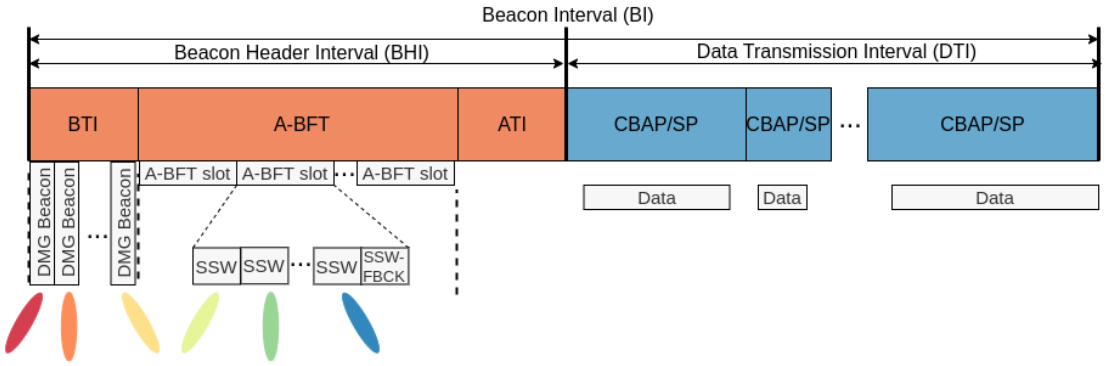


Figure 2.1: IEEE 802.11ad Beacon Interval Structure

• **DMG Channel Access Schemes:** IEEE 802.11ad defines three medium access schemes to support efficient operation at mmWave frequencies. In addition to the traditional WiFi contention channel access with Carrier Sense Multiple Access with

Collision Avoidance (CSMA/CA), mmWave WiFi also allows scheduled access in SPs where the channel is reserved for a certain user and dynamic allocation through a polling scheme. All three access schemes are supported by the ns-3 IEEE 802.11ad model. The schedule for the Time-Division Multiple Access (TDMA)-like access is done by the DMG AP, based on requests from associated STAs. The MAC low layer then ensures that communication in the SP is only allowed from the source DMG STA towards the destination DMG STA.

• **Beamforming Training Mechanisms:** IEEE 802.11ad defines two BFT procedures which aim to find the optimal beampattern for communication with a peer device. The Sector Level Sweep (SLS) procedure is mandatory and is used to find a coarse antenna sector for initial communication, while the optional Beam Refinement Protocol (BRP) phase allows further refinement of the selected sector to increase the antenna gain. Both procedures are implemented in the ns-3 model and are initiated from the `DmgWifiMac` class which controls both the frame transmission and the antenna beampattern switching. The selected beampatterns are then used by the `MacLow` class during data transmission and reception. A detailed explanation of DMG BFT is given in Chapter 3.

Additional features supported include Fast Session Transfer (FST) to the more stable 2.5/5 GHz band in cases of link breakage, relay operation to extend the communication range by allowing two devices to communicate through a third, relay device, a spatial sharing mechanism to determine communication links which can be simultaneously active and a clustering mechanism to mitigate inter-AP interference.

### 2.2.2 PHY layer operation:

A new `DmgWifiPhy` class was added to define DMG characteristics on the PHY layer and to handle the transmission and reception of PLCP Protocol Data Units (PPDUs). This includes the modelling of the DMG frame format, as shown in Figure 2.2. In addition to the typical IEEE 802.11ad fields like the packet preamble (composed of the Short Training Field (STF) and Channel Estimation Field (CEF) field), PHY header and packet payload, the DMG PPDU also optionally adds an Automatic Gain Control (AGC) and TRN field. These fields can be used to perform BFT by fast beampattern switching between the TRN subfields. There can be a maximum of 64 TRN subfields in one field, divided into up to 16 TRN units. Each unit holds 4 TRN subfields for training and one CEF for synchronization and channel estimation. The TRN subfields are composed of Golay sequences with very good correlation properties and can be used to train either the transmit or receive beampatterns of STAs. The state machines for the transmission and reception of both types of TRN fields were implemented in ns-3 along with the correct beampattern switching which is controlled by the `Codebook` class described below.

The model also implements the 3 different PHY layers defined by IEEE 802.11ad to handle operation in varying conditions. The Control PHY is designed to be robust for management communication under low Signal-to-Noise Ratio (SNR) conditions, while the Single Carrier (SC) and Orthogonal Frequency Division Multiplexing (OFDM) PHYs are meant for data communication with higher achievable data rates. More details on the different PHY layers are given in Section 2.3.

Finally, a new `DmgErrorModel` class was added for a realistic calculation of the Packet Error Rate (PER) based on the Signal-to-Interference-plus-Noise Ratio (SINR). The model uses lookup tables (LUTs) to map the Bit Error Rate (BER) to the SNR for each IEEE 802.11ad MCS. The LUTs are generated with the MATLAB *WLAN Toolbox* by simulating the transmission of IEEE 802.11ad compliant frames and calculating the BER for a set of discrete SNR values with a step size of 0.2 dB. The `DmgErrorModel` then uses linear interpolation to predict the PER for any SNR value.

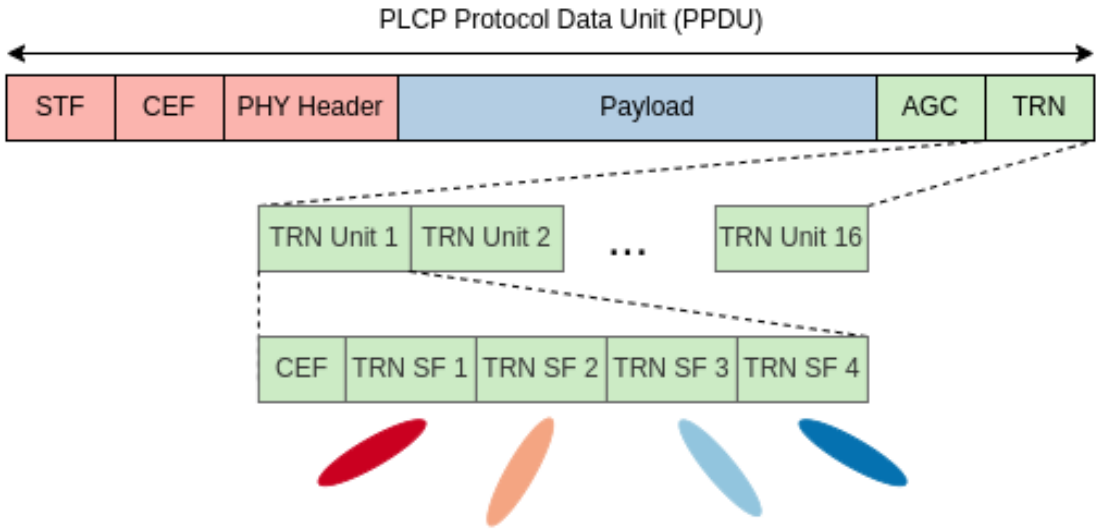


Figure 2.2: IEEE 802.11ad Frame Format

### 2.2.3 Q-D Channel Model

The standard ns-3 channel model relies on a simple propagation model which cannot capture the unique propagation characteristics of the mmWave band. As a result, the performance only depends on the geometrical distance between the nodes and the number of interfering devices. In order to enable high-fidelity simulation of mmWave networks it is necessary to capture the sparse structure of the mmWave channel and to factor in the simulation environment. The ns-3 IEEE 802.11ad model enables integration with the Q-D channel realization software developed by NIST. The software uses a mixture of ray-tracing and statistical modelling to model the mmWave channel. The first strong rays which contain most of the power are generated with ray-tracing, while the weak diffuse

components caused by rough surface scattering are generated statistically based on the material properties. The software generates Q-D trace files which contain the number of Multipath Components (MPCs) between each pair of devices in the network, and for each MPC the pathloss, phase shift, delay, and angles of arrival and departure. These Q-D files are parsed in the new `QdPropagationEngine` class to compute the received power, taking into account the phased antenna array response based on the beampattern currently used. This is determined by the `Codebook` class, described below.

### 2.2.4 Codebook

The `Codebook` class was introduced to model aspects related to beamforming to generate directional beampatterns. The codebook can contain multiple Phased Antenna Arrays (PAAs), each composed of multiple predefined sectors that allow for steering in different directions. Each sector can further have a set of custom Antenna Weight Vectors (AWVs) for optimized steering. The `Codebook` class implements the necessary functions to perform beampattern sweeping during the BFT operations, as well as steer towards the communication peer device during data communication. To support the generation of realistic beampatterns, the parametric codebook model and the *IEEE 802.11ad Codebook Generator* were developed. They characterize the antenna based on the antenna response of each element and the steering vector of the antenna array. The parametric codebook model is integrated with the Q-D channel to simulate the interactions between the radiation pattern of the PAA and the signals coming from reflections in the environment.

## 2.3 Implementation

We now present the design and the implementation details of our IEEE 802.11ay model in ns-3, which is publicly available on GitHub [60]. Figure 2.3 shows the modifications required to extend the ns-3 IEEE 802.11ad model across the different layers of the protocol stack. Details regarding the MIMO extension can be found in Chapter 5, and the Group Beamforming implementation is presented in Chapter 3. To prevent code duplication and maintain backward compatibility with IEEE 802.11ad as mandated by the standard, we chose to expand the existing DMG classes to implement Enhanced Directional Multi-Gigabit (EDMG) operation, rather than creating new EDMG classes whenever possible. This also simplifies future expansions of the model to support devices that can operate in both DMG and EDMG modes.

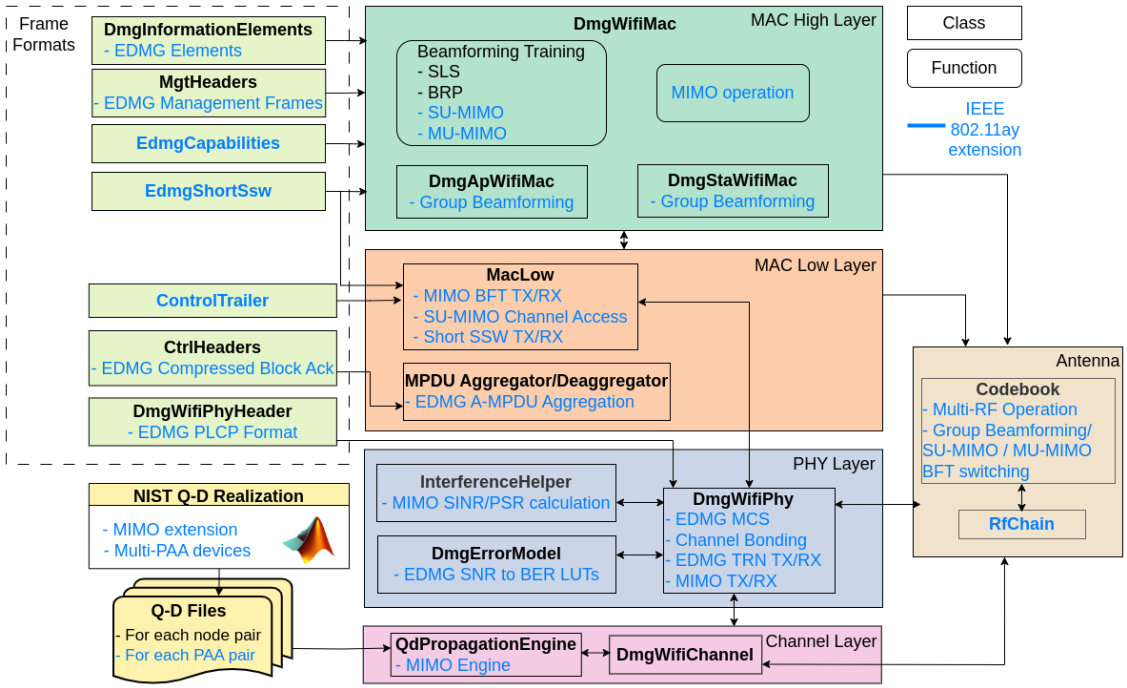


Figure 2.3: IEEE 802.11ay extension of the ns-3 IEEE 802.11ad architecture

### 2.3.1 IEEE 802.11ay Framing

Figure 2.4 depicts the EDMG frame format. To maintain backward compatibility with IEEE 802.11ad, the EDMG frame reuses both the DMG preamble and DMG header fields. Thus, the EDMG frame is divided into two parts. The first part, referred to as the Non-EDMG portion, comprises a Legacy-STF (L-STF), Legacy-CEF (L-CEF), and legacy-header fields and is recognizable by DMG devices. The second part, which is known as the EDMG portion, contains all the fields that are recognized by EDMG STAs, including the EDMG STF and CEF fields and the new EDMG headers. The legacy fields and the EDMG Header-A are present in all EDMG packets. The EDMG STF and CEF fields are only present if either channel bonding or MIMO is used to enable signal parameters and channel estimation. The EDMG Header-B is only present in Multi User (MU) transmission and contains individual information for each STA the MU packet is directed towards. The TRN field is optionally present for BFT. We provide a detailed PHY layer model for transmitting and receiving different fields in the EDMG Physical Layer Convergence Protocol (PLCP) frame including the mathematical formulas to calculate the transmission time and the construction of the PPDU based on the transmission type.

Similar to IEEE 802.11ad, IEEE 802.11ay supports three physical layer frame types: Control, SC, and OFDM. The Control PHY is dedicated to the transmission of management and control frames such as DMG beacons and BFT frames. Thus, it is

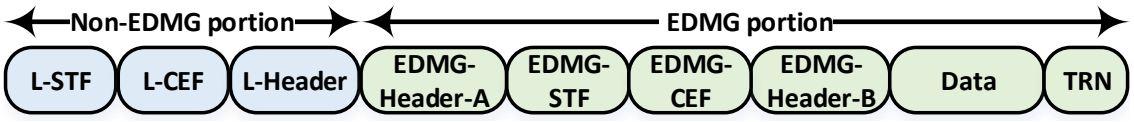


Figure 2.4: IEEE 802.11ay Frame Format

designed to be robust for communication under low SNR conditions. As a result, a control frame is first encoded using Binary Phase-Shift Keying (BPSK) modulation with a coding rate of 1/2, and then it is spread using Golay Sequence (Ga) of length 32. All frames transmitted in this mode can also be recognized and decoded by legacy DMG devices.

For data communication, either EDMG SC or EDMG OFDM can be used. The EDMG SC PHY defines an expanded set of MCSs (1 to 21) with a maximum PHY throughput of 8085 Mbit/s per spatial stream over a single channel with a normal Guard Interval (GI). IEEE 802.11ay allows for greater flexibility by also supporting transmission with a short and long GI resulting in higher and lower data rates respectively. Likewise, EDMG OFDM specifies 20 EDMG MCSs with a maximum throughput of 8316 Mbit/s. The standard mandates the support of EDMG SC mode MCSs 1 to 5 and 7 to 10 with a single spatial stream, while EDMG OFDM is optional. We support all new coding rates introduced by IEEE 802.11ay along with the different GIs.

To ensure accurate simulations, we integrate the IEEE 802.11ay SNR to BER LUTs generated by the NIST IEEE 802.11ay link-level simulator described in [62]. This allows us to update the `DmgErrorModel` class to support EDMG packet decoding.

Finally, we implement all new information elements, management frames and headers to support EDMG operation.

### 2.3.2 Channel Bonding

In IEEE 802.11ad, the 60 GHz band covers operation from 57 GHz to 64 GHz divided into four channels of 2.16 GHz. Communication at this frequency range suffers from high oxygen absorption which limits the communication range. With the growing interest in Fixed Wireless Access (FWA) deployments and the adoption of the unlicensed mmWave band for backhauling and fronthauling, the Federal Communications Commission decided to double the available bandwidth to cover 57 GHz to 71 GHz, providing a total of 14 GHz of unlicensed spectrum. The new frequency range between 64 GHz and 71 GHz does not suffer from high oxygen absorption which makes it more suitable for backhaul applications where long-range communication is needed.

Figure 2.5 shows the possible channel configurations for IEEE 802.11ay. IEEE 802.11ay supports operation in eight 2.16 GHz channels. To increase the data rate further, IEEE 802.11ay allows bonding a contiguous set of channels. A maximum of four channels

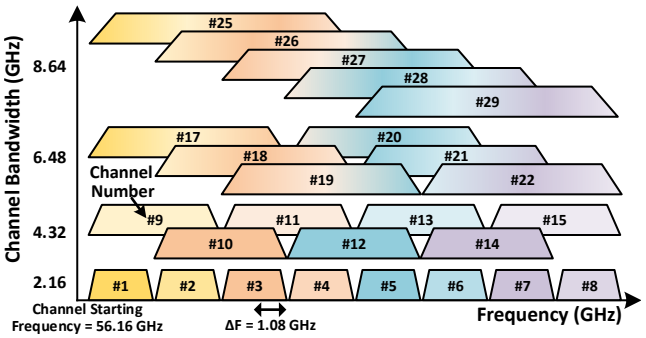


Figure 2.5: EDMG Channel Configurations

can be bonded which results in a channel width of 8.64 GHz. The standard mandates the support of two bonded channels (4.32 GHz). We implement channel bonding with all channel configurations defined in the standard as well as the corresponding channel masks for accurate interference calculations.

### 2.3.3 Frame Aggregation

Frame aggregation is a key feature used in IEEE 802.11 networks to reduce MAC inefficiency and increase the achievable data rates. Transmitting aggregated packets that contain larger payloads with single channel access allows for a reduction of channel usage time spent on MAC overhead. In mmWave WiFi where the large bandwidth enables multi-Gbps data rates the effect of MAC overhead is amplified. Time spent on MAC backoff, inter-frame spacing and control frames represents high amounts of data that could be transmitted. IEEE 802.11ay significantly increases the frame aggregation possible to enable the highest achievable data rates. Aggregate MAC Protocol Data Unit (A-MPDU) aggregation is increased from a maximum of 262 KB to 4 MB which represents an increase of over 16 times. To support this high level of aggregation IEEE 802.11ay introduces the `EDMG Extended Compressed ACK` variant that can acknowledge the reception of up to 1024 frames compared to a maximum of 64 frames in IEEE 802.11ad. We implement both the increased A-MPDU aggregation and the `EDMG Extended Compressed ACK` frame format.

### 2.3.4 Beamforming Training Support

IEEE 802.11ay introduces many new BFT mechanisms to enable more efficient and robust training and offer support for heterogeneous deployments and use cases. To support these new mechanisms it introduces two novel elements: 1) a redesigned TRN field with expanded use and flexibility and 2) a short SSW frame with lower overhead. In the following, we present their implementation:

• **EDMG TRN Field:** As discussed in Section 2.2, IEEE 802.11ad introduces a special element, called the TRN field, at the end of BRP packets to perform fast beam switching across multiple narrow beampatterns within the same packet. The TRN field is composed of multiple subfields and devices switch beampatterns between each subfield. In DMG BRP-TX packets each subfield is transmitted with a different beampattern, while in DMG BRP-RX packets the whole field is transmitted with the same beampattern and receivers switch between receive beampatterns while receiving the different subfields. Subfields contain multiple Golay sequences with good correlation properties and have a short duration of just 364 ns.

One problem with the DMG TRN field was that the standard mandates that any signal transients that occur due to the change of a beam pattern must settle by the end of the first Golay sequence. This imposes a stringent beam switching time in the order of 36 ns. Building an RF Integrated Circuit with such specifications is challenging and requires an optimized analog and digital architecture. Due to these constraints, many COTS devices either omit BRP support or implement a proprietary version with a relaxed switching time. Another deficiency is that the DMG design only allows for up to 64 beampatterns to be trained in a single packet, which might not be sufficient for large antenna arrays and for training multiple PAAs. Lastly, simultaneous transmit and receive training which is necessary for MIMO training is not supported.

To address all these issues, IEEE 802.11ay redesigned the TRN field to provide flexibility to end devices with heterogeneous hardware capabilities. To address issues with scalability, the number of TRN units that can be added to a packet was significantly increased from 16 to 255. Moreover, the number of subfields in one unit increased from 4 to up to 16. The TRN field structure was also made much more flexible to adapt to different configurations. In addition to the variable number of subfields per unit, IEEE 802.11ay introduces a variable size of the Golay sequence that can be configured by the user and additionally, in the case of channel bonding, depends on the number of continuous channels. The size of each Golay sequence $TRN\_BL$ can take one of the following values $[64, 128, 256]$ samples. If channel bonding is used, then the size of the Golay sequence is multiplied by the number of contiguous 2.16 GHz channels ($N_{CB}$). IEEE 802.11ay also defines a unique orthogonal set of Golay sequences for each space-time stream ($i_{Tx}$) to facilitate channel estimation and beamforming training for MIMO communication. To address the stringent switching requirements for the antenna array IEEE 802.11ay allows for N consecutive subfields to be transmitted with the same beampattern. Finally, it introduces a new EDMG BRP-RX/TX packet to enable simultaneous transmit and receive training in the same packet.

We implement the flexible and configurable TRN field structure for all three packet types, as shown in Figure 2.6. Figure 2.6.a shows the general structure of all TRN fields with the TRN subfield composition of 6 Golay sequences, while b, c and d show the

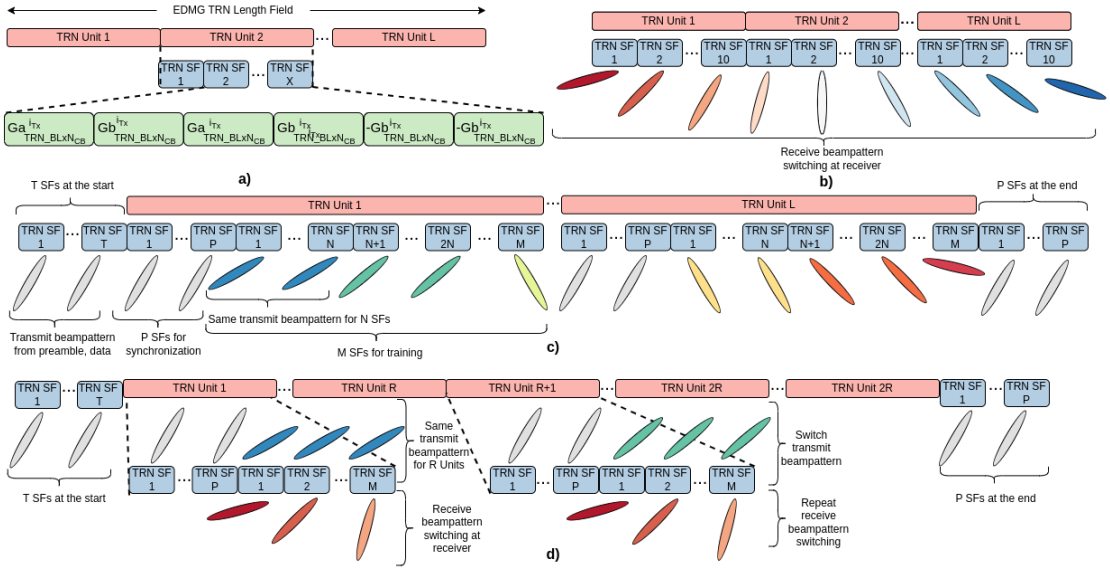


Figure 2.6: EDMG TRN Field Structure for a) General structure; b)EDMG BRP-RX; c) EDMG BRP-TX; c) EDMG BRP-RX/TX

individual structure of different EDMG BRP packets. The EDMG TRN field is composed of L TRN units. In the case of EDMG BRP-RX frames, each unit is composed of 10 Receive Training (TRN-R) subfields which are all used for receive beampattern training. In the case of EDMG BRP-TX and BRP-RX/TX frames, the TRN field begins with T subfields and ends with P subfields not meant for training and transmitted with the same beampattern as the preamble and data. These fields are meant to signal the transition between the data and the TRN field and the end of the frame. Additionally, each unit also starts with P subfields transmitted with the same beampattern as the preamble (that can be used for synchronization, gain control or channel estimation) and then M subfields used for beamtraining. In BRP-TX packets the transmit beampattern is changed after every N Transmit Training (TRN-T) subfields, while the receiver maintains a quasi-omnidirection beampattern. In BRP-RX/TX the same transmit beampattern is kept for R units, allowing the receiver to perform receive beampattern training. The next R units are transmitted with a different beampattern while the receive beampattern training is repeated, allowing for all possible combinations of transmit and receive beampatterns to be tested. The complete structure of the different types of BRP frames is explained in [7].

Additionally, we incorporated the corresponding state machines for transmitting and receiving all variants such as EDMG BRP-TX, EDMG BRP-RX, and EDMG BRP-RX/TX. In Figure 2.7 we show the state-machine for transmitting EDMG BRP-TX and EDMG BRP-RX frames, where during the transmission of EDMG BRP-RX frames the grey blocks are omitted and number of training subfields in a Unit M is set to 10. Figure 2.8 shows the corresponding receive state machine. The transmission and reception

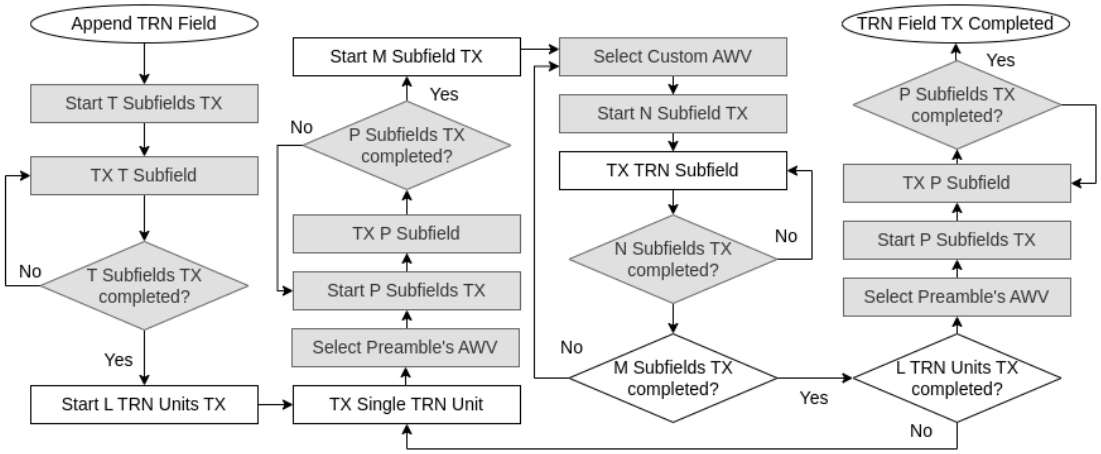


Figure 2.7: EDMG BRP-TX & EDMG BRP-RX Transmit State Machine Implementation

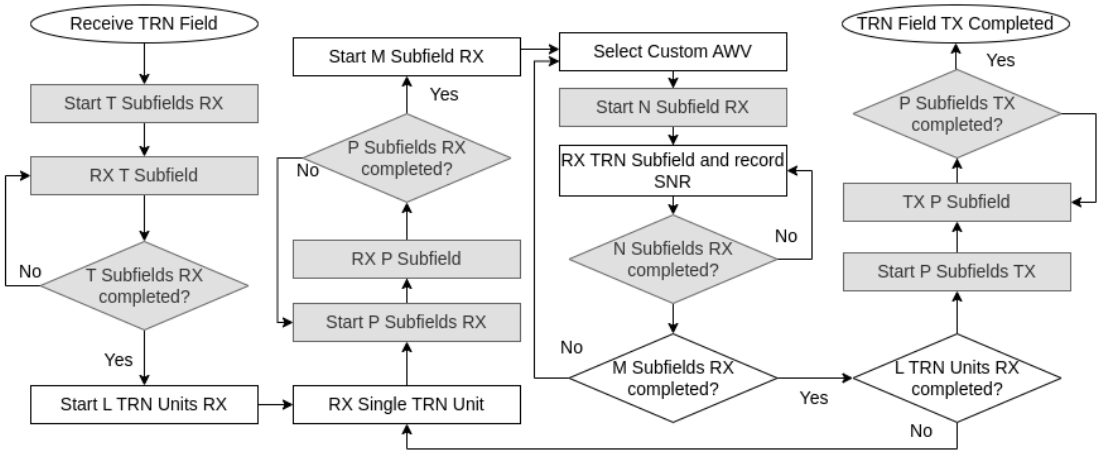


Figure 2.8: EDMG BRP-TX & EDMG BRP-RX Receive State Machine Implementation

of the TRN field are handled in the `DmgWifiPhy` class, while our `Codebook` class handles switching between AWVs at the boundary of each TRN SF. To address the expanded use of TRN fields in IEEE 802.11ay we also added the possibility to switch between antenna sectors, and not just AWVs between subfields.

• **Short SSW frame:** To speed up the BFT, IEEE 802.11ay defines a new PHY layer frame. The short SSW frame is a PHY layer frame that is 6 bytes long compared to 26 bytes for the legacy SSW which results in a 31% reduction in the transmission time. We add support for these frames by enabling the transmission of WiFi packets without a MAC header in the `MacLow` and `DmgWifiMac` classes.

The new EDMG TRN field is used to implement the novel IEEE 802.11ay Group Beamforming procedure presented in Chapter 3. Both the EDMG TRN field and the

Table 2.1: Simulations Parameters

| Parameter Name | Parameter Value |
|---|---|
| Application Type | OnOffApplication |
| Payload Size | 1472 Bytes |
| Transport Protocol | UDP |
| MAC Queue Size | 4000 Packets |
| Aggregation Type | A-MSDU and A-MPDU |
| A-MSDU Max. Size | 7935 Bytes |
| PPDU Max. Duration | 2 ms |
| Block ACK Size | 1024 Frames |
| MAC Protocol | CSMA/CA |
| Number of Transmit Sectors | 27 Sectors |
| Sector Azimuth Steering Angles | -80°:20°:80° |
| Sector Elevation Steering Angles | -45°, 0°, 45° |
| Transmit Power | 10 dBm |
| Operating Frequency | 60.48 GHz (CH2) |

Short SSW frame are used in the MIMO BFT protocols as described in Chapter 5.

## 2.4 Evaluation

In this section, we evaluate and validate our IEEE 802.11ay implementation in ns-3. All our simulation scenarios utilize the Q-D channel model. Simulation parameters are summarized in Table 2.1. All the devices in the network use a 2x8 element Uniform Rectangular Array (URA) PAA which yields a narrow beam in the azimuth plane, and a wide beam in the elevation plane.

### 2.4.1 Channel bonding and EDMG MCSs

In this simulation, we evaluate the maximum achievable throughput for the IEEE 802.11ay protocol for all the EDMG MCSs with various channel widths. Our scenario consists of two IEEE 802.11ay devices with a Line-Of-Sight (LOS) link with a distance of one meter. We configure the two devices to use the optimal beam pattern thus ensuring a high SNR value that prevents any packet loss. To eliminate beamforming training overhead, we install `DmgAdhocWifiMac` which is an experimental MAC layer implementation that facilitates studying PHY layer features without adding the complexity of the full MAC protocol. This MAC implementation allocates the whole BI for data transmission.

Figure 2.9 depicts our simulation results for EDMG SC and EDMG OFDM PHYs. To include the overhead of each layer in the protocol stack, we measure the throughput at the

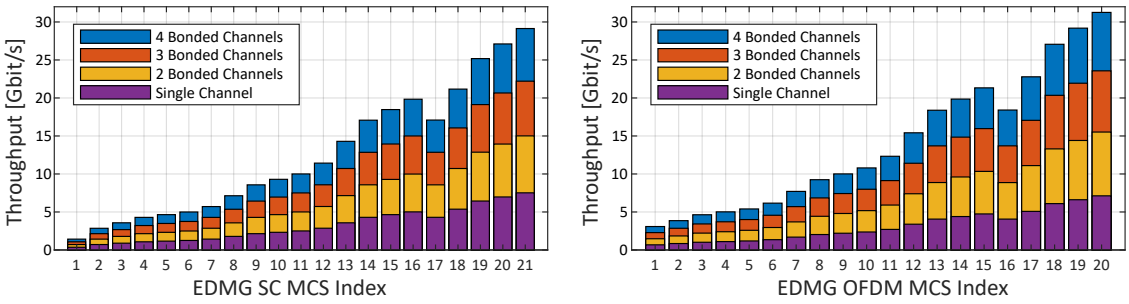


Figure 2.9: EDMG MCSs Throughput for Different Channel Sizes

application layer. We observe that the maximum achievable throughput with four bonded channels is around 29.6 Gbit/s for EDMG SC and 31.25 Gbit/s for EDMG OFDM. We notice a degradation in the throughput for EDMG-MCS-17. This is because EDMG-MCS-17 uses a 64-QAM modulation scheme with a coding rate of 1/2, which results in fewer data bits per SC block compared to EDMG-MCS-16. It is worth mentioning that this might cause issues with Rate Adaptation Algorithms (RAAs) as they would expect a monotonic increase in throughput when increasing the MCS.

The throughput obtained in this simulation considers an ideal scenario where we have neither collision on the wireless medium nor packet loss. In a real network, the throughput will be lower due to i) the overhead imposed by different channel access periods in the BI, ii) the usage of the Ready-to-Send (RTS)/Clear-to-Send (CTS) handshake protocol, and iii) frequent link maintenance through BFT in the Data Transmission Interval (DTI) access period. The impact of the latter depends mainly on the size of the codebook and the number of PAAs.

### 2.4.2 A-MPDU Aggregation

To evaluate the effect of the increased A-MPDU frame aggregation we evaluate the maximum throughput that can be achieved with a single bonded channel, varying the aggregation size. In addition to the maximum IEEE 802.11ay aggregation of 4 MB and the maximum IEEE 802.11ad aggregation of 262 KB, we also compare a medium aggregation of 2 MB and a very low aggregation of 16 KB. For this scenario, we have one AP and one STA and the maximum achievable SC, single channel data rate of 8 Gbps. We use the regular MAC implementation as one of the advantages of frame aggregation is that it reduces the MAC overhead associated with channel access. We average the results over 20 different STA locations and 2 different simulation seeds, resulting in 40 total simulations. Figure 2.10 shows the network performance in terms of the throughput Cumulative Distribution Function (CDF) measured in 0.1 s intervals (Figure 2.10a) and application delay (Figure 2.10b), comparing the different aggregation sizes. We can observe that the A-MPDU frame aggregation drastically alters the achievable performance. The low

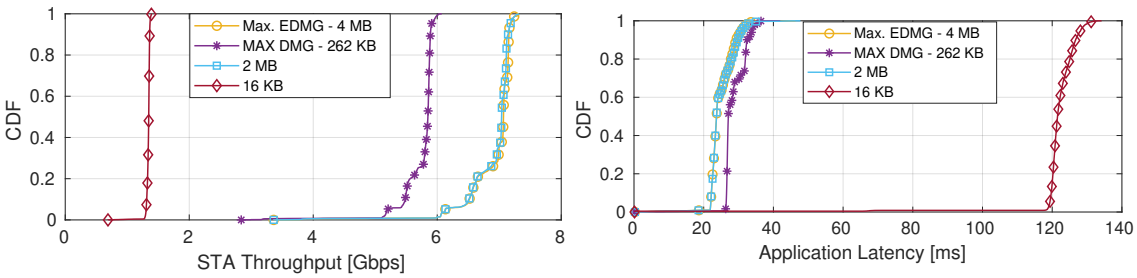


(a) Throughput CDF

(b) Application Delay

Figure 2.10: Performance for different A-MPDU aggregation sizes

aggregation of 16 KB only reaches 1.4 Gbps of traffic, compared with up to 7.3 Gbps with the maximum aggregation. Curiously, the 2 MB aggregation also reaches the same throughput. We found that this is because the maximum PPDU duration according to the IEEE 802.11ay standard is limited at 2 ms, and is already reached when aggregating 2 MB in a single channel, Single-Input and Single-Output (SISO) scenario. However, the larger aggregation can enable higher data rates when considering channel aggregation and bonding as well as MIMO transmissions. It is important to note that the achievable throughput even with the highest aggregation is approximately 0.8 Gbps lower than the theoretical throughput of IEEE 802.11ay. This is caused by the overhead of channel access with CSMA/CA, as well as the network operating overhead discussed in the previous section. We also highlight that the increased EDMG aggregation results in a gain of approximately 1.3 Gbps compared to what is possible with IEEE 802.11ad. The application delay results reveal very similar trends. Most packets with the 4 MB and 2 MB aggregation rates have a relatively low latency of 22 ms caused by the inability to reach the wanted data rate which forces packets to wait in the queue. The maximum IEEE 802.11ad aggregation has a rather similar delay with only a slight increase of approximately 5 ms. The lowest 16 KB aggregation, however, has a significantly larger delay of 120 ms, demonstrating how because of the very high data rates inmmWave, high frame aggregation is crucial to lowering the MAC overhead for efficient channel utilization.

### 2.4.3 Network scalability

Lastly, we present a first evaluation of IEEE 802.11ay performance with regard to density. For this scenario, we use the setup from Section 2.4.2 and investigate what happens when the number of STAs associated to the AP increases from 1 to 8. We keep the aggregate network load equal to 8 Gbps, distributing it equally between the STAs and using the full IEEE 802.11ay A-MPDU aggregation. We look at both downlink and uplink traffic to see how the MAC overhead and the level of contention in the network affect the performance.

