Network management and control for mmWave communications

by

Pablo Jiménez Mateo

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

Telematic Engineering

Universidad Carlos III de Madrid

Advisor: Joerg Widmer

December, 2021

Network management and control for mmWave communications

Prepared by:
Pablo Jiménez Mateo, IMDEA Networks Institute, Universidad Carlos III de Madrid
contact: pablo.jimenezmateo@imdea.org

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

This work has been supported by:

To Eva, my wife. Without her undeniable support this would not be a reality.

To my parents, Juanjo and Julia, who sacrificed themselves to give me the education I have.

To my brother, Jesús, who always encouraged me to take the difficult path.

To my grandparents, Pepita and Jesús, who told me to study (but never when to stop).

And to me! I deserve it, damn it!

Thank you!

Acknowledgements

A Ph.D. is the culmination of many years of study and effort. It has been undoubtedly one of the most challenging endeavours of my life and, even though a Ph.D is ultimately a single person objective, I was not alone at any point. I have to thank to all those now forgotten teachers from school and high school, and all the friends along the way. More importantly my family, specially my parents and my brother, have been a constant rock all this time, I deeply appreciate that. More recently my family grew, I got married to Eva my awesome wife, to whom I am specially grateful for supporting me through the last years of my Ph.D. I could not have made it without her.

I want to extend my sincere thanks to Joerg Widmer, a supervisor is probably the most important figure during a Ph.D. and he was more than up to the task. He guided me through difficult times, with patience and good advice. I also want to extend my gratitude to everyone from IMDEA Networks Institute. People feel so close to you when you are going through challenges that are similar, even if they are not the same. I have enjoyed meeting and getting to know so much people over the years, over coffee breaks and in Manzanares, I thank you all: Dr. Hany Assasa, Alejandro Blanco, Dr. Ander Galisteo, Dr. Joan Palacios, Dr. Guillermo Bielsa, Dr. Jesús Omar, Dolores García, Francisco Hervas, Dr. Claudio Fiandrino, Rafael Ruiz, Nina Grosheva, Norbert Ludant, Noelia Perez, Dr. Gines García, Dr. Narseo Vallina, Dr. Boja Genovés, Dr. Maurizio Rea, Giulia Attanasio, Constantine Ayimba, Álvaro Feal, Julien Gamba, Yago Lizarribar, Olowasegun Ojo, Pelayo Vallina, Dr. Edgar Arribas, Dr. Roberto Calvo, Neftalí González, Ricardo Padrino and Dr. Adrian Loch.

Published Content

This thesis is based on the following published papers:

[1] Pablo Jiménez Mateo, Claudio Fiandrino, Joerg Widmer. Analysis of TCP performance in 5G mm-Wave mobile networks. Published in *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 20-24 May 2019, Shanghai, China. https://doi.org/10.1109/ICC.2019.8761718

- This work is fully included and its content is reported in Chapter 3.
- The author's role in this work is the scenario design, the implementation and the analysis of the different TCP protocols.

[2] Jesus Omar Lacruz, Dolores Garcia, **Pablo Jiménez Mateo**, Joan Palacios, Joerg Widmer. mm-FLEX: An Open Platform for Millimeter-Wave Mobile Full-Bandwidth Experimentation. Published in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, June 2020. https://doi.org/10.1145/3386901.3389034

- This work is partially included and its content is reported in Chapter 5.
- The author's role in this work is focused on the Commercial-Off-The-Shelf (COTS) overhead reduction analysis based on an experimental COTS dense deployment.

[3] Pablo Jiménez Mateo, Alejandro Blanco Pizarro, Norbert Ludant, Matteo Marugan Borelli, Amanda Garcia-Garcia, Adrian Loch, Zhenyu Shi, Yi Wang, Joerg Widmer. A comprehensive study of low frequency and high frequency channel correlation. Published in *2019 International Conference on Computing, Networking and Communications (ICNC)*, 18-21 Feb. 2019, Honolulu, HI, USA. https://doi.org/10.1109/ICCNC.2019.8685565

- This work is fully included and its content is reported in Chapter 4.
- The author's role in this work is the scenario design and modelling, as well as the analysis for multiple frequencies and their correlation.

Other papers that are not included on this thesis:

[**4**] Carlos Andrés Ramiro, Claudio Fiandrino, Alejandro Blanco Pizarro, **Pablo Jiménez Mateo**, Norbert Ludant, Joerg Widmer. OpenLEON: An End-to-End Emulator from the Edge Data Center to the Mobile User. Published in *Proceedings of the 12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*, November 2, 2018, New Delhi, India. http://dx.doi.org/10.1145/3267204.3267210

- The author's role in this work is the network configuration and routing, as well as environment setup on Linux.

[**5**] Alejandro Blanco Pizarro, Norbert Ludant, **Pablo Jiménez Mateo**, Zhenyu Shi, Yi Wang, Joerg Widmer. Performance Evaluation of Single Base Station ToA-AoA Localization in an LTE Testbed. Published in *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 8-11 September 2019, Istanbul, Turkey. https://doi.org/10.1109/PIMRC.2019.8904454

- The author's role in this work is focused on the analysis and configuration of the testbed.

[**6**] Claudio Fiandrino, Alejandro Blanco Pizarro, **Pablo Jiménez Mateo**, Carlos Andrés Ramiro, Norbert Ludant, Joerg Widmer. OpenLEON: An End-to-End Emulator from the Edge Data Center to the Mobile User. Published in *Computer Communications (COMPUT COMMUN)*, December, 2019. http://dx.doi.org/10.1016/j.comcom.2019.08.024

- The author's role in this work is the network configuration and routing, as well as environment setup on Linux.

[**7**] Dolores Garcia, Jesus Omar Lacruz, **Pablo Jiménez Mateo**, Joerg Widmer. POLAR: Passive Object localization with Ieee 802.11ad using phased antenna arrays. Published in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 6-9 July 2020, Toronto, ON, Canada. https://doi.org/10.1109/INFOCOM41043.2020.9155383

- The author's role in this work is focused on the testbed design as well as the experimental evaluation.

[**8**] Jesus Omar Lacruz, Dolores Garcia, **Pablo Jiménez Mateo**, Joan Palacios, Joerg Widmer. High-speed millimeter-wave mobile experimentation on software-defined radios. Published in *GetMobile: Mobile Computing and Communications*, December 2020. https://doi.org/10.1145/3457356.3457368

- The author's role in this work is focused on the testbed design as well as the experimental evaluation.

[**9**] Maurizio Rea, Domenico Giustiniano, **Pablo Jiménez Mateo**, Yago Lizarribar, Joerg Widmer. Beam searching for mmWave networks with sub-6GHz WiFi and inertial sensors inputs: An experimental study. Published in *Computer Networks*, August 2021, Online. http://hdl.handle.net/20.500.12761/1503

- The author's role in this work is doing all the measurements with COTS Millimeter-wave (mmWave) devices as well as analyzing the results.

Abstract

Millimeter-wave (mmWave) is one of the key technologies that enables the next wireless generation. mmWave offers a much higher bandwidth than sub-6GHz communications which allows multi-gigabit-per-second rates. This also alleviates the scarcity of spectrum at lower frequencies, where most devices connect through sub-6GHz bands. However new techniques are necessary to overcome the challenges associated with such high frequencies. Most of these challenges come from the high spatial attenuation at the mmWave band, which requires new paradigms that differ from sub-6GHz communications. Most notably mmWave telecommunications are characterized by the need to be directional in order to extend the operational range. This is achieved by using electronically steerable antenna arrays, that focus the energy towards the desired direction by combining each antenna element constructively or destructively. Additionally, most of the energy comes from the Line Of Sight (LOS) component which gives mmWave a quasi-optical behaviour where signals can reflect off walls and still be used for communication. Some other challenges that directional communications bring are mobility tracking, blockages and misalignments due to device rotation. The IEEE 802.11ad amendment introduced wireless telecommunications in the unlicensed 60 GHz band. It is the first standard to address the limitations of mmWave. It does so by introducing new mechanisms at the Medium Access Control (MAC) and Physical (PHY) layers. It introduces multi-band operation, relay operation mode, hybrid channel access scheme, beam tracking and beam forming among others.

In this thesis we present a series of works that aim to improve mmWave telecommunications. First we give an overview of the intrinsic challenges of mmWave telecommunications, by explaining the modifications to the MAC and PHY layers. This sets the base for the rest of the thesis. Then do a comprehensive study on how mmWave behaves with existing technologies, namely TCP. TCP is unable to distinguish losses caused by congestion or by transmission errors caused by channel degradation. Since mmWave is affected by blockages more than sub-6GHz technologies, we propose a set of parameters that improve the channel quality even for mobile scenarios. The next job focuses on reducing the initial access overhead of mmWave by using sub-6GHz information to steer towards the desired direction. We start this work by doing a comprehensive High

Frequency (HF) and Low Frequency (LF) correlation, analyzing the similarity of the existing paths between the two selected frequencies. Then we propose a beam steering algorithm that reduces the overhead to one third of the original time. Once we have studied how to reduce the initial access overhead, we propose a mechanism to reduce the beam tracking overhead. For this we propose an open platform based on a Field Programmable Gate Arrays (FPGA) where we implement an algorithm that completely removes the need to train on the Station (STA) side. This is achieved by changing beam patterns on the STA side while the Access Point (AP) is sending the preamble. We can change up to 10 beam patterns without losing connection and we reduce the overhead by a factor of 8.8 with respect to the IEEE 802.11ad standard. Finally we present a dual band location system based on Commercial-Off-The-Shelf (COTS) devices. Locating the STA can improve the quality of the channel significantly, since the AP can predict and react to possible blockages. First we reverse engineer existing 60 GHz enabled COTS devices to extract Channel State Information (CSI) and Fine Timing Measurements (FTM) measurements, from which we can estimate angle and distance. Then we develop an algorithm that is able to choose between HF and LF in order to improve the overall accuracy of the system. We achieve less than 17 cm of median error in indoor environments, even when some areas are Non Line Of Sight (NLOS).

Table of Contents

List of Tables

List of Figures

List of Acronyms

A-BFT Association Beaforming Training

AGC Automatic Gain Control

AOA Angle of Arrival

AP Access Point

AS Angle Spread

ATI Announcement Transmission Interval

BF Beamforming

BFT Beam Forming Training

BHI Beacon Header Interval

BI Beacon Interval

BP Beam Patterns

BRP Beam Refinement Protocol

BT Beam Tracking

BTI Beacon Transmission Interval

CBAP Contention-based Access Period

CEF Channel Estimation Field

COTS Commercial-Off-The-Shelf

CRC Cyclic Redundancy Check

CSI Channel State Information

CSMA/CA Carrier-Sense Multiple Access with Collision Avoidance

DMG Directional Multi-Gigabit

DS Delay Spread

DTI Data Transmission Interval

FCC Federal Communications Commission

FPGA Field Programmable Gate Arrays

FTM Fine Timing Measurements

HF High Frequency

LF Low Frequency

LOS Line Of Sight

MAC Medium Access Control

MCS Modulation and Coding Scheme

MIMO Multiple-Input and Multiple-Output

mmWave Millimeter-wave

NLOS Non Line Of Sight

OFDM Orthogonal Frequency Division Multiplexing

OLOS Obstructed Line Of Sight

PHY Physical

RSS Received Signal Strength

RTT Round Trip Time

SLS Sector Level Sweep

SNR Signal-to-Noise Ratio

SP Service Period

SSW Sector Sweep

STA Station

STF Short Training Field

MultiLoc Multi-band Location System

ToF Time of Flight

TRN Training

TRN-R Receive Training

TRN-T Transmit Training

UFA Ultra Fast Alignment

URA Uniform Rectangular Array

WiGig Wireless Gigabit Alliance

WLAN Wireless Local Area Network

WMI Wireless Module Interface

“"If you like nerds, raise your hand. If you don’t, raise your standards"”

Violet Haberdasher

1 Introduction

Millimeter-wave (mmWave) is a key technology to reduce network latency while increasing the overall throughput. The increased traffic demand in mobile and wireless local area networks has boosted the interest in networking at mmWave frequencies, where the huge bandwidth allows for significantly higher data rates than those at sub-6 GHz frequencies. The IEEE 802.11ad amendment operates at 60 GHz with a bandwidth of 2.16 GHz, which enables data rates of up to 4.6 Gbps, while its successor, the IEEE 802.11ay amendment, can reach up to 100 Gbps through Multiple-Input and Multiple-Output (MIMO), higher order modulation schemes and channel bonding. This enables previous wired-only technologies to become wireless, such as streaming of ultra high definition video, virtual reality and wireless front and backhauling for mobile networks. However there are challenges of operating at these bands such as atmospheric and blockage attenuation, which makes omnidirectional communications unfeasible, unless very extreme conditions are meet [10]. In order to mitigate this issue mmWave uses antenna arrays composed of more than one antenna element in order to do Beamforming (BF), where the signals of each antenna element combine constructively or destructively in order to steer the energy towards the desired direction. This requires new techniques not needed at sub-6 GHz bands such as special device discovery mechanisms, beam training and beam management, especially in highly-dynamic scenarios.

1.1. Motivation and contributions

The focus of this thesis is to improve the current mmWave technologies, particularly this thesis focuses on improving TCP over mmWave, initial access, beam training, and multiband location systems. In Chapter 2 we give an in depth introduction to the new IEEE 802.11ad mechanisms and their challenges. The intrinsic characteristics of mmWave set it apart from sub-6 GHz technologies, it has very high bandwidth but the link can degrade very quickly due to blockages or even misalignments. Therefore we need to understand how current technologies perform under these new characteristics. In Chapter

3 we look at how mmWave networks perform by evaluating the behavior of the most common transport protocol, TCP. TCP is a legacy technology that gradually speeds up when there are no errors on the transmission and reacts drastically to congestions. From the point of view of TCP, a mmWave suffers from congestion problems when there is a blockage or link degradation, but usually mmWave can recover very quickly, although TCP protocols take some time to ramp up the speed. We do a thorough evaluation of different TCP protocols and scenarios and show the required parameters to improve the performance. We show that there is a need for better parameters or even new congestion protocols that are designed with the characteristics of mmWave in mind. This is only just one of the challenges of mmWave, its directionality makes device discovery and initial access a very complex process. Usually, new devices make use of multiband systems, such as 2.4 and 5 GHz in sub-6 GHz systems. This makes it possible to use out-of-band information to reduce the initial access overhead. In Chapter 4 we look into this possibility. First we do an exhaustive study of Low Frequency (LF) and High Frequency (HF) channel correlation. We consider 2 sub-6 GHz bands and 2 mmWave bands, we then find the sub-6 GHZ and mmWave combination that offers the highest correlation and then we use LF information to implement better initial access mechanism for mmWave bands. Once we have looked at the initial access, another challenge created by the directionality of mmWave communications is beam training. This mechanism consist on searching for the best combination of beams for a pair of devices. To do so, the Access Point (AP) sends beacons over each one of its sectors (between 32 and 64 in Commercial-Off-The-Shelf (COTS) devices) while the Station (STA) is listening omnidirectionally. Then they switch roles and the STA sends beacons over each one of its sectors while the AP listens omnidirectionally. Due to this complexity the mechanisms to align beams with new devices take most of the time of the connection [2, 11], which otherwise could be use for communication. In Chapter 5 we introduce a system that allows to reduce that overhead. We present mm-FLEX, an open, flexible mmWave platform. This platform allows us to implement Ultra Fast Alignment (UFA), an algorithm that completely removes the STA beam training, reducing the overhead on a factor of 8 for dense scenarios. Lastly, we again exploit sub-6 GHZ and mmWave multiband systems. Since we know there is a strong correlation we focus on improving existing mmWave location systems by using out-of-band information. Location systems are important for higher frequencies because it allows better tracking of its users. This, in turn, helps the device discovery and beam training procedures. Knowing the location of the user in real time can even help with known blockages and prevent channel degradation. In Chapter 6 we present a joint LF and HF location system. As part of this work we reverse engineer a MikroTik COTS device, and we enable Channel State Information (CSI) and Fine Timing Measurements (FTM) extraction. Using CSI and FTM data from the sub-6 GHz band in addition to the mmWave band we can overcome typical problems from a single band location system.

Namely, we can improve the accuracy of the LF location system by using HF information when in Line Of Sight (LOS), and we can use LF information to improve the HF accuracy in Non Line Of Sight (NLOS) scenarios. Additionally we do a comprehensive study on the number of antennas needed in HF to overcome the directionality limitations. Our system has an error of less than 17 cm of location error on median.

1.1.1. Software released as open source

As part of my Ph.D. I have released the following projects with an open source license:

- **MikroTik researcher tools:** This is a custom fork of OpenWRT that can be installed in any 60GHz MikroTik device. It also has a modified driver that allows to gather CSI and FTM data. This was released as part of the work included in Chapter 6. The project can be found here: https://github.com/IMDEANetworksWNG/Mikrotik-researcher-tools

- **MikroTik antenna calibration + CSI cleaning + mD-Track:** This is a full mD-Track implementation that accepts CSI measurements extracted directly from the MikroTik devices, CSI correction scripts and the instructions and scripts to perform the antenna calibration of the MikroTik devices. When used it in conjunction with the MikroTik researcher tools this enables Angle of Arrival (AOA) and ranging capabilities. This is also released as part of the work presented in Chapter 6. It can be found here: https://github.com/IMDEANetworksWNG/MikroTik-mD-Track

- **IMDEA rplidar A3 SDK for MATLAB:** A MATLAB implementation and instructions to do SLAM using the Rplidar A3. It can be found here: https://github.com/IMDEANetworksWNG/rplidar_A3_SLAM_with_MATLAB

"Because,' she said, 'when you're scared but you still do it anyway, that's brave"

Coraline, Neil Gaiman's creation

2 Background on mmWave networks

In recent years, there has been a higher demand for data-intensive wireless applications such as virtual and augmented reality, ultra wide video formats and self driving cars. This demands cannot be met with current sub 6 GHz technologies, which has brought higher frequencies into the spotlight. In 2001 the Federal Communications Commission (FCC) unlicensed 7 GHz of bandwidth between 57 and 64 GHz, which allowed operators to consider this previously unused band in order to meet their user demands. This signifies a big shift from the currently used sub 6 GHz technologies, since there are limitations that only apply to Millimeter-wave (mmWave) frequencies which range from 30 GHz to 300 GHz. mmWave bands have a much higher attenuation, even in free space, this is because this frequencies are absorbed by the gasses in the atmosphere as seen in Fig. 2.1, data from [12]. Specially, in the 60 GHz band it is the O_2 the gas which absorbs the signal. In order to mitigate this problem, mmWave devices use phased antenna arrays that enables directional communication by focusing a beam towards the desired direction. Due to the small wave length, mmWave antennas can contain a large number of antenna elements, making it possible to steer the beam towards any desired direction in both azimuth and elevation. Still, the attenuation creates some unique challenges for mmWave communications, scenarios with mobility can suffer channel degradation due to blockages, and the beams need to be trained periodically in order to ensure a good quality link. Additionally, spatial attenuation also enables high spatial reuse, enabling collision-free links in small areas. Therefore mmWave requires a full overhaul of the Medium Access Control (MAC) and Physical (PHY) layers, in a joint effort the Wireless Gigabit Alliance (WiGig) and the Wi-Fi alliances introduced the IEEE 802.11ad amendment [13, 14]. This amendment uses the wide spectrum of mmWave to provide multi-gigabit per second communications of up to 7 Gbps. IEEE 802.11ac [15] also enables multi-gigabit per second communications but achieves so by using multi-user-Multiple-Input and Multiple-Output (MIMO) whereas IEEE 802.11ad only uses 3 wide channels of 2.16 GHz each. More recently there has been a proposal for the IEEE 802.11ay amendment [16] which would enable communications of up to 100 Gbps by using MIMO at 60 GHz.

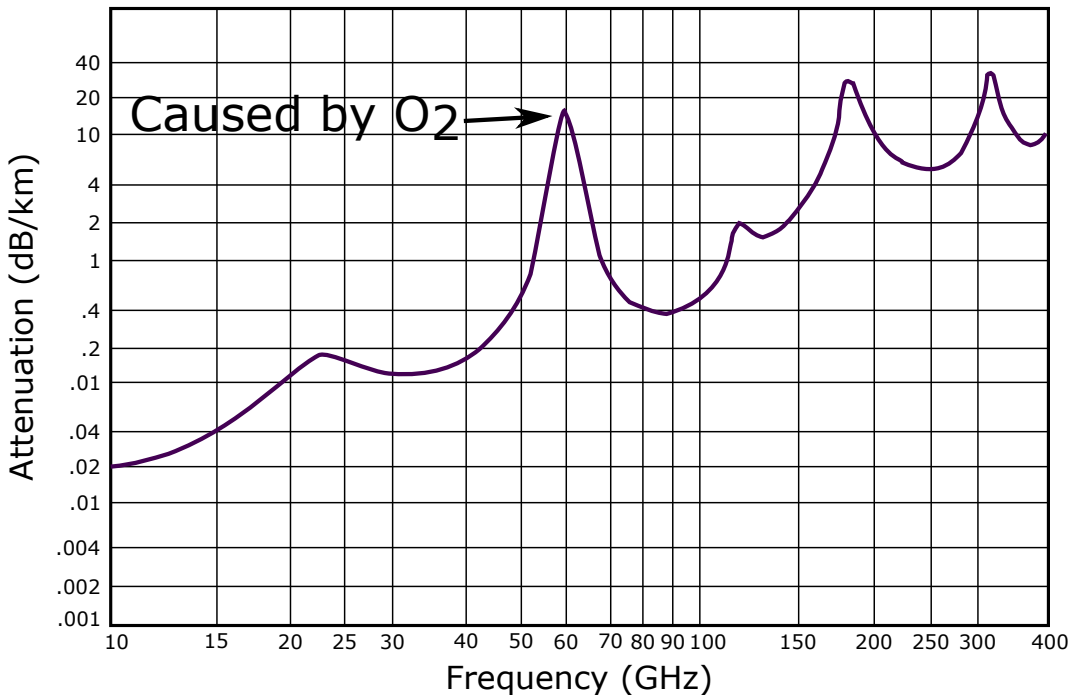

Figure 2.1: Average atmospheric attenuation for the mmWave band at 20° C, pressure of 760 mm and an H_2O concentration of 7.5 gr/m^3

However I will be focusing only in the IEEE 802.11ad amendment since my main focus was Commercial-Off-The-Shelf (COTS) that supported mmWave, and up to today the existing devices only support the IEEE 802.11ad amendment. In the following sections we will go over the necessary changes for the MAC and PHY layers and why such changes are needed. The material for this chapter is taken from [13, 17].

2.1. Physical layer

Due to the different nature of 60 GHz with respect to 6 GHZ, there are some differences in the PHY layer. The IEEE 802.11ad amendment proposes four different PHY layers, each focused on balancing the power consumption versus data rates.

- **Control PHY:** This mode is mainly used for connecting with other devices by exchanging control messages. It is mainly used during the Beamforming (BF) phase and uses Modulation and Coding Scheme (MCS)0, which is the most robust modulation, to mitigate interferences in the channel.

- **Single carrier PHY:** This is the less complex transmission option. It supports MCS from 1 to 12 and gives better choices for low powered devices such as cell phones. The speeds range from 385 Mbps (MCS 1) to 4.6 Gbps (MCS 12).

- **Low-power single carrier PHY:** This is a variation of the Single carrier

PHY which allows simpler modulation schemes. This, in turn, allows devices to further reduce power consumption.

- **Orthogonal Frequency Division Multiplexing (OFDM) PHY:** In order to achieve the fastest data rates IEEE 802.11ad introduces the OFDM PHY, with MCS 12-24 it can achieve data rates of 6.7 Gbps. This is an optional mode and vendors can decide not to implement it. Normally this PHY is found in more powerful devices, such as laptops and docking stations.

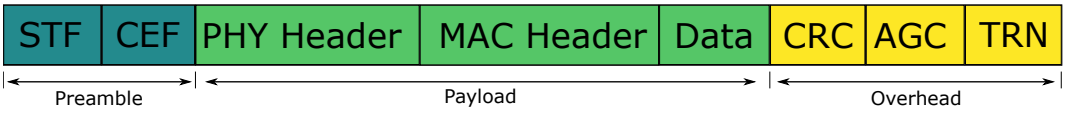

Figure 2.2: IEEE 802.11ad frame structure

Fig. 2.2 shows the structure of a IEEE 802.11ad frame. The Short Training Field (STF) and Channel Estimation Field (CEF) fields are used for Automatic Gain Control (AGC), and estimating the overall channel status. The PHY field contains the size of the packet and the PHY layer used. The MAC header and data fields contain the information of the destination device. Lastly, the Cyclic Redundancy Check (CRC) field contains a checksum and, optionally, the AGC and Training (TRN) fields are appended to aid on the BF training.

2.2. Beacon interval

In order for new wireless devices to access the medium, the IEEE 802.11ad amendment proposes a few changes to the Wireless Local Area Network (WLAN) Beacon Interval (BI). These changes are needed due to the directional nature of the mmWave communications, in sub 6 GHz communications it is enough to send one frame omnidirectionally, but in mmWave we need to send the beacons over many directions. Additionally there is a need for beamforming training, to choose a pair of aligned directional beams, which is also done during the BI. Fig. 2.3 shows the BI, which is divided between the Beacon Header Interval (BHI) and the Data Transmission Interval (DTI).

The BHI is further divided in 3 intervals, the Beacon Transmission Interval (BTI), the Association Beaforming Training (A-BFT) and the Announcement Transmission Interval (ATI). During the BTI the Access Point (AP) sends one beacon over each one of its sectors, this is further explained in Sec. 2.3. During the A-BFT the Directional Multi-Gigabit (DMG) Station (STA)s can train their transmit sectors towards the AP. Lastly, the ATI is used to exchange management frames between the AP and the STA.

During the DTI period the STAs can communicate with the AP using either Service Period (SP) or Contention-based Access Period (CBAP). During the SP the STAs have

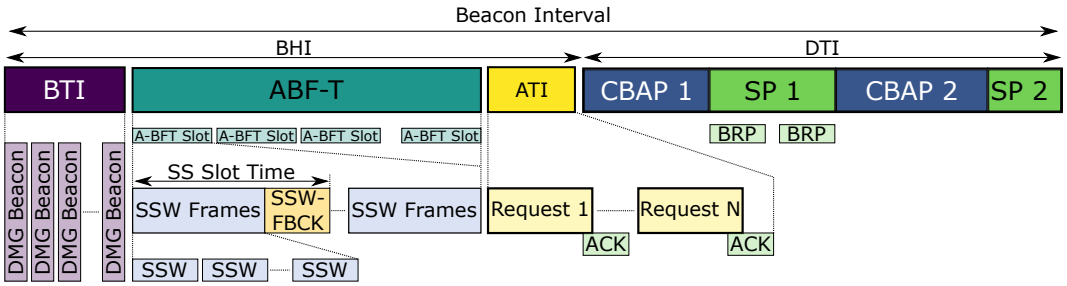

Figure 2.3: IEEE 802.11ad beacon interval

a reserved channel and can communicate with the AP directly. During the CBAP period the STAs need to contend for the channel. This methods are further explained in Sec. 2.4.

