

ANALYSIS AND PERFORMANCE IMPROVEMENT OF
CONSUMER-GRADE MILLIMETER WAVE WIRELESS
NETWORKS

by

GUILLERMO BIELSA

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

Multimedia and Communication Interuniversity PhD
Program

Universidad Carlos III de Madrid

Advisor:
Joerg Widmer

July 2019

Analysis and Performance Improvement of Consumer-Grade Millimeter Wave Wireless Networks

Prepared by:

Guillermo Bielsa, IMDEA Networks Institute, Universidad Carlos III de Madrid

contact: guillermo.bielsa@imdea.org

Under the advice of:

Joerg Widmer, IMDEA Networks Institute

Department of Signal Theory and Communications Universidad Carlos III de Madrid

This work has been supported by:



Unless otherwise indicated, this thesis is distributed under a
Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA).

Acknowledgements

Firstly, I would like to thank Dr. Joerg Widmer for his supervision, guidance, and patience through all these years. This thesis would not be possible without the opportunity that he gave me joining his research group during my bachelor studies. Secondly, I would like to thank the guidance and help received by Dr. Thomas Nitsche and Dr. Adrian Loch during the time we shared at IMDEA Networks Institute.

I would like to express my gratitude to all the collaborators that helped me developing myself during my research career: Joan Palacios, Dr. Michele Zorzi, Maria Scalabrin, Dr. Michele Rossi, Dr. Daniel Steinmetzer, Irene Tejado, and Alain Olivier. Special thanks go to Dr. Sundeep Rangan and to Dr. Marco Mezzavilla for hosting me during my stay at NYU. Moreover, I have had the opportunity to work with Dr. Paolo Casari who has been more than a great collaborator, I thank him for providing lots of useful knowledge during our time working together.

Along these lines, I would like to thank the people I met when I joined IMDEA Networks, who encouraged me to continue the adventure joining the institute as a PhD student, in addition to the rest of amazing people that I met during all these years. Without all these people and the atmosphere that they create on a daily basis, my time at the institute would not have been the same. There is no space to list all their names here, but they know who they are. I would also like to thank what allowed us to know each other better and provided us the fuel we needed to continue our research: Our beloved ping-pong table, Nesspreso, and Cola-Cao, as well as Bar Manzanares and Piratas Rock Wey.

Last but not least, I would like to thank for the support and company to all the people that surround me during my personal life. First of all, my family, who gave me support, love, and everything that I needed. Elena, for sharing her life with me making us grow together. And finally, all those people I met during my life who helped making it better.

Abstract

Millimeter-wave (mmWave) networks are one of the main key components in next cellular and WLANs (Wireless Local Area Networks). mmWave networks are capable of providing multi gigabit-per-second rates with very directional low-interference and high spatial reuse links. In 2013, the first 60 GHz wireless solution for WLAN appeared in the market. These were wireless docking stations under the WiGig protocol. Today, in 2019, 60 GHz communications have gained importance with the IEEE 802.11ad amendment with different products on the market, including routers, laptops and wireless Ethernet solutions. More importantly, mmWave networks are going to be used in next generation cellular networks, where smartphones will be using the 28 GHz band. For backbone links, 60 GHz communications have been proposed due to its higher directionality and unlicensed use. This thesis fits in this frame of constant development of the mmWave bands to meet the needs of latency and throughput that will be necessary to support future communications. In this thesis, we first characterize the cost-effective design of COTS (commercial off-the-shelf) 60 GHz devices and later we improve their two main weaknesses, which are their low link distance and their non-ideal spatial reuse.

It is critical to take into consideration the cost-effective design of COTS devices when designing networking mechanisms. This is why in this thesis we do the first-of-its-kind COTS analysis of 60 GHz devices, studying the D5000 WiGig Docking station and the TP-Link Talon IEEE 802.11ad router. We include static measurements such as the synthesized beam patterns of these devices or an analysis of the area-wide coverage that these devices can fulfill. We perform a spatial reuse analysis and study the performance of these devices under user mobility, showing how robust the link can be under user movement. We also study the feasibility of having flying mmWave links. We mount a 60 GHz COTS device into a drone and perform different measurement campaigns. In this first analysis, we see that these 60 GHz devices have a large performance gap for the achieved communication range as well as a very low spatial reuse. However, they are still suitable for low density WLANs and for next generation aerial micro cell stations.

Seeing that these COTS devices are not as directional as literature suggests, we analyze how channels are not as frequency stable as expected due to the large amount of reflected signals. Ideally, frequency selective techniques could be used in these frequency selective channels in order to enlarge the range of these 60 GHz devices. To validate this, we measure real-world

60 GHz indoor channels with a bandwidth of 2 GHz and study their behavior with respect to techniques such as bitloading, subcarrier switch-off, and waterfilling. To this end, we consider a Orthogonal Frequency-Division Multiplexing (OFDM) channel as defined in the IEEE 802.11ad standard and show that in point of fact, these techniques are highly beneficial in mmWave networks allowing for a range extension of up to 50%, equivalent to power savings of up to 7 dB.

In order to increase the very limited spatial reuse of these wireless networks, we propose a centralized system that allows the network to carry out the beam training process not only to maximize power but also taking into account other stations in order to minimize interference. This system is designed to work with unmodified clients. We implement and validate our system on commercial off-the-shelf IEEE 802.11ad hardware, achieving an average throughput gain of 24.67% for TCP traffic, and up to a twofold throughput gain in specific cases.

Published and submitted content

[1] T. Nitsche, **G. Bielsa**, I. Tejado, A. Loch, and J. Widmer (2015). Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation. In: *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT '15)*. ACM, New York, NY, USA, Article 17 , 13 pages. <https://doi.org/10.1145/2716281.2836102>

- *This paper has been partially included in Chapter 3.*
- *The author's role in this work is focused on the implementation, execution and analysis of the measurements.*

[2] **G. Bielsa** Analysis of off-the-shelf Millimeter Wave Systems with Phased Antenna Arrays. Master Diss. Universidad Carlos III de Madrid, Spain, (2016). http://eprints.networks.imdea.org/1369/1/TFM_GuillermoBielsa.pdf

- *This Master thesis is based on [1] and has been partially included in Chapter 3.*
- *The author's role in this work is focused on the implementation, execution and analysis of the measurements.*

[3] **G. Bielsa**, A. Loch, I. Tejado, T. Nitsche, and J. Widmer (2018). 60 GHz Networking: Mobility, Beamforming, and Frame Level Operation From Theory to Practice. In: *IEEE Transactions on Mobile Computing*. <https://doi.org/10.1109/TMC.2018.2875095>

- *This is a journal extension of [1] and has been fully included in Chapter 3.*
- *The author's role in this work is focused on the design, implementation, execution and analysis of the new measurements.*

[4] A. Loch, **G. Bielsa**, and J. Widmer (2016, October). Practical Lower Layer 60 GHz Measurements Using Commercial Off-The-Shelf Hardware. In: *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization* (pp. 9-16). ACM. <https://doi.org/10.1145/2980159.2980165>

- *This paper has been partially included in Chapter 3.*

- *The author's role in this work is focused on the implementation, execution and analysis of the measurements as well as some debuggin issues of the hidden monitor.*

[5] H. Assasa, **G. Bielsa**, S. Kumar Saha, P. Jiménez Mateo, A. Loch, D. Koutsonikolas, and J. Widmer. Performance Analysis of Medium Access Control and Spatial Reuse for IEEE 802.11ad Deployments. Submitted to: *Elsevier Pervasive and Mobile Computing*, [Major revision, March 19].

- *This paper has been partially included in Chapters 3 and 5.*
- *The author's role in this work is focused on the design, implementation, execution and analysis of the measurements of Section VII: Spatial Sharing.*

[6] **G. Bielsa**, M. Mezzavilla, J. Widmer, and S. Rangan (2019, June). Performance Assessment of Off-the-Shelf mmWave Radios for Drone Communications. In: *The 20th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (IEEE WOWMOM 2019)*, 10-12 June 2019, Washington DC, USA. http://eprints.networks.imdea.org/1997/1/DroneComms...3_pages_WoWMoM.pdf

- *This paper has been fully included in Chapter 3.*
- *The author's role in this work is focused on the design, implementation, execution and analysis of the aerial measurements.*

[7] **G. Bielsa**, A. Loch, and J. Widmer (2017, June). 60 GHz Range Boost: Exploiting Frequency Selectivity in Millimeter-Wave Networks. In: *A World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2017 IEEE 18th International Symposium on (pp. 1-9). IEEE*. <https://doi.org/10.1109/WoWMoM.2017.7974291>

- *This paper has been fully included in Chapter 4.*
- *The author's role in this work is focused on the design, implementation, execution and analysis of the channel measurements and frequency selectivity techniques.*

[8] **G. Bielsa**, A. Loch, and J. Widmer. Optimizing mmWave Spatial Reuse: Signal-to-Interference Aware Beamtraining. In: *Internet of Things: Smart Objects and Services*, 10 June 2019, Washington DC, USA. http://eprints.networks.imdea.org/2002/1/Optimizing_mmW_Spatial_Reuse.pdf

- *This paper has been fully included in Chapter 5.*
- *The author's role in this work is focused on the design, implementation, execution and analysis of the beamtraining techniques and measurements.*

Other publications that have not been used in this thesis are:

- [9] A. Olivier, **G. Bielsa**, I. Tejado, M. Zorzi, J. Widmer and P. Casari, Lightweight Indoor Localization for 60-GHz Millimeter Wave Systems. In: *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, London, 2016, pp. 1-9. <https://doi.org/10.1109/SAHCN.2016.7732999>
- [10] M. Scalabrin, M. Rossi, **G. Bielsa**, A. Loch and J. Widmer. Millimetric Diagnosis: Machine Learning Based Network Analysis for mm-Wave Communication. In: *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Macau, 2017, pp. 1-9. <https://doi.org/10.1109/WoWMoM.2017.7974309>
- [11] A. Loch, **G. Bielsa**, R. Santos, and J. Widmer (2017, June). mm-View: Obtaining Real-Time Lower Layer Information of Commercial Off-The-Shelf 60 GHz Hardware. In: *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, (pp. 1-3). IEEE. <https://doi.org/10.1109/WoWMoM.2017.7974330>
- [12] **G. Bielsa**, J. Palacios, A. Loch, D. Steinmetzer, P. Casari and J. Widmer, Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points. In: *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, Honolulu, HI, 2018, pp. 2384-2392. <https://doi.org/10.1109/INFOCOM.2018.8486232>
- [13] J. Palacios, **G. Bielsa**, P. Casari and J. Widmer, Communication-Driven Localization and Mapping for Millimeter Wave Networks. In: *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, Honolulu, HI, 2018, pp. 2402-2410. <https://doi.org/10.1109/INFOCOM.2018.8485819>
- [14] J. Palacios, **G. Bielsa**, P. Casari and J. Widmer, Single- and Multiple-Access Point Indoor Localization for Millimeter-Wave Networks. In: *IEEE Transactions on Wireless Communications*, vol. 18, no. 3, pp. 1927-1942, March 2019. <https://doi.org/10.1109/TWC.2019.2899313>
- [15] M. Scalabrin, **G. Bielsa**, A. Loch, M. Rossi and J. Widmer, Machine Learning Based Network Analysis using Millimeter-Wave Narrow-Band Energy Traces. In: *IEEE Transactions on Mobile Computing*. <https://doi.org/10.1109/TMC.2019.2907585>

Table of Contents

Acknowledgements	VII
Abstract	IX
List of Publications	XI
Table of Contents	XV
List of Tables	XIX
List of Figures	XXII
List of Acronyms	XXIII
1. Introduction	1
1.1. Motivation	2
1.2. Contributions	3
2. Background and Related Work	5
2.1. Transmission Characteristics	5
2.2. Beam Steering	6
2.3. Work on Practical 60 GHz Networks	8
3. Analysis of Commercial Off-The-Shelf Systems	9
3.1. Experiment Setup	12
3.1.1. Devices Under Test	12
3.1.2. Lower Layer Parameter Collection	13
3.1.3. Comprehensive Trace Collection	14
3.2. Static Analysis	15
3.2.1. Protocol Analysis	16
3.2.2. Beam patterns	21
3.2.3. Reflections	23
3.2.4. Area-Wide Coverage	26

3.2.5. Spatial Sharing	27
3.3. Mobile Analysis	29
3.3.1. Experiment Design	29
3.3.2. Mobility Results	29
3.3.3. Multiple Antenna Arrays	34
3.3.4. Link Robustness	35
3.3.5. Transient Blockage	36
3.4. Aerial Communications	37
3.4.1. Experimental Setup	38
3.4.2. Results	40
3.5. Discussion	43
4. Range Boost	47
4.1. Exploiting Frequency Selectivity	48
4.1.1. Compared Techniques	49
4.1.2. Metrics	50
4.2. Evaluation of Channel Characterization	51
4.2.1. Testbed Setup	51
4.2.2. Channel Characterization	52
4.3. Results	56
4.3.1. Range Extension	56
4.3.2. Energy Saving	59
4.3.3. Antenna Array Comparison	61
4.4. Discussion	63
5. Optimizing Spatial Reuse	65
5.1. Impact of Neighbouring Nodes Beam Selection	67
5.1.1. Measurement Setup	68
5.1.2. Beam pattern selection comparison	68
5.2. Beamtraining Mechanisms	70
5.2.1. Baseline: IEEE 802.11ad beamtraining	72
5.2.2. Weighted SIR Fairness	72
5.2.3. Argmaxmin	72
5.2.4. Power Threshold	73
5.2.5. Interference Threshold	73
5.3. Experimental Setup	73
5.3.1. Router Implementation	73
5.3.2. Methodology	74
5.3.3. Scenarios	74
5.4. Evaluation	75

5.4.1. Empty Auditorium	75
5.4.2. Office Environment: 2 Links	76
5.4.3. Office Environment: 3 Links	80
5.5. Discussion	81
6. Outlook	83
7. Conclusions	85
References	96

List of Tables

3.1. Monitored Lower Layer Parameters	15
3.2. D5000 dock, Talon, and WiHD frame periodicity	17
3.3. Link orientation related throughput	28
3.4. Static VS Mobile result statistics.	30
3.5. Performance for multiple antenna arrays on trajectories of Figure 3.16.	35
3.6. Average throughput/outage up to eight arrays on random trajectories.	35
4.1. Evaluation Scenarios	52
4.2. Frequency selectivity with distance	54
5.1. Maximum bitrate for different selected beam patterns in Figure 5.2.	70
5.2. Optimizing spatial reuse: TCP average gain	77
5.3. Optimizing spatial reuse: UDP average results	78

List of Figures

3.1. COTS analysis: Devices under test.	12
3.2. Overview of the Wilocity Monitor.	13
3.3. Frame structure of the Dell D5000 docking station system.	17
3.4. Static analysis for fixed link length.	18
3.5. Aggregation primer	18
3.6. Frame level analysis of a WiGig link operating at 60 GHz.	19
3.7. Impact of distance on link rate and transport layer throughput.	20
3.8. DVDO Air-3c WiHD frame flow.	21
3.9. Beam pattern measurements	21
3.10. D5000 Device discovery beam pattern measurements	22
3.11. Reflection patterns in conference room	23
3.12. Data transmission exclusively via a reflected path.	24
3.13. Reflection interference experiment.	25
3.14. Throughput for our area-wide grid measurement with the D5000 dock	26
3.15. Spatial sharing performance with respect to separation distance.	28
3.16. Examples of movement trajectories on our measurement grid.	30
3.17. Bitrate comparison of a trajectory for the static and the mobile case.	31
3.18. Modulation and Coding Scheme (MCS) for the static and the mobile trajectory.	32
3.19. Match of metrics in the static case compared to the mobile case.	32
3.20. Beampattern fluctuations for two trajectories.	33
3.21. Horizontal trajectories: Result for different orientations	36
3.22. Transient blockage study.	37
3.23. Aerial communications: Testbed setup.	38
3.24. Experimental setup.	39
3.25. Vertical study analysis.	40
3.26. 45° elevation study analysis.	40
3.27. 3D distance study analysis: Bitrate & MCS results.	41
3.28. 3D distance study analysis: Beam pattern selection.	42
3.29. Drone results: Wind emulation bitrate	42

4.1. Channel measurement experimental setup	51
4.2. Consecutive channel measurements in the corridor scenario at 9 m.	53
4.3. CDF of the signal strength of subcarrier 100 during one hour.	53
4.4. Channel measurements for all the positions.	54
4.5. Channel measurement in the empty room scenario at 10 m and 15 m.	55
4.6. Measured overall signal strength compared to the theoretical Friis equation.	56
4.7. Achievable non-coded throughput.	57
4.8. Achievable coded throughput	58
4.9. Achievable energy savings for different frequency selective techniques.	60
4.10. Antenna array results	62
5.1. Beamtraining example	66
5.2. Aggregated throughput on all beam pattern configurations	69
5.3. Beamtraining frame exchange	71
5.4. Spatial reuse mechanisms validation	76
5.5. Optimizing spatial reuse: Office setup	77
5.6. Optimizing spatial reuse: CDF results	79
5.7. Optimizing spatial reuse: Two link example results	80
5.8. Optimizing spatial reuse: Three link example results	81

List of Acronyms

A-BFT	Association Beamforming Training
AGC	Automatic Gain Control
AP	Access Point
BER	Bit Error Rate
BI	Beacon Interval
BTI	Beacon Transmission Interval
CDF	Cumulative Distribution Function
COTS	Commercial Off-The-Shelf
CRC	Cyclic Redundancy Check
CSMA	Carrier Sense Multiple Access
CSV	Comma-Separated Values
DTI	Data Transmission Interval
EIRP	Equivalent Isotropically Radiated Power
FST	Fast Session Transfer
GUI	Graphical User Interface
HPBW	Half-Power Beam Width
ISM	Industrial, Scientific and Medical
LOS	Line-Of-Sight
MAC	Medium Access Control
MCS	Modulation and Coding Scheme

mmWave Millimeter-Wave

OFDM Orthogonal Frequency-Division Multiple

RGF Register File

SDR Software Defined Radio

SINR Signal-to-Interference-plus-Noise Ratio

SIR Signal-to-Interference Ratio

SNR Signal-to-Noise Ratio

STA Station

TCP Transmission Control Protocol

UAV Unmanned Aerial Vehicle

UDP User Datagram Protocol

USRP Universal Software Radio Peripheral

WLAN Wireless Local Area Network

Chapter 1

Introduction

In the recent years, there has been an exponential growth of popularity in wireless devices. These devices are increasing the wireless spectrum demand (more than 20 billion connected devices are expected by 2020 [16]). At the same time, applications such as ultra-high quality video streaming or virtual reality headsets demand gigabit per second rates which up to now are impossible to deliver wirelessly. These wireless applications together with the increasing solicitation on wireless devices multiply the demand for new technology to be able to serve all these users and applications. For example, this demand is estimated to exceed 25 Gbps per square kilometer in 5G network deployments, being about 100-fold over the current capacity [17].

Up to today, Legacy Wireless Local Area Networks (WLANs), such as WiFi or Bluetooth, or cellular networks, such as LTE, work in frequency ranges that can go up to 6 GHz. These sub-6 GHz bands have to deal with the increase of throughput demand as well as managing their highly saturated wireless channels. These devices usually transmit their power in an omnidirectional way, this is, transmitting their power into every direction. This omnidirectional behaviour together with their large communication range and small penetration loss can cause significant interference with neighbouring devices belonging to other networks. In order to avoid collisions due to interference, the network needs to schedule the transmissions of each of the devices. If there is no scheduler in the network, devices will need to sense the medium to verify if there are ongoing transmissions. For such case, if the medium is busy, the devices wait for a random period of time and attempts to transmit again. This collision avoidance mechanisms decrease the communication rates due to these waiting times [18].

Looking into new and fresh ideas, the WiGig Alliance standardized the use of the unlicensed 60 GHz band to avoid the previously mentioned issues, creating the WiGig protocol in 2009 [19]. The WiGig protocol was meant for pseudo-static and single-user scenarios, such as wireless docking stations. These wireless docking stations supported the transmission of high definition video as well as being compatible with the wireless bus extension, allowing USB and Ethernet connections. With this, the vendors were able to assure a 6.7 Gbps wireless connection thanks to the use of 2.16 GHz bandwidth channels in the 60 GHz frequency band. In 2012 the

WiGig alliance merged with the WiFi Alliance, producing together the new IEEE 802.11ad amendment [20]. IEEE 802.11ad was the result of merging the whole multi-user IEEE 802.11 stack together with the 60 GHz physical layer of the WiGig protocol, allowing to have a Millimeter-Wave (mmWave) WiFi network with up to 6.7 Gbps links.

The unlicensed 60 GHz band under the IEEE 802.11ad amendment provides 6 non-overlapping 2.16 GHz bandwidth channels which support both Single-Carrier and Orthogonal Frequency-Division Multiple (OFDM). The use of higher carrier frequency in comparison with sub-6 GHz bands changes radically the wireless behavior of the systems. 60 GHz communications face very high absorption losses. For this reason communication links need to be established in non-obstructed scenarios. These 60 GHz WLAN devices usually have a maximum distance of 15 meters (absorption loss is about 22 dB higher than in legacy sub-6 GHz bands [21]). These high absorption losses produce link outages when the link gets obstructed by blockages such as columns or walls, as well as when a person crosses the Line-Of-Sight (LOS) path of the link. Conversely, as devices hardly transmit through walls, it will help with the interference reduction as well as decreasing the saturation of the wireless spectrum.

Having an increase in frequency is equivalent to a reduction in wavelength. This wavelength reduction gives the possibility of reducing the size of the antennas. 60 GHz devices equip antennas in the order of millimeters of size, allowing to stack many antennas together building what is called an ‘antenna array’ (also named ‘array of antennas’). Antenna arrays have the potential of modifying phase and amplitude of the different antenna elements creating constructive or destructive self-interference, producing the synthesis of directional beam patterns. The directivity of these antenna arrays can be used both in the transmission as well as in the reception side. This high directionality provides low-interference scenarios with high spatial reuse.

1.1. Motivation

Developing robust consumer-grade devices for mmWave wireless networks is challenging. In contrast to IEEE 802.11n/ac networks operating in the traditional 2.4 and 5 GHz bands, devices based on the recent IEEE 802.11ad standard operating in the unlicensed 60 GHz band must overcome significant hurdles. This includes handling 2.16 GHz wide channels, using directional beamforming antennas to overcome the increased attenuation at these frequencies, or dealing with a highly dynamic radio environment [22].

Dealing with such challenging component characteristics at the same time as maintaining cost effective devices for the general market is difficult. A trade-off between the characteristics of the device components and their price is needed. This trade-off leads to Commercial Off-The-Shelf (COTS) devices performing far from the theoretical models that can be found in the research community. However, there are no studies analyzing the performance of whole mmWave COTS devices, making difficult to know how these devices behave in comparison to theoretical models.

Current works in the area allow diverse assumptions, as sector-shaped or formula-based beam patterns [20, 23], as no real analysis has been provided by the community. Nevertheless, there are some well-known issues that reduce the performance of mmWave COTS systems. The relatively few antenna elements that these mmWave COTS antenna arrays mount are known to be not as good as the idealized pencil-beam sectors that literature often assumes. These imperfect low-directional antennas together with the high path loss attenuation are known to cause limited-range communications. Also, their low directivity is known to create interference with neighboring nodes, not allowing the very high spatial reuse suggested in the literature. However, these issues have not been quantified in practical deployments.

This PhD thesis aims to introduce a new point of view on mmWave communications updating the perspective that the community has on these wireless systems. In order to do so, this work presents the first characterization of real COTS 60 GHz devices in terms of frame level analysis, beam pattern measurements, interference impact, and performance under mobility. Once we have characterized these COTS devices, we explore their weaknesses and provide mechanisms to address them. We improve the two main disadvantages of these devices, which are their limited link distance and their low spatial reuse.

1.2. Contributions

We begin this thesis with an intensive analysis of mmWave COTS devices in Chapter 3, where we study the whole performance of these devices and perform a quantitative analysis of their cost-effective design. We start performing a static analysis of these systems. In there, we do a protocol analysis of these first platforms. Then, we study the directivity of these devices by measuring their beam patterns. Further, we analyze how this frequency band creates reflections and we examine how these can create interference to other devices. We also check how these devices behave in terms of distance and orientation performing an area-wide coverage and finally we investigate how these devices perform in terms of spatial reuse for different distances and orientations. The next part of Chapter 3 shows an analysis of these same COTS systems under mobility. There, we compare the previous static area-wide coverage with respect to mobility traces. We then check how the performance would be upgraded with different number of antenna arrays in the mobile device and then study the link robustness under movement and transient blockage. We finalize this analysis on mmWave COTS devices studying the performance of 60 GHz aerial links. There, we perform distance measurements for different flying heights. We also emulate different windy conditions in order to test the feasibility of these mmWave aerial communications.

Chapter 4 helps to solve the first drawn problem of mmWave technology, which is the link range limitation. In order to extend the link range of these devices, we implement frequency selective techniques into real-world 60 GHz channels. In order to do so, we first perform channel measurements in separate indoor scenarios. We measure the whole IEEE 802.11ad channel bandwidth with its OFDM characteristics. The techniques that we apply for this range extension

are subcarrier switch-off, bit loading and water filling combined with bit loading. We first analyze if the channel fits with the use of these techniques, and once this is confirmed, we apply them obtaining range extensions as well as energy savings.

Chapter 5 helps to solve the second sketched problem of mmWave technology, which is the low spatial reuse achieved by these COTS devices. In order to create parallel communications with less interference we develop new beamtraining mechanisms. Instead of using the beam pattern that provides the highest received power as the IEEE 802.11ad amendment does, we propose four different mechanisms that select the beam pattern based not only in the received power but also based on the interference created to other nodes. Each of these mechanisms is based on different optimization goals. We implement and compare all these mechanisms with COTS IEEE 802.11ad routers. For the deployment of these techniques we only need to modify the Access Points (APs) as the system can run for unmodified Stations (STAs).

We can summarize the main contributions of this thesis as:

1. **Characterization of 60 GHz COTS systems.** We focus our study into directionality, spatial reuse, impact on frame aggregation, area-wide coverage, and beamtracking and link robustness under mobility. We provide this analysis together with a free-to-download database, allowing researchers to take these hardware impairments into account [24].
2. **Feasibility study of mmWave COTS radios for aerial links.** We mount an IEEE 802.11ad laptop into a drone and measure the performance of an aerial link for different distances and different emulated wind speeds.
3. **Implementation of frequency selective techniques to extend link distances.** We measure 60 GHz real-world IEEE 802.11ad like indoor channels. To these channels, we apply frequency selective techniques achieving 50% link extension or up to 7 dB energy savings.
4. **Signal to interference aware beamtraining.** In order to improve the low spatial reuse of current mmWave COTS links, we propose different beamtraining techniques. These allow the nodes to choose their beam patterns taking into account not only received power but also the interference that can be created to other users.

Chapter 2

Background and Related Work

In this chapter, we describe the characteristics of mmWave communications and present related work in the area. mmWave frequencies have been used for commercial wireless systems for several years. These first generation systems, however, were targeting mainly static or pseudo static application scenarios, like backhaul links (for example the HXI Gigalink 6451 system used in [25]), wireless data centers [26, 27] or transmission of uncompressed high definition video data [28]. With formation of the WiGig Alliance [19], the latter use case broadened to docking station applications and finally incorporated general WiFi use cases when, in 2012, WiGig was merged into the IEEE 802.11ad amendment. This amendment defines a unified mmWave communication standard for a variety of use cases, including dense scenarios with indoor mobility [22]. The amendment adds so-called sector sweeps at regular beacon intervals. During a sweep, devices sequentially transmit beacons on each of their beam patterns. This allows nodes to discover each other, and determine the best beam pattern for communication. Within a beacon interval, devices can either use random or scheduled access. The latter avoids certain issues of random access in directional networks, such as collisions due to deafness. However, existing devices do not implement it.

2.1. Transmission Characteristics

Wireless communications in the mmWave band have distinctly different communication characteristics than those using frequencies below 6 GHz. These differences mainly result from the increased attenuation of free space propagation and signal blockage in case of obstacles. Applying the Friis free space transmission equation, we can find how the attenuation increases 20 dB when changing from 6 GHz to 60 GHz carrier frequency. This 20 dB increase on free space loss, together with the fact that atmospheric oxygen attenuates the 60 GHz band transmissions at a 10-15 dB/km regime, makes these communications unsuitable for long-range transmissions [21, 29]. This attenuation of around 20-40 dB is typically overcome by highly directional antennas, which can be used both in the transmitter as in the receiver side. Due to the

high directionality of these antenna arrays and the small wavelength, mmWave communications are highly vulnerable to blockage. Walking by a column or blocking the communication LOS can easily break the link. Antenna arrays may also be used to circumvent blockage, using a propagation path via a reflection. There are many mechanisms in the literature to try to fix transient blockages as the one seen in Beamspy [30], which is capable of predicting blockages and circumvent them. Such reflections are possible if certain highly reflective materials are present [31]. Propagation parameters and channel models for understanding mmWave propagation have been compared and modeled by various research groups and standardization bodies [32]. Further, directional communication and blockage significantly lower the amount of interference on mmWave frequencies and allow for high levels of spatial reuse.

The range of practical 60 GHz networks depends on a number of factors, such as the directivity of the antenna and the allowed transmission power. Recent work [25, 28] uses COTS hardware to determine the actual impact of path loss in real-world environments. This adds to related work which measures range using custom-built solutions [33, 34]. In both cases, the general consensus is that the range of 60 GHz networks in LOS conditions is 10 to 15 meters. In isolated cases, ranges of up to 25 meters are feasible. Such outliers may be due to, e.g., dry environments in which water absorption is limited. Still, given the limited range in the standard case, extending the link length is highly interesting.

Related work shows that directional communications play a fundamental role regarding channel characteristics, as using directional antennas reduces the multi-path effects [35]. Moreover, signals reflected off walls typically suffer strong attenuation in addition to high path loss. Still, extensive measurement campaigns show that even second-order reflections can still have a significant impact in mmWave frequencies [31, 36]. These mmWave channels are still very sparse, as independently of the number of beam directions or reflected paths available, the communication can only be established through a limited number of signal paths [37]. Designing cost-efficient phased antenna arrays to filter out reflections is feasible [38–40] but often not enough as 60 GHz COTS devices experience strong side lobes. In other words, actual commercial hardware is far from the ideal pencil-shaped beams often assumed in simulative work on 60 GHz networks [41]. Thus, reflections are likely to fall into a side lobe of the receiver, which may cause frequency selectivity. Yang et al. [42] analyze this issue in practice, and show that 60 GHz channels indeed exhibit a certain degree of frequency selectivity for beams which are not pencil-shaped. This motivates our work since it suggests that techniques exploiting such selectivity may be suitable for range extension.