Figure 2.11 shows the median network throughput with 95% confidence intervals,

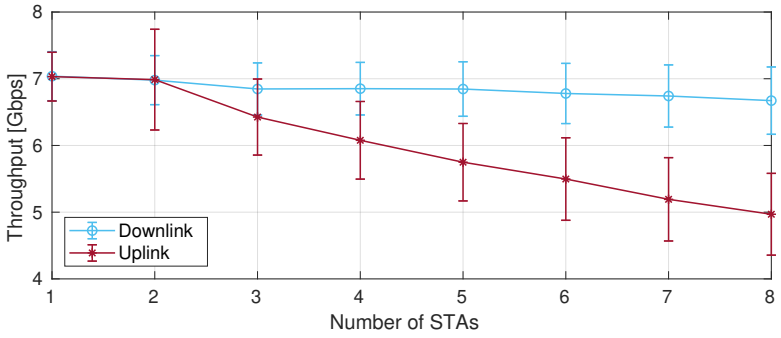


Figure 2.11: Median network throughput for different network densities

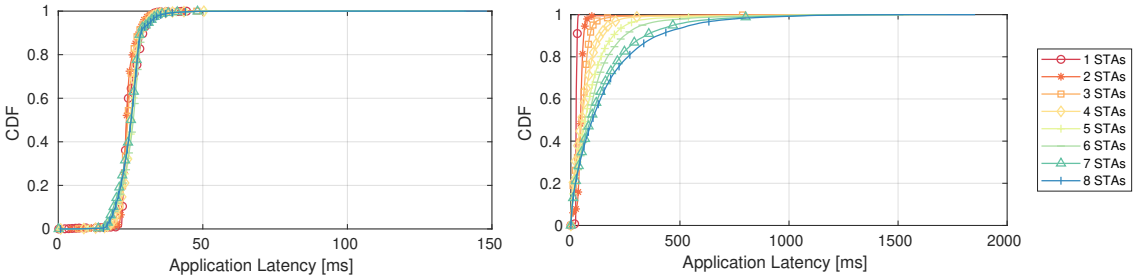


(a) Downlink (b) Uplink

Figure 2.12: Application delay for different network densities

varying the network density. We can see that with both uplink and downlink throughput the achievable data rate drops as the network size grows. The drop, however, is much more significant in the uplink case. In the downlink case, the additional MAC, channel access and BFT overhead reduce the achievable throughput by approximately 0.36 Gbps. However, since the AP coordinates the transmissions to all STAs there is no contention for medium access. In the uplink case, the STAs contend using the CSMA/CA protocol to transmit towards the AP. The increased backoff times and packet collision caused by carrier sensing failures further reduce the MAC efficiency and the achievable throughput drops to a median of 5 Gbps. Similar trends can be observed in the application delay, shown in Figure 2.12. In the downlink, the latency is contained around 20 ms for all network sizes tested. The only significant difference is that in the densest 8 STA scenario there are outliers which have increased latency up to 150 ms. This is because the AP serves all STAs equally and thus avoids large delays for all STA. In the uplink, however, we can see that the density has a larger effect on the latency and that the latency is much higher, up to 2 s. In this case, the latency consistently rises with the density as the level of contention significantly affects the delay.

We can see that even with a single AP, channel contention significantly limits the achievable throughput and latency in the network. mmWave networks are much more vulnerable to medium access problems as due to the short range and directional

communication they experience more frequent carrier sensing failures and packet collisions. More complex deployments with multiple APs introduce further challenges to ensure optimized spatial sharing. Efficient channel usage, interference management and scalability are key challenges to reaching the full potential of mmWave networks. Chapters 3 and 4 explore these problems in more detail, looking at how IEEE 802.11ay networks scale in dense deployments.

## 2.5 Conclusions and Future Work

This chapter presented our implementation of the IEEE 802.11ay standard in the network simulator ns-3. Our model builds upon the existing IEEE 802.11ad implementation, upgrading the module with a diverse set of MAC and PHY features including EDMG framing, MCSs, channel bonding, aggregation and novel BFT features like the EDMG TRN field and Short SSW packets. We demonstrated the maximum achievable throughput per spatial stream for each EDMG MCS for different channel configurations. Additionally, we looked at how frame aggregation and network density can significantly affect the network throughput and latency. The latter is a key motivation factor for this thesis and is explored in detail in the following chapters.

# 3 Beamforming Training for Dense Networks

## 3.1 Introduction

Beamforming Training (BFT) is one of the most crucial processes at Millimeter-Wave (mmWave) frequency as it determines the beamforming steering at the transmitter and receiver to align the antenna beams and thus overcome the large propagation loss experienced in the mmWave band. In dense networks, the design of the BFT protocol is even more important as a large number of devices have to be trained under a challenging, high-interference environment. Errors in the BFT result in a decrease of the link Signal-to-Noise Ratio (SNR) and can lead to a reduction in the achievable data rate and even service outage if the SNR drop is too large. Additionally, the time spent performing BFT procedures can also reduce the achievable data rate as no data can be transmitted during this time. As the number of users in the network grows, coordinating the BFT of all users and ensuring that users have timely and accurate BFT while keeping the BFT overhead low becomes quite challenging. Finding BFT protocols that scale well and are efficient in such scenarios is key to enabling high-performance dense networks.

One of the main drawbacks of IEEE 802.11ad is the large overhead and suboptimality of the Beamforming Training (BFT) process [40]. The BFT process in IEEE 802.11ay has been redesigned to be more efficient, flexible, and adaptive to cope with different deployments and heterogeneous end-devices. One of the most novel BFT procedures introduced in IEEE 802.11ay is Group Beamforming. This technique is a significant extension of the IEEE 802.11ad BFT paradigm. Instead of relying on a per Access Point (AP)-Station (STA) pair training, Group Beamforming allows simultaneous training of all STAs associated with a given AP, which drastically reduces the BFT overhead and improves the network performance. Additionally, Group Beamforming relies on Enhanced Directional Multi-Gigabit (EDMG) Training (TRN) fields for the STA training, rather than the legacy Sector Sweep (SSW) frames. This allows for a further reduction in overhead as multiple beampatterns can be trained in the same packet. Thus, Group Beamforming not only improves the BFT scalability but also provides an efficient and

simple solution to the complex problem of BFT in mobile and dense networks.

Very few papers have investigated IEEE 802.11ay BFT performance. In [40], the new BFT techniques introduced in IEEE 802.11ay are presented, with a special focus on the Beam Refinement Protocol (BRP) Transmit Sector Sweep (TXSS) (BRP TXSS) and asymmetric BFT performance. Asymmetric beamforming is also evaluated in [63] and [64]. Finally, Multiple-Input and Multiple-Output (MIMO) BFT overhead is studied in [65] and [66]. However, to the best of our knowledge, no investigation of IEEE 802.11ay Group Beamforming has been performed. A comprehensive analysis of the performance of Group Beamforming is necessary in order to understand the realistic possibilities for its future application. Moreover, this analysis can also provide valuable insights into the performance of BFT using TRN fields as compared with the legacy SSW frames and identify the benefits and drawbacks of the two approaches.

The major contributions of this chapter are as follows:

- We implement IEEE 802.11ay Group Beamforming in ns-3, based on the existing IEEE 802.11ay model [1].

- We perform a comprehensive analysis of the system-level performance of Group Beamforming against legacy IEEE 802.11ad BFT using typical IEEE 802.11ay use cases. We identify the deployments and configurations where Group Beamforming can vastly improve the network performance but also highlight situations where Group Beamforming can be detrimental to the overall system behavior.

- Based on our insights into Group Beamforming performance, we propose and evaluate different improvements. Our analysis shows that our modified version of Group Beamforming outperforms both the legacy 802.11ad BFT and the standard Group Beamforming and provides the best network performance in all scenarios we tested.

The structure of this chapter is as follows. Section 3.2 provides background on BFT in IEEE 802.11ay and its predecessor IEEE 802.11ad. Section 3.3 introduces our performance enhancements for Group Beamforming and Section 3.4 and presents our IEEE 802.11ay implementation. Section 3.5 contains our evaluation of the Group Beamforming mechanism and Section 3.6 presents a discussion of our analysis. Section 3.7 lists all the relevant state-of-the-art and finally, Section 3.8 concludes the chapter and summarizes our findings.

## 3.2 A Deep Dive into Beamforming Training

The mmWave band exhibits higher propagation loss than the sub-6 GHz band, requiring the use of beamforming to focus the energy in a specific direction and thus

increase the transmission and reception gain. This introduces the need for BFT procedures through which devices determine the optimal antenna steering directions in order to establish a directional link. The efficiency and accuracy of the BFT procedure is a key factor in the performance of mmWave networks as incorrect steering can cause link breakage. Moreover, BFT in large and dense networks is an even more complex challenge due to the dynamic wireless channel, the high-interference conditions that can disrupt the BFT and the high number of BFT procedures that need to be performed. IEEE 802.11ad and IEEE 802.11ay define several different BFT mechanisms, as explained below.

### 3.2.1 IEEE 802.11ad Beamforming Training

IEEE 802.11ad uses a two-stage process for the BFT. The first Sector Level Sweep (SLS) phase determines the initial beam patterns to enable communication, while the subsequent BRP phase allows further refinement of the beam alignment to obtain maximum gain. The SLS performs the BFT with a sequence of packets called SSW frames, each one transmitted or received with a different beampattern. Each SSW packet is used to train a single directional beampattern, and by measuring the quality of all received SSW frames, a STA can determine the optimal beampattern. The BFT requires two distinct SLSs (an initiator and responder SLS) performed sequentially to determine the optimal transmit (TX) configuration for a pair of nodes that want to communicate. The SLS procedure is simple and reliable, but it suffers from high overhead and limited scalability. Separate training is required in both communication directions and each beampattern is trained in a separate training packet. Therefore, the training duration grows with the number of beampatterns trained and the number of trainings needed grows with the network size. In order to simplify the beamtraining process, the SLS trains predefined beampatterns, referred to as sectors.

The optional BRP phase introduces special elements called TRN subfields. As discussed in Chapter 2 TRN subfields are composed of multiple orthogonal Golay sequences with good correlation properties and each subfield can be used to train a single directional beampattern. Multiple subfields are grouped in a TRN field which is added at the end of BRP training packets. The TRN fields of a packet can be used either for TX or receive (RX) training. For TX training, the transmitter switches beampatterns when transmitting each Transmit Training (TRN-T) subfield, while for RX training all Receive Training (TRN-R) subfields are transmitted with the same beampattern, but are received by switching between different beampatterns. Thus, multiple beampatterns can be trained within a single packet, reducing training overhead and improving performance, as multiple measurements can be done in a single packet in a coherent manner.

### 3.2.2 IEEE 802.11ay Beamforming Training

The IEEE 802.11ad BFT procedures provide robust BFT, however, they suffer from large overhead and do not scale well to dense and high-mobility networks. To overcome these challenges, IEEE 802.11ay introduces many changes to the BFT process of its predecessor, including support for new BFT procedures and a redesign of the TRN field. There are four approaches used for this purpose:

1. **Reducing SLS overhead and duration:** IEEE 802.11ay introduces short SSW frames, which eliminate the Medium Access Control (MAC) overhead of SSW packets, thus reducing the overall SLS overhead. It also allows the use of a partial SLS, which reduces the number of SSW packets needed to perform an SLS.

2. **Improving TRN field functionality:** IEEE 802.11ay significantly increases the number of TRN subfields that can be used in one field, and thus the number of beampatterns that can be trained in a single packet. Additionally, it introduces Receive/Transmit training (TRN-R/T) subfields that enable both TX and RX training in the same packet. Finally, it relaxes the switching time requirements for the antenna array by adding the possibility of transmitting or receiving multiple subfields with the same beampattern. A full explanation of the new TRN field structure can be found in Section 2.3.

3. **Reducing the BFT duration:** IEEE 802.11ay relies much more on TRN subfields rather than SSW packets, as the overhead is significantly lower and scales better. This includes the introduction of BRP TXSS where once a link is established, the BRP protocol is used for BFT. It is also reflected in the design of the new Single User MIMO (SU-MIMO) and Multi User MIMO (MU-MIMO) beamforming protocols which mostly train with BRP rather than SSW packets.

4. **Simultaneous multi-STAs BFT:** IEEE 802.11ay defines two new procedures, Group Beamforming and asymmetric beamforming which exploit antenna reciprocity and use TRN-R subfields to allow multiple stations to train *simultaneously*.

### 3.2.3 Group Beamforming

Group Beamforming is a promising novel technique that aims at increasing BFT efficiency. It can provide full BFT for a whole Basic Service Set (BSS) in a single step with very low training overhead. A BSS is composed of an AP and all its associated STAs.

Group Beamforming uses different approaches for training in the AP-STA and the STA-AP direction. Training in the AP-STA direction is similar to the IEEE 802.11ad

SLS phase, and reuses periodic beacon transmissions by the AP to train the AP TX sectors. These beacons are transmitted at the start of each Beacon Interval (BI) and serve for initial BSS discovery, distribution of BSS parameters and BI scheduling. To ensure that STAs all around the AP can detect the network, the beacons are transmitted sequentially in a directional manner and therefore can also function as TX SSW frames for the AP. In this way training in the AP-STA direction is done with no extra overhead, as the beacons are already a necessary part of network operation.

The novelty of Group Beamforming is in the STA-AP training which proposes adding TRN-R subfields to each beacon. All STAs in the BSS that receive the beacons can then use the TRN-R subfields to perform simultaneous RX training. Assuming antenna reciprocity at the STAs antennas, i.e., the TX and RX beampatterns are equivalent, the best STA RX beampattern will also be the best TX beampattern, which thus provides the full antenna configuration for communication with the AP. Note that this adds redundancy to the training, as STAs repeat the training with each received beacon. Since all STAs train with the same TRN-R subfields, the training overhead of Group Beamforming does not depend on the number of STAs in the network. Additionally, the TRN fields have a very low overhead - each subfield has a duration of only 437 $ns$. Whenever a change in the optimal transmit configuration for the AP has been detected, the STA informs the AP through an Information Response frame that contains the updated AP antenna configuration that should be used for communication with that STA.

Therefore Group Beamforming has a significantly lower overhead than other BFT approaches defined in IEEE 802.11ad and IEEE 802.11ay. To illustrate this, in Figure 3.1 we show a comparison of the training overhead for Group Beamforming and legacy SLS with respect to the number of STAs in the network and the antenna codebook size. In Figure 3.1 (a) we assume that the initiator and responder antenna codebooks both have 48 sectors and in Figure 3.1 (b) we calculate the overhead with one STA in the network. We can clearly see that while the SLS overhead grows with the addition of STAs in the network, the Group Beamforming overhead remains constant regardless of the size of the BSS, as it involves only the addition of TRN-R subfields to the beacons. Therefore it is a much more efficient and scalable approach for BFT in large and dense networks with multiple APs and many STAs. The jumps in the Group Beamforming overhead in Figure 3.1 (b) come from the TRN-R field structure. Namely, the TRN-R field is composed of multiple Units, each with 10 subfields for training. Therefore, every 10 additional sectors in the STA codebook require one extra TRN Unit for training. The additional growth between the jumps is because the TRN-R subfields are added to each beacon sent by the AP. Therefore, the Group Beamforming overhead also grows with the size of the AP codebook, although the growth is still much lower than the corresponding SLS overhead.

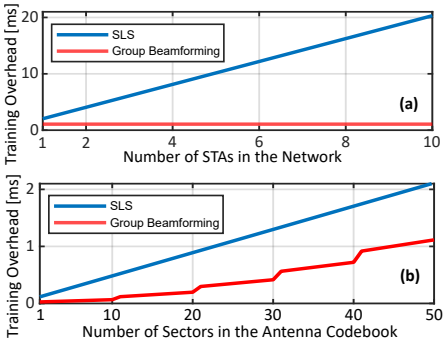


Figure 3.1: Overhead of Group Beamforming and SLS with different (a) number of STAs, (b) antenna codebook size.

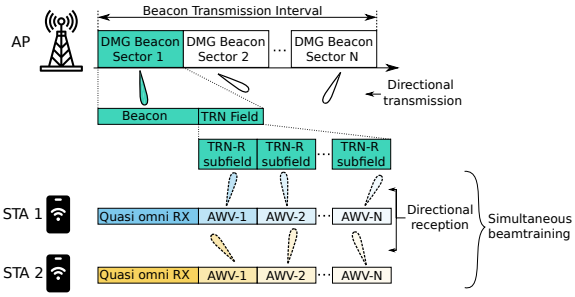


Figure 3.2: An example of Group Beamforming with 2 STAs.

Figure 3.2 shows an example of Group Beamforming training. The AP iterates over its TX beampatterns and appends TRN-R subfields to every beacon. The STAs receive the beacon with a quasi-omni RX beampattern (i.e, the widest beampattern), and then iterate over their RX beampatterns when receiving the TRN-R subfields, measuring the reception quality. Once every beacon and TRN-R subfield are received, the STAs are able to determine:

- the best TX beampattern for the AP by determining which beacon had the best reception quality.
- their best RX beampattern by determining which receive beampattern used had the best reception quality.
- assuming antenna reciprocity, their best TX beampattern and the best RX beampattern for the AP.

## 3.3 Group Beamforming performance enhancements

In addition to the standard-compliant version of Group Beamforming, we also design and implement different performance enhancements. As explained in Section 3.2.3, Group Beamforming has two different approaches for the AP and STA training. Our analysis in Section 3.5 shows that the STA training using TRN-R subfields is much more resilient to interference than the AP training using beacons, due to the ability to redundantly repeat the training with each received beacon. Therefore, we designed a version of Group Beamforming where the AP training also uses TRN-T subfields appended to beacons. In this case, the AP and STA training is done in two consecutive BIs. In the first one, the AP appends TRN-R subfields to each transmitted beacon and all STAs perform RX training and update their optimal antenna configurations. In the second BI, the AP appends

TRN-T subfields to the beacons and switches the TX beampattern when transmitting each subfield, while the STAs receive the subfields in quasi-omni mode and record the received Signal-to-Interference-plus-Noise Ratio (SINR).

We also explore the effect of directional transmission and reception on the RX training of the STAs when using TRN-R subfields. Normally, the BFT in IEEE 802.11ad and IEEE 802.11ay networks uses directional transmission and quasi-omni reception. To study how this affects the accuracy of the BFT decisions, we modified the transmissions of beacons with TRN-R subfields so that the AP switches from directional to quasi-omni transmission once the TRN field begins.

## 3.4 Implementation

We implemented IEEE 802.11ay in the system-level simulator ns-3, allowing us to evaluate high-density IEEE 802.11ay networks, taking into account different MAC layer aspects such as signalling overhead, channel access, packet collision, etc.

For our Group Beamforming implementation, we first enabled APs to add TRN-R subfields to Directional Multi-Gigabit (DMG) Beacons in the `DmgApWifiMac` class. The user can configure the periodicity of Group Beamforming by specifying the periodic interval, in number of BIs, at which TRN-R subfields are appended to the beacons. Then, in the `Codebook` and `DmgStaWiFiMac` classes we enabled STAs to sweep through their RX beampatterns while receiving the TRN-R subfields, measuring the received SINR. The STAs use these measurements to choose the best antenna configurations for communication with the AP. If they detect a change in the optimal TX sector of the AP, they send an unsolicited `Information Response` frame to the AP informing it of the change. Finally, to enable the modified versions of Group Beamforming we enabled APs to also append TRN-T subfield to DMG beacons and sweep through their TX beampatterns while transmitting the subfields as well as to switch to a quasi-omni beampattern once the TRN field starts. On the STA side it was necessary to only update the optimal AP configuration in the BI where TRN-T subfields were sent and only update its own antenna configuration when TRN-R subfields were sent.

We introduce the following parameters to configure the Group Beamforming training:

- `IsGroupTrainingInitiated`: To enable and disable Group Beamforming training.
- `TrnScheduleInterval`: This parameter controls the frequency of the Group Beamforming training as a number of BIs.
- `AntennaPatternReciprocity`: To enable support for antenna reciprocity.

- `UnsolicitedRSSEnabled`: Indicates whether the station can receive unsolicited Responder Sector Sweep (RSS) responses.

- `groupTrainingTRN-T:` To enable and disable the modified Group Beamforming training with TRN-T subfields.

- `quasiOmniEnabled:` To enable and disable the modified Group Beamforming training with quasi-omni TRN-R transmission.

## 3.5 Evaluation

We conducted a set of simulations to compare the performance of Group Beamforming and legacy SLS. For our simulations, we use the Quasi-Deterministic (Q-D) channel realization software [39], which uses a combination of ray tracing and statistical modelling to generate the channel model with all Line-Of-Sight (LOS) components and first-order reflections. We perform simulations in an indoor scenario with rectangular rooms of two sizes: a smaller room with dimensions 7.4 $m \times 13.5\ m \times 3\ m$ and a larger room with dimensions of 29.6 $m \times 54\ m \times 3\ m$. We simulate varying network densities with 1 to 8 APs and 1 to 64 STAs. The STAs are randomly placed within the room in a fixed position for the duration of the simulation. The association to the AP is according to the closest AP. However, we use load balancing to ensure an equal number of STAs per AP, meaning that in some cases STAs might not be associated with the closest AP. Unless otherwise specified, the number of STAs associated with each AP is 8. All APs are mounted on the ceiling at a height of 3 m while STAs are placed at a height of 1.2 m. Figure 3.3 shows an example network topology with 4 APs and 32 STAs. Due to the random node distribution, in some cases STAs can be close to each other and suffer from high interference. Additionally, STAs located far from the AP may experience low SNR. Finally, as mentioned, due to the greedy load balancing, some STAs may not be associated with the closest AP. These devices can suffer both from the large distance to the AP and the interference from the other close-by BSS.

All devices use 2x8 element Phased Antenna Arrays (PAAs) in the form of a Uniform Rectangular Array (URA) generated with the IEEE 802.11ad Codebook Generator [39]. This provides sufficient steering capability in the azimuth plane to separate users, as indoor scenarios do not require much separation of users in the elevation plane. The devices perform Group Beamforming or SLS each 10 BIs. In the SLS simulations, the training is initiated by the STA at the start of the Data Transmission Interval (DTI). The results are averaged over 40 simulation instances with a duration of 50 $s$, using 20 scenarios with different STA locations, each run two times with different random seeds. Additional relevant simulation parameters are shown in Table 3.1.

For the data communication, we use a rate adaptation algorithm that chooses the

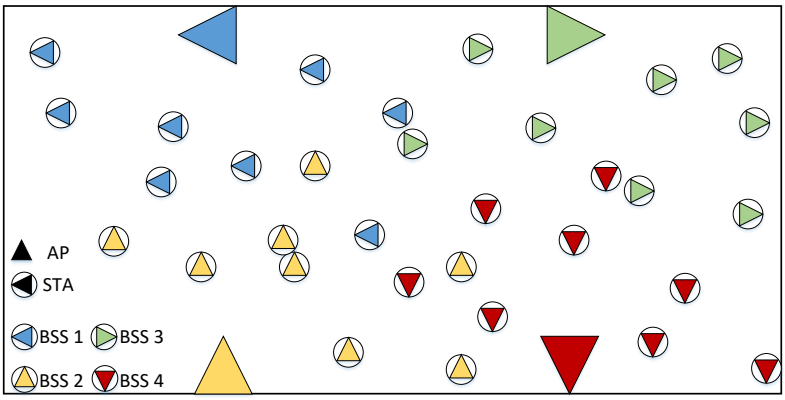


Figure 3.3: Example network topology for a 4 BSS network.

Table 3.1: Simulations Parameters for Group Beamforming evaluation

| **Parameter Name** | **Parameter Value** |
|---|---|
| Application Type | OnOffApplication |
| Payload Size | 1448 Bytes |
| MAC Queue Size | 4000 Packets |
| Frame Aggregation Type | A-MSDU and A-MPDU |
| MAC Protocol | CSMA/CA |
| RTS/DMG CTS Handshake | Disabled |
| Transmit Power | 10 dBm |
| Rx Noise Figure | 10 dB |
| Operating Frequency | 60.48 GHz |
| Number of Transmit/Receive Sectors | 48 Sectors |
| Simulation Time | 50 s |

Modulation and Coding Scheme (MCS) based on the last received packet SINR. We evaluate network performance with downlink (DL) and uplink (UL) User Datagram Protocol (UDP) traffic in order to determine both the effect that they have on the BFT and also the effect that errors in the BFT have on the throughput of the network. We vary the station data rate between 100 Mbit/s and 1 Gbit/s, allowing us to study the performance with different levels of channel saturation. Unless otherwise specified, the simulation data rate is set to 100 Mbit/s, as we found that this allowed us to study and compare networks of different densities without completely oversaturating the channel. Finally, although Group Beamforming enables communication with directional TX and RX beam patterns, we disable the directional reception and use quasi-omni reception when receiving data packets for a fair comparison with SLS.

We evaluate the BFT performance by assessing the accuracy and quality of the chosen beams. We define accuracy as the percentage of time that the optimal sector is chosen. We consider as optimal the sector that has the highest SNR in the absence of any interference, measured with directional transmission and quasi-omni reception, as this is the configuration most often used for data transmission. In addition, we look at the quality of the beams by calculating the SNR loss due to BFT errors as the difference between the maximum possible SNR of the optimal sector and the SNR of the actual sector chosen by the BFT. For this purpose, we look at pure SNR without interference. In this way, we compare the BFT quality based on the distribution of the SNR loss values and the distribution of the average SNR loss of STAs. We analyze the network performance in terms of the per STA and aggregate network throughput.

### 3.5.1 Performance in a Single BSS Network

#### 3.5.1.1 BFT performance

We begin the analysis by looking at simple scenarios with 1 AP in the network. The analysis of networks with a single AP is separated as this is the most favourable case for Group Beamforming compared to SLS.

This can be clearly seen in Figure 3.4 which shows the SNR loss due to BFT decisions for the large room with both UL and DL traffic of 100 Mbit/s and 8 STAs in the network. We show separately the SNR loss of the APs and the STAs since Group Beamforming uses different training approaches for training in the different directions. It is evident Group Beamforming is extremely accurate in choosing the optimal sector, resulting in SNR losses of zero. This is because Group Beamforming moves the BFT from the DTI, where typically STAs compete for access to the channel in a contention-based manner, to the Beacon Transmission Interval (BTI), where only the AP is allowed to transmit. Thus, the BFT is free from any interference caused by STAs within the BSS, regardless of the size of the BSS and the amount of traffic. This eliminates interference in the single

BSS case and can also improve BFT accuracy in multi-BSSs networks as STAs within our BSS are typically the closest and strongest interferers.

The few outliers that show an SNR loss slightly higher than zero come from an interesting effect where in very rare cases Group Beamforming never chooses the optimal sector. This is because Group Beamforming is performed with both the transmitter and receiver in directional mode, as described in Section 3.3. In very rare cases, this can result in different BFT decisions than the ones made with the transmitter in directional and the receiver in quasi-omni mode, as the RX beampattern affects the amplification or attenuation of different multi-path components. This can lead to sub-optimal BFT decisions since UL transmissions performed in a contention-based manner need to use directional transmission and quasi-Omni reception in order to enable the AP to receive from multiple STAs in the same contention period.

In low-interference scenarios the resulting SNR loss is low (less than 1dB in Figure 3.4) and there is no performance degradation. However, the effect can be more detrimental in dense, high-interference scenarios.

Analysing legacy SLS performance shows that interference can lead to BFT errors even in small networks. The exact SLS performance depends on the saturation of the wireless channel and we found that it degrades as the interference in the network increases. The results presented in Figure 3.4 show that the average SLS performance matches that of Group Beamforming as most BFT decisions are correct. However, we can also observe outliers with SNR loss of up 28 dB caused by SLS errors, leading to link instability and even loss of service.

Additionally, in Figure 3.4 we can also see that UL traffic causes more SLS errors than DL traffic and that the STA training is more vulnerable than the AP training. Both effects are due to the higher level of contention that STAs experience. When analyzing the effect of the room size, we found that, as expected, Group Beamforming remained perfect in both rooms. For SLS, we noticed a performance degradation in the large room. In the small room, the measured SLS accuracy was above 95%, while in the large room it went down to 75%. In addition to this, BFT errors in the large room also had a larger effect on network performance, as due to the larger node distances, even minor SNR loss can cause link outage. Lastly, we found that the highest SNR losses came from collisions with other SLS packets, rather than data packets. In these cases, the interference affects most or all SLS packets and constantly varies as the interfering STA sweeps through TX sectors, making it more difficult to find the optimal sector.

#### 3.5.1.2 Network performance

We analyze network performance in terms of achieved throughput to answer the following questions: (*i*) how do errors in BFT impact the throughput and (*ii*) how does the choice between Group Beamforming and SLS affect the overall performance of mmWave

networks?

When we analysed the network performance we found that the wireless channel is under-saturated with a data rate of 100 Mbit/s so most SLS errors are not reflected in the data throughput. In the small room, both Group Beamforming and SLS lead to optimal network performance of 797 Mbit/s and average STA throughput of 99.65 Mbit/s, with both UL and DL traffic. However, in the large room UL results, we found a slight performance degradation when using SLS. In some SLS scenarios the aggregate network throughput drops to 765 Mbit/s, while the average per STA throughput can go down to 88 Mpbs. This degradation is directly caused by errors from interference in the SLS BFT since the throughput performance in the exact same scenarios with Group Beamforming remains optimal.

#### 3.5.1.3 Performance under higher channel saturation

Next, we explore how increasing the channel saturation affects the BFT and the overall network performance. For this purpose, we analyse two additional single-BSS scenarios, one with a data rate of 1 Gbit/s (instead of 100 Mbit/s) and another with 16 STAs (instead of 8). Figure 3.5 compares the Cumulative Distribution Function (CDF) of the BFT SNR loss. We notice that the SLS accuracy drops in both new scenarios. It is also interesting to see that although the data rate increase adds much more traffic than the STA increase, the SLS accuracy is pretty similar in both cases. This is because increasing the number of STAs in the network increases both the training overhead and the collisions from carrier-sensing failure. We can also see that Group Beamforming still provides perfect BFT in both cases, while also keeping the overhead constant. The benefits of this can be seen in Figure 3.6 which shows a comparison of the aggregate network throughput with UL traffic in the large room. In both cases, the wireless channel is over-saturated and the desired data rate can not be achieved. However, in both cases, there is a gain when using Group Beamforming. This gain is on average 370 Mbit/s in the 1 Gbit/s case and 230 Mbit/s in the 16 STAs scenario and comes both from the perfect BFT and the lower overhead.

### 3.5.2 Performance in Networks with Multiple BSSs

#### 3.5.2.1 BFT performance

Networks with multiple BSSs introduce challenges to the BFT, particularly for Group Beamforming. Unlike the single BSS case, Group Beamforming is not protected from interference by STAs in other BSSs. Therefore we vary the number of BSS and the network density and study the effects on BFT performance. Figures 3.7 (a) and (b) show the SNR loss distribution for Group Beamforming and SLS respectively. The presented results are for simulations of the small room with 100 Mbit/s DL traffic and 8 STAs per AP. The trends observed also occurred in all other scenarios tested.

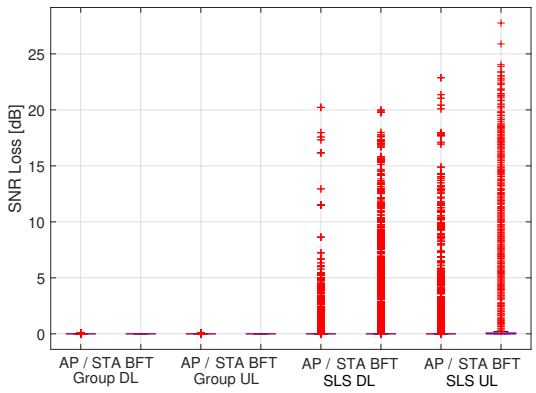


Figure 3.4: SNR loss from BFT choices, 1 AP, 8 STAs, large room.

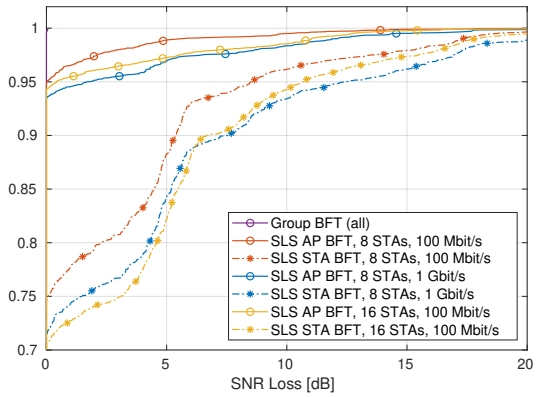


Figure 3.5: CDF of SNR loss from BFT errors, 1 AP, UL traffic, large room.

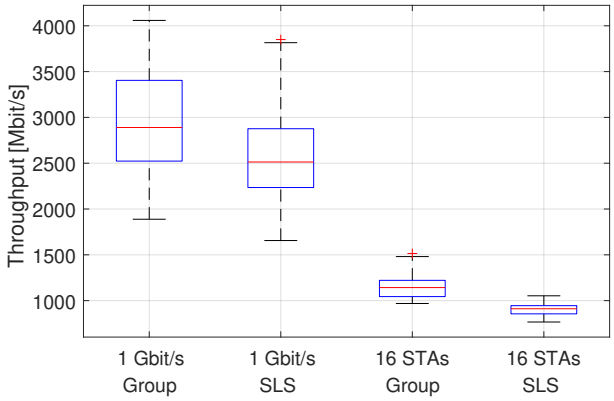


Figure 3.6: Aggregate network throughput for different BFTs, 1 AP, UL traffic, large room.

The first conclusion that can be drawn is that in high interference scenarios, the BFT procedure can be compromised, leading to the high SNR losses seen in Figure 3.7. Depending on the resulting SNR, this instability can lead to drops in the achievable data rate and even link outages. Secondly, in Figure 3.7 we observe an asymmetry in the AP and STA BFT performance, both for Group Beamforming and SLS. In order to better illustrate this asymmetry, we show in Figure 3.8 the CDF of the SNR loss for the 8 BSSs network. In the case of Group Beamforming, it is evident that the APs training is much less reliable and has the worst performance overall. For SLS, there is a less pronounced asymmetry where the STA training results in higher SNR losses. In many scenarios tested, we also found that the SLS STA training had a lower accuracy. This SLS asymmetry can be explained by the STA-initiated SLS procedure and higher channel contention that STAs experience. Additionally, most APs are located at the room's edges which exposes them to less interference. We found that this asymmetry in BFT performance is very slight and does not seem to be reflected in the network throughput performance.

In the case of Group Beamforming, however, the training difference is very large,

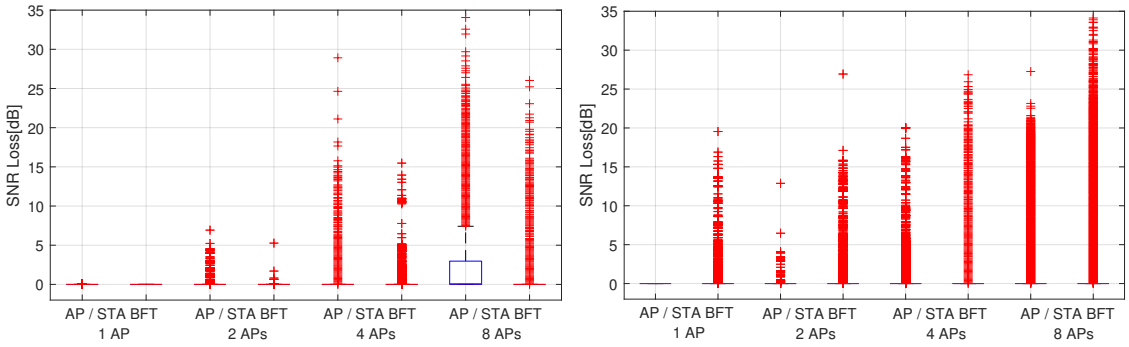


(a) Group Beamforming (b) SLS

Figure 3.7: SNR loss for different network densities, DL traffic, small room.

requiring further investigation. In Figure 3.7 it is evident that the Group Beamforming STA training has the least amount and lowest SNR losses, compared both with the Group Beamforming AP training and the SLS training. In fact, if we look at the accuracy of the STA training in Figure 3.8, we can see that it is above 97%, which shows remarkable robustness and stability. This is particularly impressive if we consider that Figure 3.8 shows a very dense, high-interference scenario, with 8 APs and 64 STAs distributed in a relatively small room. Our simulations showed the same behaviour across all other simulation scenarios. Although we found that UL traffic was slightly more disruptive, the Group Beamforming STA accuracy was always above 93%, in both the large and small room scenarios. For comparison, the worst SLS accuracy we measured was 70% and 68% for the STA and AP training, respectively. The Group Beamforming AP training, on the other hand, suffered much more in high-interference scenarios. In Figure 3.8 we see that the accuracy is 48%, far below the SLS accuracy and that this is also reflected in Figure 3.7, where the Group Beamforming AP training is the only one with an upper quartile above 0 dB. In addition to this, when we looked at the average per STA SNR loss, the Group Beamforming AP training had by far the worst performance. The mean SNR loss over the course of the simulation reached up to 13 dB, while it did not go over 5 dB for the SLS training or the Group Beamforming STA training.

This stark difference in the Group Beamforming performance comes from the fundamentally different approaches to the AP and STA BFT. The TRN-R subfields used for STA training are appended to each transmitted beacon, so STAs which receive multiple beacons perform the RX training multiple times, leading to increased reliability. Additionally, training with TRN subfields allows for all RX beampatterns to be tested in a single packet. This considerably reduces the probability that there will be significant interference variations during the training and enables obtaining coherent measurements, which can also increase the performance. This results in highly resilient BFT, which is not only robust to different types of interference but also consistently outperforms legacy SLS. In contrast to this, the AP BFT, done with directionally transmitted beacons, seems

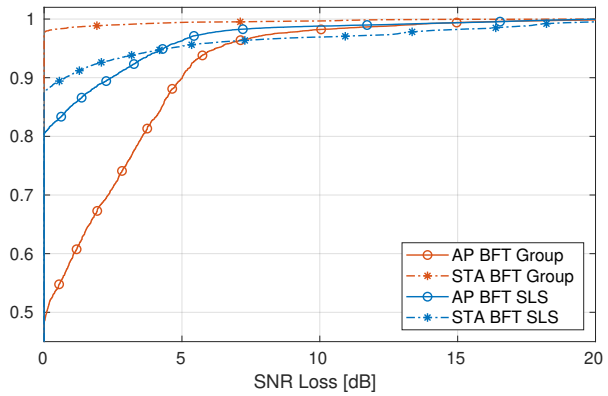


Figure 3.8: SNR loss CDF in a network with 8 APs, DL traffic, small room.

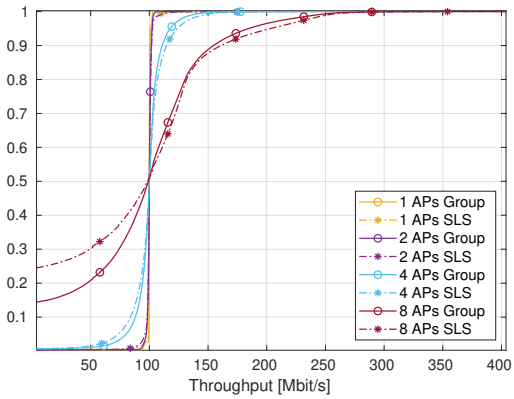


Figure 3.9: Throughput CDF for different BFTs, DL traffic, small room.

to be very sensitive to interference and more vulnerable than SLS. This is a consequence of the simultaneous training of all STAs in the BSS introduced by Group Beamforming. This greatly reduces the training overhead but also means that BFT decisions for the whole BSS are done based on the same beacon frames. Interference during the beacon transmissions will therefore affect the AP's training with all STAs in the BSS. Errors in the SLS procedure, on the other hand, are detrimental solely for the AP-STA pair performing the procedure.