2.3. Beamforming training

To comply with the IEEE 802.11ad amendment Beamforming training must be implemented. In order to overcome the high spatial attenuation, 60 GHz-enabled devices uses phased antenna arrays that are able to steer the beam towards the desired direction. Each device has a precomputed set of sectors where each sector covers a different direction. In order for two DMG devices to establish a connection they must perform the Sector Level Sweep (SLS) phase, which has two parts. During the first part of the SLS phase the Initiator sends one Sector Sweep (SSW) frame via each sector while the Responder uses an omnidirectional beam pattern to capture the frames. Once the Responder knows which is the best Initiator's transmitting sector, they both switch roles and the Responder transmits SSW frames via each one of its sectors while the Initiator listens omnidirectionally. Lastly, the Responder and the Initiator exchange information on which are their best sectors and the communication can be established, Fig. 2.4 shows an overview of that mechanism.

There are two additional steps included in the Beamforming training, although both are optional. First the Beam Refinement Protocol (BRP) allows to refine the sectors chosen during the SLS phase, they do so by using a BRP packet with either a Transmit Training (TRN-T) or a Receive Training (TRN-R) field. This mechanism allows the DMG devices to use larger sectors during the SLS phase, making it faster, and then choose a smaller sector with higher gain during BRP. Lastly, there is a Beam Tracking (BT) phase. Since the mmWave links can degrade quickly due to movement or channel degradation, IEEE 802.11ad introduces BT. BT keeps a timer that gets reset every time a packet is successfully received, if the timer expires it is assumed that the link is broken and the SLS phase is triggered again.

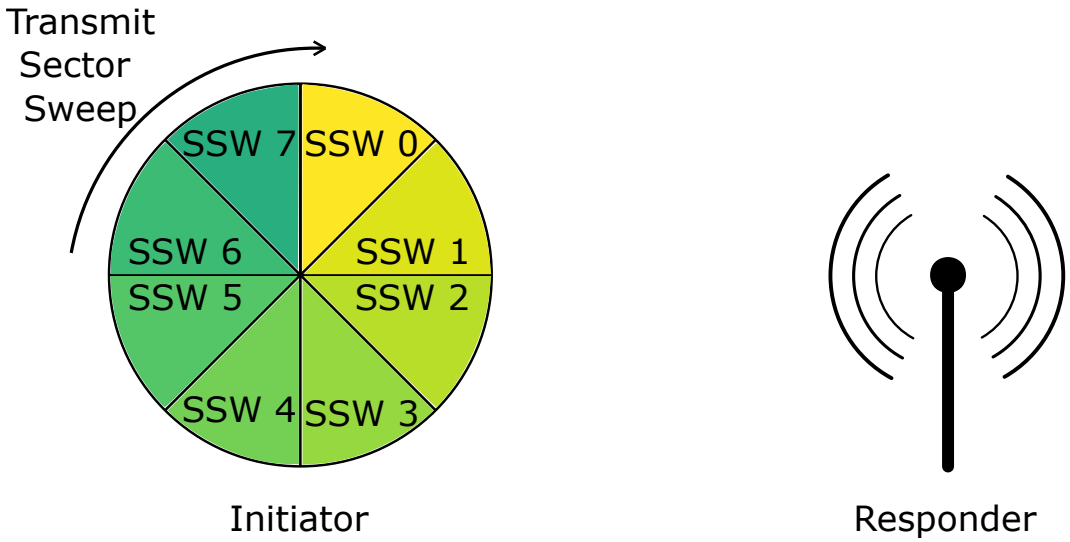

Figure 2.4: Beamforming training in IEEE 802.11ad

2.4. Channel access

IEEE 802.11ad allows STAs to access the channel in 3 different ways. In legacy WLAN technologies the medium is accessed through Carrier-Sense Multiple Access with Collision Avoidance (CSMA/CA), but IEEE 802.11ad introduces 2 new methods, SP and Dynamic channel access.

- **CSMA/CA channel access:** This medium access scheme does not require a centralized scheduler, and it has been implemented in all the previous IEEE 802.11 amendments. It is typically used because it works well with a reasonable number of devices and it is very easy to implement. However, this scheme does not work well for mmWave communications, since the communications are directional. If two DMG devices try to access the medium at the same time they would likely not hear each other and believe the medium free, causing collisions.

- **SP channel access:** During the SP the AP reserves the channel so that only 2 DMG devices can communicate to each other. This, in turn, allows the rest of the STAs to go to sleep mode to reduce power consumption and enables higher data rates since there is no risk for collision.

- **Dynamic channel access:** Lastly, IEEE 802.11ad allows STA to poll for transit requests. In turns, each STA sends the time it needs and the AP schedules a contention free period for that STA. These requests can be granted during the BI in order to obtain better utilization of the channel.

"What makes me a good student? Well, if I were a bad student,
I wouldn't be sitting here writing about it, now would I?"

Tavish DeGroot

3 Analysis of TCP Performance in 5G mmWave Mobile Networks

3.1. Introduction

The extreme propagation conditions of Millimeter-wave (mmWave) links impact the transport layer and especially the congestion control mechanisms. The role of congestion control is to regulate the amount of injected traffic in the network according to its congestion state. However, in wireless communications, traditional congestion control protocols like TCP NewReno are unable to differentiate between losses attributed to congestion and those attributed to transmission errors caused by a decay in channel quality. This well known problem is vastly exacerbated in mmWave networks because of the magnitude of rate variations that sudden transitions from Line of Sight (LoS) to Non-LoS (NLoS) cause. At the transport layer, the sender is not notified about such short-term changes and might not decrease immediately the congestion window. Instead, lower layers react immediately by lowering the Modulation and Coding Scheme (MCS) used for subsequent data transmissions to increase robustness against adverse channel conditions. Preliminary works in this area analyzed the system-level implications of mmWave channels on transport protocols [18], and have focused on understanding the performance of Multipath-TCP and the impact of link level retransmissions [19]. Additionally, Azzino et al. [20] proposed a cross-layer optimization of the congestion window, that is set according to the bandwidth-delay product by considering the latency estimated without buffering delays.

Throughout this chapter, we provide a systematic study of end-to-end performance of TCP congestion protocols in mmWave 5G mobile networks. TCP is the de facto transport protocol used by the majority of the applications such as HTTP, Skype, file transfer and email [21]. Our contribution is to identify critical design aspects of congestion control protocols and highlight transport layer performance trade-offs particular to mmWave networks. Compared to previous research in [18–20], the objective of this work is provide a comprehensive analysis, taking into consideration several loss-, delay-based and hybrid congestion protocols and studying their performance for different applications and

scenarios: i) a single user served by one AP (Section 3.4.1), ii) a single user performing handover between multiple APs (Section 3.4.2), and iii) multiple users exploiting different applications (FTP, video streaming and instant messaging) served by multiple APs (Section 3.5).

3.2. Congestion Control Protocols over mmWave Links

Channel quality in cellular networks is unpredictable for several reasons, including device mobility and handovers, frame scheduling algorithms which create burstiness and the transitions of the Radio Resource Control (RRC) state machine [22]. Congestion control protocols probe the channel to exploit the available rate by injecting packets until losses occur, or proactively adapt the injection rate to match delay- or rate-based measurements. Legacy congestion control protocols were primarily designed for wired networks assuming fixed capacity links, and hence suffer in the presence of the short term link quality variations common in cellular networks. These issues are further exacerbated in high-speed mmWave networks. Large and persistently full buffers are typical in cellular networks to help maintain a high throughput, but they lead to the bufferbloat problem, which significantly increases latency [23]. The following paragraphs revisit congestion control mechanisms, highlight challenges and identify critical issues specific to mmWave networks.

3.3. Fundamentals of TCP

TCP ensures that all the data packets are correctly received and in order. To this end, TCP requires each packet to be acknowledged by the receiver (ACK). TCP relies on a Congestion Window (CW) which determines the maximum quantity of data that can be sent within one round trip time (RTT). The CW adaptation procedure depends on the link status. During TCP slow start, the CW increases by one packet per received ACK until: (i) the *sstresh* threshold is reached, or (ii) there is a packet loss. Upon reaching *sstresh*, TCP enters the congestion avoidance phase and the CW grows, for example for the NewReno protocol, by one packet whenever a whole congestion window worth of data has been acknowledged. When the receiver obtains a packet that has a sequence number higher than the expected one, a duplicate ACK is triggered. This informs the sender that the packet with the expected sequence number has been lost and requires retransmission. If three duplicate ACKs are received, TCP assumes a packet loss: it enters in fast retransmit mode by halving the CW and continues in congestion avoidance phase. If after a given period the Retransmission Time Out (RTO), commonly set to 1 s, no ACKs were received for a packet, TCP reduces the CW to one packet and recovers from the slow start phase.

3.3.1. Congestion Controls Protocols

Congestion control protocols differ in the mechanisms utilized to adapt to the available bandwidth. According to the methodology employed to detect congestion, protocols can be attributed to three main categories: loss-, delay-based and hybrid. To obtain insightful results, a preliminary study narrowed down the initial set of protocols from twelve (BBR, CUBIC, HighSpeed, HTCP, Hybla, Illinois, New Reno, Scalable TCP, Vegas, Veno, Westwood, WestwoodPlus, and YeAH [24]) to seven. Preference to widely-known protocols and those implemented in current networks were the selection criteria. For example, Hybla was discarded being primarily used in high-latency terrestrial links or satellite links. The seven different protocols which encompass all the aforementioned categories loss-, delay-based and hybrid are: NewReno, Scalable TCP, CUBIC, Vegas, Westwood, YeAH and BBR. NewReno relies on packet losses as indication of congestion, halving the size of the CW whenever a loss is detected. Compared to the older Reno TCP, the Fast Recovery mechanism is designed to increase robustness to multiple losses in a single window. While NewReno follows a Additive Increase Multiplicative Decrease strategy, Scalable TCP implements a Multiplicative Increase Multiplicative Decrease strategy to growth the rate more quickly. The rate increase is proportional to the estimated spare bandwidth of the link. CUBIC increases the CW according to a cubic function by computing the absolute time since the last dropped packet. In the first part, the function is concave and CUBIC attempts to quickly reach the CW size registered before the last dropped packet. In the second part, the function is convex and CUBIC conservatively probes for additional bandwidth. Vegas implements congestion control on the basis of delay-measurements to proactively estimate the buffer status of the bottleneck router. The CW size increases or decreases with an Additive Increase Additive Decrease strategy by computing the difference between the actual rate and the ideal one that the TCP flow would achieve without congestion. Westwood adjusts parameters like *sstresh* and CW according to estimations of the available network bandwidth based on measurements of the average rate of received ACK packets. YeAH is a hybrid approach, which determines the CW on the basis of both losses and delay that makes it more robust to LoS-NLoS transitions. This is because YeAH employs two modes, the *fast* and the *slow* modes. The *fast* mode uses an aggressive rule in the spirit of Scalable TCP to quickly grow the CW, while the *slow* mode acts like NewReno, hence the state is determined by the estimated number of packets in the bottleneck queue. When YeAH estimates that packet buffering and queuing delays in the network are below or above predefined thresholds, it switches between *fast* and *slow* mode. The presence of NLoS enforces the *slow* mode in YeAH, which better copes with the low available bandwidth. Similarly to YeaH, TCP Bottleneck Bandwidth and RRT (BBR) also continuously measures RTT. Unlike the previous congestion control protocols, BBR is model-based. Specifically, BBR

builds a model of the network by continuously measuring the available bandwidth at the bottleneck link and the two-way propagation delay. In this way, it estimates the bandwidth-delay product and accordingly adjusts the sending rate through pacing.

3.3.2. Effects of mmWave Channel Properties on Upper Layers of the Protocol Stack

Under LoS conditions, a mmWave link can offer several Gbps of throughput [25]. During the congestion avoidance phase, TCP tries to take advantage of all of the available bandwidth. However, if the channel becomes blocked by an obstacle, the data rate drastically drops. Sudden transitions between LoS and Non-LoS make the intermediate router buffers to fill very rapidly due to the high data rate, until the fast retransmit phase starts or a RTO is triggered. When entering the fast retransmit phase, the value of the CW is halved. However, this might not be sufficient to adapt to the new low link capacity and might in turn cause a RTO. TCP then resumes with slow start to correctly adjust to the available bandwidth. Upon a RTO, the slow start threshold is reduced, which causes TCP to take longer in achieving the optimal value for the CW and obtain high throughput, especially when the channel quality improves from NLoS to LoS.

Retransmissions at lower layers of the mobile network stack, i.e., the Radio Link Control (RLC) and the MAC layer, play an essential role in maintaining throughput as they hide channel losses to TCP [26]. The RLC in Acknowledged Mode (AM) splits the TCP Protocol Data Units (PDUs) into smaller chunks (RLC PDUs) that are acknowledged by the RLC layer at the sender side. The Automatic Repeat Request (ARQ) on the RLC layer is in charge of asking for the retransmission of corrupted RLC PDUs. Hence, augmenting the RLC buffer size is a solution to compensate for high number of losses. Although the maximum number of retransmission attempts per RLC PDU is limited, this operation increases RTTs and, in turn, might trigger TCP timeouts. To reduce the number of retransmissions on the RLC layer, a Hybrid ARQ (HARQ) mechanism at the MAC layer can recover corrupt packets by soft combining two or more corrupted RLC PDUs stored in its buffer. The minimum duration of the RTO timer typically employed in current wireless networks is large and not suitable for mmWave systems. While in presence of short blockages entering in fast retransmit phase is beneficial, during extensive NLoS periods TCP is likely to trigger multiple RTOs. Thus, TCP can significantly benefit from the reduction of the waiting period before it recovers from the slow start phase.

3.4. TCP Performance Analysis

This section analyzes the performance of multiple congestion control protocols. For the simulations we use the mmWave module of ns-3 [27], which implements the complete

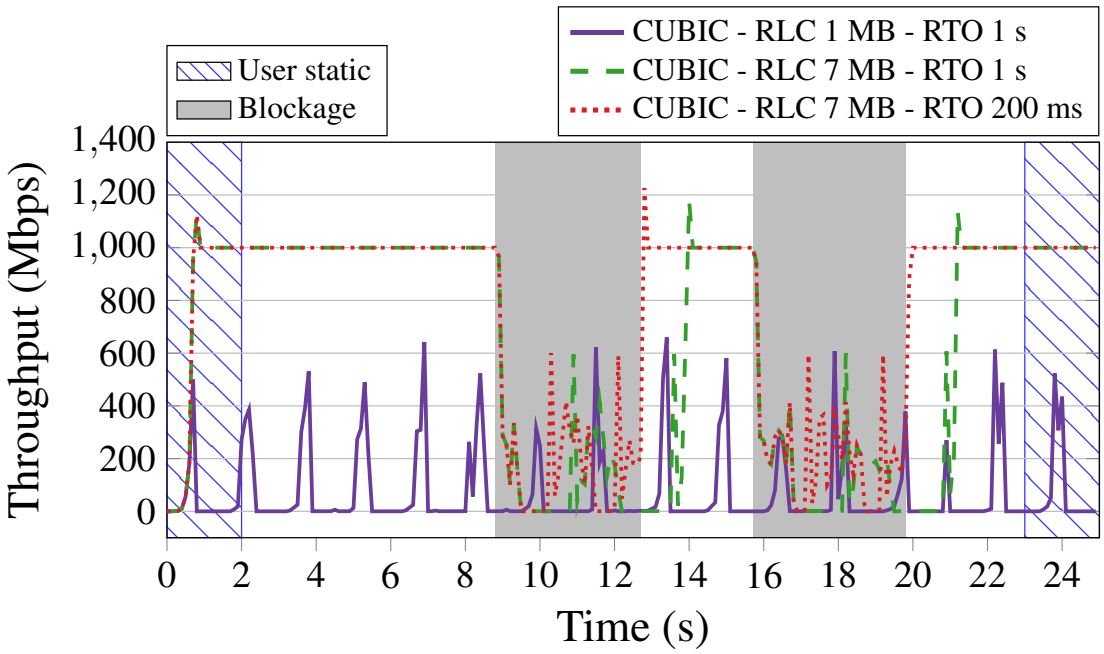

Figure 3.1: Throughput comparison with different configurations of CUBIC

protocol stack and accurately represents the underlying mmWave channel according to the 3GPP channel model.

3.4.1. Single Flow Analysis

A mmWave AP working at 1 GHz bandwidth, 30 dBm maximum transmit power and a carrier frequency of 28 GHz serves one user equipment (UE) at a distance of 50 m. The UE remains static for 2 s, then moves following a straight path parallel to the AP location at a walking speed of 1.5 m/s for 21 s and finally it stops and remains static for another 2 s. Along the path, the connection is interrupted. Three blockage scenarios are considered: i) a building creating an extensive blockage of 13 s, ii) two small buildings creating medium blockages of 4 s each, and iii) 6 small obstacles creating multiple short blockages of 0.25 s each. Scenario i) and ii) permit to analyze the performance of congestion protocols when the penetration loss remains high for a prolonged period of time. Finally, Scenario iii) extends the analysis of Scenario ii) augmenting the number of LoS-NLoS transitions and shortening their duration. In the presence of medium or extensive blockage, congestion control protocols can detect the link capacity reduction while in the presence of temporary short-term blockage this is not possible as the blockage ends too rapidly.

While in traditional mobile technologies like 3G/4G, RLC buffer sizes on the order of 1 MB were sufficient to compensate for wireless losses, in mmWave networks the size needs to be adjusted for the higher data rates. A prior analysis through experimentation in different scenarios and for all TCP protocols showed that a 7 MB RLC buffer size achieves maximum throughput and limits the bufferbloat problem. Fig. 3.1 shows the

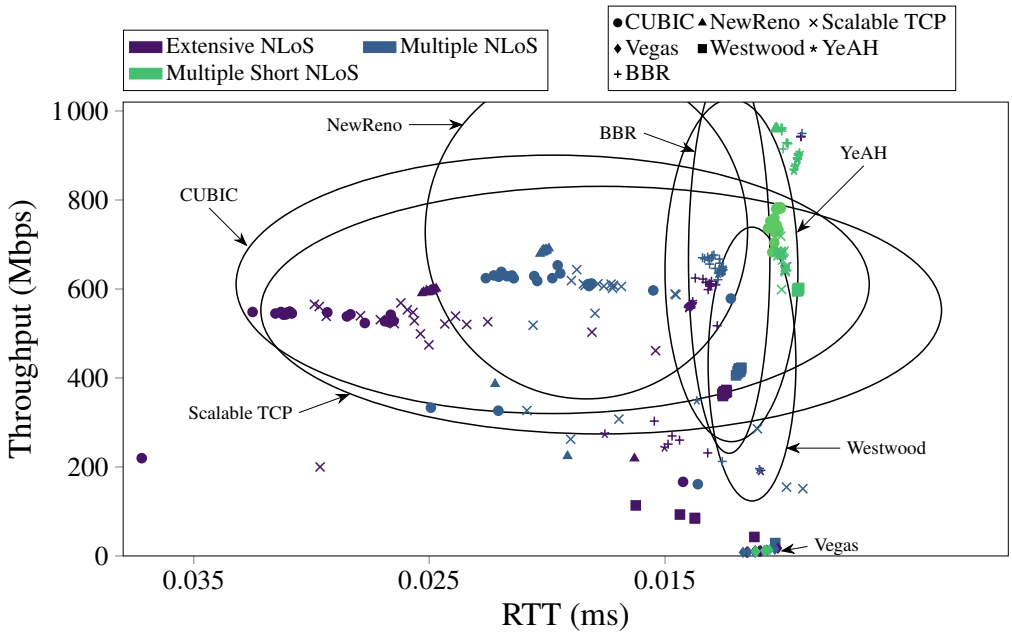

Figure 3.2: Performance of TCP congestion protocols for different scenarios. The elliptic contour defines the confidence interval with a $2-\sigma$ precision, hence the outliers are those samples outside the 95% confidence interval region.

results of this analysis. Note that the dimension is conservative with respect to the set-up of [20] and is in line with the 3GPP 36.306 TS Release 14 specification for a category 12 UE. Fig. 3.1 shows throughput of CUBIC comparing the joint RLC buffer size and RTO timer with modified setting, e.g., with an RLC buffer of 7 MB and and RTO value 200 ms, an RLC buffer of 7 MB and RTO of 1 s, and the default settings of an RLC buffer of 1 MB and RTO of 1 s. The limited RLC buffer size of the baseline setting causes CUBIC to suffer considerably from wireless losses and the protocol persistently fails to reach the congestion avoidance phase. Additionally, the duration of the RTO timer prevents CUBIC from reacting quickly. Hence, the achieved throughput in LoS phases is suboptimal. The importance of the timer becomes evident when comparing RLC versus RLC+RTO with modified setting: in NLoS periods, the latter provides a significant advantage to recover faster from NLoS-LoS transitions.

Fig. 3.2 shows performance of the different protocols after having updated the setting of RLC+RTO as previously discussed for the considered scenarios. The graph is a throughput-RTT plot, where for each scenario we take the average results and compute the $2-\sigma$ elliptic contour of the maximum-likelihood 2D Gaussian distribution. The $2-\sigma$ expression defines the level of confidence, i.e., the projection on a one-dimensional sub-space of the elliptic contour is the 95% confidence interval. The plot shows the results for individual scenarios and the overall congestion control protocol performance with different markers. On the x-axis, lower (better) RTTs are to the right. Hence, best performing protocols are on the top-right. Throughput-RTT plots highlight the variability and relative performance between protocols. The more narrow the ellipses in the axis

dimension, the more stable is the protocol in consistently achieving similar throughput or RTT performance. On the other hand, wider ellipses indicate higher variability. The ellipses' orientations define the relationship between throughput and RTT.

Both CUBIC and Scalable TCP aggressively increase their rate, with Scalable TCP implementing exponential increase/decrease dynamics, while the congestion window in CUBIC grows as a cubic function of the time elapsed since the last congestion event (see Section 3.3.1). Specifically, CUBIC's congestion window growth is fast or slow when the current window is respectively far and close to the target one. This mechanism enables prompt recovery and is particularly suitable for short NLoS periods. In presence of a LoS to medium or extensive NLoS transition, both Scalable TCP and CUBIC at first attempt to recover to the sending rate achieved during LoS. However, due to the high channel quality degradation, both protocols repeatedly attempt recovery from slow start. As a result, their performance differs considerably from one scenario to another and the protocols are very susceptible to LoS-NLoS transitions in RTT metrics.

NewReno and Westwood exhibit less variability in RTT than CUBIC and Scalable TCP. Westwood, upon detection of congestion, computes the current bandwidth-delay product and sets the slow start threshold accordingly. NewReno does not terminate the fast recovery phase until complete recovery of multiple losses and thus take longer to achieve the full link capacity. In the congestion avoidance phase, both protocols gently probe for additional bandwidth.

YeAH congestion control is based on both packet loss and RTT measurements and its objective is to limit the amount of buffering in the network (see Section 3.2). Hence, YeAH does not increases the rate further once determines that the sending rate is sufficiently high. As queuing delay computations are based on the minimum of the recently measured RTT and not the average RTT, in presence of blockage, YeAH quickly switches to slow mode which prevents recovery from slow start. BBR congestion protocol has a very similar performance to YeAH as it adapts its rate to minimize queuing over the whole network, but also achieves a lower RTT among all the scenarios. Finally, Vegas consistently achieves the lowest throughput and exhibits little variability in RTT. Indeed, the Vegas congestion control mechanism sets the data rate so that all transmitted packets can be acknowledged within the minimal RTT. Upon finding such a rate, Vegas does not attempt to increase the rate to use a higher fraction of the link capacity under-utilizing the high bandwidth at disposal in LoS significantly.

For performance evaluation, in the next sections CUBIC and YeaH are selected because of the different level of robustness to blockage. While Scalable TCP provides similar performance than CUBIC, the latter is the de-facto protocol implemented from Linux kernel v2.6.19.

3.4.2. Single Flow with Handover Analysis

Handovers transfer the ongoing data session from the current AP to another, to ensure that a UE remains connected while it is moving. While in traditional cellular networks handovers are mainly triggered by user movements out of the operational range of an AP, in mmWave networks also blockage is a probable cause for handovers. There exist several methods to trigger a handover [28]. The threshold method estimates the Signal to Interference to Noise Ratio (SINR) degradation. The *fixed* and *dynamic* Time to Trigger (TTT) generate a handover event if the SINR is below a certain threshold during a predetermined or variable time window, respectively. The objective of the timers is to reduce the ping pong effect, i.e., a UE switching between two APs in a rapid succession. The following analysis employs a fixed TTT with 200 μs time window as per 3GPP TS 36.331 V8.4.0.

Fig. 3.3(a) shows the geometry of the scenario, which mimics a typical street. Lamp posts are distributed at regular intervals of 20 m and there are two double lanes for cars next to which are pedestrian lanes of 2 m each. Three APs, deployed at a regular distance of 20 m, ensure full coverage. The UE moves at a constant speed of 1.5 m/s on the bottom pedestrian lane from left to right. The lampposts create multiple short blockages similar to the multiple NLoS scenario analyzed in Section 3.4.1. Additionally, a 13.5 m wide bus parked on the opposite lane of the user generates extensive blockage for AP2 and AP3.

Fig. 3.3(b) shows the SINR as the UE moves along the pedestrian lane. The colored backgrounds represent NLoS periods to the corresponding APs. The color of the SINR profile represents the AP to which the UE is currently attached. When an obstacle blocks the current AP, a handover event is automatically triggered. However, in the presence of short-term blockage, this decision is suboptimal as the connectivity rapidly moves from one AP to the other causing a throughput decrease (see Fig. 3.3(c)). The graph shows results for CUBIC and YeAH protocols that exhibit different levels of robustness towards blockage (see Section 3.4.1). Unsurprisingly, during rapid handover events CUBIC drops its rate and recovers quickly from slow start. YeAH, instead, switches to slow mode and reaches the link capacity slowly.

3.5. Multiple Concurrent Flows Analysis

Previous research showed that long flows in the background can significantly reduce the responsiveness of short flows [23]. This section specifically analyzes performance of TCP over mmWave channel dynamics and rates when short flows and background traffic coexist.