2.2. Beam Steering

With the central role of directional communication in mmWave systems, a device's ability to steer its antenna beams becomes essential. As the size of antennas scales with the wave length, mmWave systems can integrate antenna arrays with a high number of elements even into small

handheld devices. These antenna arrays allow electronic configuration of the antennas' beam direction and provide very high directional gain. In contrast to lower frequency beamforming mechanisms, mmWave systems usually rely on beam steering via codebooks of predefined beam patterns that implement different directions, which reduces the complexity of the beamtraining process. The way these devices perform their beamtraining, is with a periodical process where the AP tests all the possible Tx-beam patterns of its codebook with the initiator sector sweep while the STAs receive omni-directionally and then the STAs answers back in the same manner, iterating through all the possible Tx-beam patterns, including on each of these frames which was the best AP beam pattern. Then, the AP acknowledges back with the best beam pattern to be used by the STA.

While mmWave antenna arrays have been in use since first generation devices, their impact on system level network performance is not well understood, although realistic antenna patterns in mmWave cellular scenarios are being studied by means of simulations [43]. As these beam patterns do not appear to be as directional as they could be due to the use of limited codebooks, there has been large efforts from the community to improve current mmWave beamforming and beamsteering. For instance, Palacios et al. [44] propose a hybrid analog-digital adaptive beamforming technique to accelerate the link establishment relaying on a geometric strategy, resulting into time gains as well as an increase in spectral efficiency compared with the existing scanning strategies. A hybrid beamforming optimization subject to a Signal-to-Interference-plus-Noise Ratio (SINR) constrain is studied in [45]. First they consider the minimization of transmitted power to achieve a given SINR, and later they maximize the worst SINR case given a total power constraint. An other interference alignment scheme for 60 GHz communication system is proposed in [46]. The authors split their antenna array into sub-arrays where independent beamforming is performed, and then interference alignments is applied. In [47] the effective network throughput is maximized by a joint study on the overhead of aligning the Tx-beam pattern modifying its beamwidth together with the traffic scheduling on each time slot, as only a single transmission can be done per slot.

All these mentioned solutions would help boosting the global performance and the spatial reuse of mmWave dense scenarios. However, they do need to modify APs and STAs beamforming protocols, as well as needing extra control exchange. It is also important to mention that the previously exposed works are only validated by means of simulations, not taking into account the real-world cost-effective design of mmWave devices.

Another issue that mmWave communications have to deal with is the use of directional communications under user mobility. Mobility plays a key role in 60 GHz networks as even slight beam misalignments can undermine performance. By default, the IEEE 802.11ad protocol resorts to costly beamtraining when nodes move. Related work presents improvements on this behavior, such as short frames preceding data packets to assess the current state of the channel and react accordingly [48], or using the data packet preamble to slightly steer the main lobe to the sides [49]. Cacciapuoti [50] has designed a mobility-aware user association strategy for 5G-

mmWave networks based on association strategies. Other works leverage the ability of hybrid analog-digital transceivers to collect channel information from multiple spatial directions at the same time in order to probabilistically optimize the problem of mmWave link maintenance under mobility [51].

While recent 60 GHz experiment platforms provide full access to such lower layer issues [52], these are not standard compliant and are based on hardware components that are very different from the ones used in commercial devices. Despite recent advances in terms of 60 GHz link robustness, mobility remains an open issue. Existing work mitigates the impact of link blockage [30, 53] but has only limited applicability to mobile scenarios.

2.3. Work on Practical 60 GHz Networks

At the time of writing, insights into practical 60 GHz networks were limited since IEEE 802.11ad hardware was not widely available. Only two commercial wireless routers and a few mobile devices supported IEEE 802.11ad. Related work uses WiGig and WiHD devices to study the performance of 60 GHz links. For instance, Zheng et al. [25] use off-the-shelf hardware to characterize 60 GHz links, dispelling a number of common beliefs regarding mmWave communication. This includes showing that the range of such networks is large enough for outdoor communication, and that electronically steerable antenna arrays can deal with blockage as well as user motion. Still, in contrast to our work, they do not study and measure the beam patterns of off-the-shelf hardware nor do they perform area-wide measurements. Also, they focus on high-level metrics, while we analyze the frame level operation of WiGig devices and its behavior at the physical layer.

Ramanathan et al. [54] follow a different approach than [25] and, hence, us. Instead of measuring the performance of off-the-shelf devices, they use a custom-built software-defined radio that allows them to perform signal strength measurements at 60 GHz using horn antennas. Among other results, they show that different types of link breaks should be treated differently. For instance, a transmitter should change beam direction to avoid human blockage, whereas it should widen its beam to deal with mobility. Most interestingly, they show how the signal strength can be used for early detection of each type of link break. In contrast, we study actual data transmissions using hardware with electronically steerable phased antenna arrays. While [54] focuses on blockage, beam steering, and spatial reuse assuming horn antennas, we investigate topics such as reflections for range extension, data aggregation, and interference due to imperfections of commercial hardware.

Chapter 3

Analysis of Commercial Off-The-Shelf Millimeter Wave Systems

This chapter focuses on the analysis of mmWave COTS devices. In contrast to legacy Industrial, Scientific and Medical (ISM) devices where their center frequency can go up to 5.9 GHz and channel bandwidths up to 160 MHz, in the IEEE 802.11ad amendment, the carrier frequency can go up to 70.2 GHz and the channel bandwidth to 2.16 GHz. Developing electronic components such as digital to analogue converters or amplifiers capable of dealing with such large channel bandwidths, together with their corresponding channel separation is very challenging. The MAC layer of these 60 GHz devices also needs to face new problems such as beamforming and directional device discovery and initial access.

Existing work studies the individual factors that such a system must address regarding both the characteristics of 60 GHz communication—fading [33], reflections [31], frequency selectivity [36,42], and multipath effects [35,36]—as well as the design of hardware such as phased antenna arrays [40]. The challenge is that consumer-grade devices must be able to handle the resulting complexity of the overall system while still being cost-effective.

This has naturally led to the use of cost effective components in consumer-grade 60 GHz systems. For instance, current devices use electronic beam steering with phased arrays with relatively few antenna elements. As a result, while the transmitters use directional beam patterns, they do not fully achieve the 60 GHz vision of extreme pencil-beams with negligible interference. While this is qualitatively well-known, it raises a crucial question: how large is the practical, quantitative impact of such limitations? Understanding this impact is fundamental since the aforementioned limitations may undermine common 60 GHz assumptions. This in turn is key to design protocols that can reliably operate at mmWave frequencies. As a quantitative analysis of these limitations is missing, related work in this area is often based on diverse assumptions. For instance, assumptions on the shape of beam patterns range from triangle-shaped sectors as in the IEEE 802.11ad standard [20] to ideal lobes as defined by theoretical formulas [23]. Neither reflects the actual shape of practical beam patterns. This also hinders the comparison of results

among papers. Previous work on practical 60 GHz networks based on consumer-grade devices is limited to studying the impact of human blockage [28] and transmission range [25]. Also, existing studies do not make traces available for reuse.

However, 60 GHz communications are not limited for WLANs. 60 GHz communications were previously used in datacenters [26, 27] and are proposed to be used as wireless backbone links in next generation cellular or home networks [55]. This gives the opportunity of having dynamic base stations which can be moved when ever is needed. These agile base stations are also proposed to be mounted on drones for particular circumstances, such as disaster recovery mapping or other high-data rate mapping uses [56, 57]. The use of mmWave links for aerial communication presents a key technical challenge as the mmWave link is susceptible to break easily with the Unmanned Aerial Vehicle (UAV) movement. While directionality and link ruptures are issues in cellular and LAN applications, the problems are more pronounced in the aerial case due to the high motion, large angular area to cover and the power and payload limitations on the drone.

As a result, the following issues remain unanswered:

- A. **High data rates.** Which coding and modulation schemes are feasible? What is the impact of data aggregation at such high physical layer rates?
- B. **Directional communication.** How directional are consumer-grade phased antenna arrays? How large is the actual impact of side lobes?
- C. **Area-wide coverage.** What is the performance of 60 GHz devices in terms of coverage within a large room for different device orientations?
- D. **High spatial reuse.** How close can devices operate without experiencing collisions? How strong is the impact of interfering reflections and beam pattern selection?
- E. **Robustness under mobility.** How frequent are link losses? Can nodes move out of the omni-directional discovery range of a base station due to the use of directional beam patterns for data transmission?
- F. **Beamtracking under mobility.** How well can existing devices track user movement and adapt their beamsteering accordingly? How beneficial would it be if devices had multiple antenna arrays, thus avoiding shadowing by the device itself?
- G. **Transient blockage.** How does a temporary blockage affect the link? How does the speed affect related to beam pattern changes?
- H. **mmWave COTS in aerial links.** Is it feasible to mount mmWave COTS radios into drones to provide cellular connectivity? How far can these links reach? Are these links going to work under strong wind conditions?

Here, we provide answers to the questions above. To this end, we use a down-converter to overhear and analyze the communication of 60 GHz devices. This gives us unprecedented

insights into link utilization, beam patterns, and frame level operation. Moreover, we integrate this platform with a custom hardware monitor to access firmware registers that reveal lower-layer details. Our analysis allows us to determine the key limitations of 60 GHz consumer-grade hardware, which opens the door to future work on mechanisms that address these limitations. Most importantly, we make part of our measurements available to the community [24] in a format which allows other researchers to build on them. This is highly relevant, as it enables realistic and comparable results. This chapter contributions are as follows:

1. We provide a frame level analysis of 60 GHz protocols, studying parameters such as frame length and showing that existing devices only use data aggregation if a connection requires high throughput. Otherwise, they do not aggregate, even for traffic that is not delay-constrained, thus resulting in inefficient medium usage.
2. We measure the beam patterns that our devices use. We show that the patterns suffer significant imperfections, such as high irregularity and very strong side lobes of up to -1 dB compared to the main lobe.
3. We perform a comprehensive measurement campaign studying a grid of 117 measured positions with eight equally spaced device orientations in a large conference room in order to obtain the coverage capabilities of COTS devices.
4. We evaluate the impact of interference due to side lobes and reflections. We find that interfering reflections from neighboring, unaligned devices may reduce the achievable Transmission Control Protocol (TCP) throughput by more than 20% of the value that would be achieved otherwise.
5. We study the performance of commercial 60 GHz devices under movement, showing that throughput decreases on average 30% compared to the static case.
6. We show that additional arrays on a device can improve throughput up to $4\times$ and reduce outage by 62%. Moreover, we observe that nodes can easily move out of the discovery range of each other during on-going communication due to the directional beam patterns.
7. We analyze how a transient blockage affects the link. We test for different speeds of the blocking body how the link gets affected in terms of number of tested beam patterns during the blockage.
8. We test the performance of aerial mmWave links with COTS radios for different heights and distances. We then validate these aerial links emulating windy conditions.

The chapter is structured as follows. Section 3.1 provides details on the tested devices and our setup. In Section 3.2 and 3.3, we present our static and mobile analysis, respectively. Section 3.4 introduces the feasibility of ground-to-air communication based on COTS devices. Next, in Section 3.5 we discuss the insights that we gain from our results.

3.1. Experiment Setup

In the following, we present the techniques that we use to obtain the 60 GHz networking traces. This goes well beyond simple measurements, since it includes partial reverse-engineering of the commercial systems that we consider, as well as filtering and processing the data for reuse by other researchers.

3.1.1. Devices Under Test

We consider four different 60 GHz commercial devices (c.f. Figures 3.1 and 3.23). Our first device under test is the Dell D5000 wireless docking station. The D5000 dock is widely used in practical 60 GHz research since it was the first commercial device that largely implemented the mechanisms that later became the IEEE 802.11ad standard. It serves as a docking station for laptops with a compatible WiGig card. Broadly speaking, the dock acts as an AP, and the connected laptop as a station. This system uses Wilocity Wi16100 “Marlon” cards, which are the precursor model of the cards used in 60 GHz routers at the time of writing. Most interestingly, these devices allow for access to lower layer information as will be described in next Section 3.1.2 such as the beam pattern in use, control and data packet counters and the automatic gain control level, and are thus highly valuable for our analysis. After disassembling the D5000 dock, we found that it has a baseband chip connected to an upconverter and a 2x8 antenna array. The antenna covers a cone of 120° .

Our second device is a WiHD-compatible DVDO Air-3c system designed for transmitting HDMI data streams. The system consists of a transmitter and a receiver module that do not allow for any configuration. In terms of range and link stability, we found that the Air-3c performs better than the D5000 dock. Indoors, we could transmit video over 20 meters, even with misalignment and blockage on the direct path. Upon disassembly, we found on both sides of the link an antenna




Figure 3.1: Experiment hardware: (a) Talon router, (b) Dell D5000 dock, (c) DVDO Air-3c system, (d) measurement robot with Dell WiGig laptop, and (e) Acer IEEE 802.11ad P446-M laptop. The arrows mark the placement of the 60 GHz antennas on the devices.

array with 24 elements placed irregularly.

Our third device is the TP-Link Talon AD7200, which is the first commercial router operating in the 60 GHz band. It uses the Qualcomm QCA9500 chipset which comes with a 32 element antenna array. We observe that its behavior is similar to the D5000 dock. Our lab measurements reveal that both devices use very similar chipsets. Due to these similarities and the aforementioned advantage of the D5000 docking station where we can access lower layer information, we are going to focus our study on the D5000 dock with the exception of Section 3.2.5 and 3.4.

For our drone communication experiments, we use a pair of Acer Travelmate P446-M laptops to establish the 60 GHz communication. These laptops are equipped with the client-version of the Qualcomm QCA9008 chipset used in the TP-Link Talon AD7200, which in the same way, is connected to a 32 element phased-antenna array. On the time in which these drone experiments where performed, the community advanced providing the opportunity of using the open source wil6210 wireless driver [58], which provides access to lower layer parameters of the chipset. The bandwidth of all of our four devices under test is 2.16GHz, from which only 1.7 GHz is used leaving the rest as guard. They all operate on the channels defined in the 802.11ad standard. That is, their center frequencies range from 58.32 GHz to 64.80 GHz.

For the aerial experiments in Section 3.4 we use a DJI MATRICE 600 PRO drone. This drone is used to fly the COTS mmWave radio. It is capable of lifting a payload of 5 kg for approximately 20 minutes, which provides enough time to acquire the measurement traces discussed in the next section. It further provides all the GPS and IMU data of the flights for better control of the experimental setup. The aerial setup can be seen in Figure 3.23.

3.1.2. Lower Layer Parameter Collection

Apparently, the Dell drivers do not provide any access to lower layer information except for the physical layer rate in the Graphical User Interface (GUI) available to the user, which allows researchers to deduce the current MCS. However, the driver ships with a so-called “Wilocity Monitor”, which runs continuously as a Windows service. This monitor is clearly not meant for usage by the end user, and seems to be a residual of debugging activities. It is configurable via a number of readily accessible text files. Moreover, it stores the monitored physical and Medium Access Control (MAC) layer parameters to a file of Comma-Separated Values (CSV) at configurable intervals. Figure 3.2 shows an overview of the monitor and the involved files.




Figure 3.2: Overview of the Wilocity Monitor.

Monitor Operation. As shown in Figure 3.2, the monitor takes two files as input. First, the monitor configuration file specifies the interval at which lower layer parameters of the Wilocity card shall be logged. The minimum interval is 1 ms, which provides a much higher time resolution than the information shown on the device driver GUI. Among others, this file also specifies for how many days logs shall be stored before deletion, and whether the monitor shall perform a memory dump when the service is started. Second, the Register File (RGF) configuration file specifies which lower layer parameters the monitor shall log. The monitor reads any value out of a 712 KB memory area of the Wilocity card. To monitor a value in the memory, the RGF configuration must specify (a) the memory address and (b) the bit mask needed to extract the value from the memory location. The mask selects specific bits out of the memory contents. Based on the above configuration, the Wilocity Monitor generates two output files. First, it stores a memory dump of the full 712 KB memory at startup. Second, it logs the monitored values in the RGF file at the specified monitoring interval. These output files are located in a hidden system folder of the operating system.

Lower Layer Parameters. While the Wilocity Monitor provides access to a large memory area containing all kinds of lower layer information, the key problem is to know which combination of memory address and bitwise mask contains which information. From the configuration files of the monitor, it becomes evident that the design of the monitor assumes an additional CSV file which contains a mapping of the memory to human readable mnemonics. Unsurprisingly, this file is not present among the files that ship with the device driver. However, the RGF configuration file contains a number of pre-configured values which, once again, seem to be a residual of debugging activities. In Table 3.1 we show some of these values. The “INA” counters count the number of packets that the physical layer successfully detects based on the frame preamble. Further, the “CRC OK” counters show how many of the detected packets passed the Cyclic Redundancy Check (CRC). We also obtain temperature information of Baseband and Radio Chips, which may hint at energy consumption. The Automatic Gain Control (AGC) of the receiver is indicated as an attenuation, that is, the higher the logged value, the lower is the gain at the receiver. Finally, Table 3.1 includes two values which we have inferred based on measurements, namely, the MCS and the beam patterns from the laptop and docking station.

3.1.3. Comprehensive Trace Collection

To collect physical layer traces for our dataset and our frame level analysis, we use a VubiQ 60 GHz Development System [59] connected to an Agilent MSO-X 3034A oscilloscope. This setup provides traces of the analog I/Q output of the VubiQ receiver, and thus allows us to observe the individual frames on the medium. While we undersample the signal and thus cannot decode it, the traces allow us to extract the timing and amplitude of frames by processing them offline in Matlab. The VubiQ frontend supports downconversion of 1.8 GHz modulated bandwidth at the common IEEE 802.11ad/aj frequencies. Further, it has a WR-15 input, which we connect to horn antennas with different levels of directivity. For beam pattern measurements, we use a 25 dBi

Table 3.1: Monitored Lower Layer Parameters

Description	Mask Size
INA Counter Control Packets	32 bits
INA Counter Data Packets	32 bits
CRC OK Counter Control Packets	32 bits
CRC OK Counter Data Packets	32 bits
Baseband Chip Temperature	32 bits
Radio Chip Temperature	32 bits
Automatic Gain Control Attenuation	6 bits
Modulation and Coding Scheme	5 bits
Transmit Beam Pattern	6 bits
Receive Beam Pattern	6 bits

horn antenna. To obtain a wide beam pattern for protocol analysis, we use the open wave guide.

Further, we obtain transport layer throughput measurements for the D5000 dock using ‘iPerf’ [60]. Specifically, we run the iPerf server on the E7440 laptop, and connect a second laptop via Ethernet to the D5000 dock. We run the iPerf client on this second laptop, thus generating traffic on the wireless 60 GHz link. This iPerf session together with our reverse-engineer Wilocity Monitor Service on the D5000 dock, reveals the desired lower layer information. Altogether, the above setup provides detailed information of the D5000 dock at the physical, medium access control, and transport layers.

Finally, for our comprehensive performance measurements in indoor environments, we use a robotic platform that obtains all of the aforementioned traces in an automated manner on a specific set of locations within a room. In particular, we use a Kobuki Turtlebot II robot to move the E7440 laptop along a grid of points within a large conference room. The robot uses a camera to follow a line marked on the floor and stops at fixed intervals to measure performance. At each location, it repeats this measurement for eight different rotation angles, covering an overall angular range of 360°. Except for our mobility analysis, we perform all measurements while the robot is static. We place the D5000 docking station next to a wall but pointing towards the room to recreate a realistic location of an AP. As a result, we obtain a comprehensive measurement database for the entire room area and virtually any orientation of the laptop. This gives us unprecedented insights into the practical performance of 60 GHz networks, since earlier work is limited to individual locations within a room.

3.2. Static Analysis

In this section we describe the results of our static analysis. We first present our findings on protocol operation and aggregation. Second, we analyze the beamforming of the D5000 dock. Third, we investigate the impact of reflections. Fourth, we study the coverage of the D5000 in a wide area. Finally, we study the spatial sharing that these devices allow for different orientations.

3.2.1. Protocol Analysis

Measurement Setup. To gain insights into the protocol operation of our 60 GHz devices, we use the VubiQ receiver as described in Section 3.1. To overhear the frames of both the transmitter and the receiver, we use the open wave-guide of the VubiQ system. To identify which frames come from which device, we place the VubiQ down-converter behind the docking station and point it towards the notebook's lid. As a result, the down-converter receives the frames of the notebook via a direct path, and the frames of the dock via a reflection from the lid. Thus, the average amplitude of the notebook frames is larger than the one of the docking station frames, and we can easily separate them. For the wireless HDMI case, no reflector setup is needed. Since its directivity is lower, the VubiQ can easily receive each side with a different power level from a single location. We use iPerf in the Dell D5000 dock case to generate TCP traffic. This traffic is transmitted from the laptop to the docking station via the 60 GHz wireless link. For the wireless HDMI case, we transmit a high definition video signal. We use Matlab to post-process the traces, and obtain insights into frame and burst lengths. The latter allows us to analyze data aggregation.


Beacon and Discovery Frames. All of our 60 GHz devices transmit periodic control frames. Specifically, they transmit discovery frames to establish new links, and beacon frames to maintain an existing connection. Table 3.2 shows their periodicity. We observe that systems transmitting real-time video signals, such as the Air-3c, send beacon frames much more frequently than the Talon router. This allows them to maintain a very high quality of service in case of link impairments. Since the Talon router is not specifically designed for streaming, the beacon interval is larger.

Dell D5000 Dock. We observe that the communication of the Dell D5000 dock consists of three phases, namely, device discovery, link setup, and data transmission. These three phases match the Beacon Transmission Interval (BTI), the Association Beamforming Training (A-BFT), and the Data Transmission Interval (DTI) of IEEE 802.11ad, respectively. In the first stage, the D5000 dock transmits the aforementioned discovery frames. Figure 3.3(a) shows an example of such a frame. In the second stage, a complex association and beamforming process takes place. Our frame level analysis of the third stage shows that data transmission is divided into frame bursts. These bursts are equivalent to Beacon Interval (BI) in IEEE 802.11ad, and their maximum length is 2 ms. Each burst begins with two control frames, which have a different amplitude than the subsequent series of data and acknowledgment frames, as shown in Figure 3.3(b). Outside the bursts, the channel is idle except for the regular beacon exchange between the docking station and the notebook (c.f. Table 3.2).


In Figure 3.4, we show the evolution of lower layer parameters for a total of four hours in a static environment. During the experiment, we start iPerf twice for about 20 minutes. As expected, the received packets per second increase to a high value while iPerf transmits, and are close to zero otherwise, as shown in Figure 3.4(a). Interestingly, Figure 3.4(b) shows that the number of received control packets per second decreases during those periods. The underlying reason is that the Wilocity radio transmits frequent beacons to maintain connectivity while the link

Table 3.2: D5000 dock, Talon, and WiHD frame periodicity

Frame type	Periodicity
D5000/Talon Discovery Frame	102.4 ms
WiHD Discovery Frame	20 ms
D5000 Beacon Frame	1.1 ms
Talon Beacon Frame	102.4 ms
WiHD Beacon Frame	0.224 ms



(a) Dell D5000 dock device discovery frame.



(b) Dell D5000 dock frame flow.

Figure 3.3: Frame structure of the Dell D5000 docking station system.

is idle, which is unnecessary when the link is active anyhow. The AGC attenuation in Figure 3.4(c) is also closely related to the data transmission—while no data is being transmitted, the AGC attenuation increases since no signal amplification is needed. The AGC value that we show in Figure 3.4(c) is only related to directional data transmissions, that is, it does not take into account omni-directional control traffic. Also, the monitor logs a dimensionless value, which is hard to relate to an actual attenuation value in dB. Thus, we show the measurement result in “attenuation units”. The same is valid for the temperature of the radio chip in Figure 3.4(d). Still, we clearly observe how the temperature increases during data transmissions.

We also investigate the impact of data aggregation. Figure 3.5 gives an intuition on the relevance of data aggregation at very high rates such as in IEEE 802.11ad. To study this issue, we measure the length of data frames for different TCP throughput values. We control the TCP throughput by adjusting its window size in iPerf. Figure 3.6(a) depicts the Cumulative Distribution Function (CDF) of the frame lengths for each throughput value. The CDF reveals

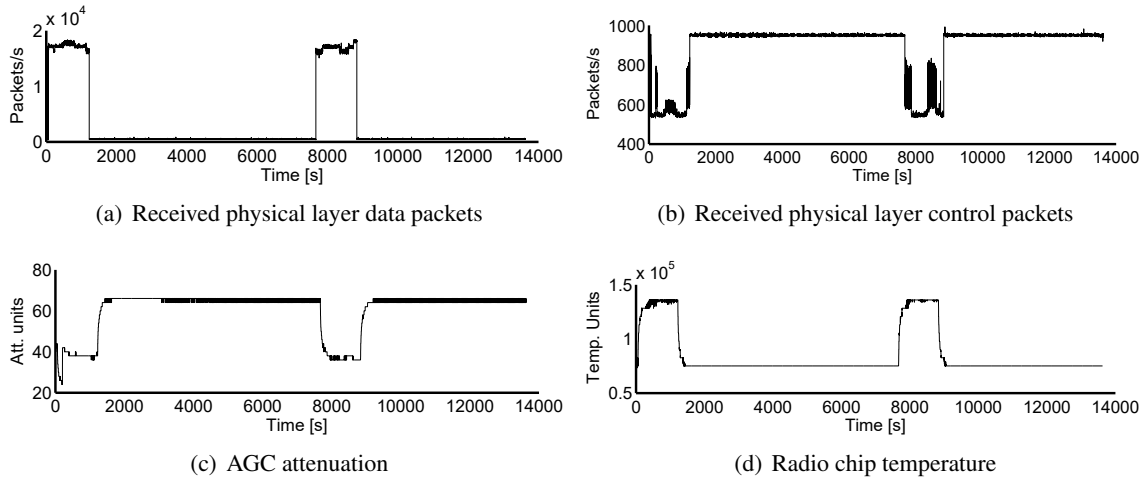


Figure 3.4: Static analysis for fixed link length.






Figure 3.5: Aggregation primer. Data bursts begin with a pair of beacons (B) which is followed by a data (d_i) and ack (a) exchange. When aggregation occurs, several data packets are combined into one frame. Light brown color represents the preamble and header of each of the data frames, while dark gray represents the payload.

that most frames are either short (around $5\ \mu s$) or long (15 to $20\ \mu s$). That is, we can divide them into two categories. The length of long frames varies more than the length of short frames, which is due to different levels of aggregation. The highest level we observe corresponds to a frame duration of $25\ \mu s$. Further, the amount of long frames increases with throughput—the higher the traffic load, the higher the level of data aggregation.



(a) CDF of the length of WiGig frames for increasing traffic load.




(b) WiGig medium usage for increasing traffic load.


Figure 3.6: Frame level analysis of a WiGig link operating at 60 GHz.

Moreover, we investigate the level of medium usage for increasing throughput values. Surprisingly, Figure 3.6(b) shows that, beyond a relatively low throughput value¹, all oscilloscope traces contained data frames. That is, the transmitter transmitted continuously. We observed that this throughput increase at constant medium usage is due to a higher level of frame aggregation. That is, WiGig can scale throughput from 171 Mbps to 934 Mbps by simply increasing the amount of data in each frame, and thus vastly reducing the overhead due to medium access. According to Figure 3.6(a), WiGig only needs to aggregate 25 μ s of data to achieve this massive improvement. For comparison, 802.11ac must aggregate 8 ms of data to achieve only a 2 \times improvement [61]. This is due to the much larger bandwidth of IEEE 802.11ad compared to IEEE 802.11ac. Any idle time due to MAC overhead translates into huge throughput losses. In addition, we verified that the transmitter used the same MCS for all traffic loads in Figure 3.6(b). That is, the throughput increase is not due to the use of a higher MCS. This is expected, since the MCS is typically not

¹We choose values in the order of kilobits per second by setting a small iPerf TCP window size to observe low link utilization.



(a) MCS with low traffic.



(b) Throughput decrease with distance.

Figure 3.7: Impact of distance on link rate and transport layer throughput.

adapted to the link load but to the link quality.

Figure 3.7(a) shows the impact of the link length on the physical layer rate. The longer the link, the lower and more unstable the data rate. While Figure 3.7(a) depicts the link rate, in Figure 3.7(b) we show the average iPerf throughput at increasing distances. We observe that the throughput is approximately stable up to a certain distance d and then falls abruptly. This distance d varies significantly for different experiments. Thus, while the individual experiments exhibit the aforementioned abrupt behavior, the average falls gradually. We do not observe results beyond 900 Mbps because the Gigabit Ethernet interface at the docking station limits the achievable throughput. This is also the case for the Talon router. We conclude that the MCS drops gradually with distance to about 1 Gbps but then falls abruptly. In our experiments, the range of the Talon router was even higher, reaching up to 40 meters in an indoor corridor scenario. Our measurements suggest that such long links are feasible due to constructive interference from reflections.

DVDO Air-3c. Next, we analyze the frame flow for the DVDO Air-3c WiHD system. We observe the same communication stages as for the D5000 dock. When analyzing the data transmission stage, we found the beacon frame frequency to be much higher than for the D5000 system.




Figure 3.8: DVDO Air-3c WiHD frame flow.