Next, we analysed the effect of the room size on BFT performance. We noticed a performance degradation in the larger room, which was stronger in lower-density networks and decreased with the number of devices in the network. As the average distance between devices is higher in the large room, there are more collisions due to carrier-sensing failure leading to BFT errors. The effect of these errors is more evident in lower-density networks where the wireless channel is not very saturated. Finally, when comparing the BFT performance with DL and UL traffic we found that UL traffic tends to be more harmful. This is due to two reasons. Firstly, in DL transmissions, there are fewer devices transmitting and therefore the collision probability is smaller. Secondly, most of the APs are located at the room's edges and cause lower interference than the STAs which are distributed inside the room. This effect is more pronounced in Group Beamforming but is also present for the legacy SLS scenarios. The bigger effect can be explained by the fact that some of the errors in legacy SLS come from collisions of the SLSs themselves and these remain unaffected by the change from UL to DL transmissions, while in Group Beamforming all errors are caused by data transmissions.

#### 3.5.2.2 Network performance

Figure 3.9 shows the CDF of the DL throughput measured in 0.1 s intervals, varying the network density and BFT method. The throughput performance here corresponds

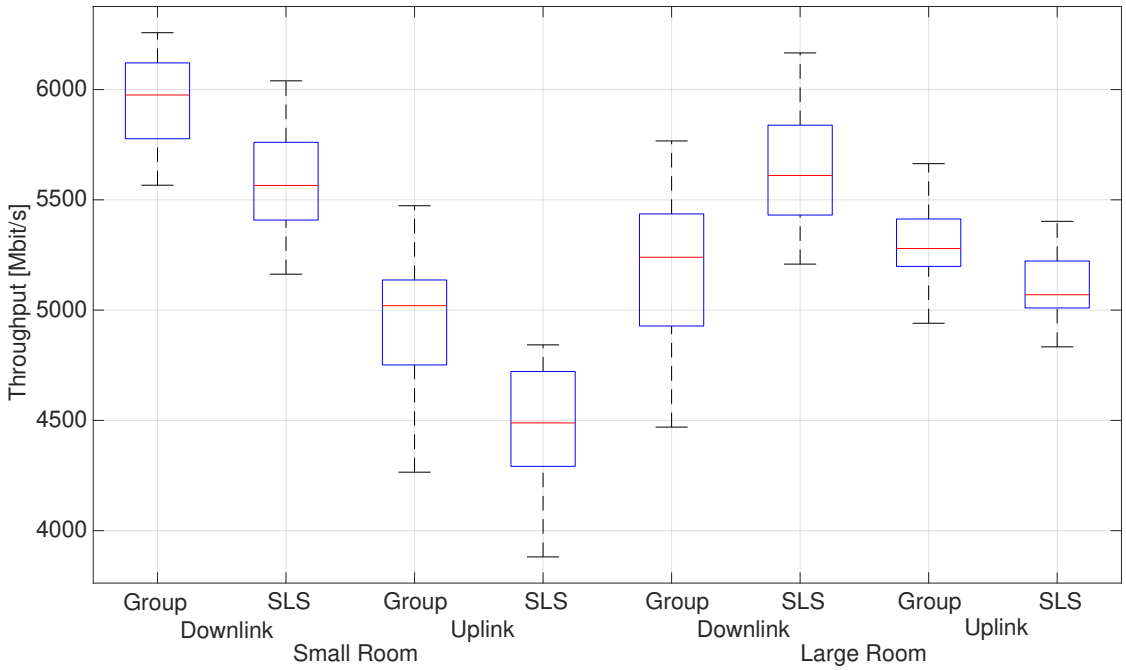


Figure 3.10: Aggregate network throughput for different BFTs, traffic type and room size, 8 APs, 64 STAs.

to the BFT performance shown in Figure 3.7. We can clearly see how increasing the wireless channel saturation affects STA performance. In the 1 and 2 AP networks, the performance is extremely close to optimal and there is no discernible difference between Group Beamforming and SLS. However, in the 4 AP network we begin to see a slight performance gain in using Group Beamforming, and this gain grows in the very dense 8 AP network. In the 4 and 8 AP networks, SLS causes more throughput values in the 0 to 100 Mbit/s interval, reflecting the time spent on SLS instead of transmitting data. In very dense networks, the extra time for data transmissions given by Group Beamforming can lead to significant increases in network performance. In Figure 3.10, we show the aggregate network throughput for the densest 8 AP, 64 STA scenario, comparing the performance with Group Beamforming and SLS. We can see that in the small room, DL traffic scenarios, Group Beamforming outperforms SLS by 400 Mbit/s on average.

This gain does seem somewhat surprising if we consider that Group Beamforming had significantly worse BFT performance. However, due to the closeness of devices in the small room, BFT errors rarely cause link outages and mostly result in a reduced MCS. Therefore, we have found that for the small room, the overhead reduction of Group Beamforming benefits network performance, despite the reduction of BFT accuracy. The overhead gains increase with the network density as SLS overhead grows, counterbalancing the BFT accuracy reduction.

As we can see from Figure 3.10, this effect is more pronounced with UL traffic, leading to an average increase of 500 Mbit/s in aggregate network throughput. Contributing to the

better UL performance is the fact that errors in Group Beamforming are overwhelmingly done for the AP sector, while the STA sector is chosen extremely accurately. When considering UL transmissions done in a contention-based manner, the STAs are in directional TX mode, while the AP is in quasi-Omni RX mode. Therefore, the wrongly chosen AP sectors have little influence on network performance, as they are only used to send back Block Acknowledgments, sent with a low MCS that is decodable even at low SNR.

We also found that in the small room the throughput performance was predominantly affected by the level of wireless medium saturation and failures in carrier-sensing caused by directional transmissions. For this reason, in these types of scenarios, Group Beamforming is a very good option due to the reduction in overhead that it provides.

In the large room, however, devices are much further apart. This results in lower optimal SNR and fewer sectors that still provide satisfactory SNR for communication. SNR losses caused by errors in BFT often mean link breakage until the device can recover from the error. In addition to this, there is more potential for spatial re-use, so the throughput is not as affected by collisions. Therefore, we found that the throughput performance in these scenarios had a significant connection to the BFT performance. To highlight the difference between operation in the small and large room, we show in Figure 3.11 a throughput comparison of two STAs that both experience severe SNR losses due to BFT errors. The network density (8 APs, 64 STAs), data rate (100 Mbit/s DL traffic) and BFT method (Group Beamforming) are the same for both scenarios and the only difference is the size of the simulated room. In both scenarios, BFT errors in the DL direction lead to drops of 10 dB to 15 dB in the achievable SNR for a significant portion of time. In the small room, we found that the errors were caused by the fact that the STA was located in the center of the room and experienced high interference from surrounding devices. In addition to this, due to collisions, some of the BFT feedback packets were not correctly received, and as a consequence, the AP was not able to update its antenna configuration in time. In the large room, due to load balancing, the STA was not associated to the closest AP and the large distance caused problems in the BFT. When observing the effect of the drop in SNR on the throughput, we can see that in the small room, the data transmissions are not affected, as the SNR is still well above the decodability limit. In the large room, however, we see that multiple BFT errors cause communication link breakage, as the SNR is too low for successful data transmissions.

We can see the effects of this in Figure 3.12, which shows the DL performance in the large room. In contrast to Figure 3.9, where we observed that Group Beamforming provided better performance than SLS, we can see that in the 4 and 8 AP scenarios, the network performance with Group Beamforming suffers due to the poor BFT performance. These were the only scenarios where we found that SLS outperformed Group Beamforming. In Figure 3.10 we can see that there is a difference of approximately

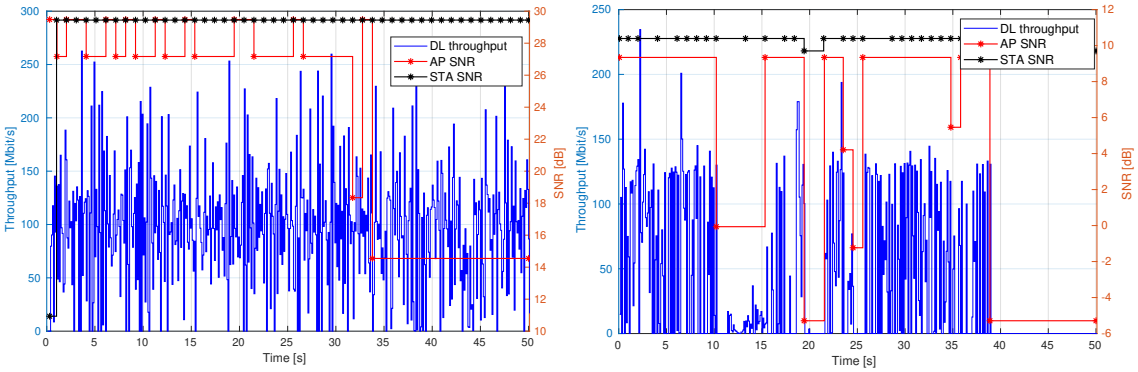


(a) Small room scenario

(b) Large room scenario

Figure 3.11: Throughput and SNR over the course of the simulation, 8 APs, DL traffic, Group Beamforming.

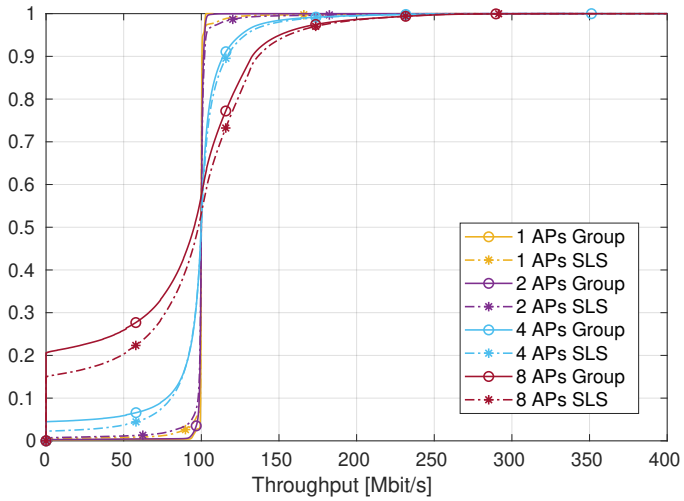


Figure 3.12: Throughput CDF for different network densities, DL traffic, large room.

400 Mbit/s in the 8 AP network. The UL results show the opposite trend, as the high accuracy of the STA training combined with the higher level of network contention result in better performance than SLS. The performance gains, however, are lower than in the small room and only go up to 200 Mbit/s. This is a combination of the more severe effects of the Group Beamforming BFT errors and the increased spatial re-use which reduces the benefit of the overhead reduction.

Similar to the BFT performance, we found that in general UL transmissions had worse throughput performance. The exception to this was the performance in the Group Beamforming, large room scenarios, with 4 and 8 APs, as can be seen in Figure 3.10. Due to the reasons discussed above, in these cases, the DL performance suffered significantly and was equal to the UL performance.

Finally, a direct comparison of the results in the small and large room for the same scenarios showed that the small room had better throughput performance due to the link

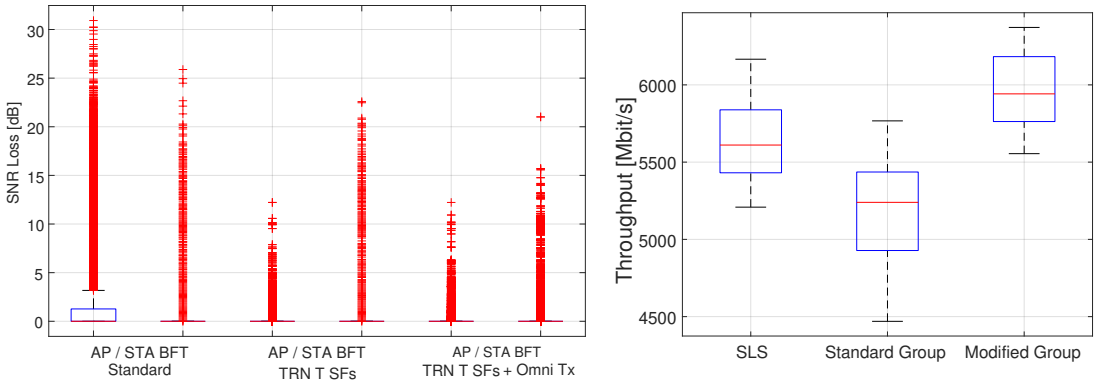


(a) SNR loss due to errors in BFT

(b) Aggregate network throughput

Figure 3.13: Performance with modified versions of Group Beamforming, 8 APs, DL traffic, large room.

breakages caused by BFT errors and the lower achieved SNRs which resulted in lower MCSs used. The exception to this were the UL scenarios with 8 APs in the network, as seen in Figure 3.10. In these cases, the wireless medium was over-saturated and therefore the higher spatial re-use in the large room caused gains in performance.

### 3.5.3 Performance with modified versions of Group Beamforming

Our performance analysis of Group Beamforming detected several causes of BFT errors, as described in the above sections. Therefore, we propose modifications that deal with these issues. The main performance limitation was the vulnerability of the AP training in high-interference scenarios. The contrast with the extremely robust STA training motivated us to implement a similar approach for the AP training. As explained in Section 3.3, we modified the AP training to use TRN-T subfields rather than rely on the beacon packets. This doubled the Group Beamforming overhead but significantly improved the training accuracy. In addition to this, we also attempted to further improve the STA training, by eliminating having both the transmitter and receiver in directional mode. Since this sometimes led to errors, we modified the APs to transmit the TRN-R subfields with a quasi-omnidirectional beampattern.

Figure 3.13 shows a comparison of the performance of Group Beamforming with the different proposed versions both in terms of the SNR loss and aggregate network throughput. The results show the performance in a network with 8 APs and DL traffic in the large room. As described in the above section, this was one of the scenarios where Group Beamforming had poor performance and was outperformed by SLS.

We can clearly see that our modifications significantly improve the BFT performance, particularly in terms of the AP training accuracy. The accuracy increased from 60 % to 95 % and the highest SNR loss measured dropped from 30 dB to 12 dB. In addition, we noticed that the overhead increase did not cause any throughput degradation, as it

was offset by the better BFT. In fact, the BFT improvements resulted in significant throughput gains, as seen in Figure 3.13 which shows an average increase of 700 Mbit/s in the aggregate network throughput. The STA BFT accuracy increase was much smaller, as the STA training was already quite close to optimal, resulting in only minor throughput gains. However, this modification also decreased the SNR losses, leading to fewer link breakages and an improvement in the quality of service.

We found that with our performance enhancements, Group Beamforming was able to outperform SLS in all tested scenarios, both in terms of BFT accuracy and achieved throughput. The accuracy rose to above 95% in all scenarios except in networks with 8 APs and UL traffic. In these scenarios the accuracy was 88% in the small room and 90% in the large room, outperforming the SLS accuracy by 20% and 13%.

## 3.6 Discussion

There are several lessons we learned from our analysis, which is the first known to evaluate the performance of Group Beamforming, both regarding BFT procedures and the performance of dense mmWave Wireless Local Area Network (WLAN) networks.

First, we found that although all mmWave networks suffer from high path loss and require directional communication, the room size determines how much the BFT accuracy affects network performance. In larger rooms, where the distance to the AP is big and there are no strong reflections, only a few sectors offer a sufficiently high SNR for communication and extremely accurate BFT is crucial for good performance. Errors in the BFT lead to link outage and it can take a long time until the STA recovers. In contrast, in smaller rooms, issues with carrier-sensing failure cause much more performance problems than BFT failure. As long as the SNR remained high enough to maintain a communication link, we found that network performance was not affected by BFT errors as even the lower MCSs can offer high data rates. However, depending on the nature of the application these errors may affect the quality of service and the user experience.

Another observation was that carrier-sensing failure can be a significant problem in dense mmWave networks that limits the network performance. In many cases, although the wireless channel can handle more traffic, it is limited by collisions caused by the failures, which are much more common in mmWave networks due to the directional communication. This is especially a problem in UL communication, where there is a lot of channel contention. More research is necessary to find solutions for this problem and improve the channel access procedure for mmWave WLAN networks. In Chapter 4 we present our contribution to this problem which uses directional reception and packet filtering to lower the packet collisions and improve the efficiency of the network.

When considering BFT performance, we found that even in very dense networks, the BFT was mostly successful with both approaches tested. BFT errors were often transient

and devices were able to recover in the next training. However, these errors could lead to significant drops in achievable SNR, leading to issues with the communication link.

The legacy SLS mechanism was shown to be quite robust to different levels of interference. However, the high training overhead affected the overall network performance and therefore this approach is not suitable for high-density networks.

In terms of Group Beamforming, our analysis showed that it is a very promising approach for WLAN networks in several ways. First, shifting the BFT to the DTI ensures complete protection from interference from STAs in the same BSS, providing perfect BFT with lower overhead. Therefore, single-BSS networks benefit significantly from its use. Moreover, the use of TRN-R subfields in Group Beamforming was proven to be extremely efficient and reliable. The drastically lower overhead offers the possibility to add redundancy and repeat the training multiple times at a fraction of the overhead needed with full packets. This resulted in extremely high BFT accuracy which outperformed the other approaches tested. We did find, however, that the simultaneous training of all STAs in the BSS has some drawbacks. As all STAs make their BFT decisions based on the same transmissions, training corrupted by interference can cause errors for all STAs in the network, leading to degradation in the performance of the whole BSS. Switching the AP training to the TRN approach, however, fixed this problem and significantly improved the performance. It is important to highlight that the standard-compliant version of Group Beamforming might have problems in high-density scenarios, particularly for DL traffic. Moreover, as typically the gains of the antennas in APs are much larger than the gains of the antennas in user devices, the AP BFT errors can cause more harm to the overall network performance. Lastly, the training overhead reduction can improve network performance as compared to legacy SLS, even when the BFT suffers more.

## 3.7 Related Work

Most existing work on IEEE 802.11ay is theoretical and either provides a protocol overview or proposes solutions to some open design issues.

In [7] the authors give an excellent overview of IEEE 802.11ay, while in [67] the authors focus on detailed technical explanations of aspects like spatial sharing, channel bonding and interference mitigation.

Some work focuses specifically on the BFT aspects of IEEE 802.11ay, such as [68], where the authors study scheduling, beamforming, and link maintenance protocols. Lastly, the authors of [40] explain the legacy IEEE 802.11ad BFT and show theoretical results on asymmetric beamforming. The authors in [69] propose a dynamic scheduling mechanism for asymmetric BFT based on SINR with random backoff. The authors in [70] focus on dense scenarios and propose an enhanced BFT which is backward compatible with IEEE 802.11ad. Their implementation divides the Association Beamforming Training

(A-BFT) slots to further improve the association. Another backward compatible BFT solution is presented in [71], where the authors extend the IEEE 802.11ad channel model to support Hybrid MIMO configurations.

Practical analysis of IEEE 802.11ay is out of the scope of these papers, which limits the potential to find new open research issues. To the best of our knowledge [72] is the only paper with system-level validation regarding IEEE 802.11ay. This paper focuses on channel bonding in IEEE 802.11ay. The authors briefly mention the implementation of an ns-3 module to validate their theoretical work, but the implementation details are unclear and the source code has not been released. The authors leave out some important considerations about the implementation, such as the BFT mechanism used. Additionally, there are many simplifications in their implementation including simple antenna models and throughput calculation from the Shannon equation. Our model is more complete as discussed in Chapter 2.

The beamforming training protocol is at the core of any mmWave system, thus, designing an efficient, scalable, and agile beamforming training mechanism has a huge impact on network performance and the perceived quality of service. For this reason, multiple measurement campaigns have evaluated the performance of IEEE 802.11ad BFT in Commercial Off-the-Shelf (COTS) devices. The authors in [73] were the first to provide an in-depth analysis on the performance and operations of WiGig compliant COTS devices, but lack an analysis of the beamforming mechanism. Authors in [74] performed a measurement campaign to extract the beam patterns of the first IEEE 802.11ad enabled COTS router from TP-LINK using their 60 GHz anechoic chamber. In addition, they implement a compressed sensing solution on the COTS routers that provides a significant reduction in training overhead compared to plain IEEE 802.11ad BFT. The authors of [75] provided a comprehensive evaluation study for the aforementioned router for various deployment settings. In their evaluation, they focused on the interplay between BFT and the rate adaptation algorithm as well as the impact of BFT errors due to mobility or blockage on the mmWave link stability. Unlike the previous papers which focused on a single device, the authors in [17], provided a comprehensive study regarding the performance of the aforementioned router in a large deployment. In particular, they assessed the impact of contention-based BFT on the performance of Transmission Control Protocol (TCP) protocol in a large network. Additionally, they evaluated the degree of spatial sharing that can be achieved using Qualcomm phased antenna arrays.

To the best of our knowledge, our work is the first to focus on IEEE 802.11ay Group Beamforming and evaluate its performance in various scenarios and network configurations.

## 3.8 Conclusions

In this chapter, we presented a study on BFT in dense networks, looking at the accuracy of the training and the effect of the training overhead on the throughput performance of the network. We decided to focus on a novel BFT technique introduced in the IEEE 802.11ay amendment called Group Beamforming as it introduced the concept of simultaneous training of all STAs in the BSS, which is very suitable for dense networks. For this purpose, we extended the existing ns-3 mmWave module to support Group Beamforming and evaluated its performance in a variety of simulation scenarios. Our evaluation showed that Group Beamforming can boost network performance by lowering the training overhead of legacy BFT procedures like SLS. We found, however, that the AP training faced challenges in high-density networks. Therefore, we implemented and tested two modifications of Group Beamforming that significantly improved its performance. With the suggested modifications, the new Group Beamforming approach was able to outperform the legacy SLS procedure in all tested scenarios, both in terms of BFT and throughput performance. Therefore, we consider it an extremely promising solution for BFT in dense mmWave networks.

# 4 SIGNiPHY: Reconciling Random Access with Directional Reception for Efficient mmWave WLANs

## 4.1 Introduction

The potential for Millimeter-Wave (mmWave) networks to provide high-throughput, low-latency performance makes them key enablers for emerging applications like Augmented Reality (AR)/Virtual Reality (VR), remote surgery, connected homes and factory automation. By including advanced physical (PHY) layer technologies like channel aggregation and Multiple-Input and Multiple-Output (MIMO) IEEE 802.11ay promises peak rates of up to 100 Gbps [7]. However, reliability and coverage remain challenging in mmWave, due to the increased attenuation and oxygen absorption [76]. In addition, multiple evaluations of Commercial Off-the-Shelf (COTS) mmWave WiFi devices [16–18, 32–34] reveal that mmWave efficiency drops significantly in multi-Access Point (AP) deployments due to problems with packet collisions and low spatial sharing. This matches our previous evaluations in Chapter 2 and 3 where we observed that the achievable data rate in IEEE 802.11ay was highly limited by the size of the network and the number of contending users. Therefore, we decided to further study channel access and multi-AP operation to find the causes of inefficiency and performance limitations.

We found that one of the main limitations for current mmWave WiFi networks came from the use of omnidirectional Beam Patterns (BPs) for reception. While highly beneficial, directional reception is incompatible with the random medium access mechanisms used in WiFi devices. Currently, mmWave WiFi devices implement legacy Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) [17,77], a contention-based random access scheme without a fixed transmission schedule. Instead, Stations (STAs) that want to communicate with the AP can transmit at *any* time, as long as they perform carrier sensing of the wireless medium to ensure that the channel is free. Since the AP does not know which user will transmit at which time, it needs to listen to *every* direction where users are located. This is why current COTS devices [16,77] only support directionality on one side of the link and use directional transmission but omnidirectional

BPs for the reception. While STAs usually only communicate with a single AP, the same issue may arise at STAs in case of multi-hop or device-to-device communication.

Omnidirectional reception harms both the single-link and aggregate network performance. The single-link performance is harmed as the lower antenna gain reduces the coverage range and achievable data rates. Additionally, it negatively affects spatial reuse, as it leaves devices vulnerable to interference coming from any direction. This makes it more difficult to successfully establish concurrent transmissions in a given area, which leads to low aggregate rates.

We also found that a particularly harmful effect was that it significantly increases the probability that a device will overhear packets that are not meant for it. While it is busy decoding the unwanted packets it might miss concurrent packets for which it *is* the intended receiver. This type of unwanted packet overhearing is a well-known problem caused by the use of a random access mechanism. Since there is no transmission schedule devices might receive packets at any time and thus need to try to decode all packets they detect. In traditional sub-6 GHz WiFi, this only lowers the energy efficiency since the device wastes energy on decoding useless packets. In mmWave WiFi, however, this limits the spatial sharing in the network by preventing concurrent transmissions. Thus, it has a much more harmful impact on network performance and can significantly limit dense networks.

Existing solutions fail to efficiently address these problems, in part as they are not specifically designed for mmWave networks. To the best of our knowledge, the only current method to enable directional reception is through mechanisms like Ready-to-Send (RTS)-Clear-to-Send (CTS) and CTS-to-self. These mechanisms protect data packets with preceding control packets which carry the transmitter identity in advance of the data. While the control packet would still need to be received with a quasi-omnidirectional BP, once the transmitter has been identified the correct directional receive BP can be used for the data packet. However, the overhead introduced due to the additional control message exchanges is non-negligible and translates into sub-optimal channel usage and low Medium Access Control (MAC) efficiency [32, 77]. Furthermore, control packets are very robust and can be overheard even by far-away nodes, which can additionally reduce the spatial re-use in the network.

Instead, in this chapter we present a mechanism that addresses these problems at the PHY layer without introducing *any* overhead, while at the same time ensuring backwards compatibility with legacy devices. SIGNaling in the PHY Preamble (SIGNiPHY) enables efficient directional communication by embedding a device identifier in the PHY packet preamble for early transmitter identification during the preamble reception, allowing for the packet payload to be received directionally. Critically, the identifier can *only* be embedded in the packet preamble, as channel estimation is done after the preamble has been received. The same BP has to be used during channel estimation and packet payload

reception, otherwise, packet decoding fails as the channel varies with the BP used. The packet preamble, however, can be received with a different (i.e., omni-directional) beam pattern without impacting the decoding.

In order for this design to work, however, preamble decoding, identifier decoding and BP switching all have to finish before channel estimation begins. To meet these stringent timing constraints, SIGNiPHY implements a highly efficient identifier decoder and a fast beam-switching mechanism once the user is identified. To ensure backward compatibility with legacy devices, SIGNiPHY retains the correlation properties of the original preamble, allowing the use of the same receiver processing blocks without *any* modification. Moreover, SIGNiPHY and legacy devices are also interoperable, as stations which can not detect the embedded identifier will simply receive the packets omnidirectionally since preamble functionalities are not affected.

The early user identification also allows SIGNiPHY to end the reception of packets which are not meant for the device. While such filtering could also be done after the MAC address decoding [78, 79], doing it during the PHY preamble is more efficient and faster.

We tested SIGNiPHY real-time operation in an FPGA-based mmWave testbed operating at 60 GHz and found that our implementation can identify the packet transmitter only 160 ns after packet detection. The user identification had an accuracy of 100% for Signal-to-Interference-plus-Noise Ratios (SINRs) higher than 7 dB and 99.5% for SINRs above -6 dB. Furthermore, we show that the packet decoding benefits from the Signal-to-Noise Ratio (SNR) boost due to directional reception and that SIGNiPHY is very efficient at coping with multiple interferers. Finally, we implement SIGNiPHY and several baseline solutions (RTS-CTS, CTS-to-Self, MAC filtering) in our IEEE 802.11ad/ay ns-3 module to evaluate its performance in dense scenarios. The results show that SIGNiPHY outperforms the other solutions and consistently improves uplink network throughput with gains of up to 230%, depending on the specific scenario and baseline against which it is compared.

To sum up, in this chapter we make the following contributions:

- SIGNiPHY is the first system that enables directional reception and early packet filtering in IEEE 802.11ad/ay WiFi networks with no additional overhead.
- We further manage to design SIGNiPHY to not modify the key characteristics of legacy systems to ensure backward compatibility and interoperability.
- We implement and validate SIGNiPHY in an FPGA-based testbed with a low-complexity architecture that makes it easy to implement in commercial devices.
- We integrate SIGNiPHY in our ns-3 IEEE 802.11ay model for validation in dense scenarios and compare the performance against different baseline solutions.

The rest of the chapter is organized as follows: Section 4.2 gives a detailed insight into the motivation for SIGNiPHY and Section 4.3 presents the SIGNiPHY protocol design. Section 4.4 then explains the SIGNiPHY preamble embedding mechanism and hardware implementation. Section 4.5 and Section 4.6 show the experimental and simulation evaluation, respectively. Section 4.7 contains a discussion of different aspects of SIGNiPHY performance, Section 4.8 gives an overview of related work and finally, Section 4.9 concludes the chapter.

## 4.2 Motivation

Although mmWave WiFi networks show great promise, multi-AP deployments where interference and collisions lead to reduced throughput and inefficient channel usage prove challenging for current COTS devices [16–18, 32–34]. The mismatch between what is envisioned by mmWave standards and what is achieved in practice lies in the fact that multi-Gbps data rates are only possible under the assumptions of 1) narrow directional beams with high gains on both sides of the link and 2) high spatial reuse with multiple concurrent data transmissions. Currently, mmWave networks can not fully implement either, leading to losses in performance. As discussed in Section 4.1 omni-directional reception and the attempted decoding of all detected packets are the main problems fundamentally limiting current networks. We now investigate them in more detail in Section 4.2.1 and analyze existing solutions in Section 4.2.2.

### 4.2.1 Performance analysis of directionality and unwanted packet overhearing

To validate our analysis of the key mmWave WiFi deficiencies we perform simulations using the ns-3 IEEE 802.11ad/ay module [1]. This is meant to complement previous conclusions drawn by performance evaluations of COTS WiFi networks such as [16–18, 32–34] by studying a dense deployment under idealized conditions, removing the effects of hardware imperfections and device-specific implementation details. In this way, we ensure that our insights are generalized and connected to essential mmWave behaviors. We use a simple indoor scenario with a dense deployment of 16 APs. Each AP has only one associated STA and can thus use directional reception. The APs have a grid distribution while STAs are placed randomly and we average the results over 50 different STA topologies. STAs transmit UDP traffic with a data rate of 300 Mbps, resulting in an aggregate load of 4.8 Gbps.

First, we evaluate the possible benefits of enabling directional reception. Figure 4.1 shows the aggregate throughput with omnidirectional and directional reception, revealing a dramatic increase of 150% in median throughput when directional reception is used.

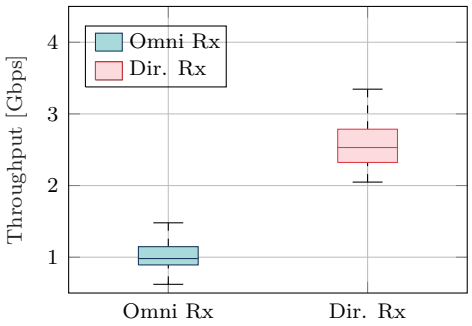


Figure 4.1: Throughput with omni and directional reception

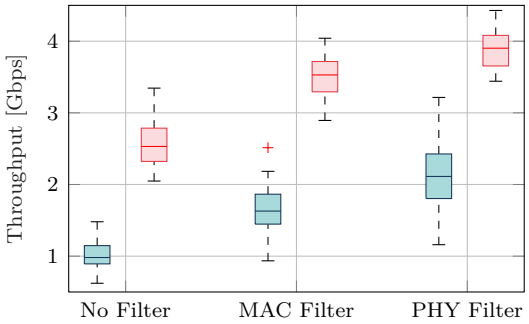


Figure 4.2: Throughput with filtering of unwanted packets

Moreover, we find that directional reception practically eliminates packet failures due to low SNR, reducing them from 10% to only 0.4%.

Next, we demonstrate the harmful effects of unwanted packet reception. For this purpose we implement a filter on the PHY layer that cuts reception of all packets not intended for the device, immediately after preamble detection, leaving it free to receive its own packets. In Figure 4.2 we show how this benefits network throughput, using both quasi-omni (blue) and directional (red) reception. We observe that in both cases performance is increased with throughput gains above 40%. Additionally, we see how directional reception and packet filtering complement each other. The directional reception enhances the spatial sharing potential and packet filtering helps realize this potential by reducing packet failures.

Our initial findings demonstrate that: a) enabling directional reception in mmWave WiFi is necessary for high network performance, b) there is a lot of potential to improve MAC efficiency and even simple solutions can have a significant impact and c) directional reception further boosts the performance of other MAC enhancements. In the following section, we explore how the deficiencies of current solutions led us to the design choices made in SIGNiPHY.

### 4.2.2 Existing solutions

The cause of both problems we study can be traced back to the usage of CSMA/CA, a random-access scheme without a fixed schedule. APs can not use a directional BP to receive packets because they are not aware *which* STAs is transmitting the packet, preventing them from steering towards them. Similarly, since they do not know *when* they will receive packets intended for themselves, they attempt to decode every packet they detect.

While one solution would be to use another channel access mechanism such as a Time-Division Multiple Access (TDMA) scheme with a pre-defined transmission schedule, in practice, proposed alternatives have been proven to be too complex, inefficient and

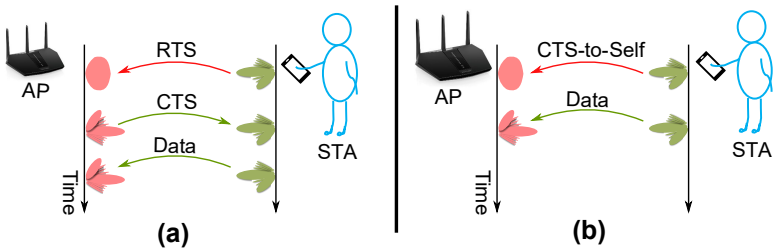


Figure 4.3: Directional reception through (a) RTS-CTS or (b) CTS-to-self

inflexible for implementation. Despite the fact that IEEE 802.11ad and IEEE 802.11ay allow for channel access with both a polling-based scheme and TDMA-like scheduled access, they are not implemented in any current COTS devices. The main issue with TDMA lies in its unsuitability for bursty or low latency traffic and the complexity of designing a schedule that implements spatial sharing. Polling, on the other hand, suffers from significant overheard which leads to sub-optimal channel usage. Finally, CSMA/CA is very simple and easy to implement and highly robust. These are desirable properties that are valuable for WiFi protocols and are worth retaining. Therefore, we looked for approaches that are designed for random-access networks and can address CSMA/CA deficiencies, such as the ones discussed below.

**RTS-CTS and CTS-to-self in mmWave WiFi:** RTS-CTS and CTS-to-self were designed to reduce WiFi collisions and as they enable user identification ahead of the transmitted data packet they can be used to enable directional reception in mmWave WiFi. Studies of COTS devices [32, 77] have found that RTS-CTS is in fact commonly used in mmWave WiFi, although not for directional reception. As they address both directional reception and MAC efficiency and they are standard compliant, we consider them as the closest baseline schemes for comparison with our work. Figure 4.3 shows RTS-CTS and CTS-to-self operation in the context of directional reception. The transmission of the RTS or CTS-to-self frame announces the intent to transmit, protecting the data by reserving the channel. By signaling the identity ahead of the data, it allows the receiver to select the receive BP before data reception.

Both mechanisms, however, have been found to limit performance due to the high overhead and inefficient airtime usage by control packets. Consequently, they can increase network latency, which might be critical for applications such as AR/VR. Analysis of RTS-CTS in mmWave WiFi has also found that it can have a negative effect on spatial sharing [32, 77]. Finally, from Section 4.2.1, we note that packet filtering based on identification from RTS or CTS-to-self packets is not fully reliable, since there is no guarantee that the next packet received is from the same source as the announcement packet and can thus lead to discarding useful packets. Therefore, it is important to analyse and carefully consider their use in mmWave WiFi, as well as propose alternatives

for enabling directional reception.

**MAC filtering:** several methods allow filtering of unwanted packets based on the MAC address. However, filtering at the MAC layer incurs a delay until the filter can be applied which reduces its usefulness. We use the simulation scenario from Section 4.2.1 to demonstrate this by comparing the performance of the optimal PHY filter, applied directly after preamble detection with a MAC filter that is applied after the reception and decoding of the MAC header. Figure 4.2 shows the comparison of the throughput results with both omnidirectional and directional reception. We can see that cutting reception early on at the PHY layer makes a difference, due to the high number of short control packets exchanged where the MAC header is a significant portion of the packet.

We find that existing solutions are inefficient as they rely on MAC identification to control PHY layer behavior such as BP selection and packet reception, which imposes either overhead or delay. To eliminate these drawbacks, the obvious solution is to instead introduce PHY layer identifiers (IDs), which allows to address both directional reception and unwanted packet reception without any additional overhead. This is the core motivation for the design of SIGNiPHY, explained in more detail in the following sections.

## 4.3 SIGNiPHY Design

As we highlight in Section 4.2, PHY layer IDs are the only approach to enabling directional reception without high control packet overhead. MAC layer signaling is simply not sufficient as it identifies the user too late in packet reception. The key limitation is that any changes to the receive BP have to finish before channel estimation starts in order to not destroy the packet decoding due to failed equalization. Figure 4.4 shows the IEEE 802.11ad packet structure. In mmWave WiFi, channel estimation is done in the Channel Estimation Field (CEF) field, immediately after packet detection. Therefore by the time that the MAC header, located in the packet payload has been decoded, it is already too late to make any changes to the BP. This implies that the PHY ID we introduce has to be embedded in the packet preamble, or more specifically in the Short Training Field (STF).

This makes the design of the PHY identification quite challenging for three reasons: 1) The packet preamble enables fundamental receive functionalities like packet detection, Carrier Frequency Offset (CFO) estimation and synchronization which must not be altered. Therefore, when modifying the preamble design we have to ensure that these functionalities are preserved. 2) The modified preamble has to have the same inherent structure and auto-correlation properties in order to ensure backwards compatibility and interoperability with legacy devices. 3) The embedded IDs must be decoded fast enough

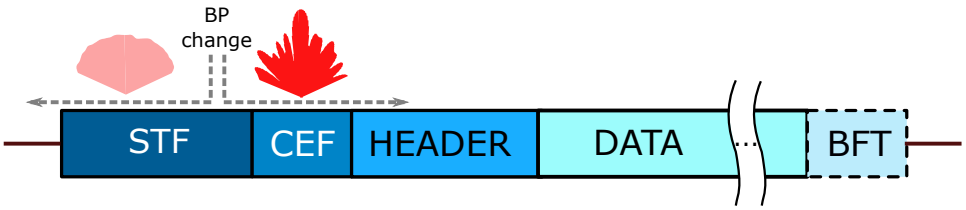


Figure 4.4: IEEE 802.11ad frame format

to allow for a BP switch before the end of the STF field.

SIGNiPHY's preamble embedding mechanism manages to accomplish all three goals and provide extremely robust ID decoding as we will demonstrate in Section 4.5. The details of how we accomplish that are given in Section 4.4, while below we elaborate on how SIGNiPHY utilizes the early PHY identification to solve key sources of mmWave WiFi inefficiency.

