

RICE UNIVERSITY

ion Wireless Systems for Joint Communication and Sensing in Millimeter-Wave and Terahe

By

Yasaman Ghasem Pour

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

APPROVED, THESIS COMMITTEE



Edward Knightly

Professor of Electrical and Computer Engineering, Rice University, Committee Chair


Ashutosh Sabharwal (Mar '50, 2020)

Ashutosh Sabharwal

Professor of Electrical and Computer Engineering, Rice University



LIN ZHONG (Mar '30, 2020)

Lin Zhong

Professor of Electrical and Computer Engineering, Rice University



Nathan Dautenhahn

Professor of Computer Science, Rice University

HOUSTON, TEXAS

May 2020

RICE UNIVERSITY

**Next-Generation Wireless Systems for Joint Communication
and Sensing in Millimeter-Wave and Terahertz Spectrum**

by

Yasaman Ghasem Pour

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Doctor of Philosophy

APPROVED, THESIS COMMITTEE:

Prof. Edward W. Knightly, *Chair*
Electrical and Computer Engineering
Rice University

Prof. Ashutosh Sabharwal
Electrical and Computer Engineering
Rice University

Prof. Lin Zhong
Electrical and Computer Engineering
Rice University

Prof. Nathan Dautenhahn
Computer Science
Rice University

HOUSTON, TEXAS
MAY 2020

ABSTRACT

Next-Generation Wireless Systems for Joint Communication and Sensing in
Millimeter-Wave and Terahertz Spectrum

by

Yasaman Ghasem Pour

The use of millimeter-wave and terahertz spectrum (30 GHz to 1 THz) for wireless communication is rapidly emerging as one of the key paradigms for future (5G and beyond) wireless systems. Millimeter-wave (mmWave) and terahertz (THz) communication have the potential to realize orders of magnitude increase in data rates due to the availability of wide spectral bands. However, the increased propagation loss necessitates directional links introducing new challenges on user mobility, blockage, and scaling to multiple users.

This thesis presents the design, implementation, and experimental evaluation of novel solutions for efficient coordination and adaptation of multi-user directional mobile networks. The key idea is to leverage the unique sensing capabilities of mmWave/THz wireless signals, namely, the ability to access large swath of spectrum flexibly, sparse scattering, and the possibility of directionality in small form factors (i.e., large antenna arrays or high-frequency antenna structures) to enhance directional networking in unprecedented ways.

In this thesis, I present the first-of-its-kind sensing-assisted spatial multiplexing system for the efficient establishment of multi-stream direc-

tional links in multi-user mmWave MIMO networks. I demonstrate how we can leverage channel sparsity, GHz-scale sampling rate, and the pre-knowledge of mm-Wave RF codebook beam patterns to sense dominant paths between two nodes and subsequently predict and mitigate the inter-stream interference by intelligently configuring beams.

Then, I introduce the first single-shot single-antenna motion sensing system in THz wireless networks proposing a novel node architecture that exploits an unexplored structure called a leaky-wave antenna. Our design allows nodes to accurately sense user motion and the environment in nanosecond timescales; hence, proactively adapt their highly directional beams under mobility or blockage. Combined, these innovations address the key challenges of directional networking. This research builds a foundation for joint communication and sensing in mmWave/THz spectrum yielding a paradigm shift in the design and development of future wireless systems.

ACKNOWLEDGEMENTS

First and foremost I would like to extend my deepest gratitude to my advisor, Prof. Edward W. Knightly, for all of his guidance, support, inspiration, and motivation. He always believed in me more than I believed in myself. He is the best mentor I could ask for during this journey, I hope that I can be a role model to my students as he has been for me.

I feel fortunate for having amazing collaborators without whom this dissertation would not look the same. Prof. Daniel Mittleman (at Brown University) enhanced my knowledge of antennas, physics of waveguides, and wave propagation through numerous discussions and brainstorming sessions. I learned a lot about 60 GHz IEEE standards and shortcomings of next-generation commodity hardware from Dr. Carlos Cordeiro (at Intel Co.). Prof. Dimitrios Koutsonikolas (University at Buffalo, SUNY) and Prof. Josep Jornet (at Northeastern University) provided resources for developing a programmable 60 GHz software-defined radio testbed and helped with organizing large-scale measurement campaigns.

I would also like to acknowledge the help and invaluable input of my thesis committee members: Prof. Ashutosh Sabharwal, Prof. Lin Zhong, and Prof. Nathan Dautenhahn. Beyond providing feedback on my thesis, they have been extremely supportive throughout my faculty application process and flexible with scheduling my defense virtually due to the Coronavirus lockdown.

I am profoundly grateful to my colleagues, labmates, and friends in Rice Networks Group (RNG): Kumail, Chia-Yi, Peshal, Vini, Adriana, Sharan, Riccardo, Xu, Oscar, Ryan, Naren, Ahmed, Zhambyl, Furqan, Keerthi, and Maryam. They provided feedback on my on-going work, practice talks, and papers, celebrated with me every success and accompanied me during the ups and downs of this journey. I am also very thankful to our group admin, Dee, for taking care of many letter submissions, paperwork, room reservations, equipment shipment, and more.

Most importantly, words cannot express how grateful I am to my parents (Zahra and Masoud) and my sister (Parisa), whom I have not visited for over 5 years due to my visa constraints and the travel ban. My academic journey would not be possible without their constant love, sacrifice, and support. Finally, I would like to thank my husband and best friend, Hamed, who always reminds me to aim high and dream big. Your love has kept me going through this process. Thank you for being my co-pilot.

*To all girls who dare to dream big and fight against the traditional biases
to turn their dreams into reality*

Previously Published Material

[Chapter 2](#) revises previous publication [1]:

Swetank Kumar Saha*, Yasaman Ghasempour*, Muhammad Kumail Haider*, et al.
X60: A Programmable testbed for Wideband 60 GHz WLANs with Phased Arrays.
[ACM WiNTECH 2017](#). **co-primary authors*

[Chapter 3](#) revises previous publications [2, 3]:

Yasaman Ghasempour and Edward W. Knightly. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. [ACM MobiHoc 2017](#).

Yasaman Ghasempour, Muhammad Kumail Haider, and Edward W. Knightly. Decoupling Beam Steering and User Selection for MU-MIMO 60 GHz WLANs. [IEEE/ACM Transactions on Networking 2018](#).

[Chapter 4](#) revises previous publications [4–6]:

Yasaman Ghasempour, Muhammad Kumail Haider, Carlos Cordeiro, Dimitrios Koutsonikolas, and Edward W. Knightly. Multi-Stream Beam-training for mmWave MIMO Networks. [ACM MobiCom 2018](#).

Yasaman Ghasempour, Muhammad Kumail Haider, Carlos Cordeiro, and Edward W. Knightly. Multi-User Multi-Stream mmWave WLANs with Efficient Path Discovery and Beam Steering. [IEEE Journal on Selected Areas in Communications 2019](#).

Yasaman Ghasempour, Claudio R. C. M. da Silva, Carlos Cordeiro, and Edward W. Knightly. IEEE 802.11ay: Next-generation 60 GHz Communication for 100 Gb/s

Wi-Fi. [IEEE Communications Magazine 2017](#).

[Chapter 5](#) revises previous publications [7, 8]:

Yasaman Ghasempour, Chia-Yi Yeh, Rabi Shrestha, Daniel Mittleman and Edward W. Knightly. Single Shot Single Antenna Path Discovery in THz Networks. [ACM MobiCom 2020](#).

Yasaman Ghasempour, Rabi Shrestha, Aaron Charous, Edward W. Knightly, and Daniel M. Mittleman. Single-Shot Link Discovery for Terahertz Wireless Networks. [Nature Communications 2020](#).

Contents

Abstract	ii
Acknowledgements	iv
Previously Published Material	vii
List of Figures	xv
List of Tables	xv
1 Introduction	1
1.1 Main Challenges of Directional Networking	2
1.2 Sensing-Enhanced Directional Networking in mmWave/THz Spectrum	4
1.3 Systems Developed	5
2 A Programmable Testbed for 60 GHz WLANs with Phased Arrays	9
2.1 X60 Testbed	11
2.2 Measurement Campaign	15
2.3 Experimental Results	18
2.4 Related Work	30
3 User and Beam Selection in Multi-User 60 GHz WLANs	32
3.1 Multi-User Architecture and Protocols	35
3.2 A Framework for Decoupling User Selection and Beam Steering . . .	41
3.3 Benchmarks and Complexity Comparison	48
3.4 Evaluation Setup: Testbeds and Trace Driven Emulation	50

3.5	Experimental Realization of Coupling in User and Beam Selection . . .	54
3.6	Receiver Separation and Array Size	61
3.7	Evaluation of Decoupled User and Beam Selection	66
3.8	Related Work	72
4	Sensing-Assisted Spatial Multiplexing in mmWave Networks	75
4.1	A Primer on IEEE 802.11ay	78
4.2	MUTE: MUlti-stream beam-Training for mm-wavE networks	96
4.3	MUTE Implementation	110
4.4	Benchmarking Algorithms	112
4.5	Experimental Results	113
4.6	Related Work	125
5	Single-Shot Single-Antenna Motion Sensing in THz Networks	128
5.1	A Primer on Leaky-Wave Antenna	131
5.2	System Architecture	132
5.3	THz Rainbow: A Unique Pattern	134
5.4	Spatial-Spectral Signatures of THz Rainbow	136
5.5	Path Sensing Module	139
5.6	Rotation Sensing Module	144
5.7	Experimental Platform	149
5.8	Experimental Results on Path Discovery	151
5.9	Experimental Results on Rotation Tracking	168
5.10	Related Work	173
6	Concluding Remarks	176
	References	178

List of Figures

2.1	Antenna array beam patterns.	12
2.2	Maps of four indoor environments and measurement locations included in our study.	17
2.3	Beam-pair heatmap for Lab Position C	19
2.4	Richness of strong beam pairs in four environments.	21
2.5	Beam-pair heatmaps for all 25×25 beam pair combinations for $5m$, $15m$, $20m$ and $25m$ distances in the corridor.	22
2.6	Average loss in SNR vs. BID for corridor positions.	23
2.7	Beam-pair heatmaps for four conference room positions.	25
2.8	(a) Average SNR, and (b) Average Throughput received in positions 1 to 10 in the conference room.	26
2.10	Impact of receiver rotation on SNR for various lobby positions.	29
3.1	Hybrid beamforming architecture of AP and users.	36
3.2	Single-Shot (S^2) User Selection in the context of analog and digital beamforming.	43
3.3	Illustrative scenario for S^2 -MAS: one AP with 16 beam patterns and four users.	45
3.4	Interference-aware Incremental (I^2) User Selection in the context of analog and digital beamforming.	46
3.5	X60 node for wideband 60 GHz WLANs with a 12-element phased array antenna.	52
3.6	WARP-based 60 GHz hardware blocks with horn antennas.	53

3.7	(a) The measurement scenario; (b) The beam pattern selected for R_1 in the SUT phase; (c) The Signal to Interference ratio (SIR) at R_2 when TX is transmitting to both R_1 and R_2 under two scenarios: (1) Analog beams are selected independent from user grouping; (2) Analog beams are selected such that the total SIR at R_1 and R_2 is maximum.	55
3.8	The Signal to Interference ratio (SIR) at R_2 when grouped with R_1 for a two-user simultaneous transmission.	59
3.9	Sum-capacity as a function of receiver separation.	62
3.10	SINR at R_1 when grouped with R_2 , which is placed in five different locations.	65
3.11	(a) Achievable sum capacity vs. number of AP's RF chains, (b) Achievable aggregate rate (including user selection overhead) vs. number of AP's RF chains.	67
3.12	Normalized capacity of S^2 and I^2 policies under LOS and NLOS link connectivities with a four-RF chain AP.	71
3.13	The sum capacity and group size of Exhaustive Decoupled search as a function of LOS probability.	72
4.1	IEEE 802.11ay packet structure.	82
4.2	EDMG TRN field structure for (a) EDMG BRP-RX packets, (b) EDMG BRP-TX packets, and (c) EDMG BRP-RX/TX packets.	84
4.3	IEEE 802.11ay beacon interval structure.	86
4.4	Example of IEEE 802.11ay multiple channel access via scheduling.	88
4.5	(a) The SISO phase of SU-MIMO beamforming, (b) The MIMO phase of SU-MIMO beamforming.	91
4.6	The MU-MIMO Beamforming (a) The SISO phase, (b) The MIMO phase.	94
4.7	Irregular beam pattern examples from X60 platform [1].	97
4.8	MUTE system architecture.	98
4.9	An example scenario with 3 dominant physical paths between the AP and client.	103
4.10	(a) An example irregular beam pattern and its average directivity, (b) <i>score</i> (θ) for path 1, and (c) <i>score</i> (θ) for path 2 in Fig. 4.9.	105
4.11	Candidate selection example.	108
4.12	The X60 platform for 60 GHz band.	111

4.13	Experimental floorplan. Square boxes represent client positions. . . .	114
4.14	Aggregate PHY rate of a two-user MIMO transmission to R_1 (fixed at position index 1) and R_2 when placed at other 11 positions.	116
4.15	Multiplexing gain of 2×2 single-user MIMO as a function of client position.	117
4.16	The SNR heatmaps for all 25×25 beam-pairs for client position indices 3, 6, 9 and 12 under LOS connectivity (top row) and blockage (bottom row).	119
4.17	LOS blockage with a wooden table.	120
4.18	Aggregate PHY rate of baseline scheme as a function of candidate set size in 2×2 single-user MIMO.	122
4.19	The number of candidate beams by MUTE in each client's location. . .	123
4.20	The aggregate PHY rate as a function of the number of spatial streams.	125
5.1	A Leaky-Wave Antenna in which the input signal emerges from the slot such that the maximum emission angle ϕ_f is related to input frequency via Eq. (5.1). The larger the input frequency, the lower the emission angle relative to the plate's axis.	132
5.2	The control plane architecture consisting of a LWA, a broadband transceiver, and a broadband RX.	133
5.3	Creation of a THz Rainbow by exciting a LWA device with a THz pulse.	135
5.4	The end-to-end channel between a broadband source and a broadband RX.	137
5.5	A schematic that illustrates the rotation of the RX waveguide and defines various parameters used in the discussion.	145
5.6	The key components of our rotation estimation model. We estimate the amount of rotation by tracking the changes in the spectral width of the received spectrum. A CW rotation causes a decrease in the high-frequency edge of received spectrum while a CCW rotation results in an increase of the low-frequency edge.	148
5.7	(a) Our custom 4×4 cm^2 LWA; (b) The ultra-broadband transceiver system and its generated THz pulse in time and frequency domains. . .	149
5.8	Our experimental setup.	150
5.9	The schematic setup for THz Rainbow characterization.	152
5.10	The LWA angle-frequency coupling, theory vs. experiment.	152

5.11	Spectrum-Angle heatmap of THz Rainbow.	153
5.12	Validating the phenomenological model of spectral-spatial behavior in THz Rainbow.	155
5.13	The baseline scheme that directly maps peaks in power spectrum to corresponding angles.	157
5.14	The histogram of estimation error in single-path settings.	158
5.15	The estimation accuracy of OPERA in single-path settings.	160
5.16	The schematic of our multipath experimental setup.	161
5.17	The estimation error of the reflected path as a function of inter-path angular separation.	162
5.18	The box plot of estimation error for more than 300 independent multipath configurations.	163
5.19	(a) Example directional radiation patterns generated by LWA; (b) Our mobility setup.	165
5.20	The SNR performance of beam adaptation when path directions are provided by OPERA.	166
5.21	The impact of OPERA's path tracking rate on the accuracy of beam adaptation.	167
5.22	The spectrum of emitted radiation vs. emission angle. a) A plot of the measured spectrum of the radiation emitted by a LWG, after excitation with a broadband input; b) Two different predictions of the spectrum-angle relation by diffraction model and ray optics.	169
5.23	(a) Rotation experimental setup; (b) Characterization of client rotation and its impact on the received spectral edges for two different values of $\phi_0 = 25^\circ, 40^\circ$	170
5.24	Evaluation of the accuracy of proposed rotation estimation technique over a range of CW and CCW measurements.	171
5.25	Maximum range of detectable rotation angles.	172

List of Tables

2.1	List of parameters logged during measurements	16
3.1	Comparing complexity or search space size of different policies.	47
3.2	<i>Exhaustive Decoupled</i> vs. <i>Exhaustive Joint</i>	60

Introduction

The ever increasing demand for ultra-fast and ultra-low-latency wireless connectivity to support emerging applications such as virtual reality and autonomous vehicles is straining the capacity of current WiFi and cellular networks [9]. Today, mmWave (30 GHz to 100 GHz) and THz spectrum (frequencies beyond 100 GHz to 1 THz) is emerging as the most promising solution due to the availability of wide spectral bands that can potentially realize data rates of up to 100 Gb/sec [6]. The use of mmWave and THz frequencies is one of the widely accepted paradigms for future (5G and beyond) wireless systems [10–12]. Even though these high-frequency systems offer numerous advantages, increased propagation loss necessitates directional communication making it fundamentally different from today’s wireless networks. Specifically, the directional nature of links brings in many new challenges in the design of next-generation wireless systems; namely, mobility, blockage, and scaling to multiple users.

This thesis explores the use of wireless mmWave/THz signals for sensing user motion and the environment to enhance directional communication in unprecedented ways. Specifically, I design systems that leverage wide spectral bands, the large antenna arrays (i.e., sensing modules), and novel high-frequency antenna structures

to extract a variety of sensing information from mmWave/THz signals traveling in indoor environments on their route toward the receiver. In particular, I show how to 1) accurately sense user motion (translation and rotation) in nanoseconds with a transmission of a wideband signal using a single high-frequency antenna structure, 2) capture sparse reflections from the environment under different node architectures, and 3) augment information across multiple directional beams and users in multi-user networks.

Subsequently, I design and experimentally evaluate solutions that exploit such information to enhance directional networking including proactive mobility adaptation, fast blockage recovery, and efficient multi-stream simultaneous transmissions in multi-user mobile mmWave and THz wireless networks. One key advantage of my solution is that it provides ubiquitous sensing by repurposing communication-dedicated hardware and wireless signals. Hence, sensing-assisted directional networking is achievable whenever and wherever we wish to establish communication links.

Finally, even though in this dissertation I focus on the sensing capabilities that are important for the control and management of directional wireless networks, my techniques can generalize to high-resolution and ubiquitous sensing of the entire environment using wideband mmWave/THz wireless signals. I believe this thesis is the first effort toward the new paradigm of unified wireless communication and sensing in the next-generation wireless system in mmWave and THz spectrum.

1.1 Main Challenges of Directional Networking

Mobility. One key challenge in exploiting directional links is that end nodes need to continually align their beams to maintain connectivity. To this end, the state-of-the-art IEEE standards for 60 GHz WLANs such as 802.11ad [13, 14] and 802.11ay [6] employ a periodic exhaustive “trial-and-error” scan, in which one node sends training

frames sequentially across all its beams while the other node adopts a pseudo-omni reception pattern to identify the strongest transmit beam. This training, when repeated at both ends, discovers the best transmit-receive beam pair for the trained nodes. However, the time overhead associated with such exhaustive training severely degrades throughput and disrupts low-latency applications. The overhead worsens for the THz network lacking the “pseudo-omni” reception of 60 GHz band, increasing the order of the beam-search space from $2N$ to N^2 for N beams at each end. Moreover, due to their inherently more directional nature (narrow beams of only a few degrees [15] as opposed to tens of degrees in mmWave beams [1, 16]), the number of possible sender-receiver beam pairs is significantly increased, rendering the overhead of exhaustive testing infeasible.

More importantly, mobility brings in a significant challenge, as beam alignment may become repeatedly lost, requiring repeated training to maintain beam alignment and incurring overhead each time. Again narrower pencil-shaped beams at THz regime would show higher sensitivity to motion necessitating an ultra-fast mobility detection and management mechanisms to achieve seamless wireless connectivity.

Blockage. The high directionality coupled with poor penetrability makes mmWave and THz transmissions vulnerable to blockage, which significantly degrades the coverage capability and makes communication links susceptible to abrupt failure [17, 18]. Most of our surrounding objects in the environment as well as the human body show poor penetration and thus can block mmWave and THz radiations [19, 20]. However, at these frequencies, most metals, smooth surfaces, and even concrete show good reflection characteristics enabling the receiver to capture secondary reflections for non-line-of-sight (NLOS) communications [21–23]. Yet, an efficient blockage recovery is not possible without techniques to sense and track the reflected paths to be used as backup steering directions in an incident of blockage.

Scalability. Today, the GHz-scale unlicensed spectrum coupled with directional antennas is able to support Gbps data rates in point to point links. Next-generation devices aim to scale to dense user populations via the concurrent transmission of multiple independent data streams, i.e., downlink multi-user MIMO. However, the key challenge to realizing such scalability via mmWave MIMO is to efficiently discover the analog beams at the Access Point (AP) and clients that support concurrent directional transmission/reception of multiple data streams and providing the maximum multiplexing gain. Establishing such multi-stream directional links can be a high overhead procedure as the search space increases with the number of spatial streams and the product of AP-client beam resolution [5, 24]. This overhead can overwhelm the channel and negate the gains of concurrent transmissions.

1.2 Sensing-Enhanced Directional Networking in mmWave/THz Spectrum

This research advances mmWave and THz Networking through sensing the user motion and the environment. In particular, I discuss how to sense user motion accurately and in real-time so that we can proactively and efficiently adapt highly directional beams at both ends; thus, completely eliminating the requirement of high-overhead beam training at mobile devices and at wireless APs. Moreover, I provide solutions on how to sense reflectors in the environment so that the mobile user and the AP can automatically switch to a reflected path for blockage recovery. Finally, by sensing spatial paths captured across different users and beams, I develop systems for efficient user and beam selection, addressing the challenge of scalability.

Sensing motion and the environment is long-standing research and various techniques have been proposed in the past, including ultrasound, radar, vision techniques,

and RF sensing with 2.4/5GHz WiFi signals [25–28]. However, in this thesis, I propose to leverage mmWave and THz wireless signals themselves as sensing modality. The key advantage of using mmWave/THz signals is that our sensing capability is ubiquitous and does not require additional hardware at wireless communication nodes. Moreover, exploiting the same frequency band for communication and sensing implicitly guarantees that the sensing resolution is proportional to the sensitivity of communication links with respect to motion. Additionally, the reflection coefficients of different materials are often frequency-dependent, which once again emphasizes the advantage of a sensing and communication co-design on the same frequency bands.

However, there are multiple challenges in extracting sensing information from wireless signals. First, as mmWave/THz waves propagate in the environment and bounce off walls, furniture, and different existing objects, they experience significant propagation and reflection loss. Hence the received signals are very weak, i.e., low SNR regime. While RF sensing in sub 6 GHz bands commonly exploit the phase or timing information of received signals, extracting such information is relatively more expensive in these high-frequency regimes requiring low noise devices and tight synchronization between transmitter and receiver which is very challenging, especially in THz bands. Finally, the directional nature of transmissions yields a limited field of view; thus, nodes might need multiple rounds of transmission to capture the entire environment resulting in high-latency in sensing and thus delay in coordination and adaptation.

1.3 Systems Developed

To overcome these challenges, I design systems that exploit the unique characteristics of wireless signals in mmWave and THz spectrum, namely the ability to access large swath of spectrum flexibly, channel sparsity (i.e., availability of a few multipath in the

wireless medium due to high reflection loss), and the possibility of directionality in small form factors (i.e., large antenna arrays or high-frequency antenna structures). Next, I'll summarize the developed systems.

1. A Programmable Testbed for 60 GHz WLANs with Phased Arrays.

Addressing the above challenges, we first need to acquire a deep understanding of how mmWave/THz waves propagate in the medium, the impact of the directional radiation pattern, and wideband channels. This requires designing software-hardware systems that allow fine-grained control over the beam configuration, provides the capability for transmission/reception of wideband signals and supports programmability over PHY and MAC layers for overhead and data-rate analysis.