Fig. 3.4 shows a more complex scenario where 10 users with different applications are served by 4 APs in a square. The scenario represents a dense small cell deployment with a

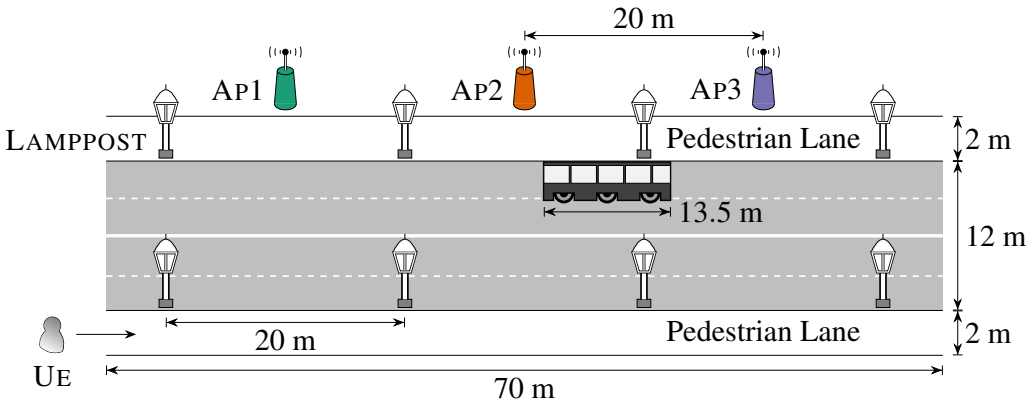

(a) Scenario

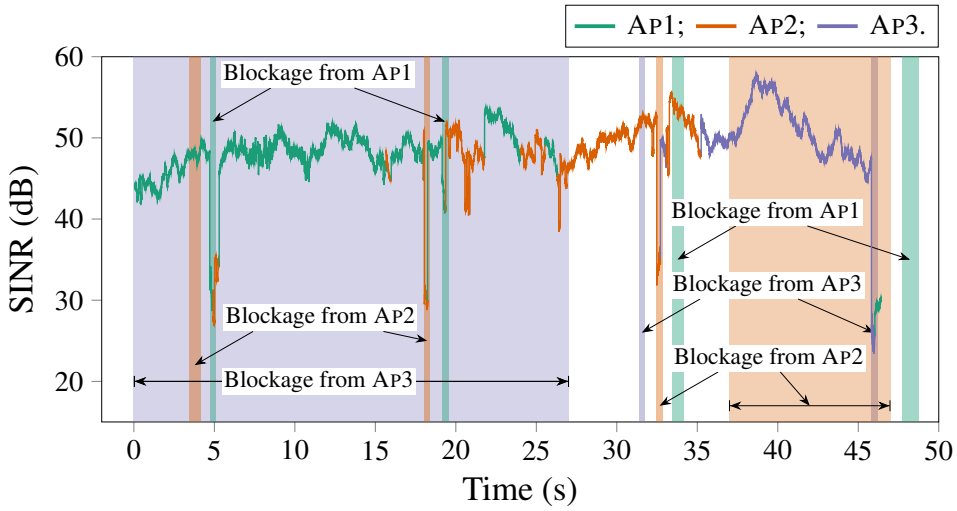

(b) SINR

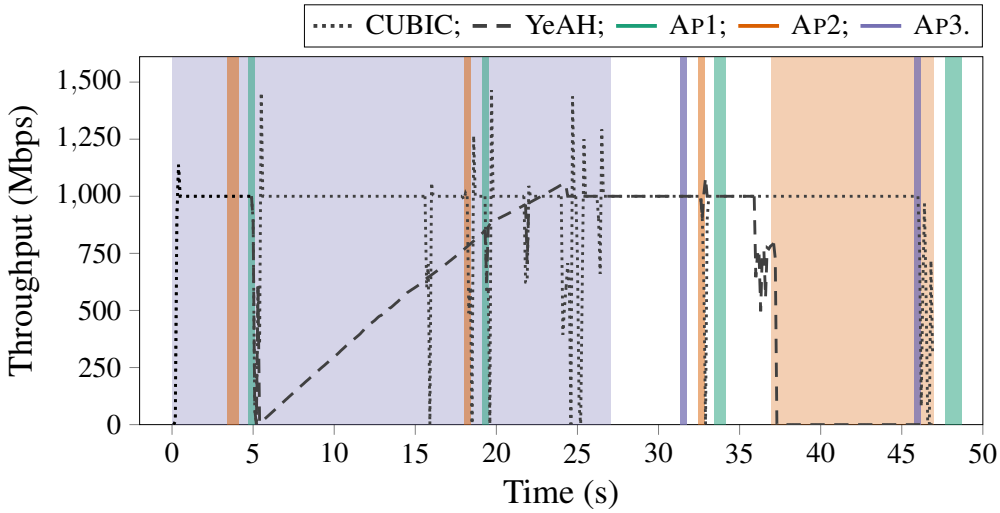

(c) Throughput

Figure 3.3: Handover analysis: scenario (a), SINR (b) and throughput (c)

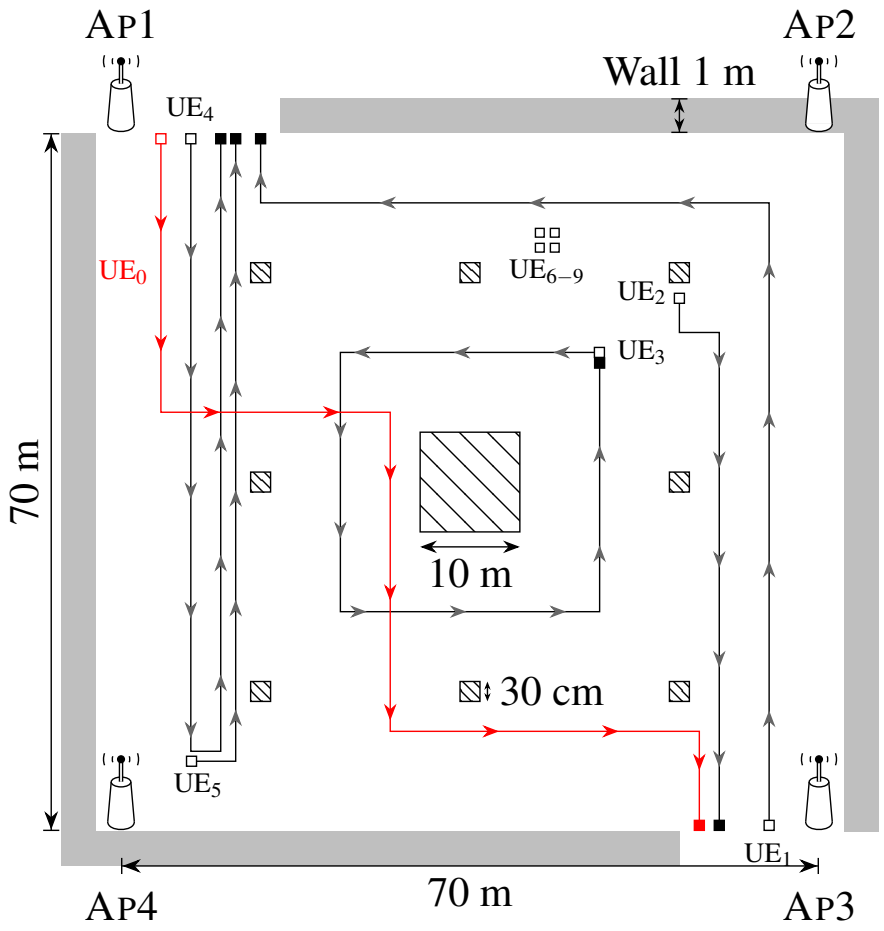

Figure 3.4: System scenario

statue in the center of the square and lampposts placed all around the statue that create blockages. UE0 has a throughput of 500 Mbps mimicking a FTP download. This flow starts at the beginning of the simulation and lasts until the end. The rest of the users are either streaming a video (UE2, UE4 and UE5) or using instant messaging applications. The video applications buffer data at 4 Mbps for 30 s and then stop their flow for 2 s before starting again. The rest of the users are using instant messaging applications or social networks, sending a few packages every few seconds. These applications are common for smartphones. Except for UE6−UE9, all UEs experience handovers while the user is walking, and eventually move the active connection to an AP that is already serving other flows. CUBIC and YeAH are the congestion control protocols used for the comparison.

Impact of Retransmissions

Fig. 3.5 compares the number of successfully received packets and retransmissions. UE0's application is the one generating the heaviest flow, hence the experiment allows to verify the performance of medium-small flows with heavy background traffic. The

complexity of the scenario also allows to verify the effect of medium flows on small ones. As UE0 moves through the environment, its connection changes from one AP to another due to handovers. Consequently, the rates of the UEs that were already connected to this AP are degraded to accommodate the incoming long flow, or their connection is moved to other APs. The combined effect of blockage and presence of the long flow is different for medium and small flows and for the two congestion control protocols. With the sole exception of UE4, YeAH consistently achieves a lower number of retransmissions. Recalling the result in Fig. 3.2, by design YeAH reacts to congestion based on both packet loss and RTT measurements. Limiting the sending rate on this basis, it achieves higher robustness than CUBIC, whose performance is strictly dictated by the duration of blockage causing packet losses. Nonetheless, for UE0 YeAH achieves data rates up to 500 Mbps similar to CUBIC. For all the others UEs and for both congestion control protocols, the achieved data rates are lower because the amount of data to be transmitted is small.

UE2 has comparatively the highest number of retransmissions with both protocols. Indeed towards the end of its movement, it competes with UE0 and the network makes UE2 to frequently handover between AP3 and AP4. UE4, although following a similar path as UE0, is not severely affected by UE0 and is served by AP4 for the entire period. Hence, YeAH outperforms CUBIC because the retransmissions are solely due to errors. Unlike UE4, UE5 is served by multiple APs during the movement and, consistent with the results obtained in Subsection 3.4.2, CUBIC is more affected than YeAH. The performance of static users UE6−UE9 is similar. They are scheduled concurrently with the medium-sized flow of UE1 and are partially affected from the long flow. YeAH provides consistent results and does not incur retransmissions, while CUBIC is more susceptible to blockages and requires more retransmissions. Specifically, CUBIC uses up to five times the RLC buffer space compared to YeAH. Unlike UE6−UE9, UE3's connectivity is persistently blocked by small obstacles, however its performance is comparable to that of static users in NLoS.

Analysis of RLC Buffer Occupancy

Fig. 3.6 analyzes the effect of the long flow on a medium-sized one. The plot shows the CDF of the contribution from UE4 to the RLC buffer occupancy of AP4 for both CUBIC and YeAH. Note that UE4 is constantly in LoS and served by AP4 during its movement. Interestingly, in absence of the long flow from UE0, both congestion control protocols lead to the same buffer occupancy. Hence, the hybrid ARQ mechanism at the MAC layer is sufficient to hide wireless errors from the RLC layer and the transport layer. In presence of the long flow in the background, this is no longer true and AP4 provides connectivity to both long and medium flows from UE0, UE4 and UE5 respectively. Consequently, the contribution of UE4 to the RLC buffer occupancy increases. Unsurprisingly, CUBIC and

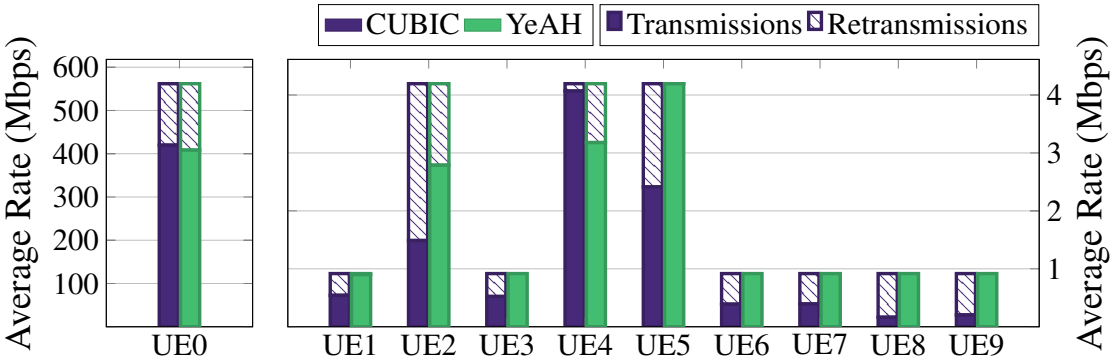

Figure 3.5: Composition of successful transmissions and retransmissions with CUBIC and YeAH

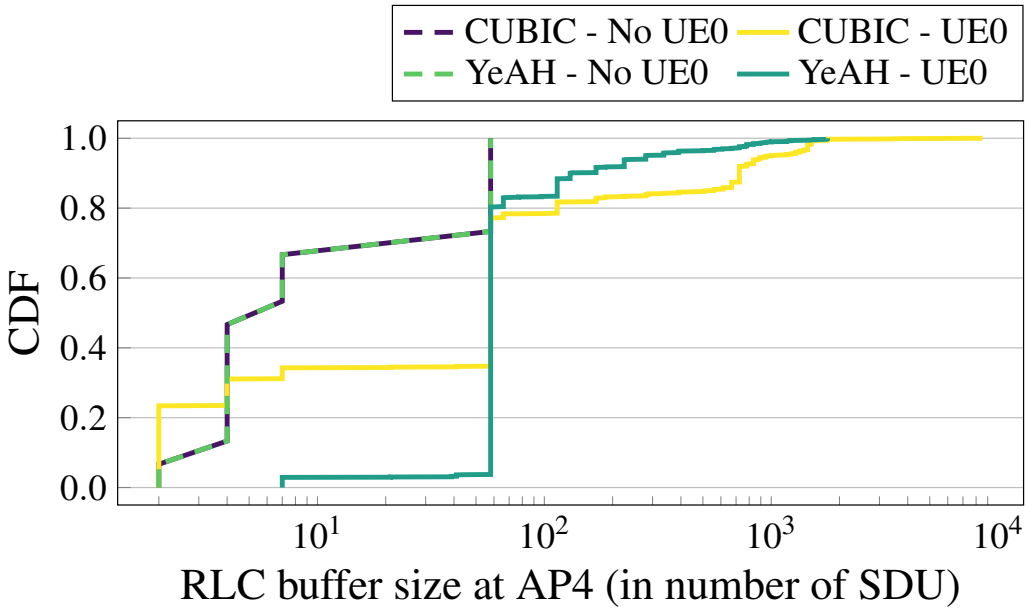

Figure 3.6: CDF of AP4 RLC buffer occupancy for CUBIC and YeAH with and without UE0 traffic in background

YeAH performance differs and CUBIC causes higher buffer occupancy on average, while achieving higher rates.

3.6. Discussion and Open Challenges

Proper setting of RLC buffer and RTO timers is essential to alleviate the impact of losses on throughput and foster prompt recovery from slow start. RLC buffer size can not be arbitrarily augmented to prevent delay increase and bufferbloat. Given the short term variations of link quality, quick recovery from the slow start phase is crucial. This can be achieved by shortening the timeout value. Further investigation on optimizing buffer size and timers is an interesting direction for future research, especially to foster ultra-low latency and reliable communications. Another promising research direction involves

cross-layer optimization of transport, MAC and PHY layers. For example, more advanced solutions for directional antennas and beamforming may potentially benefit from machine learning techniques to learn the properties of the environment and optimize resource allocation.

Congestion control protocols that aim at maximizing throughput through quick recovery like CUBIC suffer with longer NLoS periods while they perform comparatively well in the presence of short NLoS. On the other hand, the performance variations of protocols that detect congestion through hybrid mechanisms (packets loss and RTT measurements) like YeAH is minimal. This tradeoff is especially evident in presence of handovers: while CUBIC resumes from slow start upon any handover event, YeAH probes bandwidth more gently and prevents short-term throughput variations. When long, medium and small flows coexists, CUBIC always requires more transmissions than YeAH to successfully complete the data transmission. Hence, YeAH is suitable for transmission of small flows over mmWave channels. However, CUBIC outperforms YEAH in terms of achieved throughput.

3.7. Conclusions

This chapter studied the behavior of multiple congestion control protocols over mmWave networks. In contrast to traditional mobile networks, mmWave links have unique features, including high available bandwidth, high variability in channel quality and sensitivity to blockage because of high propagation and penetration loss and atmospheric absorption. Through extensive simulations, the chapter has shown aspects of packet loss-, delay-based and hybrid congestion control protocols by analyzing the impact on throughput and latency of various environments with blockages. In the next chapter, we will be focusing on how to overcome some of these limitations by reducing the Beamforming (BF) overhead.

"When you really know who you are and what you like about yourself, changing for other people isn't such a big deal"

Abed Nadir

4 A Comprehensive Study of Low Frequency and High Frequency Channel Correlation

4.1. Introduction

As explained in Chapter 1.1.1, for Millimeter-wave (mmWave) communications it is necessary to do directional communications. Therefore new Medium Access Control (MAC) mechanisms are necessary, specifically Beamforming (BF). This, unfortunately, adds a lot of latency and overhead when used in dense or mobile scenarios since the BF has to be done periodically with each Station (STA). When a STA is moved or rotated, the link quickly degrades, that may be because of obstacles or link misalignment. This causes a drop in Signal-to-Noise Ratio (SNR), a change to a lower Modulation and Coding Scheme (MCS) and even outages in the most extreme cases, which limits the channel efficiency. Commercial-Off-The-Shelf (COTS) devices usually trigger full BF training after the link has downgraded under a certain threshold, and in the worse cases this can take seconds [29, 30], although this can be reduced to milliseconds [31]. Since mmWave has gigabit per seconds communication links, even a short downtime period has an extreme impact on channel quality.

The viability and synergy of mmWave with cellular networks has been extensively investigated in the literature, with numerous works characterizing higher frequencies in isolation [32]. However, we believe that investigating both bands jointly is the key to maximize multiband system performance.

To the best of our knowledge, only a few works have analyzed in depth the correlation between Low Frequency (LF) and High Frequency (HF). Moreover, when the works are based on simulations, it is very common to abstract and simplify some real world phenomena and hardware limitations. Recent results, such as [33], show that scenarios in Line Of Sight (LOS) exhibit a high correlation in most cases. This high correlation drops significantly in Non Line Of Sight (NLOS) scenarios, where the strongest HF paths often do not match with the strongest LF ones in 60% of the cases [34].

Going one step further, the authors in [35, 36] propose inferring spatial correlation

matrices in mmWave channels from the LF Channel State Information (CSI). The estimation of the Angle of Arrival (AOA) can be used to reduce the overhead of traditional beam-steering techniques at HF. Further developing this idea, [37] builds a practical system that uses AOA measurements at LF to steer the beam in HF.

Only a few works have studied in detail the correlation between LF and HF bands, and their main conclusion is that the correlation strongly depends on the LOS versus NLOS. Consequently, in this work we perform a detailed exploration of the correlation of LF-HF channel metrics beyond AOA that may enable LF to assist HF and reduce overhead. We do so via extensive simulations with a raytracer, and by validating these simulations through measurements with a LF-HF system.

Our specific contributions are:

- We provide a thorough investigation of the correlation between several important bands LF and HF considered for 5G systems. For this we analyze different scenarios and do extensive raytracing simulations (Section 4.3).

- We consider the impact of hardware limitations. Specifically, we analyze the impact of the number of antennas on the perceived correlation (Section 4.3).

- We validate the correlation study by taking measurements in the proposed scenarios. We use two important frequencies for 5G systems, 3.5 and 61 GHz (Section 4.4).

- As an example, we show how this knowledge can be applied to optimize the time consuming beam steering process needed at HF, by beam training using LF information (Section 4.5).

4.2. Related work

In related work, researchers have carried out a large number of practical measurements campaigns to capture the differences among LF and HF channels. Such extensive campaigns are then used as a basis for channel models that cover a large range of frequencies, ranging from 6 to 100 GHz [32]. The fundamental difference among both bands is the much sparser multi-path environment at HF [38], which results in different cluster and sub-path characteristics. Diffuse scattering and diffraction also play a much more important role [32]. Besides, the delay spread becomes smaller since fewer multi-path components reach the receiver [39, 40].

Delving into the details, effects that are commonly disregarded at microwave frequencies, such as rain, can jeopardize communications at HF, where wavelength is comparable to a rain drop size. This leads to important attenuations for frequencies higher than 10 GHz, increasing as a function of frequency and rain rate. Results show

attenuations as high as 26 dB at 38 GHz under heavy rain, and even higher attenuations at 73 GHz [41, 42]. Furthermore, high penetration losses are an intrinsic characteristic of scaling up the frequency. Typical indoor office scenarios show penetration losses between 0.8 and 9.9 dB/cm, varying with the type of material (glass, closet doors, steel doors, and whiteboard walls) [43]. Moreover, HF is also sensitive to human body blockage, causing attenuations in the order of 30 dB on average [44]. Outdoor scenarios are difficult as well, as materials are generally thicker and more robust, and attenuations can increase as high as 45dB [45]. These results also hint at the high challenges of outdoor to indoor scenarios.

4.3. Simulations

In this section we present the study of the correlation among the selected LF and HF frequency bands. We first present the simulation setup and the raytracer characteristics and discuss the chosen scenarios and the metrics to assess the correlation. We then illustrate the simulation results and finally we evaluate the effects on the correlation produced by hardware complexity for various number of antennas.

4.3.1. Simulation setup

One of the pitfalls of simulation based studies can be a too high level of abstraction that the simulator provides. To mitigate this issue as much as possible, our raytracer takes into account many features of the environment such as the electromagnetic properties of the different materials in the scenario, material attenuation, the different multipath components, constructive and destructive interference at the signal level, realistic antenna models, diffraction, and so on.

4.3.2. Scenarios

Fig. 4.1 shows the two different scenarios that we have modeled in our raytracer, that are part our research facilities. This allows us to compare the simulations with real measurement data. For both scenarios, for brevity we use discuss results for one specific Access Point (AP) placement, located in the upper part in both cases, as indicated in the figure. The AP is able to transmit at 1.8, 3.5, 28 and 61 GHz. Our measurements are done at a height of 1.65 m. The gray areas on the map show the furniture that is below 1.65 m, and therefore does not block LOS.

The first scenario, Fig. 4.1 a), is a laboratory. The walls on the top and on the right are made of wood, and the left and bottom wall are concrete. The left wall has also 5 windows. Due to its small dimensions, 6 m wide, 8 m long and 3 m high, the laboratory scenario is comparatively rich in multipath even at HF. The multipath components have different characteristics depending on the material they reflect on.

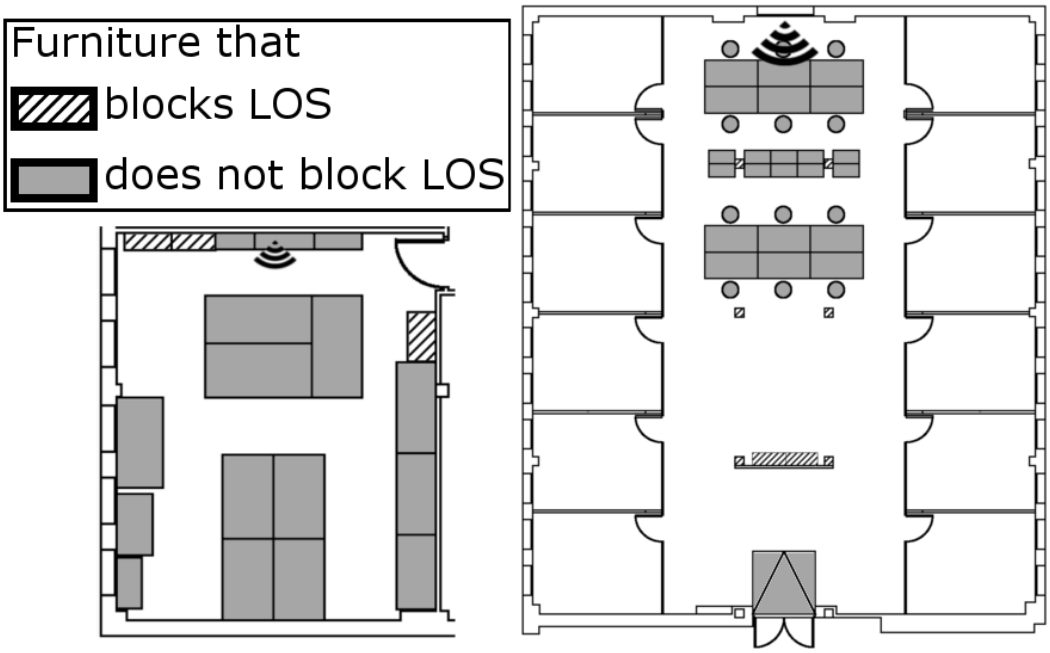

a) Laboratory b) Office area

Figure 4.1: Evaluation Scenarios

The second scenario, Fig. 4.1 b), is an office area which is 15 m wide, 19 m long and 3 m high. 12 offices surround the central part of the office area with glass walls, and the division between the offices is a thin wall of wood. This scenario has also some NLOS areas due to 4 columns in the main area and a dividing wall in the middle near the bottom.

4.3.3. Metrics

In both scenarios, simulations have been taken at measurement points that are spaced 20 cm apart. We apply a power threshold of 100 dB (receiver sensitivity) and a dynamic range threshold of 20 dB (i.e., we allow only paths that within 20 dB of the strongest path) and remove paths which do not meet those constraints. This reflects typical hardware constraints.

- **Number of equal paths**: Represents the number of paths that are geometrically identical in LF and HF. The higher the number of equal paths, the better the correlation.

- **Fraction of equal paths**: Represents the fraction of LF paths that are also found at HF, e.g., if only 5 LF paths arrive to a point but those same 5 HF paths arrive as well, we obtain a value of 1.

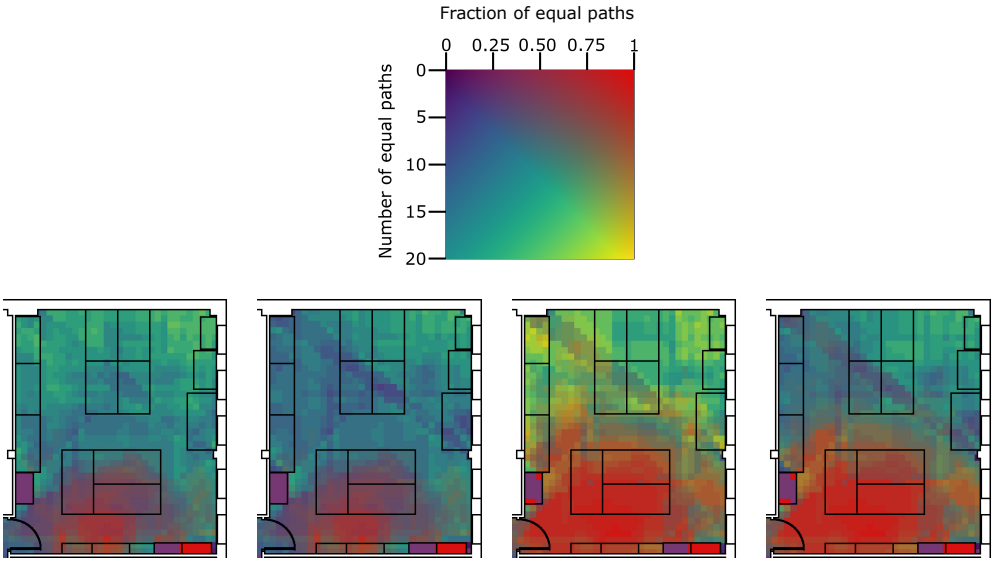

(a) 1.8 GHz and 28 GHz(b) 1.8 GHz and 61 GHz(c) 3.5 GHz and 28 GHz(d) 3.5 GHz and 61 GHz

Figure 4.2: Path study for the laboratory scenario

- **Delay Spread (DS)**: We use the definition established in [46], where the path delays are weighted by the power of each path.