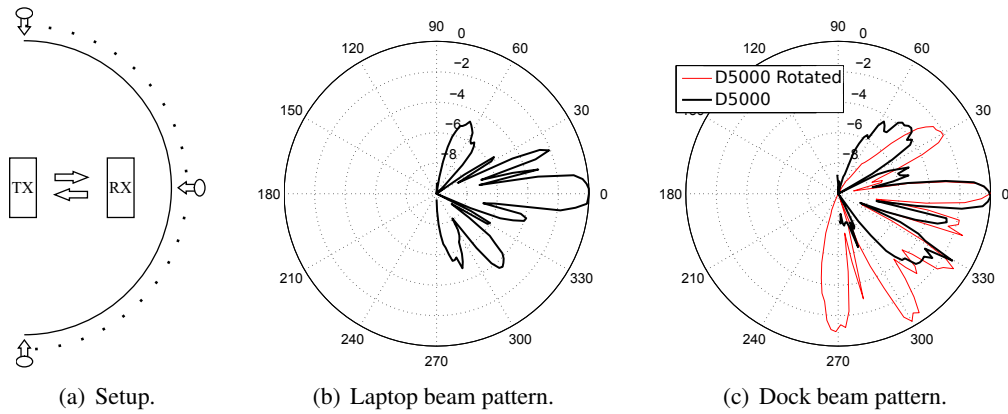


Figure 3.9: Measurement setup and beam patterns of the docking station system. Main lobe points always to zero degrees. The beam patterns are scaled such that the gain of the main lobe is at 0 dB.

Figure 3.8 shows an example frame flow for the WiHD system. Both the transmitter and the receiver emit periodic beacons in a request-response manner to maintain the link. Each pair of beacons is followed by a data packet when the transmitter has data to transmit. When no data is queued for transmission, we only observe beacon frames. The trace shows the transition from an active video transmission to an idle period. The WiHD system does not seem to perform channel sensing, which may cause significant interference.

3.2.2. Beam patterns

Measurement Setup. To measure the antenna beam patterns of our devices, we use the VubIQ system with a highly directional horn antenna. Specifically, we set up a 60 GHz link as in Figure 3.9(a) and move the VubIQ receiver along 100 equally spaced positions on a semicircle of radius 3.2 m centered at the transmitter. To avoid unwanted reflections, we perform the experiment in a large outdoor space. In post-processing, we extract the signal strength of the data frames transmitted by the device under test. We discard periodic control frames, which are transmitted with higher power and wider antenna patterns. As a result, we obtain the azimuthal plane of the beam patterns. Moreover, we also exploit the device discovery frame of the D5000 dock shown in Figure 3.3(a) to compute the beam patterns that the D5000 dock uses to discover nodes.

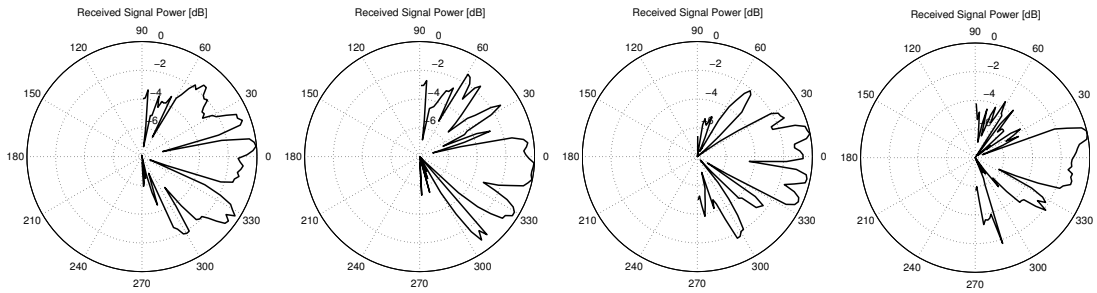


Figure 3.10: A selection of the 32 quasi omni-directional beam patterns swept by the D5000 dock during the device discovery phase.

Directional Transmission and Side Lobes. We first evaluate the beam patterns used by the D5000 dock during the data transmission stage. In contrast to the device discovery stage, the D5000 dock uses highly directional antenna configurations in order to use high MCS values. Figures 3.9(b) and 3.9(c) show the measured transmit patterns of the Dell E7440 notebook and the D5000 docking station. As expected, the patterns are highly directional with a Half-Power Beam Width (HPBW) below 20 degrees. Despite the strong signal focus, we measured significant signal energy from side lobes. These side lobes can have a transmit power in the range of -4 to -6 dB compared to the main lobe, and thus can cause major interference. In a second measurement, we rotated the notebook by 20° with respect to the laptop. Figure 3.9(c) shows the resulting docking station beam pattern as an overlay. We observe that the gain of the main lobe is 10 dB below the gain that we measure when the devices face each other. That is, beamforming towards the boundary transmission area of the antenna array significantly reduces link gain. Also, we observe a much higher number of side lobes as strong as -1 dB with respect to the main lobe.

Quasi Omni-Directional Search. Implementing omni-directional antenna patterns is a major challenge for mmWave communication [62]. However, this kind of pattern is crucial for device discovery and beamtraining [22]. The Dell D5000 system sweeps 32 different quasi omni-directional patterns during device discovery (c.f. Figure 3.3(a)). Figure 3.10 shows four out of these 32 beam patterns. The irregularities of the beam patterns are a natural result of our measurement setup—since we manually move the VubIQ receiver along the 100 measurement positions, small deviations are inevitable. Nevertheless, Figure 3.10 clearly depicts the rough shape of the lobes. While the HPBW can be as wide as 60° , each pattern contains several deep gaps that may prevent communication with devices at the corresponding angles. These gaps are due to the limitations of consumer-grade phased antenna arrays. However, probing multiple such patterns during device discovery mitigates the impact of such gaps, and allows the dock to cover the entire azimuth. The remaining 28 patterns are comparable in terms of directional focus and received signal power.

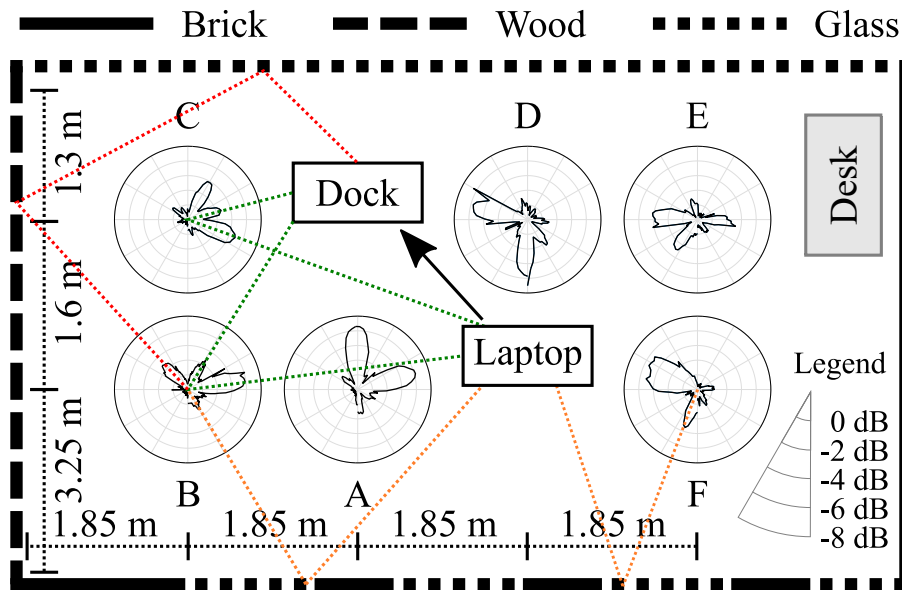


Figure 3.11: Reflection patterns in conference room. The dotted lines show examples of LOS paths as well as first and second order reflections.

3.2.3. Reflections

Reflection Analysis. Next, we analyze the impact of reflections in a realistic wireless setting. This addresses the common assumption that the 60 GHz multi-path environment is very sparse compared to the 2.4/5 GHz case, and that reflections result from quasi-optical propagation in the direction of transmission. To this end, we set up a single 60 GHz link in an empty conference room, either using the D5000 dock or the WiHD system. We then measure the energy received from all possible directions at six different locations $\{A \dots F\}$ in the room, as shown in Figure 3.11. To this end, we mount the VubIQ receiver with a highly directional horn antenna on a programmable rotation device and place it at each of the six locations. At each location, we then measure the incident signal strength for all directions and assemble the result to an angular profile. If no reflections occur, we expect to receive energy only from the direction in which the 60 GHz devices are located. For instance, at location A , we should only observe energy coming horizontally from the right and from above. Additional lobes in the angular profile indicate reflections. To analyze the impact of different materials, we perform the experiments in a room which has brick, glass, and wood walls. Figure 3.11 shows the layout.

Figure 3.11 depicts the angular profiles that we measure for the D5000 dock. In both cases, most angular patterns have at least two clearly identifiable lobes, each pointing to one side of the 60 GHz link. We observe energy from both ends of the link because the receiving end not only receives data frames but also transmits the corresponding acknowledgments. A significant number of angular patterns feature additional lobes that signal reflections off the walls. In contrast to common assumptions regarding 60 GHz communications, the lobes show a significant amount of




Figure 3.12: Data transmission exclusively via a reflected path.

incidence energy. For instance, the angular pattern at position F in Figure 3.11 has a lobe directly pointing to the lower wall. Geometrically following the reflection of the signal off the nearby window suggests that this lobe is due to the transmissions of the laptop. Further, the angular pattern at position B features a lobe pointing towards the wooden wall. This lobe arises from a second order reflection that originates at the D5000 dock, and bounces off both the glass wall as well as the wooden wall. This second order reflection has a signal strength that is comparable to the LOS path, because the gain from the antenna array in the direction of the reflection is higher than in the direction of the LOS. This compensates the additional path loss in the reflected case. This finding is highly relevant since current mmWave networking models often assume that the strong attenuation at 60 GHz results in at most first order reflections. While having second order reflections translates into better coverage, they also cause more interference. Thus, including this effect in the evaluation of 60 GHz systems is crucial.

Communication via Reflections. While the above results show that reflections are significant, we also perform a case study to analyze to what degree those reflections help to extend the coverage of a network if the LOS is blocked. To this end, we set up a link parallel to a glass wall using a D5000 dock and its corresponding laptop. Additionally, we place an obstacle in between both. Figure 3.12 depicts our setup. We then measure the angular energy profile at the receiver to verify that the LOS path is actually blocked, and that all energy arrives via the reflection off the wall. Figure 3.12 validates this—the angular energy profile does not include any lobe on the LOS. Finally, we use iPerf to measure the achievable rate on such a reflection. We obtain 550 Mbps (± 18 Mbps with 95% confidence) which is more than half of what we measure on LOS links. Thus, communication via reflections is feasible and achieves significant data rates. Earlier work [25] shows similar findings but lacks validation via an angular profile.

Interference from Reflections. The above reflection analysis allows us to assess the existence and strength of reflections. In addition, we use a second setup to determine the impact of those reflections on data transmissions. In particular, we set up two geometrically non-interfering

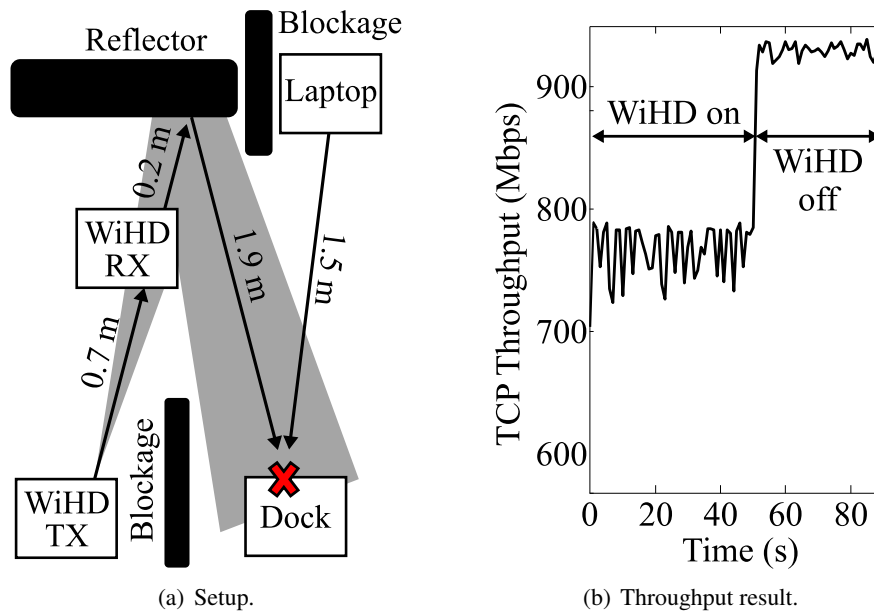


Figure 3.13: Reflection interference experiment.

60 GHz links close to a metal reflector, as shown in Figure 3.13. To eliminate the influence of side lobes on the measurement, we position shielding elements close to the WiGig devices. Further, we make sure that we do not block the reflected signal resulting from the metallic surface behind the WiHD receiver. We then analyze the coverage area of this reflection using the VubIQ transceiver to ensure that the docking station is located inside. Finally, we perform a TCP throughput measurement from the laptop to the docking station using iPerf. By switching on and off the WiHD devices, we evaluate the impact of the reflection on the WiGig connection. For the iPerf connection, we set a TCP window size of 250 KByte such that it fully loads the underlying Ethernet link, which the docking station tunnels over the 60 GHz wireless link. The WiGig tunneling manages to provide full gigabit speed for the Ethernet link. This almost completely saturates the link because the D5000 system tries to minimize the delay of the Ethernet traffic. That is, instead of aggregating data to reduce the medium usage, the transmitter sends a larger number of packets. Due to this high medium usage, we expect the link to be very sensitive to interference effects.

The results in Figure 3.13(b) show that interference indeed has a strong impact in our setup. We observe that the TCP throughput increases significantly after about 90 seconds, which is when we power off the WiHD link. The performance degradation due to the WiHD reflection is about 200 Mbps compared to the interference-free transmission. Further, we observe that the throughput fluctuates strongly for the case with interference. Most probably, these variations result from attempts of the D5000 dock to adapt to the high interference by switching among MCS levels. The average throughput reduction is about 20%, but reaches up to 33%.

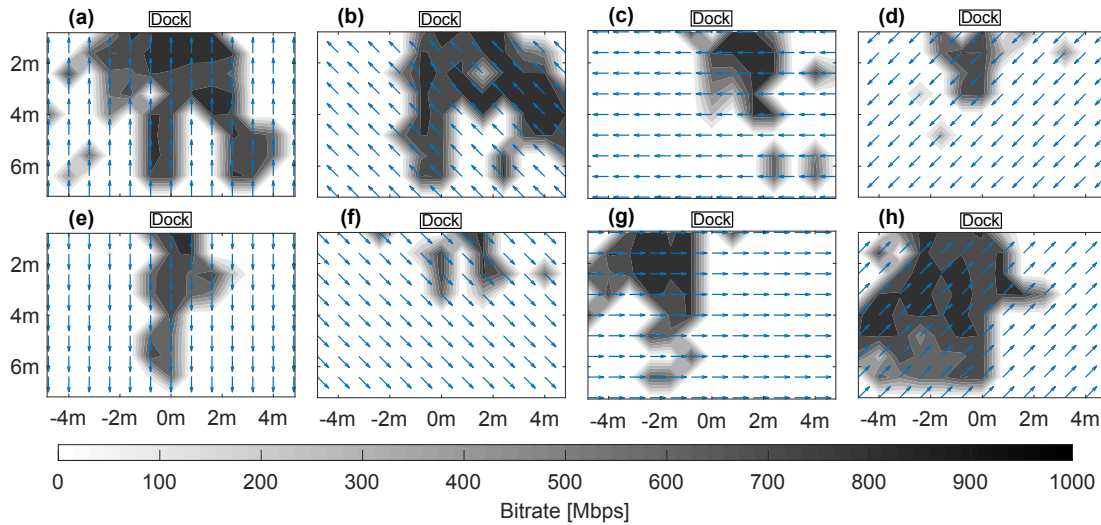


Figure 3.14: Transport layer throughput for our area-wide grid measurement with the D5000 dock. Dark areas indicate high throughput, whereas light areas represent low throughput. The arrows indicate the orientation of the E7400 laptop on the robot. The D5000 docking station is located at coordinate $(0, 0)$, as indicated on the upper part of each of the subfigures (a) to (h).

3.2.4. Area-Wide Coverage

While our above analysis focuses on specific scenarios, in the following we consider a comprehensive measurement campaign within a large conference room.

Measurement Setup. The size of the room is $21.5 \times 10.5 m^2$. As discussed in Section 3.1, we use a robot carrying the E7440 laptop to perform measurements on the physical, medium access control, and transport layers at a large number of locations and orientations within this room. The robotic automation is key to tackle the complexity of this very large measurement campaign. The measurement locations are the 117 intersections of a 13×9 grid marked on the floor. The intersections are at a distance of 0.8 meters both horizontally and vertically. At each location, the robot repeats the measurement in eight equally spaced directions to cover the full 360° azimuthal range. As a result, we obtain a comprehensive set of 936 measurements. This dataset is highly valuable to the community since it is the first systematic 60 GHz networking trace collection within a room. Researchers can easily use the traces to evaluate the performance of 60 GHz protocols regarding, e.g., rate adaptation, mobility, and handovers in indoor environments. We place the D5000 docking station at grid intersection $(0, 0)$, which we define at the midpoint of the top wall of the room, as shown in Figure 3.14. We collect a wide range of information including MCS, PER, and the identifier of the beam pattern in use.

Coverage Results. Figure 3.14 depicts the throughput results as a heatmap. As expected, the WiGig link achieves the highest throughput when the E7440 laptop is close to the D5000 dock.

However, even when the E7440 laptop is facing the D5000 dock, throughput is not symmetric across the otherwise empty room. Still, when the laptop is facing the docking station at $+45^\circ$ or -45° , such as in Figures 3.14(b) and 3.14(h), the coverage area is clearly shaped towards the right and the left half of the room, respectively. Interestingly, we still achieve significant throughput when the laptop is facing away from the D5000 dock, as depicted in Figure 3.14(e). This shows that the antenna radiates a significant amount of energy through a back lobe. Given the location of the antenna in the laptop lid, under regular usage this back lobe points towards the person using the laptop and thus may not be usable for communication. All in all, we conclude that the coverage of commercial 60 GHz hardware is highly irregular, and thus strongly differs from theoretical models. This behavior is an immediate result of the irregular shape of the beam patterns that we observe in Section 3.2.2.

3.2.5. Spatial Sharing

Measurement Setup. In order to study how regularly mmWave COTS devices share the medium given their large side lobes, we establish two sets of experiments to analyze parallel links. For these experiments, we use the TP-Link IEEE 802.11ad Talon routers, which can be used both as AP or as STA. With this devices, each link is capable of achieving between 1.4 Gbps and 1.7 Gbps of throughput during standalone operation depending on the environment, the selected MCS and channel state. The separation between the APs and the STA is of 2 m. First, we study the impact of operating two parallel links for different separation distances, seeing how for larger distances the aggregated throughput increases. Second, we study the impact of link orientation when sharing the wireless medium. The two links are separated with a distance of 15 m among them. In this measurement campaign, each AP-STA pair is placed either in a horizontal or vertical orientation, keeping maintaining a separation of 15 m between the medial of each link (leading to cases where the actual distance between the two closest interfering stations is equal to 13 m). In each setup, we have two parallel 60 GHz links where AP and STA are separated by 2 m and each STA starts a TCP session towards its respective AP (uplink scenario).

Link Separation. Figure 3.15 depicts the aggregated throughput for the two parallel links with respect to separation distance with a confidence interval of 95%. As expected, a higher separation distance results in a higher aggregated throughput for the different placements. Conversely, the smaller the separation distance, the higher the probability that the two pairs can sense each other and, as a result, start sharing the wireless medium. When the separation distance is larger than 18 m, each STA can achieve around 1.4 Gbps, the same as in standalone operation, which implies that they achieve full spatial reuse. For shorter separation distances, the two links start interfering with each other via their side lobes: Between 13 m and 19 m, the aggregated throughput decreases by 1.3 Gbps and remains around 1.5 Gbps. For shorter separation distances, the aggregated throughput varies between 1.6 Gbps and 1.2 Gbps.

Link Orientation. Table 3.3 shows the results of this second set of experiments. In the first column we can see the relative position of both links, while in the other three columns we see




Figure 3.15: Spatial sharing performance with respect to separation distance.

Table 3.3: Link orientation related throughput

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

the individual and the aggregated throughputs. The parallel link topology with communication flows in the same direction is the one with highest throughput, with an aggregated data rate of about 2.4 Gbps. This is due to the fact that the STAs do not sense each other as often as in the other scenarios. In this case, the devices are not facing each other making a balanced communication link with low interference and therefore low waiting times among transmissions. However, it can be seen that the data rate is not the same for both machines, with the left link having higher throughput due to the asymmetric beam patterns. The topology of two horizontal links where the communication flows face each other performs the worst in terms of aggregated throughput. This is due to the fact that the two STAs can overhear each others transmissions frequently and thus defer their own transmissions. In between these two, we have the rest of the cases, where depending on the amount of received power from the neighboring nodes, the probability of deferring their own transmission changes, increasing or decreasing each individual link throughput. The communication direction also matters. If the transmitter is the one interfering with the neighboring link, it will interfere more often as it will send large frames,

while the receiver will interfere less as it only sends shorter ACK frames. This effect can be seen in the horizontal link scenario or in the vertical plus horizontal scenario.

3.3. Mobile Analysis

In the following, we analyze the impact of mobility and transient blockage on consumer-grade 60 GHz devices in indoor environments. To this end, we use the comprehensive grid measurement set discussed in Section 3.2.4 as a baseline and compare it to the performance under mobility.

3.3.1. Experiment Design

Since we recorded the robot-assisted grid measurements under static and controlled conditions, they reflect the optimal performance of the link at a certain location and orientation. To ensure this, we reset the devices prior to each measurement, thus forcing them to perform a full beamsweep and find the best possible beam pattern. In contrast, in this Section we move the E7440 laptop at walking speed along a number of different paths on the grid while performing the same measurements as in the static case. Naturally, the performance is lower since the link must continuously adapt to the changing environment at each location. By comparing the static performance at the grid intersections along which we move with the corresponding mobile performance, we can assess the ability of the WiGig/802.11ad protocol to handle mobility. For instance, under mobility the E7440 laptop may not always use the best beam pattern or MCS and thus achieve suboptimal performance in terms of throughput.

We define a number of mobility paths along the grid. Figure 3.16 shows four examples of such trajectories, out of a total dataset of 15 trajectories. The arrows indicate the orientation of the E7440 laptop during the movement. In most cases the arrows follow the movement, such as in Figure 3.16(a). However, we also consider other cases for which the E7440 laptop is always facing the D5000 dock regardless of the direction of movement, such as the trajectory in Figure 3.16(b). This allows us to assess the impact of device rotation while moving. We repeat each trajectory at least three times to avoid measurement artifacts. Moreover, we record each trajectory both forward and backward but always sticking to the orientation of the E7440 laptop as depicted by the arrows in Figure 3.16. This allows us to study, for instance, whether the rate control of the E7440 laptop adapts well to both decreasing and increasing signal quality.

3.3.2. Mobility Results

Our mobility experiments show a significant performance degradation compared to the static case. As a first step, we average the performance of all trajectories compared to the corresponding static measurement data. Table 3.4 gives an overview of our results. We observe that the overall throughput in the mobile case is on average 30% lower than in the static case. At the same time, the number of control packets increases by 31.71%. The increase of control packets are related

Table 3.4: Measurement results: Static case compared to mobile case.

	Static	Mobile	Difference
Throughput	790.7 Mbps	606.8 Mbps	-30.29%
Num. of Control Pkts	603.3 pkts	883.6 pkts	31.71%
Num. of Data Pkts	24310.0 pkts	26780.0 pkts	9.22%
Ratio Control/Total	2.42%	3.19%	31.81%
Control Pkt Error Rate	0.02%	0.36%	94.11%
Data Pkt Error Rate	0.0004%	0.45%	99.92%
Laptop BP Changes	5.875	18.16	67.65%
Laptop BP Values	6	11.44	47.55%
D5000 BP Changes	5.5	19.05	71.13%
D5000 BP Values	5.625	12	53.13%

with the increase of beam pattern changes that are carried both by the laptop and the docking station which on both sides have increased around 70% in number of changes and around 50% the total number of used values. In other words, the WiGig link must invest a significant amount of overhead into maintaining the link while moving, and still suffers a significant throughput loss. Most interestingly, the PER before retransmissions for data traffic increases by three orders of magnitude from virtually zero to close to 0.5%. This shows that the link suffers strong impairments during movement, suggesting that the rate control and the beamtracking algorithms cannot keep up with pedestrian movement. We conclude that future work must focus on mitigating the aforementioned 30% performance drop.

As a next step, we study a number of individual trajectories to highlight particular effects caused by mobility. For instance, in Figure 3.17 we depict a trajectory that starts in the lower-right corner of the empty conference room and ends in front of the docking station located at position (0, 0). Both at the beginning and at the end of the trajectory, the mobile device is roughly pointing towards the docking station. However, at about half of the distance, the device follows a short horizontal trajectory, and thus points in a direction which is parallel to the docking station. In the static case, the impact of this horizontal segment is negligible. As we observe in Figure 3.17, TCP throughput remains at about 800 Mbps throughout the trajectory, and the number of control

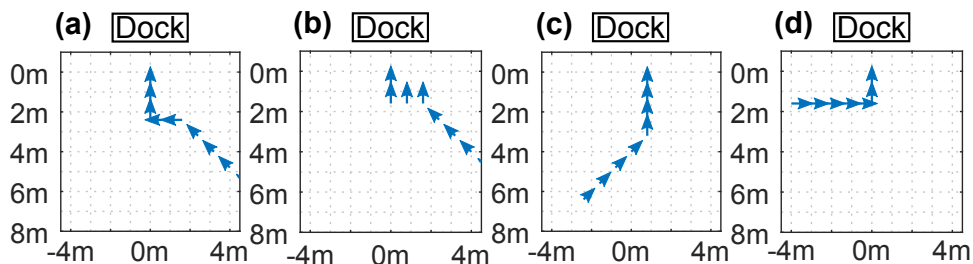


Figure 3.16: Examples of movement trajectories on our measurement grid.




Figure 3.17: Comparison of a trajectory for the static and the mobile case. In the static case, we interpolate the data between the grid locations.

packets is minimal. However, for the mobile case, we clearly see the impact of the horizontal trajectory. The throughput drops to zero for about two seconds and, simultaneously, the number of control packets shoots up to more than 6000 packets per second. That is, as soon as the device enters the horizontal segment, the connection to the docking station drops due to beam misalignment. As a result, the laptop initiates a full sector level sweep which results in the large number of control packets that we observe in Figure 3.17. The static case suggests that the orientation of the device does not hinder a reliable connection. Hence, this experiment shows that the poor performance of the connection is due to the limited beam-tracking and the inefficient beam-sweep of the device. Moreover, we observe that the 30% throughput loss in Table 3.4 is indeed just an average. The instantaneous loss can easily reach 100% as in Figure 3.17. We observed such outage in multiple of our mobility traces.

Figure 3.18 depicts a similar analysis to that of Figure 3.17 but focuses on the MCS. Specifically, the black lines show the MCS index variations for three different trajectories. The higher the MCS index, the higher the physical layer rate is. As a reference, the gray bars show the MCS index for the static case at the grid locations along which the device moves in the mobile case. Note that, while the bars have a certain width for readability, the value of each bar refers to the discrete position value on the x-axis. In contrast, the black lines represent the continuous evolution of the MCS index for the mobile case. The white diamond markers indicate the locations at which the dock, the laptop, or both devices switch beam patterns. Figure 3.18 shows that the static case often operates at higher MCS values than the mobile case. This is expected since beam misalignment is very frequent under mobility. For instance, when moving along the second trajectory in Figure 3.18, the device resorts to the most robust MCS at position five due to a suboptimal beam pattern change and does not recover until about four seconds later when it reaches position nine. During those four seconds, the device performs frequent but unsuccessful beam searches. In the third trajectory, MCS drops are present throughout the trace. However, in this case, we observe that the device continuously attempts to switch back to a higher MCS. Sometimes it succeeds for a brief interval but then immediately must fall back to lower MCSs. We also observe that MCS changes are often correlated with beam pattern changes. This

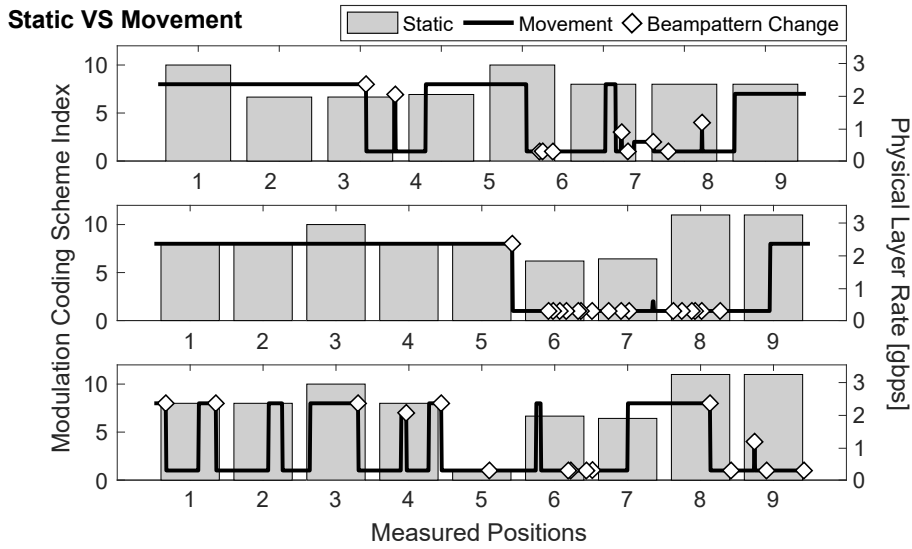


Figure 3.18: Evolution of the MCS index for three different trajectories. The gray bars represent the discrete measurements of the static grid case. The white diamonds indicate beam pattern changes in the mobile case.




Figure 3.19: Match of metrics in the static case compared to the mobile case. The CDF shows the distribution of the difference of the static MCS and the mobile MCS for all of our trajectories in Figure 3.16.

re-confirms that the impact of mobility that we analyzed at the transport layer in Figure 3.17 stems from the fluctuations at the physical layer. Interestingly, the first trajectory in Figure 3.18 depicts cases where the MCS chosen under mobility is slightly higher than in the static case. That is, the continuously changing channel conditions of the mobile case can trigger the rate adaptation mechanism to behave more aggressively than in the static case, which increases the PER.