### 4.3.1 Enabling directional reception

It is clear that the AP benefits from using omnidirectional reception when it is in idle mode since it allows it to receive packets from different STAs. However, a directional BP would be optimal for the packet decoding itself. SIGNiPHY enables a seamless transition from omnidirectional to directional reception with zero overhead. This is done by early identification of the user identity, as soon as the preamble of the packet starts being received. If the decoded ID corresponds to an associated STA, the AP performs fast switching from omnidirectional to directional reception, steering the antenna towards the identified STA. This allows the packet payload to be received with a higher antenna gain, increasing the decoding probability. If no user ID is embedded, it remains in omnidirectional reception, thus ensuring interoperability between SIGNiPHY and legacy IEEE 802.11ad/ay devices within the same network.

Interestingly, we note that this mechanism can also be used at the STA side, although this is not necessary. Since STAs only communicate with their AP, they can simply always steer towards the AP. This reduces the carrier sensing accuracy outside the BP main lobe and could increase collisions. However, carrier sensing is enhanced in the area towards the AP, which can be beneficial. Additionally, packet detection is improved by the higher antenna gain, enabling reception of low SNR packets that would be lost with an omnidirectional BP. We found that the benefits of continuous directional reception at the STA outweigh the drawbacks and thus only use SIGNiPHY for directional reception at the AP side. However, this is an implementation choice.

Lastly, we highlight that SIGNiPHY operates independently from the beamforming training and simply utilizes the identifier to BP mapping table already provided by the beam training.

### 4.3.2 PHY Packet Filtering

An extra benefit of the SIGNiPHY user ID embedding capability is the possibility of true early PHY filtering of unwanted packets. When an AP receives a packet carrying an ID that is not registered in the network, it is able to quickly react by dropping the packet and returning to idle mode to wait for the next packet. This is different from MAC filtering (as discussed in Section 4.2), since SIGNiPHY does not have to fully decode the MAC header to decide whether the packet needs to be dropped or not.

Note that even with the continuing background interference caused by the unwanted packet, it is usually possible to successfully identify incoming packets from registered STAs thanks to the high robustness of the preamble (as will be shown in Section 4.5). Once the preamble has been decoded, the directional reception enabled by SIGNiPHY reduces the gain outside the main lobe which helps to minimize the interference from the filtered packets and increases the probability of successful packet reception. In this way, SIGNiPHY improves spatial reuse and prevents nearby networks from interfering with each other. Importantly, SIGNiPHY does not increase intra-network interference, as devices still use standard carrier sensing to determine whether they are allowed to transmit. Therefore, devices do not transmit more than in standard operation, but instead, the probability of successful transmission is increased.

### 4.3.3 Protocol implementation

We have two main design considerations for SIGNiPHY: 1) ease of implementation and 2) avoiding unintended negative behavior. We ensure 1) by avoiding global optimizations and using simple distributed mechanisms that do not need coordination. In fact, the only required information is the user IDs which can be obtained with no overhead in several ways. One option is to re-purpose the Association Identifier (AID). This only requires randomizing the AID allocation and ensuring that assigned AIDs are within the SIGNiPHY ID space. Another option is to apply a known hash function to the MAC address that maps it to our ID space. In addition, backward compatibility at the protocol level is ensured, as devices always default to standard behavior in case no ID has been embedded in the preamble by legacy STAs.

Finally, handling duplicate IDs among STA of different networks is feasible by monitoring the SNR change after the switch to a directional BP and the match between the decoded ID and the MAC address. Once an AP has detected duplicate IDs, it can assign a new one to the STA in its own network, ensuring that duplicate IDs are only transient. During this transition, the AP simply uses the BP for the associated STA, rather than discarding the unwanted packet, which is no worse than receiving it omnidirectionally.

We accomplish 2) by designing SIGNiPHY to not increase transmission probabilities and instead we focus on optimizing packet reception, in order to increase spatial reuse

without adding interference. Additionally, we ensure that even in the case of incorrect ID decoding, SIGNiPHY does not cause harm to network operation. For this purpose SIGNiPHY monitors the SNR change after the switch to a directional BP, as well as the match between the decoded ID and the MAC address. This allows SIGNiPHY to estimate the probability of correct decoding and default to standard omnidirectional reception when the SNR is too low for reliable ID decoding. We also note, that our evaluation in Section 4.5 shows that SIGNiPHY ID decoding failures only happen at very low SINR at which very few packets are detected and successful packet decoding is not possible. Therefore, due to SIGNiPHY's high robustness, we find that ID decoding errors have a negligible effect in our testbed.

### 4.3.4 Initial Setup Procedure

SIGNiPHY is designed to be an enhancement to the mmWave WiFi protocols, with only minor modifications to the operation of devices. When a new STA joins the network, the only change to the initial access procedure is that the STA is assigned a SIGNiPHY ID by the AP. As discussed in Section 4.3.3, this can be done in the association process by re-using the AID or through a hash function. IEEE 802.11ad/ay specifies that before any data transmission can take place, new STAs need to perform beamforming training to determine the optimal BPs for communication and support several protocols to perform the training. SIGNiPHY does not modify the beam training procedure in any way and is compatible with all beamforming training protocols. It is only activated afterwards, during the transmission and reception of data packets between the AP and the new STA. The STA embeds the SIGNiPHY ID it was assigned in the preamble of all data packets it sends and the AP extracts the embedded ID and uses it to enable directional reception and PHY packet filtering.

### 4.3.5 Optional support for Dynamic Sensitivity Control (DSC)

SIGNiPHY early user identification can be beneficial beyond its core functionalities. Therefore, we envision that SIGNiPHY PHY signaling can be extended to support further mmWave WiFi enhancements. As an example, we use it to implement the IEEE 802.11ax DSC mechanism, as conceptually it requires SIGNiPHY functionalities. This new feature allows 802.11ax STAs to dynamically adjust their carrier sensing level for enhanced spatial reuse. For optimal performance, it is combined with the Basic Service Set (BSS) color feature that inserts an identification field for the originating BSS in the PHY header. Together, they enable STAs to use different carrier sensing thresholds for transmissions from their own and other overlapping BSS [80]. Such a mechanism can be particularly beneficial for mmWave networks as it can allow to reduce collisions within a BSS without disturbing the spatial sharing in the network. Neither BSS color nor DSC are currently

supported in mmWave WiFi, but SIGNiPHY allows us to easily implement a dynamic threshold approach by relying on SIGNiPHY IDs instead of BSS color.

## 4.4 SIGNiPHY Preamble Embedding Mechanism

As we explain in Section 4.3, directional reception is only possible if the SIGNiPHY ID is embedded in the packet preamble, specifically, in the STF. The STF of IEEE 802.11ad/ay data packets comprises 16 repetitions of $Ga_{128}$ Golay sequences and ends with a phase-inverted -$Ga_{128}$. Golay sequences have good correlation properties which ensures that the preamble is very robust and can be detected even in challenging conditions. The first key function of the STF is to enable packet detection, typically done by using normalized autocorrelation (NAC) [81, 82] to detect the repeated structure of the STF. In this way, the received signal is correlated with delayed versions of itself to detect the autocorrelation peaks of the repeated $Ga_{128}$ sequences. The STF is also used for CFO estimation and correction by measuring the phase differences across the repeated identical sequences. Finally, boundary detection to find the end of the STF can be done by identifying the phase inversion of the final Golay -$Ga_{128}$ sequence [83]. Our key insight which led to the design of SIGNiPHY is that none of the STF functionalities need knowledge of the exact preamble sequences. Instead, they all rely on the inherent structure and correlation properties of the field. Therefore, we need to keep the structure of multiple repetitions ending with a phase-inverted sequence, maintain the same correlation properties and length of the sequences, and the same total length of the field. If we keep these constraints we can modify the STF to embed the SIGNiPHY ID without affecting any of the preamble functionalities, thus maintaining interoperability with legacy devices.

To allow for user identification based on the STF it is necessary to design a set of preambles so that each user can have a unique version of the STF that is differentiated from the others. To design this set, we rely on the fundamental property of Golay sequences which states that Golay sequences can be built by concatenating shorter sequences [84]. Therefore the $Ga_{128}$ Golay sequence used in the original preamble can be built by the concatenation of 4 smaller Golay sequences with length 32 so that $Ga_{128}$ = [$Gb_{32}$ $Ga_{32}$ -$Gb_{32}$ $Ga_{32}$]. Additionally, all possible combinations of $\pm Ga_{32}$ and $\pm Gb_{32}$ also result in valid 128-bit sequences which can be detected by a legacy receiver. This gives us an original set of 256 different sequences which can be used to generate the STF field of different users. We further have to reduce this set to ensure a practical and easy implementation. One reason is that while detecting that a packet is present is very robust, determining the exact start point of a sequence is a complex and challenging task which would require channel estimation and back-propagation of the information. Therefore, as we do not know the exact detection point within the STF, we remove cyclically shifted versions of the same sequence (e.g. [$Gb_{32}$ $Ga_{32}$ -$Gb_{32}$ $Ga_{32}$] and [$Ga_{32}$ $Gb_{32}$ $Ga_{32}$ -$Gb_{32}$])

from the set to avoid mistaken identification. Additionally, while detecting phase shifts between sequences is quite reliable, distinguishing whether each sequence is positive or negative is a difficult task. Therefore, we also remove sequences which are $180^o$ shifted versions (e.g. [$Gb_{32}$ $Ga_{32}$ -$Gb_{32}$ $Ga_{32}$] and [-$Gb_{32}$ -$Ga_{32}$ $Gb_{32}$ -$Ga_{32}$]). This leaves us with a set of 38 different sequences which can be uniquely identified and assigned to different users and represent our dictionary $\boldsymbol{\mathcal{D}}$. You can find the full details for the embedding mechanism in [3].

### 4.4.1 SIGNiPHY ID decoding

Once the SIGNiPHY transmitter has embedded its identity by using the assigned set of Golay-32 sequences to build the packet preamble, it is necessary to extract the embedded information on the receiver side. This requires the SIGNiPHY ID decoder to determine which exact Golay-32 sequences were used to build the larger 128-bit sequence. For this purpose, we use a cross-correlation-based algorithm where we look at the similarity of the received sequence with candidate sequences from our dictionary $\boldsymbol{\mathcal{D}}$. The simplest option to decode the ID is brute-force parallel correlation, where the received sequence is correlated against all possible candidates from $\boldsymbol{\mathcal{D}}$. However, this approach has a high implementation cost as it requires a bank of 38 128-bit parallel correlators. In the interest of making SIGNiPHY easy to implement in commercial devices, we instead use a more lightweight approach which relies on feature identification. Namely, each sequence in $\boldsymbol{\mathcal{D}}$ can be uniquely identified if we can determine: 1) whether each of the 4 Golay-32 sequences is $Ga_{32}$ or $Gb_{32}$ and 2) the phase difference between those sequences. Figure 4.5 shows the decoding procedure. We first correlate the received signal with both $Ga_{32}$ and $Gb_{32}$ to find the correlation peaks for each of the 4 Golay-32 sequences and determine whether the maximum comes from $Ga_{32}$ or $Gb_{32}$. This gives us 4 correlation bits, where 1 signifies $Ga_{32}$ and 0 signifies $Ga_{32}$. After, we find the relative phase differences between the sequences by addition of the correlation peaks, where in-phase sequences give a maximum. This gives us 3 additional phase difference bits where 1 means in-phase and 0 out-of-phase. Finally, we map the decoded information represented by a 7-bit code to a unique entry in our dictionary $\boldsymbol{\mathcal{D}}$ by means of a lookup table (LUT), which requires little logic elements.

### 4.4.2 Hardware Implementation

When designing SIGNiPHY we tried to minimize the implementation complexity and required modifications. On the transmitter side, the only changes are the use of Golay-32 sequences and very simple logic to choose the correct sequences when building the preamble according to the assigned ID. Support for Golay-32 sequences has already been included in IEEE 802.11ay and thus SIGNiPHY needs minimal modifications. On the receiver side, all of the preamble processing blocks are kept the same, and only an

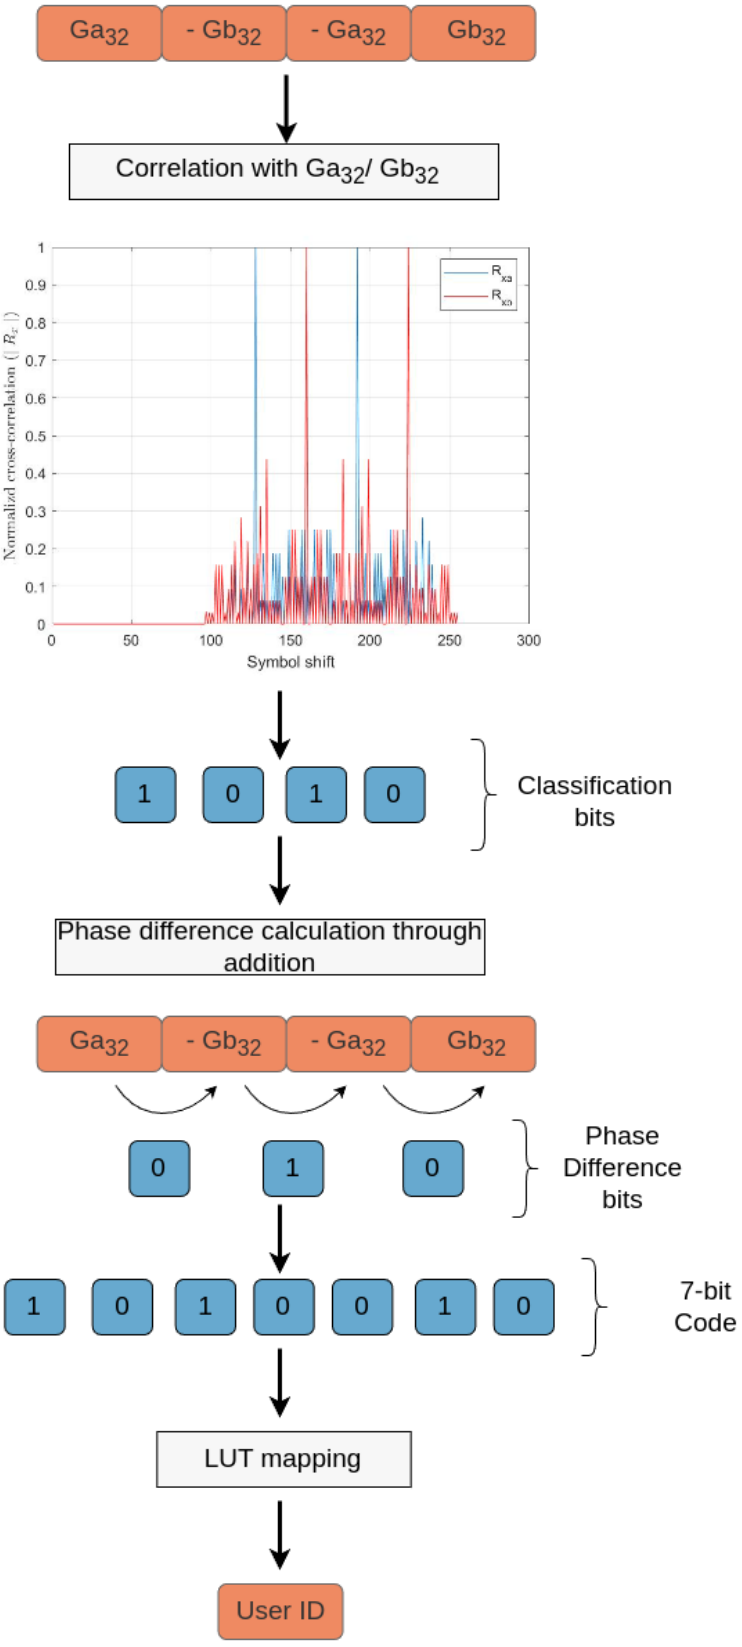


Figure 4.5: SIGNiPHY ID decoding

Table 4.1: Backwards Compatibility with IEEE 802.11ad/ay

| SNR [dB] | Packet Detection Prob. | | CFO Estimation [KHz] | |
|---|---|---|---|---|
| | 11ad/ay | SIGNiPHY | 11ad/ay | SIGNiPHY |
| -15.28 | 0.31 | 0.25 | -222.28 | -245.39 |
| -5.09 | 0.53 | 0.52 | -250.21 | -255.81 |
| -0.72 | 0.89 | 0.89 | -257.04 | -257.41 |
| 2.19 | 0.98 | 0.98 | -256.14 | -258.09 |
| 4.68 | 1 | 0.99 | -257.08 | -257.81 |
| 7.08 | 1 | 0.99 | -254.87 | -257.55 |
| 9.14 | 1 | 0.99 | -255.95 | -259 |
| 10.91 | 1 | 0.99 | -255.47 | -258 |

additional ID decoding block is needed. The ID decoding block uses a fast correlator architecture similar to [81] which can correlate with both $Ga_{32}$ and $Gb_{32}$ using the same logic elements. In addition, the ID decoding block only requires simple logic operations like addition and comparison to perform the sequence classification and phase difference calculation. All designed elements can be directly translated to commercial devices and the ID decoding block only increases the hardware resources of the preamble processing by approximately 10%, which is minor when compared to other receiver functionalities. Our implementation on a Xilinx XCZU28DR FPGA only needed 0.6% of the logic elements and was able to compute the ID in one-tenth of the duration of the STF. Finally, if the decoded ID is a registered user, the ID decoding block directs the antenna controller to initiate a switch to a directional BP by transmitting GPIO pulses to the mmWave front-end. The fast BP switching required to finish the change before the start of the channel estimation is already required in IEEE 802.11ay for beamforming training using Training (TRN) fields as discussed in Chapter 3. You can find the full details of the hardware implementation of SIGNiPHY in [3].

## 4.5 Experimental Evaluation

We implement SIGNiPHY in an FPGA testbed based on an open-source platform [36] which operates at 60 GHz and is compatible with IEEE 802.11ay, allowing us to evaluate SIGNiPHY in real scenarios.

First, we demonstrate that SIGNiPHY is backwards-compatible with legacy IEEE 802.11 devices and does not affect the preamble processing. For this purpose, the transmitter and receiver are at a distance of 4 m and we analyze different SNRs by adjusting the transmit power. We compare the performance of SIGNiPHY packets to standard IEEE 802.11ad/ay packets and show the results in Table 4.1. We found that the detection probability was practically equal when the SNR was above -5 dB, which represents the range of SNRs where data communication is possible. Below that, at very

low SNRs where packet decoding is not possible SIGNiPHY showed a minor reduction of up to 6% at an SNR of -15 dB. Likewise, the CFO estimation was practically equivalent for SNRs above -5 dB, with a difference below 2%, which goes up to 10% at -15 dB. Small variations in CFO like this do not affect the receiver decoding as they can be corrected for later on in the packet processing [85]. Lastly, when evaluating the boundary detection, we found that it was 100% accurate for both SIGNiPHY and IEEE 802.11ay packets for all SNRs tested.

Next, we evaluate the SIGNiPHY functionality under challenging conditions. First, we calculate the probability of correctly identifying the user ID under an additional interferer which constantly transmits data packets which collide with the SIGNiPHY packet preamble. This can affect the ID decoding since the guard interval of data packets is also composed of Golay sequences, making it harder to determine the SIGNiPHY sequences sent. We vary the SINR by changing the transmit power while keeping the interferer power constant. The results are presented in Table 4.2. We find that SIGNiPHY ID decoding is very robust in the SINR range where we can decode packets, with 99.5 % accuracy above -5 dB, and 100 % over 5 dB. Additionally, we measure the decoding and BP switching latency to ensure that the switch to a directional BP can happen before the end of the CEF. Packet decoding takes 0.21 $\mu s$ and the ID decoding adds just 0.15 $\mu s$ latency to the processing. The additional latency to switch to a directional BP and the time required to deliver the samples from the ADC to the FPGA logic are dependent on the hardware architecture and in our case take 0.25 $\mu s$. The rest of the STF (0.5 $\mu s$) is more than sufficient to perform CFO estimation [81].

Finally, we demonstrate how SIGNiPHY enhances packet decoding. First, we set an experiment where the receiver tries to decode a packet while under interference from an interferer with equal transmit power. Moreover, the interfering packet is the first to arrive at the receiver. While the standard decoding process would lock onto the first packet, SIGNiPHY detects that the sender of that packet is not a registered user and stops the packet reception, allowing it to receive the wanted packet later. The switch to a directional BP then increases the received signal power and allows for the packet to be decoded with zero bit errors even while still under heavy interference. We further validate that attempting to decode that packet with an omnidirectional BP is not feasible and fails. Lastly, we test SIGNiPHY in a very challenging scenario under multiple interferers. For this experiment, we have between 1-3 interferers in an indoor scenario. Both the interferers and the transmitter send packets with the same length (12 $\mu s$) with an inter-frame spacing with a random uniform distribution between 10 and 50 $\mu s$, making it likely that packets collide with one another. The interferer transmit power is the same, while the transmitter varies the power to achieve different SINR values. We compare the performance of SIGNiPHY where packets are filtered based on the embedded ID to the standard approach of decoding the first packet detected. Figure 4.6 shows the achieved throughput with the

| SINR [dB] | Decoding Acc. |
|---|---|
| -20 | 0.16 |
| -16 | 0.38 |
| -12.74 | 0.67 |
| -9 | 0.86 |
| -5.86 | 0.98 |
| 3.8 | 0.99 |
| 9.13 | 1 |
| 15 | 1 |
| 18 | 1 |

Table 4.2: SIGNiPHY ID Decoding

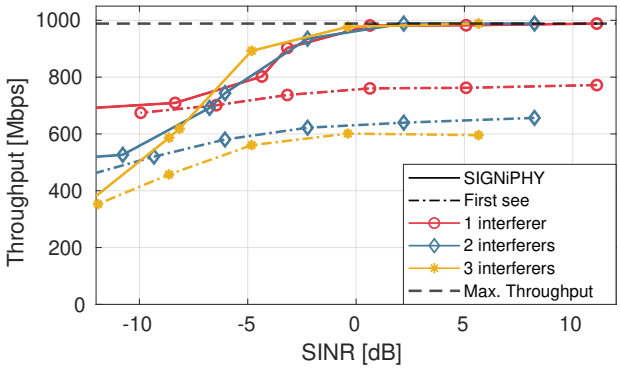


Figure 4.6: SIGNiPHY performance in a multi-interferer scenario

two approaches. We note that for a fair comparison, the first-see-first-decode device is also switching to a directional BP after the packet detection. The maximum throughput dashed line represents the throughput if all possible packets from the transmitter are decoded. It is evident that even the presence of a single interferer prevents the standard decoding approach from achieving the maximum throughput regardless of the SINR. This fundamental limitation comes because the receiver locks onto interfering packets and while it attempts to decode them misses packets from the transmitter. SIGNiPHY, on the other hand, successfully copes with multiple interferers and manages to decode all transmitted packets at an SINR of 0 dB and 80% of them at -5 dB. The full details for the experimental evaluation can be found in [3].

## 4.6 Simulation Evaluation for Large Networks

To study SIGNiPHY in dense scenarios with channel contention and spatial sharing, we use the IEEE 802.11ad/ay ns-3 module. We modify [1] to support SIGNiPHY with all features described in Section 4.3. As ns-3 does not allow re-calculation of received power within a packet, to support directional reception, we split SIGNiPHY packets in two consecutive ns-3 events. The first event carries the ID, while the second contains the packet headers and payload. If the ID is known, the best receive BP (from prior beam training) is used to receive the second event. If the ID is not known, filtering is triggered in the `StartReceiveHeader` function and reception of the second SIGNiPHY event is canceled. SIGNiPHY can be activated in both AP and STA nodes and its functionalities can be turned on or off independently for a detailed performance analysis. Finally, we also modify the RTS-CTS and CTS-to-self mechanisms to switch to directional reception after an RTS or CTS-to-self packet, to compare SIGNiPHY with these two alternatives, as well as baseline 802.11ad/ay performance with quasi-omnidirectional reception.

### 4.6.1 Simulation scenarios

We study indoor scenarios, using a rectangular room (29.6 $m \times 54\ m \times 3\ m$) as the default scenario. Three AP deployments with 4, 8 and 16 APs are considered. APs are mounted on the ceiling at a height of 3 m. Between 8 and 64 STAs with a random location and orientation are placed at a height of 1.2 m. AP association is according to distance, using load balancing to have an equal number of STAs per AP. We use the IEEE 802.11ay protocol and study uplink data transmissions, to analyze channel contention. STAs transmit with data rates of 0.3, 0.5, 1, 2 and 4 Gbps, using two Aggregate MAC Protocol Data Unit (A-MPDU) aggregation sizes: the maximum aggregation supported by IEEE 802.11ay (4 MB) and a low aggregation of 16 KB. All devices use 8x32 element uniform rectangular antenna arrays, which generate narrow beams with high gain. We use IEEE 802.11ay-compliant beamforming training, taking into account both the training overhead and BP selection errors caused by interference. Each simulation has a duration of 50 s and the results are averaged over 50 simulation scenarios with different STA locations.

### 4.6.2 Evaluation Results

We first return to the scenario from Section 4.2.1 which has high spatial sharing potential, as 16 APs are deployed in a large room, and each AP only serves one STA. We study the two A-MPDU aggregation sizes since the aggregation affects the level of contention and the achievable throughput. The data rate is set to 0.3 Gbps per STA for the low aggregation and 2 Gbps for the high, to have similar channel saturation. Figure 4.7 shows the Cumulative Distribution Function (CDF) of STA throughput measured over 2 s intervals, demonstrating the excellent SIGNiPHY performance. In the low aggregation mode, only SIGNiPHY can achieve the target data rate, with STAs maintaining the maximum throughput approximately 40% of the time. CTS-to-self only manages this 6% of the time, while with RTS-CTS and the baseline, STAs barely get half of the target data rate. The high aggregation mode shows very similar results, with SIGNiPHY showing even more stable performance and achieving the maximum throughput 60% of the time. Moreover, throughput is below 500 Mbps only 5% of the time, compared to 25% for quasi-omni and 10% for the other schemes, showing that SIGNiPHY gains are not only focused on high performing STAs but that performance is boosted equally throughout the network. It is also evident from Figure 4.7 that omnidirectional reception severely limits mmWave WiFi, leading to low achievable data rates. Additionally, the poor performance of RTS-CTS should be highlighted as even with directional reception it is not able to outperform quasi-omni IEEE 802.11ay. We found that this is not just due to the extra overhead, but also because it silenced large portions of the network, reducing the spatial sharing. STAs try to transmit a lot less often, which fundamentally limits the achievable

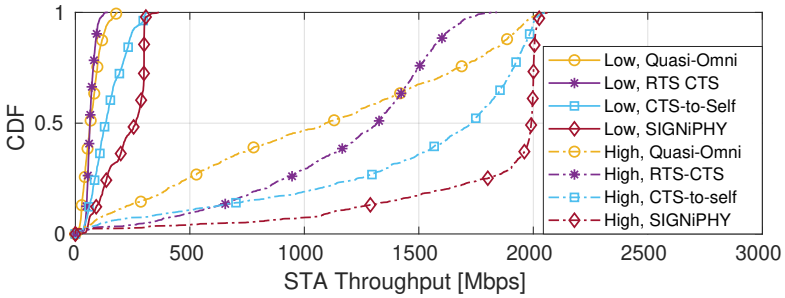


Figure 4.7: Throughput CDF for 16 APs, 16 STAs deployment

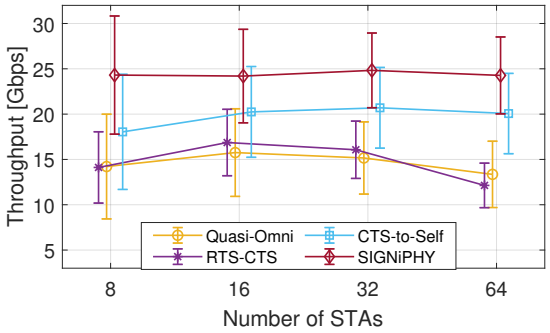


Figure 4.8: Median throughput for 8 APs deployments

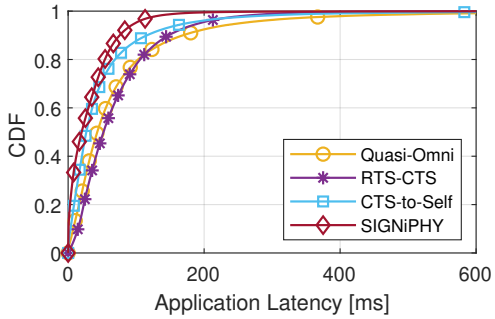


Figure 4.9: Application latency for 8 APs, 8 STAs deployment

throughput. While CTS-to-self is a better alternative, SIGNiPHY manages to have an extra gain of 50% in median throughput, since it incurs no overhead and quickly stops overhearing unwanted packets.

Since low A-MPDU aggregation severely limits the achievable throughput and over-saturated the channel, in the following scenarios we focus on the high aggregation performance, which allows us to better explore the effect of the network density and traffic load.

To test SIGNiPHY in more challenging conditions, we design a smaller 8 AP deployment, where both spatial re-use and coverage are reduced. Additionally, we vary the number of clients per AP between 1 and 8 to create more realistic deployments, with more STAs sharing the medium. We keep the offered network load at 32 Gbps, reducing the per STA throughput from 4 Gbps to 0.5 Gbps as the number of STAs increases from 8 to 64. This allows us to study how well SIGNiPHY adjusts to different network configurations. Figure 4.8 shows the median network throughput with 95% confidence intervals for the different network densities. We can see that SIGNiPHY handles density very well, with the median aggregate throughput remaining practically constant, regardless of the number of STAs. The larger variations in the 8 STAs deployment are because of the dependency on STA locations. STAs in difficult locations which have poor performance highly influence the total throughput since they carry a large load, as compared to denser deployments where the load is distributed between STAs. Although CTS-to-self also shows good

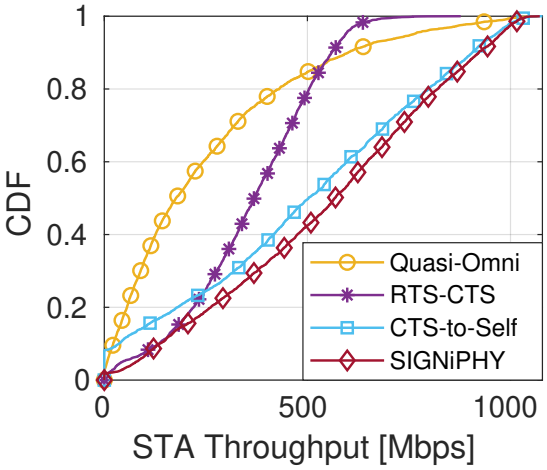


(a) Throughput CDF

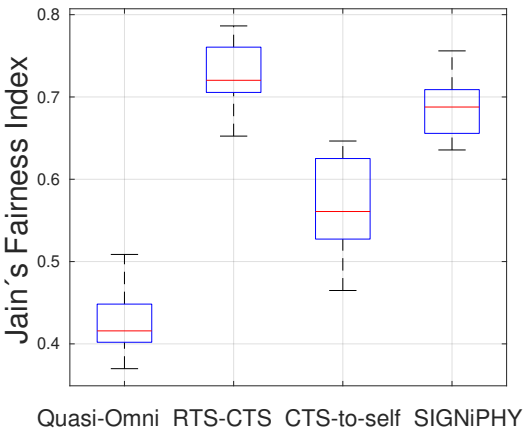


(b) Jain's Fairness Index

Figure 4.10: 8 APs, 64 STAs deployment, 1 Gbps data rate

performance, SIGNiPHY outperforms it by 20%. RTS-CTS and quasi-omni are both 50% below SIGNiPHY and furthermore they degrade as the density increases. Finally, we verify that even in dense deployments, SIGNiPHY maintains a high level of fairness. In the 64 STA deployment, STAs receive less than half of the offered throughput only 10% of the time, compared to 20% for CTS-to-self, 55% for quasi-omni and 76% for RTS-CTS.

Moreover, we demonstrate that SIGNiPHY improves another key network metric, latency. Figure 4.9 shows the CDF of the application layer latency for the 8 STAs deployment including buffering delays in the MAC queue. We see that the latency gains are especially high towards the tail of the distribution. SIGNiPHY keeps packet latency below 76 $ms$ 90% of time which is 43 $ms$ lower than CTS-to-self, the closest baseline scheme, and more than twice as low as RTS-CTS and quasi-omni. Finally, we observe that RTS-CTS overhead has the highest median latency, 32 $ms$ higher than SIGNiPHY.

While the previous scenarios overloaded the other three schemes, SIGNiPHY never operated in fully saturated conditions, with STA always achieving a median throughput of 85% or higher of the offered load. Therefore, we set a higher data rate of 1 Gbps in the 64 STA deployment to observe SIGNiPHY in fully overloaded conditions. In Figure 4.10a we see that SIGNiPHY maintains the performance gains over the three baseline schemes. Although CTS-to-self offers similar throughput for well-performing STAs, SIGNiPHY outperforms it by boosting low-performing STAs, leading to a median network throughput gain of 20%. Moreover, CTS-to-self has increased unfairness as the fraction of time STAs have 0 throughput increased to 8%, compared to 3% for quasi-omni IEEE 802.11ay. Conversely, SIGNiPHY is able to reduce it to 1%, the same as RTS-CTS. This is better illustrated in Figure 4.10b which shows Jain's Fairness Index. Quasi-omnidirectional reception leads to very low fairness as STAs far from the AP have both few transmission opportunities and high packet loss due to the low link gain. RTS-

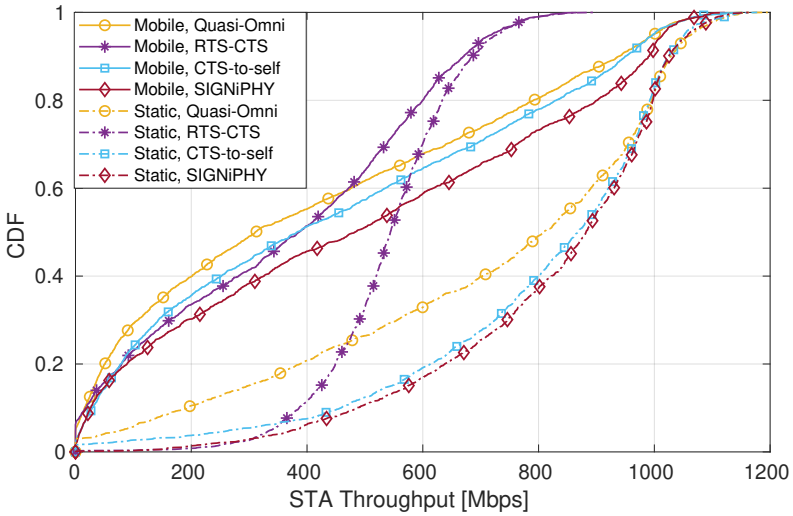


Figure 4.11: Throughput CDF in a mobility scenario

CTS increases the fairness the most, but at a significant cost in terms of achievable throughput. CTS-to-self has high throughput but also high unfairness, in particular for low-performing STAs. In contrast, SIGNiPHY further increases throughput and has much higher fairness, allowing low-performing STAs to transmit more often even under fully overloaded conditions.

Lastly, we evaluate SIGNiPHY in a scenario with user mobility. SIGNiPHY is largely independent of mobility as it is not responsible for BP selection and simply relies on the underlying beam training mechanism to adapt to movement. To demonstrate this we design a scenario with 4 APs and 32 STAs, with a per-STA data rate of 1 Gbps. In the interest of simulation time and complexity, each AP has 2 mobile and 6 static clients. Figure 4.11 shows the CDF of the per-STA throughput only for the mobile STAs, comparing the performance with a static scenario in which the STAs remain in their initial position for the duration of the simulation. We can see that while mobility has a negative effect on throughput, SIGNiPHY is not more affected than the other schemes and still outperforms them. In fact, the gain in performance compared to CTS-to-self is increased, as SIGNiPHY is more robust to interference. It is interesting to observe that quasi-omnidirectional reception, while still inefficient, is not as harmful in the case of mobility as the wide BP can sometimes prevent link misalignment. Overall, SIGNiPHY increases the aggregate network throughout by 13%, 32% and 50%, as compared to CTS-to-self, RTS-CTS and quasi-omni, respectively.

## 4.7 Discussion

SIGNiPHY was designed to enhance packet reception and improve efficiency, bringing mmWave closer to practical, real-world deployment. We have demonstrated

experimentally that it enables packet decoding under low SINR and multiple interferers, significantly enhancing the resilience and robustness of the network. In turn, the dense network simulation evaluation shows that this improves the spatial re-use in the network and improves the performance in terms of both throughput and latency by specifically helping the lowest performing users the most. SIGNiPHY can thus, be very useful for dense scenarios like Internet Of Things (IoT) or HD camera deployments. Additionally, applications which require consistently high throughput and low latency like ARVR can benefit from the enhanced packet decoding.

SIGNiPHY usage in very dense networks might run into scalability challenges as the ID space is limited to 38. Since mmWave has a short range, this should not pose a problem for most deployments, however, in extremely dense scenarios it is possible to extend the range at the cost of slightly higher hardware resources. Namely, by adding $\pm Gc_{32}$ and $\pm Gd_{32}$ sequences (included in IEEE 802.11ay) the number of valid SIGNiPHY IDs grows to 532. An additional 32-bit correlator and minor storage are required for this modification. An added benefit of an increased ID space is that it will reduce the appearance of duplicate IDs between different networks. While this is a minor, transient issue which does not affect wanted packet reception (as explained in Section 4.3.3), avoiding it can improve the filtering functionality of SIGNiPHY

Lastly, while in Section 4.6 we demonstrate that SIGNiPHY is not directly harmed by mobility, performance can still suffer in a very high mobility environment. The problem here is not specific to SIGNiPHY and instead comes from the use of directional receive BPs. Indeed, in contrast to the scenarios we have considered so far, directional reception can lead to reduced performance due to frequent link misalignment caused by user mobility. We plan to address these challenges by extending SIGNiPHY to implement a link degradation detection functionality. The idea is to measure the received power before and after the switch to a directional BP. A drop in received power instead of an increase would signal that the beams are getting misaligned due to mobility, triggering a switch back to omnidirectional reception and initiating link recovery. As the check is done within the same packet it is less likely to be affected by transient interference than mechanisms which compare the power over multiple packets. Furthermore, it can potentially react very fast to link degradation as this check can be done on a per-packet basis and while using directional transmission and reception which is more sensitive to mobility. Other planned future work includes a MIMO extension and enabling per-packet interference nulling to overcome particularly strong interferers.