- **Angle Spread (AS)**: The angle spread is given by the maximum angle difference between the direct path and the other paths in 3D. This gives a sense of directionality: the lower the spread, the more similar is the cluster of paths from the AP to that point.

- **Angle profile correlation**: The angle profile is a vector that represents the power that arrives to a point for each arrival direction. We then correlate the angle profile of LF with the one of HF using the Pearson correlation coefficient.

4.3.4. Correlation of the actual channels

Fig. 4.2 shows the number of equal paths as well as the fraction of equal paths for the laboratory scenario. This scenario has a high number of equal paths, with more than 6 paths on average for each case. Near the AP there is a low number of equal paths, but the fraction is very high, i.e., those paths are typically available both at LF and HF. Due to the proximity to the AP, the direct path is very powerful and therefore most of the reflected paths are discarded due to the 20 dB dynamic threshold. Common to all cases is also a diagonal stripe from the middle of the left wall to the bottom right corner. This stripe is caused by the column on the left wall, which blocks some of the reflections on that wall.

Regarding frequency behavior, we observe that 1.8 GHz behaves quite similar

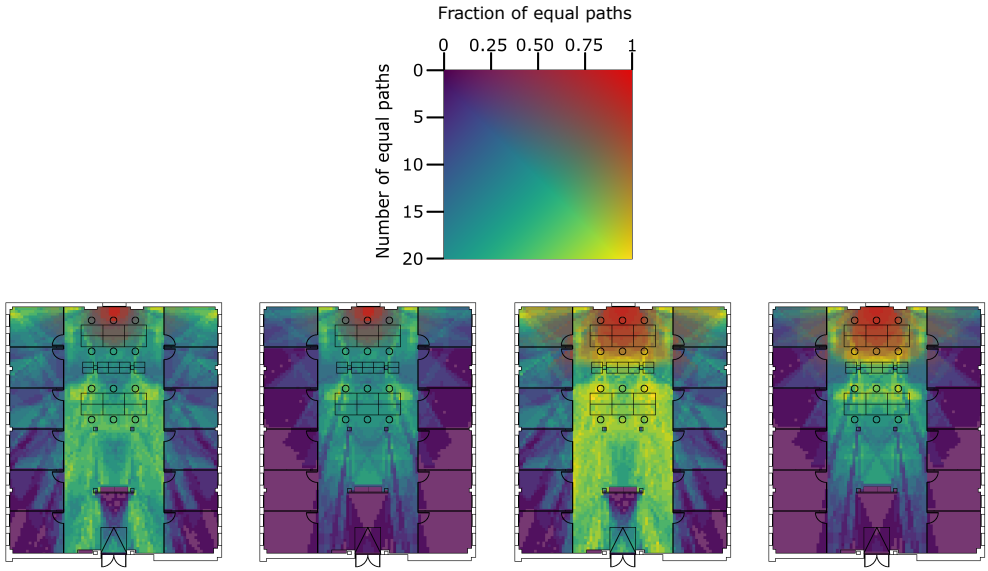

(a) 1.8 GHz and 28 GHz(b) 1.8 GHz and 61 GHz(c) 3.5 GHz and 28 GHz(d) 3.5 GHz and 61 GHz

Figure 4.3: Path study for the office area scenario

compared to the two HF frequencies. The only noticeable difference is again the reflected stripe from the left column, which is higher attenuated at 61 GHz. As expected 3.5 GHz behaves even more similar compared to HF, especially for 28 GHz. In the latter case, the fraction of equal paths is close to 1 in most cases. Besides, the area near the AP that has a low number of equal paths is bigger for 3.5 GHz than for 1.8 GHz. This is because 3.5 GHz has slightly fewer paths around that area compared to 1.8 GHz due to the higher attenuation and the dynamic threshold, and this number of paths is closer to that of HF.

Fig. 4.3 illustrates the results of the path study for the office area scenario. This scenario is heavily influenced by NLOS, and this becomes clear from the difference between 28 GHz and 61 GHz. Although both of HF frequencies are significantly attenuated, the penetration of 28 GHz through the glass walls of the offices is much better than that of 61 GHz. 28 GHz drops from 15-20 equal paths to 5-7, whereas 61 GHz looses coverage in most of the cases. In the middle part of the scenario, the actual office area, the influence of the columns is noticeable and at 61 GHz there is a visibly lower correlation behind the them. As in the laboratory scenario, the highest path fraction correlation is between 3.5 GHz and 28 GHz, this is again due to a similar attenuation between those frequencies.

Fig. 4.4 shows further of the correlation metrics for the office area scenario. We now focus on this scenario for further study due to its more complex and interesting geometry. Results in Fig. 4.4a suggest that, as expected, delay spread reduces as we scale up in frequency, due to higher penetration losses and reduction of multipath components due to attenuation. Comparing the two extreme cases, 1.8 GHz and 61 GHz, only 20% of the positions present a delay spread lower than 5 ns for the former, whereas this same

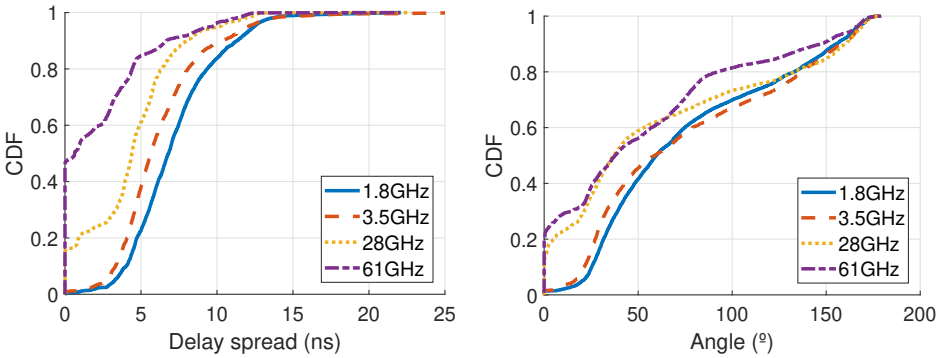

(a) Delay spread (b) Angle spread

Figure 4.4: CDFs for the office scenario

delay value represents more than 80% of the points for the latter. In terms of similarity between delay spreads, once again, 3.5GHz and 28GHz show the closest behavior. Lastly, our results show that although number of paths, power per path and maximum delay spread are always higher for all measured points for LF, somewhat counterintuitively the winner metric delay spread can be higher for HF for almost 10% of the measured points. If the cluster of paths of LF close to the mean delay value is not present in HF, and only more distant paths are present with relatively close power values, delay spread can be higher after power normalization for HF.

Fig. 4.4b illustrates the CDF for the angle spread. On average, both LF frequencies have a higher angle spread than HF due to much stronger higher-order reflections. 61 GHz has the lowest angle spread since it is the most attenuated frequency and, consequently, most reflections are lost. Also we can see that for LF the angles are evenly distributed, whereas for HF the points that are further away only receive the direct path.

The office area scenario can be divided into three main areas where DS and AS behave similarly. The LOS part in the middle area, usually presents the higher values of DS and AS, this holds true for all the frequencies. In the middle area, the values of DS quickly drop when there is a blockage, such as the pillars or the wall in the bottom half part, but since the scenario still holds its geometry on those points the AS is still high. The last area conforms the inside of the offices surrounding the middle area. This area has a higher values of DS similar to those in the Obstructed Line Of Sight (OLOS) sections of the middle part, but the AS quickly lowers as the offices are further apart from the AP.

4.3.5. Correlation of the observed channels

The previous simulation study investigated the actual correlation of the LF and HF channels. However, from a practical point of view, it is crucial to study the channels as they are perceived through the respective antennas. In particular, the number of antennas plays a vital role to differentiate paths in the spatial domain, as with N antennas we can

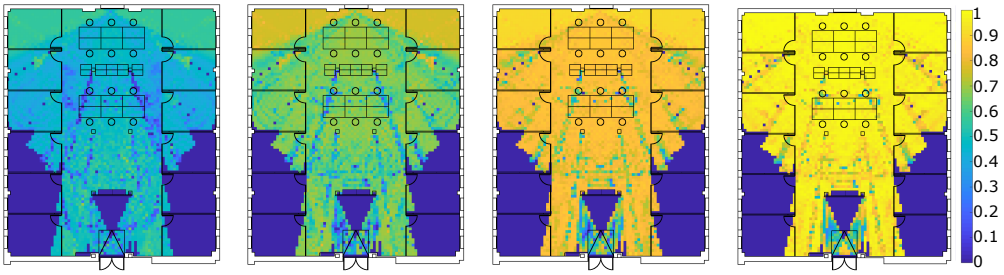

(a) 2 antennas at LF (b) 4 antennas at LF (c) 8 antennas at LF (d) 16 antennas at LF

Figure 4.5: Antenna profile correlation for 3.5 GHz and 61 GHz

theoretically differentiate up to $N-1$ paths. In addition, HF systems typically have a larger number of antennas compared to LF, as the higher frequency allows for smaller antenna elements and thus a higher number of them can be integrated in the same area. Therefore, the correlation is likely to improve when the number of antennas at LF is closer to the one at HF (for example for massive Multiple-Input and Multiple-Output (MIMO) LF systems that incorporate a large number of antennas at the base station). For this reason, this section studies the minimum number of antennas at LF required to observe a meaningful correlation. We consider 16 antennas in HF (for both transmitter and receiver) and 2, 4, 8 and 16 antennas for LF at the receiver side. For simplicity we use an omnidirectional antenna for the transmitter.

We use the angle profile, i.e. the power associated with each angle, as a metric for the correlation to analyze not only the behavior of the reflection but also the losses. To estimate the angle profile, the channel at reception, h, given by the raytracer, is multiplied by the steering vector. We define the steering vector as

$$\mathbf{s}(\theta) = [1\ e^{-j\pi\sin(\theta)}\ ....\ e^{-j\pi(N-1)\sin(\theta)}], \tag{4.1}$$

where N is the number of antenna elements. Further, we define the channel as

$$\mathbf{h} = [h_o\ h_1\ ....\ h_{(N-1)}], \tag{4.2}$$

where h_n is the channel coefficient for antenna n. The angle profile (AP) is given by

$$AP = |\mathbf{s}(\theta)^H\mathbf{h}| \quad -\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}. \tag{4.3}$$

The angle profiles are estimated for each measurement point in the open area for the frequencies of 3.5 GHz and 61 GHz and the results are illustrated in Fig. 4.5. For the case of 2 antennas at LF, the correlation is around 40%, dropping to 20% in the OLOS parts behind the pillars. The correlation quickly improves for 4 antennas at LF, averaging a 65% of correlation for the whole scenario but also dropping to 30% for OLOS. With

8 antennas, half of the antennas used for HF, the correlation goes up to 85% on LOS and around 50% on OLOS. When using 16 antennas at both LF and HF, the correlation over the whole scenario is around 90%. With this number of antennas, even in the OLOS part of the scenario the correlation is 80%, and even in NLOS, bottom middle part, the correlation goes up to 50%.

Fig. 4.6 shows the CDFs of the angle profile correlation for all the frequencies for the office area scenario. These CDFs show that every time we double the number of antennas, we improve the correlation by around 20%, and we can see that this holds true independent of the frequencies involved. Furthermore, we find that the correlation curves are very similar for a given HF frequency, i.e., there is no substantial difference whether we use 1.8 or 3.5 GHz in terms of angle profile correlation. While the multi-path channel at 1.8 GHz and 3.5 GHz differs, the actual angle profiles for a given number of antenna elements are in fact similar, resulting in very similar angle profile correlation curves. There are significant differences between 28 GHz and 61 GHz, primarily due to the lower coverage at 61 GHz.

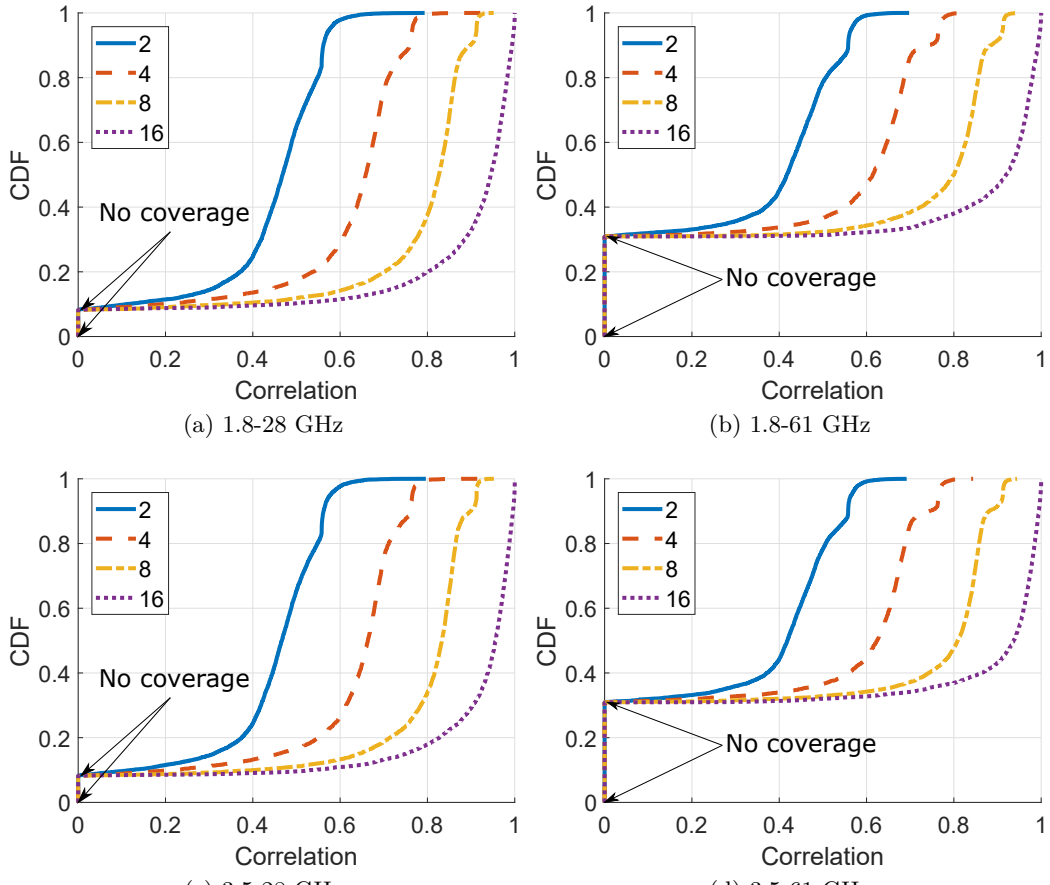

(a) 1.8-28 GHz (b) 1.8-61 GHz

(c) 3.5-28 GHz (d) 3.5-61 GHz

Figure 4.6: CDFs for the antenna profile correlation for different numbers of LF antenna elements and 16 HF antennas

4.4. Practical validation

The aim of this section is to validate the LF and HF correlation through real-world measurements in the scenarios mentioned above. For brevity, we focus on the angle profile correlation metric for this section, because it encompasses the most important information. The equipment and the procedure used for the measurements are the following:

- LF: Our set up consists of a BladeRF SDR as transmitter with an omnidirectional antenna, and a USRP X310 equipped with TwinRX daughterboards connected to a 4 element antenna array, which is in line with current mobile systems. This setup ensures phase synchronization among antennas elements. We use the MUSIC pseudo-spectrum [47] to estimate the angle profile.

- HF: We transmit a pulse signal up-converted to the HF band with an up converter (SiversIMA) which operates at 61.48GHz. The signal is transmitted by an omnidirectional antenna. On the receiver side, a horn antenna with an aperture of 7 degrees is used. Likewise, the horn antenna is connected to a SiversIMA operating at 61.48 GHz. A motor rotates the set up by 360 degrees in steps of 4 degrees, resulting in the different steering directions. We build the angle profile by computing the received power at each step.

The measured angle profiles are illustrated in Fig. 4.7 and they were taken in the office scenario for the frequencies 3.5 and 61 GHz. Values in the graph for HF range from -50 dBm to -70 dBm , with each concentric semicircle representing a step of -5 dBm, whereas for LF they range from 0dBm to -20 dBm.

Due to the reduced number of antennas at LF, the angle profile is not fine-grained, and the system can not differentiate many paths. However, values for correlation between LF-HF angle profiles are generally high for points in LOS, with values between 65 and 91%. For the bottom row, and in particular behind the dividing wall, the low receive power makes it difficult to clearly identify the main path.

The third row from the top presents peculiar behavior due to a pillar obstructing LOS. The strongest path for LF is pointing to the wall, a reflection, whereas HF's strongest path is obstructed LOS; this yields the lowest correlation values, around 15%. We observe similar behavior inside the offices, where received power is greatly attenuated. Inside the offices, LF generally finds reflections inside the room as the strongest path, and HF points to the AP but with low power. In these cases, correlation drops to values between 15% and 35%. Finally, we also took measurements with office doors closed, finding very low correlation and low power values for HF, with some coverage available only in the first two offices closest to the AP.

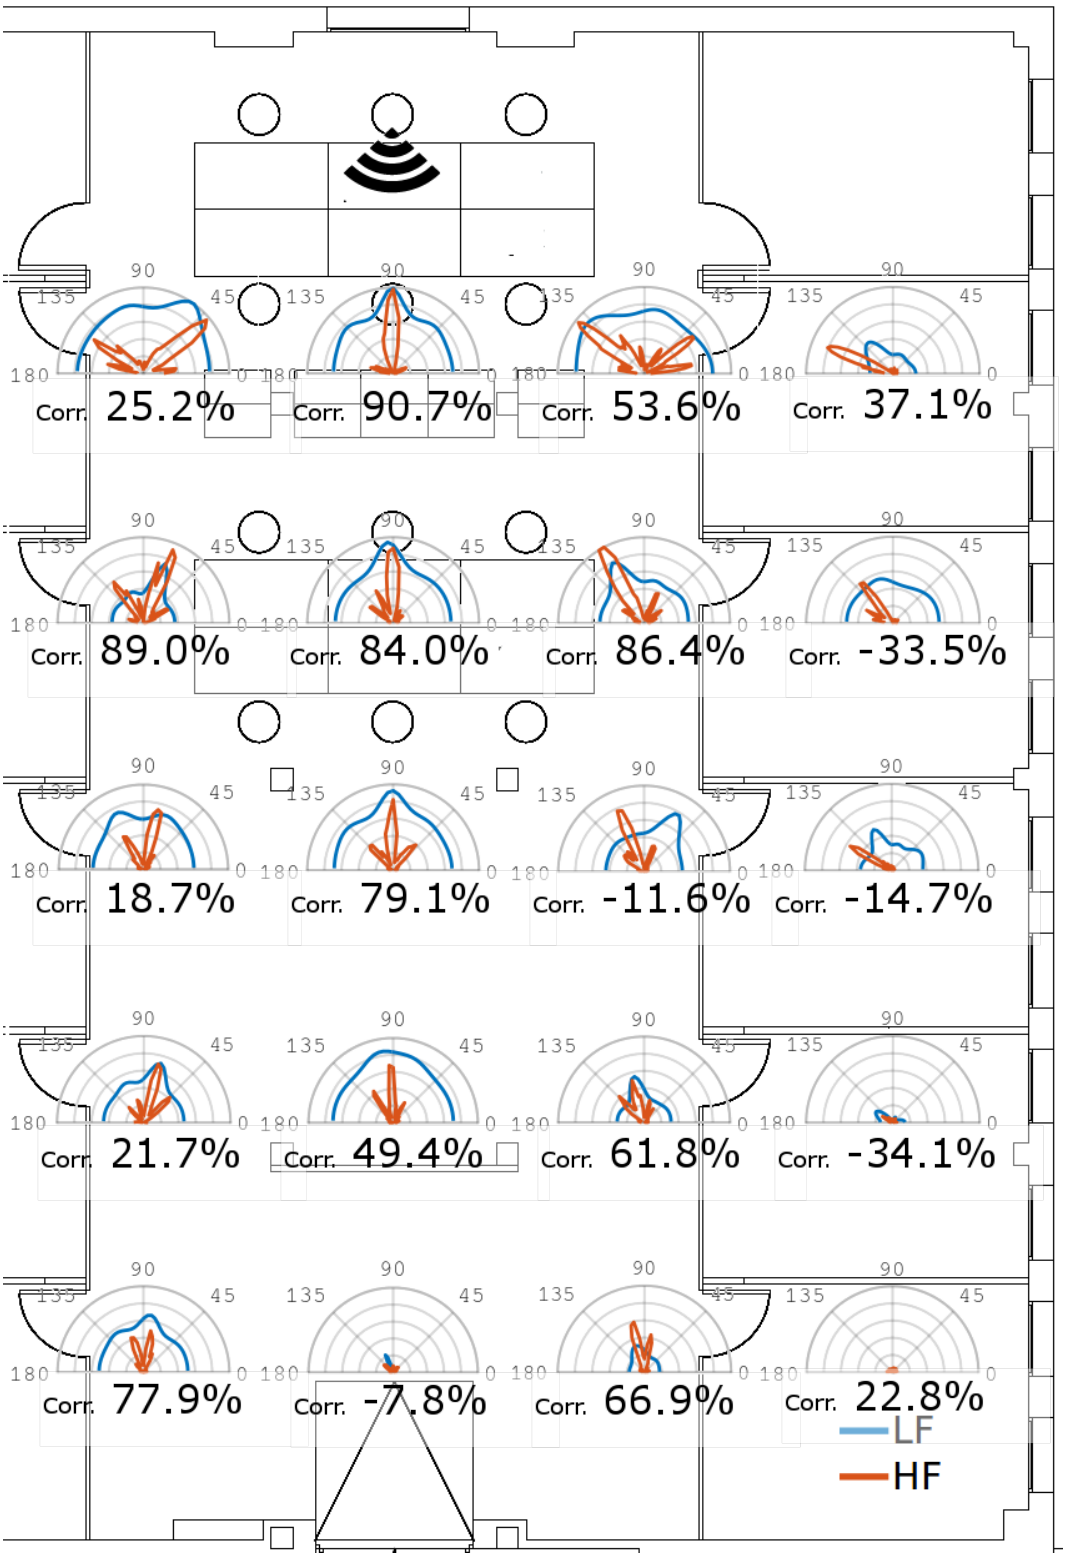

Figure 4.7: Angle profiles for LF and HF in the open area

4.5. LF assisted beam-steering

In this section we investigate how a multiband system can benefit from correlation to reduce overhead. The conventional algorithm for beam training, brute force training, probes every antenna sector and selects the one which provides the highest power. Such an approach is time consuming: it takes 20μs to check each sector and changing beam patterns takes around 1μs. To reduce the beam training overhead, we propose a system where the sector selection is performed by the angle profile information from LF. We check the sectors at HF according to the most powerful sectors of LF. For our results we take a fixed antenna width of 5°, 10° or 20°. We assess the system comparing the SNR loss of this approach to the optimal brute forcing as a function of number of sectors checked, i.e., a function of time devoted to beam training. We evaluate the system using both simulations and practical measurements. We use 4 antennas at LF and the pair of frequencies chosen are 3.5 GHz and 61 GHz.

The results are illustrated in the error bar plot in Fig. 4.8. We illustrate the losses for the 20 measurement points used in the practical measurements. The middle point of the bar is the mean of the losses, the upper whiskers represents the 90% quantile and the lower whiskers the lowest value. Although for both cases, zero losses are the most frequent value, mean losses for practical measurements are around two times higher compared to the simulations, due to hardware imperfections and other real-world effects.

Regarding differences for the different sector widths, the losses follow a similar pattern as we use wider sectors. The losses in simulations for 20° HF sector width are 0.5 dB when checking only 2 sectors around the best angle found at LF, which corresponds to 40°. With 10° and 5° sectors, the losses are 0.5 dB checking 3 and 11 sectors, corresponding to 30° and 55° respectively. Contrarily to simulations, zero losses in the practical measurements are only reached when checking all the sectors, as there always exist very rare cases where the next direction is the last one to be checked. Overall, these results show that directly using the angle given by LF already gives a very good steering direction for HF. Further, there is a considerable reduction in SNR loss when checking only one or two additional sectors, meaning that the highest power values of HF are indeed generally within the two-three most powerful sectors at LF. As a consequence, we can reduce the overhead of the beamsteering process to one third of its original time with an average loss of only 2 dB.

4.6. Conclusions

In this work we study the correlation of different frequency bands. We show that cooperation between high frequency and low frequency in multiband systems can boost system performance. Our simulation results with a raytracer for an ideal scenario show

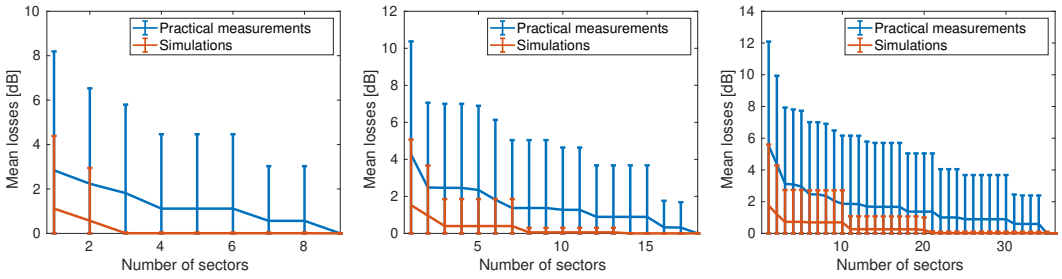

(a) 20° of aperture (b) 10° of aperture (c) 5° of aperture

Figure 4.8: Beam steering losses comparing to brute forcing

that the pair of studied frequencies with the most similar behavior is 3.5 and 28GHz. We also investigate the effect of the number of antenna elements on the perceived correlation, showing that having 4 antennas at LF and 16 at HF already suffices to exploit correlation, and scaling up the number of LF antennas to 16 gives almost full correlation. Furthermore, we carry out hardware measurements at 3.5 and 61GHz to validate the simulations. We find that for a limited number of antennas at LF, 4, angle profiles have a high correlation between 70% and 90% for LOS, whereas for NLOS, the correlation drops to values between 15% and 30%. That means that even with a small number of antennas, the strongest paths at LF correspond to the strongest paths at HF for most of the cases.

Finally, we include an application of correlation by computing what is the loss in power if the angle profile information of LF is directly applied to steer the antenna beam at HF, one of the most resource intensive procedures for mmWave systems. Results for simulations and practical measurements show we using a minimal set of sectors around the angles predicted by LF result in very low SNR loss, thus reducing the beamsteering overhead to one third of its original time with an average loss of less than 3dB.

"Science isn't about why, it's about why not!"