In Figure 3.19, we perform a similar comparison as in Figure 3.18 but for all of our trajectories. In addition to the MCS index, we also consider the beam pattern index. To this end, we perform a piecewise constant interpolation of the static grid data to obtain a continuous MCS and beam pattern value along each trajectory. We then compute the percentage of each trajectory during which the static and the mobile case result in different MCS and beam pattern indexes. In other words, we quantify how much the behavior of the devices deviates from the optimal static behavior as a result of mobility. We observe that the difference in terms of percentage is

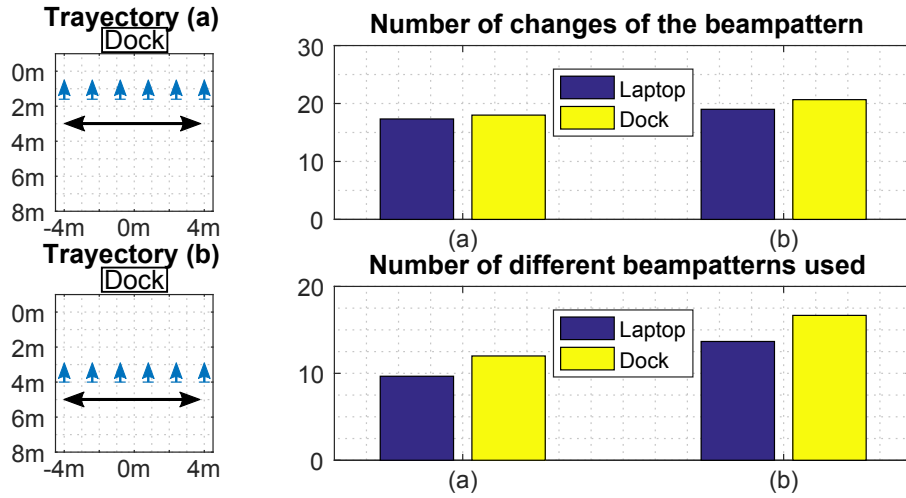


Figure 3.20: Beampattern fluctuations for two trajectories. In both cases, the laptop moves along a horizontal path, but always faces the docking station as indicated by the short arrows in the two figures on the left.

rather high. For the MCS index, the difference of static and mobile is about 70%. While the second trajectory in Figure 3.18 shows that for some specific cases the difference can be lower, generally this does not hold. Moreover, the CDF in Figure 3.19 shows that the difference is beyond two MCS index steps for half of the cases. Regarding the beam pattern index, we observe that the laptop selects a suboptimal sector roughly 70% of the time. For the docking station, this value increases up to about 80%. The reason for this slight difference between the laptop and the docking station is unclear, but is likely due to the specific implementation details of the beam-tracking and beam-searching mechanisms of these devices. While the mismatch in terms of beam pattern index is significant, it is difficult to assess the exact impact on the communication. Essentially, the devices only reveal the index number of the current beam pattern, but not to which beamshape each index number refers to. From our measurements, we concluded that successive index values do not translate into adjacent beam patterns. Even if the index numbers of the mobile and static cases do not match, the beam patterns might be similar. Still, the poor results in terms of throughput (c.f. Figure 3.17) and physical layer rate (c.f. Figure 3.18) suggest that this is not the case.

Finally, Figure 3.20 depicts a more detailed analysis of beampattern selection. We consider two trajectories (a) and (b), as in the left part of the figure. Both are horizontal, but (a) follows $y = 1.6$ m while in (b) the laptop is further away at $y = 4$ m. Moreover, the laptop is facing the docking station in both trajectories, as the arrows in Figure 3.20 indicate. We consider two metrics, namely, how often the beam pattern changes during each trajectory, and how many different beam patterns does a device use during each trajectory. Note that the former can be very large even if the latter is very small in case a device is continuously fluctuating between a

few beam patterns only. Figure 3.20 shows that the number of beam pattern changes is roughly similar for both trajectories. However, the number of beam patterns among which the device fluctuates is smaller for (a) than for (b). The reason is that the devices are closer to each other in case (a), and thus experience a higher SNR. Basic mobility mechanisms typically trigger a beam retraining when the SNR falls below a certain threshold. As expected, this happens less often for (a) than for (b). In case (a), the devices can communicate even if the main lobes of the chosen beam patterns are not pointing to each other due to the significant sidelobes of practical beam patterns. Thus, we conclude that sidelobes can be beneficial to deal with mobility since they may enable communication even in case of suboptimal beamsteering.

3.3.3. Multiple Antenna Arrays

Existing consumer-grade 60 GHz devices typically feature a single antenna array. For instance, both the Dell D5000 dock and the TP-Link Talon AD7200 have one antenna array facing forward and covering roughly 120° . This inherently limits the coverage area of the devices and may cause significant performance degradation under mobility. For instance, a device moving on trajectory (a) in Figure 3.16 suffers poor performance or even link loss (c.f. Figure 3.17) when traveling along the x-axis, that is, in parallel to the docking station. A natural solution to this issue is using two or more antenna arrays per device. In our above example, if the E7440 laptop were equipped with an antenna array facing sideways, it could switch to that array while moving along the x-axis, as depicted for trajectory (b) in Figure 3.16.







In this experiment, we emulate such additional antenna arrays on the E7440 laptop. To this end, we combine multiple measurements of the same trajectory with different device orientations. While this approach does not fully account for the antenna switching overhead, it provides a rough estimate of the benefit of using multiple antenna arrays in a device. Table 3.5 shows a comparison of the average performance for a device with two antenna arrays compared to the single array case. We observe that throughput increases on average 13.68% while the ratio of control packets to data packets decreases by 12.3%. This clearly shows that additional arrays can provide a significant benefit. Since the overall link quality is better, the PER for both control and data packets decreases, too. The difference in the case of control packets is only 10% because such traffic is transmitted at the most robust MCS defined in IEEE 802.11ad. However, for data packets, which are transmitted at higher MCSs, the PER decreases on average 56.64%. Table 3.5 also shows that MCS and beam pattern selection are more stable with two arrays. Essentially, using a second array facing sideways prevents low link quality and thus avoids that the rate and steering control algorithms cause unnecessary fluctuations.

In Table 3.6, we extend our analysis to up to eight arrays. We use random trajectories that may include areas with no coverage. As a result, the throughput values in Table 3.6 are lower than in Table 3.5. The icons show the relative placement of the antennas. As expected, the higher the number of antennas, the higher is the throughput. However, we observe that the antenna placement does play a significant role. For instance, placing a second antenna on the left side of

Table 3.5: Performance for multiple antenna arrays on trajectories of Figure 3.16.

	One Array	Two Arrays	Difference
Throughput	651.5 Mbps	740.65 Mbps	13.68%
# Control Pkt	827.9 pkts	733.6 pkts	-11.40%
# Data Pkt	25805.0 pkts	26217.5 pkts	1.60%
Ratio Control/Total	3.10%	2.72%	-12.30%
Control Pkt Error Rate	0.22%	0.20%	-10.00%
Data Pkt Error Rate	0.61%	0.26%	-56.64%
MCS Changes	19.5	18	-7.26%
Laptop BP Changes	18.25	17	-7.00%
D5000 BP Changes	20.25	18.16	-11.00%

Table 3.6: Average throughput/outage up to eight arrays on random trajectories.

						
Throughput [Mbps]	128	307	360	250	481	586
Outage [%]	79%	54%	48%	60%	33%	18%

the device results in better throughput than placing it on the right side. This is due to the case of the laptop that we use as a client. The case shadows the left side but does not block the right side. Thus, adding a second array on the left side is more beneficial. We observe an equivalent effect for outage.

3.3.4. Link Robustness

Our area-wide static measurement results in Figure 3.14 show that the E7440 laptop is not able to connect to the D5000 dock in a significant number of locations. Interestingly, during our mobile measurements we observe that the E7440 laptop stays connected at those locations if it associated to the D5000 dock while in one of covered areas in Figure 3.14. The underlying reason is that the devices use wide beam patterns during device discovery, and thus their discovery range is smaller than their communication range. In other words, devices can maintain a connection when moving into some of the light areas in Figure 3.14 because the beam patterns used for communication are more directional than the ones used for device discovery. To investigate this effect, we measure the throughput for three trajectories along the x-axis of our measurement grid. Specifically, our trajectories follow the line $y = 5$ meters, and each of them considers one out of three device orientations as seen in Figure 3.21. The three orientations of the E7440 laptop are as follows: facing the D5000 dock, facing left, and facing right. We then compare our result to the throughput that we measure in the static case (c.f. Section 3.2.4), for which we reset the connection prior to each measurement and are thus limited by the discovery range of the D5000 dock. Figure 3.21 shows our results. As expected, for the static case the laptop cannot connect to the dock at the edges of the trajectory for all three orientations. In contrast, for the mobile case

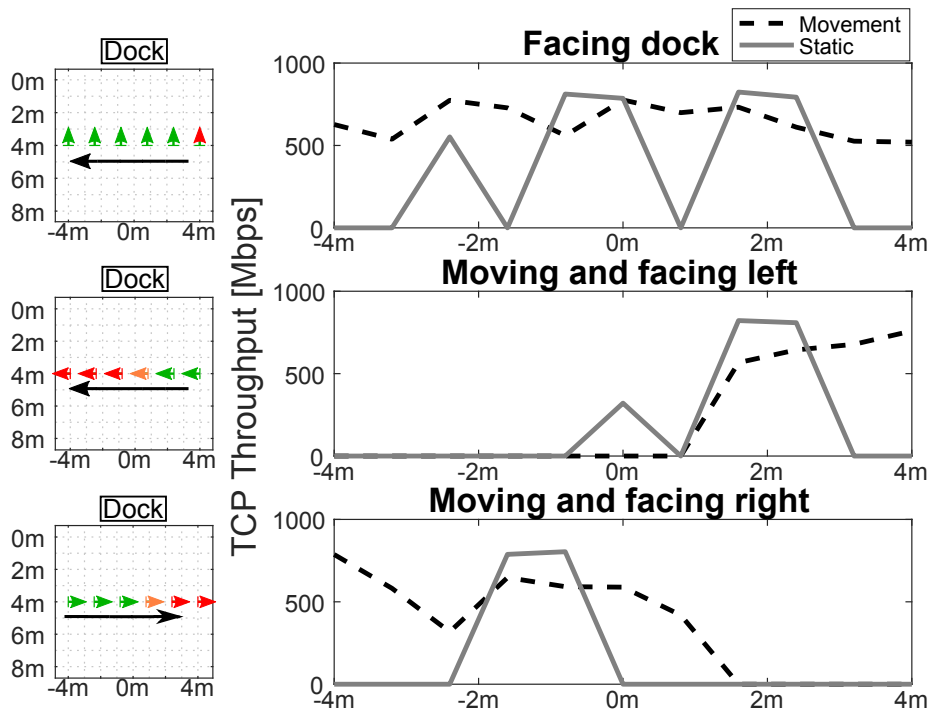


Figure 3.21: Horizontal Movement. The short arrows represent the orientation of the laptop while the long arrow represents its movement.

facing the dock, the laptop achieves a throughput of at least 500 Mbps at all locations. For the other orientations, the throughput of the mobile case decreases to zero at the locations where the dock is outside the 120° coverage area of the laptop.

3.3.5. Transient Blockage

In our last experiment, we consider a transient blockage scenario. To this end, we place the laptop and the dock at a distance of three meters and interrupt the link with a continuously moving obstacle. That is, the obstacle enters and leaves the LOS path without stopping at any time. We repeat the experiment for multiple speeds to study the impact of the blockage duration on how the nodes adjust their beam patterns to cope with blockage.

As soon as the robot enters the LOS path, the Signal-to-Noise Ratio (SNR) at the receiver drops. As a result, the nodes test different beam pattern combinations to try to avoid the obstacle. Since the obstacle moves continuously, the beam adaptation process is typically in vain—a working beam pattern combination quickly becomes sub-optimal as the robot continues its movement. Figure 3.22(a) shows the duration of this process. As expected, the process is shorter the faster the robot moves since it spends less time blocking the LOS path. However, the length of the beam adaptation process is not necessarily related to the number of beam pattern combinations that the nodes test, as shown in Figure 3.22(b). In particular, when the robot

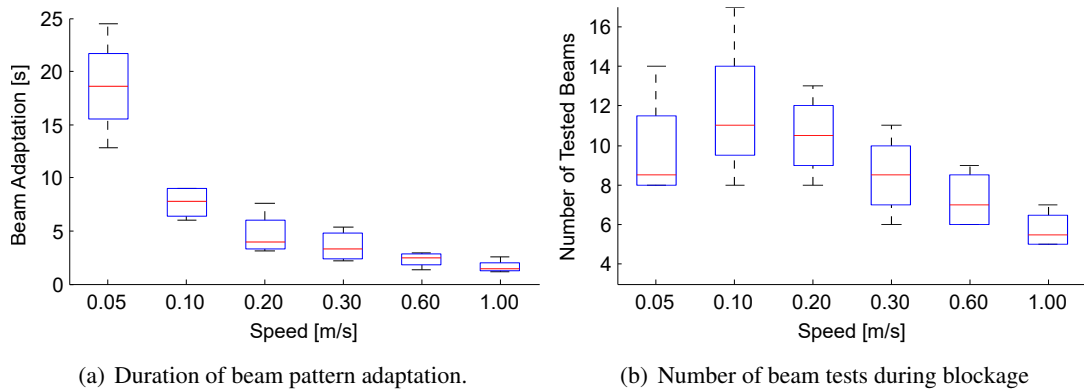


Figure 3.22: Transient blockage study.

moves at 0.05 m/s, the nodes test less beam patterns than at 0.1 m/s. This is likely because the SNR variations as the robot moves through the link are smoother, and thus the beam steering algorithm sets in less often. As speed increases, initially we observe the aforementioned increase in the number of tested beam patterns but then again a decrease. The latter is expected since the nodes simply have less time to test beams as the blockage becomes shorter. We also analyzed whether the beam patterns before and after the blockage match. Interestingly, this only happened in roughly 50% of our experiments, with no clear relation to the robot speed. This shows that a transient link blockage may induce a long-term sub-optimal beam alignment on a 60 GHz link, thus having a significant impact on performance.

3.4. Aerial Communications

There is now a growing interest in the use of mmWave frequencies for aerial to ground communication. These mmWave links to drones can, for example, be valuable for disaster recovery as well as environmental 3D and other high-data rate mapping uses [56, 57]. These drones can also give the opportunity of relaying wireless signals into very dense and crowded areas, needing coverage for an specific point in time [63]. A high level analysis of mmWave UAVs cellular networks is provided in [64]. The authors study different challenges and solutions such as beam training and tracking, the importance of the Doppler effect, spatial reuse and signal blockage. Khawaja et al. [65, 66] studied the temporal and spatial characteristics of air-to-ground mmWave propagation channels.

Our interest in mmWave COTS devices is that they can provide low-cost, mass produced radio hardware for consumer drones. In addition, due to the production volumes, the devices benefit from much smaller form factors and payload. In this section, we conduct a series of experiments to obtain some early insights into the feasibility of ground-to-air communication based on COTS devices. To the best of our knowledge, this is the first practical measurements study with drones using COTS mmWave devices.




Figure 3.23: Aerial communications testbed setup: ground station (client) with the display lid at 45° inclination, and a flying station (AP) attached to the DJI MATRICE 600 PRO drone.

3.4.1. Experimental Setup

We perform measurements with two IEEE 802.11ad laptops, having one laptop on the ground and one mounted on a drone. We create a TCP connection between these devices and analyze the achievable bitrate, as well as the chosen MCS and the selected beam pattern. The aerial node was set to AP mode and the ground station to client mode. The transmission is downlink, where the aerial node sends TCP traffic to the ground node. In order to configure the mode in each laptop we used the *wpa_supplicant* service, and to create the TCP traffic we used iPerf.

The laptop on the ground was left with the lid half-open – at 45 degrees inclination – to maximize the communication range of the antenna array into the field. The laptop mounted on the drone was left completely closed, with the antenna array pointing downwards, parallel to the ground, in order to maximize the coverage area. Figure 3.23 shows the equipment placement.

To characterize the drone-to-ground communication and see what link ranges can be achieved with mmWave COTS radios, we carry out the measurement campaigns depicted in Figure 3.24, which are:

- **Vertical:** In this test, we study the vertical limit of a mmWave aerial link flying directly over a ground AP (see Figure 3.24(a)). With this test we can obtain the maximum height an aerial mmWave link can obtain on top of a user.

- **45° elevation:** In our setup, to expand the communication range, the aerial relay should be placed at a 45° elevation angle with respect to the ground station. This 45° elevation angle represents the maximum steering angle from the flying station, corresponding to the 90° cone that a phased antenna array can cover with the most effective beam patterns, as it is placed parallel to the ground. In this test, we study how far a mobile node can maintain the mmWave link at this angular distance (see Figure 3.24(b)).

- 3D distance:** In this test, we measure the signal quality against the effective 3D distance. We test this approach for vertical heights of 3, 6, 9 and 12 meters (see Figure 3.24(c)). This way we can find out how far a drone can fly without losing connection for those specific drone heights.

- Wind:** We emulated wind by forcing the drone into minimal displacements towards different directions as seen in Figure 3.24(d). We make the drone vibrate in three independent scenarios with average speeds of 1.75, 0.50 and 0.81 meters per second, having peak movements up to 4.15, 3.19 and 2.34 meters per second, respectively (see Figure 3.24(e)). We compare this results with respect to regular wind conditions, where the drone does not shake.

For the first three scenarios, the drone is flown with a constant speed of 3 meters per second until the link breaks.

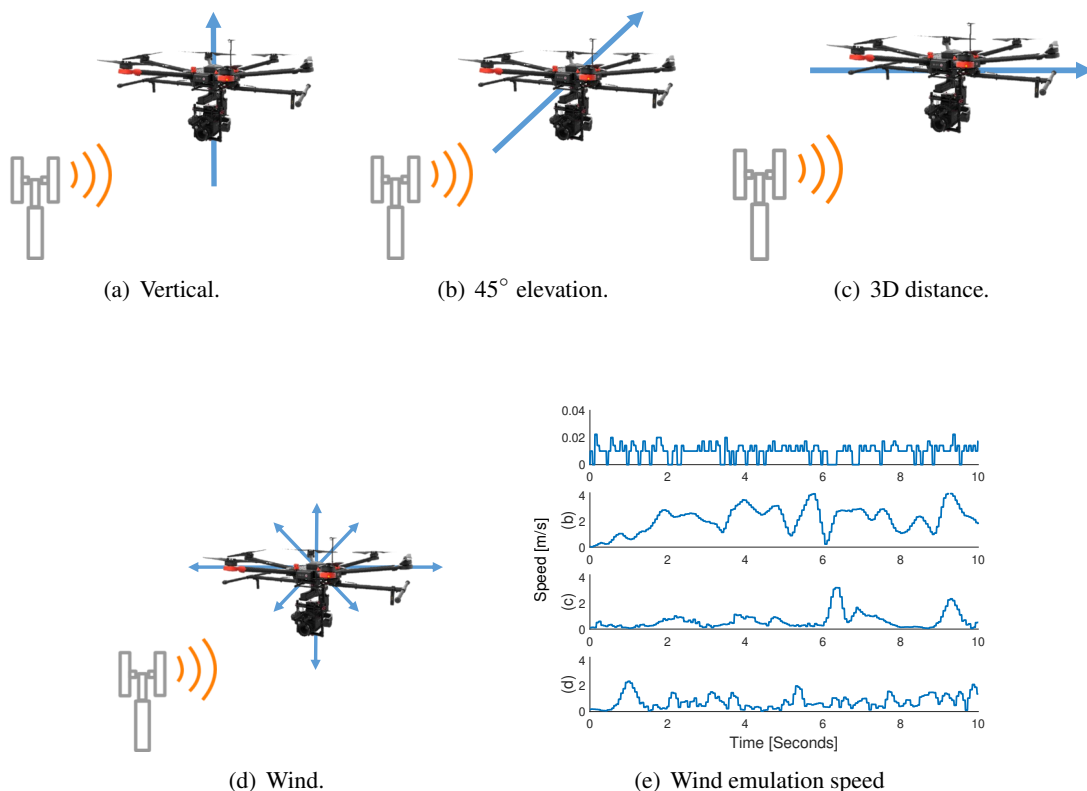


Figure 3.24: Experimental setup.

3.4.2. Results

The results from the vertical study are reported in Figure 3.25. We observe that the maximum link height is around 31 meters. For the first 20 meters, the communication throughput is mostly stable with rates of 1700 Mbps. Beyond this distance, the bitrate gradually decreases until reaching zero at 33 meters vertical distance. Figure 3.25(b) shows the beam pattern index selected during the test by the aerial node and the ground station. As the traffic is downlink, it can be seen how the drone varies from different beam patterns more often, while the ground station, which only needs to transmit acknowledgment frames, only changes its beam pattern three times. As the acknowledgment transmission uses a lower MCS than the data transmission, lower SNR values are needed, making the beam patterns to be valid for longer times, without having the need of training for better ones.

The results of the 45° elevation study are depicted in Figure 3.26. The maximum link distance achieved for this scenario is also about 31 meters. The communication rate is stable up to the distance of 25 meters, where the utilized MCS index is 8, going up to 10 in between 4 and




Figure 3.25: Vertical study analysis.




Figure 3.26: 45° elevation study analysis.




Figure 3.27: 3D distance study analysis: Bitrate & MCS results.

10 meters. When the drone flies beyond 25 meters, the rate and the MCS index go down rather abruptly. The beam pattern selection can be observed in Figure 3.26(b). We note that even though the relative angle among both stations is fixed, the drone performs several beam pattern changes during the measurement, oscillating between two main patterns—due to poor beam pattern selection mechanism and slight misalignments during the drone trajectory. In contrast, the ground station never changes its beam pattern, as explained above.

Results from the 3D distance study are reported in Figures 3.27 and 3.28, in the first showing the achieved bitrate and MCS and in the second the utilized beam pattern indexes. As expected, the drone achieves longer link distances when it flies higher, as it can serve wider areas with the use of the same beam patterns, due to the directional behavior of the antenna array, which is placed horizontally. Additionally, it can be seen that for higher drone positions, the link rate varies more, as the total distance between both stations increases. The maximum link distances depicted in these tests are equal to 14, 16.5, 22 and 24 meters for heights equal to 3, 6, 9 and 12 meters, respectively. We note that the ground station uses a set of very few beam patterns for all heights, while the beam configuration on the drone varies a lot. These variations of the drone beam pattern selection might be due to slight misalignments during the drone flight as well as poor beam pattern selection mechanism.

As the mmWave radios that we are using for this evaluation are regular laptops, their codebook is not optimized for 3D beamsteering. Once this is optimized, the communication range would be enlarged and we would see how for different angles the stations would choose different beam patterns, differently to what we see in our results, where stations jump between several beam patterns for non-related angles.

Figure 3.24(e) shows the speed of the drone versus time for each of the wind emulation measurements. Here we observe how the total magnitude of the drone movement was for each of the time intervals of the measurement. We report in Figure 3.29 the results related to the wind emulation. Here we observe the achieved bitrate versus time with a time precision of 0.1 seconds.

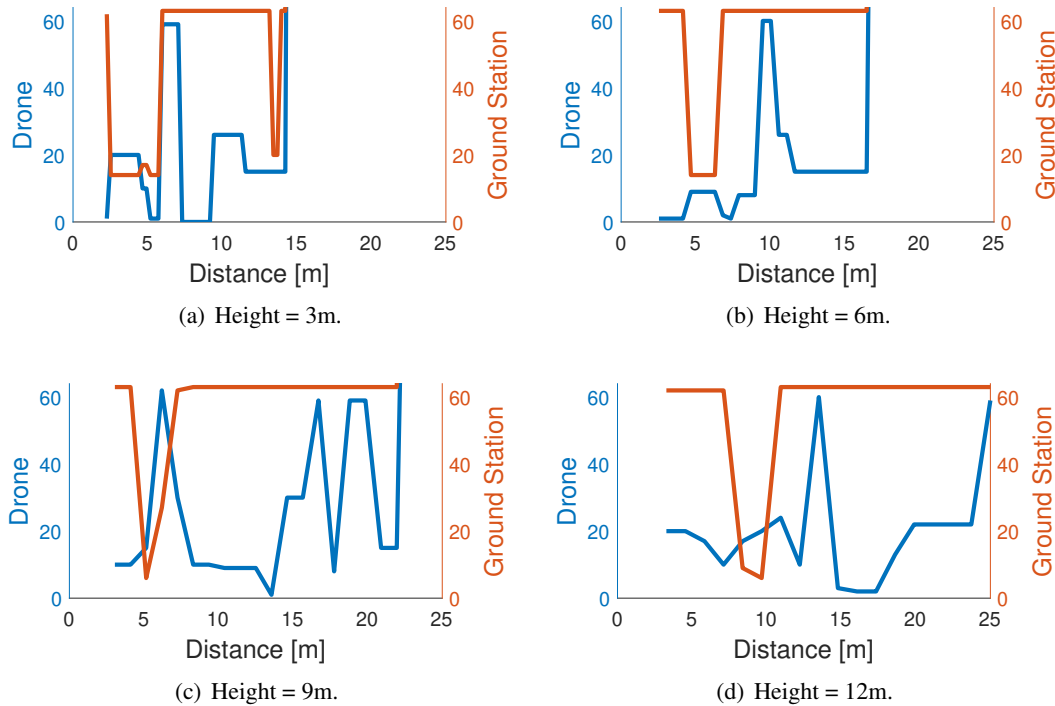


Figure 3.28: 3D distance study analysis: Beam pattern selection.

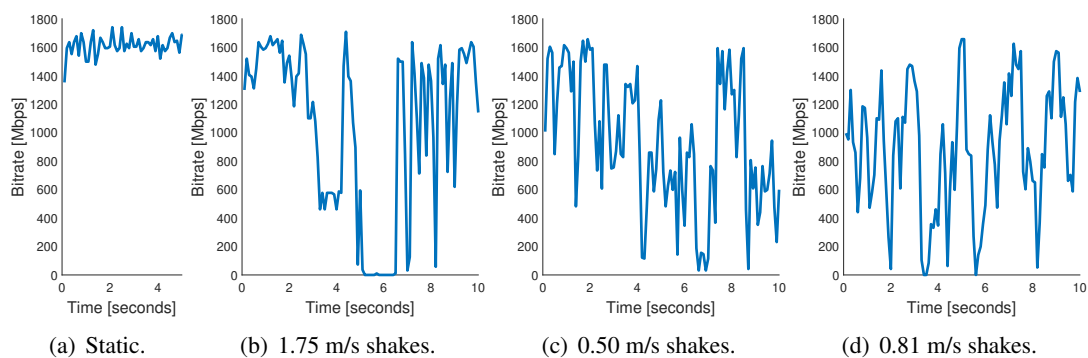


Figure 3.29: Wind emulation bitrate with respect to time. In (a) we can observe the rate achieved without emulating strong wind while in (b), (c) and (d) we can see 3 examples of strong wind emulation with average throughputs of 1000, 933 and 884 Mbps and standard deviation equal to 595, 467 and 443 Mbps. The speed versus time for each of the cases can be seen in Figure 3.24(e).

First, we recorded an example of standard wind characterization in Subfigure 3.29(a), where the drone was flying in a static point. With normal wind conditions, the drone was shaking with an average speed of 0.01 meters per second. With this steadiness, the link was able to provide a constant flow of data with small variability: an average bitrate of 1600 Mbps and standard deviation of 70 Mbps was measured in that specific location. Figures 3.29(b), 3.29(c) and 3.29(d) show three different examples of high wind emulation, corresponding with the speeds seen in the subplots (b), (c) and (d) from Figure 3.24(e). In here, we shake the drone with average speeds of 1.75, 0.81, and 0.50 meters per second, respectively, having peak movements up to 4.15, 3.19 and 2.34 meters per second. We determine the effect of the vibrations on the beam misalignment and the resulting rate degradation. Nevertheless, even though the standard deviation is high (595, 467 and 443 Mbps) for each of the cases, the system still achieves a good average throughput of 1000, 933 and 884 Mbps, respectively.

On average, the wind emulation reduces the throughput of the link by 40%, which means that the link is still capable of transmitting high speed data despite its high variability.

3.5. Discussion

We use two methodologies that allow us to gain new fundamental insights into the operation of 60 GHz networks, and which make our results stand apart from related work. First, we analyze commercial 60 GHz devices with the VubIQ system to understand their behavior at the physical layer. Second, we use a robot-assisted technique to obtain a first-of-its-kind coverage map of such consumer-grade devices. This allows us to obtain a large trace dataset that captures the behavior of 60 GHz networks. Most importantly, we partially reverse-engineer the Wilocity Monitor Service to obtain unprecedented access to a firmware interface of the `Wil6100` card. We use the resulting traces to draw general conclusions, in contrast to existing work, which mainly focuses on specific use cases. Next, we discuss the insights that we gain from the above methodologies and derive design principles for the development of 60 GHz networks.

Antenna patterns. We observe that the antenna patterns of consumer-grade 60 GHz devices not only exhibit significant side lobes due to a cost-effective design but also very much depend on the specific location of the antenna array on the device. This partly invalidates current assumptions regarding channel access at the MAC layer—in scenarios where devices with certain beam patterns do not interfere, others may cause collisions. As a design principle, we derive that 60 GHz networks should implement multiple MAC behaviors and choose the one which is most suitable for the beam patterns of the individual devices in the network. Also, devices should use more directional beam patterns, since current implementations do not exploit the gains that these antenna arrays can achieve. In order to do so, manufacturers may need to redesign the shape and placement of the antenna array within the device. At the same time, having these COTS 60 GHz devices implementing so wide patterns enables these first-of-a-class devices to have a robust link without requiring very accurate beamtraining or realignment in case of mobility. However, this

may need to be reconsidered in future hardware generations. Devices will perform poorly in case of multiple parallel links if they are not able to synthesize narrower beam patterns, generating interference and reducing spatial reuse.

Reflections. We show that 60 GHz signals may cause interference even after reflecting twice off walls. As a result, MAC layer designs which exploit the sparsity of 60 GHz signals to increase spatial reuse may incur unexpected collisions. Such designs are often based on geometric principles, that is, the beam pattern and range of a node define its area of interference. As a design principle, we derive that such protocols should extend this geometric approach to include up to two signal reflections off walls or obstacles, as we see from our measurements that these second order reflections can still be strong. Also, it is possible to take advantage of these reflections for user localization.