In Chapter 2, I present X60, the first Software-Defined-Radio (SDR) testbed for 60 GHz WLANs, featuring fully programmable MAC/PHY/Network layers, multi-Gbps rates, and a user-configurable 12-element phased antenna array. Combined these features provide an unprecedented opportunity to re-examine the most important aspects of signal propagation and performance expected from practical 60 GHz systems. Leveraging the testbed's capabilities, we conduct an extensive measurement study, looking at different aspects of indoor 60 GHz links. We find that the presence of reflective surfaces and imperfect beams generated by practical phased arrays together can result in multiple NLOS paths supporting Gbps rates [1, 29].

2. User and Beam Selection in Multi-User 60 GHz WLANs.

I demonstrate the mmWave MIMO system architecture in Chapter 3 and discuss how the achievable rate for a multi-user transmission is impacted by the selected users, analog beam steering parameters, and digital pre-coding weights. In particular, using X60 nodes, I experimentally show that the choice of users and beams are tied together such that the best choice of analog beams for the link between the AP and user u depends on the users that are grouped with user u in the multi-user downlink transmission.

Consequently, I formulate their joint optimization framework and argue that the optimum solution requires prohibitively large training time in practice. Hence, I present frameworks that decouple user selection from beam steering training; thus, enabling low-overhead and low-complexity multi-user structure for mmWave networks [2, 3].

3. Sensing-Assisted Spatial Multiplexing in mmWave Networks. In Chapter 4, I present MUTE, the first system that senses dominant paths (LOS and NLOS) between the AP and each client and leverage that information to maximize spatial multiplexing gains by efficient prediction and mitigation of inter-stream interference in 60 GHz MIMO WLANs. If nodes have ideal steering capability, i.e., narrow pencil-shaped beams, forming simultaneous links is a straightforward extension of single-user communication. However, practical mmWave systems use phased arrays that generate irregular and overlapping radiation patterns. Nonetheless, despite their irregularity, the directivity gain is known a priori in each direction as it is a deterministic function of the codebook and antenna spacing. MUTE leverages channel sparsity, GHz-scale sampling rate, and the knowledge of mm-Wave RF codebook beam patterns to sense spatial paths. Subsequently, exploiting the knowledge of the sensed paths, MUTE predicts, without over-the-air measurements, the best multi-stream beam configuration at the AP and clients that achieves the maximum multiplexing gain [4, 5].

4. Single-Shot Single-Antenna Motion Sensing in THz Networks. In Chapter 5, I introduce OPERA, the first single-shot single-antenna motion sensing platform for THz wireless networks. OPERA tracks user moves in nanosecond time-scales in order to proactively adapt highly directional THz beams to motion and blockage. To this end, I propose a novel THz node architecture in which the AP and clients are equipped with a leaky-wave antenna. This device allows traveling wave to leak from a waveguide into free-space with the unique property that the emission

angle from the waveguide is coupled to the frequency of the input signal.

Leveraging this property, I embed motion information into spectral characteristics of signals. Specifically, I input a THz pulse (equivalently, a broadband signal in frequency-domain) into the leaky-wave antenna to create unique spectral patterns simultaneously at each angle, which I termed a “THz Rainbow”. Based on the physics of EM waves and Maxwell’s Equations, I model the transmission pattern of a THz Rainbow and demonstrate how motion can be estimated based on the spectral characteristics of the received signal. Our experiments show that OPERA achieves average estimates within 2° of ground truth for LOS, reflected paths, and device rotation [7,8].

A Programmable Testbed for 60 GHz WLANs with Phased Arrays

In this chapter, we introduce **X60**, the first highly configurable software defined radio (SDR) 60 GHz testbed, featuring fully *programmable PHY, MAC and Network* layers while still allowing for ultra-wide channels and multi-gigabit data rates. Based on the National Instrument's (NI) mmWave Transceiver System [30] and equipped with a *user-configurable* 12-element phased array antenna from SiBeam, X60 nodes enable communication over 2 GHz wide channels using realistic transmit (TX) and receive (RX) beams that can be steered in real-time.

X60 offers several key advantages over other existing mmWave experimental platforms. Unlike commercial 802.11ad devices, X60 with its SDR/FPGA based architecture allows access to and complete control over the PHY and MAC layers. This not only enables experimentation that can obtain a full view of the often complex interaction among multiple layers of the networking stack, but also allows for prototyping and testing of new techniques at multiple layers. In contrast to most existing SDR mmWave experimental platforms (based on USRP/WARP), X60 provides high reconfigurability without limiting baseband bandwidth to a few hundred MHz, enabling

us to study the impact of extra wide channels supported by the 802.11ad standard. Lastly, using SiBeam’s phased array, X60 generates beam patterns that are configurable and steerable in real time, overcoming a basic limitation of horn-antenna based platforms where the beam can only be steered using a mechanical rotator and may not always be representative of the often imperfect beams generated by phased arrays.

X60’s capabilities provide an opportunity to re-examine the understanding of the most important aspects of 60 GHz WLAN signal propagation and performance. To this end, we undertake an extensive measurement campaign across four characteristic indoor environments (corridor, lab, lobby, and conference room) in a typical academic building. Enabled by the testbed’s reference implementation that uses a slotted TDD based MAC and supports multiple modulation schemes (from BPSK up to 16-QAM), we analyze various MAC performance metrics such as goodput while still having access to the underlying PHY parameters. We further study the implications of steering TX/RX beams along different directions. Our measurements encompass a range of propagation environments (dominant LOS, non-LOS only, reflections from multiple obstacles, LOS propagation with side-lobes) and TX/RX orientations.

Our major findings are as follows: *(i)* Unlike the common belief of only a few TX/RX beam-pairs achieving high SNR in the 60 GHz band, almost 15% of the total 625 possible beam-pair combinations in our setup provide at least 1 Gbps of throughput across all environments. This is primarily due to shape of imperfect beam patterns, overlap of main and side-lobes between neighboring beams, and richness of NLOS paths in the environment. *(ii)* At short range, the sender and receiver can be together off by several beam pair indices while maintaining high SNR, whereas at longer range, they can be off by only one or two indices, as the relative penalty for imperfect beam selection increases with distance. *(iii)* NLOS paths from strong reflectors can support links with comparable signal strength to LOS paths, and achieve

multi-Gbps rates. (iv) Nodal mobility scenarios such as translation and rotation can severely degrade signal strength for a fixed pair of beams being used at TX and RX ends. Further, our analysis of various beam adaptation schemes reveals that in comparison to exhaustive search, adjusting beams at either the TX or the RX side, although sub-optimal, is adequate in most cases to restore the link. Therefore, if one node (e.g., the AP) is adaptive, the other (client) may incorrectly hold on to an older beam too long without necessarily incurring link breakage.

2.1 X60 Testbed

In this section, we describe the different components of the X60 testbed. All the modules are programmed using NI LabVIEW.

2.1.1 Baseband TX/RX

Each X60 node is based on the NI mmWave Transceiver System. All modules involved in the baseband signal generation are assembled inside a NI PXIe-1085 PXI Express chassis. Most of the inter-module signaling and data transfer happens over the chassis' high-speed backplane using FIFO queues or DMA. The TX/RX chains consist of one or more high-performance FPGAs which handle the majority of the transmit/receive operations including encoding/decoding and modulation/demodulation. The FPGA outputs feed into a wideband DAC/ADC module which generates/samples the baseband signal. In addition, the chassis holds a high-end controller (host machine) running Microsoft® Windows 7. The host generates the source bits for transmission and is the sink for the receive operation. It controls different TX/RX parameters (MCS, uplink/downlink, etc.) and collects information about different parts of the TX/RX chain to allow for user-display and debugging.

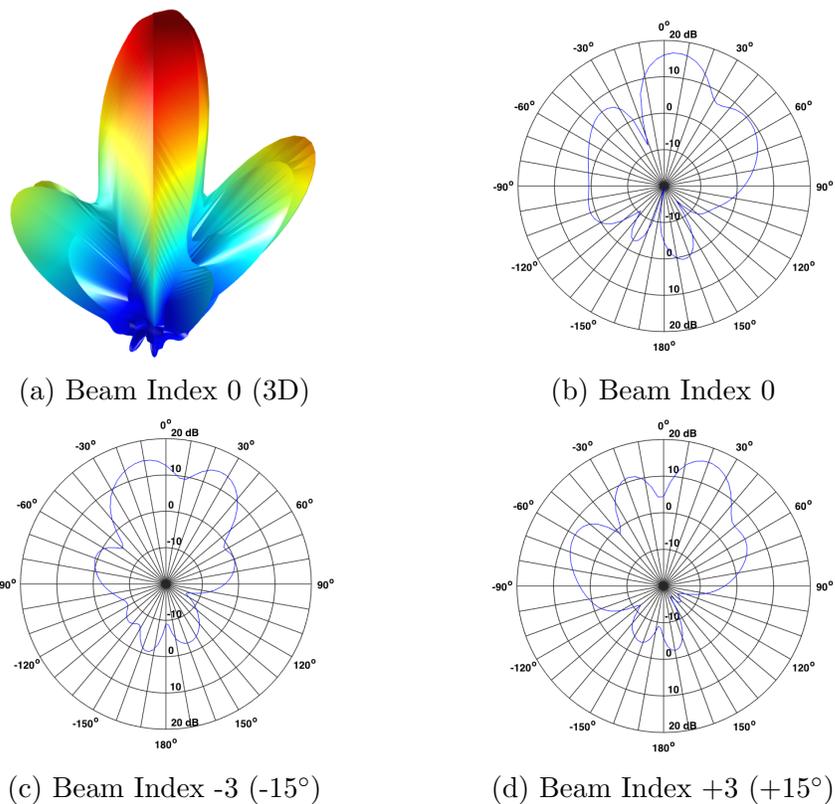


Figure 2.1: Antenna array beam patterns.

2.1.2 PHY/MAC Structure

The current reference PHY implementation allows for the following modulation and coding (Turbo codes) combinations: 1/5 BPSK, 1/4 QPSK, 1/2 QPSK, 3/4 QPSK, 1/2 16QAM, 3/4 16QAM, 7/8 16QAM, resulting in theoretical bit rates from 300 Mbps to 4.75 Gbps. Data transmission takes place in 10 *ms* frames which are divided into 100 *slots* of 100 μ s each. Both the MCS and operation type (uplink/downlink/sync) can be configured on a *per-slot* basis. A slot is made up of 92 codewords (data bit sequence after encoding), each of which has an attached CRC block. At the RX end, the throughput for a given slot is calculated by counting the number of correctly decoded codewords and multiplying by the number of actual data bits contained in each codeword (which depends on the MCS).

2.1.3 Antenna Array and Beam Patterns

The SiBeam mmWave module, on the TX path, takes as input the baseband signal (as differential I/Q), up-converts, and transmits over the air a 2 GHz wide waveform centered around one of the 802.11ad channel center-frequencies. The typical transmit power is 30 dBm EIRP at channel 2. The RX side flow is symmetric to the TX path. The in-built phased array has 24 elements; 12 each for TX and RX. The module connects to the baseband chassis over an additional dedicated control path that allows different phase values for the antenna elements through the use of codebooks. Different beams within a codebook can be switched by applying the required index into the module's registers. The phase of each antenna element can be set to one of four values: 0 , $\pi/2$, π , $3\pi/2$.

SiBeam's reference codebook defines 25 beams spaced roughly 5° apart (in their main lobe's direction). The beams cover a sector of 120° (in the azimuthal plane) centered around the antenna's broadside direction. The 3 dB beamwidth for the beams ranges from 25 to 30 degrees for TX and from 30 to 35 degrees for RX. As a result, each beam's main lobe overlaps with several neighboring beams. We refer to the beams using index range: -12 (-60°) to +12 ($+60^\circ$), with index 0 corresponding to the broadside beam.

We first computed the idealized beam patterns using COMSOL Multi-physics [31]. An antenna array integrated by 12 elements with the same size, separation, and spatial distribution was defined and the input at each antenna was set as per the defined codebook. Figures 2.1a-2.1c depict examples of 2D and 3D radiation patterns for select beam indices. These patterns highlight how, in contrast to beams generated by horn antennas, phased-array generated beams often have strong side-lobes. Moreover, as beams are steered away from the main lobe, patterns become more imperfect with even stronger side lobes and considerably weaker main lobe. For instance, compar-

ing beam index 3 (Fig. 2.1d) and 12 (Fig. 2.1b) shows how practical phased-arrays can have *non-uniform steerability* along different directions as opposed to mechanically rotated horn antenna beams. Surprisingly, beam indices equally apart from the broadside beam (e.g., +3 (Fig. 2.1d) and -3 (Fig. 2.1c) can have radiation patterns that are *not* necessarily mirror images of each other.

The remaining simulated beam patterns also exhibit similar behavior. Further, note that these particular characteristics of the beam patterns result both out of the discretization of the individual antenna element phase weights and the particular geometry that the elements are arranged in the 2D array. Nitsche et al. [32] also found the beam patterns of commercial WiGig devices to be imperfect with strong side lobes. Also, an inspection of the open source wil6210 driver [33], targeting Qualcomm 802.11ad chipsets, suggests 2 bits for phase control of the antenna elements (which allows for 4 possible phase values similar to our platform).

2.1.4 Enhancements for Measurements

We made the following modifications to the reference code to enable logging of all the required PHY/MAC parameters and to allow for more realistic measurements.

Automatic Gain Control (AGC): We implemented an AGC block running on the host machine (every 100 ms) that adjusts the receiver’s gain value based on the energy calculated from the raw I/Q samples to achieve an experimentally determined optimal target energy value that ensures best ADC operation. Through a separate set of experiments, we verified that our implementation is throughput optimal (as compared to exhaustive-search manual gain control) for different MCS and channel conditions.

Thin Control Channel: We added an external legacy WiFi radio to all four nodes to implement a reliable control path. This allows us to implement certain

features like TX-RX beam selection or MCS selection and to automate parts of our measurements with only few modifications to the existing code base, without the burden of maintaining tight timing requirements of the code running the mmWave channel. The scripts that implement this control path run on the host machine and communicate with the LabVIEW process via IPC over TCP to control parameters like MCS and beam index, and collect link metrics for further processing.

Instrumentation: We instrumented the host side LabVIEW code base to log a whole range of different parameters. Since the host is an active part of the TX/RX flow (e.g., fine synchronization operations) and needs to maintain strict timing guarantees in its generator/consumer loops, we selected different logging frequency for each parameter to minimize overhead. Some parameters including Signal Power Estimation, Noise Power Mean, Throughput, and CRC pattern are logged on a per-frame basis (every 10 ms), while others, generated in already computation-heavy timed-loops (such as RSSI, SNR, Carrier-to-Noise, Phase, and Power Delay Profile), are logged at a lower frequency (every 40 ms).

2.2 Measurement Campaign

2.2.1 Methodology

Our measurement campaign is aimed at collecting key PHY and MAC layer parameters across multiple indoor environments. Measurement locations are selected to characterize static 60 GHz channels, as well as emulate typical mobility patterns like translation and rotation to study the impact of nodal mobility.

At each location, we collect channel measurements in two steps. *(i) Beam Sweep:* The first step encompasses channel estimation for all possible beam pairs in an exhaustive search. The transmitter and receiver co-ordinate their beam switching (over

Table 2.1: List of parameters logged during measurements

Parameter	Log Interval
Received Signal Strength (RSSI)	10 ms
Signal to Noise Ratio (SNR)	10 ms
Channel Noise	10 ms
Signal Power Estimate	10 ms
Phase	40 ms
Noise Estimate	40 ms
Power Delay Profile (PDP)	10 ms
SU Channel (Re, Im, Odd, Even)	10 ms
Throughput	40 ms
CRC pattern	40 ms

the control channel) to generate all 625 (25×25) beam-pair combinations. For each beam pair, 25 frames are transmitted at MCS 0 and SNR is logged for the channel estimation slot in each frame (every $40ms$). This is a crude emulation of 802.11ad’s beam-training scheme. However, note that we only evaluate the resultant beams and not the time-efficiency of the process itself. *(ii) MCS sweep:* In the second step, we select a small subset of (TX, RX) beam pairs for which we repeat measurements at all seven achievable MCS levels. We select the three strongest beam pairs out of all 25×25 beam pair combinations based on average SNR computed during beam sweep in step (i). Further, to study the impact of selecting neighboring beams, we also include the immediate neighbors of the RX beam in each of the three (TX, RX) pairs, for a total of nine (TX, RX) beam pairs. For each MCS, we log all channel parameters for 500 frame transmissions. Table 2.1 lists the logged parameters.

2.2.2 Environments

We perform measurements in four different indoor environments in a typical office building. A detailed map of three of these environments is depicted in Fig. 2.2. The arrows at each TX and RX location indicate the physical orientation of the node i.e., the direction along the main lobe of the center beam of the phased array.

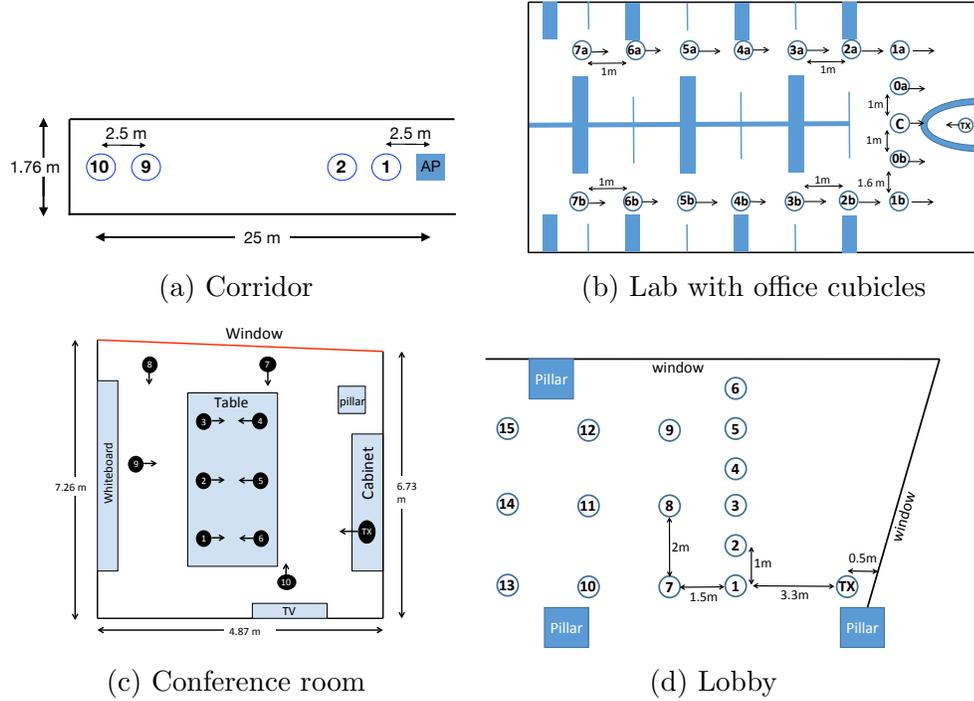


Figure 2.2: Maps of four indoor environments and measurement locations included in our study.

Corridor: The TX is fixed at one end of a 1.74 m wide corridor, at 1.23 m height. We consider 10 receiver positions varying the inter-node distance from 2.5 m to 25 m on a straight line in steps of 2.5 m, such that the RX always faces the TX. Apart from characterizing the static environment, this topology emulates receiver translation along a straight line, as it moves away from the transmitter.

Lab: The second environment is an $11.8 \times 9.2 \times 3.4 \text{ m}^3$ lab with four rows of office cubicles, with cardboard partitions and metal cabinets. The TX is fixed close to the center of the right wall at 2.05 m height, a location where 60 GHz WLAN APs would be typically deployed. We choose 17 locations in three rows for taking measurements. As shown in Fig. 2.2b, row A (locations $1a$ through $7a$) and row B (locations $1b$ through $7b$) are nearly symmetrical about the center partition, while the distance from the right wall increases from 2.9 m to 8.9 m in steps of 1 m. The front row (locations $1b$, $0b$, C , $0a$, $1a$) is selected to study the impact of increased angular

separation between TX and RX. The height of the receiver is 1.26 m at all locations, such that there is always an LOS path between the TX and the RX.

Conference Room: This is a $7 \times 4.87 \times 3.4$ m^3 room (Fig. 2.2c) with a large central table and various metallic/shiny surfaces (e.g., TV, white boards, metal cabinet, chairs, glass windows) which have been shown to be excellent reflectors in 60 GHz band [34, 35]. Hence this environment is suitable to study the impact of reflectors, non-LOS paths and side-lobes. The TX is placed in a corner at 2.23 m height, and we consider 10 different RX locations across the room.

Lobby: This is the largest open space in the building, with large glass panels as walls. To study the impact of increasing distance and angular separation, we fix the TX in one corner of the lobby, and select 15 RX locations in four rows. To study the impact of receiver rotation, we vary its orientation between -90° to 90° in steps of 15° , such that 0° corresponds to RX phased array facing the front wall, parallel to side walls. Hence the central beams of both TX and RX are perfectly aligned for 0° orientation at positions 1, 7, 10 and 13 (Fig. 2.2d).

2.3 Experimental Results

2.3.1 Beam Sweeps

In our measurements, a beam sweep captures the SNR achieved for all 25×25 possible beam pair combinations, each resulting in a distinct channel. As such, it can be used to study the distribution of strong SNR beam pairs and their mapping to the physical environment, and how the signal strength changes across different positions and environments. Therefore, we use beam sweeps as the main tool to understand different characteristics of 60 GHz links.

We represent each beam sweep as a heatmap of corresponding SNR values with

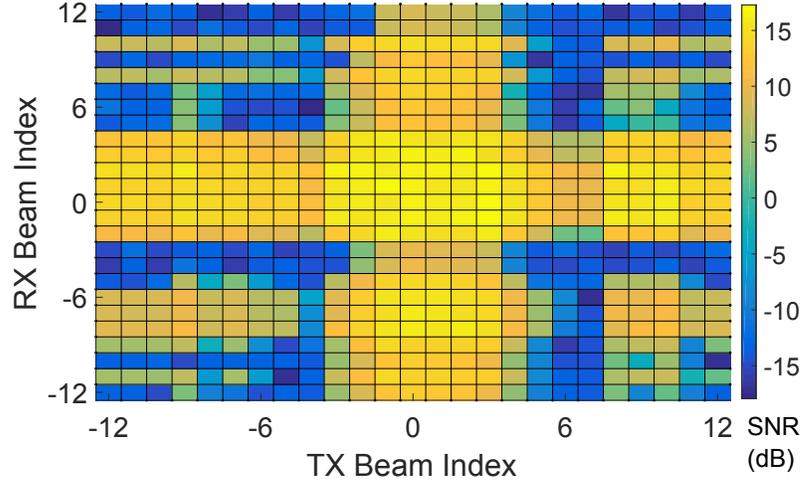


Figure 2.3: Beam-pair heatmap for Lab Position C

TX beam indices along the x-axis and RX beam indices along the y-axis. Fig. 2.3 shows the beam-pair heatmap for the center position (C) in the lab with distance of 2.3 m from TX, with yellow colored regions indicating beam pairs with the SNR of above 10 dB whereas blue regions indicate beam pairs for which the SNR is below the receive threshold (< 0 dB, determined from measurements). The central beam pair (0,0) corresponding to the LOS path between TX and RX achieves the strongest link strength. Due to overlap between neighboring beams (Sec. 2.1.3), multiple beams may include the LOS component, albeit with a different directivity gain. Hence we get a cluster of high SNR beam pairs close to the central pair. Besides the LOS central high SNR region, there are smaller clusters of beam pairs with moderate to high SNR, resulting from reflections and side-lobes. According to Fig. 2.3, for TX beam indices between -2 to 4, the received SNR is above 5 dB regardless of the RX beam choice. The TX and RX are relatively close to each other which makes any RX beam (covering -60° to 60°) to achieve high SNR provided that the TX beam is pointed towards the receiver. Similarly, when RX beam indices between -2 and 5 are used, SNR is above 5 dB for most TX beams.