Cave Johnson

5 Beam training overhead reduction

5.1. Introduction

A well known challenge of Millimeter-wave (mmWave) is how to keep the beams between two devices aligned once the connection has been established. This is normally done in the Sector Level Sweep (SLS) phase during the Beam Forming Training (BFT). Fig. 5.1 shows the SLS in IEEE 802.11ad. This procedure has to be done twice, once where the Access Point (AP) is the initiator and the Station (STA) the responder, and a second time where the STA is the initiator and the AP the responder.

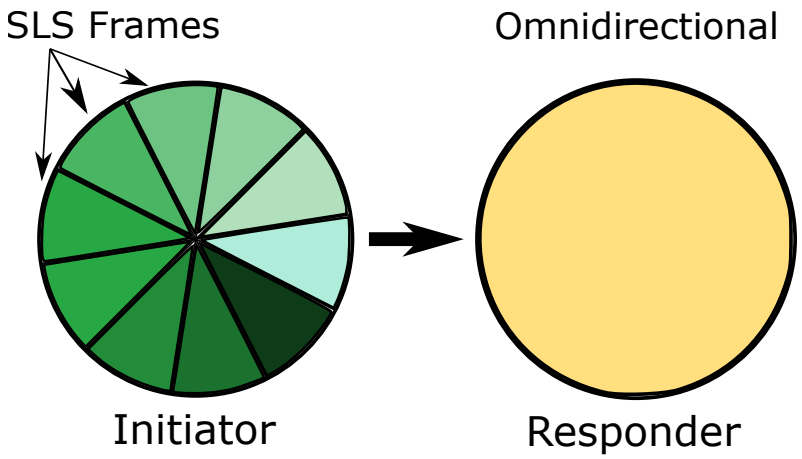

Figure 5.1: SLS in IEEE 802.11ad

In order to limit this impact, there has been some focus on exploiting the IEEE 802.11ad headers to make the Beamforming (BF) mechanism faster [2, 11]. In order to reduce the BF overhead we developed an ultra fast beam alignment in mm-FLEX [2], an open platform based on Field Programmable Gate Arrays (FPGA) which has a baseband processor with full-duplex capabilities, including a set of phased antenna arrays configurable on-the-fly. With this setup we finally have a platform where we can test modifications to the standard. With the real-time antenna reconfiguration we introduce

Ultra Fast Alignment (UFA) which eliminates the training on the station side completely. It does so by testing different receive beam patterns while receiving the preamble of a training packet during AP beam training.

As explained in Sec. 2.2, a IEEE 802.11ad frame has a preamble that includes a Short Training Field (STF) field. This field is used for symbol synchronization, frame detection and coarse synchronization. The STF consists of 48 repetitions of Golay sequences of 128 bits of length, followed by 2 extra Golay sequences to detect the end of the STF [13].

5.2. Main idea

UFA aims to cut the BFT overhead by eliminating SLS on the STA, leaving it only at the AP side. This is achieved by training the stations while they are receiving the AP transmit sector training packets. The main idea of our approach is that when you change beam receiving patterns while the AP is transmitting the STF of a certain frame, the change on the amplitude of the signal is proportional to that of the receive beam patterns. Fig. 5.2 shows the signal changes due to the receive beam pattern changes.

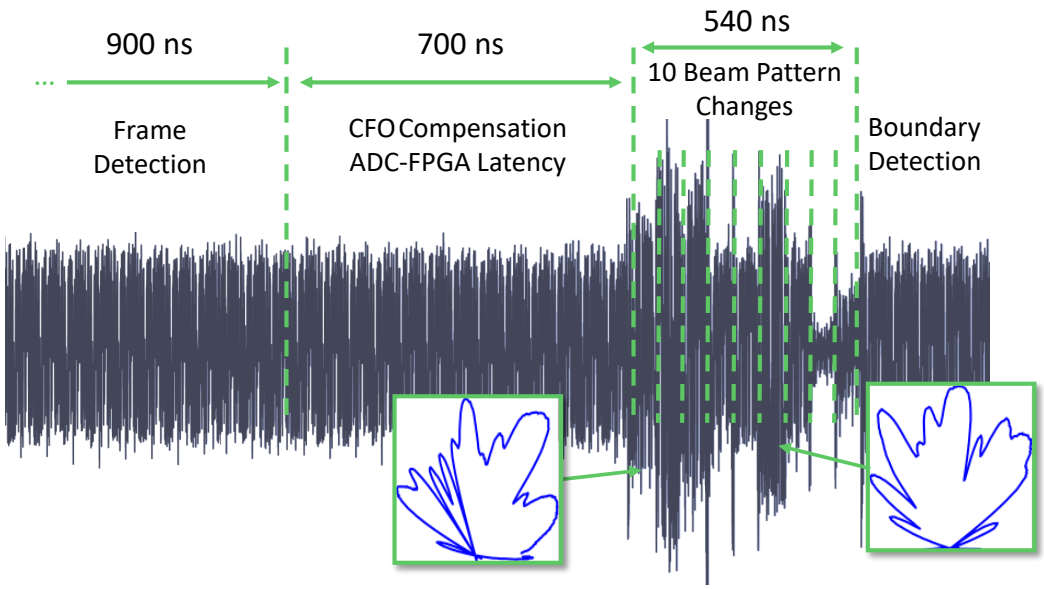

Figure 5.2: Measured STF of a control frame, showing signal changes due to beam switching at the receiver

This changes in amplitude can then be extracted from simple Received Signal Strength (RSS) measurements, which allows accurate results without increasing the complexity on the receiver. Then it is possible to estimate the Angle of Arrival (AOA) with high accuracy due to the sparsity of the mmWave channel.

5.3. Evaluation

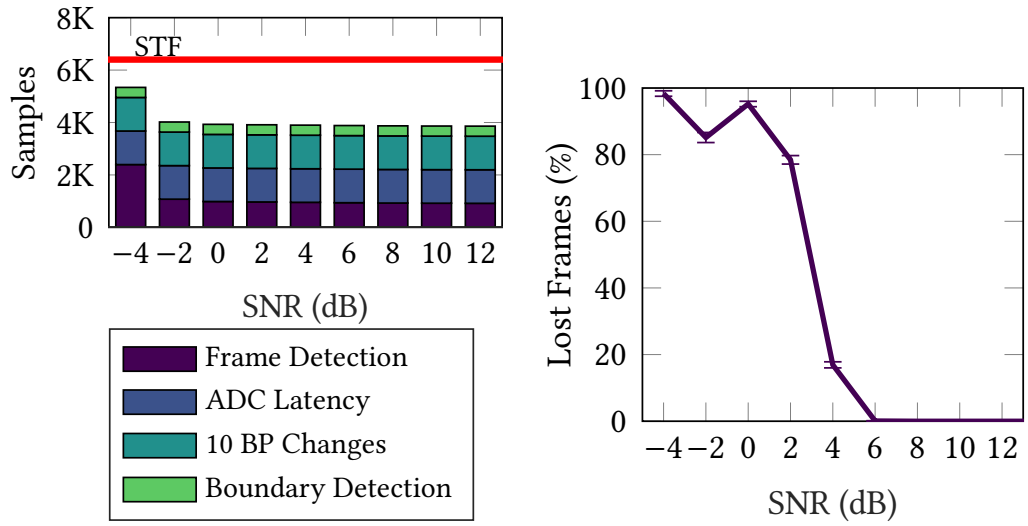

(a) Preamble processing time

(b) Fraction of undecodable SLS frames

Figure 5.3: Accuracy of UFA compared to exhaustive search

We implement the UFA algorithm in the mm-Flex platform and do a thorough evaluation. Inevitably, UFA introduces some latency in frame detection, Fig. 5.3a shows that even for low Signal-to-Noise Ratio (SNR), UFA is able to finish before the STF. That means that we can switch 10 receiving beam patterns before the STF finishes. We also do some testing to check if our approach is affected by the SNR. Fig. 5.3b shows that UFA works for all relevant SNR values. Furthermore, we do an evaluation of the AOA error and the SNR loss with respect to the exhaustive search depending on the number of beam patterns used. Fig. 5.4a shows the AOA error. With only 5 beam patterns we get a median error of 0.5°, with an error of less than 4° for 10 beam patterns for 95% of the cases. Fig. 5.4b shows the SNR loss depending on the number of beam patterns, the mean quickly goes to zero for more than 3 receive beam patterns, and we get no outliers with 5 or more beam patterns.

5.4. Overhead reduction

We carry an extensive measurement campaign with TALON AD7200 devices, which are Commercial-Off-The-Shelf (COTS) devices that implement the IEEE 802.11ad standard. Fig. 5.5b shows the overhead reduction with respect to the IEEE 802.11ad standard. It shows the measured and extrapolated beam training overhead. With 1 AP and 1 station, their overhead is approximately the same and thus IEEE 802.11ad has twice the overhead of UFA. For 4 stations, the extrapolated station overhead is 4 times

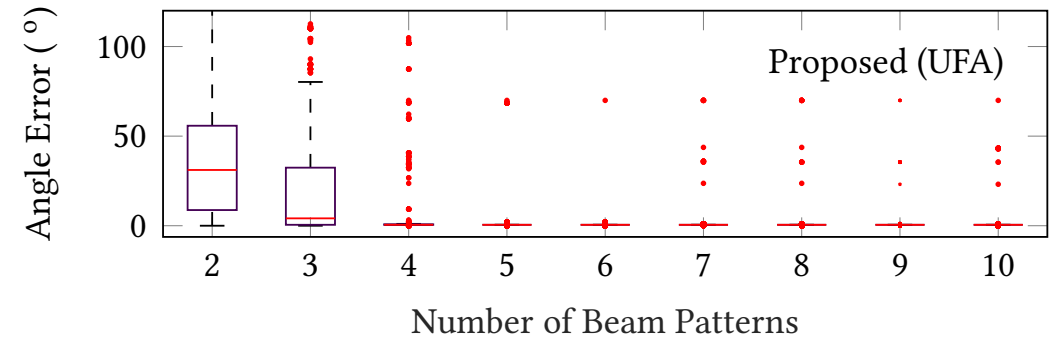

(a) Angle error for different reception beam patterns

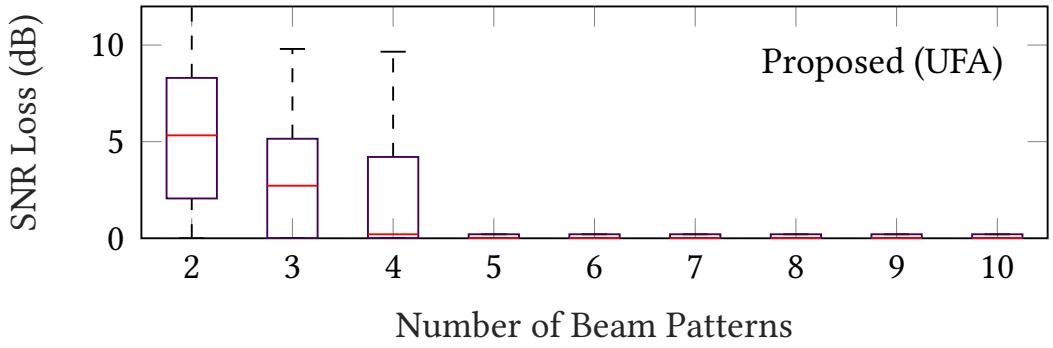

(b) SNR-loss of UFA compared to exhaustive search

Figure 5.4: Accuracy of UFA compared to exhaustive search

larger and thus IEEE 802.11ad overhead is 5.1× that of UFA (left bar). The Talon routers actually both slightly increase the DMG Beacon frequency and significantly increase the station beam training frequency so that measured IEEE 802.11ad overhead is higher by a factor of 5.9. For 7 stations, the beam training frequency of stations is reduced due to the increase in data traffic with that many stations. Here, IEEE 802.11ad overhead is around 8× that of UFA, and measured and extrapolated overhead almost match.

5.5. Conclusions

Most of the challenges of mmWave come from the high spatial attenuation and the need for directional communications. The focus of this chapter was to reduce the overhead of the BFT mechanism. In order to do so we develop an open, modular, reconfigurable testbed suitable for mmWave experiments. This platform allows us to implement and test any kind of algorithm and run it in real time. We implemented UFA, and algorithm that switches beam patterns on the STA during the packet preamble. This, in turn, reduces the overhead of the BFT since the SLS is only necessary in the AP side. We show that we can change the beam pattern up to 10 times during the STF field and that even with 5

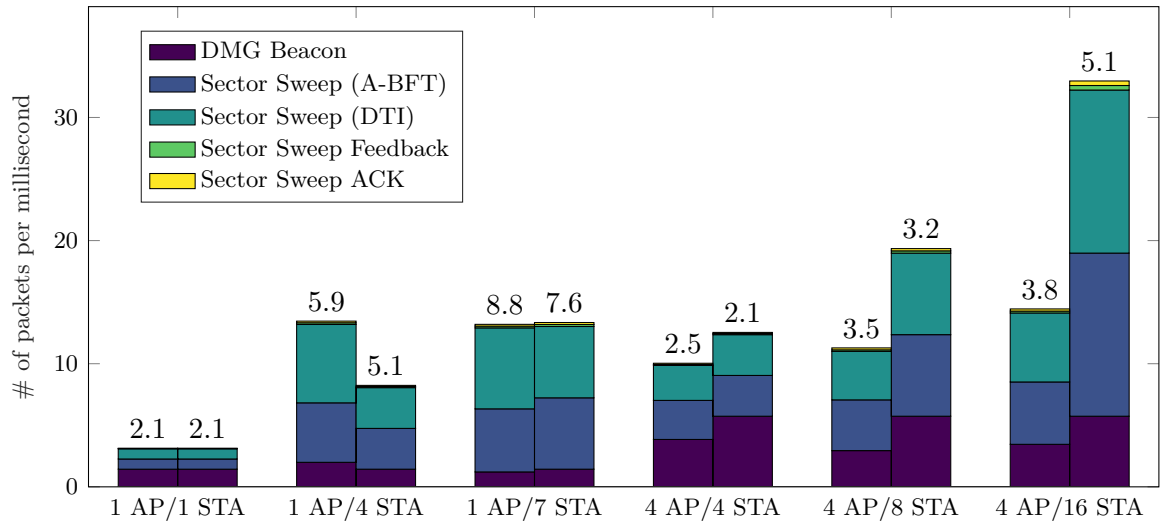

(a) Measured (left bar) and extrapolated (right bar) overhead of IEEE 802.11ad and mm-FLEX, with the overhead ratio of IEEE 802.11ad divided by mm-FLEX shown on top of the bars

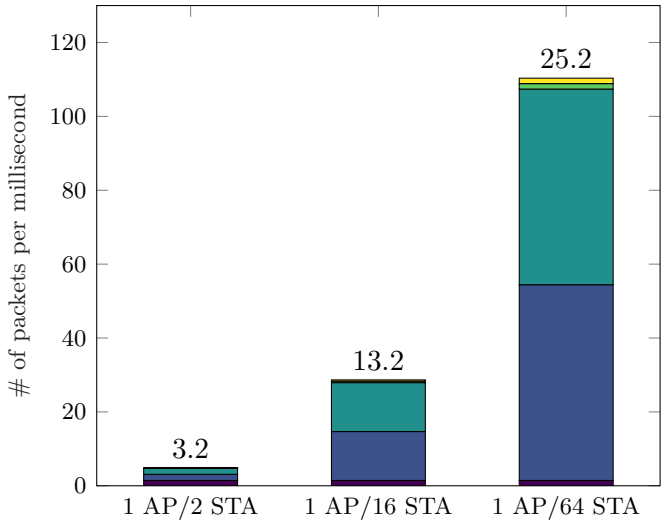

(b) Extrapolated overhead

Figure 5.5: Beam training overhead and overhead of IEEE 802.11ad compared to mm-FLEX (on top of the bars)

beam pattern switches we get less than 0.5° of error and no outliers. For dense scenarios with many STA we show that UFA reduces the overhead by a factor of 8.8 with respect to the IEEE 802.11ad standard. All the implementation of UFA is done in the STA side and therefore it is compatible with standard compliant APs.

"A great thing about doing a Ph.D. is that everyday is Saturday!
But you also work on Satudays"

Ph.D. comics

6 Augmenting mmWave Localization Accuracy Through Sub-6 GHz on Off-the-Shelf Devices

6.1. Introduction

The large number of elements in directional mmWave phased antenna arrays enables very accurate estimation of the angle of the signal propagation, and the high temporal resolution from the multi-GHz bandwidth provides excellent ranging accuracy. These characteristics make mmWave technology an ideal candidate to design highly accurate wireless location systems [48–50]. A location system estimates the Angle of Arrival (AOA) and Time of Flight (ToF) of the path between an Access Point (AP) and a client to determine the position of the client. In case this path is direct Line Of Sight (LOS), the angle and range information from AOA and ToF will lead to an accurate estimate of the position of the client. While obstructed LOS paths that pass through obstacles are often too weak to be detected, they may still provide accurate localization in some cases. However, if the path to the client is Non Line Of Sight (NLOS) via one or more reflections, both angle and range information correspond to the reflection, which is usually far from the actual client position. This can either occur when 1) only a reflected path exists, 2) the reflected path is much stronger than the LOS path since the phased array antenna is steered towards the reflection, 3) or the device orientation is such that the aperture of the antenna cannot capture the LOS path, e.g., when the client is directly pointing towards a wall. Thus, *mmWave position estimates are either very accurate, or have a high error.*

Unfortunately, the mmWave multi-path channel is typically sparse, such that only one or very few paths are available, and there is coverage from only one or very few APs. This makes it exceedingly difficult to distinguish between LOS and NLOS cases using mmWave information alone. However, when operating a sub-6 GHz location system [51–53] along with the mmWave location system, it is possible to distinguish reflections from (obstructed) LOS paths with the help of the lower frequency information. Sub-6 GHz antenna elements are omnidirectional and are thus much less affected by device

orientation, and the lower frequency leads to better coverage which results in many more visible APs. Such a lower frequency system can also provide localization whenever there is no mmWave coverage, albeit at lower accuracy due to the lower bandwidth and lower number of antennas for angle estimation.

There are few prior works study the joint use of of mmWave and sub-6 GHz bands. The most frequent use case is to reduce the mmWave beam-training overhead with the help of measurements from the sub-6 GHz band [54–57] to infer the optimal steering direction of the mmWave antenna array. The only system that combines mmWave and sub-6 GHz information for localization purposes "fuses" mmWave ToF-based ranging with sub-6 GHz AOA [58] but does not in fact provide an independent location system at either band.

As seen in Chapter 4 there exist a strong correlation between high and low frequencies. In this chapter, we propose a Multi-band Location System (MultiLoc) that combines the mmWave and sub-6 GHz bands to achieve high-accuracy device localization. Despite significant research interest in mmwave and localization, to the best of our knowledge no such multi-band location system has been proposed and implemented. MultiLoc uses the sub-6 GHz location estimate to decide whether to use mmWave information at all, and if yes which mmWave AP to use for the localization, whenever the client is covered by one or more mmWave APs. If there is no mmWave coverage or all mmWave paths are estimated to be NLOS, the sub-6 GHz location estimate is used directly. Otherwise, the location estimate from a *single* mmWave AP is used. Our system does not merge location estimates from multiple APs or frequencies since this degrades the location accuracy. In case the device integrates multiple mmWave antenna arrays with different orientations to improve coverage, MultiLoc uses the array that provides the shortest mmWave ToF estimate to a given AP, since LOS paths are characterized by a shorter length than NLOS paths.

In order to test our system at scale, we implement MultiLoc on mmWave Commercial-Off-The-Shelf (COTS) devices [59]. To this end, we port the OpenWRT embedded Linux to these devices and modify the mmWave device drivers in order to extract Channel State Information (CSI) information and enable Fine Timing Measurements (FTM) ranging. This is the first open platform that supports both mmWave CSI extraction and FTM, and we will release the corresponding code and tools to inspire future research on mmWave localization. We deploy MultiLoc in indoor and outdoor environments. Our results show that in a dense and a sparse indoor deployment with 4 or more mmWave antenna arrays per client, our system achieves a median location error of 0.13 m and 0.17 m, respectively. With only a single mmWave array, the median location error degrades to 0.65 m and 1.1 m, respectively, since in many cases, no suitable mmWave paths are available. The outdoor scenario is characterized by good LOS coverage and MultiLoc achieves a median location error of 0.17 m even with only two mmWave antenna arrays, and 0.55 m error

with a single array. Accurate location system such as this one can be used to reduce beam training overhead. By knowing the direction of the Station (STA) we can change to the beam pattern that has a higher gain towards that direction without the need of a full training.

6.2. Motivation

We examine the behavior of mmWave localization and its dependence on device orientation in more detail in an indoor scenario. We place a client at different measurement points directly in front of the AP within LOS, as shown in Fig. 6.1. The five measurement points are equally spaced (with step size 1.2 m) on a line originating at the AP, with the closest being at 4.2 m. Parallel to that line at a distance of 3.4 m there is a wall. On the AP, we measure both ToF and AOA to the client to estimate its position relative to the AP (we provide further details of the ToF and AOA estimation algorithms in Section 6.5). According to the manufacturer, the aperture of the AP and client antennas is $\pm 30°$, but we observe partial connectivity also beyond that angular range. However, due to the array geometry, angle estimation at such extreme angles is much less accurate than for angles close to the broadside direction of the antenna array.

If the client has a rotation of $0°$ with respect to the AP, i.e., it points directly at the AP, the location estimates are very accurate with errors of 0.12, 0.13, 0.08, 0.1, and 0.12 m, respectively, for the measurement points (we omit this example from the figures). For comparison, the state-of-the-art sub-6 GHz location system that we adopt for this work and based on 802.11ac obtains a median error of 0.3 m. Fig. 6.1a shows the same scenario when the clients is rotated $45°$ counterclockwise. The length and angle of each line represent the distance and the angle estimated for mmWave. For reflected paths, the actual estimated client location is given by following the solid line through the wall for the corresponding distance. The location errors are 0.12, 0.13, 5.7, 4.4 and 6.4 m. The first two points still provide the accurate localization despite the LOS path being slightly outside of the antenna aperture, since the AP does not emit a signal at an angle steep enough to reach the client via a reflection, due to its own limited aperture. For the last three points, however, the chosen path is a reflection off the wall towards which the client antenna is pointing, which significantly decreases the accuracy of the system. Fig. 6.1b shows a more extreme case where the antenna is rotated by $90°$, pointing directly at the wall. In this case, none of the antenna patterns of the array have any significant gain directly towards the AP, and for all five positions we observe reflected paths and location errors of 4.2, 6.2, 5.8, 5.6 and 5.4 m, respectively.

These examples highlight how unreliable mmWave localization can be when AP and client are not aligned. This problem is exacerbated for larger AP-client distances and more complex indoor environments. In contrast, the sub-6 GHz system is significantly

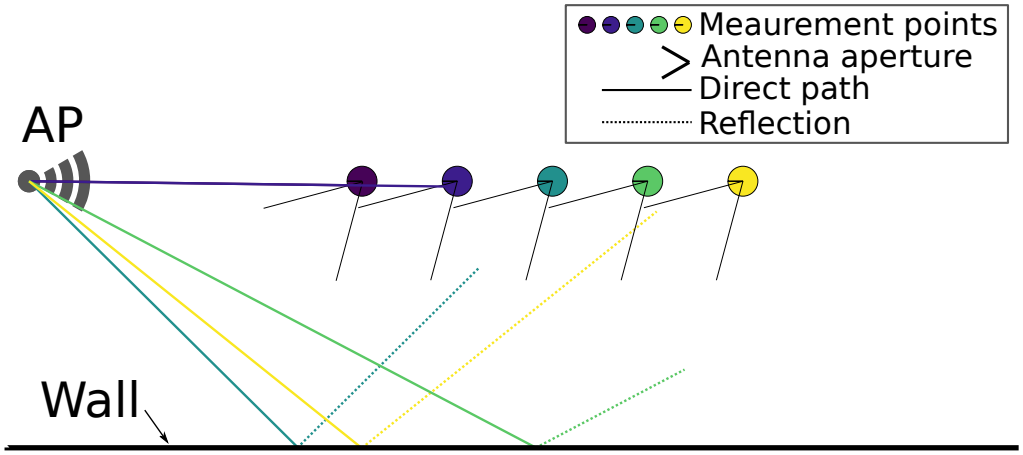

(a) Example with 45° misalignment

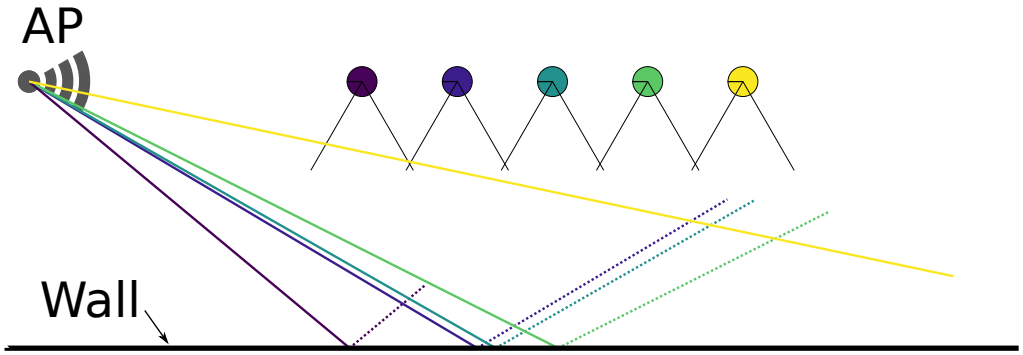

(b) Example with 90° misalignment

Figure 6.1: Estimated angles and path lengths when the client is not facing the AP

less accurate but at the same time less dependent on client orientation. This motivates the design of our multi-band location system which uses sub-6 GHz information to distinguish LOS and NLOS mmWave paths. Finally, since the mmWave location system is either very accurate or has a high error, it does not make sense to fuse the data from both bands.

6.3. Implementation

Before discussing the details of wireless localization and the design of our multi-band algorithm, we present the COTS hardware on which we implement our system and explain how we modify the device driver to access CSI and FTM.

6.3.1. mmWave system

For the experimental evaluation we use COTS devices, specifically the MikroTik wAP 60G and MikroTik wAP 60Gx3 shown in Fig. 6.2. Since the manufacturer installs a proprietary operating system and modified IEEE 802.11ad stack that provides only very

limited access, we ported OpenWRT to this platform and we replaced the stock firmware. This gave us the ability to customize the open source wil6210 Linux driver and port to this platform CSI extraction functionality and enable FTM measurements. We will release our custom fork of OpenWRT and the mmWave driver to the research community.