## 4.8 Related Work

**PHY layer signaling:** Different prior work has introduced PHY signaling mechanisms [86–92]. [86] uses correlatable sequences instead of control packets as they are

more robust and have lower overhead, however, the approach requires a bank of correlators which makes it complex to implement. [92] and [91] are not backwards-compatible as they require an extra PHY preamble which also increases the overhead. Phase-shift keying is used in [87] to embed information in the STF, however, this approach is affected by phase noise and CFO. A similar approach in [89] uses time and phase shifts for encoding, however, the accuracy of the decoding is dependent on the channel estimation. In fact, all works expect [86] require channel estimation and equalization for the decoding which makes them unsuitable for our purpose as the switch to a directional BP needs to be done before the channel estimation. Additionally, we ensure low complexity, practical implementation and backwards compatibility and we additionally are the first to focus on mmWave WiFi.

**Preventing unwanted packet overhearing:** Most prior work in this area has looked at how stopping unwanted receptions can improve energy efficiency at sub-6 GHz frequencies by allowing the device to go to sleep while the unwanted packet is transmitted. Both MAC filtering [78, 79] and RTS/CTS [93, 94] have been considered for this purpose. Additionally, PHY signalling approaches include identification in low-power mode [91], as well as security purposes [92]. Our work differs significantly from these approaches as the problem of overhearing has unique connotations in the mmWave band.

**mmWave WiFi optimization:** Recent works [33, 95–98] have considered optimizing mmWave performance in multi-AP deployments to cope with interference and increase spatial reuse. However, they approach the problem from the PHY layer perspective and attempt to optimize beam and AP selection [33, 34, 97], coordinate transmissions [95, 98] and reduce blockage [77] through global network optimization with a central controller. They are complementary to our work which focuses on improving the MAC functioning and can be jointly used with SIGNiPHY. For example, [96] mitigates interference by modifying the BP to create a null in the direction of the interferer. SIGNiPHY can provide fast interferer identification to simplify the steering of the BP null.

**CSMA/CA for directional mmWave WiFi:** Prior work on directional CSMA/CA mostly fails to take into account mmWave operation and assumes efficient communication is possible with omnidirectional antennas [99–101]. Works that are specifically targeted towards mmWave WiFi either assume the existence of a central controller [102, 103] to coordinate transmissions, require cooperation and relaying [104, 105] or introduce significant overhead [106]. Additionally, all these works propose new, modified MAC protocols with additional complexity and overhead which makes their deployment and evaluation in practical WiFi implementations difficult.

## 4.9 Conclusions

In this chapter, we looked at how to improve the efficiency of dense mmWave networks by focusing on two aspects which lead to packet failure and reduced spatial sharing - the use of omnidirectional receive BP and the overhearing of unwanted packets. We address both problems with SIGNiPHY which introduces PHY layer signalling for early user identification and allows both fast packet filtering and directional reception of the packet payload. The high antenna gain allows for improved packet decoding, as well as resilience to interference, while the filtering leads to improved spatial sharing. Our approach is backwards-compatible, with no overhead and low implementation complexity for easy integration in real systems. We evaluate SIGNiPHY experimentally for real-time operation, as well as in dense deployments using ns-3. The experimental evaluation validates the backwards-compatible design of SIGNiPHY, the robustness of the ID decoding and proves that SIGNiPHY can enhance packet decoding under challenging conditions with multiple interferers. This is complemented by the simulation evaluation which shows that SIGNiPHY improves channel access and spatial sharing in dense networks for higher throughput, lower latency and higher fairness.

# 5 IEEE 802.11ay mmWave MIMO

## 5.1 Introduction

So far in this thesis, our work has focused on Single-Input and Single-Output (SISO) Millimeter-Wave (mmWave) networks, presenting an analysis of the novel IEEE 802.11ay protocol and exploring how improved Beamforming Training (BFT) and physical (PHY) signalling can be used to enhance the operation. However, recent advancements in mmWave electronics have paved the way towards high-performance and low-cost RF Integrated Circuits with lower power consumption. This enables the development of enhanced devices with multiple Radio Frequency (RF) chains which can support advanced PHY layer techniques like Multiple-Input and Multiple-Output (MIMO). MIMO is a key technology at sub-6 GHz frequencies as it enables concurrent transmission and reception at the same time and over the same frequency and can thus increase the achievable throughput several times.

Enabling MIMO at mmWave frequencies comes with unique challenges. First, mmWave devices require the use of phased antenna arrays with a high number of antenna elements to generate narrow beams. Due to the high power consumption, cost, and implementation complexity, it is not feasible to have a separate RF chain for each antenna element [42]. Therefore a fully digital precoding architecture such as the ones used at sub-6 GHz frequencies is not appropriate and alternative analog or hybrid designs have been considered [107, 108]. The wide bandwidth introduces non-linearities and phase noise and additionally requires novel algorithms that consider frequency-selective channels [42]. Furthermore, there is a high number of parameters which need to be configured, requiring the use of new algorithms that reduce the complexity by exploiting the channel sparsity at mmWave [107, 109]. Finally, directionality also significantly affects the mmWave MIMO design. First, since analog beamforming is necessary for communication, the analog beampattern used affects the observed channel and it is challenging to determine the channel without any beamforming [42]. Additionally, all the challenges of user mobility which exist in SISO communication are exacerbated for MIMO systems, as the changes

in the channel and link breakages lead to long and complex training to reestablish the MIMO link. In fact, the complex nature of MIMO BFT protocols is a key challenge to practical implementations. They involve simultaneous transmit and receive training of multiple RF chains which makes the number of possible antenna configurations extremely high. This makes the exhaustive search BFT approaches used in SISO communication impossible to implement in MIMO architectures. Designing MIMO BFT protocols which can successfully establish independent spatial streams while keeping the complexity and length of the training low to ensure scalability is a topic of ongoing research.

The IEEE 802.11ay protocol [27] is the first attempt to standardize MIMO operation in the mmWave band. The standard includes support for both Single User MIMO (SU-MIMO) and Multi User MIMO (MU-MIMO) with up to 8 spatial streams, defining BFT protocols to establish a multi-stream link and a MIMO channel access procedure. IEEE 802.11ay includes mandatory analog beamforming through novel MIMO BFT protocols, with optional support for a hybrid architecture through an additional training phase. The complete training process has different phases and can be very lengthy, requiring practical evaluations to determine how robust it is to the unreliable nature of mmWave WiFi, how the overhead affects the performance and how well it can establish multiple independent spatial streams. The latter is particularly relevant for the fully analog design, due to the sparsity of the mmWave channel. At the time of writing, no commercial devices exist which implement IEEE 802.11ay MIMO, making real-world protocol evaluations unfeasible. Therefore, we believe that the ns-3 IEEE 802.11ay module can be a valuable tool for research in this important technology.

In this chapter, we present the first standard-compliant IEEE 802.11ay mmWave MIMO implementation. We expand the Quasi-Deterministic (Q-D) channel realization software to support MIMO channels to realistically capture the effects of inter-stream interference. We then add a new MIMO engine to calculate the received power when signals are transmitted with multiple RF channels and update the PHY layer to support the decoding of MIMO packets. Next, we implemented the full standard-compliant analog BFT protocols for both Single User (SU) and MU-MIMO, trying to provide insight into the performance and to allow for easy modification of the training parameters. As hybrid beamforming is optional in the standard we wanted to explore the viability of a fully analog architecture for mmWave MIMO. Moreover, as digital precoding is currently not supported by commercial devices, we leave hybrid beamforming for future work. Finally, we implement the channel access procedure to allow for the simulation of SU-MIMO data communication. We present a validation of the MIMO BFT algorithms which reveals how the selection of good training candidates is crucial to the performance and that trade-offs exist between the quality of the selected beams and the duration of the training. Additionally, we look at the effect of user mobility, as well as how many streams can be supported by the mmWave channel. Finally, we demonstrate that analog precoding can be

sufficient to establish independent beams even with a small separation between the Phased Antenna Array (PAA) and to enable communication with high data rates. However, the achievable Signal-to-Interference-plus-Noise Ratio (SINR) varies significantly based on the user position and even minor movements can cause significant drops in throughput. Additionally, as the number of PAAs grows it can be challenging to find independent beams which have sufficient signal strength using solely analog beamforming.

The rest of the chapter is organized as follows: In Section 5.2 we provide background on MIMO operation as defined in IEEE 802.11ay and in Section 5.3 we present the details of our implementation. Section 5.4 shows the results from our performance evaluation and Section 5.5 concludes the chapter.

## 5.2 MIMO operation in IEEE 802.11ay

In IEEE 802.11ad, even though a Directional Multi-Gigabit (DMG) Station (STA) can have multiple PAAs connected to its RF chain, only a single PAA can be used at a time which results in a single stream transmission, providing diversity gain but not multiplexing gain. Instead, IEEE 802.11ay supports MIMO for a multi-fold increase in throughput. IEEE 802.11ay supports concurrent transmission and reception of up to eight spatial streams at the same time and over the same frequency. IEEE 802.11ay implements two MIMO variants: SU-MIMO allows transmitting and receiving multiple spatial streams (up to eight) between two devices, whereas, with MU-MIMO, an Access Point (AP) can transmit different spatial streams to multiple users (up to 8) at the same time. SU-MIMO communication can happen in both the uplink and downlink traffic direction, while MU-MIMO is only enabled in the downlink as it requires synchronization of the transmission by the AP.

The standard mandates the support of analog RF precoding for MIMO communication. In this mode, PAAs can synthesize a narrow beam pattern to create a spatial channel for each stream. However, generating pencil beam patterns with low inter-stream interference is not always feasible depending on the quality of the phase shifters and the geometry of the PAA. To this end, IEEE 802.11ay also supports a hybrid analog and digital architecture to compensate for the deficiencies of analog beamforming through digital precoding and achieve higher MIMO gains.

To successfully establish independent streams, it is crucial to minimize the inter-stream interference to achieve sufficient per-stream SINR for data decoding. To this end, IEEE 802.11ay introduces MIMO BFT. MIMO BFT is a very challenging task since an exhaustive evaluation of all the possible PAA stream configuration combinations is not viable in real-world MIMO implementations. For example, a small codebook with 27 predefined sectors in a 2x2 MIMO setup would already require testing over half a million combinations.

To overcome this problem IEEE 802.11ay decided to decouple MIMO BFT in two phases, the SISO phase and the MIMO phase. The SISO phase aims to find the optimal SISO BFT for every SISO transmit/receive PAA pair of the MIMO communication. Even though these results do not provide an estimation of the inter-stream interference, they can be used to identify a promising subset of candidates to evaluate in the MIMO phase. In the subsequent MIMO phase, the different transmit and receive MIMO candidate combinations are tested and the actual MIMO performance with inter-stream interference is measured. As mentioned above, an optional third phase can enable hybrid beamforming by introducing digital precoding with spatial mapping matrices.

Finally, after the successful establishment of a MIMO link, IEEE 802.11ay introduces channel access procedures for MIMO data communication. The main difference to SISO links is that MIMO requires directional transmission and reception to cope with the inter-stream interference and enable sufficient spatial separation of the links. As discussed in Chapter 4, due to the lack of a transmission schedule in WiFi networks directional reception needs advanced notice of the transmission so that the receiver can set up the correct receive antenna configuration. IEEE 802.11ay uses the Ready-to-Send (RTS)/Clear-to-Send (CTS) or CTS-to-self mechanisms for this purpose.

## 5.3 Implementation in the IEEE 802.11ad/ay ns-3 model

In this section, we present the implementation details of the MIMO extension to the IEEE 802.11ad/ay model. Figure 2.3 shows the required modifications across the different layers of the protocol stack.

### 5.3.1 MIMO Q-D Channel Generation

In order to support MIMO multi-stream transmission we needed to update the Q-D channel model [39] to support devices with multiple transceiver chains. Since the signal propagates through different paths, each transceiver chain may experience different copies of the signal. Moreover, the variation of the phase and the amplitude of each path among transceiver chains is fundamental to understanding the performance of a mmWave MIMO system. Hence, for the design and evaluation of IEEE 802.11ay MIMO, we provide an accurate channel description for each Tx-Rx chain pair. The updated NIST Q-D Channel Realization Software [61] enables the generation of a 3-D multi-point to multi-point Double Directional channel Impulse Response (CIR) (DDIR) providing the details of the magnitude, phase, and time of arrival, Direction of Departure (DOD), and Direction of Arrival (DOA) of individual propagation paths between multiple points in space. These points in the 3D space approximate the focal point of the antennas, i.e. the points from which the wireless signal is generated or collected. For IEEE 802.11ay devices with multiple PAAs, the points in space can be considered the center of each PAA.

To obtain the multi-point to multi-point DDIR a full 3D ray-tracing model captures the geometrical properties of the channel for each point-to-point pair. This allows a realistic channel representation that includes spatial correlation. Since some physical phenomena, e.g. rough surface scattering, cannot be easily modeled in a deterministic manner, several mmWave channel measurement campaigns have been conducted to integrate the deterministic channel description with realistic stochastic models [110]. The DDIR is converted to the system level channel applying the antenna array model [39] for each transceiver and re-sampling at the system rate.

### 5.3.2 MIMO Operation

We extended the `QdPropagationEngine` class to include a MIMO engine that handles the calculation of the received signal power whenever a transmission is initiated with more than one active PAA. Our approach avoids the scheduling of multiple events for the different streams transmitted to guarantee the same simulation scalability as SISO. On the transmitter side, a single transmission event is scheduled and the transmit power is allocated equally between the transmit PAAs. On the receiver side, the MIMO engine uses the MIMO Q-D channel realizations provided by the National Institute of Standards and Technology (NIST) Q-D Channel Realization Software to calculate the received signal power for each pair of active transmit and receive PAAs. The `DmgWifiPhy` class then receives a list of RX signal powers and handles the event reception according to the type of MIMO transmission.

In the case of SU-MIMO data communication, a packet decoding operation is scheduled as explained in Section 5.3.4. However, for Beam Refinement Protocol (BRP) packets transmitted during the MIMO BFT procedures, a different approach is necessary. The standard specifies that these packets are transmitted using spatial expansion, i.e., a single space-time stream is mapped to all active transmit chains with a relative cyclic shift between the different chains. This allows the receiver to separate signals coming from the different transmit PAAs and removes unintended beamforming effects. For simplicity, in our implementation, the effect of spatial expansion is modeled by only decoding the stream with the highest received power as we assume that the cyclic shift diversity is sufficient to remove the interference from the other received streams. The decoding of the packet then follows the standard SISO procedure. The Training (TRN) field of the BRP packets is also transmitted in MIMO mode and is composed of orthogonal waveforms. This orthogonal design allows the training of multiple transmit and receive antennas simultaneously by extracting the TRN subfield of each stream without any interference. Therefore, for MIMO TRN subfields, we can calculate the Signal-to-Noise Ratio (SNR) of each received stream. These values are calculated without taking into account any inter-stream interference and are equivalent to multiple SISO transmissions. Additionally, we add the possibility to calculate the SINR values of each TRN subfield. These values are

calculated by adding the received power from the other active TX antennas as inter-stream interference. We use the SNR values in the SISO phase of SU-MIMO BFT in order to get accurate measurements for the SISO performance, and we use the SINR later in the MIMO phase of SU-MIMO and MU-MIMO BFT to evaluate the effects of inter-stream interference.

Finally, we updated the `Codebook` class to support simultaneous activation of multiple PAAs. This includes a new structure which separates the `RfChain` representation in a separate class to enable different antenna configurations in the different RF chains and to control the activation and deactivation of the individual chains. Additionally, in the `Codebook` class we added new functions to control the simultaneous sweeping with multiple antennas during MIMO BFT.

### 5.3.3 MIMO Beamforming Training

As described in Section 5.2, IEEE 802.11ay introduces MIMO BFT to enable the simultaneous training of multiple transmit and receive PAA. The training has two phases - the SISO phase which determines the optimal SISO antenna configuration for each PAA trained and the MIMO phase tests different MIMO antenna configurations with to determine the optimal beam configuration, taking into account the inter-stream interference. As the number of possible MIMO antenna configurations is too large to perform an exhaustive search, the SISO phase is needed to provide information about possible candidate configurations with good performance. Based on the input of the SISO phase, a selection algorithm then needs to reduce the number of MIMO configurations which will be tested.

The selection of candidates to test in the MIMO phase is implementation-specific and not defined by IEEE 802.11ay. Thus, for the selection of transmit beam pattern combinations, we developed a custom approach based on [111], which suggests assigning a joint-beam score to different beam pattern combinations and selecting the top $K$ combinations as MIMO phase candidates. In our implementation, the joint-beam score is the sum of the individual transmit beam pattern SNRs obtained in the SISO phase. We decided to implement the selection algorithm in the simulation script to allow for easy extension to other algorithms without the need to modify the core ns-3 classes. The list of transmit candidates generated by the selection algorithm is trained in the MIMO phase to evaluate their performance. Each candidate is comprised of a TX configuration for each PAA involved in the MIMO training.

For the receive beam pattern training, we use a different approach because of our observation that the measurements at one RX PAA are independent of the configuration of the other RX PAAs. Therefore, instead of testing specific RX combinations, it is possible to just test each RX sector once and then, in post-processing, determine the performance of different combinations by combining the measurements taken at the different PAAs. We

thus implement a simultaneous sweeping with all PAAs across all sectors for the receive training in the MIMO phase. This greatly improves the scalability as the overhead of the receive training is only determined by the number of predefined sectors in the codebook and does not increase with the number of PAAs being trained.

Additionally, in the MIMO phase, we implement an option to refine the beam selection by testing different Antenna Weight Vectors (AWVs) for each sector. As accurate estimation of the inter-stream interference is crucial to this phase, if this option is activated, all possible combinations of transmit AWVs are tested. The number of combinations increases exponentially with the number of active PAAs and therefore this option improves the accuracy of the chosen beams but reduces the scalability of the MIMO phase training.

After the MIMO phase is completed, it is necessary to rank the performance of the different combinations tested and determine the optimal MIMO configuration. To this end, we choose the combinations that maximize the minimum per stream SINR as it maximizes the probability that multiple spatial streams can be established.

It is important to note that in our implementation, we make no assumptions about the transmit and receive PAA pairs that establish the streams. Instead, all possible pairs are tested and the optimal combination is selected. Additionally, we added traces to allow the user to obtain the full set of SISO and MIMO phase measurements, as well as the chosen lists of TX candidates by our selection algorithm. In this way, the user can gain insights into the MIMO performance and evaluate the MIMO BFT algorithms.

We implemented standard-compliant SU-MIMO and MU-MIMO BFT algorithms. IEEE 802.11ay specifies that the SISO Feedback can be obtained from a previous SISO BFT or an optional new SISO Transmit Sector Sweep (TXSS) can be performed. In both algorithms, we choose to support the SISO TXSS subphases to guarantee up-to-date SISO Feedback, as in this case the training is executed just before the MIMO phase. Additionally, the MIMO phase can be non-reciprocal or reciprocal, depending on whether the STAs involved in the training support antenna pattern reciprocity, i.e., the transmit antenna configurations are the same as the receive antenna configurations. For now, we support the non-reciprocal MIMO phase as it must be supported by all MIMO capable STAs and can also be used in reciprocal scenarios. Below we discuss the specifics of the SU-MIMO and MU-MIMO algorithms we implemented.

#### 5.3.3.1 SU-MIMO Beamforming Training

The SU-MIMO BFT algorithm enables training between two SU-MIMO capable devices. It includes training of the transmit and corresponding receive antenna configurations for both devices involved, which means that after the conclusion of the BFT SU-MIMO communication can be established in both directions.

Figure 5.1 shows our SU-MIMO BFT algorithm implementation. As explained above,

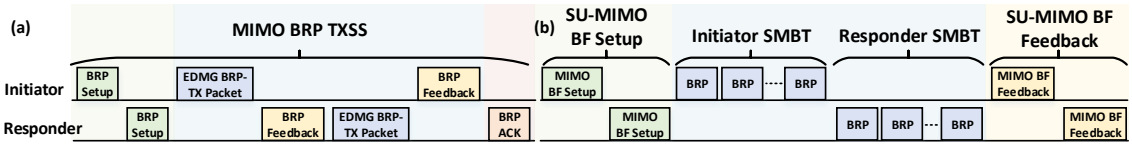


Figure 5.1: SU-MIMO Beamforming Training Phases: (a) SISO Phase; (b) MIMO Non-reciprocal Phase

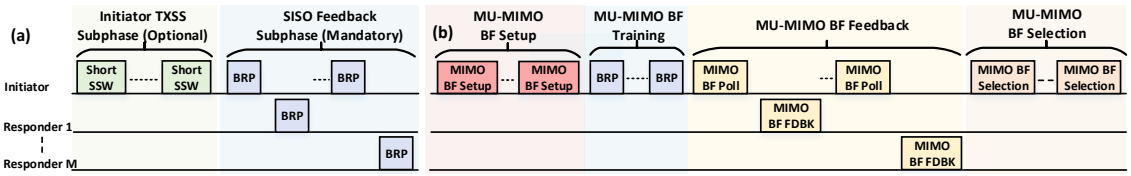


Figure 5.2: MU-MIMO Beamforming Training Phases: (a) SISO Phase; (b) MIMO Non-reciprocal Phase

it includes both the full SISO phase with the training subphases and the non-reciprocal MIMO phase.

In the SISO phase, only transmit training is performed using BRP packets with Transmit Training (TRN-T) subfields transmitted and received with multiple active PAAs. As explained in Section 5.3.2, the orthogonal design of the MIMO TRN field in these packets allows us to determine the SNR values of each transmit chain without considering any inter-stream interference. In this way, multiple PAAs can be simultaneously trained which significantly reduces the training duration and increases the scalability as the number of PAAs being trained increases.

The MIMO phase, on the other hand, involves both transmit and receive training of MIMO combinations. This is done with BRP packets with TRN-R/T subfields, which enable simultaneous transmit and receive training. The same transmit configuration is kept for as many TRN Units as the Responder has requested for receive training. During the reception of these Units, the Responder switches the RX configuration at the start of each TRN subfield. As we explained in Section 5.3.2, in this phase we record the calculated SINR values that allow us to estimate the inter-stream interference.

#### 5.3.3.2 MU-MIMO Beamforming Training

The MU-MIMO protocol, shown in Figure 5.2, is conceptually very similar to the SU-MIMO BFT protocol , with two main differences. First, during the MU-MIMO BFT an Initiator trains with multiple Responders from a MU group, requiring a modification of the Feedback phases to a poll and response format. Second, IEEE 802.11ay only defines MU-MIMO transmissions in the downlink direction and performs only transmit training for the Initiator and receive training for the Responders.

Additionally, the transmit training in the SISO phase is performed with Short Sector

Sweep (SSW) packets transmitted and received in SISO mode, instead of MIMO TRN-T subfields. This is because the Initiator is training with multiple Responders and it is not possible to guarantee that all of them will be able to receive the BRP packets. In order to reduce the training time, the new short SSW frames are used, instead of legacy SSW frames. The short SSW frame is a PHY layer frame that is 6 bytes long compared to 26 bytes for the legacy SSW which results in a 31% reduction in the transmission time.

The MU-MIMO training is performed using TRN-R/T subfields, similar to SU-MIMO. However, it requires an additional subphase called MU-MIMO BF Selection, where the Initiator informs the MU group of the Responders and optimal MIMO configurations that have been selected for MU-MIMO communication. This allows the Responders to use the correct receive configuration when MU-MIMO transmissions take place.

### 5.3.4 SU-MIMO Channel Access Procedure and Data Transmission

IEEE 802.11ay defines various methods for MIMO channel access before data transmission. As MU-MIMO data transmission is left for future work, we only discuss the SU-MIMO implementation. We implement the RTS/DMG CTS mechanism where a control trailer is added to the RTS and DMG CTS frames. The control trailer contains signaling regarding the SU-MIMO configuration used for data transmission, allowing the STAs to set up the transmit and corresponding receive antenna configurations previously trained.

Moreover, for the data transmission, we extend the `DmgWifiMac`, `MacLow`, `DmgWifiPhy` and `InterferenceHelper` classes to support transmission and decoding of MIMO packets. In the `Interference Helper`, we calculate the per stream SINR values that take into account the inter-stream interference and use this to determine the per-stream packet success rate. Analogous to the calculation of the chunk success rate, the success rate for the packet is equivalent to the multiplication of the per-stream Packet Success Rates (PSRs).

## 5.4 Evaluation

### 5.4.1 SU-MIMO Beamforming Training Validation

The first scenario to validate our SU-MIMO implementation consists of one AP and one STA, each equipped with two PAAs separated by 3cm along the x-axis, deployed in a 5m $\times$ 10m $\times$ 3m room as depicted in Figure 5.4. Each PAA is connected to a separate RF chain which allows for a maximum of two spatial streams.

Figure 5.3 depicts the results from the different phases of our SU-MIMO BFT algorithm between the AP (TX) and the STA (RX). The SISO phase measurements in Figure 5.3 (a) show the SNR of the different transmit sectors from both TX PAAs

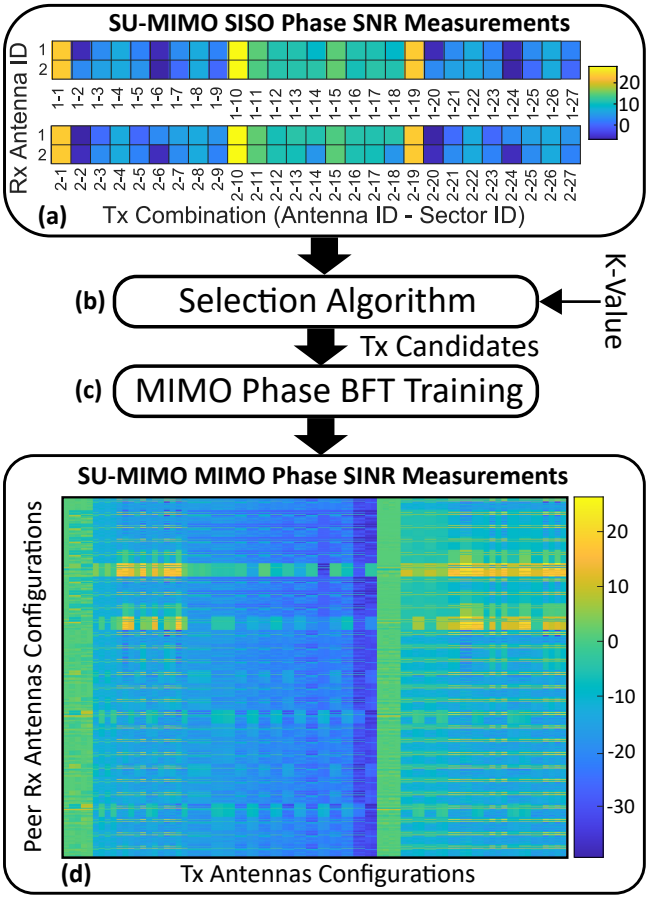


Figure 5.3: MIMO Beamforming Training Results

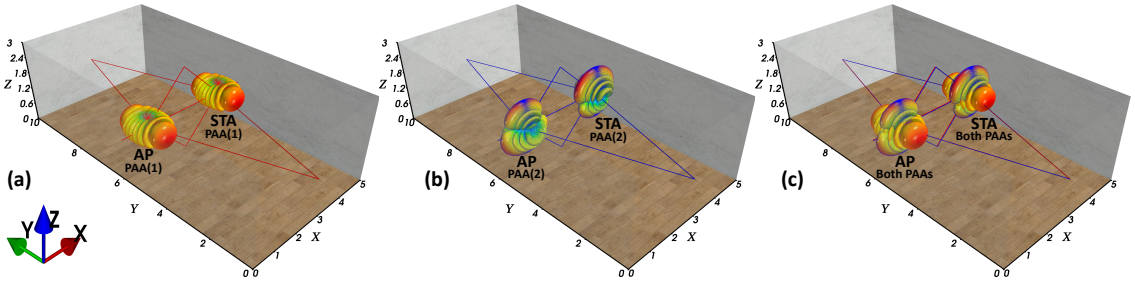


Figure 5.4: SU-MIMO Beamforming Training Qualitative Results: (a) PAAs Beam Patterns Corresponding to Stream 1; (b) PAAs Beam Patterns Corresponding to Stream 2; (c) Combined PAAs Beam Patterns for Streams 1 and 2

measured at both RX PAAs. Since the PAAs separation distance is small, we can observe that the SNRs from the same transmit sector at both receiver's PAAs are very similar in most cases. The SISO results then serve as input to our selection algorithm that selects the top $K$ combinations as shown in Figure 5.3 (b). The list of $K$ candidates is tested in the MIMO phase shown in Figure 5.3 (c), which results in a set of SINR measurements.

For this scenario, we use the top $K$=85 combinations tested, as we observed that this value offers a good compromise between scalability and accurate SU-MIMO configuration. In Figure 5.3 (d) we present a heatmap of the minimum per stream SINR for each tested candidate. On the x-axis, we show the different TX candidates according to their ranking by the selection algorithm, the first column representing the candidate with the highest joint SNR. On the y-axis, we present the different receive combinations tested. As explained in Section 5.3.3, we can determine the SINR for all possible receive combinations and we present them sequentially (the bottom row representing (RX PAA 1 - Sector 1, RX PAA 2 - Sector 1) and the top row representing (RX PAA 1 - Sector 27, RX PAA 2 - Sector 27). We can see that the highest-ranked candidates (leftmost columns) experience low SINR. For these candidates, both PAAs have beam patterns that utilize the Line-Of-Sight (LOS) path as it gives the highest SISO SNR. However, when used for MIMO communication, such a combination results in high inter-stream interference due to the small PAA separation. The candidates with the lowest measured SINR similarly suffer from high inter-stream interference, but additionally have lower received signal strength as they utilize reflected paths. This shows the significance of the MIMO phase, as the optimal SISO configurations can sometimes result in poor MIMO performance. Additionally, we can observe no obvious pattern in the SINR measurements for the different configurations tested. This implies that it can be extremely challenging to predict the MIMO performance from the SISO Feedback and that the selection of good candidates for the MIMO phase is crucial to the overall functioning of the MIMO BFT algorithms. As mentioned in Section 5.3.3, our implementation was designed to be able to evaluate the effect of different selection algorithms and can therefore be of crucial interest to study mmWave MIMO behavior. Finally, we observe two high SINR areas. The first area, located in the top left half allows for SINRs of around 15 dB. However, by testing a higher number of candidates we discover a second high SINR area in the top right half of the map with more optimal antenna configurations that can achieve SINRs above 20 dB.

Figure 5.4 shows a visualization of the best SU-MIMO configuration chosen by our BFT algorithm. We can see that the first stream established, shown in Figure 5.4 (a), utilizes the reflections from the front and back walls and has very low gain for the LOS path and the reflections from the side-walls and the ceiling/ground. The second stream, shown in Figure 5.4 (b), utilizes precisely those links and receives very low interference from the front and back wall reflections. The resulting combination shown in Figure 5.4 (c) has very high per stream SINR of 23.52 dB and 39.25 dB respectively, validating that our BFT algorithm can successfully determine good antenna configurations for MIMO communication.

Finally, after the BFT is completed, we validate our SU-MIMO data transmission implementation using the output of the MIMO Phase BFT to set up the transmit and receive antennas. The large SINR experienced by the two streams enables the use

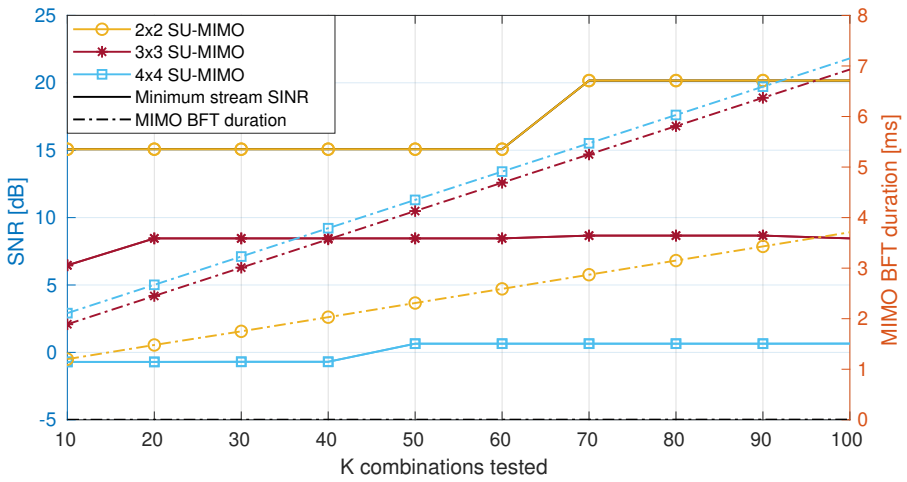


Figure 5.5: SU-MIMO BFT Training duration and achieved SINR depending on K

of Enhanced Directional Multi-Gigabit (EDMG)-Single Carrier (SC) Modulation and Coding Scheme (MCS)-21 (8 Gbit/s). We observe an aggregate throughput of around 14 Gbit/s, validating the multi-stream transmission implementation.

To investigate further the performance of SU-MIMO BFT we use a second scenario where the STA begins at a distance of 3 m from the AP and moves in a straight line towards the AP until it is 10 cm away. The STA moves with an average pedestrian speed of 1.5 m/s. In this scenario, we study devices with 2, 3 and 4 PAAs to look at how many spatial streams can be supported by the mmWave channel. Additionally, we consider an antenna separation of 3, 5 and 7 cm to look at how it affects the spatial diversity of the links.

First, we look at the trade-off between the duration of the MIMO BFT and the SINR that can be achieved. For this purpose, we analyse the SINR in the initial position of the STA and look at the performance setting $K$ between 10 and 100. As an example, we look at an antenna separation of 3 cm and we show the minimum per-stream SINR as it limits the data throughput that can be achieved. Figure 5.5 shows the SINR and duration of the BFT for different values of $K$. First, we can see that for the 2x2 MIMO case, similar to the previous scenario, testing a number of combinations larger than 60 results in a gain of approximately 5 dB. This is enough to increase the MCS from EDMG SC MCS 16 to EDMG SC MCS 20 for a gain of approximately 4 Gbps of achievable throughput. The higher SINR comes at the cost of a long training time of over 3 ms for a single STA which represents significant overhead for the operation of the network. When increasing the number of antennas, we can observe that the achievable SINR drops significantly and that training more combinations does not provide a significant gain in SINR. The first cause of this is that the mmWave channel is sparse, making it challenging to establish independent spatial streams using solely analog beamforming. The intra-

stream interference poses a fundamental limitation to the achievable SINR, requiring more complex hybrid beamforming to successfully establish independent streams. Additionally, the number of possible antenna combinations grows exponentially with the number of PAAs that need to be trained. Therefore, it is necessary to train a much larger number of combinations to discover new beam configurations with a higher SINR.

Finally, we observe that there is an interesting trend in the training time as the number of PAAs grows. There is a significant increase between the training time for 2x2 and 3x3 SU-MIMO, and an additional minor increase between 3x3 and 4x4 SU-MIMO. Interestingly, we found that the increase was not related to the training time, but rather the feedback time. In fact, the training time is not related to the number of antennas that need to be trained, as in the SISO phase all antennas are trained at the same time and in the MIMO phase the number of combinations tested is solely determined by $K$. However, the number of measurements taken in the SISO phase grows with the number of antennas which increases the size of the feedback packet and thus the total duration of the training. Due to restrictions on the size of the control feedback packet, there is an upper bound on how many measurements can be sent back. We found that this limit is already reached when training 3 PAAs, which is why there is not a large increase when the number of PAAs grows to 4. The minor increase we observe is from the feedback for the MIMO phase. For the MIMO phase we only feedback the top antenna configurations and the size of each configuration grows with the number of antennas trained.

Lastly, we validate our SU-MIMO BFT algorithm in case of user mobility. Figure 5.6 shows the theoretically achievable data rate over time, based on the SINR enabled by the beamforming training. We compare the performance with 2, 3 and 4 PAAs per device and 3 cm, 5 cm and 7 cm antenna separation. To cope with mobility we initiate SU-MIMO BFT once every 10 Beacon Interval (BI), or each 1.1 s, setting $K = 100$. The timing of the BFT is shown with black vertical lines on the figure. The STA moves with a speed of 1.5 m/s and reaches the final position at 2.010 s. As a reference, we also show the achievable data rate maintaining a SISO link. We can see with SISO communication we can consistently use the highest MCS to transmit data, resulting in a constant 8.085 Gbps throughput. As the STAs moves in a straight line towards the AP, SISO communication always uses the LOS link. Thus, mobility does not cause problems for the BFT and in fact, the achievable SNR constantly grows throughout the simulation.

SU-MIMO communication, on the other hand, is very sensitive to mobility as it uses Non-Line of Sight (NLOS) links which can vary as the STA changes position. We can see that in Figure 5.6a, 5.6b and 5.6c there are time intervals where the achievable throughput falls below what is possible with a SISO link, showing the complexity of maintaining MIMO communication in a mmWave network. This is most pronounced in the case of 4x4 SU-MIMO, where due to the high inter-stream interference most of the time it is difficult to maintain a link and only the lowest MCS 1 can be used. We can also see that

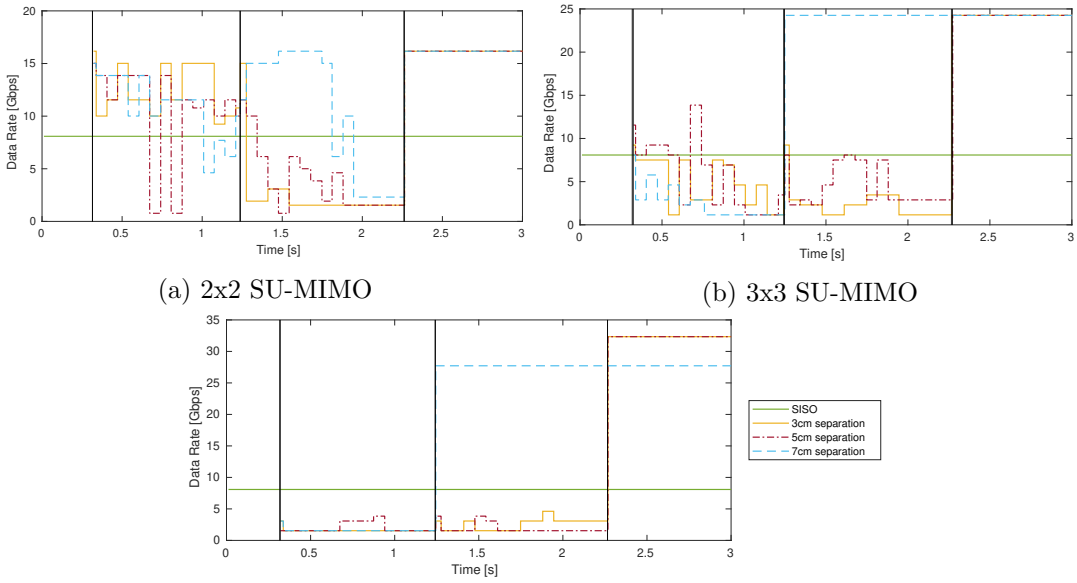


(a) 2x2 SU-MIMO

(b) 3x3 SU-MIMO

(c) 4x4 SU-MIMO

Figure 5.6: Achievable data rate over time

more frequent BFT can help cope with mobility as the data rate tends to rise after the BFT. However, due to the high overhead of MIMO BFT, this might cause disruptions to the data traffic. With regard to the antenna separation, we observe that in most cases the higher separation increases the diversity of the channel and enables higher data rates. Finally, we notice that in all scenarios once the STA has reached its final position and the last BFT is done there is a large jump to the highest achievable data rate. We found that this is because, in the final position, the STA is extremely close to the AP and has very high signal strength. This allows for communication with a high SNR with multiple diverse beampatterns. Thus a more diverse set of beams can be chosen, resulting in a lower inter-stream interference.