2.3.2 Richness of Strong Beam Pairs

mmWave channels are expected to be sparse due to the higher path loss and penetration loss [34]. Therefore, we would expect to see only a few physical paths, including LOS and NLOS, between the TX and RX. A particular beam pattern captures a number of these paths and applies different directivity gain. Hence, the provided signal strength by a particular (TX, RX) beam pair depends on the number of captured paths, their link budget, and beam directivity. The isolation of LOS/NLOS paths in order to measure the richness or sparsity of 60 GHz channels is not feasible with our platform; however, in this subsection, we explore the richness of *strong beam pairs*. In particular, we study how many beam pairs can achieve at least 1 Gbps data rate in different environments. From our data set, we found that 10 dB SNR is sufficient for achieving 1 Gbps data rate. Hence, we define a strong beam pair as a beam pair that provides at least 10 dB SNR. We count the number of such strong beam pairs for each position in the corridor, conference room, lab, and lobby (see Fig. 2.2). Fig. 2.4 shows the average (over all measurement positions) ratio of the number of strong beam pairs over the total 625 possible beam combinations for the four environments. E.g., 0.2 in the y-axis means that on average 125 beam combinations (out of a total of 625 beam pairs) can provide at least 1 Gbps data rate.

First, Fig. 2.4 reveals that the fraction of strong beam pairs is highest for the conference room and lowest for the corridor. The materials in the conference room such as metallic cabinet, white board, and glass windows are known to be good reflectors for millimeter waves [34] providing several reflected paths, while there are no strong reflectors in the corridor. Further, the high error bars show that in any given environment, the number of strong beam pairs between two nodes highly depends on the RX position, its relative orientation with the TX, and the distance between them.

Second, the average ratio of strong beam pairs is above 0.13 for all four environ-

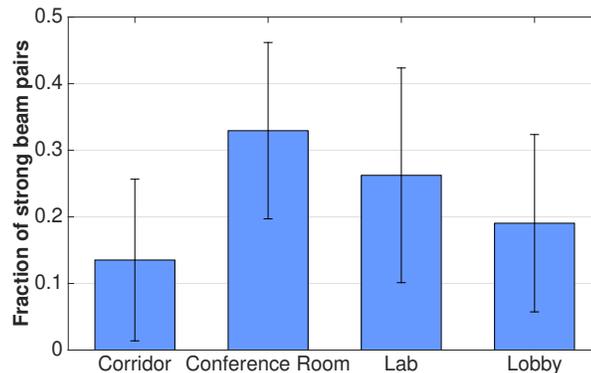


Figure 2.4: Richness of strong beam pairs in four environments.

ments, i.e., more than 80 beam pairs provide at least 1 Gbps of throughput. This shows that, in contrast to the common belief, there are several beam pairs that are able to provide Gbps data rates for 60 GHz communication. This result is caused by the shape of the imperfect beam patterns in use which have side-lobes as well as overlap (Fig. 2.1); thus, a physical LOS/NLOS path can be captured by multiple beams. The richness of strong beam pairs implies that beam training/adaptation algorithms might be able to avoid exhaustive time-consuming search through all beam combinations to find the best beam pair. Another implication is that interference between simultaneous transmissions may not be negligible in 60 GHz.

2.3.3 Relative Strength of Neighboring Beams

In principle, an exhaustive search over all possible beam combinations is required to discover the highest signal strength beam pair. However, the associated training overhead may be prohibitively high, especially in case of mobile links. Hence, it is sometimes desirable to adopt lower overhead strategies which search over a subset of beam pairs. E.g., the 802.11ad standard specifies an initial coarse level search with quasi-omni beams at one end, followed by beam refinement for only a subset of beam pairs. Here, we evaluate the significance of selecting the highest strength beam pairs,

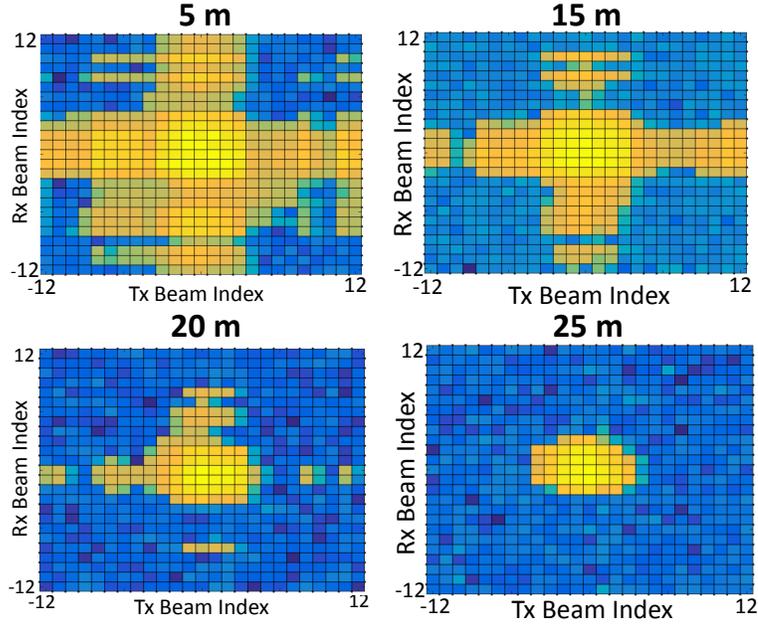


Figure 2.5: Beam-pair heatmaps for all 25×25 beam pair combinations for $5m$, $15m$, $20m$ and $25m$ distances in the corridor.

and quantify the loss in signal strength if a sub-maximal pair is selected during the training process.

In particular, we are interested in studying the impact of distance and multiple paths (from reflections and side-lobes) on the relative strength of neighboring beams. The corridor data set encompasses both these scenarios, as illustrated by beam-pair heatmaps for $5m$, $15m$, $20m$ and $25m$ positions (Fig. 2.5). We make two key observations. *(i)* For the closest position ($5m$), beam 0 and its two nearest neighbors at the TX side achieve (> 10 dB) SNR for all RX-side beams and vice versa. This is due to reflections off of side walls from the narrow corridor and the side lobes, resulting in a strong channel. *(ii)* As the TX-RX distance increases, the high SNR region shrinks and includes only the central beam pairs at $25m$. This is because the impact of reflections from side walls becomes less pronounced with distance. By geometry, the azimuth angle for first-order reflection paths (strongest NLOS components) from either wall reduces from 19° at $5m$ to less than 4° at $25m$. As such, the angular

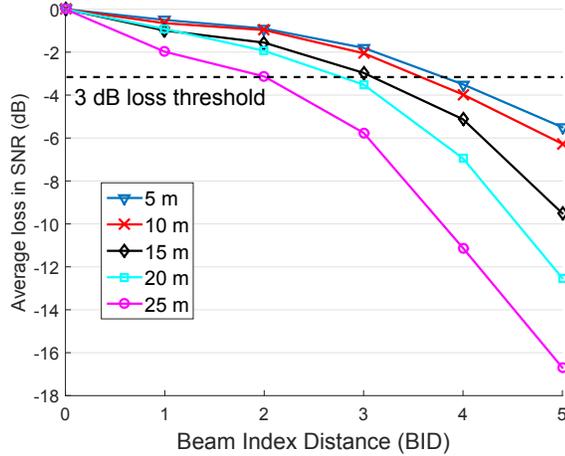


Figure 2.6: Average loss in SNR vs. BID for corridor positions.

separation between LOS and the strongest NLOS component decreases and only the central beams include these paths. Hence, more distant RX positions will require a larger search space to discover high strength beam pairs.

To further quantify the impact of selecting sub-maximal beam pairs, we analyze the loss in SNR as we move away from the highest strength beam pairs. We associate a distance metric with each beam pair (i, j) ; Beam Index Distance (BID). If (T, R) is the beam pair with highest SNR, we define BID as $(|T - i| + |R - j|), \forall i, j \in [-12, 12]$. E.g., $BID=1$ indicates a difference of one beam index, either in TX or in RX beam. Fig. 2.6 shows SNR loss in dB vs. BID for five different positions (at distances 5m, 10m, 15m, 20m, 25m) in the corridor. Since multiple beam pairs can have the same BID for some highest signal strength pair, we plot average SNR loss over all such pairs. Since $BID=0$ indicates the maximal strength pair, SNR loss is 0 in this case for all distances.

The figure reveals that for all positions, SNR drops monotonically with increase in BID (i.e., for beam pairs farther and farther away from the maximal pair). However, at 5m, beam pairs with $BID \leq 2$ are still within 1dB of the maximal pair, indicating only a small loss in link strength for selecting these sub-maximal beam pairs. More-

over, for $BID \leq 4$, the loss in SNR is still within 3dB (i.e., 50% of highest achievable SNR).

As the TX-RX distance increases, SNR decreases more rapidly with BID , indicating a greater degradation in relative strength of neighboring beam pairs. Furthermore, fewer beam pairs on average are within 3dB of the maximum possible signal strength for greater inter-node distances. Hence, the gain in signal strength is higher if an exhaustive search is performed for longer TX-RX distances, whereas for short distances selecting sub-maximal beam pairs can still yield high signal strength.

2.3.4 Performance of NLOS links

mmWave signals experience attenuation due to reflection and thus the link budget for NLOS components is expected to be lower compared to the LOS path [34]. Here, we measure and compare the achievable SNR and throughput with and without the presence of the LOS path. In particular, we want to explore the feasibility of Gbps scale throughput via reflections in the absence of LOS path. To this end, we consider the conference room since it has many reflectors such as whiteboard and TV screen (see Fig. 2.2c). The RX orientation in positions 4, 5, 6 on the table is such that there cannot be a LOS path between the TX and RX (the back of the phased array is blocked and there is no back lobe). On the other side, the RX sees the LOS path when located at positions 1, 2, 3 on the table.

Fig. 2.7 depicts the beam-pair heatmaps for positions 2 to 5. First, by comparing this figure with Fig. 2.3 and Fig. 2.5, we observe that more beam pairs provide positive SNR values due the better reflection in the conference room and shorter distance. The cluster of high SNR beam pairs for position 2 and 3 maps to the physical LOS path between two nodes confirming that LOS path was present for these positions. Similarly, the cluster of high SNR beam pairs for position 4 and 5 and the map of

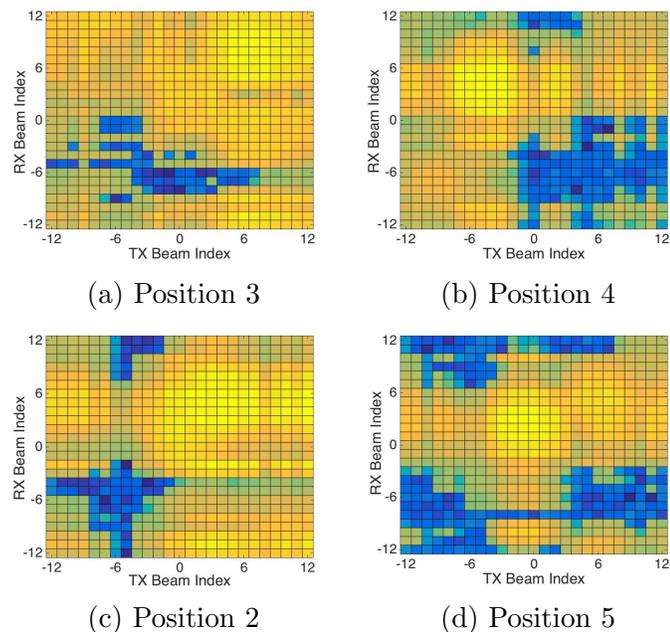


Figure 2.7: Beam-pair heatmaps for four conference room positions.

the conference room (Fig 2.2c) suggest that these beam pairs include a reflected path from the whiteboard.

Next, we measure the highest achievable SNR and throughput through the best beam-pair for each receiver position facing the transmitter (1 to 3) or reverse facing the transmitter (4 to 6). We depict the average SNR and Throughput in Fig. 2.8a and Fig. 2.8b, respectively. Note that MCS 4 (1/2 16 QAM) was used for modulation since it provides the highest throughput in all positions. Fig. 2.8a reveals that one can achieve 17-18 dB SNR, even in the absence of a LOS path. Furthermore, the throughput is close to 1.9 Gbps with and without the LOS path. Hence, we conclude that the SNR and throughput values for reflected paths can be as high as for the LOS path in real 60 GHz systems.

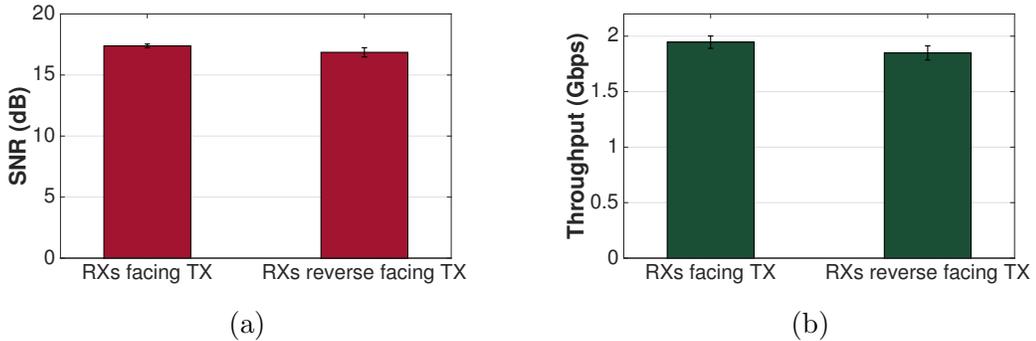


Figure 2.8: (a) Average SNR, and (b) Average Throughput received in positions 1 to 10 in the conference room.

2.3.5 Beam Misalignment and Nodal Mobility

In Sec. 2.3.2, we discussed the richness of strong beam pairs across different environments, and saw that any of these beam pairs, if identified by the training procedure, can establish a Gbps 60 GHz link. However, the alignment of the selected beams may subsequently be lost due to nodal mobility, which may lead to a degradation in signal strength or may even break the link, depending on the extent of mobility [36]. To study the impact of mobility on misalignment of selected beams and the subsequent loss in signal strength, we perform controlled experiments in the lobby isolating two key types of mobility, translation and rotation.

Lateral Translation: First we consider the scenario where the orientation of both TX and RX remains fixed; however, a change in receiver position results in a change in the relative angle between the two nodes. For this, we consider positions 1 through 6 in the lobby for a fixed RX orientation (0°). These positions emulate a path taken by a node as it moves perpendicularly to the TX in steps of 1m. Further, we consider three possible adaptation strategies by both nodes to adjust their beams. (i) Fixed beams i.e., both nodes keep using the same beams throughout the experiment. For this, we consider beam pair (0,0) which is the strongest at initial position. (ii) TX and RX adaptation, when both nodes perform an exhaustive search at each position

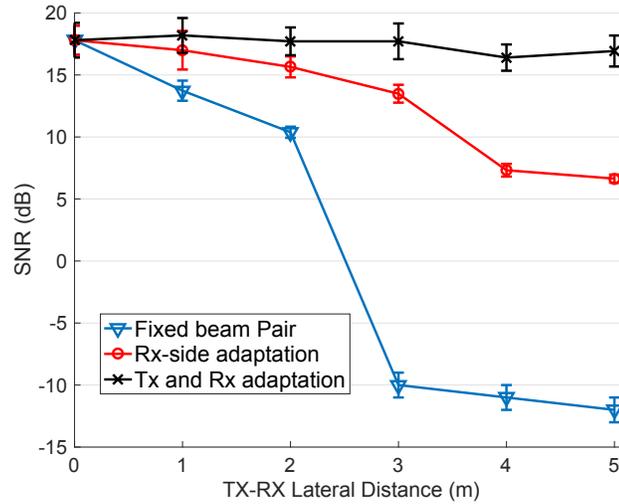


Figure 2.9: Beam adaptation strategies for lateral translation.

and re-select the strongest beam pair. *(iii)* RX-only adaptation, when only the RX locally adapts its beam to maximize link strength, while the TX beam remains fixed.

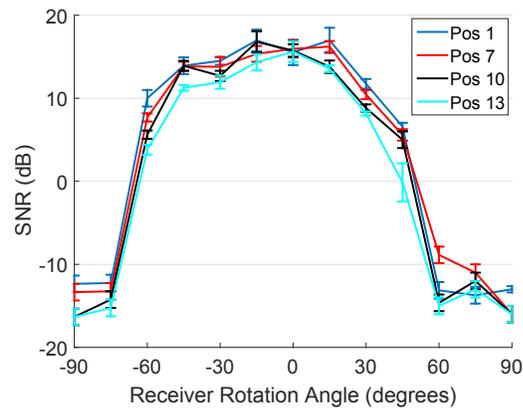
Fig. 2.9 plots SNR vs. TX-RX lateral distance for the three aforementioned strategies across lobby positions 1 through 6. For the case when TX and RX beams remain fixed to (0,0), the link strength decreases monotonically from position 1 through 6, as the relative angle between the two nodes increases from 0° at Position 0 to 60° at Position 6. In fact, SNR drops below the receive threshold after Position 3 when the relative angle is 40° . Note that the relative angle at Position 3 is still outside the beamwidth of the main-lobe for beam 0 at the receiver. The high SNR at this position results from a side-lobe of beam 0. However, for positions 4 through 6, the link cannot be sustained for beam pair (0,0) due to higher angular separation between TX and RX, which illustrates the significance of adapting beams for 60 GHz links in response to mobility.

In case of beam adaptation, SNR remains nearly constant across all positions for strategy-*(ii)*, when an exhaustive search is repeated at each position. This is the ideal scenario for beam adaptation, and depicts highest achievable SNR for the mobile receiver. For RX-only adaptation (strategy-*(iii)*), the search space only spans

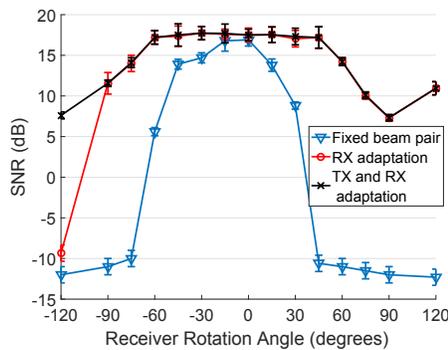
RX side beams and hence this strategy incurs much lower overhead than exhaustive search in the first case. However, in this case, only the RX-side beam has maximum alignment with the TX, while the TX beam remains fixed at 0. This results in SNR degradation as the TX-RX angular separation increases from Position 1 to Position 6. Despite this loss in TX-side alignment, Fig. 2.9 reveals that SNR achieved with RX-only adaptation is significantly better than that with no-adaptation strategy, and a link is sustained across all positions. This illustrates that a local search at the receiver, although sub-optimal, may be sufficient to maintain a directional 60 GHz link while avoiding exhaustive training.

Rotation: To analyze the impact of receiver rotation, we consider the change in signal strength of the central beam pair (0,0) for positions 1, 7, 10 and 13 which are in front of the TX (Fig. 2.2d). Fig. 2.10a shows SNR vs. receiver angle such that for all positions, 0° corresponds to a perfect alignment between TX and RX, and hence beam pair (0,0) achieves maximum SNR. As the RX rotates on either side, the SNR decreases sharply due to misalignment of RX beam 0. We also observe that the SNR for counter-clockwise rotation remains steady over a larger range of angles before dropping below 0 dB. This is due to asymmetrical radiation pattern of beam 0, as discussed in Sec. 2.1.3, which results in higher gain for counter-clockwise rotation. Moreover, as the inter-node distance increases, the SNR degrades for all orientations. However, the impact of rotation is more pronounced than that of increasing distance. This shows that a 60 GHz link, corresponding to a fixed beam pair (selected during training process), is highly susceptible to misalignment due to nodal rotation, and slight rotation can result in multiple dBs of SNR loss.

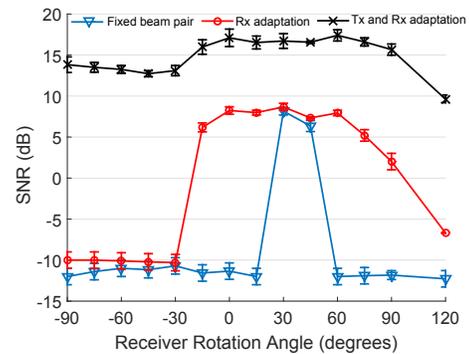
To improve link budget, the beams at either the TX or the RX need to be adjusted. Here we compare the three beam adaptation strategies, i.e., *fixed beam pairs*, *TX and RX adaptation*, and *RX adaptation only*, as we did for the lateral translation case.



(a) Rotation with fixed beam pair (0,0)



(b) Beam adaptation for Position 10



(c) Beam adaptation for Position 12

Figure 2.10: Impact of receiver rotation on SNR for various lobby positions.

For this, we consider two different receiver positions; Position 10 which is directly in front of TX (6.3m apart) and Position 12 which is roughly 30° to the right of TX (7.5m apart). Thus the latter position captures impact of both angular separation and receiver rotation.

Fig. 2.10b shows that for Position 10, RX-side adaptation achieves similar SNR as TX-and-RX adaptation. This is because the RX is directly in front of TX and is fixed, hence beam 0 is the best TX beam for all orientations of the receiver. However, unlike lateral translation, the maximum achievable SNR diminishes for higher RX angles on either side. This is a consequence of non-uniform angular spread of beam patterns and diminishing directivity gain of beam indices farther from the central

beam, a limitation of practical phased array antennas. Further, the increase in SNR for 120° RX angle depicts the impact of side-lobes, since for this orientation the main lobes of all receiver beams are misaligned with the TX. This is an example scenario of side-lobes generating additional paths to provide resilience to receiver mobility, an effect that cannot be observed in the case of horn antenna based systems.

For Position 12 (Fig. 2.10c), the angular separation between TX and RX further degrades signal strength compared to Position 10. For cases with fixed beam (0,0) and RX-only adaptation, the highest SNR is achieved for receiver angle 30° , since for this orientation the RX has maximum alignment with the TX. However, the SNR is low across all angles even with RX-only adaptation, since the TX is still using beam 0, which is misaligned with the receiver due to an angular separation of 30° . When beam adaptation is used at both TX and RX, the TX-side beam also becomes aligned with the RX, adding another 5-7 dB of SNR gain and making the highest signal strength similar to that observed for Position 10. These two examples illustrate the importance of identifying different mobility scenarios for protocol design, since adaptation strategies are highly dependent on the type of mobility.

2.4 Related Work

Initial experimental studies of 60 GHz in indoor environments focused on measuring and modeling channel propagation characteristics using dedicated channel sounding hardware (e.g., [34, 37–40]).

The recent work in this domain has largely been driven by observations and models derived from measurements with platforms that implement narrow-band transmissions instead of wideband [41–44], or/and horn antennas instead of phased arrays [2, 36, 45, 46]. While initial studies based on such platforms have provided valuable insights into mmWave propagation, such setups cannot capture the effects associated

with wideband transmissions, e.g., as specified in 802.11ad. For instance, past works relied on SNR measured over few hundred MHz of bandwidth to estimate rates by looking up a receiver sensitivity table. More importantly, the use of horn antennas masks the effects of imperfect beam-patterns, side-lobes, and non-uniform steerability, all typical features of beams realized through practical phased arrays. It is important to understand the impact of these artifacts as they directly affect mmWave link characteristics, interference, or spatial-reuse.

The only SDR testbeds capable of wideband transmission with phased arrays are OpenMili [47] and the testbed in [48]. OpenMili nodes are based on an off-the-shelf FPGA processor supporting a channel width of 1 GHz. These nodes are equipped with electronically steerable 60 GHz four-element phased arrays, with 2 possible values for each element’s weight. The testbed in [48] uses eight-element phased arrays but operates in the 24 GHz band. In contrast, X60 nodes have twelve-element phased arrays, 4 discrete possible phase values per element, and support a 2 GHz channel width, enabling higher rates and higher resolution experiments.

Many works have also explored the performance of WirelessHD or WiGig hardware available commercially [32, 49–53]. These devices offer the chance to understand the often complex inter-play between higher layers of the network stack and WirelessHD/WiGig directional PHY. However, they reveal rather limited information about and allow no control over the PHY/MAC parameters or the weights of the antenna elements of the phased array. These limitations, combined with proprietary rate and beam adaptation techniques, often make it hard for researchers to understand the causes of the observed performance. Further, the closed source firmware (which implements most of the lower MAC and PHY functionality) that these devices ship with limits the possibility of prototyping any new protocols.