6.3.1.1. Hardware

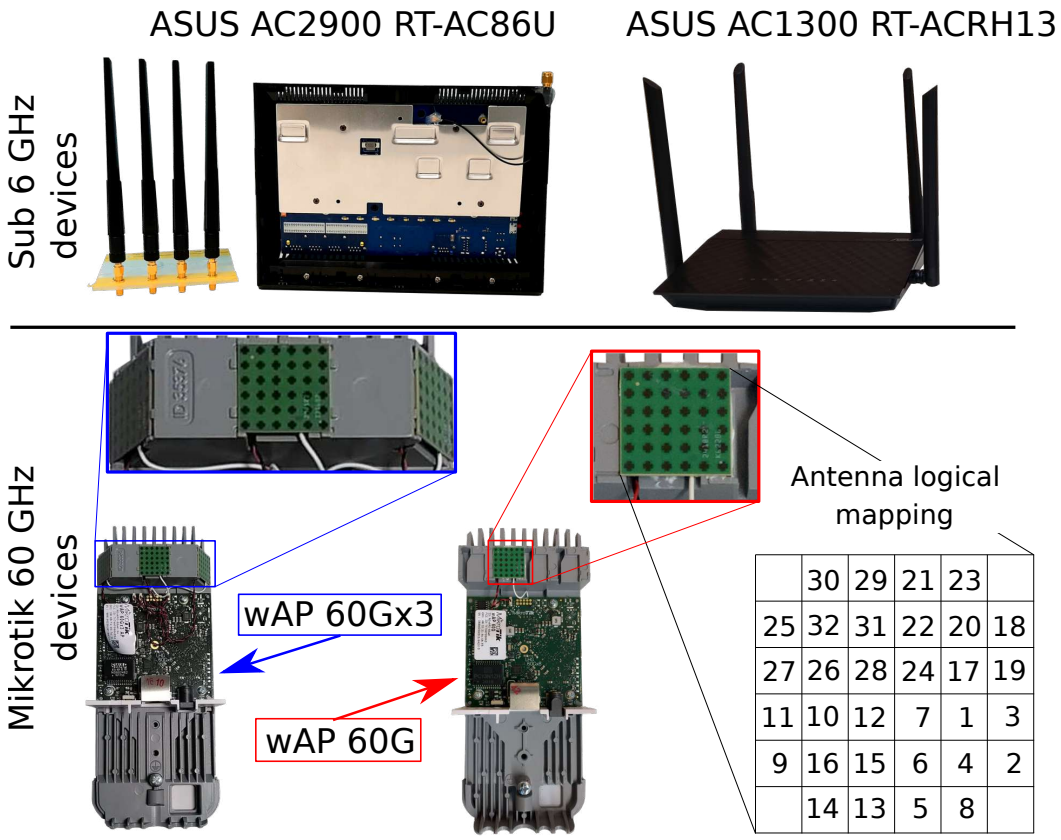

Figure 6.2: Sub-6 GHz and mmWave hardware

The MikroTik wAP 60G devices come with a Uniform Rectangular Array (URA) of 6 by 6 antenna elements and 60° aperture. The MikroTik wAP 60Gx3 uses three of such arrays arranged at different angles to provide a combined aperture of 180°. Unfortunately, as of now it is not possible to select which antenna to use and the CSI and FTM data are always gathered from the antenna in the middle, and we thus use these devices interchangeably. The same antenna has been partially reverse engineered [60] in order to have further control over it by using a custom FPGA implementation for the control path. Instead, our approach does not require additional hardware modifications which makes it inexpensive and easily accessible. The first challenge we faced was to map each

logical antenna element to its physical position on the antenna array, which is important to estimate the AOA. To this end, we establish a link between a pair of MikroTik devices while capturing CSI measurements and we successively cover each antenna element with tin foil. When covered, an antenna element becomes noisy, and we can find the logical element by observing the changes in CSI magnitude of the antenna elements. The mapping is shown in Fig. 6.2. The antenna spacing is 0.58 times the wavelength at 60.48 GHz [60]. This antenna spacing creates spatial ambiguity since it is more than half the wavelength, and we further discuss how to deal with it in section 6.5. Since Qualcomm uses 32 bits to define the Beam Patterns (BP) of their devices, only 32 antennas can be addressed, and for this reason the 4 antennas in the corners are unused. Old versions of the MikroTik wAP 60G use the Sparrow B0 chipset, whereas recent versions come with a Sparrow D0 chipset.

6.3.1.2. Software

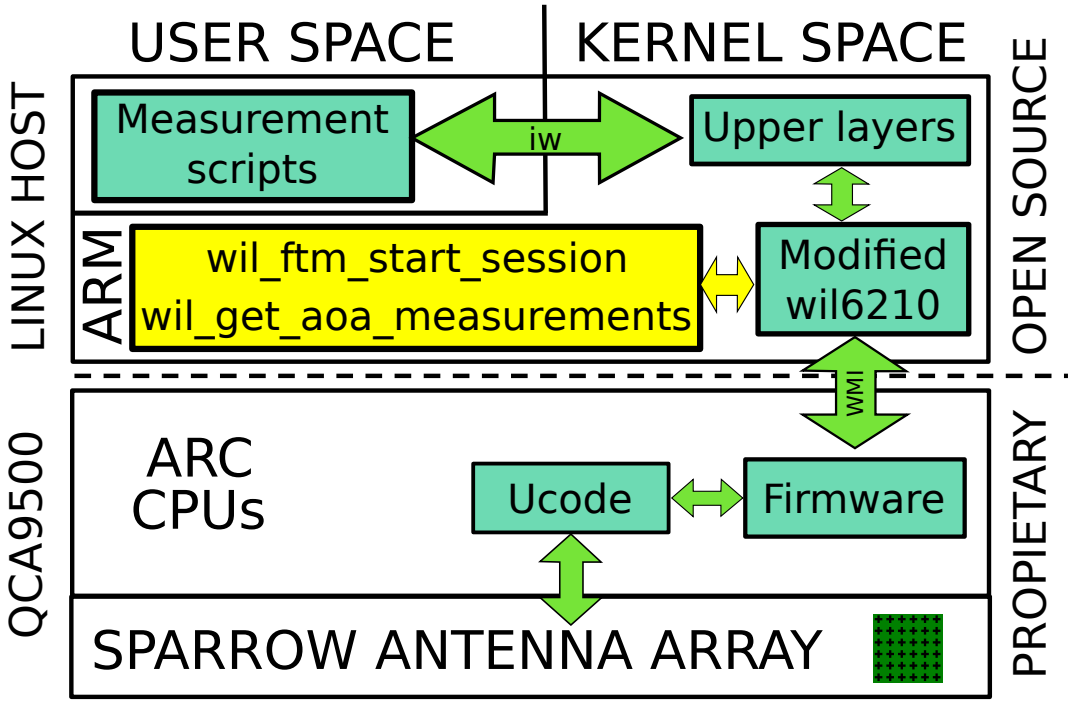

Figure 6.3: mmWave system overview modifications

MikroTik devices use a proprietary operating system called RouterOS based on a heavily modified 3.3.5 Linux Kernel with a unique binary called *nova* that bundles all functionality. This OS is very limited, and only very basic operations are enabled. In order to do research with these devices, we created our own custom OpenWRT fork with 60 GHz support. Our original development was done with first generation devices which have a Sparrow B0 chipset, which has an official Qualcomm firmware that supports it. However,

current MikroTik wAP 60G and 60Gx3 APs come with a Sparrow D0 (QCA9500), for which no official Qualcomm firmware is available. After some trial and error we found that it is possible extract the required firmware from the package wireless-6.40.9-arm.npk available on the MikroTik website, which contains the wil6210-d0.fw firmware. Newer packages are encrypted and are not compatible with the Linux driver. Furthermore, not all the firmwares have all functionality enabled (in particular CSI and FTM), and thus finding this compatible firmware was very important.

Despite having a firmware that works for newer chipsets, there is very little documentation on how to use the capabilities. All the communication from the operating system to the chipset and vice versa is done through Wireless Module Interface (WMI) calls. Most of those calls are defined in a file called wmi.h which is part of the wil6210 driver, the official Linux driver for IEEE 802.11ad communications. This file hints at CSI and FTM commands, but provides no information on how to invoke them. By carefully checking the different WMI structures, we were able to create the functionality to enable the calls. We then integrated the new commands with our wil6210 driver, and modified them to be able to pass through calls from and return information to user space, as shown in Fig. 6.3. To this end, we ported the CSI extraction functionality available for the old chipset on the old Talon AD7200 devices [49] to our OpenWRT image. FTM funcionality was previously not available, and to the best of our knowledge this is the first time that FTM has been enabled in mmWave COTS devices. A CSI measurement only requires one WMI command with few parameters, whereas for FTM we need to start a specific FTM session with a well defined set of parameters. After automating the parameter testing, we found that the only parameters that successfully start such a session are the ones listed in chapter 7.

6.3.2. Sub-6 GHz system

For the sub-6 GHz localization implementation based on IEEE 802.11ac, we use two COTS devices as was done in other prior work [61], since to date no freely available system supports both CSI and FTM extraction in the same device. Fig. 6.2 shows the devices we use. The *CSI device* is an ASUS AC2900 RT-AC86U with a Broadcom BCM4366E chipset. We use the Nexmon CSI extractor tool from [62] to obtain CSI for 4x4 MIMO and 80 MHz IEEE 802.11ac frames. In this configuration the extracted CSI matrix has a size of 4x4x256 complex values, where 256 corresponds to the number of subcarriers of the 80 MHz channel. As *FTM device* we use an ASUS AC1300 RT-ACRH13 router which supports FTM requests out of the box and a laptop with an Intel Dual Band Wireless-AC 8260 wireless interface as initiator, as in [63]. We apply some modifications to the Linux driver of the card to start the FTM session and extract the FTM measurements.

6.4. System calibration

Unfortunately, a range of preliminary measurements reveal that some calibration is needed to handle hardware imperfection that affect the CSI and ToF values, which would considerably reduce the accuracy of the AOA and ToF estimates.

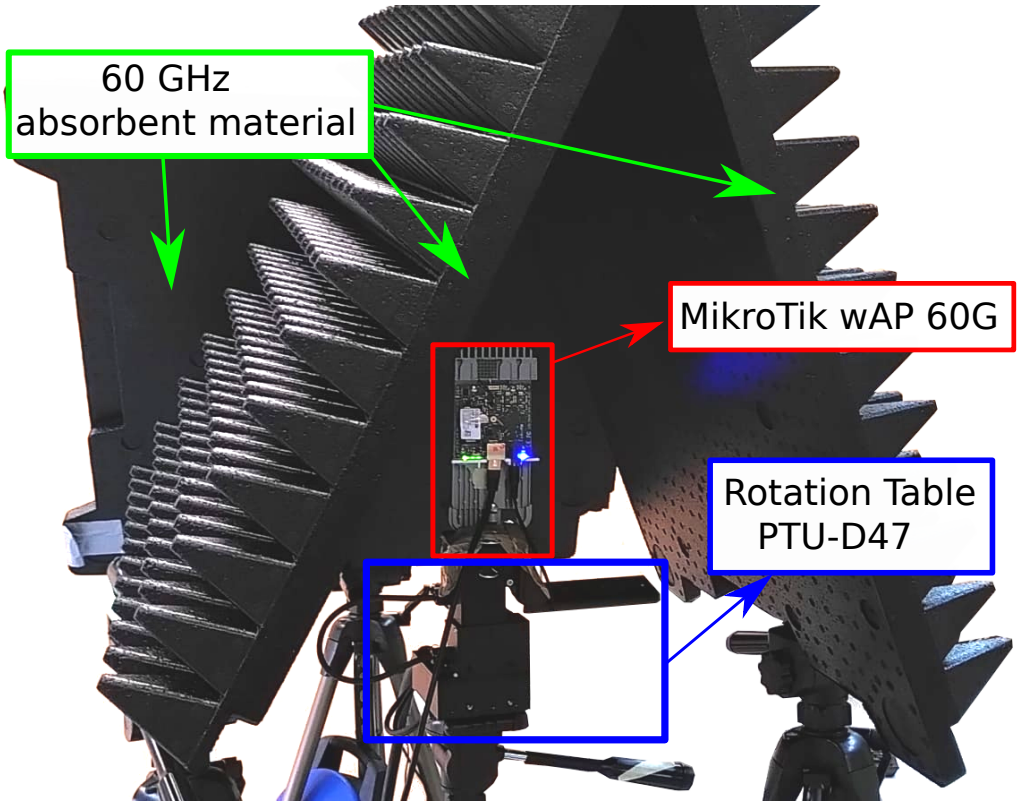

Figure 6.4: Calibration setup

6.4.1. 60 GHz CSI

A CSI measurement contains a phase and amplitude value for each of the 32 antenna elements. Note that for all devices, antenna elements 31 and 32 only report noise, which leaves a total of 30 antenna elements to work with. The value reported by the driver is an integer between 0 and 1024 which corresponds to a phase between 0 and 2π. For every CSI measurement, we convert the 30 CSI values into a matrix form of 6x6 with the same layout as in fig. 6.2, to convert the logical to the corresponding physical antenna elements in the rectangular array. Further analysis of the CSI data reveals that we have to handle the phase added by the oscillator to every antenna in the array and some systematic phase shifts that are WiFi frame dependent, in order to obtain a meaningful AOA estimate.

• **Phase offsets.** Since the devices have one radio-frequency (RF) chain, there is one common oscillator to tune to the central frequency. However, the internal circuitry and

layout of the device introduce phase offsets at each antenna element, which need to be removed for accurate AOA estimation. To this end, we mount an AP/client pair in a shielded environment at 0° azimuth and elevation angles, as seen in Fig. 6.4. This ensures that the direct path itself does not introduce any phase delays and the shielding removes any multipath effects that may distort the phase. We then measure the reference CSI that corresponds to the phases offsets. To remove them from actual CSI measurements, we perform the element-wise Hadamard division between measured CSI data and the reference CSI. These offsets are constant over power-cycles and for different devices with the same antenna type, and thus this calibration has to be done only once.

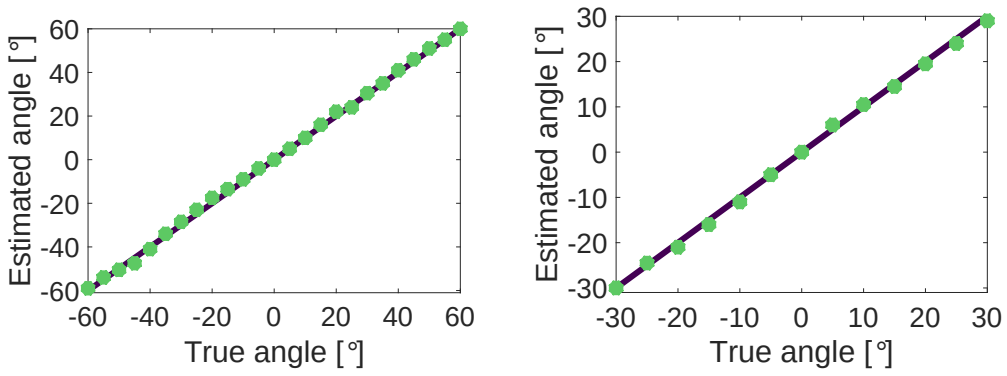

(a) Azimuth (b) Elevation

Figure 6.5: Estimated azimuth and elevation angles (green) compared to ground truth (black line)

We verify the calibration using a rotation table that rotates the device in steps of 5° both in the azimuth within ±60° and elevation angles within ±30°. As shown in fig. 6.5, the estimated angles closely follow the ground truth, with very low angle errors below 2°. This measurement also shows that under ideal conditions, accurate angle estimation is possible far beyond the official antenna aperture of ±30°.

• **Systematic phase shifts.** We also observe systematic phase shifts that affect all antenna elements and that appear randomly over the WiFi packets, as shown in fig. 6.6a. These phase jumps remain the same as can be seen from the two Gaussian distributions in fig. 6.6b, one which corresponds to the correct phase which is centered at 0 and another one centered between $3\pi/8$ and $\pi/2$. We can thus compensate the phase shifts as they follow a clear pattern, rather than discard them. To do so, we apply a Gaussian mixture model to the phase values over a set of WiFi packets to model the two Gaussian distributions, one for the correct CSI phases and the other for the phase-rotated CSI. We can then shift the mean of the less populated Gaussian distribution to center it on the correct phase value.

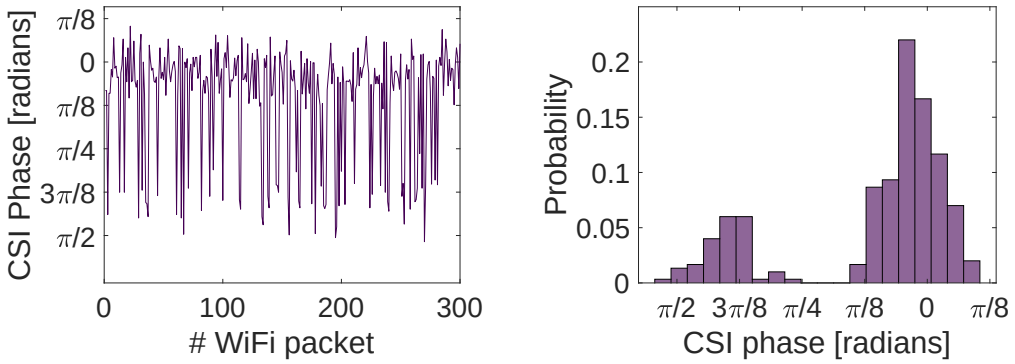

(a) Phase evolution (b) Phase histogram

Figure 6.6: Example of the systematic phase shifts for a specific antenna in the array

6.4.2. 60 GHz ToF

We further observe an offset when capturing FTM measurements. To calibrate the FTM, we use the same setup as for CSI and set the routers at 2.5 m of distance. We then take 100 measurements and use the difference to the true distance to correct the offset. Also this calibration only needs to be done once per device.

6.4.3. sub-6 GHz CSI

The oscillator at sub-6GHz introduces a phase uncertainty at each of the four RF chains, which we have to remove it to compute an accurate AOA. We connect each transmit port from the AP to each receive port at the client over cables of the same length [53], which ensures that the signal arrives at each receiver chain at the same time. The reported CSI values thus contain the constant phase offsets.

6.4.4. sub-6 GHz ToF

Finally, the sub-6 GHz FTM measurements have an offset which we determine and correct by taking outdoor measurements at known distances.

6.5. Multi-band localization

We now explain the details of the MultiLoc localization techniques. We first describe the AOA estimation for both mmWave and sub-6 GHz then the ToF. Finally, we discuss how to combine the measurements from several APs and for both frequency bands to derive an accurate location estimate.

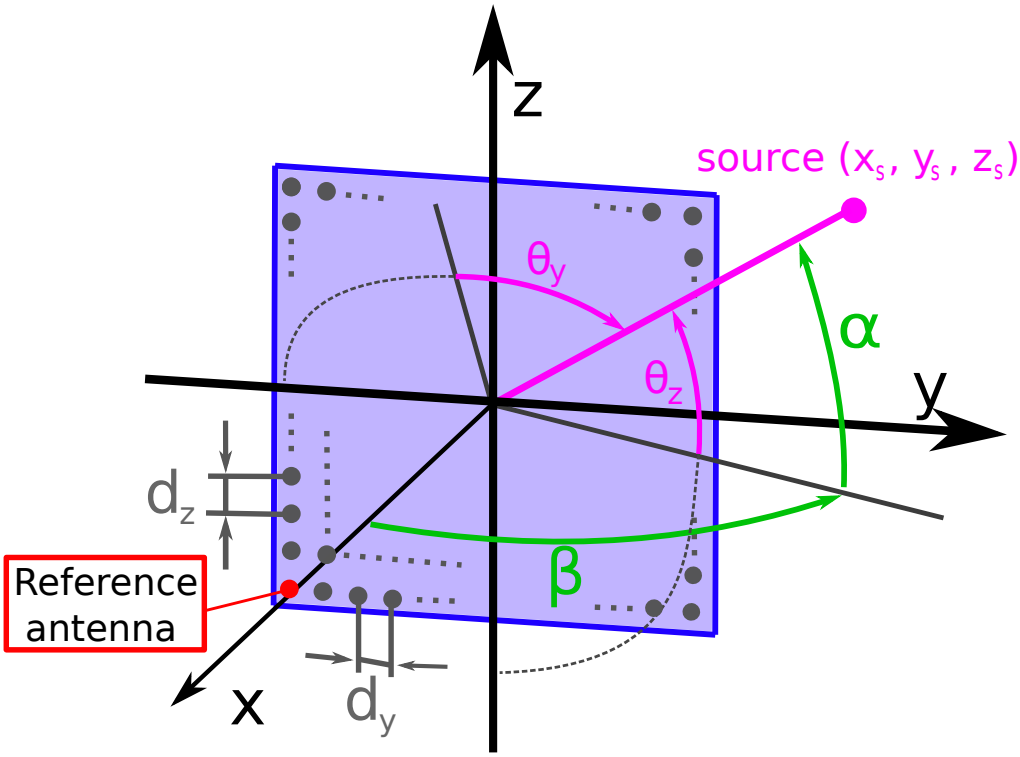

Figure 6.7: Layout of the Uniform Rectangular Array.

6.5.1. Angle of arrival

We consider a wireless system where the transmitter and receiver have a single RF-chain connected to a URA with MxN antenna elements. The signal sent by the transmitter propagates through the multipath channel along P different paths and arrives at the receiver. In the following we refer to fig. 6.7 where the transmitter (the source) is located at (x_s, y_s, z_s) and the URA lies in the $x = 0$ plane. We characterize the paths from the source as they appear at the centre of the URA using the following parameters:

- **Complex attenuation** γ of the signal over the path.
- **Elevation angle** α between the path and the plane $z = 0$.
- **Azimuth angle** β between the path and the plane $y = 0$.

In order to extract the three parameters, we first need to characterize how the phase rotation in the elements of the URA changes over z and y. We denote by θ_z and θ_y the angles that characterize the antennas at the Z and Y axis, see fig. 6.7. We formulate the phase rotation introduced by the signal on every antenna element according to the reference antenna element as:

$$a_{m,n} = md_z sin(\theta_z) + nd_y sin(\theta_y), \tag{6.1}$$

where m goes from 0 to $M-1$, n goes from 0 to $N-1$, and $(m,n) = (0,0)$ is selected as

the reference antenna. Thus, the vector that contains the phases for the antennas across the Z direction is defined as:

$$\boldsymbol{\phi}(\theta_z) = [1, e^{-j2\pi d_z \sin(\theta_z)/\lambda}, ..., e^{-j2\pi d_z (m-1) \sin(\theta_z)/\lambda}], \tag{6.2}$$

where λ is the wavelength of the RF signal. Similarly, for the antennas across the Y direction we have:

$$\boldsymbol{\phi}(\theta_y) = [1, e^{-j2\pi d_z \sin(\theta_y)/\lambda}, ..., e^{-j2\pi d_z (n-1) \sin(\theta_y)/\lambda}] \tag{6.3}$$

Both angles θ_z and θ_y depend on the the direction of the path, which is characterized by the elevation and azimuth angles. From fig. 6.7 we can see that θ_z corresponds to elevation α, so that $sin(\theta_z)$ is equal to $sin(\alpha)$. Angle θ_y, instead, depends on both the elevation and azimuth according to the relation $sin(\theta_z) = sin(\alpha)cos(\beta)$

With the parameters above, we can express the channel by summing all the contributions given by the P paths as follows:

$$\mathbf{H} = \sum_{p=0}^{P-1} \boldsymbol{\phi}^{\mathrm{T}}(\theta_{z,p}) \gamma_p \boldsymbol{\phi}(\theta_{y,p}) + \mathbf{N} \ , \tag{6.4}$$

where $\mathbf{N}$ is white a Gaussian noise matrix of size MxN and $(\cdot)^T$ is the Transpose operator.

6.5.2. AoA estimation at mmWave

The channel described above contains the path parameters for all the paths, but for localization we are interested in the direct path between the client and the AP. Estimating the path parameters of superimposed signals is not trivial, since they interfere with each other when they arrive at the receiver. To separate the paths and accurately extract the parameters of the direct path, we follow the approach of mD-track [53]. This algorithm estimates the path parameters iteratively by reconstructing the strongest path and its parameters and then subtracting it from the received signal, so that successively weaker paths can be estimated. Once it obtains an initial estimate for all the paths, it applies several rounds refinements. Note that mD-track was designed to work with a uniform linear array. We extend the algorithm to adapt it to the mmWave URA.

Initial estimation. We iteratively determine the contribution of the strongest path, estimate its parameters, reconstruct it, and then subtract it from the received signal. The result of the subtraction is the channel residual, with which we then estimate the second strongest path's contribution, and so on. We iterate multiple times until the parameters of all the significant paths are estimated.

To compute the path parameters of the p-th path, we apply a matching projection to the channel residual based on the path parameters. We formulate the matching projection

for the p-th path as:

$$z_p(\theta_z, \theta_y) = \boldsymbol{\phi}^*(\theta_z)\mathbf{H}^{\mathrm{r}}_{\mathbf{p}}\boldsymbol{\phi}^{\mathbf{H}}(\theta_{\mathbf{y}})\,, \tag{6.5}$$

where $(\cdot)^*$ and $(\cdot)^H$ are the conjugate and Hermitian operator, respectively, and H^{r}_p is the channel residual. The residual for $p = 0$ is the original channel $\hat{\mathbf{H}}^{\mathrm{r}}_0 = \mathbf{H}$. The projection is calculated over all possible combinations of θ_z and θ_y, and the estimated parameters are the ones that maximize the projection:

$$\hat{\theta}_{z,p}, \hat{\theta}_{y,p} = \arg\max_{\theta_z,\theta_y} z_p(\theta_z, \theta_y)\,. \tag{6.6}$$

The attenuation is thus estimated as

$$\hat{\gamma}_p = \frac{1}{M \cdot N} z_p(\hat{\theta}_{z,p}, \hat{\theta}_{y,p})\,, \tag{6.7}$$

and we can reconstruct the channel for path p as

$$\hat{\mathbf{H}}_p = \boldsymbol{\phi}^{\mathrm{T}}(\hat{\theta}_{z,p})\hat{\gamma}_p\boldsymbol{\phi}(\hat{\theta}_{y,p})\,. \tag{6.8}$$

The residual at iteration p is given by

$$\mathbf{H}^{\mathrm{r}}_p = \mathbf{H} - \sum_{p'=0}^{p-1} \hat{\mathbf{H}}_{p'}\,. \tag{6.9}$$

The initial estimation ends when all the P paths are estimated.

Iterative path parameter refinement. The initial estimates may contain errors since the subtraction step may leak information from stronger paths to weaker paths and vice-versa. To do so, we apply several rounds of refinement by summing the residual to the reconstructed path so that we can iterate over this path and re-estimate its parameters another time. By adding the reconstructed path to the residual we get the noisy path as follows:

$$\hat{\mathbf{H}}^n_p = \hat{\mathbf{H}}_p + \mathbf{H}^{\mathrm{r}}_p \tag{6.10}$$

We re-estimate the path parameters for the noisy path as in the initial estimation rounds and the resulting path parameters will contain less interference than before. We also re-estimate the residual for every path so that the residual also contains less interference. We continue iterating until the difference between the path parameters obtained in the last two rounds is lower than a configurable threshold.