### 5.4.2 MU-MIMO Beamforming Training Validation

In this scenario, we deploy one EDMG AP and two STAs in the same room as depicted in Figure 5.7. The AP is equipped with two RF chains, each connected to a separate PAA, while the two STAs have a single PAA. As a result, the AP can transmit two spatial streams, allowing communication with two users at the same time. The STAs have a high angular separation between them of 120°with respect to the AP. In Figure 5.7 we show the optimal MU-MIMO configuration chosen by our algorithm. We can see that the high spatial separation between the STAs allows us to have two streams that utilize different multi-path components. The resulting per stream SINRs of 33.8 dB and 23.4 dB are very high and will be sufficient for MU-MIMO communication with high data rates.

Finally, we investigate the performance for different angular separations between the

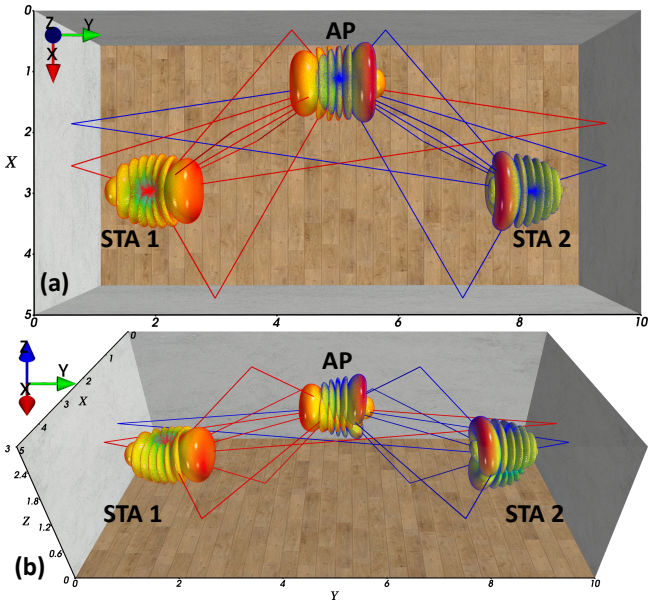


Figure 5.7: MU-MIMO Beamforming Training Qualitative Results (a) Top View; (b) Side View

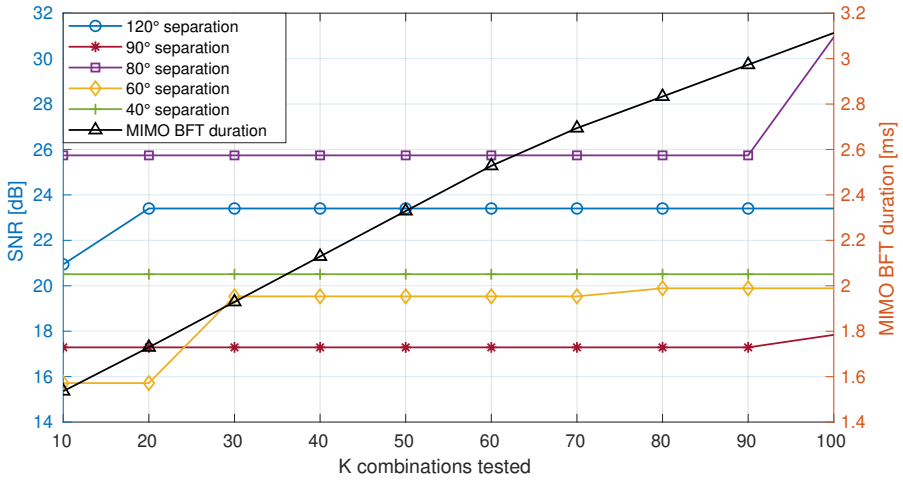


Figure 5.8: MU-MIMO BFT Training duration and achieved SINR depending on K

STAs and different values of $K$. Figure 5.8 shows the minimum achievable SINR as well as the BFT duration with regard to $K$. We can see that compared to SU-MIMO, it is easier to achieve high SINRs even with a low value of $K$, as the separation between different STAs is much larger than between different PAAs on the same STA, allowing for higher spatial diversity. Interestingly, we found that a higher angular separation does not always result in higher achievable SINR due to the limited ability to steer the beampatterns in all directions which causes variations in the signal strength, as well as the beampattern sidelobes which affected the intra-stream interference.

## 5.5 Conclusion

In this chapter, we presented the significant extension of the IEEE 802.11ad/ay module to support MIMO operation, which included extensions of the channel model, the PHY layer and many new functionalities in the Medium Access Control (MAC) to support the novel MIMO BFT protocols. The goal was to preserve the scalability of the simulations by only scheduling a single event per MIMO packet. The implementation of the MIMO protocols revealed their complexity and length, with the candidate selection algorithm playing a crucial role in the performance due to the impossibility of an exhaustive search of all possible configurations. Our evaluation demonstrated the viability of a fully analog beamforming architecture, as we were able to successfully establish multiple communication streams in both a SU-MIMO and MU-MIMO setup. However, we also found trade-offs between the quality of the established beams and the length of the training, difficulty in maintaining the links in case of user mobility and challenges in finding sufficiently separated analog beams for SU-MIMO with more than 2 PAAs. In the following chapter, we explore an alternative multi-RF architecture which can address issues like resilience and robustness with a simpler BFT and system model design.

# 6 In-Band Multi-Connectivity for Improving mmWave Network Resilience

## 6.1 Introduction

The exponential rise of wireless devices and new, high-performance applications continuously increase demands from wireless networks. The need for higher capacity, lower latency and higher reliability demand innovative designs and an evolution of current networks. The large bandwidth of Millimeter-Wave (mmWave) networks has enabled extremely high data rates and low latency, making them key enablers for 5G and future 6G applications. However, the increased attenuation, sensitivity to blockage and need for directionality lead to problems with low coverage ranges, unstable links and reduced reliability [10, 11]. The challenges of robustness and reliability present a significant challenge for the wide deployment of mmWave technology. State-of-the-art mmWave electronics design is enabling the development of multi-Radio Frequency (RF) chain devices that can push mmWave networks to satisfy the highest performance demands. So far, their use has mostly been considered in traditional Multiple-Input and Multiple-Output (MIMO) architectures, with multiple independent streams to a single Access Point (AP). We investigated the performance of the standardized mmWave WiFi protocol in Chapter 5. Due to the sparsity of the mmWave channel, however, successfully establishing multiple independent streams for a MIMO setup can be challenging, with a long and complex beamtraining process. Additionally, this approach only aims to increase the link data rate and can not prevent link blockage and loss of service, which are arguably much more relevant issues for mmWave networks where the large bandwidth already allows for high data rates even with a Single-Input and Single-Output (SISO) connection.

An alternative network architecture has been proposed by cell-free networks as a way to improve coverage and capacity and provide uniform service for all users [47, 48]. In cell-free networks, a large number of APs are densely distributed and jointly and simultaneously serve clients. The idea is that by maintaining simultaneous connections with multiple APs we can exploit the spatial diversity against shadow fading and improve

network coverage. Additionally, by removing the division of the network into cells and having APs jointly serve clients, cell-free networks can provide more uniform service with a user-centric approach [112]. As the client simultaneously communicates with multiple APs, it is necessary to have tight coordination and synchronization across the whole network to enable phase-coherent signal processing across the geographically distributed APs [113]. This remains a key challenge for the further development of cell-free networks.

More practical approaches to exploiting link diversity have been standardized in 3GPP as multi-connectivity. Various multi-connectivity architectures have been considered, from simultaneous connections in the sub-6 GHz and mmWave band to multi-Radio Access Technology (RAT) diversity. In-band multi-connectivity similar to cell-free has been introduced as Coordinated Multi-Point (CoMP) where the network is divided into disjointed clusters that jointly serve users in their joint coverage area. Depending on the number of device RF chains, the user can either establish multiple communication paths with different APs or only communicate with one at a time, switching between links to maximize performance. CoMP has lower implementation complexity compared to a full cell-free network, as synchronization is only required between a cluster of APs in the same geographical area. However, this network-centric approach can result in sub-optimal coverage as compared to a cell-free design. In addition, joint transmission from multiple APs still requires phase-coherent processing, which can be challenging to implement.

Applying in-band multi-connectivity to mmWave networks can be particularly beneficial. In addition to the spectral efficiency benefits, enabling users to simultaneously connect to multiple APs can address the key challenges of resilience and coverage, reducing link breakages. Prior studies demonstrate that CoMP enhances coverage and link reliability in mmWave networks [44, 45]. The impact of interference and blockage in a mmWave CoMP system was also experimentally studied in [49, 50], showing its ability to suppress interference and its robustness against link blockage. A full cell-free design or CoMP with phase-coherent processing, however, can be extremely challenging to practically implement at mmWave frequencies, due to the tight synchronization requirements and need for global channel state information if full interference suppression is desired. The use of narrow directional beams in mmWave introduces the opportunity for a different multi-connectivity design that relies on local spatial separation of signals for interference management. In contrast to solutions that rely on synchronization, this approach can be implemented in a decentralized manner with very low AP coordination. This design can also incorporate the benefits of cell-free networks - as there is no need for tight AP synchronization it is possible to have a user-centric design with no cell division, where users can be served by any AP in the network. In such network implementations, the local beamtraining which selects the analog beams becomes the key performance factor, as it needs to ensure sufficient interference suppression.

In this paper, we implement and evaluate a multi-connectivity mmWave network to

show it can enhance the overall performance, but especially the reliability of mmWave networks. Our goal is to evaluate the benefit of investing more resources in a multi-RF architecture to improve the reliability and stability of mmWave networks. Unlike most previous work that studies multi-connectivity from an information-theoretic or signal processing perspective, we look at how such a system would work when deployed with a full protocol stack and what network designs are needed for its functioning. For this purpose we use the network simulator ns-3 and implement a multi-RF, multi-connectivity mmWave network, comparing its performance with a standard small cell mmWave deployment. Our system design focuses on simple and practical implementation, taking advantage of the spatial sharing properties of mmWave to design networks with a very low level of network coordination. To avoid the need for phase-coherent processing and PHY layer cooperation, in our system, all links are considered interference and we rely only on local analog beamforming to perform interference suppression. Thus, each AP decodes the signals from the user without cooperation with the other APs. Furthermore, we also do not coordinate channel access on the Medium Access Control (MAC) layer, allowing for asynchronous, independent transmissions. While performance could be enhanced by using scheduled access to lower interference and collisions, this necessitates building a global schedule that implements spatial sharing, which is complex and costly, especially when considering multi-connectivity. Instead, we chose a simpler approach that still validated the benefits of multi-connectivity for mmWave. Therefore, we base our implementation on mmWave Wi-Fi [27] as Wi-Fi networks represent the perfect examples of dense distributed networks without PHY and channel access coordination.

Our work proves that this network design is viable and can outperform a traditional single-connectivity design. We observe large and consistent gains in reliability, reducing the time spent with no service from up to 20% to 0% in all simulation scenarios we ran, and improvements in the overall network throughput of up to 30%. Interestingly, we observe that mmWave multi-connectivity has a bigger benefit in denser scenarios, where the spatial diversity allows to avoid blockage and constant packet collisions from nearby nodes. Lastly, we study how network design choices such as the beamforming training method affect the performance, exploring different ways to enhance the multi-connectivity implementation.

The rest of this chapter is organized as follows: In Section 6.2 we present our system model design and in Section 6.3 we show an evaluation of the performance with ns-3 simulations. Section 6.4 gives an overview of prior work and Section 6.5 concludes the chapter.

## 6.2 System Model and Design

Cell-free and multi-connectivity CoMP architectures exploit spatial diversity by enabling simultaneous service by multiple APs. The key challenge in implementing these designs at mmWave frequencies lies in the tight synchronization requirements for phase-coherent processing. However, by leveraging the spatial multiplexing property of directional communication it is possible to design simpler multi-connectivity implementations which instead rely on analog beamforming for interference suppression. The goal is to use beams that allow for high spatial reuse so that multiple links that are sufficiently spatially separated can be simultaneously active, making the local beamtraining the crucial performance factor. Such an approach also removes the need to design transmission schedules that incorporate spatial sharing, which is complex to do in a dynamic environment where the interference varies not only with the device locations but also with the transmit and receive beampatterns used. In this way, it is possible to design future small cell mmWave deployments to incorporate distributed features from current Wi-Fi.

Therefore, we base our design on mmWave Wi-Fi. We consider indoor scenarios with very dense AP deployment to ensure coverage from multiple APs. All multi-connectivity devices have multiple Phased Antenna Arrays (PAAs) and each PAA is connected to an RF chain. Stations (STAs) use the RF chains to associate and communicate with different APs simultaneously, splitting the data uniformly between the RF chains. To focus on the benefits of multi-connectivity spatial multiplexing only, we consider that APs serve only 1 client per RF chain. APs are only loosely coordinated by a central controller that manages association and packet allocation. This allows us to integrate the user-centric approach of cell-free networks, where clients can connect to any AP in the network, regardless of cell limits. Transmissions occur without any synchronization or cooperation on the MAC or PHY layers, relying only on analog beamforming for interference management. Thus, concurrent packets received at different RF chains on the same device are considered (self-)interference and we rely on local training to achieve good spatial separation to successfully receive them.

Our objective is to explore the viability of in-band mmWave multi-connectivity with local beamforming with a design that minimizes network synchronization for a more practical deployment. In this way, our results represent a lower bound for multi-connectivity, as hybrid beamforming with digital precoding, as well as additional network synchronization or channel access coordination, would improve interference management and result in higher data rates. We compare the multi-RF multi-connectivity system with a standard single-RF implementation where each STA is only associated with one AP at a time. As we are primarily interested in increasing mmWave network resilience, and not in increasing the per-link throughput, we believe this is a better comparison than a

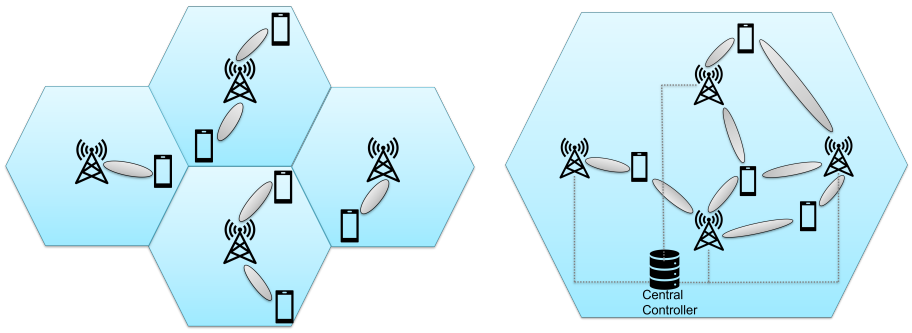


(a) (b)

Figure 6.1: (a) Single-Connectivity vs (b) Multi-Connectivity Multi-RF Implementation

MIMO system. Additionally, we study data rates below the maximum capacity of a single mmWave Wi-Fi channel, meaning that there is no need to use MIMO to serve the offered traffic.

Instead, we evaluate the trade-off between the cost of a multi-RF device and the improvements in reliability that can be gained. Figure 6.1 shows a comparison of the two implementations. In the single-RF case, STAs have multiple PAAs, but only one is active and connected to the RF chain through which all data gets transmitted. We explore two options on how to select the active PAA: 1) random selection or 2) choosing the PAA that gets coverage from the highest number of APs to maximize the probability of good coverage. In the multi-connectivity, multi-RF case, APs jointly serve STAs within each AP's coverage, enabling STAs to establish multiple directional links with multiple serving APs. The number of directional links is upper-bounded by the number of RF chains of the STA. All APs are connected to a central controller via reliable fronthaul links, enabling AP selection for each STA. In this way, STA benefit from the overlapped coverage of multiple APs.

### 6.2.1 Antenna Model

We consider devices with 4 PAAs, each oriented in a different direction, as shown in Figure 6.2a. In the multi-connectivity implementation, each PAA is connected to a separate RF chain, while in the single-RF case only one PAA is active at a time. We use a 2x8 element Uniform Rectangular Array (URA) architecture to generate narrow beams to spatially separate signals. Figure 6.2b shows examples of the generated directional beampatterns. We notice that the beampatterns are back baffled, meaning that the PAAs do not generate a response on the back side of the antenna. This means that STAs can lose connection if there is no available AP in their field of view. During communication, we use both directional transmission and reception, as they are crucial both for communication and interference management.

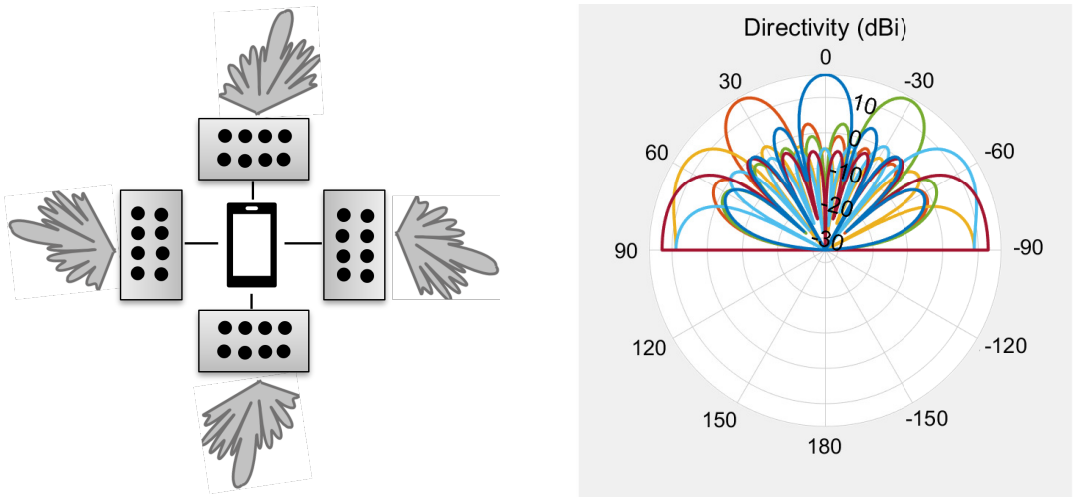


(a) (b)

Figure 6.2: (a) PAA orientation (b) Directivity of the antenna beampatterns

### 6.2.2 Beamforming Training for Multi-Connectivity Operation

Local analog beamforming training is performed independently by each STA. In the training, devices test different beampatterns from a pre-determined codebook to find the optimal one for communication. We are considering dense networks where each device has multiple PAAs that need to be trained, which can result in very high training overhead. To avoid excessive overhead, we investigated different training methods, focusing on efficiency, scalability, and sensitivity to interference. In the end, we implement a mechanism based on the Group Beamforming principle evaluated in Chapter 3, adapting it to our needs. The key change was to make the training user-initiated as users require training with multiple APs. Thus, a user-centric approach can reduce overhead and improve efficiency and accuracy. We train each device PAA individually and do not take into account the inter-PAAs interference within a device when selecting the beams. We note that the performance can be improved by jointly training the PAAs to minimize the interference between them; however, this adds complexity and overhead. Not only is the number of possible combinations of receive and transmit beampatterns too large to do an exhaustive search but AP coordination is needed to measure the joint interference. In Chapter 5 we show the length and complexity of such a beamtraining mechanism. As the goal of this work is to evaluate a system with minimal coordination we leave joint-PAAs training for future work.

In our modified Group Beamforming, a STA initiates the training by transmitting multiple packets, each with a different directional beampattern. APs that receive the packets measure the Signal-to-Interference-plus-Noise Ratio (SINR) of each packet, determining the quality of each trained beampattern. Thus, each AP can decide which of the trained transmit beampatterns the STA should use to communicate with them. The

novel aspect of Group Beamforming is that the STA additionally appends an element called a Training (TRN) field to each training packet. The TRN field has multiple subfields and each subfield is composed of Golay sequences which have good correlation properties. This allows multiple additional beampatterns to be trained in the same packet, by fast beampattern switching between subfields. APs that receive the packets use a different receive beampattern for each subfield. By measuring the SINR of each subfield they can determine the optimal beampattern for reception from the STA. If we consider antenna reciprocity, meaning that transmit and receive beampatterns are equivalent, we can get the full transmit and receive configuration for the STA and all APs. As multiple beampatterns are trained in the same packet, this approach minimizes the training overhead. Moreover, by relying on receive rather than transmit training we can both train many APs at the same time and reduce the training interference since APs do not transmit during the training.

This method results in low overhead and high accuracy. However, in some dense scenarios, it suffers from errors caused by interference, particularly on the STA side. We observed the same behavior in Chapter 3 where an enhancement that used TRN training for both APs and STAs resolved most errors. Therefore, we also implement and evaluate this enhanced version, to observe how the beamforming training accuracy affects network performance in interference-challenged scenarios.

### 6.2.3 Association

We consider that the association is determined by the central controller based on the maximum received power between two devices. Each AP RF chain only serves a single client for even load distribution and clients associate to a different APs with each RF chain. STAs try all APs within their field of view, in descending order of the received power, until they find an available AP. Some STAs (or STA RF chains in the case of multi-RF multi-connectivity devices) will not have a connection because they are in a bad location with no available APs in their field of view. To try to ensure minimal service for each client, priority during the selection is given to the clients with the lowest maximum received power. The same association process is used for the single-RF implementation, with the distinction that STAs only associate to one AP with the single RF chain connected to the active PAA.

## 6.3 Performance Evaluation

### 6.3.1 Simulation Scenarios

Unlike previous work on cell-free and multi-connectivity which mostly looks at analytical performance bounds simplifying practical protocol or training aspects, we

Table 6.1: Simulations Parameters

| Parameter Name | Parameter Value |
|---|---|
| Transport Protocol | UDP |
| Aggregation Type | A-MSDU and A-MPDU |
| A-MSDU Max. Size | 7935 Bytes |
| A-MPDU Max. Size | 4 194 303 Bytes |
| MAC Protocol | CSMA/CA |
| Number of Codebook Sectors | 30 Sectors |
| Transmit Power | 10 dBm |
| Operating Frequency | 60.48 GHz |
| Simulation Time | 50 s [1] |

consider a more complete implementation that takes into account aspects like channel access, packet collisions and interactions between APs. For this purpose, we expand our IEEE 802.11ad/ay ns-3 implementation to support multi-connectivity, specifically, our system model as described in Section 6.2. This involves the creation of STAs that can associate to different APs with each RF chain and independently communicate with them. In addition, we enable STAs to use the Group Beamforming mechanism for our user-initiated beamforming training. Finally, we generate a variety of simulation scenarios to evaluate multi-connectivity under different conditions.

We study indoor environments using the Quasi-Deterministic (Q-D) channel realization software [39]. The default scenario is a rectangular room ($7.4\ m \times 13.5\ m \times 3\ m$) and we also study a larger room ($29.6\ m \times 54\ m \times 3\ m$). We simulate dense deployments with 10, 20 and 40 APs for sufficient coverage. To isolate the effects of multi-connectivity the number of STAs is reduced to 2, 4 and 8 respectively. Devices are randomly placed in a fixed position and we average the results over 50 simulations with different device locations. We study both downlink and uplink traffic, with varying per-STA data rates of 800 Mbps, 1.6 Gbps, 3.2 Gbps and 6.4 Gbps. Other simulation parameters are listed in Table 6.1.

### 6.3.2 Simulation Results

We analyse multi-connectivity under different levels of network load and interference by varying both the network density and data rate. As discussed in Section 6.2, we compare the performance to a single-RF design, using two strategies to select the active PAA - randomly or to maximize coverage. Analysing both the individual and aggregate throughput we identified three network states of low, medium and high channel load with similar behaviour and performance trends, discussed below.

[1] We also tested longer simulation times (2 min, 5 min) and found no changes in the statistics of the performance as the effects we observe happen on the sub-second level.

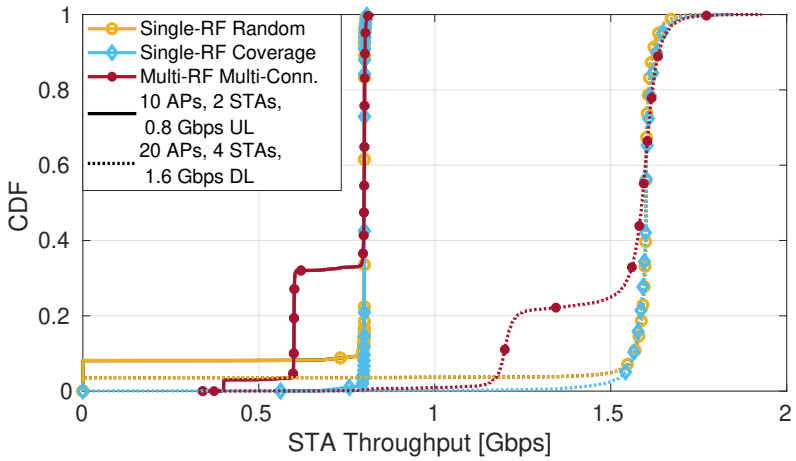


Figure 6.3: STA Throughput CDF, low network load.

#### 6.3.2.1 Low Network Load

In our environment, an aggregate network load of up to 6.4 Gbps under-saturated the wireless medium, as the maximum capacity of a single 2.16 GHz IEEE 802.11ay wireless channel using the highest single-carrier Modulation and Coding Scheme (MCS)-21 is 8 Gbps. Channel access procedures and packet collisions lower the effective achievable rate, however, traffic up to 6.4 Gbps does not overload the network and causes no congestion.

Figure 6.3 shows the Cumulative Distribution Function (CDF) of the per-STA throughput for two such scenarios: a network of 10 APs, 2 STAs and a per-STA uplink data rate of 0.8 Gbps and a network of 20 APs, 4 STAs and a per-STA downlink data rate of 1.6 Gbps. The scenarios were shown as an example to demonstrate the identical performance trends which are also observed in all other tested scenarios below 6.4 Gbps. First, the targeted data rate is achieved by a large majority of STAs in both the single-RF and multi-RF deployments. However, we also observe an interesting pattern for multi-connectivity, where approximately 20% of STAs only achieve 75% of the targeted rate. This reflects the increased risk that one of the 4 active RF chains is blocked or does not have coverage. Multi-connectivity performance in this case can be improved by redistributing the data from the blocked chain to the other chains, with the added complexity of detecting the loss of service and determining how the data should be distributed. In the single-RF case, this probability is much lower. However, the harm if it happens is much larger - when the RF chain is blocked, the STA completely loses service. We can see that with random PAA selection, this happens approximately 8% of the time in the 2 STA network and 3% in the 4 STA network. In low-load scenarios, selecting the PAA to maximize coverage helped us cope with these issues. Finally, we observe that the traffic direction does not affect the performance. As each AP only serves a single STA, uplink transmissions do not have added channel contention and the performance is equivalent to downlink scenarios.

**Lessons learned**: When the network is lightly loaded, the main challenge is coverage. In this case, needing coverage for multiple RF chains from different directions can lead to lower achievable data rates as there is simply no AP available in the field of view. However, loss of coverage in the single-RF case is more harmful as it leads to complete loss of service. There are scenarios where multi-connectivity pays off, but in others, it can have disadvantages. This demonstrates the complexity and sensitivity of mmWave networks and why they require tailored networking approaches.

#### 6.3.2.2 Medium Network Load

From the tested rates, we found that scenarios with an aggregate load of 12.8 Gbps represented medium network load, where the channel is saturated and spatial sharing is necessary to achieve the desired data rate. Therefore, the performance varies based on the network topology, traffic parameters and especially the location of interfering devices. We can see this in Figure 6.4a, which shows the throughput CDF in two medium-load scenarios - a network with 40 APs, 8 STAs and a per-STA downlink data rate of 1.6 Gbps and a network with 10 APs, 2 STAs and a per-STA uplink data rate of 6.4 Gbps. We can observe clear differences in multi-connectivity performance in the two scenarios. We found that this data rate represented a transition region where performance depends on the network configuration, with the first scenario (solid lines) representing the worst multi-connectivity performance and the second (dotted lines) the best one. All other tested configurations with an aggregate rate of 12.8 Gbps were between these two scenarios. In the first scenario, the behaviour is similar to the low-load case, as multi-connectivity performance is slightly disturbed by the higher probability of blocked RF chains. However, as the single-RF performance is also slightly degraded by interference, all three deployments have similar performance.

The second scenario is significantly different. Here we begin to see how a multi-connectivity design can improve mmWave networks. By distributing the load among 4 RF chains, the multi-connectivity implementation manages to cope with interference better and especially to ensure uninterrupted service. STAs receive half of the desired data rate only 5% of the time, compared with 10% and 15% in the single-RF deployments. We also highlight how in the single-RF case where the active PAA is randomly chosen, STAs are left without a network connection 9% of the time due to no coverage or high interference. Multi-connectivity, instead, manages to offer a much more uniform service to users, which leads to a gain of over 0.8 Gbps in the average per-STA throughput. Figure 6.4b presents the aggregate network throughput for this scenario, showing that not only does the multi-connectivity improve the median aggregate by approximately 2 Gbps, but the gains are achieved by boosting the lowest-performing scenarios. This was a consistent trend we observed, where even when multi-connectivity had lower average performance, it had higher lower performance bounds. This validates both the benefit for

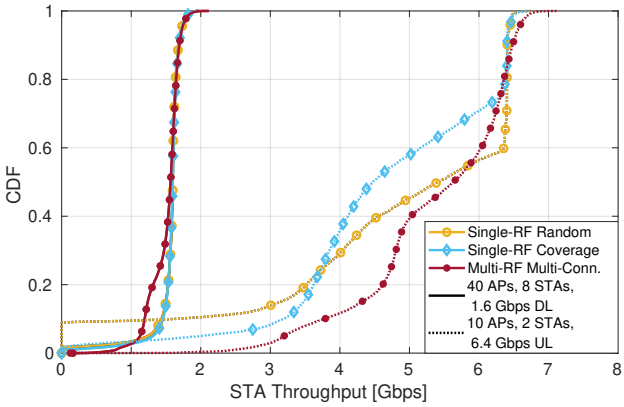


(a) STA throughput CDF

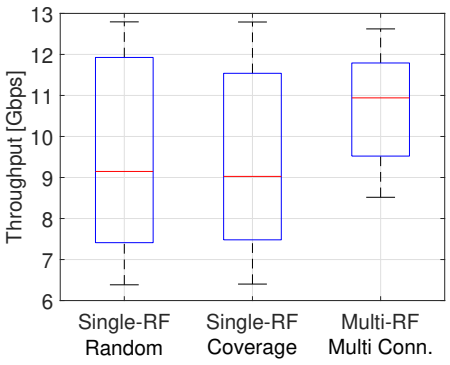


(b) Aggregate network throughput, 10 APs, 2 STAs, 6.4 Gbps UL

Figure 6.4: Performance in medium network load scenarios.

increased reliability and the ability to provide more uniform service to all the clients in the network.

Another interesting trend is when comparing single-RF performance with active PAA selection to maximize coverage and randomly. In Figure 6.4a we can observe that although selecting the PAA to maximize coverage still reduces the probability that the STA loses coverage, the performance degrades above 3.7 Gbps. Here, the random PAA selection performs much better and STAs can achieve the wanted data rate an additional 16% of the time. We can see in Figure 6.4b that the resulting aggregate network throughput is equal in both cases. The reason for this poor performance when choosing the PAA to maximize coverage is that it can inadvertently select a PAA that is transmitting towards a zone with multiple APs and thus high interference in the uplink direction. Therefore, we found that this PAA selection had poor performance in uplink scenarios, an unexpected performance trend that showed the complexity of interaction in dense mmWave networks.

Finally, we found that multi-connectivity was also affected by the traffic direction, with an opposite trend. In this case, downlink traffic caused interference problems since the STA is receiving with all 4 RF chains simultaneously. While analog beamforming proved to be surprisingly robust, high-interference environments were still challenging and performance suffered, as compared to uplink scenarios. To overcome this, it is necessary to perform more complex beamforming training, which attempts to minimize the interference between the RF chains of the STA or to use hybrid beamforming by adding digital precoding.

**Lessons learned:** At higher network loads, interference begins to affect the performance. In this case, spatially distributing the load across multiple RF chains leads to gains in performance. In addition, it's important to consider issues like inter-RF interference in multi-RF chain systems, or selecting APs from low interference zones for single-RF systems.

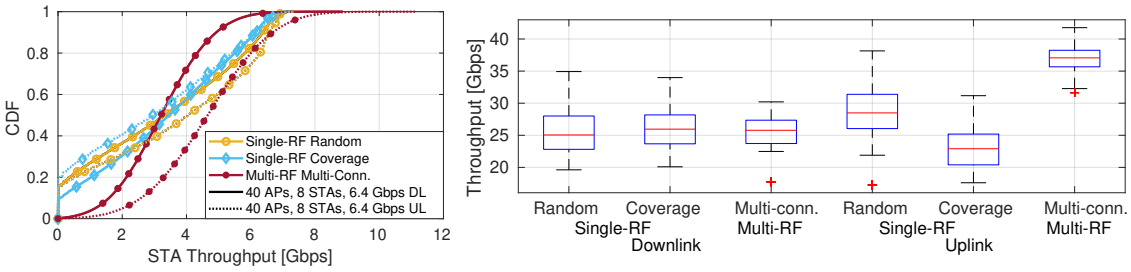


(a) STA throughput CDF

(b) Aggregate network throughput

Figure 6.5: Performance in high network load scenarios.

#### 6.3.2.3 High Network Load

The highest network loads of 25.6 Gbps and 51.2 Gbps oversaturated the wireless medium, pushing the network to operate in overloaded mode. In these cases, the desired data rate was extremely difficult to achieve due to high interference. The performance in these scenarios matched the observations of the second scenario with medium network load. There were two main trends, with a difference between the uplink and downlink direction, as seen in Figure 6.5 which shows the performance in a network with 40 APs, 8 STAs and a per-STA data rate of 6.4 Gbps.

We can observe that in the uplink case, the multi-RF multi-connectivity deployment is able to significantly outperform the single-RF deployments, with the highest gains in the low throughput areas below 2 Gbps, leading to an average aggregate network throughput gain of over 10 Gbps (30%). Crucially, even in the highest-interference scenarios, our multi-connectivity implementation ensures that STAs never lose connectivity and the per-STA throughput is always above 0.5 Gbps. This shows a remarkable improvement in the stability and reliability of mmWave networks, where link breaks and service interruptions are identified as key challenges. This is noticeable in the single-RF implementations, where STAs experience outage up to 15% and 20% of the time. We can also see that in this case, selecting the active PAA to maximize coverage consistently results in suboptimal performance. Due to the oversaturation of the wireless channel, interference is the dominant problem and selecting the active PAA to maximize coverage no longer lowers the probability that the STA will not experience breaks in service.

Figure 6.5 also shows downlink performance in high-load scenarios. Multi-connectivity maintains its advantage in the low throughput areas below 3 Gbps, ensuring that STAs can communicate with at least one RF chain at all times. However, the simultaneous reception with all 4 RF chains exposes STAs to too much interference and prevents them from achieving as high throughput as is possible with a single-RF. The resulting aggregate network throughput achieved with the single and multi-RF deployments is approximately the same. As discused above, digital precoding or a combining process from the received signals can help address this issue, although it comes at the cost of increased complexity.

**Lessons learned:** In high interference scenarios the spatial diversity of multi-

connectivity can lead to high gains in throughput. In addition, service outage becomes a crucial issue for performance and multi-connectivity can provide continuous coverage and greatly increase mmWave network resilience.

#### 6.3.2.4 Large room scenarios

In addition to our default environment, we also study a 4 times larger room. This is a rather interesting scenario because there are two conflicting effects. On one hand, the larger room allows for more spatial sharing since devices are further apart which can improve network performance. On the other hand, since devices are further apart this can lead to lower Signal-to-Noise Ratio (SNR) between a client and its AP, lowering the achievable MCS and consequently user data rate. We found that in different scenarios the trade-off between these two effects varied, leading to better or worse performance as compared to our baseline, smaller room, scenario. In Figure 6.6 we show the throughput CDF for exemplary scenarios with low, medium and high network load. A very interesting trend was a clear difference between the multi-connectivity implementation and the single-RF deployments.

In the multi-connectivity case, performance in the small and large rooms is rather similar, as seen in Figure 6.6a. When the network load is low, the performance is equal. In some medium and high-load scenarios, the performance is better in the large room due to the increased spatial sharing, as in the medium-load scenario in Fig 6.6a. However, in most cases, the performance slightly degrades in the large room, due to the larger distance between STAs and APs, as can be observed in the high-load scenario. This is consistent with previous results that show that the main challenge in our multi-connectivity system is the simultaneous reception at the different RF chains without any interference suppression at the PHY or channel access optimization at higher layers. Since that interference is not significantly affected by the increased spatial sharing, the larger link distances predominately affect the performance negatively. Lastly, these results demonstrate how multi-connectivity architectures offer consistent service in different environments, especially in terms of reliability and connectivity.

In single-RF deployments (Figure 6.6b and 6.6c), however, we observed both that the effect of the room size was much larger and that performance was consistently improved in the large room. In this case, the reduced interference from neighbouring nodes and the increased spatial sharing allowed for larger performance gains, up to 25 %. This was particularly pronounced when the PAA was chosen to maximize coverage. The issues with interference in the downlink were significantly improved, leading to an equal performance in the two single-RF implementations. In addition, since single-RF implementations require coverage for fewer RF chains, the likelihood that one will have to associate to an AP at a large distance is lower than in the multi-RF scenarios. Therefore the larger distances between nodes affected the performance less.