User and Beam Selection in Multi-User 60 GHz WLANs

Together, millimeter-scale wavelength and GHz-scale bandwidth available at 60 GHz enable large antenna arrays, high directionality, and the potential to achieve high throughput. Today, the 60 GHz WLAN standard IEEE 802.11ad [13] supports rates up to 6.7 Gbps by transmitting to a single client at a time using analog beam steering which is discretized via predefined codebooks. In this chapter, we target scaling 60 GHz WLAN capacity via *multi-user* transmission by creating opportunities for an AP to simultaneously transmit to multiple users with limited inter-user interference. In particular, we make the following contributions.

First, we consider a 60 GHz system architecture in which the AP has multiple Radio Frequency (RF) chains used to support simultaneous transmission to multiple clients. Moreover, the AP has many more antennas than RF chains such that each RF chain can be precoded at baseband and independently steer beams at radio frequency, often termed digital beamforming and analog beam steering respectively. We show how the achievable rate for a multi-user transmission is impacted by the selected users, analog beam steering parameters, and digital pre-coding weights, and consequently,

we formulate their joint optimization. However, since the optimal solution cannot be realized in practice, we propose *decoupling* user selection and beam steering training as a simpler and lower overhead protocol structure. In this structure, each client undergoes beam steering training independent of potential grouping or multi-user transmission.

Second, we introduce a two-class framework for design and evaluation of user selection schemes in the decoupled structure. The first class, termed *Single-Shot (S^2) User Selection*, performs only based on the information obtained in beam steering training, i.e., without further measurements and in a “single shot.” We design a policy in this class termed *S^2 Maximum beam Separation* which exploits the idealized beam pattern corresponding to each user’s codebook entry selected during beam training. Namely, as the name suggests, it attempts to minimize inter-user interference by selecting a group of users with minimum overlap in (idealized) beams. Here, digital training follows user selection to enable zero-forcing to mitigate residual interference. The second class, termed *Interference-aware Incremental (I^2) User Selection*, performs user selection and digital training via several rounds, with users added incrementally after AP acquisition of empirical interference information at each round. We design two policies in the class: the first adopts the same beam separation strategy as above, but also has an interference test before addition of each user; the second employs a “partitioned multi-test” in which users are partitioned according to ranked beam training information, and users in the same partition compete to be part of the selected user group. While the S^2 class is inherently simpler and lower-overhead, I^2 strategies can potentially achieve higher rates due to the additional interference information they acquire.

Third, we explore the coupling between user selection and beam selection in practical 60 GHz systems using X60 nodes consist of National Instruments’ (NI) millimeter-

wave transceiver system equipped with a 24-element phased array antenna that can be steered in real-time and communicate over 2 GHz wide channels. We perform measurements and experimentally confirm that the choice of users and beams are tied together such that the best choice of analog beams for the link between the AP and user u depends on the users that are grouped with user u in the multi-user downlink transmission. Although X60 offers several key advantages over horn-antenna based platforms, it only supports a pre-defined analog codebook consisting of 25 beams with fixed beamwidth. Therefore, we also implement a testbed using a steerable 60 GHz RF-fronted combined with the software-defined radio platform WARP. This testbed utilizes mechanically steerable horn antennas to emulate 802.11ad phased-array, and can be configured with antennas with different beamwidths. Moreover, for comparison purposes, we define two benchmarking algorithms which use exhaustive search to characterize the maximum achievable rate with joint user-beam selection and the decoupling structure.

Finally, we perform an extensive measurement and simulation study. We begin with a baseline case of two simultaneous users and show how the best choice of analog beams depend on the user grouping. Then, we explore the performance loss due to decoupling beam steering from user selection with variable number of available users and AP's number of RF chains. We use exhaustive search to compare the decoupled structure with joint optimization and find that despite its improved simplicity and lower complexity, decoupling beam steering and user selection incurs less than 5% capacity loss with four RF chains at the AP. However, the performance gap increases as the group size increases and when only few users are available in the system. Then, we experimentally characterize the spatial multiplexing gains of 60 GHz WLANs and find that two receivers cannot share an analog beam (i.e., codebook entry), even if digital precoding attempts to remove inter-client interference. Surprisingly, in such

cases, zero forcing yields lower empirical sum capacity than no digital precoding. Nonetheless, we show that reducing beamwidth via a larger number of antennas not only helps prevent users from sharing the same beam, but also boosts the relative gain of zero forcing for residual interference cancellation via improved SINR. Next, we explore scaling AP's number of RF chains and streams. We show that while the performance of I^2 is able to increase with the number of RF chains (as does exhaustive search), S^2 performance *degrades* when too many streams cause excessive inter user interference that cannot be countered via zero forcing. Further, we study the case that a LOS path is not available and a reflected NLOS path is required. We find that while the lower SNR inherent to NLOS users makes them less tolerant to interference, up to 12 spatial streams can be multiplexed for NLOS, LOS, and any mix of the two.

The remainder of this chapter is organized as follows: Section 3.1 provides our multi-user architecture and explains the joint and decoupled structures for user selection and beam steering. Section 3.2 describes S^2 and I^2 platforms for decoupled user selection. Section 3.3 presents our benchmarks and compares them with decoupled user selection policies in terms of complexity. Section 3.4 presents our evaluation setup and Section 3.5 studies the coupling between beam selection and user selection in practical 60 GHz systems using phased array antenna. Section 3.6 investigates the spatial multiplexing gains of analog/digital beamforming. Section 3.7 evaluates S^2 and I^2 frameworks. Finally, we discuss related work in Section 3.8.

3.1 Multi-User Architecture and Protocols

In this section, we first describe the system architecture required to support multi-stream transmission. Next, we present the protocol and capacity implications of a decoupled vs. joint beam steering and user selection methodology.

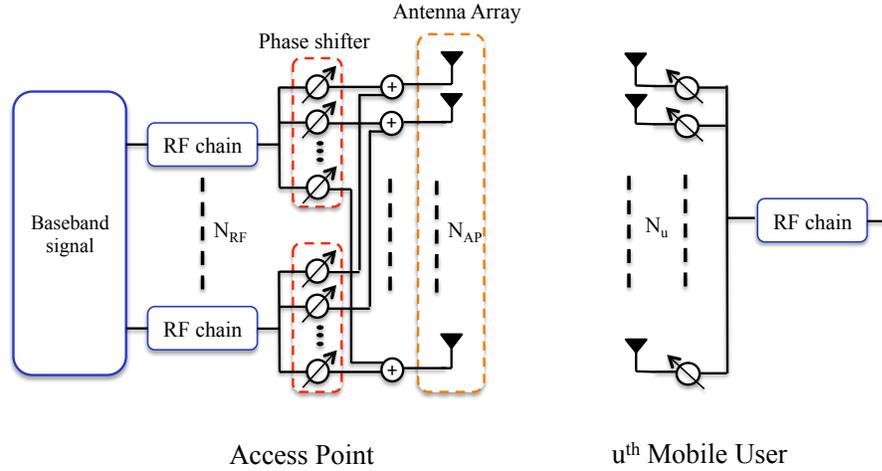


Figure 3.1: Hybrid beamforming architecture of AP and users.

3.1.1 System Architecture

60 GHz WLANs employ analog or RF beamforming to improve SNR [54, 55]. We interchangeably use the terms beam steering, RF beamforming, and analog beamforming to refer to application of different phase delays to different antenna elements in the RF domain. The IEEE 802.11ad standard supports such beam steering, but limits the AP to transmit to a single user at a time [13]. Consequently, both the AP and client require only a single RF chain for digital baseband processing. In contrast, the next 60 GHz WLAN standard IEEE 802.11ay, will supporting concurrent spatial streams from the AP to multiple clients [6]. This requires multiple RF chains at the AP (at least one per stream) whereas the clients require only a single RF chain. Moreover, the AP's plurality of RF chains can also be used for digital pre-coding at baseband to compliment analog beam steering. Precoding schemes such as zero-forcing can be exploited to minimize or ideally cancel the inter-user interference of a multi-user simultaneous transmission [56].

Such hybrid analog/digital beamforming has been studied in the literature with various architectural design [57, 58]. Fig. 3.1 depicts an example architecture to support multi-stream transmissions. As shown, the AP (left) is equipped with N_{RF}

RF chains, each capable of transmitting an independent data stream to a different client, and each can be precoded as described above. After modulation (not shown), the data streams can be independently steered with the depicted phase shifters, and the resulting signal is mapped to an array of N_{AP} antennas. The transmitted signal from each antenna element is the superposition of multiple data streams with different phase delays. Hence, as long as each data stream is transmitted via a fixed analog beam pattern (i.e., set of phase delays), the overall multi-stream transmission is similar, no matter how data streams are allocated to different RF chains. The AP can simultaneously transmit to N_{RF} out of U backlogged users, with the u^{th} mobile user (right) equipped with N_u antennas. Users have only one RF chain to allow for a lower cost and more energy efficient design compared to the AP, and therefore each user can receive only one spatial stream at a time.

3.1.2 Achievable Multi-User Rate

The achievable sum-rate is the key metric to evaluate user selection schemes as it can incorporate gains from analog beam steering and digital zero forcing as well as detrimental effects of residual inter-user interference.

Let $w_{u,tx} \in \mathbb{C}^{N_{AP} \times 1}$ be the transmit RF beamforming vector at the AP corresponding to u^{th} user, and $w_{u,rx} \in \mathbb{C}^{N_u \times 1}$ be the receive RF combining vector at u^{th} user. Assume that the AP intends to serve users in the group G simultaneously. Further let $F_{BB} = [f_n^{BB}]_{n \in G}$ be the $|G| \times |G|$ digital baseband precoder. Then, the received signal for user $u \in G$ after hybrid analog/digital beamforming can be written as

$$y_u = w_{u,rx}^* \sum_{n \in G} (H_u F_{RF} f_n^{BB} s_n + n_u), \quad (3.1)$$

where $F_{RF} = [w_{n,tx}]_{n \in G}$, $n_u \sim N(0, \sigma^2 I)$ is the Gaussian noise, and s_n is the transmitted symbol for n^{th} user. In Eq. (3.1), H_u is the $N_u \times N_{AP}$ dimensional channel

between the AP, transmitting with N_{AP} antennas, and user u .

Given the received signal in Eq. (3.1), the achievable sum-rate of the downlink simultaneous transmission to user group G , denoted as $R_{sum}(G)$, can be formulated as

$$R_{sum}(G) = \sum_{u \in G} \log_2 \left(\frac{\sigma^2 + \frac{P}{|G|} \sum_{n \in G} |w_{u,rx}^* H_u F_{RF} f_n^{BB}|^2}{\sigma^2 + \frac{P}{|G|} \sum_{n \in G, n \neq u} |w_{u,rx}^* H_u F_{RF} f_n^{BB}|^2} \right), \quad (3.2)$$

where P is the fixed transmit power of the AP and $|G|$ is the cardinality of user set G .

We therefore incorporate beam steering via the F_{RF} and $w_{u,rx}$ terms and zero-forcing via the f_n^{BB} term, as they impact sum-rate of multi-user transmission to user group G . Consequently, a user grouping protocol can be viewed as targeting to select users and their transmission parameters to maximize sum-rate with minimum training and grouping overhead.

3.1.3 Joint User and Beam Selection: the Optimal Approach

In general, the sum-rate above can be maximized by finding the optimal subset of users (G^*), optimal analog (RF) transmit beams (F_{RF}^*), RF receive beams $\{w_{u,rx}^*\}_{u \in G}$, and optimal digital (baseband) precoders (F_{BB}^*) that solve

$$\begin{aligned} \{G^*, F_{RF}^*, \{w_{u,rx}^*\}_{u \in G}, F_{BB}^*\} = & \\ & \arg \max \sum_{u=1}^U R_u(G, F_{RF}, w_{u,rx}, F_{BB}) \\ \text{s.t. } & [F_{RF}]_{:,u} \in F, u = 1, 2, \dots, U, \\ & w_{u,rx} \in W, u = 1, 2, \dots, U, \\ & |G| \leq N_{RF}, \end{aligned} \quad (3.3)$$

where the first two constraints ensure that RF beamforming and combining vectors need to be selected from the codebook F and W , respectively. RF phase shifters can only take quantized angles; hence, the analog beamforming/combining vectors can only take certain values which are stored in the finite-size pre-defined RF codebooks. The third constraint guarantees that the number of selected users is not greater than the number of AP's RF chains.

The solution to the problem in (3.3) requires the user set to be jointly selected with RF beamforming/combining vectors which yields a search over the entire $\sum_{m=1}^{N_{RF}} \binom{U}{m} (F^m \times W^m)$ space of all possible user and beam combinations. Furthermore, the digital precoder F_{BB} needs to be jointly designed with the analog beamforming/combining vectors. In practice, this needs the feedback of the channel matrices H_u , $u = 1, 2, \dots, U$. Hence, the direct solution of this sum-rate maximization requires prohibitively large training and feedback overhead. Nonetheless, we use exhaustive search in simple scenarios as a performance benchmark (cf. Section 3.3.1).

3.1.4 Decoupled User and Beam Selection

Instead of joint selection of users and beamforming weights, we employ a low-complexity structure for decoupling beam steering and user selection. The key technique is to first find analog beam steering parameters between the AP and each client, independent of potential grouping or multi-user transmission. Analog beam steering is the essential component of link establishment between the AP and clients and should precede user selection or any data exchange. Consequently, we term such training *Single-User Training (SUT)* since RF vectors are chosen based on the individual AP to client channels. In general, RF beamforming/combining vectors can be found by

solving the following optimization problem

$$\{w_{u,tx}^*, w_{u,rx}^*\} = \arg \max_{\substack{\forall w_{u,tx} \in F \\ \forall w_{u,rx} \in W}} \|w_{u,rx}^* H_u w_{u,tx}\|. \quad (3.4)$$

The knowledge of $w_{u,tx}^*$ and $w_{u,rx}^*$ is sufficient to establish a directional link between the AP and user u . The complexity cost of solving Eq. (3.4) is $O(N^2)$, where N is the codebook size; even suboptimal standardized beam training algorithms, e.g., beamforming training in 802.11ad, require complexity of $O(N)$ [13]. The large overhead associated with RF beamforming makes it undesirable to repeat for every multi-user transmission unless the selected RF beams are not reliable anymore due to channel variations. Therefore, we decouple RF beam steering and user selection and perform the latter for each AP transmission and the former only as required.

By decoupling beam steering training from multi-user considerations, we simplify the problem in (3.3) to a user selection problem. As a benefit of rich scattering, almost any randomly selected user group can achieve a multiplexing gain in sub-6 GHz networks [59]. However, 60 GHz channels lack the rich scattering properties and have different system architecture. This fundamental difference necessitate designing new user grouping strategies for 60 GHz networks.

Finding the digital weights is not a major challenge in the decoupled methodology since for a given selected user group and their beam steering vectors, the AP can obtain the digital beamforming weights (i.e., F_{BB}) by feeding back pilot measurements for a single user at a time, much as in conventional multi-user systems below 6 GHz. The optimal method of computing digital weights known as Dirty Paper Coding is impractical due to high computational complexity. Instead, we utilize Zero-Forcing Beamforming (ZFBBF), a sub-optimal yet simple method of computing digital weights.

3.2 A Framework for Decoupling User Selection and Beam Steering

3.2.1 Overview

We define two complimentary classes for user selection: Single-Shot (S^2) and Interference-aware Incremental (I^2) User Selection. Both procedures employ the decoupled methodology, i.e., they first perform SUT to discover the RF beamforming/combining weights.

(i) Single-Shot (S^2) User Selection. This class of user selection schemes groups users solely based on information acquired in SUT. We name this class single-shot since user selection is performed in one epoch without further channel sounding or feedback exchange. The rationale is that the sparse-scattering nature of 60 GHz channels makes it possible to mitigate inter-user interference solely via beam steering. Thus, by utilizing the reports from SUT, single-shot schemes choose users with the lowest beam overlap and rely on ZFBF to cancel residual inter-user interference.

(ii) Interference-aware Incremental (I^2) User Selection. This class selects users via a multi-round procedure in which each round includes AP acquisition of empirical interference information and the achievable sum-rate *after* ZFBF for users included in the round. In contrast to S^2 , I^2 employs incremental user selection in which users are added one by one in each round. The AP tests and trains one or more users per round, and incorporates the net positive effect of adding a user with the determinant of additional inter-user interference. While I^2 user selection can potentially achieve higher sum-rate due to additional multi-user interference measurements, it requires larger overhead compared to S^2 approaches.

3.2.2 Single-Shot (S^2) User Selection

3.2.2.1 Overview

User selection and training procedures that are members of the S^2 can be characterized by the timeline depicted in Fig. 3.2. The figure shows that the Single-User Training (Phase 1) is performed initially or as needed according to client or environmental mobility, but typically at a slower time scale than packet transmission. When the AP has packets queued for multiple clients and can operate a multi-user transmission, it performs user selection using channel information only from the prior Single-User Training. The selected users can then be digitally trained for ZFBF. In particular, the three phases are as follows

SUT (Phase 1). The beam training procedure in 802.11ad consists two steps [13]. First, the AP transmits beam training frames and sweeps through all beam patterns of codebook F while the receiver adopts a quasi-omni pattern. This results in selecting the transmit RF beamforming vector which provides highest signal strength at the user, i.e., $w_{u,tx}$ for user u . Then, the user sweeps through beam patterns of codebook W while the AP is in quasi-omni mode in order to find the RF combining vector ($w_{u,rx}$). Finally, the AP and trained user exchange a feedback frame to finalize the selected beams. The AP and user adopt their selected RF vectors for this transmission. Hence, the directional single-user channel can be extracted from this feedback as

$$h_u^{SU} = w_{u,rx}^* H_u w_{u,tx}, \quad u = 1, 2, \dots, U. \quad (3.5)$$

User Selection (Phase 2). In this class, user selection needs to be completed before digital precoding since digital beamforming weights are computed for a specific set of users to cancel or mitigate their inter-user interference. Therefore, user selection is the intermediate step between the analog and digital precoding as depicted in

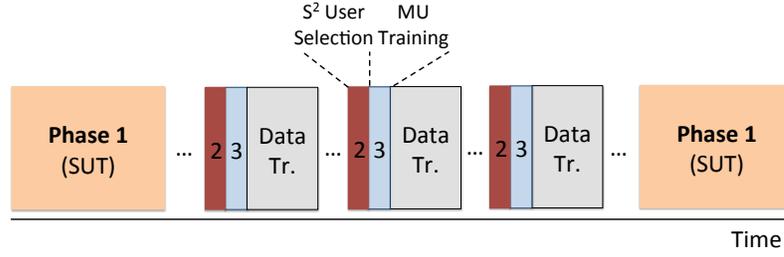


Figure 3.2: Single-Shot (S^2) User Selection in the context of analog and digital beamforming.

Fig. 3.2. Available information for S^2 user selection includes not only the above Phase 1 information, but also the AP system state. Namely, the AP knows the hardware configuration of itself and its users such as the number of transmit antennas (N_{AP}), number of RF chains (N_{RF}), the number of users (U), and the finite-size RF codebooks (W and F).

Thus, we define the family of S^2 user selection to include all schemes that rely on the above information

$$\mathcal{G} = f(N_{RF}, U, N_{AP}, W, F, \{w_{u,tx}\}_{u=1}^U, \{h_u^{SU}\}_{u=1}^U), \quad (3.6)$$

where \mathcal{G} is the set of users that are co-scheduled (grouped). In Eq. (3.6), N_{RF}, U, N_{AP}, W, F are system state information, $\{w_{u,tx}\}_{u=1}^U$ and $\{h_u^{SU}\}_{u=1}^U$ are beam training information. Thus, a user selection mechanism that does not require any extra information is a member of this class.

Multi-user Zero-Forcing (Phase 3). In order to cancel the residual inter-user interference of the co-scheduled (selected) users, the AP can exploit common digital precoding schemes (e.g., ZFBF). Thus far, the AP knows selected user group \mathcal{G} and RF beamforming vectors from Phase 1 and Phase 2. Therefore, the directional channel between the AP and user group \mathcal{G} can be measured through a conventional channel sounding procedure. The measured channel depends on the choice of users

and beams as follows

$$\begin{aligned}\bar{h}_u &= w_{u,rx}^* H_u [w_{u,tx}]_{u \in \mathcal{G}}. \\ \bar{H}_{\mathcal{G}} &= [\bar{h}_u]_{u \in \mathcal{G}}.\end{aligned}\tag{3.7}$$

Applying ZFBF scheme, the digital weights (i.e., $F_{BB}(\mathcal{G})$) will be

$$F_{BB}(\mathcal{G}) = \bar{H}_{\mathcal{G}}^* (\bar{H}_{\mathcal{G}} \bar{H}_{\mathcal{G}}^*)^{-1}.\tag{3.8}$$

Once user selection and digital training are complete, the AP is ready to transmit independent streams to users in \mathcal{G} simultaneously.

3.2.2.2 S^2 Maximum beAm Separation (S^2 -MAS)

We define an exemplary single-shot user selection strategy which groups users based on the selected beams, targeting maximum beam separation. The relative beam separation of two users can be realized via their RF beamforming vectors selected in SUT. In particular, we define $BeamID_u$ associated with user u as the index of its RF beamforming vector (i.e., $w_{u,tx}$) among columns of AP's codebook. We elaborate on $BeamID_u$ and illustrate S^2 -MAS strategy via an example.

Fig. 3.3 depicts an example scenario in which the AP generates 16 beam patterns with 16 antenna elements based on the codebook proposed in [60]. Four users and their associated beam IDs are shown in this picture. Every beam pattern has two strong main lobes and multiple side lobes. For illustration, assume we intend to group an additional user with User 1 via S^2 -MAS selection policy for a two-user simultaneous transmission. As shown in Fig. 3.3, the second beam pattern (red) is the directional beam that corresponds to User 1 (i.e., $BeamID_1 = 2$). S^2 -MAS selects User 2 with $BeamID_2 = 13$ to be grouped with User 1 since it has the maximum

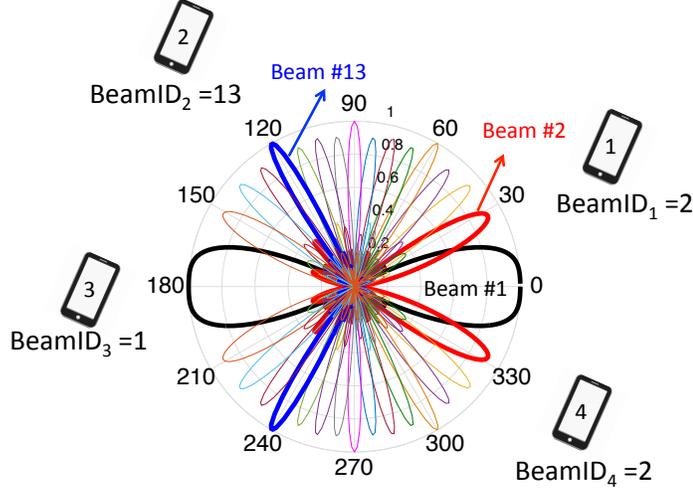


Figure 3.3: Illustrative scenario for S^2 -MAS: one AP with 16 beam patterns and four users.

Beam ID distance with User 1 compared to others. As we observe, the outcome of this selection policy depends on the beam patterns (or RF codebook) at the AP and not the physical angular distance between users; for instance, User 3 has the highest angular separation with User 1 ($\sim 120^\circ$) but small Beam ID distance and is not selected via S^2 -MAS. In general, S^2 -MAS strategy groups $N_{RF} - 1$ users with the primary user in one shot such that the sum of pairwise beamID distance (between any two users) is maximum. The primary user is the one selected based on any fairness mechanisms in the IEEE 802.11 DCF protocol based networks. Note that after the user selection phase, the AP and users will undergo digital training to best cancel out any residual inter-user interference as described above.