Interference and spatial reuse. We observe significant interference in scenarios with two or more 60 GHz links, even in cases for which the parallel operation of the links should be feasible. This also includes interference via reflectors. As a result, spatial reuse drops dramatically in the systems that we analyze. The underlying reason is that devices need to carrier sense each other, leading to backoff and increasing the channel idle time. Orientation of links also have a strong impact on this spatial reuse. As a design principle, we derive that more directional patterns should be used to improve spatial reuse, as well as implementing medium access control approaches that take this into account and are thus more efficient than traditional carrier sensing.

Aggregation. Our analysis reveals that 60 GHz aggregation provides large throughput gains at much smaller timescales than legacy systems. For instance, while IEEE 802.11ac achieves a $2\times$ gain with a frame length of 8 ms (c.f. Table 1 in [61]), IEEE 802.11ad achieves a $5.4\times$ gain by aggregating only up to $25\ \mu\text{s}$ due to its very high data rates. Hence, the impact of aggregation on delay is much smaller for IEEE 802.11ad, while reducing dramatically channel usage time. As a design principle, we derive that the frame length should not only depend on the desired throughput and delay, but also on how many nodes share the medium. If many nodes share it due to, e.g., wide beam patterns, a higher aggregation level helps to provide channel time for all nodes.

Range. We show that the range of 60 GHz links often varies significantly between experiments (c.f. Figure 3.7(b)) due to, e.g., differences in temperature and humidity. As a design principle, we derive that devices may need to adjust their transmit power to control interference even in quasi-static scenarios, such as wireless backhaul networks. Performing power control will also lead to energy savings when the links are close to each other.

Coverage. Although 60 GHz communication experiences less multi-path effects than 2.4 GHz or 5 GHz communication, our area-wide measurements reveal that coverage cannot be modeled based on geometric principles only. Even in an empty room, we do not observe a cone of coverage as often suggested in earlier work. The underlying reasons are (a) the irregular beam patterns of consumer-grade 60 GHz hardware, and (b) the impact of antenna placement. For instance, the antenna of the D5000 dock is located at the right side of the device case. As a result, in Section 3.2.4 we observe that the left part of Figure 3.14 tends to suffer from less coverage than

the right part. This suggests partial shadowing due to the device case. The same issue would arise with external adapters (e.g., IEEE 802.11ad USB dongles), since their orientation would also be critical for coverage. Thus, we conclude that coverage models for 60 GHz networks must not only take into account the room geometry and obstacles, but also the physical characteristics of the transmitter and the receiver. This means that the AP placement and orientation play a very important role for such 60 GHz wireless networks. Placing APs on a wall or in a corner instead of in the middle of the room may result in better performance due to the placement of the antenna array in current devices and their difficulty for beam steering into extreme angles. Placing an AP in the center of a room requires full 3D beam steering in azimuth and elevation angle, whereas current phased arrays typically have limited steerability particularly for the elevation angle.

Mobility. The performance gap in terms of the impact of indoor mobility on throughput is on average 30% for IEEE 802.11ad and traffic below 1 Gbps. Beyond beam pattern misalignments, the underlying reasons include MCS fluctuations, higher PER, and 30% higher control traffic overhead. While our setup does not allow us to infer how this performance gap behaves beyond 1 Gbps due to the use of Gigabit Ethernet, we expect it to be more pronounced, requiring the development of faster beamtracking techniques, such as the one in [49].

Multiple antenna arrays. Our coverage analysis shows that existing 60 GHz devices radiate energy in unexpected directions, such as backwards or sideways. However, connections in such directions are strongly impaired and may often cause significant harm due to strong MCS and beam pattern fluctuations. Hence, even if transferring communication to a second antenna array on the device involves a certain overhead and is not strictly necessary, we expect it to pay off. Moreover, since the device already has significant information regarding the angle at which its communication partner is located, a full beamtraining is unnecessary. Also, manufacturers should take this into account when designing their devices, as the antenna placement and the number of antennas plays an important role in terms of performance.

Discovery range. We observe that the discovery and communication ranges of the Dell docking station system differ significantly. This suggests that the dock and the laptop use quasi-omni directional beam patterns during the discovery phase to keep beamtraining overhead to a minimum, and directional beam patterns during the data communication phase to improve performance. The IEEE 802.11ad standard recommends such quasi-omni beamtraining for efficiency. Instead of requiring each side to perform a full beam sweep for each of the sectors of the other side, the standard allows one side to remain in quasi-omni mode during the sweep. The process is repeated for each side, which means that beamtraining only requires two full beam sweeps. A natural limitation of this approach is that nodes cannot perform beamtraining if both nodes need to use a directional beam pattern to be in range. To achieve both efficient beamtraining and large discovery range, future work should focus on more efficient beam sweeping. For instance, nodes could perform partial sweeps based on the expected location of their communication partner, or they could leverage signals received via side-lobes. Given the strong side lobes of consumer-grade 60 GHz antennas that we observe in Section 3.2.2, the latter

is a highly promising approach for practical network deployments.

Drone Communications. We show the feasibility of using drones as mmWave relays. Being this, for the best of our knowledge, the first work with COTS 60 GHz radios mounted on drones, we demonstrate that this application is indeed possible and very promising for future cellular applications. Even though the link range is still limited (up to 24 meters extension for our experimental results) our aerial links look promising. Also, the limited range achieved in this work should improve with next-generation mmWave radios, as the ones that we used for this study are the first-commercial IEEE 802.11ad radios to be in the market and, as shown before, these COTS systems still have room for improvement. We show that for the measured static position with low-wind conditions, the achieved TCP throughput is stable with an average throughput of 1600 Mbps, making this drone relays a perfect candidate for cellular mmWave micro cells as well as for applications such as natural disaster link recoveries or High-Definition video streaming for disaster mapping. However, when emulating high wind conditions the achieved throughput varies a lot, making this drone relays not optimum for low-jitter applications such as the previous mentioned High-Definition video streaming for disaster mapping, as the wireless link could cause some jitter and delays on the video transmission. However, it would still be valid for next-generation cellular micro cells, as the throughput is still large for user applications that only need regular Internet access where the delays on the link are not so critical.

Overall, we can conclude that this technology, still being in an early stage, is going to make a big difference in future wireless deployments. We have seen how these devices perform great in short range indoor communications as well as in outdoor aerial links. The directivity and the throughput that these devices achieve, still being far from what theoretical models promised years ago, is a great upgrade with respect to previous wireless systems in the market. Next-generation devices will provide better beamforming capabilities which will help to achieve longer link distances and more efficient spatial reuse, improving these devices into higher performance COTS systems. The current status of this technology, together with the advancements that manufacturers will provide in future revisions of the hardware, will give to the users the possibility of having multi-gigabit per second wireless rates while achieving nearly to zero interferences.

Chapter 4

Range Boost

As previously discussed, 60 GHz network deployments are focused on short-range communications. The range of 60 GHz networks depends on a number of factors, such as the directivity of the antenna and the allowed transmission power. Our measurements from Section 3.2.1 reveals maximum link distances of 17 meters (i.e. Figure 3.7(b)). In other works, the general consensus is that the range of 60 GHz networks in LOS conditions is between 10 to 15 meters [25, 28, 33, 34]. In isolated cases, ranges of up to 25 meters are feasible. Such outliers may be due to, e.g., dry environments in which water absorption is limited.

Given the limited range in the standard case, extending the link length is highly interesting as there will be cases where the users will need longer link distances, not being able to maintain connectivity. For such cases, 60 GHz networks may need to fall back to legacy ISM IEEE 802.11 using Fast Session Transfer (FST). FST consists of the process where IEEE 802.11ad devices decide to fall back from the 60 GHz band to the 2.4 or 5 GHz band due to poor channel characteristics [20]. This has a disastrous impact on throughput since such legacy networks are orders of magnitude slower than IEEE 802.11ad. Otherwise, 60 GHz networks must resort to relays as included in the IEEE 802.11ad amendment [20]. These relays have been demonstrated in [67, 68] and have been widely researched in terms of cross-layer modeling. For example, Singh et al. [69] propose a cross-layer modeling methodology and novel multihop MAC architecture for efficient 60 GHz wireless transmissions. Qiao et al. [70] proposes an algorithm to exploit the spatial capacity of mmWave with a concurrent transmission scheme for multihop networks.

In order to avoid the use of relays, in this chapter we propose the application of different frequency selective techniques to extend the range of 60 GHz communications. In order to do so, we measure real-world indoor 60 GHz band channels and study their behavior as distance increases. We focus on the channels defined in the IEEE 802.11ad standard, that is, each channel has a bandwidth of 2.16 GHz. Moreover, we consider the OFDM parameters as in the standard. We also study both the practical performance for the case of horn antennas, as well as for phased antenna arrays in our testbed. Our results show that, indeed, channels exhibit frequency selectivity. Based on this, we compare the standard operation of 60 GHz networks with techniques

that do take into account this frequency selectivity.

We can summarize the contributions of this chapter as:

- We measure real-world 60 GHz indoor channels, showing that they do exhibit significant frequency selectivity.
- We show that frequency selectivity becomes stronger the longer the link, or the wider the beam pattern, since for both cases reflections become more likely.
- We analyze how schemes that exploit frequency selectivity can extend the range of a link by up to 50%.
- We demonstrate that the above techniques can also exploit frequency selectivity to save up to 7 dB transmit power.

The chapter is structured as follows. Section 4.1 gives an overview on exploiting frequency selectivity, introducing the techniques and metrics that we use in this chapter. Section 4.2 provides details on the tested setups as well as the channel characterization. In Section 4.3 we present our results and analysis for range extension and energy savings, as well as we validate these mechanism with the use of an antenna array. Finally, in Section 4.4 we discuss the insights that we gain from our results.

4.1. Exploiting Frequency Selectivity

Typical communication systems use the same MCS for the whole channel indistinctly of the different channel conditions that can be along the channel bandwidth. What this means is that when the link suffers from frequency selectivity, the global performance of the connection decreases. Even though only some fraction of the channel experiences low SNR values, the rest of the channel has to adapt and lower their MCS in order to have all the frequency carriers working under the same conditions, decreasing the global performance of the system. In the same way, existing 60 GHz wireless networks ignore frequency selectivity. Most implementations use single-carrier schemes as defined in WiGig [19] and the IEEE 802.11ad amendment [20]. This means that each physical layer symbol spreads over the entire 2 GHz channel bandwidth, regardless of potential variations of the channel in the frequency domain. While the standard also allows for OFDM, all subcarriers are treated equally in terms of MCS thus not taking into account frequency selectivity.

Since nodes must use highly directive antennas to overcome the high path loss at those frequencies, reflections are expected to be rare and thus frequency selectivity should be limited. In other words, due to the lack of a rich multi-path environment as in systems using omni-directional antennas, channels should be flat. However, 60 GHz channels often behave differently as shown in Chapter 3, where antenna beam patterns are not as pencil-shaped as theory typically assumes

or indoor environments feature a significant number of reflections. Hence, the signal is likely to reach the receiver not only via the LOS path but also via reflections. While the number of reflections is much lower compared to lower frequency systems, frequency selectivity occurs to a certain degree in mmWave systems [42]. Related work shows that this directionality plays a fundamental role regarding channel characteristics as using directional antennas reduces multi-path effects [35]. However, for short strong links, this does not play a significant role. In most cases, the difference in terms of power between the weakest and strongest subcarrier is limited. Thus, the complexity of techniques that take into account frequency selectivity (e.g., bitloading, waterfilling, or subcarrier switch-off) does not pay off. However, such techniques are beneficial for weak, long links. As channel quality worsens with distance, some subcarriers may not be able to sustain communication while others can still operate. According to the current design of 60 GHz networks, this would break the link entirely.

The above behavior points to the following symbiosis. While techniques exploiting frequency selectivity only provide limited benefit for short links, for which 60 GHz networks anyhow provide very large throughput, such techniques can improve performance significantly for long links, which is precisely where 60 GHz networks have strong limitations. However, for these short links we can take advantage of these techniques in order to obtain power savings of up to 7 dB while in the case of weak long links we are able of extending their range for around 50% of their maximum link distance.

4.1.1. Compared Techniques

We consider an IEEE 802.11ad network using OFDM as defined in the standard. Thus, nodes occupy a bandwidth of 2.640 GHz using 512 subcarrier in total. Out of those, 336 subcarriers transport data, 16 are pilots, and the rest is nulled, using 1.88 GHz out of the channel bandwidth of 2.640 GHz [20]. Each subcarrier has a bandwidth of roughly 5.15 MHz. The compared techniques that we use in this piece of work are:

Baseline. In IEEE 802.11ad, the transmitter uses the same MCS with even energy distribution for all subcarriers. We use this baseline to assess the performance of the different techniques exploiting frequency selectivity. It is the standard technique to have a communication link between devices, but its robustness is limited. If a link has one or more weak subcarriers, those subcarriers force all other subcarriers to use a lower modulation and coding scheme, resulting in poor performance.

Subcarrier Switch-Off. This mechanism is based on using only powerful sub-carriers in order to achieve a high communication rate, not using the subcarriers that experience poor channel conditions [71–73]. Specifically, our implementation of subcarrier switch-off selects the combination of subcarriers and MCS value which maximizes throughput. That is, we do not just switch-off subcarriers that are not able to transmit data even at the lowest MCS, but also switch off the ones which could transmit data but would limit the overall performance. Similarly to our baseline, all subcarriers use the same MCS.

Bit Loading. OFDM systems that support bit loading can use a different MCS value for each subcarrier. As a result, subcarriers with a high SNR are mapped to high MCS values, while the ones with low SNR are mapped to low MCS values. In contrast to subcarrier switch-off, all usable subcarriers are used.

Waterfilling. OFDM implementations typically send the same amount of power uniformly across all subcarriers. Power loading allows transmitters to distribute the transmit power unevenly among subcarriers, such as in water-filling [74–76]. Hence, the transmitter can, for instance, allocate more power to poor subcarriers to mitigate deep fades [77,78]. In contrast, waterfilling distributes the power according to the channel response in order to maximize the spectral efficiency of the system. This means that waterfilling allocates more power to subcarriers which are less noisy, and less power to subcarriers which anyhow have a low SNR. Such an approach only is beneficial if combined with bit loading or subcarrier switch-off, since otherwise the subcarriers with low SNR would limit the performance of the system. In our case, we use waterfilling along with bit loading. As a result, the best subcarriers can transmit data using very high MCS values. We expect waterfilling to perform best out of all our techniques.

4.1.2. Metrics

In our evaluation, we obtain the per-subcarrier SNR that we can achieve for each of the techniques in Section 4.1.1. While this already provides an intuition of the performance of each technique, the actual gain only becomes evident in terms of throughput. Related work often approximates throughput simply using the well-known Shannon-Hartley theorem, which provides the capacity C of the link based on its SNR and bandwidth B as $C = B \cdot \log_2(1 + SNR)$. However, this does not take into account the threshold effects that result from using a discrete set of modulation schemes or coding rates. In order to take this into account, we use two metrics:

Non-Coded Bitrate. We call ‘non-coded bitrate’ the throughput that results from rounding down the spectral efficiency of each subcarrier to the spectral efficiency of the next-lowest MCS. For instance, if we obtain a spectral efficiency value of 3.2 bits/s/Hz for a certain subcarrier, we round it down to 2 bits/s/Hz since it is the next-lowest power of two. This would correspond to using 4-QAM on that subcarrier. When computing the non-coded bitrate, we consider modulation values up to 64-QAM, which is the highest modulation defined in the IEEE 802.11ad amendment [20]. Since this metric does not take into account channel coding, it allows us to study the impact of the MCS thresholds only. These thresholds often determine the gain of a technique compared to our baseline even if a technique provides an SNR improvement, the gain may be zero if it is not enough to reach the next MCS.

Coded Bitrate. Our second metric is the so-called ‘coded bitrate’, which takes into account both the specific modulation and the coding values as defined in IEEE 802.11ad. Essentially, the standard defines a guideline regarding which MCS to use for which SNR. This allows us to convert the SNR directly into the actual data rate at the physical layer, taking into account all the details defined in the standard such as the overhead due to headers and preambles. Compared to

our first metric, the coded bitrate results in lower throughput values due to the impact of coding. This impact is particularly strong for weak subcarriers, since IEEE 802.11ad resorts to codes with high redundancy to overcome the channel impairments. Conversely, the impact on strong subcarriers is relatively low.

4.2. Evaluation of Channel Characterization

In the following, we study the performance of the techniques presented in Section 4.1 in practice. First, we describe our testbed, we then show the feasibility of these techniques for 60 GHz networks and then finally discuss our results.

4.2.1. Testbed Setup

We use a Universal Software Radio Peripheral (USRP) X310 Software Defined Radios (SDRs) to generate the baseband signal at the transmitter, and an other one as the receiver to decode the received signal. As the USRP daughterboards do not allow mmWave transmission or reception, we use the Sivers IMA FC1005V/00 converters to do the mixing of the baseband signal into the 60 GHz channel, both at the transmitter and receiver. In order to control and run the measurements we use a server machine with GNUradio connected to both USRPs. For the bulk of our experiments, we use 20° horn antennas with a gain of 25 dB for both transmission and reception. A schematic of our testbed setup can be seen in Figure 4.1. However, we also compare our results to the performance achieved with a practical phased antenna array in Section 4.3.3.

Since our USRPs generate a signal at -18 dBm, we set the gain of the converter to 20 dB such that overall we comply to the maximum Equivalent Isotropically Radiated Power (EIRP) as regulated by the US and EU authorities. Due to the limited bandwidth of the USRP daughterboards, we cannot capture the entire bandwidth of an IEEE 802.11ad channel at once. Thus, we scan one subcarrier at a time. To avoid channel variations while we scan the channel, we perform all experiments in controlled environments. In Section 4.2.2 we validate that the channel is actually static across multiple measurements, as well as inspect how frequency selective the




Figure 4.1: Setup. Server controls the Rx and Tx SDRs connected to the 60 GHz converters.

Table 4.1: Evaluation Scenarios

	Width	Length	Measured distances
Empty Room	6 m	18 m	$d \in [2.5, 5, 10, 15]$ m
Lab	6 m	18 m	$d \in [2.5, 5, 8.7, 11]$ m
Corridor	2.5 m	25 m	$d \in [2.5, 6, 9, 15]$ m

channel is. We measure the impact of frequency selectivity in three different scenarios, namely, a corridor, a lab, and an empty room. Table 4.1 provides an overview of their dimensions. We expect these scenarios to provide very different propagation environments. The small dimensions of the corridor should benefit reflections off walls. In the lab, walls should play a much less important role but reflections from furniture should be high. Finally, in the empty hall we expect only few reflections. Due to the arrangement of each of the locations, we cannot perform measurements at exactly the same distances in all three locations. For instance, in the lab we are limited by the fixed layout of the tables. Table 4.1 also contains the measured distance-positions for each of the setups.

To characterize the channel, we transmit a probe signal with a duration of three seconds for each subcarrier. Once the measurement finishes, we combine all subcarriers to obtain the channel response of the entire channel bandwidth. Without loss of generality, we consider the third channel defined in the IEEE 802.11ad standard, which goes from 61.56 GHz to 63.72 GHz. Based on our measurements, we obtain the SNR for each subcarrier, which is the input that we need for the frequency selective techniques that we study. Next, we compute the metrics described in Section 4.1.2. For our first metric, we choose the modulation of each subcarrier such that we do not exceed a Bit Error Rate (BER) of 2.4%. This is in accordance to widespread methods for determining the sensitivity of a device. We then add up the rate of each subcarrier to obtain the total link capacity. For our second metric, we select the MCS for each subcarrier based on the measured SNR.

4.2.2. Channel Characterization

Before applying frequency selective techniques into our mmWave link, we need to know if the channel characteristics fit to these. First, we validate that channels are stable in an environment without mobility. Next, we analyze how frequency selectivity varies with distance. Then, we evaluate how well the signal strength matches the free-space path loss formula.

4.2.2.1. Channel Stability

In Figure 4.2 we show two subsequent measurements of our 60 GHz channel within a time interval of three hours. We observe that the channel is identical in both measurements. Hence, our measurement technique based on scanning one subcarrier at a time is valid. Figure 4.2 also depicts a significant frequency selectivity, since subcarriers 0 to 100 suffer from significant

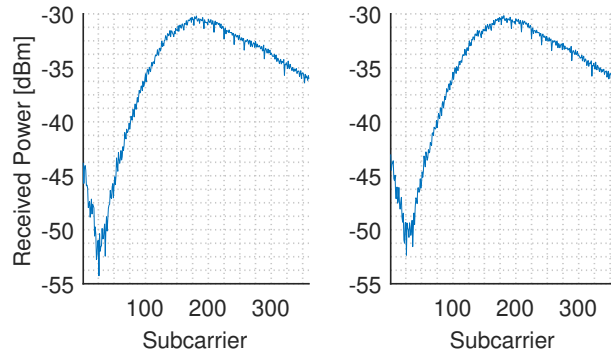


Figure 4.2: Consecutive channel measurements in the corridor scenario at 9 m.




Figure 4.3: CDF of the signal strength of subcarrier 100 during one hour.

fading compared to the rest of the channel. The difference between the strongest and the weakest subcarriers is up to 22.25 dB.

Further, in Figure 4.3 we show the CDF of the signal strength of a single subcarrier over a time interval of one hour for the Lab scenario at distance 2.5 m. We depict subcarrier 100, but the result is equivalent for all other subcarriers. As expected, the subcarrier exhibits high stability—since the channel frequency fading only depends on a few multi-path components, the channel remains constant as long as the environment does not change. In contrast, lower frequency networks are prone to much more fluctuations since the receiver receives many more multi-path components.

4.2.2.2. Frequency Selectivity

Next, we discuss the frequency selective behavior of communication in the 60 GHz band. Figure 4.4 shows the detailed channel measurements for each scenario and link distance. Further, Table 4.2 provides an overview on the difference in dB between the best and the worst subcarrier for each of the channels in Figure 4.4. For all scenarios, we observe that selectivity is small for the shortest distance. This is expected, since transmitter and receiver are placed at only 2.5 m, which means that any reflection off the surrounding walls or obstacles is highly unlikely to fall into the beam of the receiver. In general, Figure 4.4(c) exhibits the strongest frequency-selective behavior (22.25 dB), and Figure 4.4(a) the lowest (less than 5 dB). This is expected due to the

Table 4.2: Frequency selectivity with distance

Empty Room		Lab		Corridor	
2.5 m	4.27 dB	2.5 m	5.28 dB	2.5 m	5.09 dB
5 m	4.86 dB	5 m	6.13 dB	6 m	6.95 dB
10 m	8.80 dB	8.7 m	5.64 dB	9 m	22.25 dB
15 m	4.95 dB	11 m	6.31 dB	15 m	7.50 dB

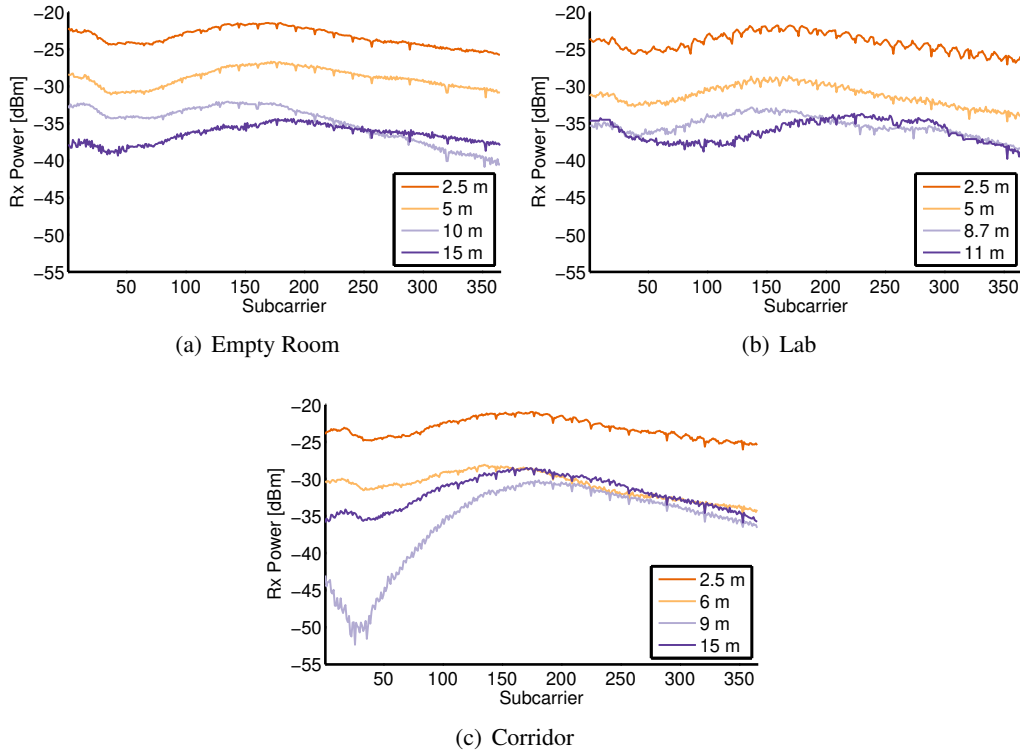


Figure 4.4: Channel measurements for all the positions.

narrow nature of the corridor, and the wide space of the empty room, which limits reflections. Interestingly, Figure 4.4(b) also exhibits low frequency selectivity (5 dB). While the lab scenario is prone to reflections off furniture and metallic objects, blockage of such reflections is also more likely compared to our other scenarios. For instance, in the lab scenario we place the transmitter and the receiver on the lab tables. As a result, the tables themselves as well as the chairs in the lab often block potential ground reflections.

For increasing distances, we observe that selectivity increases since it becomes more probable that the receiver receives reflections. However, this behavior is highly location dependent. For instance, at a distance of 10 m in the empty room, selectivity increases up to 8.80 dB but then decreases again to 4.95 dB for 15 m. This suggests that at 10 m a strong multi-path component geometrically hits the receiver beam, but misses it at 15 m. Similarly, selectivity increases up to 22.25 dB for the corridor case at a distance of 9 m. Hence, we conclude that frequency selectivity plays a significant role in indoor 60 GHz scenarios. While it does not necessarily increase with




Figure 4.5: Channel measurement in the empty room scenario at 10 m and 15 m.

distance, we observe that the selectivity is lowest for short links.

Moreover, frequency selectivity in the 60 GHz band may also result in higher performance for a subset of subcarriers as distance increases. For instance, we observe such a behavior in Figure 4.5, which depicts channel measurements at 10 and 15 m in our empty room scenario. As expected, the received power is generally lower for 15 than for 10 m, but for subcarriers 250 to 350 the opposite occurs. The remaining subcarriers roughly follow the free-space path loss model since up to subcarrier 250 the shape of the curve is approximately the same for both distances. This highlights that the impact of reflections in the 60 GHz band can be much stronger than often assumed in the literature. The overall result in Figure 4.5 clearly contradicts the free-space path loss model, which should thus not be used as a basis for, e.g., simulating communications and networks operating in the 60 GHz band.

4.2.2.3. Signal Strength

After seeing the previous results from Figure 4.5, where some subcarriers receive higher power when they are further away, in this section we want to study how the received power varies for the different positions and the different scenarios that we considered. Figure 4.6 shows the resulting received power and how this compares to free-space path loss propagation formula, using the first measurement as a reference. We can see that, in general, the different measurements adjust to the theoretical model. However, we also observe small variations in virtually all of them. This is due to the existence of reflections, and how they increase or decrease the received signal power depending on whether reflections add constructively or destructively at the receiver. The most striking result can be seen for the corridor measurements, since at 15 m the received power is more than 7 dB above the theoretical result due to the constructive addition of reflections. This is surprising, given that the measured positions at 6 and 9 m adjust very well to the theoretical model. We observe a similar behavior at 15 m in the empty room case, even though the impact of reflections is limited in that scenario. This again validates our findings in Section 4.2.2.2, that is, that the theoretical free-space path loss model is not valid in 60 GHz indoor environments.

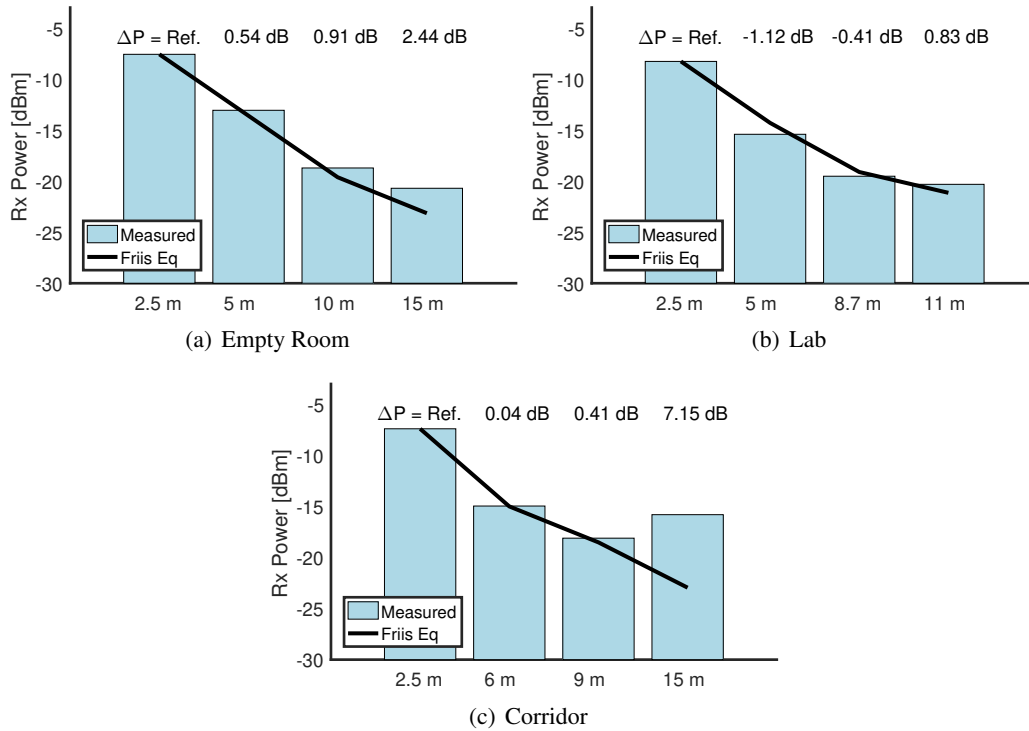


Figure 4.6: Measured overall signal strength compared to the theoretical Friis equation.