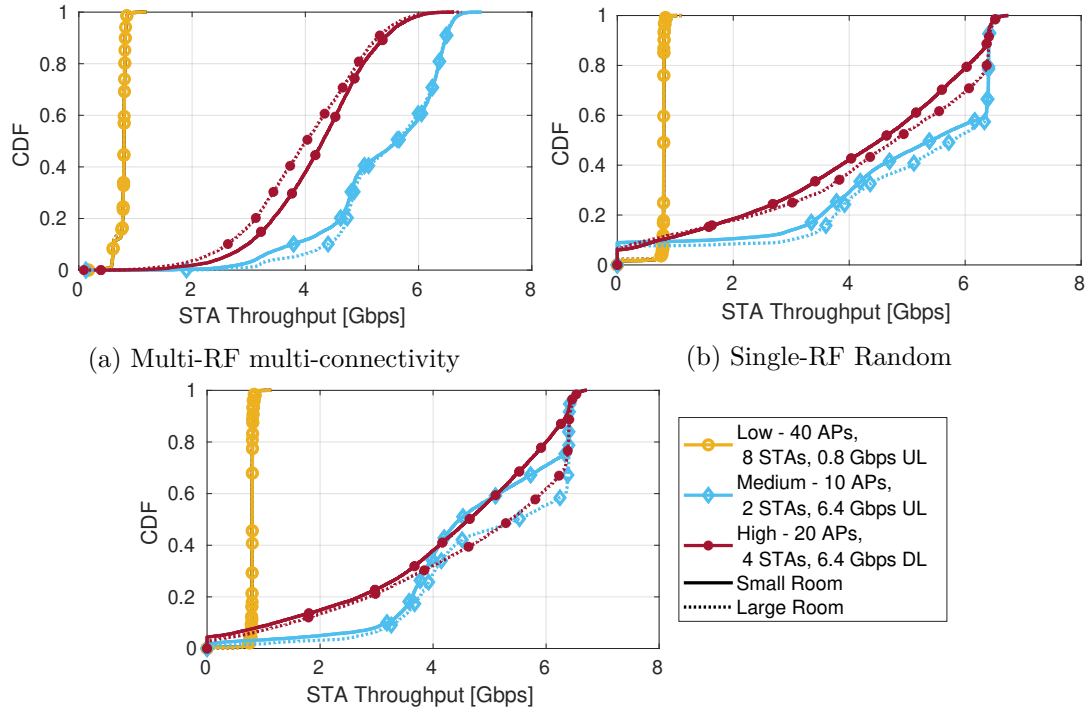


(a) Multi-RF multi-connectivity

(b) Single-RF Random

(c) Single-RF Coverage

Figure 6.6: Throughput CDF comparison of small and large room scenarios.

**Lessons learned:** When considering the deployment size there exists a trade-off between the spatial sharing potential and the average SNR of the links. Multi-connectivity deployments were shown to be less sensitive to environment changes, as the spatial diversity of the links helps ensure consistent and uniform service.

#### 6.3.2.5 Beamforming training improvements

Once we analyzed multi-connectivity in various scenarios and validated the benefits to reliability and resilience, we further tested various options that may improve the performance. This included trying to improve the carrier sensing thresholds, selecting the beampattern not only based on the measured SNR but also on the achieved throughput to maximize spatial sharing, as well as looking at different Rate Adaptation Algorithm (RAA) mechanisms. Ultimately, however, none of these factors had any major impact on performance. This was due to the fact that the main limitation of our multi-connectivity design came from the interference between the multiple RF chains that had to simultaneously receive data.

The one attempted improvement that had a positive impact was an enhancement to the beamforming training. Our previous work has shown that in high-interference scenarios the Group Beamforming principle can lead to wrong beampattern choices if the training is corrupted by interference. We observed the same behaviour in this

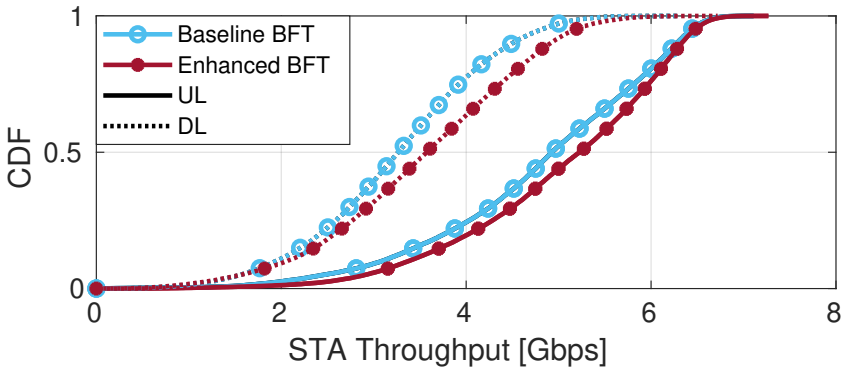


Figure 6.7: Throughput CDF with enhanced BFT, large room, 40 APs, 8 STAs, 6.4 Gbps.

implementation, with the majority of beam selection errors on the STA side. Therefore, we used the modification from Chapter 3 which uses the TRN approach for training both APs and STAs. We found that this was successful at fixing most errors, and furthermore resulted in gains in throughput. The largest gains were in the densest scenario with 8 STAs and 40 APs, where interference caused the most problems with the beamforming training. In addition, the gains were slightly larger in the larger room scenarios, where selecting the correct beampattern is crucial to overcoming the larger distance between devices. Figure 6.7 shows a comparison of multi-connectivity performance in the large room with a per-STA data rate of 6.4 Gbps.

**Lessons learned:** Beamforming is a crucial factor in our multi-connectivity design, as the beampattern configuration affects both the signal strength and the received interference. Improving the analog beamforming, or using hybrid architectures, particularly to cope with intra-RF interference can greatly improve performance.

## 6.4 Related Work

Previous work on multi-connectivity most commonly uses an analytical approach by deriving theoretical performance bounds. In this context, different works demonstrate the benefits of multi-connectivity to mmWave networks reliability [43–46]. [43] focuses on the degree of multi-connectivity to prevent outage, while in [44] a blockage-aware beamformer design is presented. The ability of multi-connectivity to enhance the robustness against link blockage is experimentally verified in [10,50]. Not much work has been done, however, on protocol aspects of multi-connectivity networks, including evaluation scenarios that consider MAC and network layer procedures. In addition, these works all consider a phase-coherent processing approach to joint transmission, requiring tight AP synchronization. Similarly, work on cell-free mmWave networks [114–117] focuses on PHY layer aspects like power control and pilot allocation, as well as the beamforming architecture. The design and performance of hybrid beamforming is studied, as they also use a phase-coherent processing design. Recently, practical aspects like handover [118] and association [119]

have been investigated. Our work differs significantly from all of these works, both in the design which relies on analog beamforming for spatial separation and our evaluation methodology which uses a full-stack multi-connectivity implementation in ns-3.

## 6.5 Conclusions

In this chapter, we present a network design for multi-connectivity mmWave networks. Our main goal was to show how adding extra resources in terms of RF chains can help solve the key challenges of reliability and resilience. Unlike a standard MIMO approach which only aims to increase the overall spectral efficiency, multi-connectivity designs allow to simultaneously maintain links with multiple APs and thus have increased coverage and resilience to blockage. Moreover, mmWave networks are particularly suited to multi-connectivity designs, as they can be implemented with a much lower level of network synchronization. The use of narrow directional beams allows for interference management through spatial separation of signals, rather than phase-coherent signal processing. We validate this in our design, which is based on the mmWave Wi-Fi protocol and relies solely on analog beamforming for interference management, with no PHY or channel access coordination. Unlike most previous work on multi-connectivity, we implement our design in ns-3 for a performance evaluation with a full protocol stack. Our results show that multi-connectivity maintains extremely high reliability even in high-interference scenarios. In fact, clients could always maintain a connection to at least one AP, contrasted with up to 20% of time spent without service in a standard single-RF deployment. In addition, distributing the load across multiple RF chains also allowed for a gain of up to 30% of aggregate network throughput in dense scenarios and high-interference scenarios. Furthermore, multi-connectivity provides a more uniform service to users, with consistent performance in varying conditions. However, issues with interference caused by simultaneous reception at the multiple RF chains sometimes limited the achievable throughput in the downlink, for high-load scenarios. In low-load scenarios, on the other hand, finding multiple serving APs for each client sometimes led to issues with coverage and bad connection. We found that as the network size and traffic grew the benefits of spatially distributing the load were more apparent, however, so did the issues with interference in the downlink. Therefore, in future work, we will look at how interference-aware beamtraining, digital precoding and increased network synchronization can help overcome these issues to fully realize the potential of multi-connectivity for mmWave networks.

# 7 IEEE 802.11bf Directional Multi-Gigabit (DMG) Sensing in ns-3

## 7.1 Introduction

The previous chapters of the thesis looked at different aspects of Millimeter-Wave (mmWave) in a traditional communication architecture. Future wireless networks, however, can expand their functionalities to support features like WiFi sensing which leverages the ubiquitous WiFi signals to perform environmental sensing. The goal is to analyze how the wireless signals reflect from different objects in the environment and deduce the presence and location of objects, as well as patterns and changes in movement. The mmWave band is particularly suited to sensing tasks as the wide signal bandwidth offers high resolution and range, enabling high-precision sensing tasks like localization, gesture recognition and respiration monitoring. Additionally, enabling environmental sensing can enhance network operation, for example by detecting blockages before they occur or performing localization-based beampattern selection.

WiFi sensing can overcome some of the challenges of traditional sensing methods like cameras and LiDAR. For one, it does not require dedicated sensors and instead reuses the existing WiFi infrastructure which lowers the cost and provides ubiquitous deployment. Additionally, WiFi signals can reach Non-Line of Sight (NLOS) areas and are not affected by environmental factors like light or smoke. Finally, it can address privacy concerns as it does not capture images and instead relies on signal analysis. A variety of previous work [51–55, 58] has demonstrated the potential of WiFi sensing.

The most efficient form of WiFi sensing is represented by the Joint Sensing And Communication (JSAC) concept which combines the communication and sensing tasks into a unified framework designed to jointly optimize the performance of both tasks. This requires integrating WiFi sensing in the standardized WiFi operation to design simultaneous communication and sensing procedures. Such standardization support can address many of the current limitations of WiFi sensing applications like interoperability between devices from different vendors and the ability to coordinate multiple devices

when performing sensing tasks. Recognizing the wide interest in the research community and industry, the Task Group IEEE 802.11bf (TGbf) was established in September 2020 to design an amendment to the IEEE 802.11 standard that integrates JSAC into the IEEE 802.11 framework. The resulting IEEE 802.11bf draft standard [59] defines sensing operation in both the sub-7 GHz band and mmWave bands, referred to as the *sensing procedure* and the *DMG sensing procedure*, respectively. The DMG sensing procedure relies on the Medium Access Control (MAC) and physical (PHY) features of the IEEE 802.11ad and IEEE 802.11ay standards, incorporating the specifics of operation in the mmWave band.

Evaluation of the novel IEEE 802.11bf standard is still rather limited. An overview of the foundational elements of IEEE 802.11bf is given in [120]. More practical aspects of integrating sensing in WiFi are presented in [121], focusing only, however, on the sub-7 GHz range. DMG sensing procedures are evaluated in [122] by using a link-level simulator to evaluate the walking speed detection of people in indoor environments.

MAC and network aspects of IEEE 802.11bf systems remain unexplored and may prove crucial to the real-world performance of JSAC. Sensing accuracy requires precise timing of the sensing frames which can be challenging in WiFi networks due to the unreliable, best-effort nature of the standard. Sensing packets may contend with regular data communication for channel resources, underscoring the importance of efficient resource allocation and scheduling strategies. Ensuring that the sensing and communication tasks can coexist without negatively affecting each other is a key challenge for JSAC systems. This requires full-stack system analysis which takes into account medium access and network procedures.

In this chapter, we present a first step in this direction by expanding our ns-3 mmWave module to support DMG sensing operation. This allows us to examine the impact of DMG sensing traffic on communication effectiveness, as well as the reliability of sensing packets in a best-effort WiFi network. The main contributions of this chapter are:

- The first system-level simulation performance evaluation of IEEE 802.11bf DMG Sensing.
- An analysis of DMG sensing varying the sensing type and operational parameters, highlighting the trade-offs between sensing and data performance.
- Identifying configurations that can ensure seamless data communication and efficient sensing.

The rest of this chapter is organized as follows. Section 7.2 offers an overview of DMG sensing operation as defined by IEEE 802.11bf. Section 7.3 presents the ns-3 model extension to incorporate DMG sensing and Section 7.4 presents an evaluation of the performance. Finally, Section 7.5 summarizes our findings.

## 7.2 Overview of the DMG sensing procedure

The goal of the DMG sensing procedure is to obtain measurements to estimate characteristics such as range, velocity and movement of objects through the PHY and MAC features of DMG Stations (STAs). The IEEE 802.11bf amendment outlines several phases to coordinate STAs and perform the necessary measurements. This section provides an overview of the steps to perform DMG sensing.

IEEE 802.11bf supports 6 different types of DMG sensing to address different sensing applications:

- *Monostatic sensing*: Involves only a single STA/Access Point (AP) which both transmits and receives the sensing signals.
- *Bistatic sensing*: Involves two STAs/APs - one transmits the sensing signal and the other receives it.
- *Coordinated monostatic sensing*: An advanced form of monostatic sensing where an AP coordinates multiple STAs to perform synchronized monostatic sensing.
- *Coordinated bistatic sensing*: An advanced form of bistatic sensing where an AP coordinates multiple STAs to perform synchronized bistatic sensing.
- *Multistatic sensing*: Involves multiple STAs and one AP - the AP transmits the sensing signal and multiple STAs receive it.
- *Passive sensing*: Performs sensing measurements by re-using the DMG beacon frames rather than dedicated sensing frames.

IEEE 802.11bf defines two primary roles for STAs - the initiator and the responder of the DMG sensing session. Furthermore, each of them can be the transmitter or receiver, indicating who is transmitting and receiving the sensing signals.

Figure 7.1 shows the different phases of the DMG sensing procedure. The first step is the *DMG Sensing Capabilities exchange*, through which STAs discover which sensing configurations and features are supported by other STAs. To avoid extra overhead, the sensing capabilities can be communicated in the existing beacon and association frames. Additionally, announce or probe frames can be used for an explicit exchange of sensing capabilities. Once the sensing capabilities are exchanged a *sensing session* can be established between the initiator and responder. The following *DMG sensing measurement session* phase is used to establish the sensing operational parameters for the sensing. This includes the duration and timing of sensing packets, as well as determining which STA will act as the transmitter and receiver.

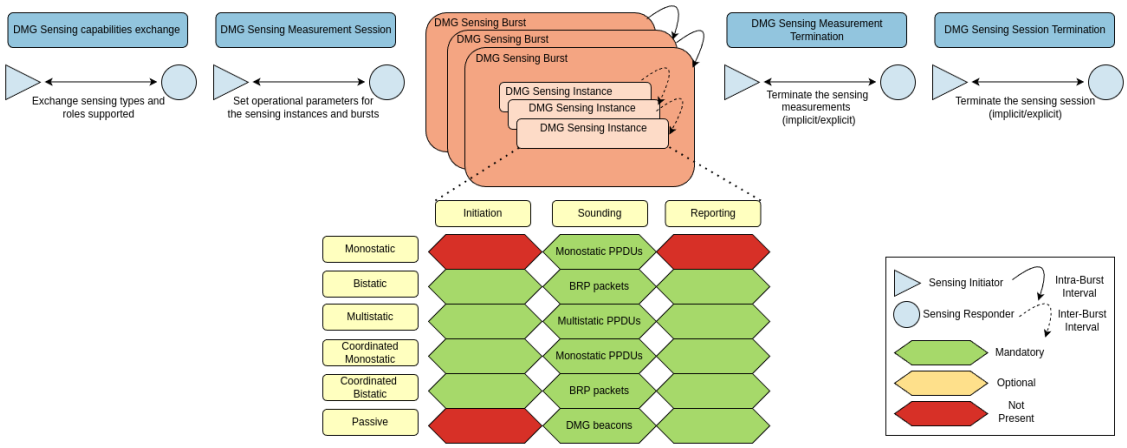


Figure 7.1: Overview of the DMG sensing procedure

Once the STAs agree on the sensing parameters, the sensing measurements can begin. IEEE 802.11bf organizes the measurements in two phases - the *DMG Sensing Burst* and the *DMG Sensing Instance*. Each DMG Sensing Burst is composed of multiple Sensing Instances during which measurements are taken. Taking several measurements in close succession allows for Doppler estimation to determine the velocity of moving targets. The duration between two consecutive Sensing Instances is referred to as the *intra-burst* interval. This parameter determines the range and resolution of velocities which can be detected during the sensing session. The *inter-burst* interval, on the other hand, specifies the time between two consecutive DMG Sensing Bursts. Measurements over several different Bursts enable long-term tracking of targets, e.g., to find the movement trajectory. Additionally, slow-moving targets can require measurements over a longer time duration over several DMG Bursts.

Each DMG Sensing Instance can be composed of an *initiation*, *sounding*, and *reporting* stage, with only the sounding stage being obligatory in all types of sensing sessions. The presence of the other two phases depends on the sensing type. The initiation phase allows synchronisation of the STAs and verifies their availability. Monostatic sensing only involves one STA which makes this phase unnecessary. Passive sensing likewise does not require initiation as it has no dedicated sensing transmissions. Coordinated Monostatic/Bistatic and Multistatic sensing, however, require initiation as they include multiple responders. Finally, in the case of bistatic sensing, the initiation is only necessary if the sensing initiator is the receiver of the sensing frames. In this case, the initiator sends a trigger frame to the responder to request the transmission of a sensing frame.

In the sounding phase, the sensing transmitter sends sensing packets to enable the receiver to perform measurements. Most sensing types use DMG/Enhanced Directional Multi-Gigabit (EDMG) Beam Refinement Protocol (BRP) PLCP Protocol Data Units (PPDUs) to perform the sensing. As discussed in Chapter 2, BRP packets contain an

additional element called a Training (TRN) field which enables fast beamforming training. As they allow transmission and/or reception with multiple beampatterns within the same packet, BRP packets allow to illuminate different parts of the environment with rapid angular scans. Therefore, IEEE 802.11bf re-uses the BRP packets to perform sensing tasks. Additionally, it defines a new EDMG Multistatic Sensing PPDU which builds upon the BRP PPDU to enable multistatic sensing. The key challenge is to ensure that multiple receivers can collect measurements from a single PPDU. Due to the need for directional communication, the PPDU has to be transmitted with a certain beampattern and it is not always guaranteed that all receivers will be able to correctly decode the packet. To overcome this challenge the Multistatic PPDU contains an additional Sync field. The Sync field contains several TRN subfields transmitted with the optimal beampattern for each intended responder and whose timing is known to the responders. This allows receivers to synchronize and receive the TRN field of the packet even if they were not able to decode the packet itself, by receiving the TRN subfields from the Sync field. Similarly, the TRN field contains $P$ TRN subfields for each responder that help with phase and frequency tracking. In the case of monostatic sensing, it is possible to employ different waveforms for sensing, as long as compatibility with existing mmWave WiFi devices is maintained. Finally, passive sensing uses the standard DMG beacon frames to perform sensing and does not require dedicated sensing PPDUs.

After the sensing measurements have been taken, it is sometimes required to feedback the measurement data in the reporting phase. This phase is only required when the sensing initiator is the sensing transmitter and does not perform the sensing measurements. Depending on the sensing type and the capabilities of the STAs there are different types of reports which can be sent:

- Channel State Information (CSI) measurements: Transmitting the raw measured data does not require any processing on the receiver side and provides the most comprehensive data set, but is quite high in overhead and can thus significantly affect the data communication.

- Pre-processed sensing data: To lower the overhead of the reporting phase, IEEE 802.11bf allows to process the sensing measurements at the receiver and then report the processed data back to the initiator, with two main report formats:

  – DMG Sensing Image Report: In this case, Doppler processing is performed to determine the range, velocity and/or angle of arrival/departure. The overhead of the report depends on the dimensionality of the format.

  – DMG Sensing Targets Report: This report requires more sophisticated processing to identify sensing targets and report some of their characteristics to minimize the sensing overhead. The trade-off here is always the complexity of the processing required at the receiver versus the reporting overhead.

Finally, the end of the sensing measurement is signalled by the *DMG Sensing Measurement Termination* which can be done explicitly by a BRP frame exchange or implicitly on the expiration of a timer. Similarly, the *DMG Sensing Session Termination* concludes the whole sensing session and can be implicitly or explicitly signalled.

IEEE 802.11bf allows for sensing packets to be transmitted in a scheduled manner using Service Periods (SPs) or in a contention-based manner in Contention-Based Access Periods (CBAPs). Scheduled access can help ensure the timing of the sensing packets, but might affect the data communication more severely by reserving the channel during the sensing instances.

## 7.3 Implementation

For the initial implementation and evaluation of IEEE 802.11bf we focused on bistatic sensing. We chose bistatic sensing as it is compatible with current Commercial Off-the-Shelf (COTS) devices since it uses existing frame formats (unlike multistatic sensing), but still requires coordination between multiple STAs (unlike monostatic sensing) and allows configuration of the sensing parameters (unlike passive sensing). Similarly, we implemented sensing with contention-based channel access in CBAPs as it is currently the only one supported by COTS devices. By focusing on the IEEE 802.11bf features which can be implemented without significant modifications we not only simplify the simulation process but also in a future step facilitate the validation of the performance from the system-level simulations in actual hardware. In future work, we plan to extend the implementation to cover other sensing types, as well as sensing with scheduled access.

To preserve backwards compatibility with IEEE 802.11ad/ay we expand the existing `DmgWifiMac` class to support the new sensing procedures. Modifications are also needed to the `DmgWifiPhy` and `Codebook` classes to handle the correct transmission and reception of sensing packets including the antenna beampattern switching during the TRN field. We further add a new `DmgSensingCapabilities` class that implements all information elements and packet frames necessary for the sensing procedure signalling.

To study how the priority of sensing traffic affects the performance we added two new packet queues - `DmgSensingHighTxop` and `DmgSensingLowTxop`. `DmgSensingHighTxop` sends the sensing traffic with higher priority than data traffic, while `DmgSensingLowTxop` uses lower priority than data. This priority is with regard to the contention inside STAs between the data and sensing packets they have to transmit. To give sensing traffic priority in contention with other STAs in the network, we can further configure the packet queue to lower the inter-frame spacing when accessing the channel. We refer to these three priority modes as High, Medium and Low.

In case the sensing initiator is the sensing transmitter, there is no need for an initiation phase and the sensing packet is transmitted directly using the standard Carrier

Sense Multiple Access with Collision Avoidance (CSMA/CA) channel access. When the sensing initiator is the sensing receiver an initiation packet is necessary to request the transmission of a sensing packet from the responder. The standard does not specify the details regarding channel access in this case. In our implementation, the initiator uses standard channel access when transmitting the initiation packet, but additionally uses the Network Allocation Vector (NAV) mechanism to reserve the channel for the full duration of the sensing instance. This allows the sensing responder to directly transmit the sensing packet without needing to again contend for channel access. Finally, we also pause the transmission of any other packets at the initiator while it waits for the sensing packet to arrive, to maximize the probability of successful reception.

To enable easy configuration of the sensing parameters our implementation uses configuration text files. The first file contains the timing for each sensing session initiated in the simulation and the name of the configuration file for the session. The configuration file for the session then contains all details regarding the sensing type, duration, TRN field used, channel access, etc. This enables the user to easily modify the sensing parameters without the need to modify the ns-3 model. The parsing of the sensing files and the initiation of the sensing sessions is done by a new `DmgSensingHelper`.

Finally, we enable comprehensive tracing to gain insight into the sensing performance. We record the timing of each sensing packet, along with whether the packet was successfully decoded or not and the receive Signal-to-Noise Ratio (SNR). For the failed packet we further record the reason for the failure (e.g. low SNR, the receiver is busy and drops the packet, the packet time-to-live expires). This allows us to analyse the sensing performance in terms of the percentage of lost packets and the latency of the sensing packets. The existing data traffic traces allow us to analyse the communication performance for a full system evaluation of IEEE 802.bf networks.

## 7.4 Performance Evaluation

### 7.4.1 Simulation scenarios

We evaluate the performance in a set of simulations conducted in an indoor rectangular space with dimensions 30 $m \times 54\ m \times 3\ m$. The results are averaged over 30 different scenarios, varying the location of the devices. Each simulation has a duration of 30 s where both sensing and data traffic is transmitted. To better observe the effects of the sensing traffic we focus on scenarios with only a single AP and study a network with only a single STA and thus no channel contention, as well as 8 STAs for significant contention. We study both downlink and uplink data traffic, with a low data rate of 100 Mbps and a higher load of 500 Mbps.

Sensing traffic is transmitted throughout the whole simulation and sensing is

performed with a single STA. The baseline sensing configuration has 64 packets per burst with an intra-burst interval of 300 $\mu s$ and an inter-burst interval of 34 ms. We perform wide sweeping with 15 transmit and 15 receive beampatterns using a Receive/Transmit training (TRN-R/T) field, resulting in a large sensing packet with a duration of approximately 150 $\mu s$. This sets the sensing overhead to 50% within the intra-burst interval and approximately 30% of the overall simulation. We study both the sensing configuration where the initiator (in our case the AP) is the transmitter and where it is the receiver of the sensing packets. The latter case has a higher overhead as it is necessary to send an initiation packet for each sensing instance. We compare both scenarios with a network with only data traffic to observe how the sensing affects communication. Lastly, we study all three levels of sensing traffic priority described in Section 7.3 for further insight into how the sensing and data traffic interact with one another.

We study communication performance in terms of data throughput, while for the sensing packets, we analyse both the loss percentage and latency. This allows us to gain insight into how accurately the sensing applications will perform.

### 7.4.2 Evaluation Results

We begin the analysis with the simpler scenarios with only a single AP and a single STA in the network. In this case, there is no contention with other devices and we can see the effects of contention within the device between the sensing and data traffic. Additionally, when the sensing and data traffic are transmitted in different directions (i.e the STA sends uplink data to the AP, but the AP is sending sensing packets toward the STA in the downlink) there is a possibility for packet collisions.

In this scenario, we found that data communication is not affected by the additional sensing traffic with either the lower 100 Mbps or the higher 500 Mbps data rate as the mmWave channel has the capacity for Gbps of throughput in the absence of contention. While extremely high data rates will cause throughput degradation, those are not configurations which are suitable for deploying JSAC. Therefore, we focus on lower data rates and in this first scenario study the sensing traffic performance.

First, in Figure 7.2a we show the sensing packet loss, for both the case when the sensing initiator (the AP) is the transmitter and when it is the receiver of sensing packets and with uplink data traffic of 100 Mbps and 500 Mbps. In the case of downlink traffic, the loss is always 0 as the AP coordinates both the data transmissions and the initiation of the sensing instances. Uplink traffic, however, can sometimes cause packet loss as both the AP and the STA try to access the medium at the same time resulting in a collision between the data and sensing packet. The results in Figure 7.2a are presented for the high sensing traffic priority, however, the trends observed apply to all priorities. We can see that the loss percentage increases with the data rate, as the overall traffic in the network

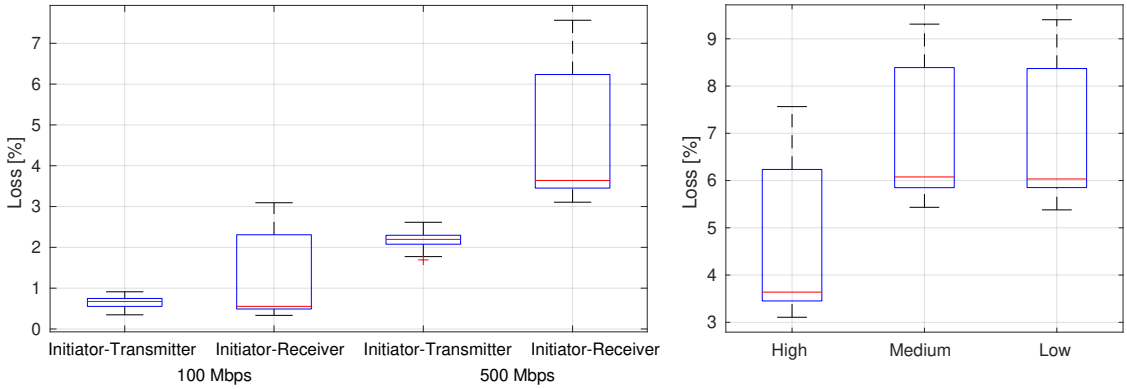


(a) High sensing traffic priority

(b) Initiator-Receiver, 500 Mbps

Figure 7.2: Sensing packet loss, 1 AP, 1 STA, uplink traffic

grows which increases the probability of collision. Additionally, we can see that when the initiator is the sensing receiver there is a higher probability of loss. We found that this is because the sensing overhead is larger in this case, resulting again in higher overall traffic in the network. We also highlight that even a relatively low data rate of 500 Mbps can disturb the sensing with the loss in certain scenarios reaching 7%. This demonstrates how the unreliable nature of WiFi affects sensing accuracy as packet loss is an inherent part of the functioning of WiFi networks due to the uncoordinated channel access. Lastly, in Figure 7.2b we show the effect of the sensing packet priority. The results shown are for the configuration where the initiator is the receiver and the data rate is set to 500 Mbps, however, we saw the same effect in all other tested setups. We can see that high priority reduces the loss from a median of 6% to 3.6%, representing a 66% reduction compared to medium and low priority. As explained in Section 7.3, medium priority only affects the internal contention within a device, and thus it does not affect the loss percentage which comes from contention with other devices. High priority, however, affects the contention with other devices and was thus able to reduce the sensing packet loss.

Lastly, we look at the sensing packet latency due to transmission, channel access and queuing delays as it can profoundly affect the accuracy and resolution of the sensing procedures. We found that the baseline latency when no other devices are contending for the channel was 43 $\mu s$ when the initiator was the sensing packet transmitter and 75 $\mu s$ when the initiator was the sensing packet receiver. The latter is significantly higher due to the overhead of the initiation packet. When using the high sensing traffic priority this latency was reduced by 5 $\mu s$, due to the lower inter-frame spacing for channel access. Figure 7.3a shows the Cumulative Distribution Function (CDF) of the sensing packet latency for the lower 100 Mbps data rate, showing the effect of the traffic and sensing packet direction. For easier comparison, we show only the latency up to 1 ms, however, in Figure 7.3b when the initiator is the receiver and traffic is in the uplink the latency

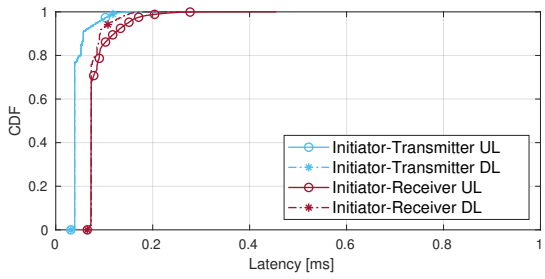


(a) 100 Mbps, high priority

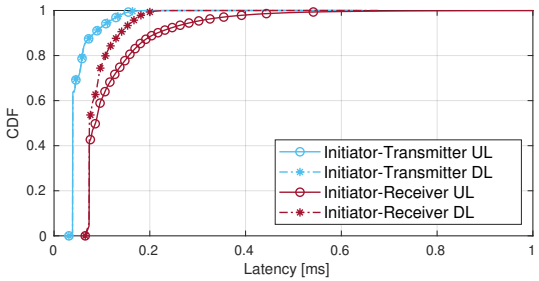


(b) 500 Mbps, high priority

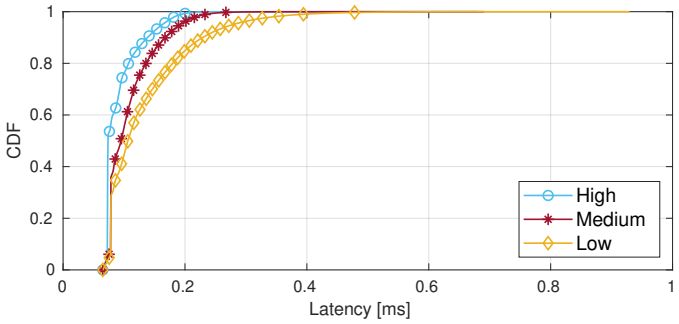


(c) 500 Mbps downlink traffic, initiator as receiver

Figure 7.3: Sensing packet latency with 1 AP, 1 STA

reaches up to 3.7 ms. We can see that the majority of the sensing packets have the baseline latency and do not experience any delay due to queuing or channel contention. The percentage of such packets is higher when the initiator is the sensing packet transmitter at approximately 80% compared to approximately 70% when the initiator is the sensing packet receiver due to the higher sensing overhead. The rest of the packets suffer minor additional delays below 1 ms, however, we can see that overall the sensing is not affected by the latency of the packets. The higher data rate of 500 Mbps shown in Figure 7.3b shows similar trends with most packets reaching the receiver with delays below 1 ms. The exception is when the initiator is the sensing packet receiver and traffic is transmitted in the uplink direction. In this case, some packets experience higher delays of up to 3.7 ms, which will affect the accuracy and resolution of the sensing. These delays are caused by queuing delays due to the larger network traffic. Additionally, there is contention for the channel between the AP which is transmitting the initiation sensing packets and the STA which is transmitting uplink traffic. This can lead to larger contention windows for the CSMA/CA protocol, as well as packet collisions which lead to re-transmissions and longer queues. While still functioning well most of the time, we begin to see how even with a single STA in the network there can be minor disruptions to the sensing operation due to the unreliable nature of WiFi.

Finally, when looking at the effect of the sensing traffic priority we found that in the 100 Mbps scenarios, high priority yielded a lower latency as devices could access the

channel faster, while medium and low priority were equivalent due to the relatively low data rates. In the scenarios with the higher data rate of 500 Mbps, however, we observed that when the traffic was in the downlink direction all three priorities had a different performance, as seen in Figure 7.3c. In these scenarios both the sensing and data traffic is initiated from the same device and therefore the medium priority lowers the latency as it allows for sensing packets to be transmitted before data packets when both queues have packets to send. In the 100 Mbps scenarios, this does not happen as the low data rate means that the traffic queue is empty the vast majority of the time. In the uplink cases, on the other hand, the sensing and data traffic are initiated from different devices, which means that the medium priority which only affects the internal contention within the device, similarly does affect the performance.

In the next step of our analysis, we want to see how channel contention between multiple STAs will affect the performance by increasing the number of STAs associated with the AP to 8. Sensing is still performed only with a single STA, however, all 8 STA exchange data traffic with the AP, with the same data rates of 100 Mbps and 500 Mbps. In this scenario, not only is the overall traffic in the network higher but additionally, STA contend for channel access in the uplink, which can lead to significant problems with collisions and MAC efficiency, as we have explored in Chapters 2, 3 and 4 of this thesis.

The first major difference was that we found that the sensing traffic affected the data performance in scenarios with uplink traffic. While downlink scenarios where the AP coordinates all transmissions allowed for data communication to be uninterrupted, in the uplink we saw a degradation of the data transmissions. Figure 7.4 shows this in terms of the aggregate network throughput for a data rate of 100 Mbps (Figure 7.4a) and 500 Mbps (Figure 7.4b). We can see that for the lower data rate only the sensing configuration where the initiator is the receiver is affected, with a relatively mild reduction of up to 50 Mbps in aggregate network throughput. The higher data rate, however, shows much more drastic effects. In this case, both sensing configurations cause a degradation in the performance of 20 % and 40 % for the initiator as sensing transmitter and receiver case respectively. Without sensing, most STAs achieve the set data rate with a median per-STA throughput of 492 Mbps. Adding sensing traffic, however, pushes the network into overload leading to collisions and packet losses. In terms of the priority of the sensing traffic, we found that the high priority mode caused a slight degradation in the data throughput. This is to be expected, as giving priority to the sensing traffic prevents data packets from being transmitted at the same time. While the effect was minor in most scenarios, in the configuration where the initiator was the sensing packet receiver with 500 Mbps uplink traffic we saw a drop of approximately 200 Mbps in aggregate network throughput. This shows some of the trade-offs in the configuration of JSAC networks, as enhancing the sensing traffic performance can lead to problems for the communication part of the network.

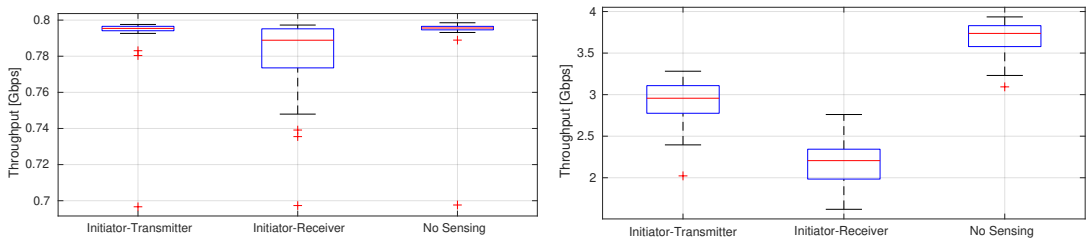


(a) 100 Mbps

(b) 500 Mbps

Figure 7.4: Uplink aggregate data traffic throughput with 1 AP, 8 STAs, high priority

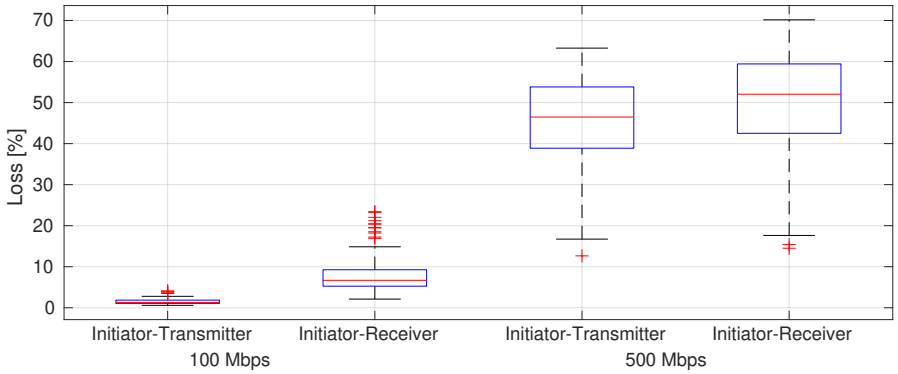


Figure 7.5: Sensing packet loss, 1 AP, 8 STA, uplink traffic, high sensing traffic priority

We observe very similar trends in terms of the sensing performance, both in the latency and loss of sensing packets.