3.2.3 Interference-aware Incremental (I^2) User Selection

3.2.3.1 Overview

I^2 user selection employs interference-aware user addition with a measurement and feedback round for addition of each user. Namely, I^2 adopts a multi-round procedure

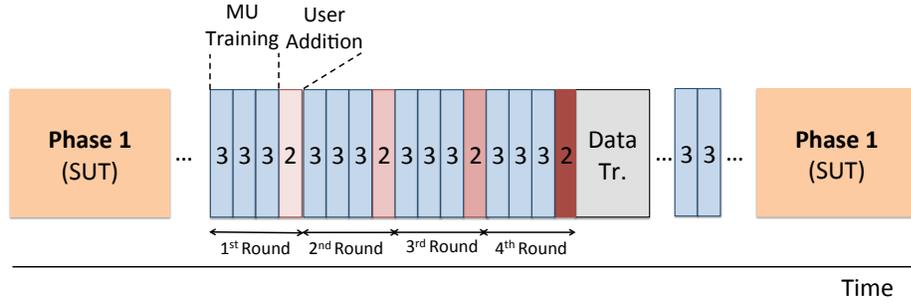


Figure 3.4: Interference-aware Incremental (I^2) User Selection in the context of analog and digital beamforming.

for testing and training each additional user, incorporating the net positive effect of adding a new user with the detriment of additional inter-user interference.

Fig. 3.4 depicts the I^2 mechanism in which SUT is completed as the first phase similar to single-shot procedure. Subsequently, user selection and digital training are performed via several rounds (e.g., Fig. 3.4 shows four rounds). In each round, the AP acquires empirical interference information of a set of *candidate users* from which only one user will be added to the multi-user group in that round (in Fig. 3.4, three potential users are assumed in each round). Hence, the AP estimates the achievable sum-rate when a potential user is grouped with already selected users. Let G_{i-1} be the user group at the end of $(i-1)^{th}$ round and n be the index of the user which provides the highest sum-rate among potential users in round i . If the collected interference information shows that $R_{sum}(G_{i-1} \cup \{n\}) > R_{sum}(G_{i-1})$, the AP proceeds to the next iteration and adds user n to the group (i.e., $G_i = G_{i-1} \cup \{n\}$). Otherwise, the user addition is terminated and the AP transmits to user group G_{i-1} . While incremental algorithms potentially provide higher data rates due to measurement-based interference-aware addition of users, complexity and overhead time for feedback is increased compared to single shot.

Table 3.1: Comparing complexity or search space size of different policies.

Algorithm	Beam selection	User selection	Total search space
Exhaustive Joint	-	-	$\sum_{m=1}^{N_{RF}} \binom{U}{m} (F^m \times W^m)$
Exhaustive Decoupled	$U \times (F \times W)$	$\sum_{m=1}^{N_{RF}} \binom{U}{m}$	$U \times (F \times W) + \sum_{m=1}^{N_{RF}} \binom{U}{m}$
$S^2 - MAS$	$U \times (F \times W)$	0	$U \times (F \times W)$
$I^2 - MAS$	$U \times (F \times W)$	$O(N_{RF})$	$U \times (F \times W) + O(N_{RF})$
$I^2 - PM$	$U \times (F \times W)$	$(\frac{U}{N_{RF}})(N_{RF} - 1)$	$U \times (F \times W) + (\frac{U}{N_{RF}})(N_{RF} - 1)$

3.2.3.2 Example User Selection Policies

The members of I^2 class differ in how they select the set of candidate users in each round. Here, we define two exemplary strategies.

(i) I^2 Maximum beam Separation (I^2 -MAS). We define I^2 -MAS as an incremental user selection strategy, in which an additional user is added in each round as follows. Like S^2 -MAS, I^2 -MAS targets adding user which has maximum beam. In contrast, I^2 -MAS adds users one by one and not all in one-shot. This enables the AP to evaluate the achievable sum-rate of multi-user transmission as the user group grows. User selection is terminated when adding one more user decreases the achievable sum-rate due to excessive inter-user interference or serving another user is impossible due to lack of RF chains. At this stage, ZFBF training has already been performed for all users in the group and transmission can occur.

(ii) I^2 Partitioned Multi-test (I^2 -PM). We define I^2 -PM as an incremental algorithm, in which for each round, multiple users are tested/trained and the one that maximizes the sum-rate will be added to the existing group. The algorithm first sorts users in descending order according to the norm of their single-user channels (i.e., h_u^{SU} for user u which is available after SUT and formulated in Eq. (3.5)). Then, I^2 -PM partitions users into N_{RF} partitions and labels them from 1 to N_{RF} such that partition 1 includes $\lceil U/N_{RF} \rceil$ users with the highest channel norms, partition 2 with the second highest, etc. For any pre-selected first user, the algorithm finds the par-

tion index which contains this user (denote this index as i). I²-PM tests all users in the $(i + 1)^{\text{th}}$ partition by measuring their channel state information to calculate their induced interference and their achievable sum-rate when grouped with already selected user(s). Next, the user from the partition with the maximum sum-rate is selected and the algorithm proceeds by updating $i \equiv i + 1 \pmod{N_{RF}}$ and continuing the same procedure. This incremental addition can happen at most $N_{RF} - 1$ times and, like I²-MAS, user addition terminates when the sum-rate would degrade even for the best tested user. While I²-PM tests multiple users in each round, the test only involves $\lceil U/N_{RF} \rceil$ users belonging to one partition, thus reducing complexity compared to exhaustive testing of all users in each round. Therefore, the total number of tests for I²-PM is $O(U)$.

3.3 Benchmarks and Complexity Comparison

In this section, we first introduce two benchmarking algorithms that provide the maximum (i.e, upper-bound on) achievable sum-capacity via joint and decoupled structures. Then, we compare the complexity of all introduced policies along with the benchmarking algorithms.

3.3.1 Benchmarking Algorithms

For evaluation purposes, we introduce two benchmarking algorithms for joint and decoupled user and beam selection.

(i) Exhaustive Decoupled. In this algorithm, we adopt Single-User Training (SUT) to find the RF beamforming/combining vectors corresponding to each user. However, for user selection, the algorithm goes over $\sum_{m=1}^{N_{RF}} \binom{U}{m}$ possible user combinations. For a test user group G_t , it computes zero-forcing weights based on Eq.

(3.8) and then calculates $R_{sum}(G_t)$ according to Eq. (3.2). The algorithm records the achievable rate of all possible user groups and picks the one with maximum sum-rate. Since the overhead of collecting all users' channel linearly scales with number of users, *Exhaustive Decoupled* is inefficient for dense user population scenarios. However, we study this algorithm for comparison purposes as it provides an upper-bound for the achievable sum-rate via decoupling beam steering and user selection.

(ii) Exhaustive Joint. We formulated the general problem of user selection, RF beamforming and digital precoding in (3.3). The optimal solution of (3.3) given a fixed digital precoding scheme (e.g., zero-forcing) yields to an exhaustive search over all possible user-beam tuples. We call this algorithm *Exhaustive Joint* since it searches through all different (user, beam) combinations. The search space includes $\sum_{m=1}^{N_{RF}} \binom{U}{m} (F^m \times W^m)$ distinct combinations where F and W are AP's and users' codebooks, respectively. Implementation of *Exhaustive Joint* algorithm may not be practical in real scenarios due to the high computational complexity; however, comparing it with *Exhaustive Decoupled* algorithm determines the performance loss due to decoupling beam steering and user selection. We quantify and analyze this performance loss in Section 3.5.2.2.

3.3.2 Complexity

To explore the tradeoff between complexity and optimality, we summarize the search space size for described S^2 and I^2 policies as well as benchmarking algorithms in Table 3.1. Search space size captures both time complexity and computational complexity, i.e, a larger search space implies larger overhead time (to collect the necessary information to construct to search space) and higher computational costs (to find the best choice in the search space). In the decoupled structure, as explained in Section 3.2, the AP and each user discover the best transmit and receive beam which

provide the highest received signal strength for directional single-user transmission. This is accomplished by searching through all beam patterns in the pre-determined transmit and receive codebooks. Hence, the decoupled approaches, including Exhaustive Decoupled, have the search space size of at most $U \times (F \times W)$ for beam selection, where U is the user population size, F is the transmit codebook size and W is the receive codebook size. S²-MAS is a single-shot policy inducing zero user grouping search while I²-MAS adopts an incremental approach with maximum $(N_{RF} - 1)$ rounds and testing users with maximum beam separation with already selected users in each round. Hence, the exact number of users tested in each round is scenario-dependent; however, the search space size is $O(N_{RF})$. Similarly, I²-PM has maximum $(N_{RF} - 1)$ rounds and searches through $\frac{U}{N_{RF}}$ candidate users in each round, as shown in Table 3.1. One can compute the overhead ratio of two schemes, e.g., S²-MAS and Exhaustive Joint, using Table 3.1. For instance, when $U = 10, m = 3, |F| = 16, |W| = 4$, this ratio will be around 0.002% implying that S²-MAS incurs only 0.002% of Exhaustive Joint's overhead for user and beam selection.

3.4 Evaluation Setup: Testbeds and Trace Driven Emulation

We implement the key components of multi-user MIMO transmission in 60 GHz WLAN. In order to evaluate different parameters such as beamwidth and shape of of analog beam and number of simultaneous streams, we exploit two different testbeds, X60 with phased array antennas which was introduced in Chapter 2, and WARP testbed with horn antennas.

3.4.1 X60 Testbed

We use the testbed setup from [1] which is a configurable software defined radio based 60 GHz testbed and shown in Fig. 3.5. X60 nodes are equipped with National Instruments' (NI) millimeter-wave transceiver systems and 24-element phased antenna arrays. X60 nodes enable communication over 2 GHz wide channels via realistic predetermined beam patterns that can be steered in real-time, overcoming the inherent limitation of horn-antenna based platforms where the beam is steered using a mechanical rotator and may not always be representative of the often imperfect beams generated by phased arrays. The built-in codebook consists of 25 beam patterns spaced roughly 5° apart (in their main lobe directions) and covering a sector of -60° (corresponding to beam index 1) to 60° (corresponding to beam index 25) around the antenna's broadside direction. The 3 dB beamwidth of the beams is around 30 degrees and the idealized beam patterns can be found in [1] and also in Chapter 2. One limitation is that the beamforming codebook or beam patterns are not programmable; hence, we cannot evaluate the impact of antenna array size, beamwidth or similar factors that require configuring the analog codebook. Moreover, X60 nodes have single RF chains and do not support simultaneous transmission/reception of multiple data streams or digital beamforming. Instead, we emulated zero-forcing using the channel traces collected via X60.

3.4.2 WARP testbed

We use the testbed setup from [36, 45] enhanced as follows: our testbed consists of commercial mm-wave transceivers from the VubIQ 60 GHz development system, two WARP v1 boards, and circuits for signal adjustment outlined in Fig. 3.6. The mm-wave transceivers are capable of communicating in 57-64 unlicensed band with up to 1.8 GHz modulation bandwidth. Using WARP-lab [61] and VubIQ control panels,

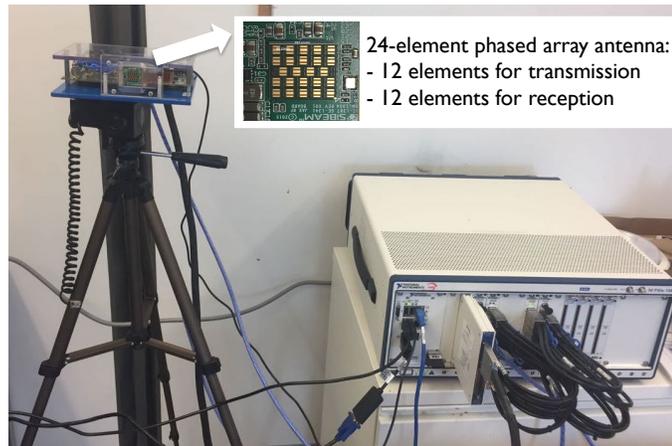


Figure 3.5: X60 node for wideband 60 GHz WLANs with a 12-element phased array antenna.

we apply the transmit module as the AP and the receive module as the client. Random binary data is generated via WARPLab and encoded using BPSK. The VubIQ module converts the signal to 60 GHz band and horn antennas provide directional transmission emulating phased array antennas. To achieve different beamwidths, we configure 7° , 20° and 80° horn antennas. In the receiver's VubIQ module, the signal is received by a horn antenna and is downconverted to analog I/Q baseband. In order to collect Received Signal Strength (RSS) for different locations and antenna orientations, we use mechanical motors and DC microstep driver with a motion control setup connected to the transceivers to steer the beams with sub-degree accuracy.

Using this platform, we can measure RSS of a point to point transmission. However, WARP board is equipped with a single RF chain and does not allow for multi-user concurrent transmissions. Analog beamforming is emulated with the help of mechanical motors which allows the transmitter to steer. The advantage of WARP-based platform over X60 is the possibility of changing the beamforming gain by exploiting different horn antennas whereas beamforming codebook of X60 is fixed. Similar to X60, direct implementation of digital beamforming which requires the transmitter to

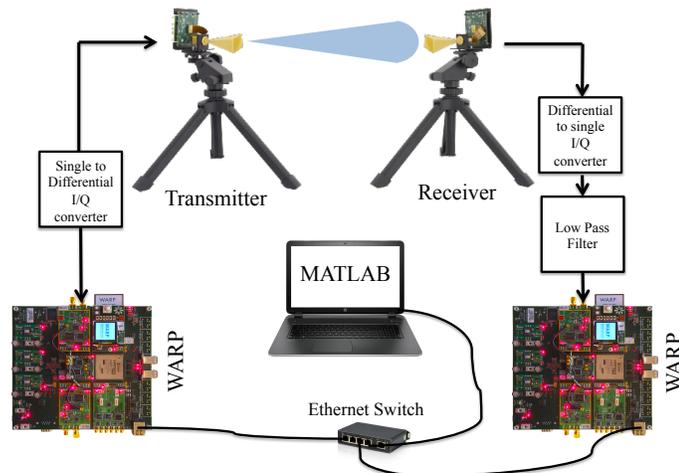


Figure 3.6: WARP-based 60 GHz hardware blocks with horn antennas.

obtain channel state information is still impossible. However, we can obtain channel matrix by trace-driven emulations based on the millimeter wave channel models from the literature [35, 62, 63]. These channel models are designed based on measurement campaigns in different LOS and NLOS environments. We compute the channel matrix (H) between any two nodes via these channel models and compute the analog beamforming or digital weights. To confirm the integrity of the obtained channel matrix, we use our 60 GHz testbeds to measure the main components of the channels such as RSS, Angle of Departure (AoD) and Angle of Arrival (AoA) of the LOS path and strong reflected paths in a similar scenario (i.e, same distance and orientation). We perform over 10,000 measurements varying receiver location, antenna orientation, and RF beamwidth, and all data sets will be available online upon publication. We observe an average error of 7.02% in RSS values by directly comparing measurements and simulations. We present results using simulations that incorporate measurement data for parameter setting as well as channel traces as feasible.

3.5 Experimental Realization of Coupling in User and Beam Selection

In this section, we first show how user selection and beam selection are coupled via an experimental study. Then, we evaluate the performance loss due to decoupling user and beam selection. We analyze how this performance loss depends on the number of users (U), and maximum user group size (N_{RF}) in subsections 3.5.2.1 and 3.5.2.2, respectively.

3.5.1 Coupling of User and Beam Selection

Here we experimentally explore the inherent coupling of beam selection and user selection in real 60 GHz systems employing phased array antennas to generate directional beam patterns. In particular, given two users, we show how the optimum choice of analog beams would change if the AP knows these two users are grouped for simultaneous reception. To this end, we conduct over-the-air measurements with X60 nodes (described in Section 3.4).

Scenario. We deploy a scenario as depicted in Fig. 3.7a in which the TX is fixed and oriented toward west. We assume a two-user system in which the first user (labeled as R_1) is fixed, 3.3m apart from the TX and pointing toward east. The second user (R_2) can be placed in any of the 13 marked positions in Fig. 3.7a and is always orientated toward east. We employ X60 nodes as transmitter and receivers. As explained in Section 3.4, X60 nodes have a 12-element phased antenna arrays and a codebook consisting of 25 beam patterns, with 3dB beamwidth of around 30 degrees, spaces roughly 5° apart (in their main lobe directions) and covering a sector of -60° (corresponding to beam index 1) to 60° (corresponding to beam index 25) around the antenna's broadside direction.

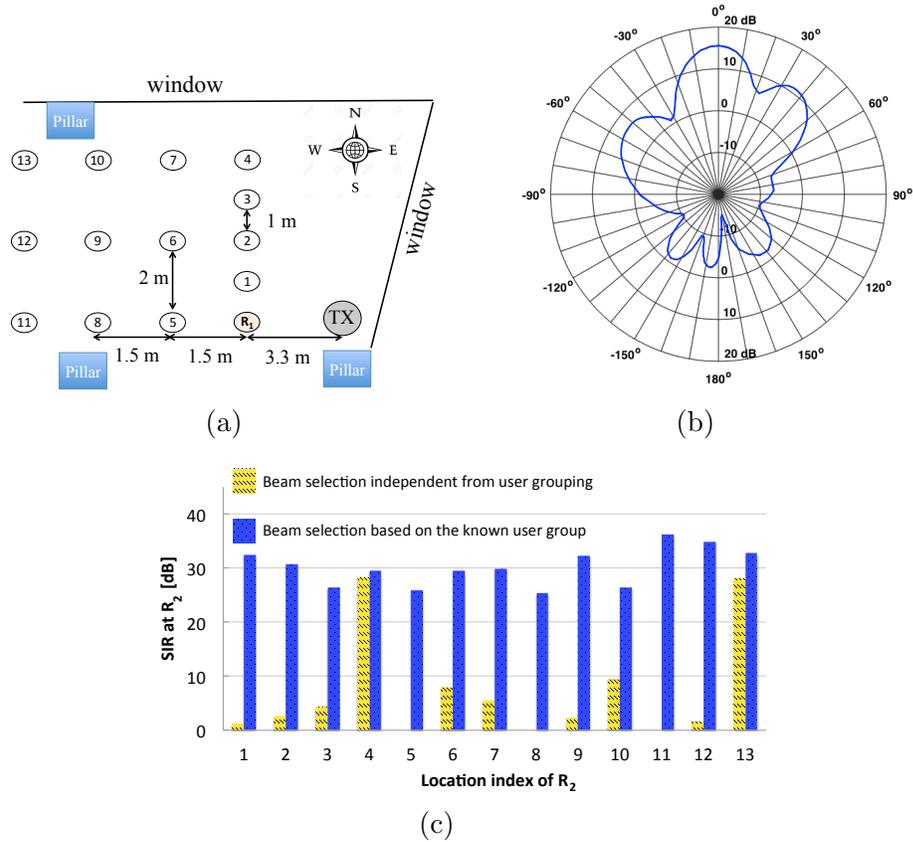


Figure 3.7: (a) The measurement scenario; (b) The beam pattern selected for R_1 in the SUT phase; (c) The Signal to Interference ratio (SIR) at R_2 when TX is transmitting to both R_1 and R_2 under two scenarios: (1) Analog beams are selected independent from user grouping; (2) Analog beams are selected such that the total SIR at R_1 and R_2 is maximum.

First, without having any user grouping knowledge, nodes find the TX/RX analog beams that maximize the received SNR for the single-user links $TX - R_1$ and $TX - R_2$. To this end, we exhaustively search through all 25 beams in TX codebook and 25 beams of RX codebook and capture a 25×25 matrix of SNR values such that the value in m^{th} row and n^{th} column corresponds to the received SNR when transmitter and R_1 adopt beam index m and n , respectively ($m, n = 1 : 25$). Hence, similar to SUT, we select the beam pair that provides the highest SNR value for directional single-user transmission. Since TX and R_1 are pointing toward each other, SUT results in the selection of beam index 13 (with main lobe direction of 0°) in the transmit side as

well as the receive side. Fig. 3.7b shows the beam pattern of beam index 13 with main lobe direction 0 degree and two side-lobes with directions 40 and -60 degrees. Similarly, we place R_2 in 13 different positions (marked in Fig. 3.7a) and find the beam pair that achieves the maximum SNR for the single-user link between TX and R_2 .

Then, knowing that R_1 and R_2 are grouped for a two-user transmission, we perform post-processing calculations on the obtained SNR matrices to find the best beams for TX , R_1 and R_2 that maximizes total Signal to Interference (SIR) ratio. Assume (m_1, n_1) is an arbitrary beam pair for link $TX - R_1$ and (m_2, n_2) is an arbitrary choice for for link $TX - R_2$, then we estimate SIR at R_1 and R_2 as follows:

$$SIR(R_1) = SNR_{R_1}(m_1, n_1)/SNR_{R_1}(m_2, n_1) \quad (3.9)$$

$$SIR(R_2) = SNR_{R_2}(m_2, n_2)/SNR_{R_2}(m_1, n_2) \quad (3.10)$$

where SNR_{R_1} and SNR_{R_2} are the collected 25×25 SNR matrices for R_1 and R_2 , respectively. Hence, for a given position of R_2 , we calculate $SIR(R_1)$ and $SIR(R_2)$ for any beam combinations and select the beams that maximize $SIR(R_1)+SIR(R_2)$. We refer to this method as joint beam selection since it includes the impact of inter-user interference in selecting beams for directional two-user transmission.

Fig. 3.7c plots SIR at R_2 for different location index of R_2 under two different scenarios: (i) when employing RF beams that were selected via SUT independent from user grouping (yellow bars); and (ii) when employing beams that were jointly selected for R_1 and R_2 as described above (blue bars).

First, we observe that when R_2 is placed at location 5, or 8, or 11, SIR at R_2 is 0 dB when analog beams are selected independent from choice of users. The reason is, at these locations, R_2 and R_1 are placed in the same line pointing toward the transmitter. Thus, the transmitter selects the same analog beam, i.e., beam index

13, for both R_1 and R_2 . This results in interference power being equal to the signal power and SIR to be 0 dB. In contrast, knowing that R_1 and R_2 are in the same multi-user group, the transmitter is able to find “separate beam” transmission via NLOS path and still provide above 25 dB SIR at R_2 .

Second, we observe that via single-user beam selection, the achievable SIR at R_2 increases with higher spatial separation between R_1 and R_2 causing lower inter-user interference, e.g., when R_2 is moving Fig. 3.7c shows a similar trend in SIR as the spatial separation increases by changing the position index of R_2 from 8 to 10 and from 11 to 13. However, position indices 5-6-7 does not follow this trend. The reason is behind the transmit antenna beam pattern used for R_1 . As explained earlier, beam index 13 with the pattern depicted in Fig. 3.7b is chosen at TX to establish a directional point to point link with R_1 . This beam has a side lobe at 40 degrees which is pointing toward location index 7 with relative angle $\tan^{-1}(4/4.8) = 39.8^\circ$ with the transmitter. This side-lobe increases the interference level at R_2 when locating at position 7 and cause R_2 to experience low SIR even with high spatial separation with R_1 . The same side-lobe also slightly decreases the SIR of R_2 at location index 10.

Third, 3.7c reveals that the performance of single-user beam selection in inter-user interference mitigation highly varies for different R_2 locations. In particular, assume we have 13 users (other than R_1) which are positioned in the 13 marked locations in Fig. 3.7a. In such case, a single-user beam selection approach with a “bad” user group, e.g., R_2 at position index 5 and R_1 , would result in low multi-user rate due to the excessive inter-user interference. However, for a “good” choice of users, e.g., R_2 at position index 4 and R_1 , high multi-user rate can be achieved even via using analog beams that have been selected independent from choice of users. Hence, the *Exhaustive Decoupled* algorithm would suggest the user at position index 4 or 13 to be grouped with R_1 .

Finally, Fig. 3.7c reveals that the joint selection of beams can successfully provide > 25 dB SIR at R_2 disregard of its distance to TX or spatial separation with R_1 . In most cases, the SIR gap between two methods is high due to the large (around 30 degree) beamwidth and presence of strong side-lobes causing significant inter-user interference. On the other side, beam selection based on SIR guarantees that chosen beams provide high SIR and data rate, consequently. We further explore this performance gap under different scenarios, e.g., different beamwidth and receiver separation, in the remainder of this chapter. Furthermore, we will discuss whether digital precoding (e.g., zero-forcing) can eliminate interference and reduce this gap.

Finding: The selection of beams and users are tied to each other such that an analog beam that maximizes the received signal strength at particular user for a single-user transmission may experience more than 25 dB degradation in SIR when that user is grouped with another user for a two-user transmission due to the excessive inter-user interference.

3.5.2 Performance Loss of Decoupling

To quantify the performance loss caused by decoupling user selection from beam selection, we explore the performance gap between *Exhaustive Joint* and *Exhaustive Decoupled* algorithms. This loss depends on the number users with downlink data (i.e., number of available users), and the number of RF chains at AP (i.e., the maximum group size).