4.3. Results

In this section we show how the different frequency selectivity techniques can help extending the range of a 60 GHz network or providing energy savings.

4.3.1. Range Extension

Next, we study the range extension that we achieve when using the techniques discussed in Section 4.1.1. Specifically, we compute the throughput according to our metrics in Section 4.1.2 for each of the locations in Table 4.1. Figure 4.7 shows the achievable non-coded throughput while Figure 4.8 shows the achievable coded throughput. The figures show that all mechanisms perform significantly better than the baseline. We observe a clear relation between the received power in Figure 4.6 and the throughput, as expected. Further, we find a clear relation of the throughput and the selectivity of the channels listed in Table 4.2 for the different mechanisms. Figures 4.7 and 4.8 depict the impact of coding on the rate achievable on a link. As discussed in Section 4.1.2, the impact is stronger for weak links that require high redundancy to operate than for strong links. In all of our three scenarios we observe that bit loading is the technique that provides the highest bit rate, particularly for impaired links. This holds true both for the case when bit loading is used with a power distribution based on water filling and for the case with a homogeneous power distribution. As a result, bit loading increases the range at which

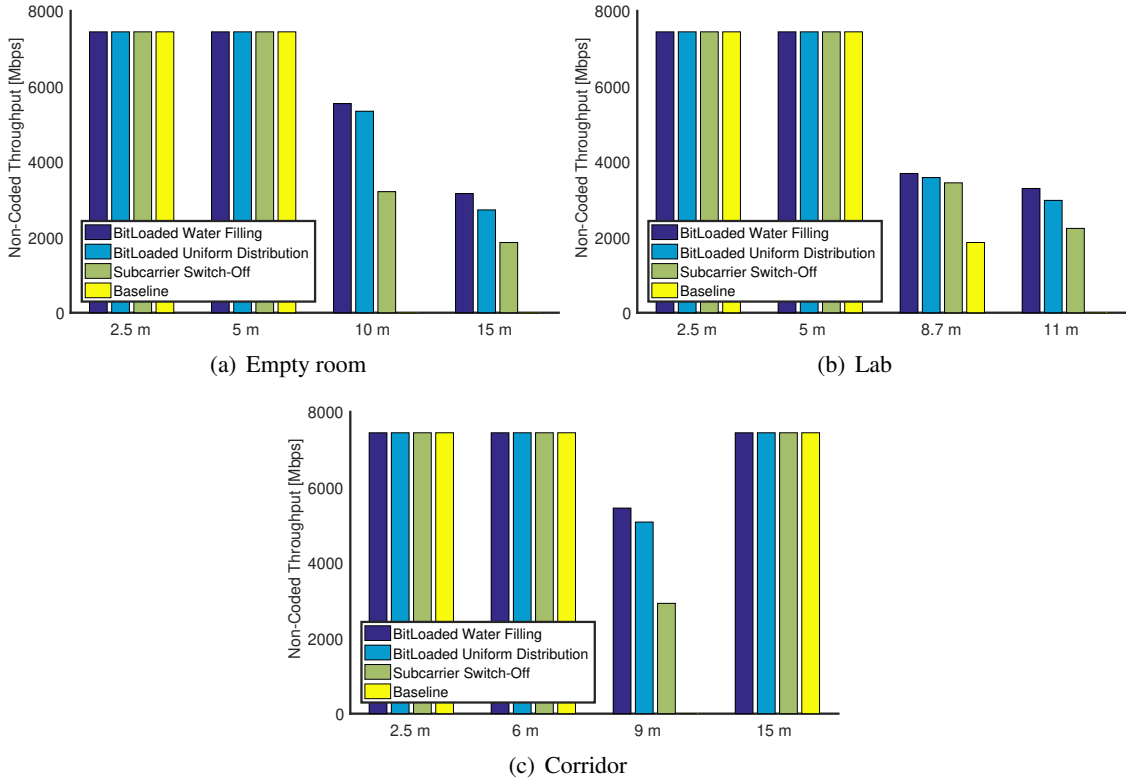


Figure 4.7: Achievable non-coded throughput for increasing distances and different frequency selective techniques.

communication is possible. Subcarrier switch-off is also able to extend the range of the link for all three scenarios. However, it only achieves a lower data rate since having to use the same MCS on all subcarriers limits its performance.

Specifically, the limited range of the 60 GHz link prevents communication beyond 5 m on our empty room scenario. For 10 and 15 m, our three mechanisms are able to establish a link, whereas the baseline mechanism fails. Exploiting frequency selectivity, we are capable of transmitting data up to 15 m with a non-coded rate of at least 2000 Mbps and a coded rate of at least 630 Mbps for this scenario. That is, we triple the range of the link. As expected, our techniques based on bit loading achieve even better performance. In our lab scenario, the gain in terms of range is lower since the baseline scheme is able to transmit data up to a distance of 8.7 m in the non-coded case. Still, our frequency selective techniques are able to achieve significantly higher rates, and transmit data up to a distance of 11 m. Moreover, in this scenario we observe that the impact of coding limits the range of the baseline down to only 5 m, since the BER is too high beyond that distance. In contrast, the frequency selective techniques that we consider still reach the aforementioned link distance of 11 m. Finally, the results for our corridor scenario differ significantly from the lab and the empty room cases. As listed in Table 4.6, the corridor is the most frequently selective scenario out of the settings that we consider in this work, featuring up

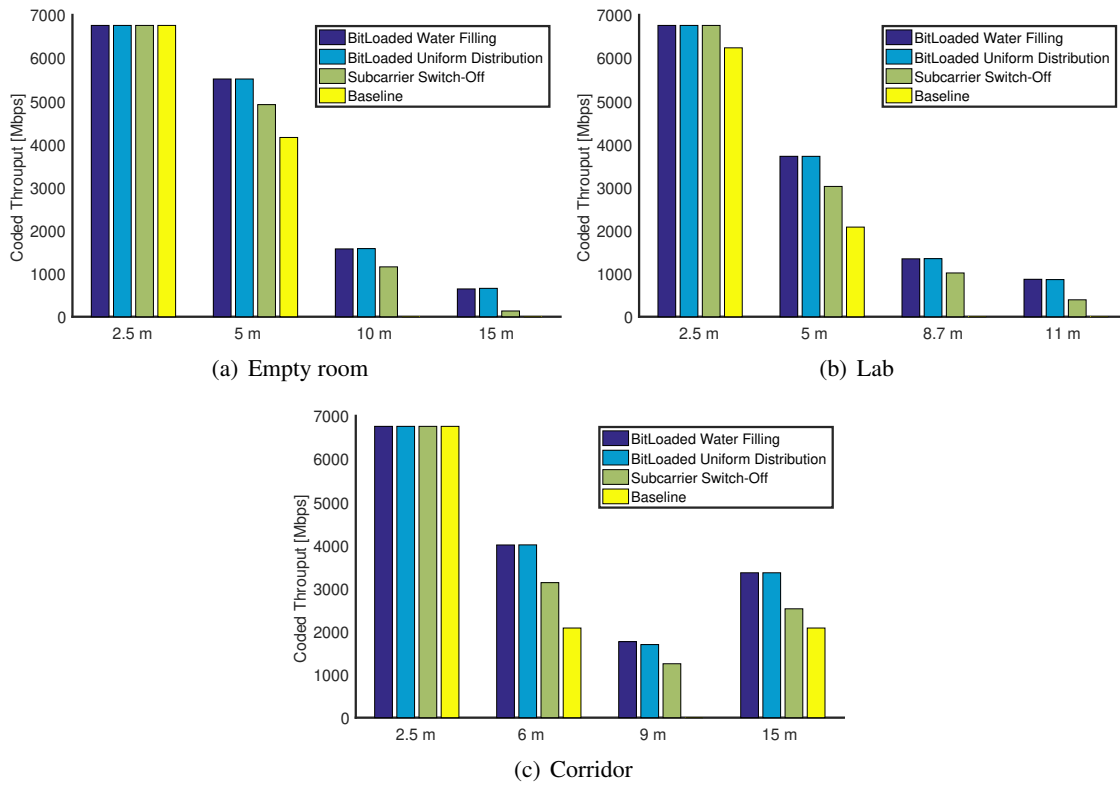


Figure 4.8: Achievable coded throughput for increasing distances and different frequency selective techniques.

to 22 dB difference in terms of received power among subcarriers at 9 m. As a result, the baseline mechanism is unable to establish a link at that position, while our other techniques can operate normally. Most interestingly, at 15 m the overall received energy increases (c.f. Figure 4.6) and the channel becomes more flat in the frequency domain, enabling the baseline scheme to operate again. The underlying reason is that the destructive reflection at 9 m disappears when moving to the next position further away, resulting in a flat channel response. Moreover, constructive reflections at 15 m boost the received power, allowing for similar coded throughput rates at 15 m and at 6 m. All in all, we observe that frequency-selective techniques can improve the range of a 60 GHz link on average by 50%, ranging from a 26% range increase in the lab scenario to $3\times$ longer links in the empty room case. Moreover, such techniques improve the achievable data rate significantly in all of our scenarios. For the coded case, range and rate improvements are even larger since the impact of coding limits the performance of the baseline. Results are roughly in the same order of magnitude but, for instance, the range improvement in the lab case increases from 26% to $2\times$. Hence, we conclude that frequency-selective techniques can be very beneficial for 60 GHz communication.

4.3.2. Energy Saving

Next, we analyze the performance of the techniques described in Section 4.1.1 for a range of different power margins. This provides insights regarding potential energy savings when using them. Figure 4.9 depicts the achievable rate for positions one and two in each of our scenarios. We compute the rate for a range of transmit power values. A power saving of 0 dB stands for transmission at the maximum power that regulators allow. Negative power saving values represent lower transmission power values. For instance, a power saving of -10 dB means that the transmitter is transmitting 10 dB below the maximum allowed transmission power. The depicted rate values are for the coded case (c.f. Section 4.1.2).

Figure 4.9 shows that frequency selectivity plays a role in terms of energy saving, since each of the scenarios performs differently. Regarding the individual techniques, we observe that both bit loading approaches perform almost identically across all scenarios. That is, the impact of water filling compared to a uniform power distribution is limited. The results for those two cases actually differ by tens of megabits per second, but this difference is not visible in Figure 4.9 because the scale of the y-axis ranges from 0 Mbps to 7000 Mbps. The similarity in terms of results of both bit loading approaches is due to the large bandwidth of the system. Given the large number of subcarriers, the amount of energy that waterfilling spreads among subcarriers is small. As a result, it is typically not enough to reach the next MCS threshold, and thus results in the same throughput. Finally, Figure 4.9 is consistent with previous sections, since we observe that our bitloading techniques generally outperform subcarrier switch-off. Only for the scenarios for which the SNR is anyhow very high, the results of both approaches match.

Regarding the individual scenarios, for the measured position at 2.5 m in the empty room case we observe that all of our techniques could achieve the maximum throughput with a transmit power at least 3 dB below the maximum allowable transmit power. If the link does not need to operate at the maximum rate, further power savings are possible. The subcarrier switch-off technique occasionally matches the baseline but requires on average slightly less transmit power to achieve the same rate. In contrast, the bit loading techniques consistently outperform the baseline, requiring on average 2 dB less to achieve the same rate. At 5 m, gains with respect to the baseline increase since the larger distance results in more impairments on certain subcarriers. Specifically, the bit loading techniques require 3 dB less transmit power than the baseline. Also, the subcarrier switch-off approach is now consistently better than the baseline, with a difference of 2 dB in terms of transmit power. Most interestingly, Figure 4.9(b) shows that the baseline cannot operate beyond a power saving of 12 dB but our other techniques can transmit data up to a power saving of 17 dB. Next, Figure 4.9(d) depicts the potential energy saving at 2.5 m in the lab scenario. In this case, bit loading can yield a saving of up to 7 dB compared to the baseline. That is, bit loading can achieve the same rate than the baseline with less than a quarter of the transmit power. Subcarrier switch-off performs similarly, resulting in 6 dB power saving. The reason behind such large power savings is that for the particular case in Figure 4.9(d), many subcarriers are operating close to the threshold of the next MCS level when using the baseline. Thus, a slight

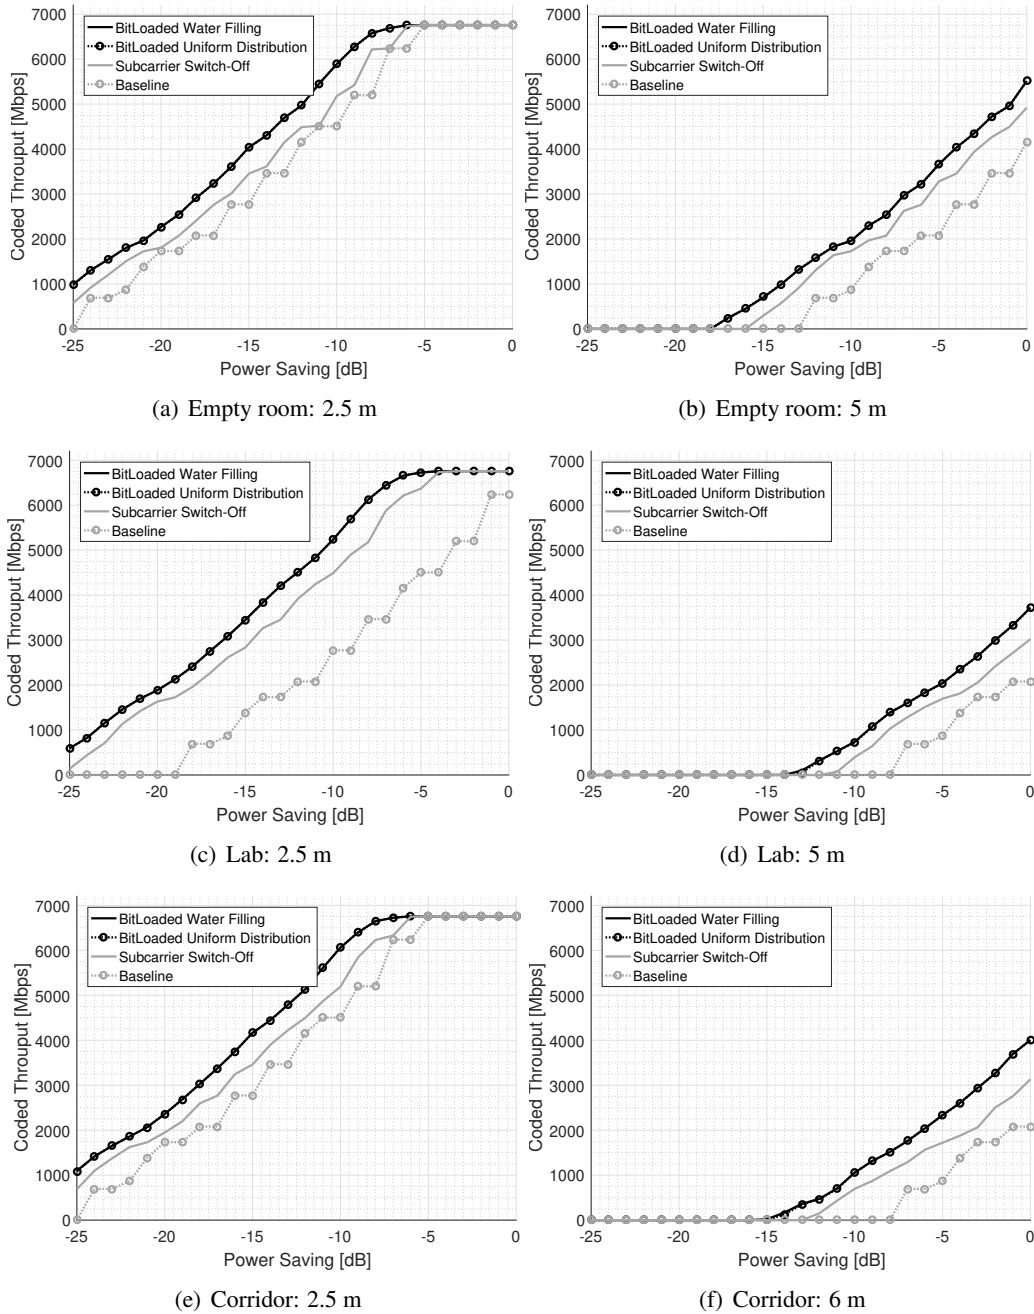


Figure 4.9: Achievable energy savings for different frequency selective techniques.

improvement based on our techniques allows many subcarriers to switch to a higher MCS, thus achieving much higher throughput. This does not occur for the next position in the lab scenario (c.f. Figure 4.9(d)), resulting again in similar gains as for the empty room scenario. Concretely, bit loading can save on average 3 dB transmit power, while subcarrier switch-off only saves 2 dB. Finally, in the corridor scenario we obtain similar results as for the empty room case. While subcarrier switch-off is consistently better than the baseline at the first position, the difference is just 1 dB. In contrast, bit loading achieves again a 3 dB improvement. As expected, the rates at the second position are significantly lower due to the increased distance. In this case, bit loading achieves 4 dB power savings and subcarrier switch-off 2 dB. All in all, we conclude that bit loading is key to achieve significant energy savings. While energy distribution techniques such as waterfilling provide some benefit, their impact is minimal in the case of very wide channels. Subcarrier switch-off is beneficial, too, but performs consistently worse than bit loading. Our above results are for the coded throughput case only. We observe a similar general behavior for the non-coded throughput. However, when coding is not taken into account, gains may be inflated for very low SNRs, as these weak links contribute to the throughput while no MCS is available that can operate at these SNRs. (c.f. Section 4.3.1).

4.3.3. Antenna Array Comparison

Our results in Section 4.3 are based on a 60 GHz testbed equipped with horn antennas, which focus the transmit power into a given direction. However, such antennas are bulky and not electronically steerable, which makes them unsuitable for commercial hardware. Instead, commercial 60 GHz devices typically use phased antenna arrays which may result in more unexpected reflections, which in turn may lead to stronger frequency selectivity. In order to assess whether this effect has a significant impact on our analysis, in this section we present a first-of-its-kind study regarding the frequency selective behavior of 60 GHz channels when using phased antenna arrays.

Unfortunately, phased antenna arrays for experimental purposes were not commercially available at the time of writing. This hindered us from carrying out the bulk of our experiments on such a testbed. However, in cooperation with the IMEC Research Center in Belgium, we were able to perform detailed channel measurements on a phased antenna array prototype. In particular, we use an antenna array with 2×8 elements and measure channels with a bandwidth of 1.8 GHz, similarly to our experiments in Section 4.3. We place the transmitter and the receiver in a lab at a distance of 4 m. We then design two different beam patterns using the antenna array. Our first beam pattern has an ideal shape, that is, very narrow and with negligible sidelobes. We can generate such a beam pattern based on the advanced capabilities of the phased array beamforming transceiver module from IMEC [79]. However, commercial devices cannot generate such directional beam patterns, as seen in Chapter 3, so we design a second beam pattern featuring a wider main lobe and higher sidelobe energy level.

Figure 4.10 shows the results achieved with these antenna array measurements. We clearly




Figure 4.10: Antenna array results

observe the impact of beamwidth and sidelobes—while the potential power savings for the narrow case are about 4 dB, savings increase to 9 dB for the case with significant sidelobes. We again confirm that frequency selectivity is mainly caused by reflections as we obtain a difference of 5 dB with the use of different beam patterns. Hence, this validates our assumptions regarding the impact of sidelobes. Moreover, the frequency selectivity of our phased antenna array measurements is similar to the case with horn antennas. This means that the gains that can be achieved with phased antenna arrays when using techniques that exploit frequency selectivity are going to be equivalent or even higher than our results in Section 4.3.

In the case of using COTS systems, we can confirm that the gains of using frequency selectivity techniques are going to be larger than the ones seen in this chapter. The antenna array used in this section is capable of synthesizing very clean and narrow beam patterns while reducing the number of sidelobes. However, COTS devices, as the ones analyzed in Chapter 3, experience a large number of sidelobes which will create very rich multipath components, creating a very frequency selective channel.

We can thus conclude that frequency selectivity plays a significant role on COTS 60 GHz devices and that these techniques could be used for enlarging mmWave link distances or save transmitted energy.

4.4. Discussion

Our channel measurements question common assumptions in the recent 60 GHz networking literature. Such related work often deals with the 60 GHz band from a networking perspective, thus often resorting to assumptions such as flat channels and free-space path loss models. However, early work on propagation characteristics [42] that deals with the 60 GHz band from a physical layer perspective suggests that such assumptions are not realistic. In this work, we bring together both perspectives by considering typical indoor scenarios for 60 GHz network deployments based on IEEE 802.11ad. We show (a) that the free-space path loss model is not accurate in such deployments, and (b) that the impact of frequency selectivity becomes evident as soon as the distance between transmitter and receiver is beyond a few meters.

From this work, we derive two crucial insights. First, we observe that, while the throughput improvements that we can achieve when using frequency selective techniques may be limited, the gains in terms of range extension are significant for low SNRs. We can extend the range of a 60 GHz link by up to $3\times$. In other words, frequency selective techniques can operate when the basic physical layer in IEEE 802.11ad cannot. This results in a highly beneficial symbiosis—the longer the link, the stronger the impact of the very high path loss in the 60 GHz band but also the higher the benefit of frequency selective techniques. Our second insight is that frequency selective techniques help to achieve a uniform coverage area. As discussed in Section 4.3, basic approaches may not be able to establish a link at a distance d_1 due to strong frequency selective fading on certain subcarriers, but may be able to operate at distance $d_2 > d_1$ if the channel response becomes flat again. That is, the area of coverage may contain areas with no coverage at all even if no obstacle blocks the link. We show that frequency selective techniques can help to homogenize the coverage area of a 60 GHz transmitter.

Although we focus in OFDM systems, it is important to mention that the link outages that we experience during this study are also going to happen in single-carrier devices. Single-carrier devices also adapt their transmission characteristics to the weakest portion of their channel, limiting their transmission range. In this chapter we want to encourage vendors and developers to adapt their single-carrier 60 GHz systems to OFDM. Once these devices support OFDM they should implement the frequency selective techniques exposed here in order to extend their link range and save energy in their transmissions. The 10 dB energy savings that we demonstrate in this chapter would be even more useful in battery-driven devices like smartphones.

Chapter 5

Optimizing Spatial Reuse

mmWave devices are known to have directional beam patterns that can achieve very high spatial reuse allowing many simultaneous links in the same area, providing almost interference-free communications. In contrast, the omni-directional transmissions and rich multi-path environment in legacy ISM bands require a MAC layer to avoid concurrent transmissions to prevent collisions. In practice, however, the difference between mmWave and lower frequency bands is not that pronounced. For simplicity and to lower manufacturing costs, current mmWave COTS use antenna arrays with beam shapes that are much wider and with many more side lobes than theory suggests, as seen in previous Chapter 3. These create enough interference so as to prevent spatial reuse in most practical scenarios. Furthermore, current IEEE 802.11ad devices use a set of predefined beam patterns from a fixed codebook and do not adapt the patterns to the current RF environment.

Beamtraining is the process where mmWave nodes select the most suitable beam pattern to be used for communications, which in the case of IEEE 802.11ad, is the one that achieves the highest SNR. In this chapter, we propose a novel approach for beamtraining that takes into account not only the SNR, but the possible interference that a beam pattern incurs. To this end, we design a centralized coordinated system that chooses the most efficient beam pattern for each STA that improves throughput for all the users. Our system requires only changes to the APs and works with unmodified clients. We use two alternatives for the beam patterns: 1) modifying the choice of Tx-beam pattern while using the default omni-directional Rx-beam pattern of current COTS devices and 2) modifying also the Rx-beam pattern to be equal to the chosen Tx-beam pattern.

Figure 5.1 shows a simple example consisting on two APs and two STAs. Each of the STAs can use two different beam patterns (black and light blue), which determine the power with which the signals are received at both APs. Note that the beam patterns are examples of actual beam patterns used by these devices [80].

The problem that we want to solve in Figure 5.1 is the following: the power arriving from STA_1 to AP_1 (named A) and the power arriving from STA_2 to AP_2 (named D) should be maximized, while the interference, which is the power arriving from STA_1 to AP_2 (named B) and

the power arriving from STA₂ to AP₁ (named C) should be minimized. Following this example, STA₁ should use the light blue beam pattern. Even though the black beam pattern results in a slightly higher power towards AP₁, it also generates much more interference at AP₂. This interference is negligible if STA₁ uses the light blue beam pattern. On the other hand, STA₂ should use the black beam pattern, as this is the one that results in higher power at AP₂ and lower interference at AP₁.

To solve the problem of beam pattern selection, we investigate four different mechanisms (detailed in Sec. 5.2) and compare them with the baseline IEEE 802.11ad protocol:

- **IEEE 802.11ad (baseline):** The 802.11ad standard does not take interference into account and simply selects the beam patterns maximizing A and D.
- **Weighted SIR Fairness:** This mechanism takes into account all the combinations of beam patterns of all STAs and choose the ones maximizing the sum of Signal-to-Interference Ratios (SIRs), i.e., $(A - C) + (D - B)$. This is optimal for total channel capacity calculated according to the Shannon-Hartley theorem.
- **Argmaxmin:** From all the combinations of beam patterns, the ‘Argmaxmin’ mechanism selects the combination that maximizes the minimum SIR, i.e., the one where $\min(A - C, D - B)$ is largest. This provides fairness by maximizing the rate of the weakest link.
- **Power Threshold:** For each STA, this mechanism determines the set of beam patterns for which A is larger than a given threshold (corresponding to a minimum desired rate or MCS), and from this set selects the beam pattern that minimizes the interference B. This is repeated independently for all STAs.




Figure 5.1: Beamtraining example: A and D wants to be maximized while B and C minimized.

- **Interference Threshold:** This mechanism checks for which beam patterns the interference B is below a threshold (typically selected to be the carrier sense threshold to enable concurrent communication), and from those selects the one that maximizes A . This is repeated for all STAs to find the beam pattern that maximizes D given that C must be below the threshold.

The main contributions of this chapter are the following:

- We first analyze how the selection of different beam patterns in a 2-link scenario may affect the aggregated throughput.
- We design four beamtraining algorithms to increase aggregated throughput in dense deployments by improving spatial reuse.
- We implement the different algorithms in COTS devices using alternatively the default omni-directional Rx-beam pattern and the modified Rx-beam pattern to be the equal to the Tx-beam pattern. We modify the firmware of the devices to extract all the information needed to select and fix the utilized beam patterns.
- We then analyze the performance of the different algorithms, first in a proof of concept scenario and then evaluate them in a real-world environment. There we test how the different algorithms behave for different positions and link combinations in an open-space office. We check both TCP and User Datagram Protocol (UDP) as well as uplink, downlink and bidirectional traffic.

This chapter is organized as follows. We first study the impact of neighboring nodes beam pattern selection in Section 5.1. We then explain the different beamtraining mechanisms in more detail in Section 5.2. The experimental setup, including the router implementation, network configuration, methodology and scenarios are presented in Section 5.3. The evaluation of the different algorithms and scenarios is carried out in Section 5.4 and discussed in Section 5.5.

5.1. Impact of Neighbouring Nodes Beam Selection

As previously discussed, the IEEE 802.11ad beamtraining mechanism is purely based on the SNR achieved by the different beam patterns of the two devices forming the link, without taking into account any other devices with which they share the wireless medium [20]. This beam pattern selection mechanism is optimal when there is no spectrum sharing among different networks, but can lead to strong interference when different networks share the wireless medium. In this set of measurements, we manually select the transmit beam patterns and compare the overall network performance with respect to the default selected beam pattern. We do this first round of measurements in order to motivate the possibility of developing more efficient beam pattern selection mechanisms in terms of spatial reuse.

5.1.1. Measurement Setup

In order to study this phenomenon we carry the following study: we set up two parallel links and iterate among all the possible beam pattern configurations from both STAs and APs. We repeat these measurement iterations for separation distances of 2, 5 and 15 meters, first iterating among the different transmit beam patterns while using the default omni-directional receive beam pattern, and then we iterate over both transmit and receive beam patterns (which for simplicity we select to be the same, assuming that the channel is reciprocal). This measurement is realized in a symmetric way, meaning that each of the diagonal items use the same beam pattern from an interference and power point of view, as they share the same relative angles. This way we simplify the beam pattern selection measurements reducing the number of combinations and thus the overall measurement time. Doing so, we iterate among all the pair-wise possible beam pattern configurations from the codebook, using the same beam pattern in the left link's AP and the right link's STA, at the same time that we fix the same beam pattern in right link's AP and the left link's STA, resulting in a total of 1296 beam pattern combinations per measured link distance.

5.1.2. Beam pattern selection comparison

Figure 5.2 shows the resulting aggregated throughput for each of the beam pattern combinations, while Table 5.1 compares the default throughput with respect to the maximum achieved rate among the different measurement configurations. In the figure, the columns from left to right show the distances 2 m, 5 m, and 15 m. A throughput lower than 1000 Mbps is represented with the same dark blue and a throughput larger than 2500 Mbps is represented with the same yellow color. White represents the cases where link cannot be established. At a first glance, it can be seen that the default beam pattern selection mechanism can be improved, as in every scenario it is possible to achieve a larger throughput when selecting beam patterns different from the default configuration. When distances are short, the gains are not very large as nodes usually sense each other's transmission, due to the large interference power at other stations independently of the beam pattern they use. However, this changes when the distance between the links increases, as we can see that selecting a different beam pattern can avoid interference and boost the throughput, even doubling it for the case of 5 m distance. Fixing the receive beam pattern can further improve the spatial sharing. In this case, not only the interference emitted towards other nodes decreases, but also the interference observed through the station's receive beam pattern is reduced. This is extremely important in the case where the links are close to each other.