Once the iterative refinement finishes, we extract the path parameters of the direct path. Since the path parameters do not contain any temporal information as in an OFDM system, we select the direct path as the one with highest power: we express the power of every path as $power_p = ||\gamma||^2$. We denote the direct path as p_{dp}. Since we are interested

in the azimuth angle to localize, we need to extract it from the $\theta_{z,p_{dp}}$ as follows:

$$\hat{\theta} = \text{asin}(\sin(\theta_{z,p_{dp}})/\cos(\theta_{y,p_{dp}}))\,. \tag{6.11}$$

6.5.3. AoA estimation at sub-6 GHz

We use the same iterative procedure also for sub-6 GHz. Here we deal with OFDM MIMO systems with uniform linear arrays. For every path we can extract the attenuation, the angles of arrival and departure, and the path length, which corresponds to the time delay of the path. The AOA estimator behaves similarly since we extract all the paths and their parameters. The strategy to select the direct path, instead, is different, since for sub-6 GHz we can choose as direct path the one that actually arrives first. We skip the details due to space constraints and we refer interested readers to [53].

6.5.4. Spatial ambiguity

Spatial ambiguity occurs when the spacing between consecutive elements of the array is larger than half of the wavelength (λ), which is the case with the 60 GHz antenna's spacing of 0.58 λ. Here, several AOAs may produce the same phase delays at the antennas. To remove the ambiguity and select the correct AOA, we first identify the possible AOAs and we then choose the correct one based on the AOA estimate at sub-6 GHz. We choose as mmWave AOA the one that minimizes the angular difference with respect to the sub-6 GHz AOA. At sub-6 GHz there are no ambiguities because the antenna spacing is exactly half of the wavelength.

6.6. FTM ranging

IEEE P802.11-REVmc [64] describes the FTM procedure that can be used by two FTM capable devices to estimate their distance without clock synchronization. The device that starts the FTM session is the *Initiator*, and the other is the *Responder*. The Initiator starts the FTM session by sending a FTM request frame, which embeds the parameters with the configuration of the FTM session as in Table 1. The Responder then has 10 ms to acknowledge the request and confirm that the FTM session is started. An FTM session contains one or more bursts that in turn are composed of one or more measurements. A burst is a specific period of time during which the channel is reserved for FTM measurements, so that all the measurements performed in the same burst are consecutive. Fig. 6.8 shows an example with one burst consisting of 4 measurements. After the FTM request has been acknowledged, the Responder sends the first FTM packet with t_1 and t_4 set to 0. Then, for the rest of the measurements the Responder sends an FTM frame with the local clock timestamp of when it is sent (t_1), as well as a

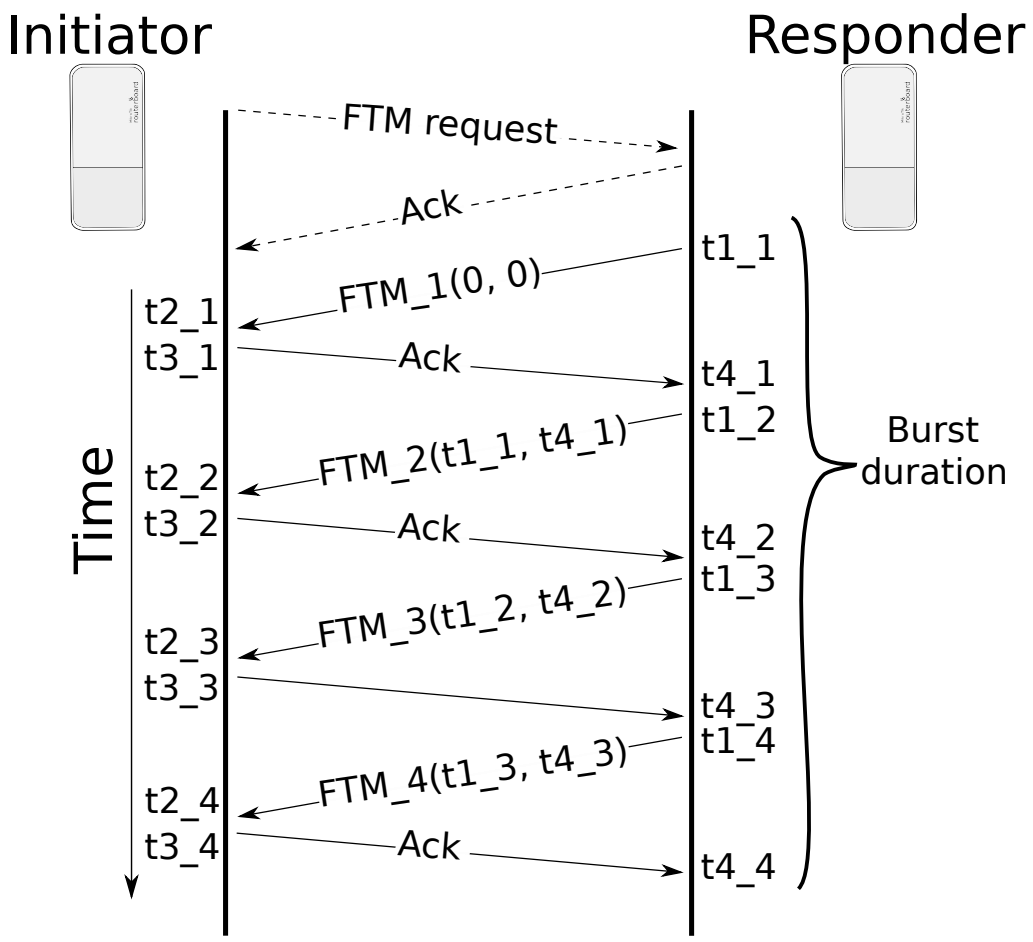

Figure 6.8: Example of an FTM session with 1 burst composed of 4 measurements

local clock timestamp of when the last Ack was received $(t-4)$. Thus, in FTM frame N we obtain information about the measurement $N-1$ for $N>1$. This is repeated until all measurements are done. Since there is a full round trip, the offset between the unsynchronized clocks is compensated: the offset on one way is cancelled by the opposite offset on the other way. We can estimate the Round Trip Time (RTT) as follows, where n is the number of measurements:

$$RTT = \frac{1}{n}\sum_{x=1}^{n}((t_4(x) - t_1(x)) - (t_3(x) - t_2(x)))$$

We can then determine the distance in meters using the speed of light c as $\hat{d} = c \cdot RTT/2$. This procedure is the same for mmWave and sub-6 GHz.

6.7. Location estimation

Under ideal conditions when AP and client are in LOS, the accuracy of mmWave is always better than that of sub-6 GHz. In contrast, the mmWave measurements are affected by large errors in case of NLOS. Hence, MultiLoc needs to determine which of the mmWave or sub-6 GHz measurements to use. We describe in the following the principles of MultiLoc.

First of all, the a-th AP computes the location of the client using its own known position $\mathbf{x}_a$ as

$$\hat{\mathbf{p}}_a = \mathbf{x}_a + \hat{d}_a \begin{bmatrix} \cos\hat{\theta}_a \\ \sin\hat{\theta}_a \end{bmatrix}, \tag{6.12}$$

We denote $\hat{\mathbf{p}}_{a,lf}$ and $\hat{\mathbf{p}}_{a,hf}$ for the computed location using the sub-6 GHz and mmWave information, respectively. All the location information is gathered at a central controller that computes the client position. Before showing how MultiLoc merges the mmWave and sub-6 GHz location estimates, we first explain how we localize at mmWave and sub-6 GHz bands individually.

6.7.1. Standalone mmWave location system

Typically, no more than one mmWave AP will cover a specific room. The large path attenuation when traversing obstacles makes it unlikely to have good multi-AP coverage. It is also possible to have more than one antenna array on each device, which in turn help with the angular orientation. The first challenge is to pick the best array. For for each point we obtain K position estimates, where K is the number of arrays on the client. We then compute $\hat{\mathbf{p}}_{a,hf,k}$ which is the estimated position by mmWave from the AP a and the array k.

For each of these measurements, we estimate the distance and pick the combination of array and AP which minimizes the estimated distance. Since mmWave behaves quasi-optically, the shortest distance is most likely to point to the LOS path and to the best AP if there is more than one within range. The position estimate at mmWave is $\hat{\mathbf{p}}_{hf} = \hat{\mathbf{p}}_{a',hf,k'}$, where a' and k' are the AP and array that result in the lowest distance estimate.

6.7.2. Standalone sub-6 GHz localization

The sub-6 GHz band has better coverage than mmWave and is more likely to penetrate through obstacles. Therefore, the sub-6 GHz localization system has to deal with more position estimates since more APs are able to localize the client than with mmWave. Obstructed paths may provide useful position information but the location estimates from the farthest APs to the client are more likely to be unreliable. We thus compensate

the unreliable estimates by weighting the location estimate from every AP based on the estimated distance. We give the weight values of 1 and 0 to the APs which have the minimum and maximum estimated distance, respectively. For the remaining APs, we compute their weights by a linear interpolation between the estimated distances and the minimum and maximum distances. The position at sub-6 GHz is computed by the weighted centroid approach as follows:

$$\hat{\mathbf{p}}_{lf} = \frac{\sum_{a=0}^{A-1} \hat{\mathbf{p}}_{a,lf} \cdot w_a}{\sum_{a=0}^{A-1} w_a}, \tag{6.13}$$

where w_a is the weight given by the a-th AP and A is the number of available APs.

6.7.3. MultiLoc localization

MultiLoc combines the best from both worlds, the outstanding localization precision of mmWave and the availability of sub-6GHz. MultiLoc provides the localization of mmWave if only if it is reliable, if not MultiLoc takes the sub-6GHz positioning. To do so, MultiLoc applies three steps. (1) MultiLoc takes the best array from the client for localization. However, the best array at the client may be still unreliable because it is pointing to an obstacle instead of the AP, thus (2) MultiLoc applies a sub-6GHz and mmWave discrimination. Finally, (3) MultiLoc selects the best AP. We explain in details the MultiLoc steps below.

(1) Best array selection. As explained before, the client may have several antenna arrays pointing in different directions. We are interested in determining the array that points in the direction of the AP, to avoid the large bias in the localization that arises when MultiLoc takes the AOA and ToF of a reflection. To do so, we assume that the array with the lowest estimated distance is the correct array since the direct path is the shortest path from the client to the AP and is most likely to correspond to the LOS path.

(2) Sub-6GHz/mmWave discrimination. This is the most critical step of MultiLoc. Selecting the mmWave information from a reflection leads to high localization errors, but at the same time, whenever choosing the sub-6 GHz location even though the mmWave one is correct, MultiLoc looses accuracy. To this end, MultiLoc measures the distance between position estimates at mmWave and sub-6 GHz per AP. If the distance between estimates is high, this is a good indication that mmWave localization has failed and did not select a LOS path, since that sub-6 GHz is less precise but more robust. In contrast, if this distance is low, mmWave localization is likely to be accurate and MultiLoc should use it rather than the sub-6 GHz information. We set a threshold th_d, and if the distance between the estimates is lower than the th_d, MultiLoc takes the mmWave estimate and otherwise it uses sub-6 GHz. We fix th_d to 5 m. We denote as $\hat{\mathbf{p}}_a$ the position estimate after applying discrimination for the a-th AP.

(3) Best AP selection. Having chosen between sub-6 GHz and mmWave for each

individual AP estimate, MultiLoc has to determine the best AP for localization from all the AP candidates. Note that there may be APs with NLOS in the set of candidates, and taking one of those APs to localize will create large errors. Therefore, MultiLoc has to take the AP with the best channel conditions. We observed that the lowest values of the estimated distance are highly correlated with the most accurate localization estimates. Thus, we selected the AP with the lowest estimated distance at sub-6GHz, denote by AP as a_{best}. We based this selection on sub-6GHz because is agnostic to the client rotation since the mmWave ToF of the AP in LOS may contain outliers due to the wrong rotation of the client. Thus, the final position estimate of MultiLoc is given by $\hat{\mathbf{p}}_{\text{MultiLoc}} = \hat{\mathbf{p}}_{a_{best}}$.

6.8. Experimental evaluation

To further investigate the impact of the high directionality and the high attenuation of mmWave, we run an extensive measurement campaign in two different scenarios, an indoor and an outdoor one. In each scenario, every AP operates on both the sub-6 GHz and 60 GHz band. Fig. 6.9a shows the map of the indoor environment that comprises six rooms and one corridor. While there are five mmWave APs that we indicate with the numbers of the rooms, only one AP can be in LOS at each measurement point even though some rooms are covered by more than one AP: for instance, the bottom-left corner of room 2 is covered by AP 2 and 7. Fig. 6.9b uses the same indoor environment but with a more sparse deployment of only two APs that ensure at least sub-6 GHz coverage everywhere (not show in the figures). We observe a surprising degree of NLOS mmWave even through walls, and for the 5 AP setup we have ubiquitous mmWave coverage even in rooms without mmWave AP. Finally, Fig. 6.10 shows the setup for the outdoor environment, where four APs cover the whole measurement area. For the indoor environment we take 110 measurement points, and we collect 16 additional points for the outdoor one, for a total of 126 measurement points. At each point we collect CSI and FTM measurements for both sub-6 GHz and 60 GHz.

6.8.1. Experimental setup

We use the same client setup for both indoor and outdoor scenarios. We mount 2 MikroTik wAP 60G devices back to back on top of a rotation table model PTU-D47. This table allows for rotations of up to 200° so to cover the full 360° azimuth we need two devices. We then take 8 different rotations at 45° steps, and each client connects to each AP in turn. We take 300 CSI measurements and 10 FTM measurements per rotation and AP. This allows us to carry out an in-depth analysis for different client rotations, and enables us to emulate different numbers of mmWave subarrays.[1] For sub-6 GHz we use 2

[1]We tried but unfortunately were not able to obtain an Airfide AP that directly integrates 8 of these antenna arrays. However, there should not be any difference between the Airfide setup an our emulated

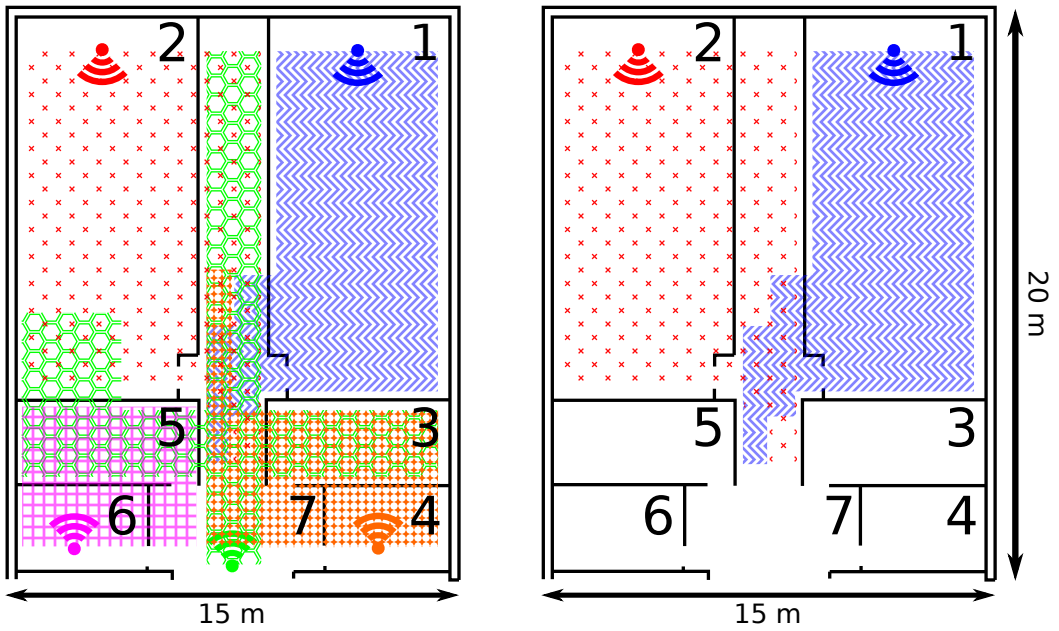

(a) 5 APs (b) 2 APs

Figure 6.9: Indoor 60 GHz AP coverage map

different routers, one to capture the CSI and a second one to capture the ToF.

6.8.2. Rotation analysis

In order to expand on the intuitions of section 6.2, we perform an extensive rotation analysis in the indoor scenario. Fig. 6.9 shows the coverage map for that scenario. We have divided the analysis of the indoor scenario in two separate sets to better explain the challenges to overcome, i.e., one set with positions and rotations where one of the APs is in LOS, and one set with the positions and rotations where no AP is in LOS for AOA and ToF ranging. As explained in Section 6.2, mmWave is very susceptible to antenna misalignment when the client array is not directly facing the AP. Therefore, we quantify the estimation error according to array orientation, where the geometrical oracle indicates the array that is most closely aligned with the direct line between AP and client, and the misalignment in degrees indicate by how much the other orientations deviate from that direct line. Fig. 6.11 shows the Empirical CDF for all the cases. In particular, subfigures (a) and (b) show the results for AOA in LOS and NLOS, respectively, and (c) and (d) show ToF in LOS and NLOS. The figures include all estimates from all APs, not just the ones that would be chosen by MultiLoc.

From Fig. 6.11a with AOA error for indoor LOS, we observe that if we choose the optimal rotation, there is less than 2.75° of angle estimation error for 90% of the cases. The

8 antenna arrays other than causing a longer measurement duration.

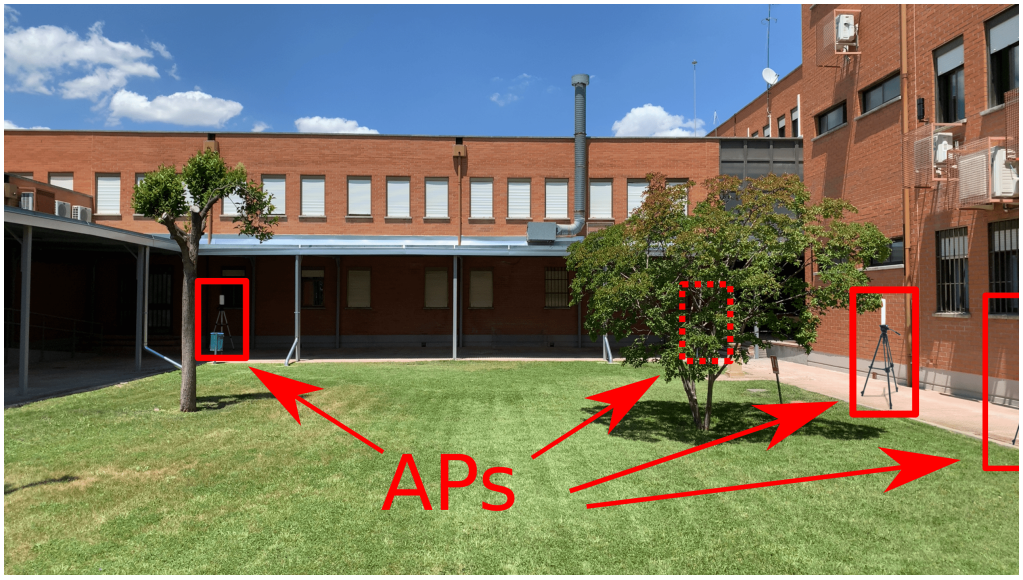

Figure 6.10: Outdoor scenario

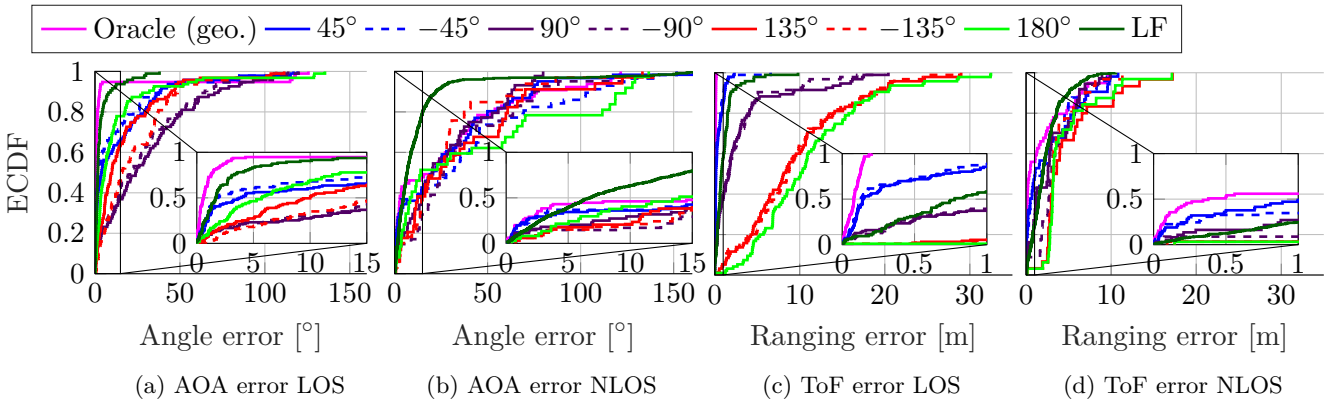

(a) AOA error LOS (b) AOA error NLOS (c) ToF error LOS (d) ToF error NLOS

Figure 6.11: AOA and ranging errors for the indoor scenario

next best rotation is 180° with an error of less than 33° for 90% of the cases since it gets the direct path from behind and estimates a good angle with the sign changed, although 45° has a better median of 1.8° of error. The worst estimations are obtained when the antenna arrays are perpendicular to each other, i.e., ±90°. The default beampatterns adopted by the MikroTik devices in these conditions favor pointing to the right side of the device, which can be seen also for the two rotations of ±45°. As a comparison, results for sub-6 GHz in the same conditions reveal less than 10° of error for 90% of the cases. It is also noticeable that the AOA errors for the oracle rotation jump from less than 5° to more than 120°, which is due to the spatial ambiguity as it may take the incorrect AOA. As explained in section 6.5.4, we can use sub-6GHz information to resolve this ambiguity, and thus MultiLoc itself has a maximum error of 6.04°. Concerning the ToF ranging, Fig. 6.11c shows the estimated distance error for the same points. The oracle rotation has an

error of less than 0.20 m for the whole scenario, and even for two important deviations like ±45° we obtain less than 1.2 m of error for 90% of the cases. The worst case here is for 180°, as it always captures a very long reflected path. In comparison, sub-6 GHz reports less than 2 m error for 90% of the cases. We next summarize the NLOS case in Fig. 6.11b. We can easily see that the AOA estimated for mmWave is affected by huge errors for all rotations, while sub-6 GHz always performs better and reports less than 23° of error for the 90% . Finally, we show in Fig. 6.11d the estimated distance error for the NLOS case. mmWave still yields better results than sub-6 GHz for 80% of the cases, with a median of 0.4 m for the oracle rotation. Although the corridor has only one 60 GHz AP in LOS, all of the other NLOS APs still provide coverage to some parts of it. Next to the doors to the rooms, for example the door in room number 1, the coverage mainly comes from reflections, but the top part of the corridor has also coverage from AP 2 via obstructed LOS directly through the wall.

For brevity we only briefly discuss the outdoor scenario, which is simpler than the indoor one since there is no significant NLOS. As in the indoor scenario, we have less than 3.5° of AOA error for 90% of the cases. Here, the two offsets ±45° perform better than 180° as there are no close walls that can reflect that beam. Once more the two offsets ±90° are the worst. With sub-6 GHz we get less than 7° error for 90% of the cases. As for the distance, the oracle rotation reports less than 0.55 m for all the points, followed by offset −45° with an error of less than 0.42 m for 90% of the cases. For sub-6 GHz we obtain an error of less than 4.8 m for all the points.

In summary, the performance of mmWave is excellent under ideal conditions but quickly degrades even for small changes in device orientation. Sub-6 GHz is relatively agnostic to rotation due to its omnidirecional nature and has less path attenuation when traversing obstacles and walls.

6.8.3. Location results

As explained in Section 6.7, MultiLoc can overcome the limitations of standalone mmWave location systems by identifying when mmWave estimates are accurate and when they are not. We now analyze the joint location accuracy based on the number of antenna arrays that are available for mmWave, for 1, 2, 4 and 8 antenna arrays at the client side. In the case of 1 array, we provide the localization errors per measurement point for each of the 8 possible rotations. For the case of 2 arrays back-to-back, we have 4 possible rotations. In the case of 4 arrays, the arrays are placed in a diamond shape, which leave two possible rotations. Finally, for the 8 arrays case, there is only one possible rotation, and we thus include only one localization error per measurement point. We compare MultiLoc to the standalone mmWave location system (HF), which selects between its available antenna arrays and the available APs based on minimum ToF. We also show the standalone sub-6 GHz location system (LF), which weights the importance of the APs

based on signal strength and then does a weighted average to estimate the position.

6.8.4. Indoor scenario with 5 APs

The first scenario has 5 60 GHz and sub-6 GHz multi-band APs in rooms 1, 2, 4 and 6 as well as in the corridor 7. This dense deployment only has NLOS in rooms 3 and 5 as seen in Fig. 6.9a where we take 14 measurement points. In this setting, MultiLoc has a median error of 0.65, 0.2, 0.13 and 0.13 m for 1, 2, 4 and 8 antenna arrays, respectively. This accuracy mostly comes from the standalone mmWave system which has the same median error of 0.13 m except for 1 array, with a median error of 3.1 m, where relies on reflections in many of the cases. For comparison, the standalone sub-6 GHz system has a median error of 0.98 m. The corridor, even though is in LOS, is a very complicated area due to the high level of multipath. Most of the errors of the standalone sub-6 GHz system come from this area. In contrast, mmWave has much less multipath due to the narrow antenna beams and the higher signal attenuation, which leads to higher accuracy. Lastly, even tough there are pure NLOS areas, MultiLoc is able to achieve an error of less than 0.5 m for all the cases. In this setting, MultiLoc is able to improve over the sub-6 GHz system by a factor of 4.5 for all the cases.

6.8.5. Indoor scenario with 2 APs

This scenario mimics a typical office environment, where the workspaces are located in rooms 1 and 2. The only APs are in rooms 1 and 2 to provide as many users as possible with mmWave connectivity, as seen in fig. 6.9b. We measure the same 110 points as in the 5 APs scenario, with the difference that there is only mmWave LOS in rooms 1 and 2, and obstructed LOS in the corridor, whereas the rest of the rooms have NLOS coverage from sub-6 GHz.