3.5.2.1 Impact of number of users

Here, we explore the impact of number of available users on the performance gap between decoupled and joint selection frameworks. We employ the same node deployment setup of Fig. 3.7a using X60 nodes as TX and RXs. As before, we want to

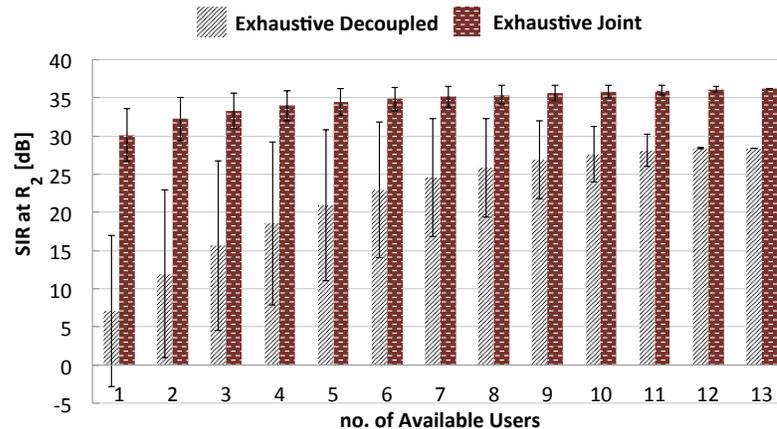


Figure 3.8: The Signal to Interference ratio (SIR) at R_2 when grouped with R_1 for a two-user simultaneous transmission.

group another user with the primary user R_1 for a two-user simultaneous reception. We put 13 other users in the marked positions Fig. 3.7a. However, the TX may not have downlink data for all users; thus, only users which downlink traffic should be considered for user grouping. We again use SIR metric to directly evaluate and compare the inter-user interference. Fig. 3.8 depicts SIR at R_2 as a function of number of available users (i.e., users with downlink data). For example, when number of available users is 1 only one user (out of 13) has downlink data and is selected to be grouped with R_1 under both *Exhaustive Joint* and *Exhaustive Decoupled* algorithms; however, these algorithms may employ different analog beams. Fig. 3.8 reveals that the performance gap between *Exhaustive Joint* and *Exhaustive Decoupled* algorithms decreases as number of available users increase. The high standard deviation for the decoupled strategy indicates high performance dependency on the location of available users. In contrast, *Exhaustive Joint* successfully degrades inter-user interference and provides above 30 dB SIR disregard of the number of available users and their relative positions.

3.5.2.2 Impact of user group size

Here we compare the sum-rate of *Exhaustive Joint* and *Exhaustive Decoupled* algorithms after applying zero-forcing for different user group sizes.. *Exhaustive Joint* searches through all user-beam tuple combinations while *Exhaustive Decoupled* performs SUT separated from user selection. Thus, we investigate the performance loss incurs by decoupling beam steering from user selection for different user group sizes.

Scenario. We place 20 users in random locations within a 12×12 m^2 indoor environment. The AP is equipped with $N_{AP} = 24$ antennas and $N_{RF} = 2$ RF chains while users have only one RF chain, one antenna and have LOS connectivity with the AP. We refer to the achievable sum-rate of *Exhaustive Decoupled* and *Exhaustive Joint* as R_d and R_j , respectively. We repeat the same experiment for $N_{RF} = 3$ and 4.

We compare the performance of these two algorithms in Table 3.2 in which the first column shows the scenario and the second column reports the percentage of R_d over R_j . We have simulated 100 different random configurations of users and counted the number of observations in which *Exhaustive Decoupled* provides the optimal user-beam set. We report this count ratio over total number of observations as $Prob(R_j = R_d)$ in the third column.

First, we observe that *Exhaustive Decoupled* achieves more than 95% of the optimal sum-rate with four simultaneous users. However, the performance loss slightly increases as the number of RF chains at the AP increases since the joint methodology is able to reduce the excessive inter-user interference by jointly configuring analog beams and users to be served. Due to computational limitations, we could not inves-

Table 3.2: *Exhaustive Decoupled* vs. *Exhaustive Joint*.

Scenario	$\frac{R_d}{R_j}\%$	Prob ($R_j = R_d$)
$N_{RF} = 2$	98.26	0.63
$N_{RF} = 3$	98.06	0.52
$N_{RF} = 4$	95.79	0.41

tigate scenarios with more than 4 RF chains.

Second, Table 3.2 shows that while $Prob(R_j = R_d)$ is close to 0.5, the ratio of $\frac{R_d}{R_j}$ is more than 95%. While the decoupled approach may not provide the exact optimal set of user-beam tuples, it still can provide close-to-optimal data rate, yet with much lower complexity and overhead time. In such cases, the *Exhaustive Decoupled* algorithm successfully selects the best user combination but the chosen analog beams are off by one or two indices.

Finding: Decoupling beam steering and user selection results in 5% capacity loss compared to joint user-beam selection for a four-user simultaneous downlink transmission. However, the performance loss increases as the group size increases or when only few users have downlink traffic.

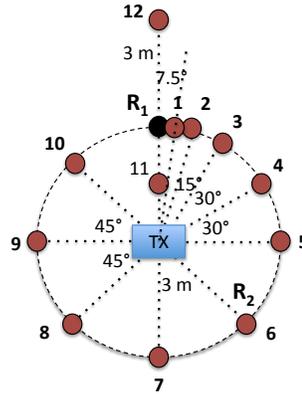
3.6 Receiver Separation and Array Size

In this section, we experimentally characterize the spatial multiplexing gains of beam steering with and without ZFBF. We consider two receivers and investigate empirical capacity as a function of receiver separation as well as array size.

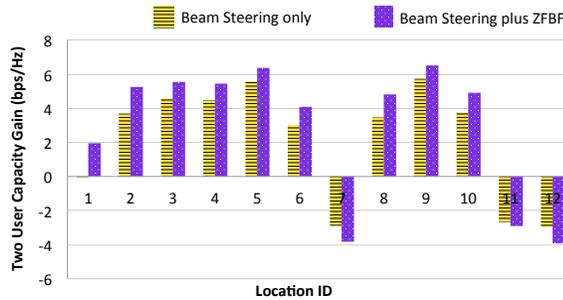
3.6.1 Receiver Separation

Below 6 GHz, multi-user transmission with spatial multiplexing was demonstrated to be achievable even with sub-wavelength receiver separation [59]. However, [59] exploits increased multi-path than is available at 60 GHz and employs one RF chain per *antenna*. Here we explore the required receiver separation for the case of only one RF chain per *user*.

Scenario. We designed a scenario depicted in Fig. 3.9a consisting of an AP with two RF chains simultaneously transmitting to two receivers. The first receiver, R_1 , is



(a) Evaluation scenario



(b) Sum capacity as a function of location

Figure 3.9: Sum-capacity as a function of receiver separation.

at a fixed location, whereas the second receiver, R_2 , is placed at 12 different locations as labeled in Fig. 3.9a. In these experiments, the AP always has LOS components to both receivers. We perform Single-User Training (SUT) using the beam patterns in Fig. 3.3 and find the empirical capacity of single-user transmission to R_1 . This empirical capacity is computed based on the RSS at R_1 using beams chosen in SUT. We validated the RSS of single-user transmission to R_1 and R_2 (in all 12 locations) through measurements using 20° horn antenna at the transmitter. For each of the location IDs in Fig. 3.9a, we perform simultaneous transmission to R_1 and R_2 via beam steering with and without ZFBF. In beam steering only, we skip the the digital training and only apply RF beamforming/combining vectors found via SUT. When turning on ZFBF, the two RF chains are pre-coded with weights computed based

on Eq. (3.8). Although measuring channel matrix and zero-forcing weights is not possible due to hardware limitations, we validated the main components of channels such as RSS, AoA and AoD of the LOS path.

Fig. 3.9b depicts the capacity gain of two-user simultaneous transmission over single-user transmission to only R_1 for the different locations of R_2 . As a baseline, single-user transmission to R_1 achieves 5.8 bps/Hz rate in average. First, the figure shows that when the two receivers are in the same beam as in locations 11 and 12, ZFBF cannot mitigate inter-user interference such that it is preferable to transmit to only one user than to transmit to two users when they are both within the same analog beam. Likewise, the same effect occurs for the backlobe, as location 7 with 180° angular separation also yields a capacity decrease (also cf. Fig. 3.3). Note that here, no strong NLOS path is available to R_1 and R_2 . If they were, the AP might be able to create “separate beam” transmission via NLOS path, even if users are along the same LOS path, and still obtain positive capacity gain from multi-user transmission.

Second, in the cases in which two receivers have the same selected RF beams (locations 7, 11, and 12), employing ZFBF for multi-user transmission is *worse* than purely analog beam forming. The reason is that ZFBF is a sub-optimal precoding scheme which works best in high SINR regimes. (Optimal precoding, Dirty Paper Coding, is difficult to implement due to computational complexity.) When the selected RF beams for two receivers have sufficient overlap, the linear processing at the user will not be sufficient to overcome the high path loss plus the inter-user interference.

Here, beam patterns have half power beam width (HPBW) of 15° . Thus, the angular separation between R_2 and R_1 is greater than the $0.5 \times HPBW$ for locations 2–6 and 8–10 and equals to $0.5 \times HPBW$ for location 1. Third, Fig. 3.9b shows that when the angular separation between two receivers is greater than the $0.5 \times HPBW$,

ZFBF provides an average capacity boost of 27% over beam steering without ZFBF. Indeed, the highest capacity gain belongs to locations 5 and 9 with minimum beam overlap with R_1 .

Finding: Two receivers cannot share an analog beam (less than $0.5 \times \text{HPBW}$ angular separation) even when they also perform inter-user interference cancellation via zero forcing. Indeed, zero forcing was detrimental compared to no digital precoding as it was an ineffective mechanism in the low SINR (high interference) regime.

3.6.2 Antenna Array Scaling and Beamwidth

The width of the beam patterns generated by an analog codebook depends on the number of antenna elements with larger arrays generating higher gain with more focused beams. IEEE 802.11ad allows beams as narrow as 3° [13]. Narrow and highly directional beams not only increase received signal strength at the intended receiver(s), but may also help suppress interference at other receivers. Previously, we observed that in two-user LOS scenario with 15° HPBW, ZFBF can achieve 27% capacity boost provided that receivers are in separate beams. Here, we investigate the impact of beam width (equivalently, number of antenna elements) on the performance gain of ZFBF to mitigate the inter-user interference.

Scenario. We employ the same node deployment setup of Fig. 3.9a. However, instead of finding capacity gain, we study the received SINR at R_1 in a two-user transmission to directly study the impact of ZFBF on the interference reduction. The AP uses 16-element (as above) or 64-element arrays (i.e., $N_{AP} = 16$ or 64) and beam steers with and without ZFBF. The 16-element antenna array can generate beam patterns with approximately 15° HPBW while the 64-element array makes 4° beams. We repeat this experiment for different locations of R_2 (location ID 1 to 5 as depicted in Fig. 3.9a). In all topologies, R_1 and R_2 have LOS connectivity with

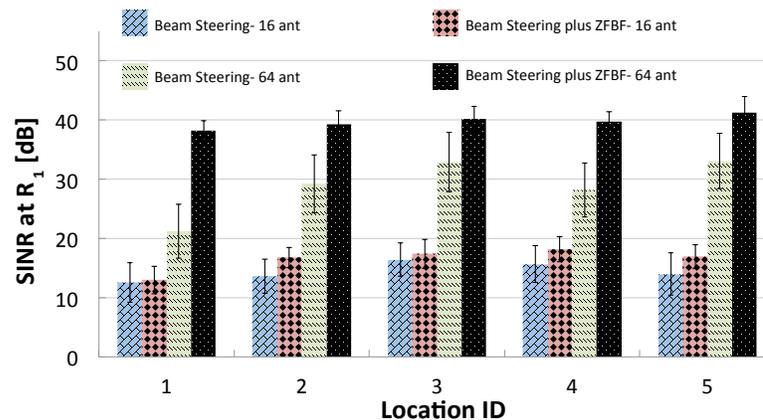


Figure 3.10: SINR at R_1 when grouped with R_2 , which is placed in five different locations.

the transmitter. Fig. 3.10 shows the SINR variation at R_1 when grouped with R_2 , in which R_2 is placed in five locations (as in x-axis).

First, as expected, beam steering with 64 antenna elements provides higher antenna gain, and higher SINR consequently, compared to 16-elements. Even though R_1 is fixed, its SINR is vulnerable to change significantly depending on the location of R_2 when the AP only employs beam steering and forgoes digital training via ZFBF. Furthermore, narrower beams are more vulnerable to SINR change as the figure shows larger error for beam steering only with 64 antennas compared to the 16 antennas.

Second, Fig. 3.10 reveals that ZFBF provides greater SINR boost when beam steering employs a 64-element array (4° beams). Specifically, in location 1, ZFBF achieves $2\times$ SINR gain with 64 antenna element and $1.04\times$ gain with 16 antenna elements. On one hand, increasing the number of antenna elements at the AP reduces the inter-user interference by generating more focused beams; hence, making less opportunity for ZFBF to boost the SINR by canceling out the residual inter-user interference. On the other hand, larger antenna arrays provide higher spatial diversity which can be further exploited by ZFBF to boost SINR. The figure shows that the latter outweighs the former such that ZFBF can provide greater SINR boost, and

consequently capacity boost, even if highly directional beams are adopted at the transmitter.

Finding: ZFBF can provide greater SINR boost, and consequently higher capacity gain, as the transmitter is equipped with larger antenna arrays or equivalently more directional beams in a multi-user transmission, even having the same number of RF chains.

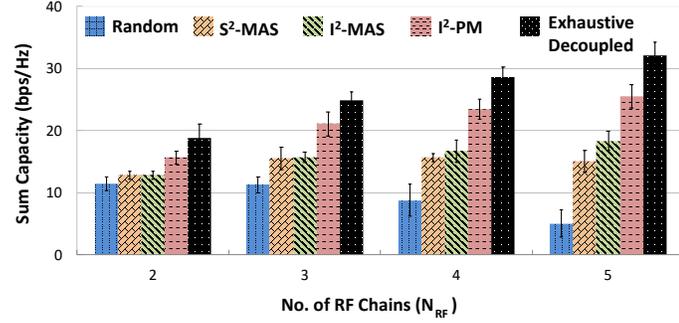
3.7 Evaluation of Decoupled User and Beam Selection

In this section, we evaluate S^2 and I^2 strategies, first focusing on the LOS case and investigating the impact of scaling group-size. Subsequently, we compare the performance of user selection strategies in LOS and NLOS connectivity and study the impact of link connectivity on the spatial multiplexing.

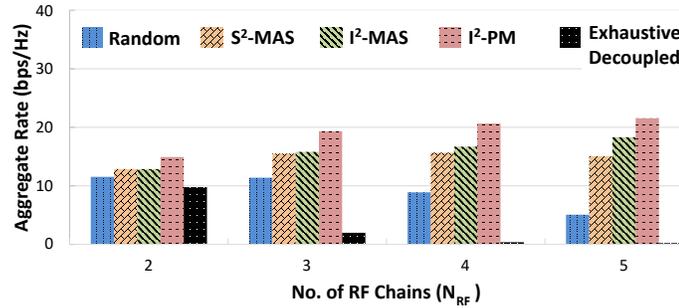
3.7.1 Scaling Group Size (LOS only)

We evaluate S^2 -MAS as a single-shot user selection policy and I^2 -MAS and I^2 -PM as interference-aware incremental policies. For comparison purposes, we also present results for *Exhaustive Decoupled* and *Random* user selection.

Scenario. We deploy ten different node setups with each having 40 users randomly placed in a $6 \times 6 m^2$ indoor environment with the AP in the center and a LOS path between the AP and each user. We equip the AP with $N_{RF} = m$ RF chains, so that it can multiplex up to m data streams to m users (m can be 2, 3, 4, and 5 in this experiment). The AP has a codebook consists of 32 beam patterns while each user is equipped with a single RF chain and 4 beams. For each configuration, we perform Single-User Training (SUT) and employ the selected RF beams for 40



(a)



(b)

Figure 3.11: (a) Achievable sum capacity vs. number of AP's RF chains, (b) Achievable aggregate rate (including user selection overhead) vs. number of AP's RF chains.

sequential user groupings and data transmissions. In our setup, nodal and environmental mobility is negligible such that SUT information is reliable, even when used at a later time. We let user i be the prime (pre-selected) user in multi-user transmission i and perform *Random*, S^2 -MAS, I^2 -MAS, I^2 -PM and *Exhaustive Decoupled* user selection algorithms to group other user(s) with user i . In all cases, we apply ZFBF to suppress inter-user interference and find the average empirical sum-capacity of the simultaneous transmission to the selected group.

Fig. 3.11a depicts empirical sum-capacity averaged over all sub-topologies for different number of AP's RF chains. First, we observe that *Random* is the only technique with decreasing sum-capacity as the number of RF chain increases. This is because *Random* aggregates as many users as RF chains, N_{RF} , without considering their mu-

tual interference. While for 2 RF chains, random user selection often selects users in separate beams, with 3, 4, and 5 simultaneous users, this is increasingly unlikely. Moreover, despite having an increased number of RF chains to support digital beamforming, ZFBF is unable to correct for the poor user grouping, as was also the case with two RF chains in Section 3.6. Hence, *Random* selection achieves 61% of the maximum sum-capacity (via *Exhaustive Decoupled*) with 2 RF chains, and only 16% with 5 RF chains.

Second, the S^2 -MAS policy yields only a marginal improvement with an increasing number of RF chains beyond 2, and even slightly degrades sum capacity with 5 RF chains. Like *Random*, S^2 -MAS always uses the maximum group size (the number RF chains at the AP) and employs ZFBF to cancel out any residual interference. Yet unlike *Random*, the S^2 -MAS policy mitigates inter-user interference by selecting users according to their RF beams (and not randomly). Unfortunately, for larger groups, the interference can be excessively high: With users randomly located in the environment, there is often no user in a position which is sufficiently angularly separated from already selected users. While for 2 RF chains, this probability is low, with 3, 4, and 5, this becomes increasingly problematic such that with 5 RF chains, S^2 -MAS achieves only 47% of the capacity of Exhaustive Decoupled approach. Furthermore, Fig. 3.11a shows that I^2 policies I^2 -MAS and I^2 -PM, never lose sum-rate due to an additional RF chain at the AP. By design, I^2 schemes add users sequentially and evaluate the net effect of adding another user before data transmission.

Next, we incorporate the user selection overhead and compare the achievable aggregate rate of different user selection policies in Fig. 3.11b. Note that all decoupled user selection strategies have the same beam training overhead; hence, in this figure, we show the aggregate rate based on the sum-capacity (presented in Fig. 3.11a) and user selection overhead only. We discussed the user selection search space size

corresponding to each strategy in Table 3.1. Translating search space size to overhead is standard dependent. Here, we use IEEE 802.11ay parameters (e.g., preamble length) and approximate the time overhead to inquiry one user to be $0.15ms$. Therefore, the total user selection overhead can be calculated according to Table 3.1. In Fig. 3.11b, the aggregate rate is simply found by dividing the number of bits sent in one beacon interval ($100ms$ in IEEE 802.11ay) to the total time (overhead + $100ms$ of data transmission time). Fig. 3.11b reveals that *Exhaustive Decoupled* loses to all other strategies due to overwhelming user selection overhead, despite providing higher sum-capacity. We can observe that *Exhaustive Decoupled* rate exponentially degrades with increasing number of RF chains due to exponential growth in the search space size. With larger group sizes (e.g., $N_{RF} = 5$), *Random* and *Exhaustive Decoupled* both perform poorly, one due to excessive interference and the other due to excessive overhead.

Finding: While both S^2 and I^2 policies perform nearly equally well with two RF chains, their performance gap with exhaustive search increases as the number of RF chains at the AP increases. However, if incorporating the user selection overhead, S^2 and I^2 easily outperform due to the large search space size of *Exhaustive Decoupled* approach.

3.7.2 LOS vs. NLOS Connectivity

Here we consider scenarios in which LOS path is unavailable due to blockage and therefore the AP must connect via a reflected path.

Scenario. We adopt the same node deployment setup except that we block the LOS path from the AP to each user. Consequently, the AP finds a reflected path (e.g., off of a table or wall) during SUT. Fig. 3.12 depicts the empirical capacity for an AP with 4 RF chains. In addition to the NLOS scenario, it also depicts LOS

as a baseline. The figure is normalized to the capacity of *Exhaustive Decoupled* for the respective scenarios. Fig. 3.12 indicates that all user selection schemes perform closer to *Exhaustive Decoupled* in LOS than NLOS: In the NLOS scenario, because the received signal strength is lower, there is less residual tolerance for inter-user interference. Therefore, it is more important to apply a user selection strategy which selects users with minimum inter-user interference in such scenarios, and all policies end up far from what can be achieved by *Exhaustive Decoupled* in the NLOS case.

As observed in Fig. 3.12, with NLOS, S^2 -MAS and I^2 -MAS can only achieve 23% and 30% of *Exhaustive Decoupled* sum-capacity. These strategies select users according to their beam separation to mitigate interference at receivers. In NLOS case, the RF beam is mostly locked into the strongest NLOS path. In such scenarios, the high beam separation does not necessarily mean low inter-user interference. For example, in an extreme example, two users which are co-located may see two strong NLOS paths (e.g., off two walls) from the AP. While in SUT, each can lock its RF beam to a different NLOS paths, both users still experience significant interference from the other NLOS path. This would not be an issue in LOS scenario since the RF beam is mostly locked to the LOS path in SUT which is significantly stronger than any NLOS path. Thus, beam separation based policies perform poorly in NLOS scenarios.

Even I^2 -PM which provides 82% of *Exhaustive Decoupled* capacity in LOS connectivity, achieves only 42% of the maximum sum-rate in NLOS. I^2 -PM searches through a subset of users (i.e., $\lceil U/N_{RF} \rceil = 10$ users here) in each round. Considering the lower RSS of NLOS users, it becomes increasingly likely that I^2 -PM cannot find a user (out of 10 users) which improves the sum-capacity and makes it terminate the multi-round procedure.

Finding: User grouping algorithms that select users according to RF beam separa-

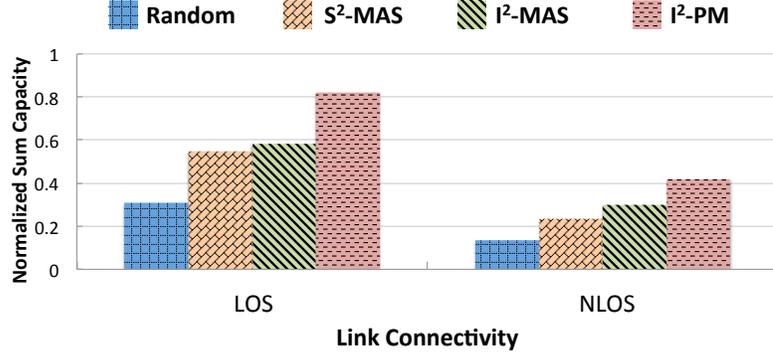


Figure 3.12: Normalized capacity of S^2 and I^2 policies under LOS and NLOS link connectivities with a four-RF chain AP.

tion are vulnerable to performance loss in NLOS environments in which the dominant LOS path is blocked. NLOS users can tolerate less inter-user inference as their RSS is reduced due to lack of a LOS path.

3.7.3 Maximum Multiplexing Potential: Mixing LOS and NLOS

We now investigate the impact of link connectivity on the spatial multiplexing potential quantified via two metrics: (i) maximum achievable capacity via *Exhaustive Decoupled* search; (ii) the maximum number of users that can be simultaneously served. To find the impact of link connectivity, we define a new factor termed *LOS Probability* which characterizes the fraction of users that have a LOS path to the AP.