It can be seen that the throughput is very uniform among the beam pattern combinations for a distance of 2 m, as the amount of interference causes the stations to carrier sense with each other. Nevertheless, some gains can be achieved when fixing the transmission beam pattern. In addition, changing the omni-directional beam pattern into a directional one gives the possibility of increasing the throughput around by 700 Mbps. At 5 m, there are some beam patterns that

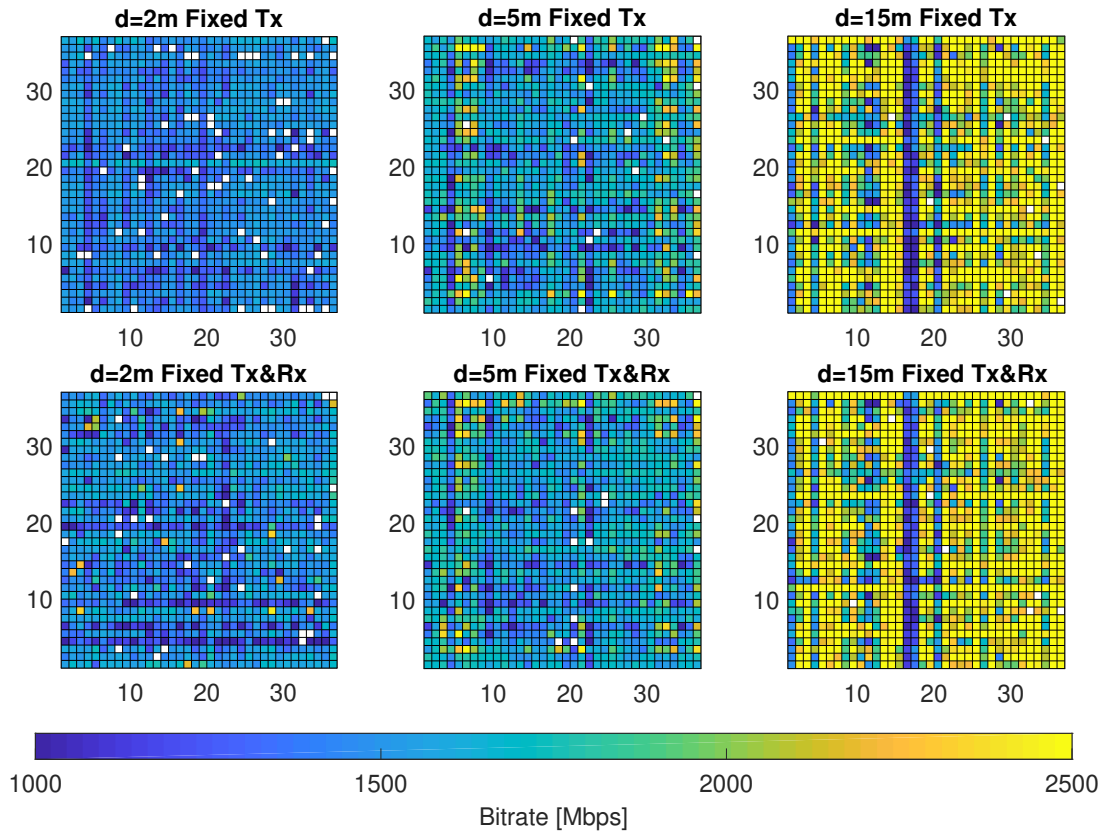


Figure 5.2: Aggregated throughput for each of the beam pattern combinations: the top row uses omni-directional Rx-beam pattern, while the bottom row modifies both Tx- and Rx-beam pattern.

provide significant gains compared to the default. For example we can see higher throughputs in the columns corresponding to beam patterns 5, 6 and 7. These beam patterns transmit less power at the angle where the neighboring stations are located. In contrast, columns 4, 9 or 22 show how the bitrate decreases due to the high power transmitted towards the neighbors. Fixing the receiving beam pattern does not significantly affect the performance of the link in this case, and the gains are similar to the ones achieved when fixing only the transmit pattern. For both cases, the maximum throughput can be doubled, providing gains of 1400 Mbps. For distance equal to 15 m, we can see many cases with large gains, and some combinations where the link quality is degraded. Specifically, it can be seen that beam patterns 16, 17 and 20 are suboptimal, having 1400 Mbps less aggregated throughput. Beam pattern with index 20 is the one used by the default configuration, showing that even though the received power in the single link case is large, the interference created is also large. For this link distance, fixing the receiving beam pattern can increase the link throughput by another 100 Mbps.

Table 5.1: Maximum bitrate for different selected beam patterns in Figure 5.2.

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

It is also important to note that, given the asymmetric beam patterns of these COTS devices, the behavior of the beam pattern selection is also asymmetric, so a pairwise switch of the beam patterns leads to different rates. Even though we have bi-directional communication in the link, the acknowledgments require less time than the payload, making the interference-free operation more important in the transmitter node than in the receiving one.

These measurements confirm that the default beamtraining mechanism can be improved as it is not optimum when different links share the collision domain, motivating the different beamtraining techniques that we show and test during this chapter.

5.2. Beamtraining Mechanisms

In this section we explain the different beam training algorithms developed for COTS mmWave devices. Instead of using the strongest beam pattern to reach their APs, the STAs' beam patterns will be chosen to balance the received power at the AP and the interference created to other links. Furthermore, we aim to maximize the aggregated throughput while taking into account fairness, to avoid possible starvation for some of the STA. The algorithms work with unmodified STAs, requiring only a network of APs with centralized control. This way, any device entering the network can take advantage of this.

In order to establish directional communications following the IEEE 802.11ad standard, the nodes have to perform the beamtraining procedure [20, 22]. This process can be seen in Figure 5.3(a) and consists on the following: first, the AP tests all the possible Tx-beam patterns of its codebook with the initiator sector sweep, while the STAs listen with omni-directional Rx-beam pattern. Then, the STAs reply similarly, sending replies over all their possible Tx-beam patterns as the responder sector sweep. These replies also include which one was the best AP beam pattern. Finally, the AP acknowledges which is the best beam pattern to be used by the STA. After this short process, both ends of the link know which are the beam patterns leading to the highest SNR.

Our system needs to do an extra step in order to obtain all the information from the different STAs. This extra step can be seen in Figure 5.3(b) and consists on the following: All of the STAs follow the IEEE 802.11ad beam training process as previously described, sending frames through the different beam patterns. In parallel, all the APs overhear these frames and store the received power from every STA. Once all STAs finish their beam training, the APs will exchange among

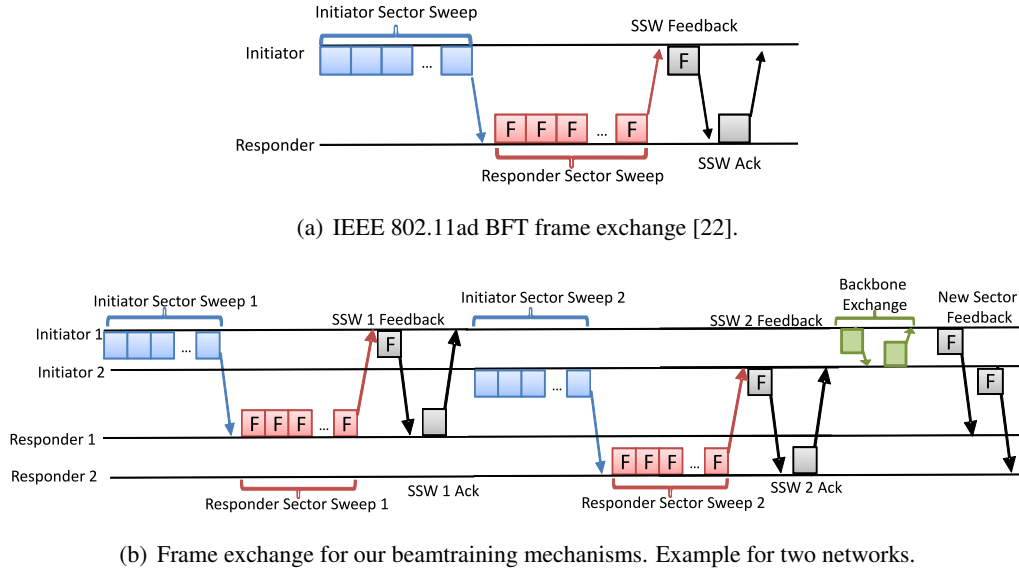


Figure 5.3: Beamtraining frame exchange Mechanisms. Frames marked with ‘F’ represent feedback frames.

them the received power information, allowing them to know how much their STAs are interfering with other APs. Once the APs compute which are the most suitable beam patterns for their STAs, they send feedback to their STAs specifying which beam pattern to use.

To explain our algorithms we use the following notation:

- \mathcal{D} is the set of beam patterns for each STA.
- P_{xn}^i is the power (in decibel units) received by AP_n from STA_x using the beam pattern i , with $i \in \mathcal{D}$.
- A is the set of APs.
- S_n is the set of STAs connected to AP_n , where $n \in A$.
- b_x^i is a binary variable that indicates whether STA_x is using the beam pattern i , $i \in \mathcal{D}$. b_x^i is equal to 1 if x selects the beam pattern i , and 0 otherwise. Only one beam pattern can be used at a time, i.e., $\sum_{i \in \mathcal{D}} b_x^i = 1, \forall x$.
- B is a matrix with matrix elements b_x^i , indicating a specific configuration of the beam patterns used by each of the stations.
- \mathcal{B} is the set of all possible beam pattern allocations B .
- B_x represents the x -th row of B , i.e., the vector of b_x^i for STA_x and $i \in \mathcal{D}$.

5.2.1. Baseline: IEEE 802.11ad beamtraining

As described previously, when an AP and a STA want to establish a connection, they will use the Tx-beam pattern that provides the highest power. The chosen beam pattern B_x^* is given by:

$$B_x^* = \arg \max_{B_x \in \mathcal{B}_x} \sum_{i \in \mathcal{D}} b_x^i P_{xn}^i \quad (5.1)$$

for all $n \in A$ and $x \in S_n$.

5.2.2. Weighted SIR Fairness

This mechanism chooses the most efficient beam pattern in terms of SIR. It computes all the possible combinations of beam patterns among all the STAs and selects the combination B^* resulting in the highest SIR:

$$B^* = \arg \max_{B \in \mathcal{B}} \sum_{n \in A} \sum_{x \in S_n} \left(\sum_{i \in \mathcal{D}} b_x^i P_{xn}^i - \sum_{m \in A \setminus \{n\}} \sum_{y \in S_m} \sum_{j \in \mathcal{D}} b_y^j P_{ym}^j \right) \quad (5.2)$$

This mechanism could lead into STAs switching off, as sometimes the result that maximizes Eq. (5.2) will have some links with very low SNR, meaning those links would disconnect. To obtain a fair allocation, we apply the α -fairness method as seen in Eq. (5.3), where α is chosen between 0 and 1. The higher the value, the more fair the allocation. In our implementation we choose an α equal to 0.99.

$$B^* = \arg \max_{B \in \mathcal{B}} \sum_{n \in A} \sum_{x \in S_n} \left(\sum_{i \in \mathcal{D}} b_x^i P_{xn}^i - \sum_{m \in A \setminus \{n\}} \sum_{y \in S_m} \sum_{j \in \mathcal{D}} b_y^j P_{ym}^j \right)^{1-\alpha} \quad (5.3)$$

5.2.3. Argmaxmin

This mechanism maximizes the SIR of the weakest link among all the links. It computes the SIR for all the possible beam pattern combinations and then selects the beam pattern combination B^* that maximizes the minimum of the link SIRs.

$$B^* = \arg \max_{B \in \mathcal{B}} \left(\min_{\substack{x \in S_n \\ n \in A}} \left(\sum_{i \in \mathcal{D}} b_x^i P_{xn}^i - \sum_{m \in A \setminus \{n\}} \sum_{y \in S_m} \sum_{j \in \mathcal{D}} b_y^j P_{ym}^j \right) \right) \quad (5.4)$$

5.2.4. Power Threshold

In wireless communication systems, whether a certain MCS can be used depends on the SNR, i.e., achieving a certain minimum rate requires the SNR to be higher than the corresponding threshold. We build a mechanism that uses such a threshold in order to determine a set of beam patterns satisfying a power threshold \bar{P} , and then selecting the one interfering the least with the neighboring nodes (Eq. (5.5)). If there is no group of beam patterns that satisfy the threshold for a given \bar{P} , we decrease \bar{P} such that $\exists i \in \mathcal{D} \mid P_{xm}^i > \bar{P}$.

$$B_x^* = \arg \min_{B_x \in \mathcal{B}_x} \sum_{b \in A \setminus \{n\}} \sum_{y \in S_b} \sum_{i \in \mathcal{D} \mid P_{xn}^i > \bar{P}} b_x^i P_{ym}^i \quad (5.5)$$

for all $n \in A$ and $x \in S(n)$. Since this selection could imply link starvation, again we apply the α -fairness method as in Sec. 5.2.2, resulting in

$$B_x^* = \arg \min_{B_x \in \mathcal{B}_x} \sum_{b \in A \setminus \{n\}} \sum_{y \in S_b} \sum_{i \in \mathcal{D} \mid P_{xn}^i > \bar{P}} (b_x^i P_{ym}^i)^{1-\alpha} \quad (5.6)$$

This technique minimizes interference while guaranteeing a certain minimum data rate.

5.2.5. Interference Threshold

To maximize the spatial reuse we need to find an interference threshold \bar{I} that is below the carrier sensing level of other APs. We select the group of beam patterns that satisfy the threshold \bar{I} which would allow parallel links, and then choose the one that maximizes the rate. If this threshold cannot be satisfied, we apply the default beamtraining mechanism.

$$B_x^* = \begin{cases} \arg \max_{B_x \in \mathcal{B}_x} \sum_{i \in \mathcal{D} \mid P_{xm}^i < \bar{I}, m \in A \setminus \{n\}} P_{xn}^i, & \text{if } \exists i \in \mathcal{D} \mid P_{xm}^i < \bar{I} \\ \arg \max_{B_x \in \mathcal{B}_x} \sum_{i \in \mathcal{D}} b_x^i P_{xn}^i, & \text{otherwise} \end{cases} \quad (5.7)$$

for all $n \in A$ and $x \in S_n$

5.3. Experimental Setup

In this section, we first discuss the implementation details, from the custom firmware installation to the beam pattern selection. We then present the methodology as well as the different scenarios in which we test our mechanisms.

5.3.1. Router Implementation

We use the TP-Link Talon AD7200 router for the implementation of these techniques. This router is equipped with the QCA9500 60 GHz chipset which comes with a 32 element antenna

array and a predefined codebook of 34 Tx-beam patterns.

With the default firmware running on the QCA9500 60 GHz chipset we cannot extract or set any information regarding the beamtraining or beam patterns. We thus use the Talon Tools framework [81] to:

1. Obtain beamtraining information: We can access the IEEE 802.11ad driver and record the received power for each of the beam patterns received by the router, together with the MAC address of the transmitting node.
2. Set the Tx-beam pattern to be used: We can fix the corresponding index of the utilized beam pattern in the QCA9500 chipset.
3. Modify the omni-directional Rx-beam pattern to one of the predefined beam patterns from the codebook: We can load one beam pattern from the codebook and write it to the chipset to be used as Rx-beam pattern.

This way, the mmWave APs can share how much energy they are receiving from each of the different Tx-beam patterns from each of the different STAs, and choose the pattern which is best for each STA to use according to our algorithms.

5.3.2. Methodology

The use of our mechanisms in a real world system is simple. The only constraint is that the different APs are connected and have a minimum processing power, which is usually satisfied for example in managed WLAN deployments. Otherwise, the APs can share their data wirelessly.

For our experimental analysis, we have connected all our APs to a LAN and carried all the beamtraining mechanisms and control of the routers from a central server. In order to measure the performance of the different algorithms, we perform 30 second iPerf measurements for TCP and UDP traffic. We also modify the traffic flow direction, having uplink, downlink and bidirectional communications. We will first run the measurements with the default beamtraining, then iterate for the different sets of beam patterns that our mechanisms provide, fixing first the Tx-beam pattern with omni-directional Rx-beam pattern and then fixing both Tx and Rx-beam patterns.

Given the Carrier Sense Multiple Access (CSMA) used by IEEE 802.11ad networks, we have simplified our system to have AP-STA pairs (rather than having multiple STAs per AP). An AP with several STAs would lead to similar results given the multiplexing among different STAs, but would significantly increase the complexity of our deployment.

5.3.3. Scenarios

All the algorithms explained in Section 5.2 are tested in two different scenarios. First, we validate the mechanisms in an empty auditorium and then check their behavior in a real open-space office environment.

5.3.3.1. Empty Auditorium

To validate our mechanisms and evaluate their performance we chose an empty auditorium of size 11 x 21 meters where we place two parallel links. For both links, the distance between AP and STA is equal to 3.2 m and the distance between the links increases from 1.6 to 9.6 meters with steps of 1.6 m for a total of 6 measured positions.

5.3.3.2. Office Environment

Our second environment is a open-space office environment that is in active use, with size 7.4 x 13.5 meters as shown in Figure 5.5. The APs are placed under the ceiling of the room while the STAs are placed on top of regular desks in the working places of the office. In this environment we first do an extensive study of the mechanisms with two parallel links (using AP_1 and AP_2) and then study a few examples with three parallel links (where we add the AP_0 – STA_0 link).

5.4. Evaluation

In this section we evaluate the behavior of the four different mechanisms discussed in Section 5.2. We first validate the algorithms in a simple setup and evaluate a more complex office scenario. For the latter we study both setups with two and three parallel links.

5.4.1. Empty Auditorium

The performance of the four different algorithms in the auditorium setup can be seen in Figure 5.4. As expected, in general the gains for the different algorithms increase together with the link distance, implying that with larger distances it is easier to achieve spatial reuse.

We also observe that UDP traffic obtains higher gains than TCP traffic; gains which are enlarged even more when choosing the Rx-beam pattern to be the same as the Tx-beam pattern. In contrast, for TCP traffic happens the opposite, and higher gains are achieved when only the Tx-beam pattern is fixed and the omni-directional pattern is used for reception.

Comparing the different algorithms in this validation measurement set, we observe the following characteristics for each of the different mechanisms:

- **Weighted SIR Fairness:** This mechanism achieves large gains for UDP traffic when the links are far apart. It also achieves high gains for TCP traffic but only with directional Rx-beam pattern, as with omni-directional Rx-beam pattern the gains are negligible.
- **Argmaxmin:** This mechanism provides some level of fairness for the different protocols. It has similar gains for TCP and UDP traffic, both for directional and omni-directional Rx-beam pattern. There is an exception for the longest measured distance, where the gains increase up to 166% and 132% for UDP.

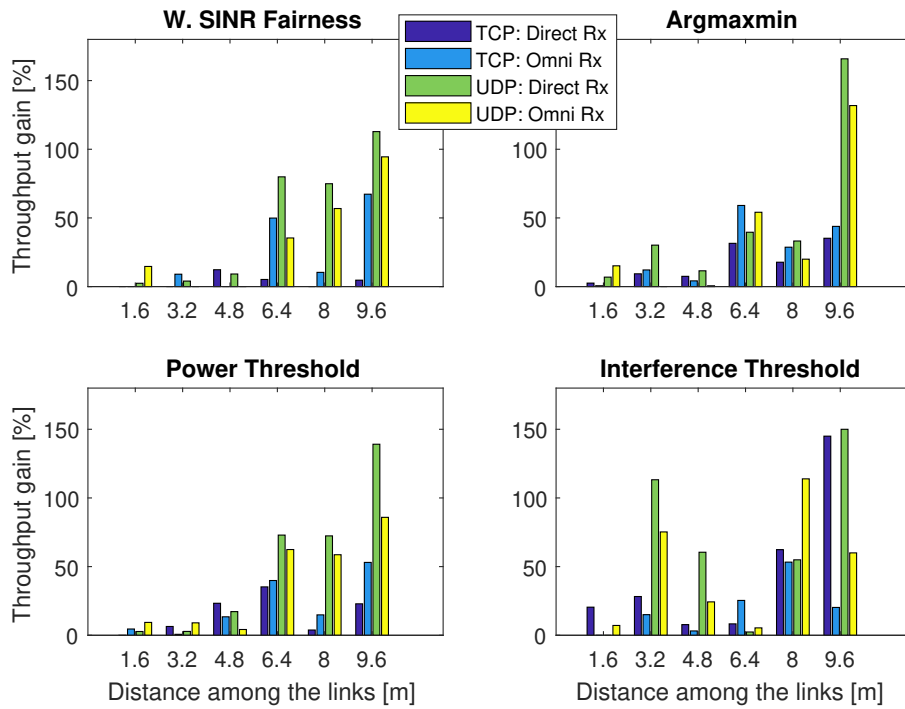


Figure 5.4: Mechanisms validation result: Empty Auditorium. Throughput gains for TCP and UDP traffic, when fixing Tx- and Rx-beam pattern and when fixing only Tx-beam pattern with the default omnidirectional Rx-beam pattern.

- Power Threshold:** For this scenario, this mechanism is the most promising and stable one of the four studied. Gains increase together with the distance. For UDP traffic we achieve higher gains with directional Rx-beam patterns, and for TCP we achieve higher gains with the omni-directional Rx-beam pattern.

- Interference Threshold:** This mechanism provides the largest gains among the algorithms when the links are close, but it does not have a reliable behavior for longer distances. It obtains very large gains for UDP traffic for a distance of 3.2 m, but then these gains decrease for the measurements taken at 4.8 m and 6.4 m. The gains increase again for 8 m and 9.6 m. The performance of this algorithm is also inconsistent with respect to directional versus omni-directional Rx-beam pattern.

5.4.2. Office Environment: 2 Links

After validating our measurements in the simple auditorium scenario, we test our algorithms in the more realistic office environment shown in Figure 5.5.

We summarize all the measurement results in Tables 5.2 and 5.3 for TCP and UDP traffic. Table 5.2 shows the average gain for each of the algorithms for TCP traffic. From this table we see that the most efficient beamtraining algorithm is the ‘Power Threshold’ algorithm, which achieves the maximum average gain value of 24.67% for bidirectional communications with the

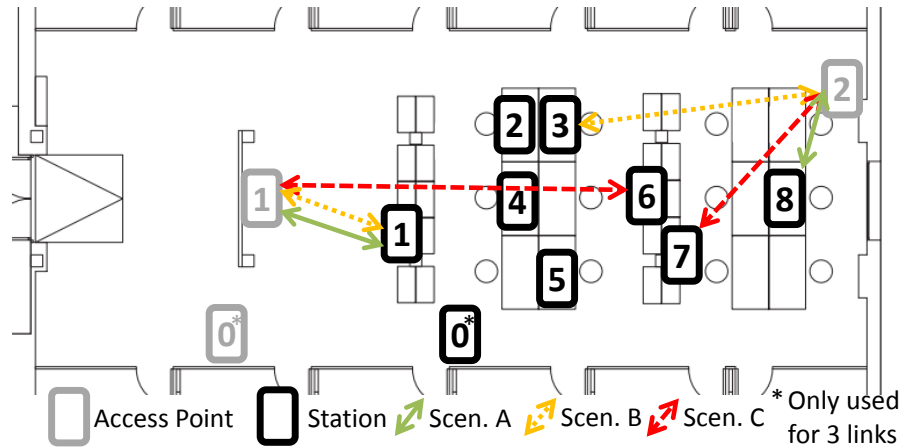


Figure 5.5: Office environment setup: Grey boxes represent APs while black boxes STAs. Scenarios A, B and C represents high, medium and low gain scenarios. AP₀ and STA₀ marked with a star are only used for the three link setup.

Table 5.2: Average gain for TCP traffic: Bidirectional (BD), Downlink (DL) and Uplink (UL) traffic, directional Rx-beam pattern (Direc Rx) and omni-directional Rx-beam pattern (Omni Rx).

	W. SIR Fair.	Argmaxmin	Power Th.	Int. Th.
BD Direc Rx	7.01%	7.65%	24.67%	22.71%
BD Omni Rx	13.86%	23.83%	21.56%	12.76%
DL Direc Rx	0.83%	5.19%	7.72%	9.94%
DL Omni Rx	5.00%	9.03%	1.81%	8.16%
UL Direc Rx	-2.53%	1.36%	2.67%	-1.45%
UL Omni Rx	-7.91%	4.81%	5.26%	-6.25%

use of directional Rx-beam pattern. In fact, for all of the studied algorithms the highest gains are achieved in the bidirectional case. In bidirectional communications, the channel access times become critical as the usage of the medium increases due to the long frames from both nodes (as opposed to the regular data-ACK patterns of uni-directional traffic), obtaining higher gains with the use of our algorithms.

We can also see in Table 5.2 that the ‘Argmaxmin’ together with the ‘Power threshold’ are the only two methods which have positive gains for all the communication directions. In contrast, ‘Weighted SIR Fairness’ and ‘Interference Threshold’ have some sub-10% losses for uplink traffic.

Table 5.3 summarizes the measurements for UDP traffic. As for the case of TCP traffic, the best performing algorithm is again the ‘Power Threshold’, having a maximum average gain of 23.70% of aggregated throughput. It is the only algorithm that has positive gains for all traffic directions with UDP traffic, as the others have some performance issues with uplink traffic, and some 1% losses in downlink.

In order to see how the algorithms behave for the different location combinations, Figure 5.6

Table 5.3: Average gain for UDP traffic: Bidirectional (BD), Downlink (DL) and Uplink (UL) traffic, directional Rx-beam pattern (Direc Rx) and omni-directional Rx-beam pattern (Omni Rx).

	W. SIR Fair.	Argmaxmin	Power Th.	Interf. Th.
BD Direc Rx	14.89%	5.48%	18.41%	1.62%
BD Omni Rx	11.37%	8.35%	23.70%	2.60%
DL Direc Rx	15.36%	-0.07%	15.14%	4.64%
DL Omni Rx	15.47%	3.50%	16.79%	-0.89%
UL Direc Rx	-0.37%	-2.55%	6.68%	-3.58%
UL Omni Rx	-2.90%	-1.42%	7.61%	-13.25%

shows the CDF for the different combinations in the bidirectional communication case. Here we can see that there are some link combinations where the gains are very high, but others where there are losses. The mechanism that finds highest gains is the ‘Argmaxmin’ mechanism, but in turn, it is also the algorithm that has higher losses. It is thus a ‘risky’ algorithm whose performance very much depends on the specific link configuration.

As depicted in the previous tables 5.2 and 5.3, we confirm that the best algorithm is ‘Power Threshold’, as it has gains for all the different link combinations and traffic types with the exception of two outliers: when using a directional Rx-beam pattern with UDP traffic and when using omni-directional Rx-beam pattern for TCP (as seen in the CDFs from Figure 5.6). These gains are due to the nature of the algorithm, which will always choose beam pattern combinations providing a sufficiently high MCS while minimizing interference. If the interference cannot be reduced while maintaining the threshold, the nodes carry out regular beamtraining without considering interference.

From the CDFs seen in Figure 5.6, there are link combinations where the algorithms work efficiently, and others where they do not. In Figure 5.5 where we show how the STAs spread around the scenario, three links are marked, one as good, called Scenario A (AP_1 connected with STA_1 and AP_2 connected with STA_8), one sub-average, called Scenario B (AP_1 connected with STA_1 and AP_2 connected with STA_3) and one as bad, called Scenario C (AP_1 connected with STA_6 and AP_2 connected with STA_7). In Figure 5.7 we can see three examples of link combinations where we can compare the aggregated throughput and the gain of each of the mechanisms with respect to the default IEEE 802.11ad beamtraining mechanism for TCP traffic. For the good combination, named A, all the algorithms have gains greater than 20%, with the maximum gain of 53% for the ‘Power Threshold’ algorithm for omni-directional Rx-beam pattern. This algorithm is followed by the ‘Argmaxmin’ with omni-directional Rx-beam pattern, achieving a gain of 46%. The other two algorithms, ‘Weighted SIR Fairness’ and ‘Interference Threshold’ achieve gains of 21% and 24%. With directional Rx-beam pattern, gains increase to around 40% for all the mechanisms, with ‘Weighted SIR Fairness’ performing best and ‘Argmaxmin’ performing worst. With this network configuration, it is easy to obtain gains as the STAs are far apart from each other, facilitating the selection of beam patterns that do not interfere.




Figure 5.6: CDF of the throughput gains for bidirectional TCP and UDP traffic, when using directional Rx-beam pattern and when using the default omni-directional Rx-beam pattern.

Scenario B shows a below-average combination, resulting from the combination of AP_1 connected with STA_1 and AP_2 connected with STA_3 , achieving gains of 13% and 10% for the Power and Interference threshold techniques respectively, for directional Rx-beam pattern. For the omni-directional case we have gains of 22% for the ‘Power Threshold’ mechanism and 21% for ‘Argmaxmin’.

On the right hand side of Figure 5.7, we can see a connection example with low gains, named C. Here, ‘Weighted SIR Fairness’ has losses of 21% when for directional Rx-beam pattern, while it achieves 3% gains for omni-directional Rx-beam pattern. ‘Argmaxmin’ achieves 24% gains with respect the default IEEE 802.11ad mechanism for directional Rx-beam pattern and 12% for omni-directional. ‘Interference Threshold’ achieves gains of 13% and losses of 11% for directional and omni-directional Rx-beam pattern, respectively. ‘Power Threshold’ has gains of 14% for both cases, making it once again the most stable algorithm. For this network configuration it is difficult to obtain gains as the STAs are close to each other, and most of the beam patterns generate too much interference to obtain spatial reuse.

With these examples we can see that when the STAs are very close to each other the gains that




Figure 5.7: Aggregated throughput for three measured positions with two parallel links, showing the percentage gain with respect to the default aggregated throughput.

can be achieved are very low since no matter which beam pattern is chosen, interference is high. In contrast, links that are far apart can achieve very high gains, as it is easier to find patterns that minimize interference to neighboring nodes.

5.4.3. Office Environment: 3 Links

We now set up a third mmWave link in our scenario. Figure 5.8 shows how the mechanisms behave for three parallel TCP links for three different link configurations for bidirectional traffic. For these cases, AP_0 connects with STA_0 , AP_1 connects with STA_8 , and AP_1 connects with STA_1 , STA_2 or STA_4 as Scenarios A, B and C, respectively.

Again, similarly to the results shown in Section 5.4.2 for 2 parallel links, we can see that the ‘Power Threshold’ algorithm outperforms the others for most of the cases. For scenario A, where AP_1 is connected to STA_1 , we have gains of 76% for all the mechanisms for directional Rx-beam pattern except for ‘Weighted SIR Fairness’, which only has 20% gain. For omni-directional Rx-beam pattern, the gains decrease for the three best mechanisms, having now gains of 38%, 46% and 31% respectively, whereas the gain obtained by the ‘Weighted SIR Fairness’ method increases to 31%.