We can see in Figure 7.5 which shows the sensing packet loss for the uplink scenarios with both data rates and sensing configurations, that the loss is significantly higher than in the scenarios with a single STA. This is not only due to the increase in overall network traffic, but also as the channel contention is much higher leading to more collisions. We also see in the 100 Mbps case a large difference between the two network configurations. While the loss when the initiator is the sensing transmitter is always lower than 5 %, it can go up to 20 % when it is the receiver. This large gap is partially caused by the higher overhead of the second sensing configuration. Additionally, in this case, the AP virtually reserves the medium through the use of NAV in the initiation packet, thus protecting the upcoming sensing packet from the STA. Virtual channel reservation requires the receiver to decode the MAC header of the packet to read the NAV. In contrast, the physical channel reservation when transmitting a packet can function even if the receiver is not able to decode the packet, solely by detecting the transmit power on the medium. Therefore, the virtual channel reservation used in the second sensing configuration has more of a potential to fail and result in additional collisions. In the 500 Mbps data rate case, the difference between the two configurations is much smaller, as both setups have extremely high loss percentages of approximately 50 % on average and up to 70 % in

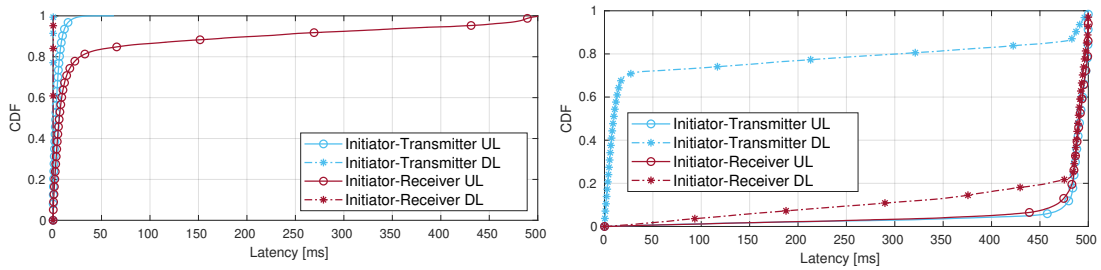


(a) 100 Mbps (b) 500 Mbps

Figure 7.6: Sensing packet latency with 1 AP, 8 STAs, high priority

some cases. We found that here the majority of packets were not caused by collisions in the wireless channel. Instead, packets were dropped on the transmitter side as their time-to-live of 500 ms expired due to extremely long queuing delays. In fact, this problem even occurred in some downlink scenarios, where the losses to due to channel collisions are 0 %, as the AP coordinates all transmissions.

Lastly, in Figure 7.6 we show the CDF of the sensing packet latency which matches the conclusions drawn from the previous performance metrics. In the 100 Mbps case (Figure 7.6a), the sensing application works mostly well and packets arrive without significant delays. The exception is the sensing configuration where the initiator is the sensing receiver with uplink data traffic, where we see that approximately 20 % of packets experience very high delays of up to hundreds of ms. This reflects moments of congestion in the network where devices can not access the channel for long periods. We found that the delays were exacerbated by the large aggregation enabled by IEEE 802.11ay, which led to long data packets that occupied the medium for up to 2 ms. We see in Figure 7.6b that increasing the data traffic leads to a drastic increase in latency. Sensing with the initiator as the transmitter of sensing packets and downlink traffic was the only configuration where most sensing packets experienced a reasonable latency of several ms. In all other tested configurations, the majority of packets had a latency of approximately 500 ms, revealing a complete breakdown of the sensing procedure.

### 7.4.3 Discussion

Our work represents the first system-level analysis of DMG JSAC with the novel IEEE 802.11bf protocol. We found many interesting insights into the practical performance of such a system when it uses the contention-based channel access of current WiFi. First, JSAC is only viable in lightly loaded communication networks, as the sensing performance drastically degrades when the medium becomes saturated, leading to extreme delays and packet loss. Running a sensing application in such scenarios is not viable and the sensing traffic can additionally harm the data performance. Second, similar to the conclusions drawn in other parts of this thesis, dense networks with a lot of channel

contention are quite challenging for applications like sensing which require reliability in packet transmission. In such scenarios, even with a relatively low load of the network, a certain percentage of sensing packets will be negatively affected by packet collisions and large queuing delays. Depending on the requirements of the sensing application, such degradation might not be acceptable and alternative implementations such as scheduled access should be considered. In lightly loaded networks, however, we found that communication and sensing applications could coexist successfully without harming each other. Our results demonstrate both the potential of JSAC mmWave WiFi networks and the need for further research in enhancing their implementation.

When analysing the sensing parameters, we found that the sensing configuration where the initiator (in our case the AP) is the receiver of sensing packets was much more vulnerable to performance degradation. The reason why was a combination of the higher overhead due to the need for an initiation packet and the virtual medium reservation using NAV. This led to a higher probability of collision with data packets which further increased the network load due to the data re-transmissions and need for larger aggregation and longer packets. Finally, the priority of sensing packets also affected the performance. Giving sensing packets priority, particularly when contending with other devices, was shown to reduce the latency and loss, in some cases by 50 %. However, in cases of network overload, it caused a reduction in the achievable data throughput. Therefore, such strategies should be carefully considered, taking into account the requirements of both the sensing and data traffic to minimize the negative effects.

## 7.5 Conclusion

In this chapter, we presented an initial study of a JSAC DMG system based on the novel IEEE 802.11bf protocol. Our work is the first to study the full system performance of such a network by implementing sensing procedures in the mmWave WiFi ns-3 module. As an initial step, we focused on sensing procedures that are easily supported by current devices, implementing bistatic sensing with contention-based channel access. We found that sensing should be considered in low-loaded networks to ensure that the sensing packets have high reliability and low latency. In such scenarios, both types of traffic can coexist without negative effects. Highly loaded networks, however, resulted in both degradation of the communication performance and a complete breakdown of the sensing application due to extreme delays and packet losses. It is important to note, that in cases of high contention, even data rates far below the theoretical channel capacity can cause such poor performance due to the challenges of contention-based channel access in mmWave WiFi. Finally, we saw large differences between different sensing configurations, with the sensing traffic priority and particularly the roles of the initiator and responder affecting the sensing and also data performance. In future work, we plan to expand our

implementation and provide a comprehensive analysis of the full IEEE 802.11bf protocol. This will allow us to determine the modes of operation which allow for JSAC operation and design resource allocation strategies which optimize the performance.

# 8 Conclusions

Millimeter-Wave (mmWave) technology is foreseen as a key enabler for emerging high-performance applications, utilizing the large bandwidth available in the high-frequency range to provide multi-Gbps, low-latency wireless connectivity. The past decade has seen tremendous progress in enabling communication at mmWave frequencies, using a variety of physical (PHY) and Medium Access Control (MAC) techniques to cope with the challenging propagation environment. This has led to wide standardization support, the development of Commercial Off-the-Shelf (COTS) devices and the first real-world deployments. Moving from predominantly single-link research architectures to practical multi-link networks, however, has revealed the unique challenges which mmWave introduces to networking protocols. The goal of this thesis is to provide insight into the functioning of dense mmWave networks, finding the causes of inefficiency and providing practical solutions to enhance their operation.

The majority of the work in this thesis was supported by the simulation tool presented in Chapter 2. Research into dense mmWave networks, especially from a protocol perspective, can be challenging due to the lack of research tools that enable detailed insight and fast implementation of novel solutions. Our implementation of the IEEE 802.11ay protocol in the network simulator ns-3 allowed us to study state-of-the-art mmWave WiFi networks with a detailed MAC and PHY implementation and a high-fidelity channel model. We have used this model throughout the thesis to both investigate performance trends and behaviours in current networks and implement novel solutions. Our model is continuously updated with new features and is freely available to the research community.

The first part of the thesis focused on current mmWave networks, looking at different operational aspects and performance trends. Chapter 2 contains initial insights into the scalability of IEEE 802.11ay networks which demonstrate how throughput and latency degrade in dense networks, particularly when Stations (STAs) contend for channel access in the uplink. In Chapters 3 and 4 we then analysed in detail two key problems for dense mmWave networks and designed solutions to address them.

Chapter 3 focuses on Beamforming Training (BFT) as a crucial factor in mmWave

communication and looks at how accuracy and overhead scale in dense networks. For this purpose, we evaluated a novel solution from IEEE 802.11ay called Group Beamforming which can address scalability and density as it allows for all STAs associated with the same Access Point (AP) to be simultaneously trained. This approach keeps the BFT overhead constant as the network size grows, making it much more scalable than legacy BFT protocols which rely on per-pair training. We found that by reducing BFT overhead, Group Beamforming could lead to significant throughput gains in dense networks. However, the accuracy of the training suffered in high-interference environments leading to wrong beampattern selection and drops in the Signal-to-Interference-plus-Noise Ratio (SINR). To overcome these challenges we designed enhanced versions of Group Beamforming which increased the accuracy by up to 35% and enabled throughput gains of up to 0.7 Gbps. Our enhanced Group Beamforming was able to outperform legacy BFT in all tested scenarios, demonstrating how efficient and scalable BFT protocols are an important part of mmWave operation in large deployments.

Chapter 4 addresses two other sources of inefficiency in mmWave WiFi. First, the use of omnidirectional receive beampatterns which leads to lower antenna gains and higher received interference, thus affecting both the single-link and network performance. Second, the overhearing of unwanted packets which limits the spatial re-use in the network by keeping devices busy decoding packets for which they are not the recipient. Our proposed solution, SIGNaling in the PHY Preamble (SIGNiPHY), represents a PHY signalling mechanism that embeds the user identifier (ID) in the PHY preamble of the packets. Identifying the transmitter early in the packet reception allows for the correct receive beampattern to be used to receive the data payload and to filter any unwanted packets. SIGNiPHY introduces no overhead and is backwards-compatible and interoperable with legacy devices. It has a practical, lightweight design to ensure it can be easily integrated into commercial devices. SIGNiPHY was evaluated in ns-3, as well as a hardware testbed and was shown to enhance packet reception for improved resilience to interference and noise, leading to improved spatial re-use and significant gains in throughput, latency and fairness.

The rest of the thesis looked beyond the capabilities of current devices to future possibilities for mmWave technology. In particular, Chapters 5 and 6 looked at two different architectures for mmWave with advanced devices with multiple Radio Frequency (RF) chains. Chapter 5 presented our Multiple-Input and Multiple-Output (MIMO) implementation as standardized by IEEE 802.11ay, with a focus on the MIMO BFT protocols. We demonstrate that a fully analog architecture is a viable option for mmWave MIMO, although the BFT overhead is significant due to the lengthy simultaneous training of multiple transmit and receive beampatterns. Our BFT algorithm was able to find up to 4 independent streams for a multifold increase in throughput. The challenge, however, is that the achieved SINR dropped significantly at larger distances, particularly when

the number of streams was larger than 2 due to the sparsity of the mmWave channel. Additionally, mmWave MIMO was shown to be very sensitive to movement and even minor changes in position could lead to link misalignment and breakage, thus contributing to the key issues of robustness and reliability in mmWave. Therefore, in Chapter 6 we presented an alternative design based on the multi-connectivity paradigm which utilizes the multiple RF chains to enhance the resilience of mmWave networks. This is done by enabling STA to simultaneously connect to multiple APs which increases the resilience to any single link failing. In contrast to most previous work which relies on coherent signal processing across the network with a high implementation complexity, we instead proposed a purely analog design that uses only spatial separation of the signals for interference management. We found that a multi-connectivity network with local analog beamforming could successfully increase the reliability of the network and provide uninterrupted service. Additionally, spatially distributing the load across multiple APs was found to have throughput benefits too, as the spatial diversity improved the resilience to interference for gains of up to 30%.

Finally, in Chapter 7 we implement the novel IEEE 802.11bf protocol which aims to standardize Joint Sensing And Communication (JSAC) for WiFi networks. JSAC represents an evolution of wireless networks that enables environmental sensing through reflections of communication signals in a single system. mmWave networks are particularly interesting for sensing applications as the large signal bandwidth enables high-precision, high-resolution target identification. We integrated sensing procedures into IEEE 802.11ay operation to look at how the communication and sensing traffic affect each other. We found that the unreliable nature of WiFi can affect the timing and successful delivery of sensing packets which degrades the detection accuracy. When the network is highly loaded, sensing is not possible due to long queuing delays and the additional sensing traffic harms data communication in the network. However, when operating in lightly loaded conditions, both communication and sensing procedures can successfully operate in parallel even when relying on the contention-based channel access. As the last contribution of this thesis, Chapter 7 expands on the possible applications of mmWave technology, moving from an enabler of emerging applications to directly implementing them in a novel JSAC architecture.

This thesis investigates how mmWave technology affects various aspects of networking protocols by studying different proposed deployment architectures. We design solutions that improve network operation, however, many open issues remain. Localization-based BFT mechanisms can enhance the performance and provide fast beampattern adaptation in cases of mobility or blockage. Optimizing association and handover mechanisms can enhance spatial reuse and reduce packet collisions. Hybrid beamforming algorithms can address challenges with interference and channel sparsity in multi-RF networks. Finally, JSAC opens new possibilities for mmWave networks with many open questions regarding the sensing implementations and configurations which optimize the joint performance.

We hope that the insights from this thesis can be used to design new, innovative solutions that bring mmWave networks closer to widespread deployment, realising the potential for next-generation, ultra-high performance wireless connectivity.

# References


[1] H. Assasa, N. Grosheva, T. Ropitault, S. Blandino, N. Golmie, and J. Widmer, "Implementation and evaluation of a WLAN IEEE 802.11ay model in network simulator ns-3," *Workshop on ns-3*, pp. 9–16, 2021.

[2] N. Grosheva, H. Assasa, T. Ropitault, P. J. Mateo, J. Widmer, and N. Golmie, "A Comprehensive Analysis and Performance Enhancements for the IEEE 802.11ay Group Beamforming Protocol," *2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 194–203, 2022.

[3] N. Grosheva, S. P. Deram, J. O. Lacruz, and J. Widmer, "SIGNiPHY: Reconciling random access with directional reception for efficient mmWave WLANs," *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, 2023. [Online]. Available: https://doi.org/10.1145/3581791.3596860

[4] N. Grosheva, R. Hersyandika, J. Widmer, and S. Pollin, "In-Band Multi-Connectivity with Local Beamtraining for Improving mmWave Network Resilience," *Proceedings of The 26th International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM'23)*, 2023. [Online]. Available: https://doi.org/10.1145/3616388.3617529

[5] D. Naboulsi, M. Fiore, S. Ribot, and R. Stanica, "Large-Scale Mobile Traffic Analysis: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 18, pp. 124–161, 2016.

[6] FG-NET-2030, "Network 2030 - a blueprint of technology, applications and market drivers towards the year 2030 and beyond," *ITU*, 2019.

[7] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi," *IEEE Communications Magazine*, pp. 186–192, 2017.

[8] T. Nitsche, C. de M. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. Widmer, "Ieee 802.11ad: directional 60 ghz communication for multi-gigabit-per-second wi-fi [invited paper]," *IEEE Communications Magazine*, vol. 52, pp. 132–141, 2014. [Online]. Available: https://api.semanticscholar.org/CorpusID:206455123

[9] D. García, J. O. Lacruz, P. J. Mateo, and J. Widmer, "POLAR: Passive object localization with IEEE 802.11ad using phased antenna arrays," *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, pp. 1838–1847, 2020.

[10] T. T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[11] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling," in *Measurement and Modeling of Computer Systems*, 2015.

[12] M. Giordani, M. Mezzavilla, and M. Zorzi, "Initial Access in 5G mmWave Cellular Networks," *IEEE Communications Magazine*, vol. 54, pp. 40–47, 2016.

[13] M. L. Attiah, A. A. M. Isa, Z. Zakaria, M. K. Abdulhameed, M. K. Mohsen, and I. A. Ali, "A survey of mmWave user association mechanisms and spectrum sharing approaches: an overview, open issues and challenges, future research trends," *Wireless Networks*, vol. 26, pp. 2487–2514, 2020.

[14] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, "Improved Handover Through Dual Connectivity in 5G mmWave Mobile Networks," *IEEE Journal on Selected Areas in Communications*, vol. 35, pp. 2069–2084, 2016.

[15] G. Bielsa, A. Loch, and J. Widmer, "Optimizing mmWave Spatial Reuse: Signal-To-Interference Aware Beamtraining," *2019 IEEE 20th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, pp. 1–6, 2019.

[16] 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," *IEEE SECON*, pp. 1–9, 2018.

[17] 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 *2018 IEEE 19th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*, Chania, Greece, 2018.

[18] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Wireless Networks," *ACM MobiCom*, pp. 28–41, 2017.

[19] H. Assasa, A. Loch, and J. Widmer, "Packet mass transit: Improving frame aggregation in 60 GHz networks," *2016 IEEE 17th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 1–7, 2016.

[20] Y. Ren, W. Yang, X. Zhou, H. Chen, and B. Liu, "A survey on TCP over mmWave," *Comput. Commun.*, vol. 171, pp. 80–88, 2021.

[21] IEEE 802.11 working group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band." *IEEE Standard 802.11ad*, 2012.

[22] "TP-LINK Talon AD7200," [Online]. Available: https://www.tp-link.com/us/home-networking/wifi-router/ad7200/.

[23] "wAP 60G," [Online]. Available: https://mikrotik.com/products/group/60-ghz-products.

[24] "Netgear Nighthawk X10," [Online]. Available: https://www.netgear.com/home/wifi/routers/ad7200-fastest-router/ .

[25] "Dell Wireless Docking Station WLD15," [Online]. Available: https://www.dell.com/support/kbdoc/en-us/000129322/the-dell-wireless-docking-station-wld15-wigig-capable?lwp=rt .

[26] "ASUS ROG Phone II ," [Online]. Available: https://rog.asus.com/phones/rog-phone-ii-model/ .

[27] IEEE 802.11 working group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications-Amendment 2: Enhanced Throughput for Operation in License-Exempt Bands Above 45 GHz," *IEEE Standard 802.11ay*, 2021.

[28] "MikroTik Cube 60Pro ac," [Online]. Available: https://mikrotik.com/product/cube_60pro_ac.

[29] 3rd Generation Partnership Project, "3GPP Release-15," *http://www.3gpp.org/release-15/*, 2016.

[30] M. Mezzavilla, M. Zhang, M. Polese, R. D. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-End Simulation of 5G mmWave Networks," *IEEE Communications Surveys & Tutorials*, vol. 20, pp. 2237–2263, 2017.

[31] A. Narayanan, M. I. Rochman, A. Hassan, B. S. Firmansyah, V. Sathya, M. Ghosh, F. Qian, and Z.-L. Zhang, "A Comparative Measurement Study of Commercial 5G mmWave Deployments," *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, pp. 800–809, 2022.

[32] S. Aggarwal, M. Ghoshal, P. Banerjee, D. Koutsonikolas, and J. Widmer, "802.11ad in Smartphones: Energy Efficiency, Spatial Reuse, and Impact on Applications," *IEEE INFOCOM*, pp. 1–10, 2021.

[33] T. Wei and X. Zhang, "Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support," *ACM MobiCom*, pp. 42–55, 2017.

[34] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks," *ACM MobiCom*, pp. 257–271, 2018.

[35] 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," *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, 2015.

[36] J. O. Lacruz, R. R. Ortiz, and J. Widmer, "A Real-Time Experimentation Platform for Sub-6 GHz and Millimeter-Wave MIMO Systems," in *ACM MobiSys*, 2021, pp. 427–439.

[37] H. Assasa and J. Widmer, "Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3," in *Proceedings of the 2016 Workshop on ns-3*, Seattle, WA, USA, 2016.

[38] ——, "Extending the IEEE 802.11ad Model: Scheduled Access, Spatial Reuse, Clustering, and Relaying," in *Proceedings of the Workshop on ns-3*, Porto, Portugal, 2017.

[39] 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*, Florence, Italy, 2019.

[40] C. R. C. M. Da Silva, J. Kosloff, C. Chen, A. Lomayev, and C. Cordeiro, "Beamforming Training for IEEE 802.11 ay Millimeter Wave Systems," in *2018 Information Theory and Applications Workshop (ITA)*. San Diego, CA: IEEE, Feb. 2018, pp. 1–9.

[41] 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," *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018.

[42] R. W. Heath, N. G. Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, "An Overview of Signal Processing Techniques for Millimeter Wave MIMO Systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, pp. 436–453, 2015.

[43] M. Gapeyenko, V. Petrov, D. Moltchanov, M. R. Akdeniz, S. D. Andreev, N. Himayat, and Y. Koucheryavy, "On the degree of multi-connectivity in 5g millimeter-wave cellular urban deployments," *IEEE Transactions on Vehicular Technology*, vol. 68, pp. 1973–1978, 2019.

[44] D. Kumar, J. Kaleva, and A. Tölli, "Blockage-aware reliable mmwave access via coordinated multi-point connectivity," *IEEE Transactions on Wireless Communications*, vol. 20, no. 7, pp. 4238–4252, 2021.

[45] D. Maamari, N. Devroye, and D. Tuninetti, "Coverage in mmwave cellular networks with base station co-operation," *IEEE Transactions on Wireless Communications*, vol. 15, no. 4, pp. 2981–2994, 2016.

[46] H. Wang, K.-W. Huang, and T. A. Tsiftsis, "Base station cooperation in millimeter wave cellular networks: Performance enhancement of cell-edge users," *IEEE Transactions on Communications*, vol. 66, pp. 5124–5139, 2018.

[47] H. Q. Ngo, A. E. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-Free Massive MIMO: Uniformly great service for everyone," *2015 IEEE 16th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 201–205, 2015.

[48] ——, "Cell-Free Massive MIMO Versus Small Cells," *IEEE Transactions on Wireless Communications*, vol. 16, pp. 1834–1850, 2016.

[49] G. R. MacCartney and T. S. Rappaport, "Millimeter-wave base station diversity for 5g coordinated multipoint (comp) applications," *IEEE Transactions on Wireless Communications*, vol. 18, no. 7, pp. 3395–3410, 2019.

[50] R. Hersyandika, J. Sanchez, Y. Miao, S. Pollin, and F. Tufvesson, "Measurement-based blockage and intra-cluster interference analysis in mmwave multi-point connectivity networks," in *EmergingWireless '22*. New York, NY, USA: Association for Computing Machinery, 2022, pp. 12–17.

[51] H. Abdelnasser, K. Harras, and M. Youssef, "A Ubiquitous WiFi-Based Fine-Grained Gesture Recognition System," *IEEE Transactions on Mobile Computing*, vol. 18, no. 11, pp. 2474–2487, 2019.

[52] P. Falcone, F. Colone, A. Macera, and P. Lombardo, "Localization and tracking of moving targets with WiFi-based passive radar," in *2012 IEEE Radar Conference*, 2012, pp. 0705–0709.

[53] C. Wu, F. Zhang, B. Wang, and K. J. Ray Liu, "mmTrack: Passive Multi-Person Localization Using Commodity Millimeter Wave Radio," in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020, pp. 2400–2409.

[54] S. Mosleh, J. B. Coder, C. G. Scully, K. Forsyth, and M. O. A. Kalaa, "Monitoring Respiratory Motion with Wi-Fi CSI:Characterizing performance and the BreatheSmart Algorithm," *IEEE Access*, pp. 1–1, 2022.

[55] F. Wang, F. Zhang, C. Wu, B. Wang, and K. J. R. Liu, "ViMo: Multiperson Vital Sign Monitoring Using Commodity Millimeter-Wave Radio," *IEEE Internet of Things Journal*, vol. 8, no. 3, pp. 1294–1307, 2021.

[56] C. B. Barneto, E. Rastorgueva-Foi, M. F. Keskin, T. Riihonen, M. Turunen, J. Talvitie, H. Wymeersch, and M. Valkama, "Millimeter-Wave Mobile Sensing and Environment Mapping: Models, Algorithms and Validation," *IEEE Transactions on Vehicular Technology*, vol. 71, pp. 3900–3916, 2021.

[57] F. Wang, S. Zhou, S. Panev, J. Han, and D. Huang, "Person-in-WiFi: Fine-Grained Person Perception Using WiFi," *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5451–5460, 2019.

[58] L. Storrer, H. C. Yildirim, M. Crauwels, E. I. P. Copa, S. Pollin, J. Louveaux, P. De Doncker, and F. Horlin, "Indoor Tracking of Multiple Individuals With an 802.11ax Wi-Fi-Based Multi-Antenna Passive Radar," *IEEE Sensors Journal*, vol. 21, no. 18, pp. 20 462–20 474, 2021.

[59] IEEE 802.11 working group, "IEEE P802.11bf/D2.0 Draft 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 Amendment 2: Enhancements for Wireless LAN Sensing," *IEEE Standard 802.11bf*, 2023.

[60] H. Assasa and N. Grosheva. (2021) A Collection of Open-source Tools to Simulate IEEE 802.11ad/ay WLAN Networks in Network Simulator ns-3. [Online]. Available: https://github.com/wigig-tools/ns3-802.11ad

[61] A. Bodi, S. Blandino, N. Varshney, J. Zhang, T. Ropitault, M. Lecci, P. Testolina, J. Wang, C. Lai, and C. Gentile. (2021) The NIST Q-D Channel Realization Software. [Online]. Available: https://github.com/wigig-tools/qd-realization

[62] N. Varshney, J. Zhang, J. Wang, A. Bodi, and N. Golmie, "Link-Level Abstraction of IEEE 802.11ay based on Quasi-Deterministic Channel Model from Measurements," in *2020 IEEE 92nd Vehicle Technology Conference (VTC 2020-Fall)*, Victoria, B.C, Canada, 2020.

[63] A. Maltsev, A. Pudeyev, I. Bolotin, and O. Bolkhovskaya, "Asymmetric Links Beamforming for mmWave Overlay in LTE-Based Heterogeneous Network," in *2017 IEEE 86th Vehicular Technology Conference (VTC-Fall)*, 2017, pp. 1–5.

[64] Y. Kim, S. Lee, and T. Ropitault, "STS Adaptation for Beamforming Training of Asymmetric Links in IEEE 802.11ay-based Dense Networks," in *IEEE Vehicular Technology Conference (VTC2020-Spring)*, 2020, pp. 1–6.

[65] M. Kim, T. Ropitault, S. Lee, and N. Golmie, "Efficient MU-MIMO Beamforming Protocol for IEEE 802.11ay WLANs," *IEEE Communications Letters*, vol. 23, no. 1, pp. 144–147, 2019.

[66] M. S. Kim, T. Ropitault, S. Lee, N. Golmie, H. Assasa, and J. Widmer, "A Link Quality Estimation-Based Beamforming Training Protocol for IEEE 802.11ay MU-MIMO Communications," *IEEE Transactions on Communications*, vol. 69, no. 1, pp. 634–648, 2021.

[67] P. Zhou, K. Cheng, X. Han, X. Fang, Y. Fang, R. He, Y. Long, and Y. Liu, "IEEE 802.11ay based mmWave WLANs: Design Challenges and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1654–1681, 2018, arXiv: 1803.07808.

[68] C. Chen, O. Kedem, C. R. C. M. da Silva, and C. Cordeiro, "Millemeter-Wave Fixed Wireless Access Using IEEE 802.11ay," *arXiv:1907.00082 [cs]*, Jun. 2019, arXiv: 1907.00082.

[69] Y. Kim, S. Lee, and T. Ropitault, "Adaptive Scheduling for Asymmetric Beamforming Training in IEEE 802.11ay-based Environments," *IEEE Wireless Communications and Networking Conference (WCNC)*, 2019.

[70] P. Zhou, X. Fang, Y. Fang, Y. Long, R. He, and X. Han, "Enhanced Random Access and Beam Training for Millimeter Wave Wireless Local Networks With High User Density," *IEEE Transactions on Wireless Communications*, vol. 16, no. 12, pp. 7760–7773, 2017.

[71] F. Fellhauer, N. Loghin, D. Ciochina, T. Handte, and S. Brink, "Low Complexity Beamforming Training Method for mmWave Communications," *2017 IEEE 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2017.

[72] L. Lei, B. Li, M. Yang, and Z. Yan, "System Analysis and Performance Evaluation for the Next Generation mmWave WLAN: IEEE 802.11ay," in *IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*. Qingdao: IEEE, Sep. 2018, pp. 1–6.

[73] 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*, vol. 18, no. 10, pp. 2217–2230, Oct. 2019.

[74] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '17, New York, NY, USA, 2017.

[75] 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)*, Hong Kong, China, 2018.

[76] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, 2014.

[77] D. Zhang, P. S. Santhalingam, P. H. Pathak, and Z. Zheng, "Characterizing Interference Mitigation Techniques in Dense 60 GHz mmWave WLANs," *IEEE ICCCN*, pp. 1–9, 2019.

[78] B. Balaji, T. B. Reddy, and B. S. Manoj, "A Novel Power Saving Strategy for Greening IEEE 802.11 Based Wireless Networks," *IEEE GLOBECOM*, pp. 1–5, 2010.

[79] B. Feng, C. Zhang, H. Ding, and Y. Fang, "PhyCast: Towards Energy Efficient Packet Overhearing in WiFi Networks," *IEEE ICC*, pp. 1–6, 2018.

[80] L. Lanante and S. Roy, "Performance analysis of the ieee 802.11ax obss_pd-based spatial reuse," *IEEE/ACM Transactions on Networking*, vol. 30, no. 2, pp. 616–628, 2022.

[81] C.-Y. Liu, M.-S. Sie, E. W. J. Leong, Y.-C. Yao, C.-W. Jen, W.-C. Liu, C.-F. Wu, and S.-J. Jou, "Dual-Mode All-Digital Baseband Receiver With a Feed-Forward and Shared-Memory Architecture for Dual-Standard Over 60 GHz NLOS Channel,"

*IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, no. 3, pp. 608–618, 2017.

[82] J. O. Lacruz, D. Garcia, P. J. Mateo, J. Palacios, and J. Widmer, "mm-FLEX: An Open Platform for Millimeter-Wave Mobile Full-Bandwidth Experimentation," in *ACM MobiSys*, 2020, pp. 1–13.

[83] Y.-S. Huang, W.-C. Liu, and S.-J. Jou, "Design and implementation of synchronization detection for ieee 802.15.3c," in *2011 International Symposium on VLSI Design, Automation and Test*, 2011, pp. 1–4.

[84] C.-C. Tseng and C. Liu, "Complementary sets of sequences," *IEEE Transactions on Information Theory*, vol. 18, no. 5, pp. 644–652, 1972.

[85] D. Garcia Marti, J. O. Lacruz, P. Jimenez Mateo, J. Palacios, R. Ruiz, and J. Widmer, "Scalable Phase-Coherent Beam-Training for Dense Millimeter-wave Networks," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2021.

[86] E. Magistretti, O. Gurewitz, and E. W. Knightly, "802.11ec: Collision avoidance without control messages," *IEEE/ACM Transactions on Networking*, vol. 22, no. 6, pp. 1845–1858, 2014.

[87] J. Classen, M. Schulz, and M. Hollick, "Practical covert channels for WiFi systems," in *IEEE CNS*, 2015, pp. 209–217. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7346830

[88] H. Rahbari and M. Krunz, "Full Frame Encryption and Modulation Obfuscation Using Channel-Independent Preamble Identifier," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 12, pp. 2732–2747, 2016. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7495006

[89] Z. Zhang, H. Rahbari, and M. Krunz, "Adaptive Preamble Embedding with MIMO to Support User-defined Functionalities in WLANs," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2021. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9476997

[90] ——, "Expanding the Role of Preambles to Support User-defined Functionality in MIMO-based WLANs," in *IEEE INFOCOM*, 2020, pp. 1191–1200.

[91] X. Zhang and K. G. Shin, "E-MiLi: Energy-Minimizing Idle Listening in Wireless Networks," *IEEE Transactions on Mobile Computing*, vol. 11, no. 9, pp. 1441–1454, 2012.

[92] A. Bartoli, J. Hernández-Serrano, M. Soriano, M. Dohler, A. A. Kountouris, and D. Barthel, "Secure Lossless Aggregation Over Fading and Shadowing Channels for

Smart Grid M2M Networks," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 844–864, 2011.

[93] S. Biswas and S. Datta, "Reducing Overhearing Energy in 802.11 Networks by Low-power Interface Idling," in *IEEE IPCCC*, 2004, pp. 695–700.

[94] M. Ergen and P. Varaiya, "Decomposition of Energy Consumption in IEEE 802.11," *IEEE ICC*, pp. 403–408, 2007.

[95] S. Jog, J. Wang, J. Guan, T. Moon, H. Hassanieh, and R. R. Choudhury, "Many-to-Many Beam Alignment in Millimeter Wave Networks," in *NSDI*, 2019, pp. 783–800.

[96] S. Madani, S. Jog, J. O. Lacruz, J. Widmer, and H. Hassanieh, "Practical Null Steering in Millimeter Wave Networks," in *NSDI*, 2021, pp. 903–921.

[97] Y. Yang, A. Zhou, D. Xu, S. Yang, L. Wu, H. Ma, T. Wei, and J. Liu, "mmMuxing: Pushing the Limit of Spatial Reuse in Directional Millimeter-wave Wireless Networks," in *IEEE SECON*, 2020, pp. 1–9.

[98] W. Feng, Y. Wang, D. Lin, N. Ge, J. Lu, and S. Li, "When mmWave Communications Meet Network Densification: A Scalable Interference Coordination Perspective," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 7, pp. 1459–1471, 2017.

[99] R. R. Choudhury, X. Yang, N. H. Vaidya, and R. Ramanathan, "Using Directional Antennas for Medium Access Control in Ad Hoc Networks," in *ACM MobiCom*. Association for Computing Machinery, 2002, pp. 59–70.

[100] R. R. Choudhury and N. H. Vaidya, "Deafness: a MAC Problem in Ad Hoc Networks When Using Directional Antennas," *IEEE ICNP*, pp. 283–292, 2004.

[101] M. Takai, J. Martin, R. L. Bagrodia, and A. Ren, "Directional Virtual Carrier Sensing for Directional Antennas in Mobile Ad Hoc Networks," in *ACM MobiHoc*, 2002, pp. 183–193.

[102] M. X. Gong, R. Stacey, D. Akhmetov, and S. Mao, "A Directional CSMA/CA Protocol for mmWave Wireless PANs," *IEEE WCNC*, pp. 1–6, 2010.

[103] I. K. Son, S. Mao, M. X. Gong, and Y. Li, "On frame-based scheduling for directional mmWave WPANs," *IEEE INFOCOM*, pp. 2149–2157, 2012.

[104] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding-Royer, and M. J. W. Rodwell, "Blockage and directivity in 60 GHz wireless personal area networks: from cross-layer model to multihop MAC design," *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 8, pp. 1400–1413, 2009.

[105] Q. Chen, J. Tang, D. T. C. Wong, X. Peng, and Y. Zhang, "Directional Cooperative MAC Protocol Design and Performance Analysis for IEEE 802.11ad WLANs," *IEEE Transactions on Vehicular Technology*, vol. 62, pp. 2667–2677, 2013.

[106] A. Akhtar and S. C. Ergen, "Directional MAC protocol for IEEE 802.11ad based wireless local area networks," *Ad Hoc Networks*, vol. 69, pp. 49–64, 2018.

[107] A. Alkhateeb, O. E. Ayach, G. Leus, and R. W. Heath, "Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, pp. 831–846, 2014.

[108] S. Han, C. lin I, Z. Xu, and C. Rowell, "Large-scale antenna systems with hybrid analog and digital beamforming for millimeter wave 5G," *IEEE Communications Magazine*, vol. 53, pp. 186–194, 2015.

[109] M. Malloy and R. D. Nowak, "Near-Optimal Adaptive Compressed Sensing," *IEEE Transactions on Information Theory*, vol. 60, pp. 4001–4012, 2013.

[110] M. Lecci, M. Polese, C. Lai, J. Wang, C. Gentile, N. Golmie, and M. Zorzi, "Quasi-deterministic channel model for mmwaves: Mathematical formalization and validation," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6.

[111] F. Fellhauer, N. Loghin, D. Ciochina, T. Handte, and S. ten Brink, "Low Complexity Beamforming Training Method for mmWave Communications," in *2017 IEEE 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2017, pp. 1–5.

[112] G. Interdonato, E. Björnson, H. Q. Ngo, P. K. Frenger, and E. G. Larsson, "Ubiquitous cell-free Massive MIMO communications," *EURASIP Journal on Wireless Communications and Networking*, vol. 2019, pp. 1–13, 2018.

[113] E. Björnson and L. Sanguinetti, "Making Cell-Free Massive MIMO Competitive With MMSE Processing and Centralized Implementation," *IEEE Transactions on Wireless Communications*, vol. 19, pp. 77–90, 2019.

[114] M. Alonzo, S. Buzzi, A. Zappone, and C. D'Elia, "Energy-Efficient Power Control in Cell-Free and User-Centric Massive MIMO at Millimeter Wave," *IEEE Transactions on Green Communications and Networking*, 2019.

[115] G. Femenias and F. Riera-Palou, "Cell-Free Millimeter-Wave Massive MIMO Systems With Limited Fronthaul Capacity," *IEEE Access*, vol. 7, 2019.

[116] N. T. Nguyen, K. Lee, and H. Dai, "Hybrid Beamforming and Adaptive RF Chain Activation for Uplink Cell-Free Millimeter-Wave Massive MIMO Systems," *IEEE Transactions on Vehicular Technology*, vol. 71, 2020.

[117] J.-C. Guo, Q. Yu, W. Sun, and W.-X. Meng, "Robust Efficient Hybrid Pre-Coding Scheme for mmWave Cell-Free and User-Centric Massive MIMO Communications," *IEEE Transactions on Wireless Communications*, 2021.

[118] C. D'Andrea, G. Interdonato, and S. Buzzi, "User-centric Handover in mmWave Cell-Free Massive MIMO with User Mobility," *2021 29th European Signal Processing Conference (EUSIPCO)*, pp. 1–5, 2021.

[119] R. Hersyandika, Q. Wang, and S. Pollin, "Association in Dense Cell-Free mmWave Networks," *ICC 2021*, pp. 1–6, 2021.

[120] C. Chen, H. Song, Q. Li, F. Meneghello, F. Restuccia, and C. Cordeiro, "Wi-Fi Sensing Based on IEEE 802.11bf," *IEEE Communications Magazine*, vol. 61, no. 1, pp. 121–127, 2023.

[121] F. Meneghello, C. Chen, C. Cordeiro, and F. Restuccia, "Toward Integrated Sensing and Communications in IEEE 802.11bf Wi-Fi Networks," *IEEE Communications Magazine*, vol. 61, no. 7, pp. 128–133, 2023.

[122] S. Blandino, T. Ropitault, C. R. C. M. da Silva, A. Sahoo, and N. Golmie, "IEEE 802.11bf DMG Sensing: Enabling High-Resolution mmWave Wi-Fi Sensing," *IEEE Open Journal of Vehicular Technology*, vol. 4, pp. 342–355, 2023.