Scenario. We deploy 10 different node setups in which 32 users are placed randomly in a $6 \times 6 m^2$ environment. The AP is equipped with 32 RF chains and its RF codebook consists of 32 beams (i.e., $U = N_{RF} = 32$). Users have a single RF chain with four beam patterns. Therefore, the AP has a sufficient number of RF chains to serve all users concurrently. For each node deployment topology, we vary the *LOS Probability* from zero to one by blocking a subset of paths, zero indicating that all

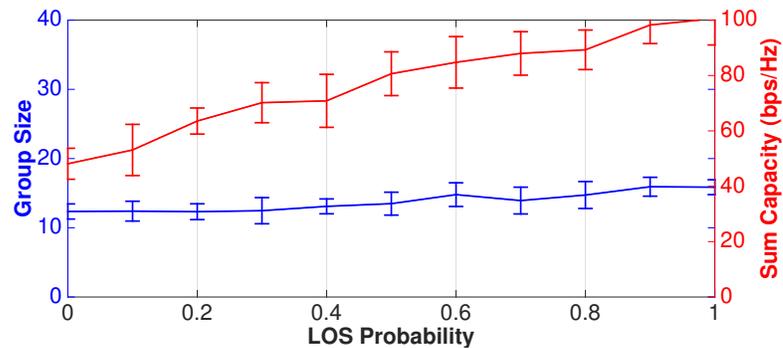


Figure 3.13: The sum capacity and group size of Exhaustive Decoupled search as a function of LOS probability.

users have only NLOS paths.

Fig. 3.13 depicts the number of simultaneous users (group size) and the empirical sum-capacity vs. *LOS probability*. The figure reveals that even with only NLOS (LOS probability of 0), the AP is capable of serving 12 users for multi-user concurrent transmission, i.e., 13 users would have yielded lower capacity due to excessive interference. While the number of simultaneous users varies slightly with LOS probability, we observe that the empirical sum-capacity of *Exhaustive Decoupled* increases roughly linearly with LOS probability. Namely, LOS users have higher RSS and per-user data rate, so that even with the same group size, the achievable sum-rate in LOS scenario would be almost two times of the NLOS.

Finding: Even though NLOS users experience reduced SNR, 12 spatial streams can be multiplexed via Exhaustive Decoupled search for LOS, NLOS and any mix of the two.

3.8 Related Work

Multi-user Beamforming. Prior work on multi-user beamforming in millimeter-wave networks focuses on developing low-complexity algorithms that can achieve near-optimal capacity performance [4,57,64,65]. Such work does not address user selection,

and instead aims to maximize the capacity of simultaneous transmission to a given group of users by adjusting digital and analog weights. Consideration of user selection would create a dependency in which the choice of RF beamforming affects user selection, and consequently, the digital precoder to cancel out the inter-user interference of the selected users. In contrast, in this work, we do not design new analog/digital beamforming algorithms; instead, we design and analyze user and beam selection procedures that can be applied to a broad class of beamforming algorithms.

User Selection. Extensive prior work has addressed user grouping in MU-MIMO systems which operate below 6 GHz. Example results include user grouping based on channel state information and/or expected transmission time [66–68]. Likewise, other work has targeted user grouping without channel state information by exploiting the rich scattering propagation in indoor environments below 6 GHz [69, 70]. In contrast, we consider both a different frequency band and node architecture: 60 GHz channels lack the rich scattering properties observed below 6 GHz [35]; moreover, as described in Section 3.1, we consider that each RF chain has multiple antennas available to it which can be controlled via analog beam steering weights such that the SUT phase considered here does not exist in the aforementioned prior work. Availability of multiple antenna elements per single RF chain makes the acquisition of channel state information to incur more overhead and thus makes the user grouping based on channel state information to be impractical for many scenarios. Authors in [71] suggest a channel-based user grouping protocol for 60 GHz WLANs such that the AP adopts an incremental strategy to select users for a simultaneous multi-user transmission. However, analog beams for each user are selected independently; hence, this user selection is in fact an example of I^2 user selection class.

Joint User-Beam Selection. Prior work studied the problem of joint user and analog beam selection in two contexts. First, in a multi-cell millimeter-wave

network wherein each cell is equipped with a set of analog beams. A key problem in such networks is to determine the set of users each cell should serve and the group of beams it should employ. However, the choice of beams at any user is coupled with the choice of serving cell; likewise, the user association to cells depends on the employing beams at each cell [72–74]. These prior works formulated the problem of user-beam selection in multi-cell millimeter-wave networks maximizing a network utility function under practical considerations. However, in this chapter, we consider the problem of joint user and beam selection for downlink multi-user transmission in a single cell and experimentally compare it with the decoupled framework wherein the selection of beams precedes user grouping.

Very few works have studied the problem of joint user and beam selection in the context of hybrid beamforming in multi-user single-cell systems. For example, authors in [75] formulated a joint user and beam selection problem based on the power and AoA of the LOS path for every user. This formulation is a simplified version of equation (3.3) in Section 3.1 in which instead of collecting the channel state information, the AP only acquires the power and the AoA information for every user and selects users with very different strongest AoA peaks in one group. To reduce the overwhelming overhead associated with collecting power and AoA for every user, authors in [75] suggested a multi-carrier system assigning different carriers to different users. Instead of direct optimization of choice of users and beams in a complex multi-carrier system, we proposed decoupling user and beam selection and experimentally showed that the incurred performance loss is less than 5% for a four-RF chain system.

Sensing-Assisted Spatial Multiplexing in mmWave Networks

In the previous chapter, we discussed the fundamentals and frameworks for user and beam selection in multi-user directional networks. In this chapter, we introduce MUTE, Multi-stream beam-Training for mm-wavE networks, a novel system that enables efficient simultaneous transmissions to a given group of users via sensing spatial paths across different analog beams. MUTE has two key components: *(i)* a path sensing module that identifies dominant paths between the AP and each client under the commercial RF front end and phased array architecture; and *(ii)* a diversity steering module that integrates the knowledge of RF codebook and the discovered paths to form multi-stream beam configurations over diverse or ideally orthogonal paths, such that undesired channel correlations are minimized; thus, the spatial multiplexing gain is maximized.

Our design is motivated by two key observations: *(i)* Unlike legacy MIMO in sub 6 GHz bands that are privileged with high multiplexing gain as a result of rich scattering, mmWave channels are sparse, i.e., only a few dominant Line of Sight (LOS) and reflected Non-LOS (NLOS) paths characterize the channel between any

two nodes [34, 76]. Furthermore, exploiting an analog beam may impact the mmWave channel as it amplifies certain paths and weakens others. Hence, multiplexing independent streams should avoid common paths as it will otherwise incur throughput degradation due to channel correlation and inter-stream interference. (ii) While a “perfect” non-overlapping set of beam patterns would ensure that use of different beam codebooks would yield non-overlapping paths, practical 60 GHz beams generated via phased array antennas have irregular beam patterns [1, 16]. Nonetheless, despite their irregularity, the directivity gain is known *a priori* in each direction as it is a deterministic function of the codebook and antenna spacing. We exploit these two properties, combined with GHz-scale sampling rate to design MUTE. In particular, we make the following contributions:

First, MUTE repurposes beam acquisition sweeps, which occur periodically in 60 GHz WLANs to establish and maintain a directional link between the AP and each client, to estimate the Power Delay Profile (PDP) of each beam with *zero* additional overhead. In contrast to sub-6 GHz bands, the GHz-scale sampling rate and sparsity of 60 GHz channels provide the unique capability to obtain the high-resolution PDP. While PDP reveals the presence of multiple paths as well as their relative timing, it does not convey any direction information and cannot solely identify orthogonal paths across different clients. Thus, MUTE next couples the known radiation patterns over the suite of irregular beam patterns with PDP measurements for each pattern in order to infer the direction of each path. Lastly, MUTE leverages these direction inferences to construct a candidate set of transmit and receive beams over diverse or ideally orthogonal paths for a multi-stream transmission which will be further trained.

Leveraging the design of MUTE, we present the *first experimental exploration of MIMO beam steering* in mmWave networks. We implement key components of MUTE on *X60*, a programmable testbed for wideband 60 GHz WLANs with electronically-

steerable phased arrays [1], and modify it to access link-level statistics such as SNR, channel, and PDP in fine-resolution. We collect channel samples (in time and frequency domains) from over-the-air measurements in an indoor setting and subsequently perform trace-driven emulations. Our key findings are as follows:

It may seem that high SNR beams are always good candidates for multi-stream beam steering as they focus the signal energy towards the intended receiver and that digital precoding can mitigate or ideally cancel any residual interference. However, our results reveal that this is not the case with practical beam patterns since high SNR beams might share a dominant path causing high channel correlation as a result of sparse scattering. This is due to the irregularity of beam patterns generated by phased arrays including the presence of partial overlap among different beams in the RF codebook as well as strong side lobes that cause a particular path to be captured via multiple beams, yet with different directivity gains. We show that in such cases digital precoding methods such as zero-forcing are of little help. That is, digital precoding cannot compensate for a *bad* choice of analog beams that obtain low stream separability in the analog domain.

In contrast, MUTE achieves 90% of the maximum aggregate rate for both single-user and multi-user MIMO, with only 0.04% of the training overhead compared to exhaustive search. In particular, MUTE maximizes stream separability in the analog domain by selecting a candidate subset of beams that capture diverse or ideally orthogonal paths. Although this candidate beam selection itself does not require any additional signaling overhead, discovering the final choice of beams out of this candidate set requires further training with overhead proportional to the candidate set size. MUTE targets that the candidate set size be in the order of the number of LOS/NLOS paths, which is small due to sparse scattering in 60 GHz band. Hence, MUTE approximates the PHY throughput of exhaustive search, while searching over

only a few beams with diverse paths.

The rest of this chapter is organized as follows. Section 4.1 provides a primer on IEEE 802.11ay, the standard that allows for multi-user MIMO in 60 GHz WLANs. Section 4.2 presents the design of MUTE. Section 4.3 introduces the experimental platform. Our benchmarking algorithms are introduced in Section 4.4. Section 4.5 describes the experimental evaluation and Section 4.6 reviews the related work.

4.1 A Primer on IEEE 802.11ay

While local and personal wireless technologies have greatly evolved in the last few years [14, 77], new applications and continued usage growth demand greater throughput and reliability with lower latencies. AR/VR applications, mobile offloading, high-bandwidth connectivity to multiple TV and monitor displays, and indoor and outdoor wireless backhaul are just a few applications that require new wireless technologies. To meet the demanding requirements of such diverse applications, the IEEE 802.11 Task Group ay (802.11ay) was formed in 2015 to define PHY and MAC amendments to the 802.11 standard that enable Wi-Fi devices to achieve 100 Gbps using the unlicensed mm-Wave (60 GHz) band at comparable ranges to today's commercial 60 GHz devices based on the 802.11ad standard¹.

As illustrated in previous chapters, 60 GHz transmissions must be directional to take advantage of beamforming gains and cope with increased path loss (e.g., 22 dB for 10 meters) and other propagation losses compared to sub 6 GHz bands. As discussed in [14], in order to support highly directional transmissions, in place of quasi-omni, IEEE 802.11ad redefined fundamental principles of Wi-Fi systems and incorporated innovative techniques and procedures to overcome unique challenges associated with

¹According to the official IEEE 802.11 Working group timeline, the draft of the IEEE 802.1ay amendment is scheduled to go for a vote (WG letter ballot) in November 2017.

mm-Wave propagation. IEEE 802.11ad supports transmission rates of up to 8 Gbps using single-input-single-output (SISO) wireless transmissions over a single 2.16 GHz channel.

IEEE 802.11ay, the next-generation Wi-Fi standard for the 60 GHz band, increases the peak data rate to 100 Gbps through supporting multiple independent data streams and higher channel bandwidth, among other advancements, while ensuring backward compatibility and coexistence with Directional Multi-Gigabit (DMG) stations (STAs). We use the terms DMG and Enhanced DMG (EDMG) stations to refer to devices that can support features of IEEE 802.11ad and IEEE 802.11ay standards, respectively.

Channel Bonding and Aggregation. The band allocated to unlicensed use around 60 GHz has approximately 14 GHz of bandwidth, which is divided into channels of 2.16, 4.32, 6.48, and 8.64 GHz bandwidth. The channel center frequencies for the 2.16 GHz channels are: 58.32, 60.48, 62.64, 64.80, 66.96, and 69.12 GHz for channel numbers 1 through 6, respectively. Unlike IEEE 802.11ad, which only allows for single (2.16 GHz) channel transmission, 802.11ay includes mechanisms for channel bonding and aggregation; in channel bonding, a single waveform covers at least two contiguous 2.16 GHz channels, whereas channel aggregation has a separate waveform for each aggregated channel. IEEE 802.11ay mandates that EDMG STAs must support operation in 2.16 GHz channels as well as channel bonding of two 2.16 GHz channels. Channel aggregation of two 2.16 GHz or two 4.32 GHz (contiguous or non-contiguous) channels and bonding of three or four 2.16 GHz channels is optional.

Directional MIMO Communication. In typical DMG implementations, one or more phased arrays are driven by a single Radio Frequency (RF) chain and thus only a single data stream is transmitted at a time. Therefore, the multiple antenna elements used by DMG STAs only provide beamforming gain but not multiplexing gain. To achieve both beamforming and multiplexing gain, IEEE 802.11ay defines

new mechanisms to enable MIMO operation including both Single-User MIMO (SU-MIMO) and downlink Multi-User MIMO (MU-MIMO). The maximum number of spatial streams per station is eight, and downlink MU-MIMO transmission can be made to up to eight stations.

EDMG STAs may also use digital pre-coding at baseband to compliment analog beamforming to minimize or ideally cancel the inter-stream interference in MIMO transmissions. Such hybrid analog/digital beamforming architectures have been studied in the literature employing single-polarized phased antenna arrays [2, 58]. IEEE 802.11ay also supports stations with dual-polarized antenna arrays. The use of polarization is of great value in mm-Wave communications since, for example, it allows for diversity gains and spatial multiplexing in LOS environments. To obtain spatial and polarization separation, signal streams must be independently steerable and be transmitted and received with different polarizations. Experimental results have shown that practical phased antenna arrays can have a cross polarization discrimination factor of approximately -24 dB [6]. Therefore, in a 2x2 SU-MIMO configuration with dual-polarized antenna array, both streams can operate under LOS conditions with orthogonal horizontal and vertical polarizations. In general, the number of streams that a given MIMO link supports is determined by different factors including the environment, the directivity of the antenna used, and on whether antenna polarization is exploited.

4.1.1 IEEE 802.11ay Physical Layer (PHY) Overview

Building upon the DMG PHY, IEEE 802.11ay defines a new PHY specification that includes both single carrier (SC) and orthogonal frequency division multiplexing (OFDM) modulations. As described in this section, to support MIMO transmissions and channel bonding while guaranteeing backward capability, a new packet structure

is defined in IEEE 802.11ay. The EDMG packet contains new fields necessary to support the additional capabilities defined for EDMG stations, as well as a redefined training (TRN) field that is more flexible and efficient than the one defined in IEEE 802.11ad.

4.1.1.1 EDMG Packet Format

A single packet format is defined for the three EDMG PHY modes: SC, OFDM, and control. This packet is shown in Fig. 4.1 with all of its possible fields. Not all fields are transmitted in an EDMG packet: Fields are included depending on whether the packet is used for single channel or channel bonding operation, for SISO or MIMO transmission, and if it is used for beamforming training/tracking.

To enable backward compatibility, the first portion of an EDMG packet, referred to as non-EDMG portion, is defined to be recognizable by DMG stations. The L-STF (legacy-short training field) and L-CEF (legacy-channel estimation field) are compatible with the preamble defined in IEEE 802.11ad, and enable detection of the packet and acquisition of carrier frequency and timing. The L-Header field is the same as the header field in an IEEE 802.11ad packet, with the exception that some of its bits are re-defined.

The second portion of an EDMG packet, referred to as the EDMG portion, includes fields that are only recognized by EDMG STAs. The EDMG-Header-A field carries information required to interpret EDMG packets, including bandwidth, Modulation and Coding Scheme (MCS), and number of spatial streams. The EDMG-STF and EDMG-CEF fields enable EDMG stations to estimate various signal parameters and the channel when channel bonding and/or MIMO are utilized. The EDMG-Header-B is only included in MU-MIMO packets.

The non-EDMG portion of the packet together with the EDMG-Header-A is trans-

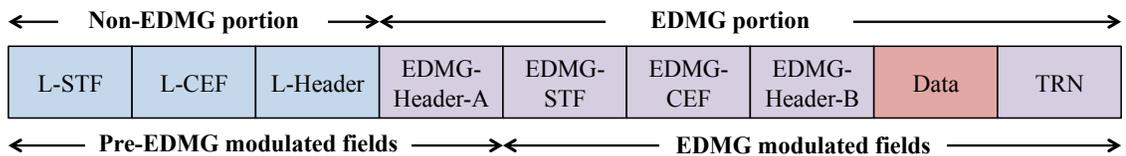


Figure 4.1: IEEE 802.11ay packet structure.

mitted with the IEEE 802.11ad rate. When channel bonding is utilized, these fields are transmitted in duplicate mode over the channels being bonded. In MIMO transmissions, identical copies of these fields are transmitted in each stream with different cyclic shifts. Due to these characteristics, the first four fields of an EDMG packet are said to be pre-EDMG modulated. The remaining fields, which are said to be EDMG modulated, are transmitted in bonded and/or MIMO mode when these features are employed.

4.1.1.2 Data Field Format

The data field consists of the payload data and possible padding. The bits to be transmitted are padded with zeros if necessary, scrambled, encoded, and modulated according to an EDMG MCS. Then, symbols are grouped and each group is prepended by a modulated Golay sequence, forming a block. The SC block size consists of $512 \times N_{CB}$ symbols for both SISO and MIMO transmissions, where N_{CB} is the number of utilized 2.16 GHz channels. The control mode, which corresponds to MCS 0 in both DMG and EDMG PHYs, enables low SNR operation prior to beamforming with BPSK modulation and a spreading factor of 32. The other MCSs defined in IEEE 802.11ay for the SC mode are based on BPSK, QPSK, 16 QAM, and 64 QAM modulations and LDPC codes with rates of $1/2$, $5/8$, $3/4$, $13/16$, and $7/8$. The achievable data rate for each MCS index depends on the number of spatial streams, N_{ss} ($1 \leq N_{ss} \leq 8$), and the number of 2.16 GHz channels, N_{CB} ($1 \leq N_{CB} \leq 4$),

used in the transmission of the packet.

4.1.1.3 Training (TRN) Field Format

The TRN field enables transmit and receive beamforming training and is appended to packets used in a beam refinement protocol (BRP). BRP is a process in which a station can improve its antenna configuration for transmission and/or reception. The TRN field was redesigned in IEEE 802.11ay to increase efficiency and make it configurable based on the characteristics of the particular beamforming training procedure being executed.

The “basic unit” of the TRN field is the TRN subfield, which is composed of 6 Golay complementary sequences. By concatenating a variable number of TRN subfields, a TRN-Unit is formed. The format of a TRN field is shown in Fig. 4.2. The TRN field is composed of a variable number of TRN-Units, which is defined by the parameter EDMG TRN Length, denoted by L in Fig. 4.2.

In a BRP procedure used for receiver training, all TRN subfields are transmitted with the same Antenna Weight Vector (AWV) as the data field. Such packets are referred to as EDMG BRP-RX packets in [78] and their TRN field structure is shown in Fig. 4.2a. As defined in [78], AWV is a vector of weights describing the excitation (amplitude and phase) for each element of an antenna array. This configuration allows the receiver to switch AWVs when receiving the different TRN subfields and thus searches for an improved antenna configuration setting.

In a BRP procedure used for transmitter training, the transmitter uses different AWVs in the transmission of the TRN field while the receiver uses the same AWV in its reception. The TRN field structure for transmit training is shown in Fig. 4.2b. As shown in this figure, three parameters define the format and length of a TRN-Unit used for transmit beamforming training (referred to as EDMG BRP-TX packets

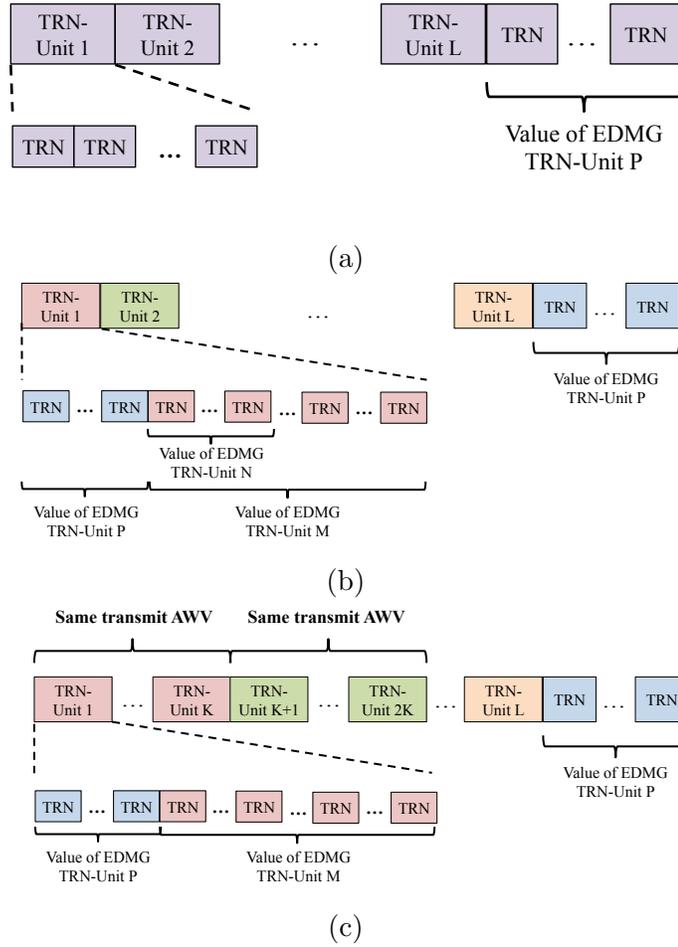


Figure 4.2: EDMG TRN field structure for (a) EDMG BRP-RX packets, (b) EDMG BRP-TX packets, and (c) EDMG BRP-RX/TX packets.

in [78]): EDMG TRN-Unit P, EDMG TRN-Unit M, and EDMG TRN-Unit N, which are referred to as P, M, and N in what follows for ease of notation. In a TRN-Unit, the first P TRN subfields are transmitted with the same AWV as the data field. Therefore, the receiver may use such TRN subfields to maintain synchronization and estimate the channel. In the transmission of the remaining M TRN subfields of a TRN-Unit, the transmitter may change AWV at the beginning of each TRN subfield. In order to improve the robustness of the beamforming training process, of the last M TRN subfields of a TRN-Unit, more than one consecutive TRN subfield may be transmitted

with the same AWV. The number of consecutive TRN subfields transmitted with the same AWV is N .

To enable simultaneous training of the transmitter and receiver, a given TRN-Unit may be re-transmitted a number of times. In this case, the same AWV is used in the transmission of the last M TRN subfields of a given TRN-Unit, and the same TRN-Unit is repeated a number of times. Such packets are referred to as EDMG BRP-RX/TX packets in [78] and are shown in Fig. 4.2c. The number of TRN-Units transmitted with the same AWV is given by the parameter RX TRN-Units per Each TX TRN-Unit, which is noted as K in Fig. 4.2c. In such packets, the value of N is not applicable. For all EDMG BRP packets, following the transmission of all TRN-Units, there are P repetitions of the TRN subfield to allow the receiver to track frequency offset for the last transmitted TRN-Unit.

4.1.2 IEEE 802.11ay Medium Access Control Layer

This section describes the main changes made to the IEEE 802.11ad MAC layer specification to support MIMO transmission and multi-channel operation.