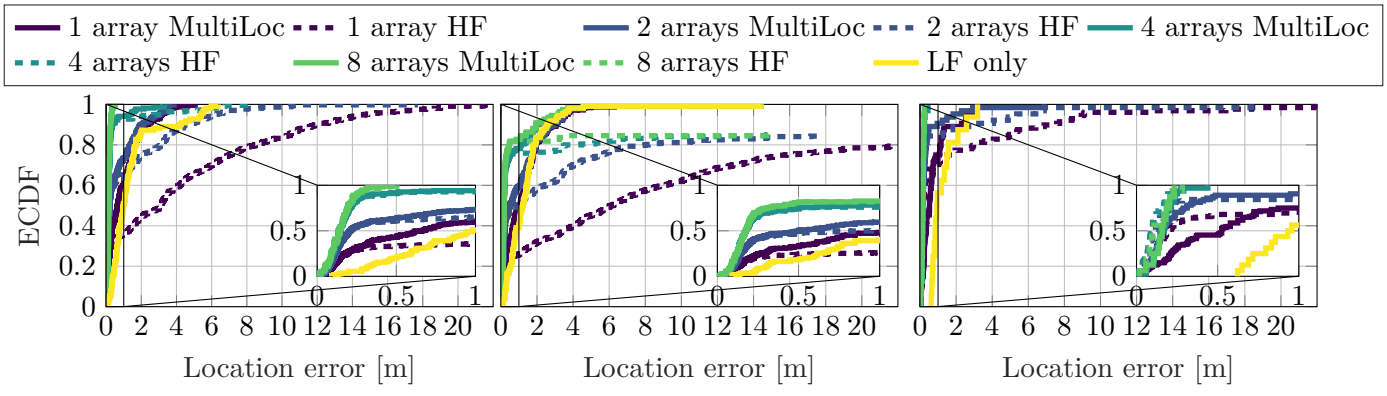

(a) Indoor scenario with 5 APs (b) Indoor scenario with 2 APs (c) Outdoor scenario

Figure 6.12: Location error for all the scenarios with different antenna array configurations

We summarize results in fig. 6.12b which shows for MultiLoc a median error of 1.1, 0.5, 0.17 and 0.17 m respectively for the case of 1, 2, 4 and 8 antenna arrays per AP. There, we also see that the standalone sub-6 GHz location system has a median error of 1.2 m. These results highly improve the accuracy of the mmWave standalone location system for 1 and 2 antenna arrays that has a median error of 6 and 0.9 m. For 4 and 8 antenna arrays the results remain the same, which is expected since MultiLoc is able to correctly identify mmWave as being the most accurate. With 1 antenna array, the standalone mmWave location system depends on reflections for the majority of the cases, but just by adding a second antenna array it is able to have a direct path in the LOS rooms for most of the cases. With 4 and 8 antenna arrays, MultiLoc is able to find a direct path in all the LOS rooms and even in some parts of the corridor where there is through-wall coverage from APs 1 and 2. Overall, MultiLoc is able to take advantage of mmWave in (obstructed) LOS areas and sub-6 GHz in NLOS, reducing the median error by a factor of 5.8 for the single array mmWave system compared to a standalone system.

6.8.6. Outdoor scenario

The outdoor scenario does not have as many reflections as the indoor one since the walls are much further away. This means that for most of the cases we either have a direct path or a very large reflection, which can be challenging for the standalone mmWave location systems. Every time the system chooses a reflection, MultiLoc is able to correctly identify it and fall back to sub-6 GHz. For the most challenging setting, MultiLoc is able to improve the accuracy of the single array mmWave system by 7.6 times compared to the standalone system.

6.9. Related work

• **COTS.** There has been previous work on mmWave COTS devices in the past. Authors in [65] do a extensive measurement campaign that shows the limitations of COTS devices. Authors in [66] use reverse engineering techniques to enable compressive path tracking into the Talon AD7200, the predecessor of the MikroTik wAP 60G used on this work. Authors in [67] also focus their work on the Talon AD7200, freeing it from the static codebook by computing the beam patterns in real time.

• **Multi-connectivity systems.** Many authors have combined sub-6 GHz and mmWave bands in different systems, most of the time with a focus on reducing the mmWave beam training time using sub-6 GHz information. Authors in [68] demonstrate a strong correlation between high and low frequency bands by collecting real sub-6 GHz measurements at 1.8 and 3.5 GHz and simulating the mmWave behaviour at 28 and 61 GHz. They hence validate the possibility to create a sub-6 GHz assisted mmWave beam steering mechanism. Authors in [57] also focus on sub-6 GHz assisted mmWave

beam steering and link maintenance. They experimentally demonstrate how to use sub-6 GHz ToF and an external gyroscope for estimating the position of the AP. In [56] authors simulate one sub-6 GHz antenna with 4 elements and one in mmWave range with 32 and show how to improve mmWave beam search using sub-6 GHz information. In [69] they show through simulation how to transform the spatial correlation matrix from sub-6 GHz to 60 GHz in order to reduce mmWave training overheads. On a similar line is [70] where the transformation is determined using deep-learning techniques. Authors in [71] show with simulations how to detect mmWave path blockages using sub-6 GHz Received Signal Strength (RSS) information. [55] introduces experimentally a mmWave link adaptation technique using sub-6 GHz information aimed at finding the best mmWave beam and detecting blockages. Similar to all these techniques, [72] introduces an AI-based technique for reusing CSI data from sub-6 GHz eNodeB to configure mmWave links between high speed trains and the mobile network. A different approach to reduce the overhead of mmWave beam training is proposed in [73] where only mmWave data gathered on a multiply connected mobile device is used for retraining each single mmWave link instead of reusing sub-6 GHz information. Authors in [74] review techniques that can improve mmWave IoT communications by taking advantage of location awareness and multi-connectivity operation.

• **Testbeds.** There has also been some effort in mmWave testbeds, authors in [60] present M-Cube, an open source 60 GHz testbed with up to 8 antenna arrays. They focus on reversing the control path for 60 GHz antennas, such as the one from the MikroTik devices, in order to have better control over the antenna. [75] shows OpenMili a 60 GHz reconfigurable testbed that makes use of custom antenna arrays. The platform X60 is presented in [76] which is a closed source 60 GHz testbed based on FPGAs, lastly authors in [77] propose MIMORPH an open source multiband sub-6 GHz and 60 GHz testbed based on FPGAs that supports real time operations.

• **Location systems.** There is also a lot of interest in general purpose location systems. Authors in [78] present a sub-6 GHz location system with sub-meter accuracy by using ToF and AOA from COTS devices. Another sub-6 GHz location system based on AOA and ToF is [79] which achieves decimeter-level accuracy with a single access point. And lately there has been special interest in mmWave location systems, due to its higher spatial and temporal resolution. Authors in [80] locate devices using Angle Difference Of Arrival (ADOA) which shows submeter accuracy for a small scenarios with 4 LOS APs. Time-Difference of Arrival (TDOA) is, instead, experimentally demonstrated in [81] using non-synchronized anchors for locating with submeter accuracy non cooperating clients. In [82] authors use a single mmWave RF chain and virtual anchors with previous knowledge of the beam codebook. They use a horn antenna with Vubiq hardware and demonstrate very promising results for short distances. Authors in [50] show a passive location system that is capable of detecting multiple people in a range of around 4 meters. [49] presents

a mmWave based system that can be used for predicting handovers and for location purposes. Over a dense deployment with 7 APs in LOS it reports a median error of around 0.7 m. [58] is a multiband location system. It uses sub-6 GHz for AOA and 60 GHz for ranging. In contrast to all the other mentioned location systems, MultiLoc can work with more realistic APs deployment (1 per room in indoor), is fully implemented on COTS devices and it has been deployed in very large scenarios.

6.10. Conclusions

In this chapter, we demonstrated that the accuracy of stand-alone mmWave location systems highly depends on the antenna array orientation. To address this problem, we introduced MultiLoc, a dual band location system where we use the sub-6 GHz technology to i) decide which of the mmWave APs and which of the available antenna arrays should be chosen when localizing a target and to ii) directly determine the position of the client when the mmWave location estimate is unreliable. We carried out an extensive measurement campaign for both bands to evaluate MultiLoc. The results show that our system is the first to achieve less than 17 cm median location error with off-the-shelf devices in indoor environments. To inspire future research on mmWave and multi-band localization, we will release the location system and the CSI and FTM extraction tools, as well as the dataset of our measurements campaign.

"Given the pace of technology I propose we leave the math to the machines and go play outside"

Calvin, Bill Watterson's creation

7 Conclusions

In this thesis we have presented a series of works that aim to improve Millimeter-wave (mmWave) systems. We have correctly identified the intrinsic challenges and limitations of higher frequencies, namely that due to the high spatial attenuation we need to do directional communications. This, in turn, has its own challenges such as initial access, link tracking, blockage detection, etc. This needs substantial changes from sub-6 GHz telecommunication, there is a need for new Medium Access Control (MAC) and Physical (PHY) mechanisms. All this new mechanisms are introduced by the IEEE 802.11ad amendment and they incur in high overheads.

In Chapter 2 we give an overview of the new mechanisms introduced by the IEEE 802.11ad amendment and what challenges they introduce. Chapter 3 takes a look at how the new characteristics of mmWave affect the performance of legacy TCP protocols. TCP cannot distinguish between congestions caused by packet loss or by channel degradation, therefore we do an extensive study of many TCP congestion protocols and define a set of parameters that improve the quality of mmWave communications over TCP even for mobile scenarios. Then in Chapter 4 we look at the existing correlation of mmWave bands with sub-6 GHz bands. We exploit that correlation and propose an algorithm to reduce the initial access in mmWave by using sub-6 GHz information. This reduces the initial access overhead to one third compared to the IEEE 802.11ad standard. Then in Chapter 5 we look at reducing the overhead of beam tracking. In order to do that we implement an open platform based on Field Programmable Gate Arrays (FPGA) that allows us to implement custom protocols in the 60 GHz band. We use this platform to implement an algorithm that completely removes the need to beam train on the Station (STA) side, by changing beam patterns while the Access Point (AP) is sending the packet preamble. Lastly, in Chapter 6 we present a multiband location system implemented in Commercial-Off-The-Shelf (COTS) devices. We manage to reverse engineer some MikroTik 60 wAP devices, develop a custom Linux port and enable the Channel State Information (CSI) and Fine Timing Measurements (FTM) extraction. Then, in conjunction with a sub-6 GHz system we develop a location system that has a median error of 17 cm for an indoor

scenario, even in Non Line Of Sight (NLOS).

Overall, my contributions to the state-of-the-art are the following:

- An overview of how TCP protocols behave over mmWave links and which parameters need to be changed to maximize channel efficiency.
- A comprehensive study of multiband systems showing the existing correlations, applied to initial access overhead reduction.
- A system to reduce the beam training overhead, compatible with standard IEEE 802.11ad devices.
- A multiband location system based in COTS devices. Which allows further reduction of the beam training by enabling simultaneous location and communication.
- A custom Linux fork to enable mmWave experimentation in COTS devices, allowing the extraction of CSI and FTM for the first time.

Appendices

Table 1 shows the parameters used to start an FTM session on the MikroTik router. These are the only parameter values that the QCA9500 firmware accepts. Our wil6210 driver modification has those parameters hardcoded in order to avoid user errors. The resulting values are the T1, T2, T3 and T4 from the FTM algorithm in picoseconds, as explained in Section 6.6.

Parameter	**Value**	**Comment**
session_cookie	0	
aoa_type	0	Phase only
n_peers	1	Only one peer
peers[0] parameters		
freq	60480	
mac_addr	-	Peer MAC address
flag	0x2	ASAP measurement
secure_token_id	0	
aoa_burst_period	0	Only one burst
peers[0].params parameters		
meas_per_burst	4	Only 4 measurements
burst_period	0	Only one burst
num_of_bursts_exp	0	Only one burst
burst_duration	2	250 μs

Table 1: Parameters to start a 60 GHz FTM session

References

[1] P. J. Mateo, C. Fiandrino, and J. Widmer, "Analysis of TCP performance in 5G mm-Wave mobile networks," *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 2019.

[2] J. O. Lacruz, D. Garcia, P. J. Mateo, J. Palacios, and J. Widmer, "mm-FLEX: An Open Platform for Millimeter-Wave Mobile Full-Bandwidth Experimentation," *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020.

[3] P. J. Mateo, A. B. Pizarro, N. Ludant, M. M. Borelli, A. Garcia-Garcia, A. Loch, Z. Shi, Y. Wang, and J. Widmer, "A comprehensive study of low frequency and high frequency channel correlation," *2019 International Conference on Computing, Networking and Communications (ICNC)*, 2019.

[4] C. Andrés Ramiro, C. Fiandrino, A. Blanco Pizarro, P. Jiménez Mateo, N. Ludant, and J. Widmer, "OpenLEON: An End-to-End Emulator from the Edge Data Center to the Mobile User," *Proceedings of the 12th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, 2018.

[5] A. Blanco, N. Ludant, P. J. Mateo, Z. Shi, Y. Wang, and J. Widmer, "Performance Evaluation of Single Base Station ToA-AoA Localization in an LTE Testbed," *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2019.

[6] C. Fiandrino, A. Blanco Pizarro, P. Jiménez Mateo, C. Andrés Ramiro, N. Ludant, and J. Widmer, "OpenLEON: An End-to-End Emulator from the Edge Data Center to the Mobile User," *Computer Communications*, 2019.

[7] D. Garcia, J. O. Lacruz, P. Jiménez Mateo, and J. Widmer, "POLAR: Passive Object localization with Ieee 802.11ad using phased antenna arrays," *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020.

[8] J. O. Lacruz, D. Garcia, P. Jimenez, J. Palacios, and J. Widmer, "High-speed

millimeter-wave mobile experimentation on software-defined radios," *GetMobile: Mobile Computing and Communications*, no. 4, 2021.

[9] M. Rea, D. Giustiniano, P. Jiménez Mateo, Y. Lizarribar, and J. Widmer, "Beam searching for mmWave networks with sub-6GHz WiFi and inertial sensors inputs: An experimental study," *Computer Networks*, 2021.

[10] A. Loch, A. Asadi, G. H. Sim, J. Widmer, and M. Hollick, "mm-Wave on wheels: Practical 60 GHz vehicular communication without beam training," *2017 9th International Conference on Communication Systems and Networks (COMSNETS)*, 2017.

[11] A. Loch, H. Assasa, J. Palacios, J. Widmer, H. Suys, and B. Debaillie, "Zero Overhead Device Tracking in 60 GHz Wireless Networks Using Multi-Lobe Beam Patterns," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '17. New York, NY, USA: ACM, 2017. [Online]. Available: http://doi.acm.org/10.1145/3143361.3143395

[12] "Millimeter Wave Propagation: Spectrum Management Implications," https://transition.fcc.gov/Bureaus/Engineering_Technology/Documents/bulletins/oet70/oet70a.pdf, accessed: 2021-08-23.

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

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

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

[16] Y. Ghasempour, C. R. C. M. Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: 60 GHz Communication for 100 Gb / s Wi-Fi," *IEEE Communications Magazine*, 2017. [Online]. Available: http://ieeexplore.ieee.org/abstract/document/8088544/

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

[18] M. Zhang, M. Mezzavilla, R. Ford, S. Rangan, S. Panwar, E. Mellios, D. Kong, A. Nix, and M. Zorzi, "Transport layer performance in 5G mmWave cellular," in *Proc. IEEE INFOCOM WKSHPS*, Apr 2016, pp. 730–735.

[19] M. Polese, R. Jana, and M. Zorzi, "TCP and MP-TCP in 5G mmwave networks," *IEEE Internet Computing*, vol. 21, no. 5, pp. 12–19, Sep 2017.

[20] T. Azzino, M. Drago, M. Polese, A. Zanella, and M. Zorzi, "X-TCP: a cross layer approach for TCP uplink flows in mmwave networks," in *Proc. Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net)*, Jun 2017, pp. 1–6.

[21] J. F. Kurose and K. W. Ross, "Computer networking: a top-down approach," *Addison Wesley Computing*, 2013.

[22] Y. Zaki, T. Pötsch, J. Chen, L. Subramanian, and C. Görg, "Adaptive congestion control for unpredictable cellular networks," in *Proc. ACM SIGCOMM*, 2015, pp. 509–522.

[23] S. Alfredsson, G. D. Giudice, J. Garcia, A. Brunstrom, L. D. Cicco, and S. Mascolo, "Impact of TCP congestion control on bufferbloat in cellular networks," in *Proc. IEEE WoWMoM*, Jun 2013, pp. 1–7.

[24] A. Afanasyev, N. Tilley, P. Reiher, and L. Kleinrock, "Host-to-host congestion control for TCP," *IEEE Communications Surveys Tutorials*, vol. 12, no. 3, pp. 304–342, Third Quarter 2010.

[25] H. Shokri-Ghadikolaei, C. Fischione, G. Fodor, P. Popovski, and M. Zorzi, "Millimeter wave cellular networks: A MAC layer perspective," *IEEE Trans. on Communications*, vol. 63, no. 10, pp. 3437–3458, Oct 2015.

[26] R. Kumar, A. Francini, S. Panwar, and S. Sharma, "Dynamic control of RLC buffer size for latency minimization in mobile RAN," in *Proc. IEEE WCNC*, Apr 2018.

[27] M. Mezzavilla, M. Zhang, M. Polese, R. Ford, S. Dutta, S. Rangan, and M. Zorzi, "End-to-end simulation of 5G mmwave networks," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2237–2263, Third Quarter 2018.

[28] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, "Improved handover through dual connectivity in 5G mmwave mobile networks," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 9, pp. 2069–2084, Sep 2017.

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

[30] O. Abari, H. Hassanieh, M. Rodriguez, and D. Katabi, "Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections," in *The 15th ACM Workshop*, 11 2016.

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

[32] M. Peter *et al.*, "Measurement campaigns and initial channel models for preferred suitable frequency ranges," 03 2016.

[33] P. Kyösti, I. Carton, A. Karstensen, W. Fan, and G. F. Pedersen, "Frequency dependency of channel parameters in urban los scenario for mmwave communications," in *10th European Conference on Antennas and Propagation (EuCAP)*, 2016, pp. 1–5.

[34] A. O. Kaya, D. Calin, and H. Viswanathan, "28 GHz and 3.5 GHz wireless channels: Fading, delay and angular dispersion," in *IEEE Global Communications Conference (GLOBECOM)*, 2016, pp. 1–7.

[35] A. Ali, N. González-Prelcic, and R. W. Heath, "Estimating millimeter wave channels using out-of-band measurements," in *Information Theory and Applications Workshop (ITA), 2016.* IEEE, 2016, pp. 1–6.

[36] N. Gonzalez-Prelcic, A. Ali, V. Va, and R. W. Heath, "Millimeter-wave communication with out-of-band information," *IEEE Communications Magazine*, vol. 55, no. 12, DECEMBER 2017.

[37] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *IEEE Conference on Computer Communications (INFOCOM).* IEEE, 2015, pp. 2416–2424.

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

[39] R. J. Weiler, M. Peter, T. Kühne, M. Wisotzki, and W. Keusgen, "Simultaneous millimeter-wave multi-band channel sounding in an urban access scenario," in *9th European Conference on Antennas and Propagation (EuCAP)*, May 2015.

[40] G. Lovnes, J. J. Reis, and R. H. Raekken, "Channel sounding measurements at 59 GHz in city streets," in *5th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Wireless Networks - Catching the Mobile Future.*, Sep 1994, pp. 496–500 vol.2.

[41] H. Xu, T. S. Rappaport, R. J. Boyle, and J. H. Schaffner, "Measurements and models for 38 GHz point-to-multipoint radiowave propagation," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 310–321, March 2000.

[42] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[43] J. Ryan, G. R. MacCartney, and T. S. Rappaport, "Indoor office wideband penetration loss measurements at 73 GHz," in *IEEE International Conference on Communications Workshops (ICC Workshops)*, May 2017, pp. 228–233.

[44] G. R. MacCartney, S. Deng, S. Sun, and T. S. Rappaport, "Millimeter-wave human blockage at 73 GHz with a simple double knife-edge diffraction model and extension for directional antennas," in *IEEE 84th Vehicular Technology Conference (VTC-Fall)*, Sept 2016, pp. 1–6.

[45] H. Zhao, R. Mayzus, S. Sun, M. Samimi, J. K. Schulz, Y. Azar, K. Wang, G. N. Wong, F. Gutierrez, and T. S. Rappaport, "28 GHz millimeter wave cellular communication measurements for reflection and penetration loss in and around buildings in new york city," in *IEEE International Conference on Communications (ICC)*, June 2013, pp. 5163–5167.

[46] P. Kyösti, J. Meinilä, L. Hentila, X. Zhao, T. Jämsä, C. Schneider, M. Narandzi'c, M. Milojevi'c, A. Hong, J. Ylitalo, V.-M. Holappa, M. Alatossava, R. Bultitude, Y. Jong, and T. Rautiainen, "IST-4-027756 WINNER II D1.1.2 v1.2 WINNER II channel models," 02 2008.

[47] H. L. Van Trees, *Optimum array processing: Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 2004.

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

[49] J. Palacios, P. Casari, H. Assasa, and J. Widmer, "LEAP: Location Estimation and Predictive Handover with Consumer-Grade mmWave Devices," *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, 2019.

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

[51] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti, "Spotfi: Decimeter level localization using wifi," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, 2015, pp. 269–282.

[52] J. Xiong and K. Jamieson, "Arraytrack: A fine-grained indoor location system," in *10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*, 2013, pp. 71–84.

[53] Y. Xie, J. Xiong, M. Li, and K. Jamieson, "mD-Track: Leveraging Multi-Dimensionality in Passive Indoor Wi-Fi Tracking," *The 25th Annual International Conference on Mobile Computing and Networking*, 2019.

[54] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: Mm-Wave beam steering without in-band measurement," *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015.

[55] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Wireless Networks," *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, 2017.

[56] N. Gonzalez-Prelcic, A. Ali, V. Va, and R. W. Heath, "Millimeter-Wave Communication with Out-of-Band Information," *IEEE Communications Magazine*, vol. 55, no. 12, 2017.

[57] M. Rea, D. Giustiniano, G. Bielsa, D. De Donno, and J. Widmer, "Beam Search Strategy for MillimeterWave Networks with Out-of-Band Input Data," *2020 Mediterranean Communication and Computer Networking Conference (MedComNet)*, 2020.

[58] N. Maletic, V. Sark, J. Gutiérrez, and E. Grass, "Device localization using mmWave ranging with sub-6-assisted angle of arrival estimation," in *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*. IEEE, 2018, pp. 1–6.

[59] P. Jiménez Mateo, A. Blanco, F. Gringoli, and J. Widmer, "Augmenting mmWave Localization Accuracy Through Sub-6 GHz on Off-the-Shelf Devices," *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies, CoNEXT 2021*, 2021.

[60] R. Zhao, T. Woodford, T. Wei, K. Qian, and X. Zhang, "M-cube: A millimeter-wave massive MIMO software radio," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020.

[61] M. Rea, T. E. Abrudan, D. Giustiniano, H. Claussen, and V.-M. Kolmonen, "Smartphone positioning with radio measurements from a single wifi access point," in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, 2019, pp. 200–206.

[62] F. Gringoli, M. Schulz, J. Link, and M. Hollick, "Free your CSI: A channel state information extraction platform for modern Wi-Fi chipsets," in *Proceedings of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, 2019, pp. 21–28.

[63] M. Ibrahim, H. Liu, M. Jawahar, V. Nguyen, M. Gruteser, R. Howard, B. Yu, and F. Bai, "Verification: Accuracy evaluation of WiFi fine time measurements on an open platform," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 417–427.

[64] Many, "IEEE Approved Draft Standard for Information technology–Telecommunications and information exchange between systems - Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE P802.11-REVmc/D8.0, August 2016*, pp. 1–3774, 2016.

[65] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, M. Hollick, and et al., "Fast and Infuriating: Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware," *2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2018.

[66] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices," *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, 2017.

[67] J. Palacios, D. Steinmetzer, A. Loch, M. Hollick, and J. Widmer, "Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices," *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018.

[68] P. J. Mateo, A. B. Pizarro, N. Ludant, M. M. Borelli, A. Garcia-Garcia, A. Loch, Z. Shi, Y. Wang, and J. Widmer, "A Comprehensive Study of Low Frequency and High Frequency Channel Correlation," *2019 International Conference on Computing, Networking and Communications (ICNC)*, 2019.

[69] A. Ali, N. Gonzalez-Prelcic, and R. W. Heath, "Estimating millimeter wave channels using out-of-band measurements," *2016 Information Theory and Applications Workshop (ITA)*, 2016.

[70] M. Alrabeiah and A. Alkhateeb, "Deep learning for mmwave beam and blockage prediction using sub-6 ghz channels," *IEEE Transactions on Communications*, vol. 68, no. 9, 2020.

[71] E. M. Mohamed, M. A. Abdelghany, and M. Zareei, "An Efficient Paradigm for Multiband WiGig D2D Networks," *IEEE Access*, vol. 7, 2019.

[72] L. Yan, X. Fang, X. Wang, and B. Ai, "AI-Enabled Sub-6-GHz and mm-Wave Hybrid Communications: Considerations for Use with Future HSR Wireless Systems," *IEEE Vehicular Technology Magazine*, vol. 15, no. 3, pp. 59–67, 2020.

[73] E. Rastorgueva-Foi, O. Galinina, M. Costa, M. Koivisto, J. Talvitie, S. Andreev, and M. Valkama, "Networking and positioning co-design in multi-connectivity industrial mmw systems," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, 2020.

[74] E. S. Lohan, M. Koivisto, O. Galinina, S. Andreev, A. Tölli, G. Destino, M. Costa, K. Leppänen, Y. Koucheryavy, and M. Valkama, "Benefits of positioning-aided communication technology in high-frequency industrial iot," *IEEE Communications Magazine*, vol. 56, no. 12, 2018.

[75] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan, "OpenMili: A 60 GHz software radio platform with a reconfigurable phased-array antenna," in *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, 2016, pp. 162–175.

[76] S. K. Saha, Y. Ghasempour, M. K. Haider, T. Siddiqui, P. De Melo, N. Somanchi, L. Zakrajsek, A. Singh, R. Shyamsunder, O. Torres *et al.*, "X60: A programmable testbed for wideband 60 ghz wlans with phased arrays," *Computer Communications*, vol. 133, 2019.

[77] J. O. Lacruz, R. Ruiz Ortiz, and J. Widmer, "A Real-Time Experimentation Platform for sub-6 GHz and Millimeter-Wave MIMO Systems," *The 19th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys 2021)*, 2021.

[78] A. B. Pizarro, J. P. Beltran, M. Cominelli, F. Gringoli, and J. Widmer, "Accurate ubiquitous localization with off-the-shelf IEEE 802.11ac devices," *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021.

[79] D. Vasisht, S. Kumar, and D. Katabi, "Decimeter-level localization with a single WiFi access point," *NSDI'16: Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, Mar 2016.

[80] J. Palacios, P. Casari, and J. Widmer, "JADE: Zero-knowledge device localization and environment mapping for millimeter wave systems," *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, 2017.

[81] F. Ricciato, S. Sciancalepore, F. Gringoli, N. Facchi, and G. Boggia, "Position and velocity estimation of a non-cooperative source from asynchronous packet arrival time measurements," *IEEE Transactions on Mobile Computing*, vol. 17, no. 9, 2018.

[82] J. Chen, D. Steinmetzer, J. Classen, E. Knightly, and M. Hollick, "Pseudo Lateration: Millimeter-Wave Localization Using a Single RF Chain," *2017 IEEE Wireless Communications and Networking Conference (WCNC)*, 2017.