For scenario B, where AP_1 is connected to STA_2 , the ‘Argmaxmin’ and ‘Power Threshold’ algorithms again outperform the others. In this case, ‘Argmaxmin’ gives the highest throughput for directional Rx-beam pattern, having a gain of 138% while ‘Power Threshold’ has a gain of 110%. For omni-directional Rx-beam pattern, the ‘Power Threshold’ algorithm gives the highest gain of 128%.

In the last scenario, C, where AP_1 is connected to STA_4 , we can see that the highest gain is again for the ‘Power Threshold’ mechanism for directional Rx-beam pattern, resulting into a gain




Figure 5.8: Aggregated throughput for three measured positions with three parallel links, also showing the percentage gain with respect to the default aggregated throughput.

of 71%. For omni-directional Rx-beam pattern, we achieve a maximum gain of 56% with the ‘Argmaxmin’ mechanism.

As for the case of two links, the algorithm that consistently delivers good performance is the ‘Power Threshold’ algorithm, whereas other algorithms might have higher gains in some specific configurations but underperform in others.

Our three link measurements shows that the aggregated throughput decreases when more than two links work in parallel in the same collision domain as more nodes have to share the medium whenever spatial reuse is not possible.

5.5. Discussion

During the development and validation of these mechanisms we found different insights regarding the possibilities of spatial sharing for 60 GHz COTS devices.

IEEE 802.11ad beam pattern selection. We have shown how beam pattern selection affects parallel links in terms of spatial reuse. IEEE 802.11ad devices base their beamtraining mechanism on selecting the beam pattern that provides the highest SNR. However, this beam pattern does not need to provide the most efficient network behaviour when sharing the wireless medium with other networks as it might cause large interferences into neighbouring nodes, while other beam patterns could achieve larger aggregated throughput.

Weighted SIR Fairness leads to the highest Signal-to-Interference ratio. This mechanism results into the optimum channel capacity for a wireless communication system given by the Shannon-Hartley capacity theorem. But reality is different for our system, as the IEEE 802.11ad standard performs carrier sensing. This implies that even if we could achieve very high Signal-to-Interference ratios resulting in very high MCS, if a node receives interference above the carrier

sense threshold it needs to wait until the channel is idle before it can access the medium and transmit. This beamtraining mechanism can lead to rate losses, because the beam pattern that the algorithm chooses might still interfere too much for spatial reuse while at the same time reducing the link's SNR. This occurs for half of the cases, which have sub-25% bitrate losses. However, gains are achieved the rest of the time when using this mechanism.

Argmaxmin chooses the beam pattern combination in a way that the SIR of worst link in the whole scenario is maximized. This selection can lead us to an average poor scenario with poor connections among the links. Due to this, the different nodes might use worse MCS values and may still carrier sense with other nodes. In contrast, the default IEEE 802.11ad algorithm will carrier sense, but it will use stronger link budgets resulting in higher MCS and bitrates. Nevertheless, the gains of this mechanism are usually positive.

Power Threshold is the most reliable mechanisms from the ones developed. Here, we select a given group of beam pattern that provide a sufficiently high SNR and thus MCS, and from those select the one that produces the least interference. This way, the MCS is often equal to or not much below the one selected in the default beamtraining mechanism, but with the advantage that the devices will carrier sense less often.

Interference Threshold can fail due to the low directional beam patterns. While this mechanism should provide very efficient communications, the low-directionality of the current beam pattern codebook makes very difficult to find a beam pattern capable of providing good communication conditions while having zero-interference with neighboring nodes. Also, as we developed our beamtraining mechanisms to work for unmodified clients, it is impossible to achieve 100% spatial reutilization as only one link direction is optimized.

Other important insights that we discovered during this study are the following. For our 2-link measurement scenarios, we have an upper-bound of 3000 Mbps of aggregated throughput, which, given that the Talon AD7200 routers can achieve individual rates of up-to 2000 Mbps, means that we never achieve full spatial reuse. As a consequence, there is still a large performance gap to improve spatial sharing for mmWave devices, even though we have shown gains larger than 100%. Related with this previous observation, we have seen that the aggregated throughput for our three parallel link scenario is lower than the one we have with two parallel links. Despite some spatial reuse, the increased overhead due to the carrier sensing of a larger number of nodes and the corresponding medium access delays, decrease the overall efficiency of medium use. Nevertheless, our algorithms obtain very large gains for this 3-link scenario.

Chapter 6

Outlook

Overall, this thesis shows the behavior of current mmWave COTS devices together with two simple techniques to improve their main limitations which are their short link range and their restricted spatial reuse. This first-generation mmWave COTS analysis is highly relevant as next-generation mmWave devices will be based on them, so it is decisive to know them well in order to exploit their advantages as well as improve their weak points.

An important step forward in the 60 GHz mmWave industry is going to be the introduction of new WLAN mmWave devices under the IEEE 802.11ay amendment into the market and the deployment of mmWave wireless backbone links (as AT&T has proposed with the AirGig Project [55]). IEEE 802.11ay is the evolution of the current IEEE 802.11ad. This new mmWave WLAN amendment will support hybrid beamforming, channel aggregation and MU-MIMO, resulting into up to four independent streams with up to 8.64 GHz bandwidth which are capable of achieving 44 Gbps, resulting in a maximum of 176 Gbps when taking all the streams together [82]. AT&T's AirGig combines LTE technology with mmWave communications in order to provide internet in rural areas achieving hundreds of megabit per second connections with the use of 'self-install' equipment.

The aforementioned soon-to-be-deployed technologies encourage the study that we perform in this thesis, as these new devices should take into account all the previous-generation COTS design problems and try to solve them. Implementing the techniques exposed in this thesis to enhance the wireless communication range and spatial reuse, would also help in next-generation mmWave wireless deployments as we will later explain. The other characteristic that these devices should improve is the user mobility model, as the movement of a user may damage the link causing large throughput losses which will get worse when these mmWave chipsets are integrated into Virtual-Reality headsets or hand-held devices such as smartphones.

With the arrival of the IEEE 802.11ay amendment, devices will be able to aggregate up to 4 channels creating a bandwidth of 8.64 GHz. With such a large bandwidth, there is a large likelihood that the channel is going to be frequency selective. This frequency selectivity is going to be larger than in IEEE 802.11ad devices as the bandwidth can increase up to four times. These

large bandwidths provide the opportunity of using the techniques seen in Chapter 4 to extend the link distances or achieving energy savings. We also expect very dense deployment scenarios with the arrival of this new amendment, as these chipsets will be used in smartphones, tablets, and laptops. These dense scenarios are going to cause large interference among them if they do not improve their beamforming or their beamtraining capabilities. This is why we encourage these devices to adapt to the beamtraining techniques exposed in Chapter 5, not only taking into account the received power but also taking into account the interference that may be created.

mmWave backhaul links such as AirGig [55] are going to need very precise antennas. As these links are going to be very stable over time, they need to have very fine beamforming capabilities in order to avoid interferences with other devices. This very fine beamforming is far from the beamforming capabilities that we have seen in our COTS devices. However, if these next generation devices do not allow the synthesis of pencil-shape beam patterns, they will need to improve their beam pattern selection in order to reduce interferences with neighboring nodes. As these are static links, they need the best coverage with minimum interference as possible as they will be deployed for a long time without needing to retrain the link often. This is why they should implement techniques such as the ones seen in Chapter 5 to have the highest MCS possible while decreasing the interference avoiding medium sharing. The only time when these devices will need to train is when they are moved from one location to another or when environmental conditions displace them (as could happen with very heavy winds). If one of these AirGig devices fail, we have proved that an aerial mmWave link could fulfill its duty.

Continuing our work, one should analyze the co-channel interference that might be created when two links working in the same area are using adjacent channels. This needs to be studied as this kind of interference is very probable to exist in such large bandwidth devices and in such case, every device in dense deployments might share the wireless medium decreasing the throughput dramatically. If this co-channel interference happens, mmWave devices should equip better RF-filters to avoid transmitting power out of their desired bandwidths. If devices do not improve their beam pattern synthesis or their beam pattern selection this channel sharing will be even enlarged.

If we compare this thesis with the ideal performance of the mmWave technology, the most important part that vendors should improve are their antenna arrays. Having a better antenna array would result in better beamforming capabilities that would result in less problematic and higher efficiency networks. With an ideal antenna array capable of creating ultra-directional beam patterns we would not need to look for techniques to extend the link range or save energy consumption as these would manage their energy focusing all the transmission power into the receiving node. At the same time, we would not need to look for mechanisms to avoid interference, as these high directional patterns would make non-interference communications as they would not have sidelobes. In order to create better antenna arrays, manufacturers should include more antenna elements as well as higher-definition phase shifters. Real-time beamforming devices would also implement better beam patterns as devices relying in beam pattern codebooks will always be more limited in terms of beam pattern shapes.

Chapter 7

Conclusions

In this thesis, we perform the first mmWave COTS devices characterization, looking for the real characteristics of these devices and once identified their two main weaknesses, we target them to improve their global performance. This thesis is divided into seven chapters. The first one consists of the introduction, where we do a general view on what this thesis targets, how it is motivated, and the main contributions. Next we can find background and related work. Then, we continue with our analysis of mmWave COTS devices, showing the low range and the low spatial reuse that they achieve given their limited manufacturing cost. After this analysis, we show how to increase these devices link distances. In the fifth chapter, we show how to increase the low-spatial reuse thanks to different beamtraining mechanisms. The sixth chapter consists of a general outlook of this work together with the future work and finally, we arrive at this conclusion chapter.

Going into details, Chapter 3 presents an in-depth analysis of consumer-grade off-the-shelf 60 GHz systems [1–3, 5, 6]. In order to perform this analysis, we overhear the frame level communication with a mmWave downconverter, as well as extract the low layer parameters from a hidden monitor in these COTS devices [4]. Our goal is to investigate the impact of the cost-effective designs of such devices. While these effects are often qualitatively well-known, we quantify them in order to provide a practical intuition on how important they are. This contributes crucial insights for the design of 60 GHz networking protocols. In particular, we investigate the impact of data aggregation, reflections, beam patterns, beamtraining, interference, and mobility. We find that data aggregation yields throughput gains of up to $5.4\times$. Further, beamforming often results in strong side lobes due to the limitations of consumer-grade antennas. Regarding reflections, we show that not only first-order but also second-order reflections occur. As a result, we observe significant interference in scenarios with two or more 60 GHz links that in principle should not interfere with each other. In mobile scenarios, throughput degrades on average about 30% for link loads below one gigabit-per-second compared to the static case. Beyond beam pattern misalignments, the underlying reasons include MCS fluctuations, higher packet error rates, and 30% higher control traffic overhead. We also show how drone communications can rely into mmWave links. We carried out measurements to determine the maximum range that IEEE

802.11ad COTS mmWave aerial links can achieve as well as we emulate strong wind conditions. We observed that these radios can be pushed to a vertical distance of 31 meters and to a 3D distance of 24 meters when the mobile mmWave link is moving on the horizontal plane at a height of 12 meters. Strong winds may jeopardize specific use cases such as real-time high resolution video streaming due to high link variability, but it would still be useful for less latency sensitive applications such as providing Internet access or for file sharing as the throughput only decreases by up to 40%. From this chapter, we can conclude that, while some common 60 GHz assumptions hold (e.g., data aggregation), others become critical for consumer-grade devices (e.g., impact of side lobes).

Chapter 4 shows how frequency selectivity can help mitigate the strong range limitation from 60 GHz indoors networks [7]. Instead of using techniques such as bit loading, power loading, and subcarrier switch-off to improve throughput, we leverage them to reduce the impact of poor subcarriers at the edge of the coverage of a 60 GHz AP. To this end, we exploit the theoretically unnoticeable frequency selectivity in indoor 60 GHz channels. Although 60 GHz networks use highly directive antennas which greatly reduce the impact of multi-path components, we show practically that typical indoor environments do feature a significant number of such reflections. This results in an unexpectedly strong frequency selective behavior which becomes particularly clear for link lengths beyond 2.5 meters. Our practical testbed results show that using techniques that exploit this frequency selectivity can increase range by up to 50% in many of the scenarios that we consider. Moreover, this effect allows for transmit power savings of up to 7 dB.

In Chapter 5 we propose novel beamtraining mechanisms in order to extend spatial reuse on mmWave COTS dense deployments [8]. IEEE 802.11ad devices choose the strongest communication beam pattern from a predefined codebook without taking into consideration other links that can be in the same collision domain. This, together with the non-ideal shapes of their beam patterns, prevents efficient spatial sharing among the nodes, reducing aggregated and individual throughput. In order to experiment how two links in the same collision domain can enhance their behavior only due to use of different beam patterns, we setup two parallel links at different distances and cycle among all the beam patterns from their codebooks, showing that the aggregated throughput can go up to two times if selecting preferable beam patterns than the default ones [5]. Seeing this room for improvement, we implement four different beamtraining schemes to select more efficient beam patterns, not only in terms of received power but also in terms of interference. This system is implemented on the AP side and is transparent for the STAs, making possible for any unmodified STA joining the wireless network to benefit from the system. It provides aggregated throughput gains of up to 100% for some cases and average gains of 25% for two links with TCP and bidirectional transmissions. We implement and test our four beamtraining mechanisms in mmWave IEEE 802.11ad COTS devices. We first validate the implementation and mechanisms in a simple scenario and then test them in a real world open-area office environment. In this second scenario, we test our algorithms for 2 and 3 parallel links.

Wrapping up we can conclude that this work has provided a new perspective on how first-generation mmWave COTS devices behave. We show the main advantages and disadvantages of these devices, as well as how to solve some of these disadvantages. We show that their limited link range can be improved with the use of frequency selective techniques. We also analyze that new beamtraining mechanisms can be implemented in these systems in order to take into account interference from other nodes, resulting in better spatial reuse. A number of open issues that need to be clarified remain open, as the co-channel interference or how mmWave networks would behave in very dense scenarios. Development and implementation of better antenna arrays should also be taken into account in future studies.

References

- [1] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, “Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation,” in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '15. New York, NY, USA: ACM, 2015, pp. 17:1–17:13. [Online]. Available: <http://doi.acm.org/10.1145/2716281.2836102>
- [2] G. Bielsa, “Analysis of off-the-shelf Millimeter Wave Systems with Phased Antenna Arrays,” 2016.
- [3] 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*, pp. 1–1, 2018.
- [4] A. Loch, G. Bielsa, and J. Widmer, “Practical Lower Layer 60 GHz Measurements Using Commercial Off-The-Shelf Hardware,” in *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*. ACM, 2016, pp. 9–16.
- [5] H. Assasa, G. Bielsa, S. Kumar Saha, P. Jiménez Mateo, A. Loch, D. Koutsonikolas, and J. Widmer, “Performance analysis of medium access control and spatial reuse for ieee 802.11ad deployments,” in *Elsevier Pervasive and Mobile Computing*.
- [6] G. Bielsa, M. Mezzavilla, J. Widmer, and S. Rangan, “Performance assessment of off-the-shelf mmwave radios for drone communications,” in *2019 IEEE 20th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2019.
- [7] G. Bielsa, A. Loch, and J. Widmer, “60 GHz range boost: Exploiting frequency selectivity in millimeter-wave networks,” in *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2017, pp. 1–9.
- [8] G. Bielsa, A. Loch, and J. Widmer, “Optimizing mmwave spatial reuse: Signal-to-interference aware beamtraining,” in *Internet of Things: Smart Objects and Services*, 2019.

- [9] A. Olivier, G. Bielsa, I. Tejado, M. Zorzi, J. Widmer, and P. Casari, "Lightweight Indoor Localization for 60-GHz Millimeter Wave Systems," in *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, June 2016, pp. 1–9.
- [10] M. Scalabrin, M. Rossi, G. Bielsa, A. Loch, and J. Widmer, "Millimetric diagnosis: Machine learning based network analysis for mm-wave communication," in *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2017, pp. 1–9.
- [11] A. Loch, G. Bielsa, R. Santos, and J. Widmer, "mm-View: Obtaining real-time lower layer information of commercial off-the-shelf 60 GHz hardware," in *A World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2017 IEEE 18th International Symposium on*. IEEE, 2017, pp. 1–3.
- [12] G. Bielsa, J. Palacios, A. Loch, D. Steinmetzer, P. Casari, and J. Widmer, "Indoor Localization Using Commercial Off-The-Shelf 60 GHz Access Points," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 2384–2392.
- [13] J. Palacios, G. Bielsa, P. Casari, and J. Widmer, "Communication-driven localization and mapping for millimeter wave networks," pp. 2402–2410, April 2018.
- [14] J. Palacios and G. Bielsa and P. Casari and J. Widmer, "Single- and multiple-access point indoor localization for millimeter-wave networks," *IEEE Transactions on Wireless Communications*, vol. 18, no. 3, pp. 1927–1942, March 2019.
- [15] M. Scalabrin, G. Bielsa, A. Loch, M. Rossi, and J. Widmer, "Machine learning based network analysis using millimeter-wave narrow-band energy traces," *IEEE Transactions on Mobile Computing*, pp. 1–1, 2019.
- [16] M. Hung, "Leading the iot, gartner insights on how to lead in a connected world," *Gartner Research*, pp. 1–29, 2017.
- [17] X. Ge, S. Tu, G. Mao, C.-X. Wang, and T. Han, "5G Ultra-Dense Cellular Networks," *IEEE Wireless Communications*, vol. 23, no. 1, pp. 72–79, 2016.
- [18] E. Ziouva and T. Antonakopoulos, "CSMA/CA performance under high traffic conditions: throughput and delay analysis," *Computer communications*, vol. 25, no. 3, pp. 313–321, 2002.
- [19] C. Hansen, "WiGiG: Multi-gigabit wireless communications in the 60 GHz band," *IEEE Wireless Communications*, vol. 18, no. 6, 2011.

- [20] IEEE, “Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band,” *IEEE Std 802.11ad-2012*, 2012.
- [21] P. Schmulders, “Exploiting the 60 GHz Band for Local Wireless Multimedia Access: Prospects and Future Directions,” *IEEE Communications Magazine*, vol. 40, no. 1.
- [22] T. Nitsche, C. Cordeiro, A. Flores, E. Knightly, E. Perahia, and J. Widmer, “IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi,” *IEEE Communications Magazine*, vol. 52, no. 12, 2014.
- [23] Y. M. Tsang, A. S. Y. Poon, and S. Addepalli, “Coding the Beams: Improving Beamforming Training in mmWave Communication System,” in *Proc. of the 2011 IEEE Global Telecommunications Conference (Globecom)*, 2011.
- [24] G. Bielsa, A. Loch, I. Tejado, T. Nitsche, and J. Widmer. (2017) IEEE 802.11ad Trace Database. [Online]. Available: <http://wireless.networks.imdea.org/60-ghz-networking-trace-dataset>
- [25] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, “Demystifying 60GHz Outdoor Picocells,” in *Proc. of ACM Mobicom’14*, 2014.
- [26] Y. Katayama, K. Takano, Y. Kohda, N. Ohba, and D. Nakano, “Wireless data center networking with steered-beam mmwave links,” in *2011 IEEE Wireless Communications and Networking Conference*, March 2011, pp. 2179–2184.
- [27] H. Vardhan, N. Thomas, S. Ryu, B. Banerjee, and R. Prakash, “Wireless data center with millimeter wave network,” in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, Dec 2010, pp. 1–6.
- [28] X. Tie, K. Ramachandran, and R. Mahindra, “On 60 GHz Wireless Link Performance in Indoor Environments,” in *Proc. of PAM’12*, 2012.
- [29] H. Liebe, P. Rosenkranz, and G. Hufford, “Atmospheric 60-GHz oxygen spectrum: New laboratory measurements and line parameters,” *Journal of quantitative spectroscopy and radiative transfer*, vol. 48, no. 5-6, pp. 629–643, 1992.
- [30] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, “BeamSpy: Enabling Robust 60 GHz Links Under Blockage,” in *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, 2016, pp. 193–206.
- [31] H. Xu, V. Kukshya, and T. S. Rappaport, “Spatial and temporal characteristics of 60-GHz indoor channels,” *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 3.

- [32] T. S. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang, "Overview of Millimeter Wave Communications for Fifth-Generation (5G) Wireless Networks—With a Focus on Propagation Models," vol. 65, no. 12, Dec 2017, pp. 6213–6230.
- [33] N. Moraitis and P. Constantinou, "Indoor channel measurements and characterization at 60 GHz for wireless local area network applications," *IEEE Transactions on Antennas and Propagation*, vol. 52, no. 12, 2004.
- [34] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling," in *Proc. of ACM SIGMETRICS'15*, 2015.
- [35] T. Manabe, Y. Miura, and T. Ihara, "Effects of antenna directivity and polarization on indoor multipath propagation characteristics at 60 GHz," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 3, 1996.
- [36] T. Zwick, T. Beukema, and H. Nam, "Wideband channel sounder with measurements and model for the 60 GHz indoor radio channel," *IEEE Transactions on Vehicular Technology*, vol. 54, no. 4, 2005.
- [37] T. S. Rappaport, E. Ben-Dor, J. N. Murdock, and Y. Qiao, "38 GHz and 60 GHz angle-dependent propagation for cellular peer-to-peer wireless communications," in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 4568–4573.
- [38] A. Lamminen, J. Saily, and A. Vimpari, "60-GHz Patch Antennas and Arrays on LTCC With Embedded-Cavity Substrates," *IEEE Transactions on Antennas and Propagation*, vol. 56, no. 9, 2008.
- [39] X.-P. Chen, K. Wu, L. Han, and F. He, "Low-Cost High Gain Planar Antenna Array for 60-GHz Band Applications," *IEEE Transactions on Antennas and Propagation*, vol. 58, no. 6, 2010.
- [40] S. B. Yeap, Z. N. Chen, and X. Qing, "Gain-Enhanced 60-GHz LTCC Antenna Array With Open Air Cavities," *IEEE Transactions on Antennas and Propagation*, vol. 59, no. 9, 2011.
- [41] S. Singh, R. Mudumbai, and U. Madhow, "Interference Analysis for Highly Directional 60-GHz Mesh Networks: The Case for Rethinking Medium Access Control," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, 2011.
- [42] H. Yang, P. Smulders, and M. Herben, "Frequency Selectivity of 60-GHz LOS and NLOS Indoor Radio Channels," in *Proc. of VTC Spring'06*, vol. 6, 2006.
- [43] M. Rebato, L. Resteghini, C. Mazzucco, and M. Zorzi, "Study of Realistic Antenna Patterns in 5G mmWave Cellular Scenarios," *arXiv preprint arXiv:1802.01316*, 2018.

- [44] J. Palacios, D. D. Donno, D. Giustiniano, and J. Widmer, "Speeding up mmwave beam training through low-complexity hybrid transceivers," in *2016 IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, Sept 2016, pp. 1–7.
- [45] Wu, Sau-Hsuan and Lin, Ko-Yen and Chiu, Lin-Kai, "Hybrid beamforming using convex optimization for sdma in millimeter wave radio," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on.* IEEE, 2009, pp. 823–827.
- [46] J. Zhao and D. Liu, "An interference alignment scheme for 60 GHz millimeter-wave communication system," in *Vehicular Technology Conference (VTC Fall), 2012 IEEE.* IEEE, 2012, pp. 1–5.
- [47] H. Shokri-Ghadikolaei, L. Gkatzikis, and C. Fischione, "Beam-searching and transmission scheduling in millimeter wave communications," in *2015 IEEE International Conference on Communications (ICC)*, June 2015, pp. 1292–1297.
- [48] M. K. Haider and E. W. Knightly, "Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation," in *Proc. of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2016, pp. 61–70.
- [49] A. Loch, H. Assasa, J. Palacios, J. Widmer, H. Suys, and B. Debaillie, "Zero Overhead Device Tracking in 60 GHz Wireless Networks using Multi-Lobe Beam Patterns," in *Proc. of the 13th ACM Conference on Emerging Networking Experiments and Technologies*, 2017.
- [50] A. S. Cacciapuoti, "Mobility-Aware User Association for 5G mmWave Networks," *IEEE Access*, vol. 5, pp. 21 497–21 507, 2017.
- [51] J. Palacios, D. De Donno, and J. Widmer, "Tracking mm-wave channel dynamics: Fast beam training strategies under mobility," in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, May 2017, pp. 1–9.
- [52] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan, "OpenMili: A 60 GHz Software Radio Platform with a Reconfigurable Phased-array Antenna," in *Proc. of the 22Nd Annual International Conference on Mobile Computing and Networking*, 2016, pp. 162–175.
- [53] L. Simić, J. Arnold, M. Petrova, and P. Mähönen, "RadMAC: Radar-enabled Link Obstruction Avoidance for Agile Mm-wave Beamsteering," in *Proc. 3rd Workshop on Hot Topics in Wireless*, 2016, pp. 61–65.
- [54] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling," in *Proc. of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2015, pp. 71–84.

- [55] AT&T. Project airgig gets closer to initial commercial deployment. [Online]. Available: <https://about.att.com/story/2018/airgig.html>
- [56] M. Mezzavilla, M. Polese, A. Zanella, A. Dhananjay, S. Rangan, C. Kessler, T. S. Rappaport, and M. Zorzi, "Public safety communications above 6 GHz: Challenges and opportunities," *IEEE Access*, vol. 6, pp. 316–329, 2018.
- [57] I. Guvenc, W. Saad, M. Bennis, C. Wietfeld, M. Ding, and L. Pike, "Wireless communications, networking, and positioning with unmanned aerial vehicles [guest editorial]," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 24–25, 2016.
- [58] Qualcomm Atheros. (2016). [Online]. Available: <https://wireless.wiki.kernel.org/en/users/drivers/wil6210>
- [59] VubIQ. V60WGD03 60 GHz Waveguide Development System. [Online]. Available: <http://www.pasternack.com/60-ghz-development-systems-category.aspx>
- [60] iPerf - The TCP, UDP and SCTP network bandwidth measurement tool. <https://iperf.fr/>. [Online]. Available: <https://iperf.fr/>
- [61] S. Byeon, K. Yoon, O. Lee, S. Choi, W. Cho, and S. Oh, "Mofa: Mobility-Aware Frame Aggregation in Wi-Fi," in *Proc. of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, 2014, pp. 41–52.
- [62] K. Hosoya, N. Prasad, K. Ramachandran, N. Orihashi, S. Kishimoto, S. Rangarajan, and K. Maruhashi, "Multiple Sector ID Capture (MIDC): A Novel Beamforming Technique for 60-GHz Band Multi-Gbps WLAN/PAN Systems," *IEEE Transactions on Antennas and Propagation*, vol. 63, no. 1, 2015.
- [63] E. Arribas, V. Mancuso, and V. Cholvi, "Fair Cellular Throughput Optimization with the Aid of Coordinated Drones," in *IEEE INFOCOM 2019 WKSHPs - MiSARN*, 2019, pp. 1–6.
- [64] Z. Xiao, P. Xia, and X.-G. Xia, "Enabling UAV cellular with millimeter-wave communication: Potentials and approaches," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 66–73, 2016.
- [65] W. Khawaja, O. Ozdemir, and I. Guvenc, "UAV air-to-ground channel characterization for mmWave systems," in *Proc. IEEE VTC*, 2017.
- [66] W. Khawaja, O. Ozdemir, and I. Guvenc, "Temporal and spatial characteristics of mm wave propagation channels for uavs," pp. 1–6, May 2018.
- [67] "Facebook's terragraph project:," <https://terragraph.com/>.

- [68] G. H. Sim, A. Asadi, A. Loch, M. Hollick, and J. Widmer, "Opp-relay: Managing directionality and mobility issues of millimeter-wave via d2d communication," *IEEE COMSNETS*, 2017.
- [69] S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell, "Millimeter Wave WPAN: Cross-Layer Modeling and Multi-Hop Architecture," in *Proceedings of the IEEE INFOCOM 2007-26th IEEE International Conference on Computer Communications*. IEEE Computer Society, 2007, pp. 2336–2340.
- [70] J. Qiao, L. X. Cai, and X. Shen, "Multi-hop concurrent transmission in millimeter wave wpans with directional antenna," in *2010 IEEE International Conference on Communications*.
- [71] J. Gross, M. Emmelmann, O. Puñal, and A. Wolisz, "Dynamic Single-user OFDM Adaptation for IEEE 802.11 Systems," in *Proc. of ACM MSWiM*, 2007.
- [72] O. Puñal, H. Escudero, and J. Gross, "Power loading: Candidate for future WLANs?" in *Proc. of WoWMoM'12*, 2012.
- [73] T. Nitsche and J. Widmer, "Sub-carrier Switch Off in OFDM-based wireless local area networks," in *Proc. of SECON'13*, 2013.
- [74] D. P. Palomar and J. R. Fonollosa, "Practical algorithms for a family of waterfilling solutions," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, 2005.
- [75] E. N. Onggosanusi, A. M. Sayeed, and B. D. V. Veen, "Efficient signaling schemes for wideband space-time wireless channels using channel state information," *IEEE Transactions on Vehicular Technology*, vol. 52, no. 1, 2003.
- [76] F. Rey, M. Lamarca, and G. Vazquez, "Transmit filter optimization based on partial CSI knowledge for wireless applications," in *Proc. of ICC'03*, 2003.
- [77] Z. Song, K. Zhang, and Y. L. Guan, "Joint bit-loading and power-allocation for OFDM systems based on statistical frequency-domain fading model," in *Proc. of VTC Fall'02*, 2002.
- [78] E. Bedeer, O. A. Dobre, M. H. Ahmed, and K. E. Baddour, "A Systematic Approach to Jointly Optimize Rate and Power Consumption for OFDM Systems," *IEEE Transactions on Mobile Computing*, vol. 15, no. 6, 2016.
- [79] G. Mangraviti, K. Khalaf, Q. Shi, K. Vaesen, D. Guermandi, V. Giannini, S. Brebels, F. Frazzica, A. Bourdoux, C. Soens, W. V. Thillo, and P. Wambacq, "13.5 a 4-antenna-path beamforming transceiver for 60ghz multi-gb/s communication in 28nm cmos," in *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, 2016.

-
- [80] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in o-the-shelf IEEE 802.11 ad devices," in *Proc. ACM CoNEXT*, 2017.
- [81] D. Steinmetzer, D. Wegemer, and M. Hollick. Talon tools: The framework for practical iee 802.11ad research. [Online]. Available: <https://seemoo.de/talon-tools/>
- [82] Y. Ghasempour, C. R. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 186–192, 2017.