4.1.2.1 Beacon Interval

IEEE 802.11ay organizes access to the medium in Beacon Intervals (BIs), similar to 802.11ad. Fig. 4.3 illustrates a typical BI consisting of two main access periods: Beacon Header Interval (BHI) and Data Transmission Interval (DTI). The BHI enables beam training of unassociated DMG and EDMG STAs and network announcements through a sweep of multiple directionally transmitted frames. The BHI is further subdivided into three sub-intervals: 1. **Beacon Transmission Interval (BTI)** used by the AP or the personal basic service set control point (PCP) for transmission of beacon frames; 2. **Association Beamforming Training (A-BFT)** used

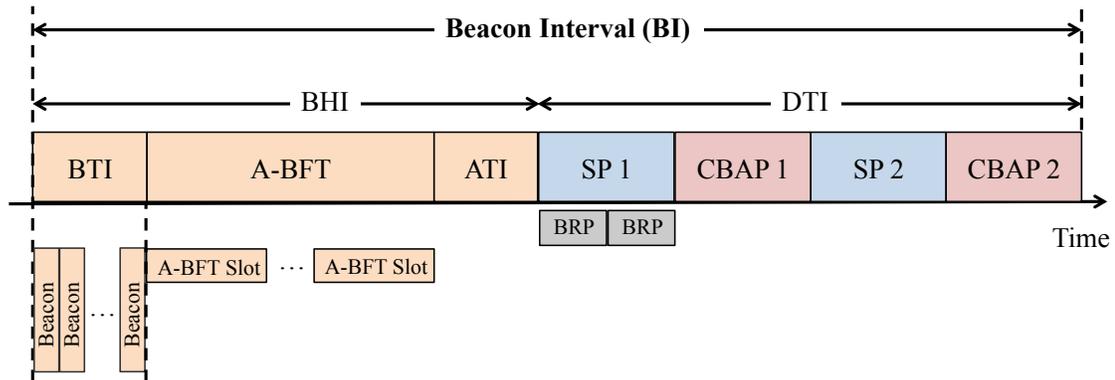


Figure 4.3: IEEE 802.11ay beacon interval structure.

by DMG/EDMG STAs to train their receive antenna configurations; 3. **Announcement Transmission Interval (ATI)** used for management frame exchange between the AP/PCP and beam-trained stations.

A-BFT is slotted (up to 8 slots for 802.11ad) and stations randomly choose one of the slots for transmitting their sector sweep (SSW) frames; consequently, collisions may occur when more than one STA choose the same slot. To accommodate a larger number of STAs attempting access during A-BFT, IEEE 802.11ay supports up to 40 A-BFT slots in each BI.

The BHI is followed by the DTI, which facilitates different types of medium access for data transmission and beamforming training. In the DTI, data frames can be exchanged either in contention-based access periods (CBAPs) or scheduled service periods (SPs) for contention-free communications. A description of IEEE 802.11ad channel access rules can be found in [14].

4.1.2.2 Multiple Channel Access

EDMG transmissions always involve a primary channel with 2.16 GHz bandwidth to maintain compatibility with 802.11ad. EDMG stations can occupy secondary 2.16

GHz, 4.32 GHz or 6.48 GHz channels which might be adjacent or non-adjacent to the primary channel. EDMG STAs must be capable of performing physical and virtual carrier sensing (i.e., network allocation vector or NAV) in the primary channel and at least energy detection in the secondary channels. Physical carrier sensing is a measurement of the received signal strength of an incoming Wi-Fi signal preamble, whereas energy detection determines if the medium is busy by measuring the total energy received at the station, regardless of whether it is a valid Wi-Fi preamble or not. Both methods compare the measurements with a pre-defined threshold.

To enable the coexistence of DMG and EDMG STAs, network announcement and management frames need to be transmitted through the primary channel; hence, the BHI is present on the primary channel. IEEE 802.11ay supports the presence of A-BFT on secondary channels to provide more slots for contention-based transmission of SSW frames in dense use cases. Transmissions within the DTI (CBAP or SP) can use more than one channel or be performed over a bonded channel. IEEE 802.11ay supports channel access over multiple channels through scheduling and within Transmission Opportunity (TXOP). With scheduling, the AP/PCP specifies the channel width for the following DTI, whereas in TXOP, STAs expand their bandwidth opportunistically when secondary channels are idle. Next, we elaborate on these two approaches.

Multiple Channel Access Through Scheduling. The AP/PCP can allocate aggregated and bonded channel(s) using the EDMG Extended Schedule Element (ESE), which can be transmitted by the AP in the BTI. Fig. 4.4 depicts four possible channel allocations. When the used channels are adjacent, both channel bonding and channel aggregation are possible (allocation #1); however, channel aggregation is the only option for allocation #2 since it includes non-adjacent channels. An EDMG AP/PCP can also schedule allocations over different channels overlapping in time

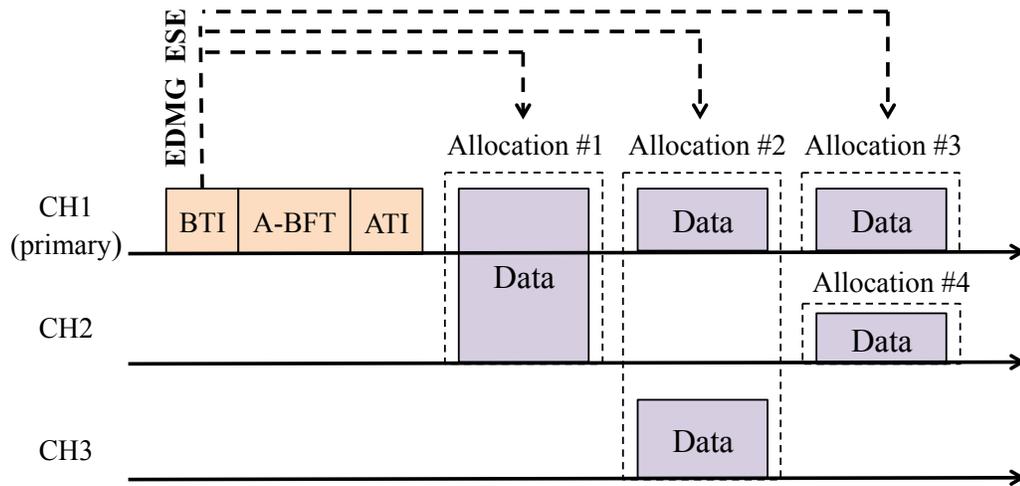


Figure 4.4: Example of IEEE 802.11ay multiple channel access via scheduling.

(e.g., allocations #3 and #4). As shown in Fig. 4.4, allocation #4 does not include the primary channel. Such allocations that do not include the primary channel are limited to a single 2.16 GHz channel. If the allocation is a CBAP that does not include the primary channel, the STAs must perform full carrier sensing in the secondary channels. In addition, when the AP/PCP is the transmitter or receiver in an allocation, the allocation must include the primary channel so that the transmission would not be hidden from the DMG STAs.

Multiple Channel Access Through TXOP. Medium access in CBAPs follows Enhanced Distributed Channel Access (EDCA) rules in which the AP/PCP and STAs obtain a transmission opportunity (TXOP) either by winning an instance of EDCA contention or by receiving a Grant frame. When EDMG STAs and AP/PCP support multiple channel widths, an EDCA TXOP is obtained based solely on the activity of the primary channel, i.e., if the primary channel is idle. However, the width of the transmission is determined by the occupancy status of the non-primary channels. Hence, the TXOP initiator monitors the status of its primary channel during the PIFS (Point Coordination Function Inter-frame Space) interval immediately preceding the expiration of the backoff counter to determine which secondary channels are idle.

Once the TXOP initiator finds the primary and secondary channels to be idle, it performs an RTS/DMG CTS exchange to inhibit collision on the secondary channel(s). The TXOP initiator sends RTS frames in the primary and secondary channels if they are determined to be idle to determine the available bandwidth at the responder. The TXOP responder then transmits a DMG CTS in the idle channels in order to help the TXOP initiator determine which channels are available for data transmission.

4.1.2.3 MIMO Channel Access

To perform a MIMO transmission, the transmitter must obtain a TXOP. To this end, EDMG STAs that support MIMO maintain physical and virtual carrier sensing and perform the backoff procedure. The MIMO channel is said to be idle when all the MIMO transmit antennas intended to be used in the TXOP (determined by MIMO beamforming protocols defined in Section 4.1.3) are sensed to be idle for a period of PIFS before the backoff timer reaches zero. In this case, the EDMG STA is permitted to obtain a TXOP for a SU-MIMO transmission. In downlink MU-MIMO, the AP/PCP is the transmitter, which needs to obtain a TXOP.

Before the transmitter accesses the channel, it must indicate to one or more EDMG STAs its intention to transmit an SU-MIMO or a MU-MIMO packet to them. To this end, the EDMG transmitter can send a RTS frame, a DMG CTS-to-self frame, or a Grant frame to the intended EDMG STAs. This frame indicates whether the following transmission is SU-MIMO or MU-MIMO and also the antenna configuration to be used. The receiving EDMG STAs can infer the operating channel number and bandwidth from this frame and respond with a DMG CTS frame or an ACK frame. This response frame confirms the availability of the STA for MIMO reception and protects it from hidden STAs.

4.1.3 IEEE 802.11ay Beamforming Protocol

Beamforming (BF) training is used to determine the appropriate transmit and receive antenna configurations for a pair of stations through two sub-phases [14]: (1) Sector Level Sweep (SLS) enables communication between two participating stations at the control mode rate or higher MCS. Normally, the SLS phase provides only transmit beamforming training. (2) BRP enables receive training and iterative refinement of the AWW at both participating STAs. IEEE 802.11ay includes several new beamforming training protocols, including SU-MIMO/MU-MIMO beamforming training, BRP transmit sector sweep, and beamforming for asymmetric links. These BF procedures are performed in the DTI after STAs have an established link. In this section, we focus on MIMO beamforming and a description of other IEEE 802.11ay BF procedures can be found in [78].

4.1.3.1 SU-MIMO Beamforming

The SU-MIMO BF protocol determines transmit and receive antenna configurations for simultaneous transmission of multiple spatial streams between two SU-MIMO capable EDMG STAs. The SU-MIMO beamforming protocol consists of two consecutive phases: SISO phase and MIMO phase.

1) SISO Phase: In this phase, both STAs collect the necessary feedback for possible candidate sectors, some of which are then used in the following MIMO phase. The station that initiates the beamforming training is called the initiator, and the other the responder. All transmissions in this phase use the DMG control mode to extend the range. Fig. 4.5a depicts the SISO phase, which comprises three sub-phases: an optional initiator transmit sector sweep (I-TXSS), an optional responder transmit sector sweep (R-TXSS) sub-phase and a mandatory SISO feedback sub-phase.

The optional I-TXSS and R-TXSS subphases allow the responder and initiator to

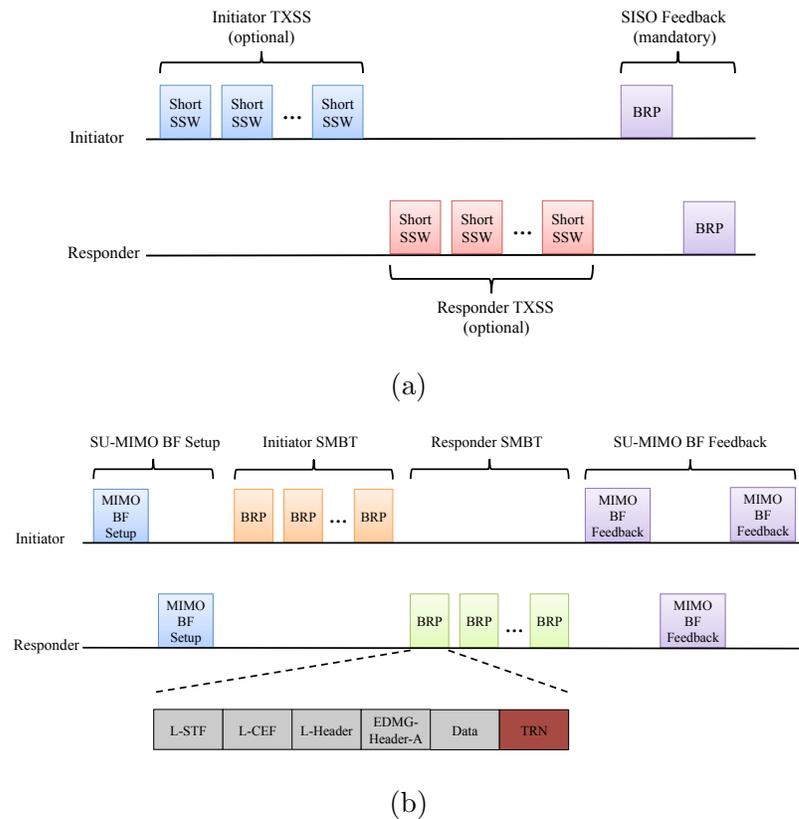


Figure 4.5: (a) The SISO phase of SU-MIMO beamforming, (b) The MIMO phase of SU-MIMO beamforming.

estimate the received SNR value for the different sectors being trained. These two sub-phases may be skipped and information obtained in the immediately preceding TXSS can be used instead. During the I-TXSS, the initiator transmits short SSW packets using different sectors while the pairing STA receives with a quasi-omnidirectional pattern. This procedure enables the responder to create an ordered list of best transmit sectors for the initiator according to their corresponding estimated SNR values. Similarly, in the R-TXSS, the initiator creates an ordered list of the best sectors based on their estimated SNR values.

In the mandatory SISO feedback sub-phase, the stations exchange BRP frames with TXSS feedback information. The BRP frame transmitted by the initiator contains a list of sector identifiers and corresponding SNR values of the transmit sectors

trained in the last R-TXSS. Similarly, the BRP feedback frame transmitted by the responder reports a list of sector identifiers and corresponding SNR values estimated in the last I-TXSS.

2) MIMO Phase: The MIMO phase enables the simultaneous training of transmit and receive sectors for each DMG antenna (e.g., phased antenna array). The MIMO phase, depicted in Fig. 4.5b, comprises four mandatory sub-phases: an SU-MIMO BF setup, an initiator SU-MIMO BF training (SMBT), a responder SMBT, and an SU-MIMO BF feedback sub-phase.

First, in the SU-MIMO BF setup sub-phase, both the initiator and responder select a subset of candidate transmit sectors per DMG antenna based on the SNR values provided in the SISO phase. The two STAs announce candidate sectors through exchanging MIMO BF setup frames. Since these sectors need to be further trained in SMBT, MIMO BF setup frames contains the number of BRP frames to be transmitted in the following initiator SMBT (or responder SMBT) sub-phase, and the order that candidate sectors will be trained in each BRP frame. Furthermore, the MIMO BF setup frame indicates a decision maker for each link (from the initiator to the responder and vice versa) defined as the STA responsible for determining the final transmit and receive antenna configurations for SU-MIMO transmissions.

In initiator SMBT, multiple transmit configurations of the initiator and multiple receive configurations of the responder are trained. This becomes possible when the initiator transmits EDMG BRP-RX/TX packets appending the TRN field. As explained in Section 4.1.1, the TRN field structure is defined by the parameters “TRN length (L)” and “RX TRN-Units per Each TX TRN-Unit (K)” (see Fig. 4.2c). Here, K should be set to the number of candidate receive sectors (at responder) and L should be set to the number of initiator’s candidate transmit sectors multiplied by the number of responder’s candidate receive sectors. Similarly, the responder performs

the responder SMBT sub-phase by sending EDMG BRP-RX/TX packets to evaluate multiple transmit configurations of the responder and the receive configurations of the initiator.

Finally, two STAs exchange two or three (depending on the decision makers) MIMO BF feedback frames. If the initiator is the decision maker for the responder link (from responder to initiator), its MIMO BF feedback specifies NR best transmit and receive combinations discovered on the prior responder SMBT sub-phase. Otherwise, the initiator includes the training feedback (from the responder SMBT) and lets the responder decide on the suitable antenna configurations. Similarly, the responder sends a MIMO BF feedback frame which either contains the NI best transmit and receive sectors combinations or feedback of the initiator SMBT sub-phase. Additionally, if the responder is the decision maker for the responder link, the NR best sector combinations are also included in this feedback frame. When the responder is the decision maker for the initiator link, the beamforming procedure terminates here since both STAs have already determined their antenna configuration for SU-MIMO operation. Otherwise, the initiator sends another MIMO BF feedback frame to announce the NI best transmit and receive sector combinations. Note that the selection of the best sector combinations is implementation dependent; however, in general, only one transmit and one receive sector is selected per DMG antenna.

4.1.3.2 MU-MIMO Beamforming

MU-MIMO BF enables an initiator and a group of responders to determine appropriate antenna configurations for simultaneous transmission of multiple data streams with minimum inter-stream interference. Here, we describe the BF procedure for a given multi-user (MU) group; mechanisms of forming such a group are discussed in [77]. IEEE 802.11ay only supports downlink MU-MIMO transmission; hence,

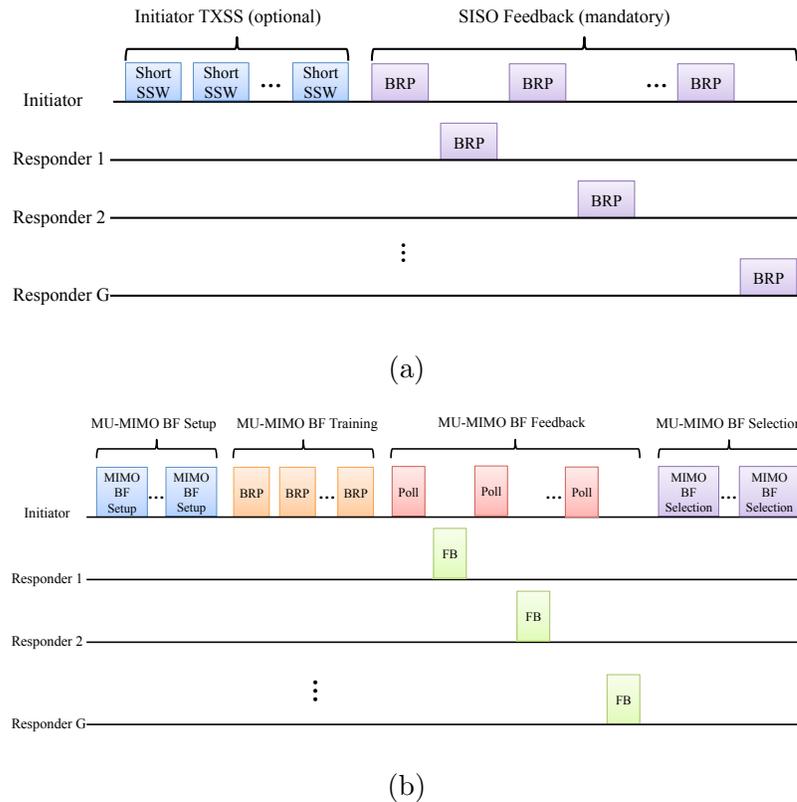


Figure 4.6: The MU-MIMO Beamforming (a) The SISO phase, (b) The MIMO phase.

the initiator (AP/PCP) only trains the transmit antenna configurations per DMG antenna while the responders only train their receive antenna configurations. The MU-MIMO beamforming protocol is started and controlled by the initiator, which is always the decision maker. Fig. 4.6 depicts the MU-MIMO BF procedure consisting of SISO and MIMO Phases.

1) SISO Phase: The SISO phase (Fig. 4.6a) starts with an optional initiator TXSS sub-phase and is followed by a mandatory Feedback sub-phase. As described in Section 4.1.3.1, the initiator performs TXSS by sending short SSW packets from different transmit sectors of each of its DMG antennas. During this time, the stations in the MU Group use quasi-omni pattern and measure the link quality of each transmit sector. In the SISO Feedback sub-phase, the initiator polls every station in the MU

group via sending a BRP frame. The polled STA responds with a list of sectors per each transmit DMG antenna and their corresponding quality indicators (e.g., measured SNR values).

2) MIMO Phase: The MIMO phase is depicted in Fig. 4.6b and is comprised of four consecutive sub-phases, MU-MIMO BF setup, MU-MIMO BF training, MU-MIMO BF feedback, and MU-MIMO BF selection sub-phase.

First, the initiator sends a MIMO BF setup frame to each intended responder. Based on the feedback provided by the SISO phase and to reduce training time, the initiator may select only a subset of STAs in the initial MU group. This frame specifies the selected responders, the transmit sectors for each DMG antenna that will be trained, and the order of transmission/training. The setup frame is transmitted to all intended responders employing their best known directional configuration; hence, the initiator may repeat its transmission several times to ensure reception by all intended responders.

Second, in the MU-MIMO BF training sub-phase, the initiator sends BRP frames similar to initiator SMBT sub-phase of SU-MIMO BF. To sweep antenna configurations throughout a frame, TRN fields are appended to the BRP frames. Next, the initiator polls each remaining intended responder for its BF feedback. The MIMO BF feedback transmitted by each responder contains the list of the initiator's transmit DMG antennas/sectors, each with its corresponding responder's receive DMG antenna/sector and the associated signal quality.

Finally, with the help of the obtained feedback, the initiator selects and announces a set of recipient STAs along with their antenna configurations. The set of selected STAs does not have to be the same as the initial MU group or the intended responders of the MU-MIMO BF training sub-phase.

4.2 MUTE: MUlti-stream beam-Training for mm-wavE networks

In this section, we describe the design of MUlti-stream beam-Training for mm-wavE networks (MUTE) that aims to provide the best analog configurations for downlink MIMO transmissions.

4.2.1 Design Overview

60 GHz channels lack rich scattering, i.e., a few dominant paths fully characterize the channel [34]. An analog beam acts as an amplifier, boosting the strength of certain paths within its main lobe (and side lobes) and weakening the others. As a result, beams that cover the same physical paths have highly correlated channels; transmitting multiple data streams using such beams hinders multiplexing gains. Hence, MUTE targets selection of analog beams that capture diverse or ideally orthogonal paths to the intended receivers. With idealized analog beam patterns, i.e., non-overlapping pencil-shaped beams without side lobes, knowing the received SNR via each beam corresponds to knowledge of path direction. In other words, if a particular beam provides high SNR, we can infer that it captures a path whose direction is within its main lobe. However, beam patterns generated by phased arrays are highly irregular and may even have multiple equally strong lobes [1] as depicted in Fig. 4.7a. If using such beams at the AP achieves high SNR at a client, it is hard to conclude the number of paths captured by each lobe. Hence, selection of this beam pattern (i.e., codebook entry) to exploit one of these lobes might preclude use of another codebook entry that directs energy along the other lobe (e.g., the beam pattern shown in Fig. 4.7b), as it causes interference in multi-stream transmissions. Note that irregularity in the beam pattern is a byproduct of small array size and limited phase levels

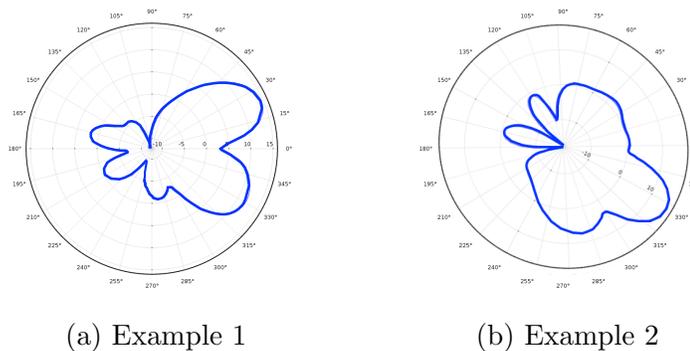


Figure 4.7: Irregular beam pattern examples from X60 platform [1].

at the phase shifters. We observe the same order of irregularity with commodity off-the-shelf 802.11ad devices with 32 antenna elements [16]. Although next-generation 60 GHz devices tend to have even more antennas, it seems that the beam patterns are still far from the “perfect” pencil-shaped patterns due to complexity, size, and power consumption constraints.

We propose MUTE, a MIMO beam steering protocol to select a set of analog beams with diverse or ideally orthogonal paths. Fig. 4.8 depicts the MUTE system architecture. MUTE runs a background process on the training frames received during the initial beam acquisition phase, which occurs periodically to establish and maintain a directional link between the AP and every client. Thanks to the high sampling rate at 60 GHz band, we can estimate the high-resolution PDP corresponding to each beam. Unfortunately, solely adding PDP information does not solve the MIMO beam steering problem. Namely, while PDP does reveal the presence of multiple paths as well as their relative timing, it does not provide any direction information.

While practical beam patterns are highly irregular, they nonetheless have beam-forming gain that is known a priori in each direction. MUTE couples radiation pattern knowledge over a suite of irregular patterns with PDP estimates for each pattern to infer the direction of each path. In other words, by weighting each PDP to the known directional gain for that beam pattern, we can narrow down the direction interval

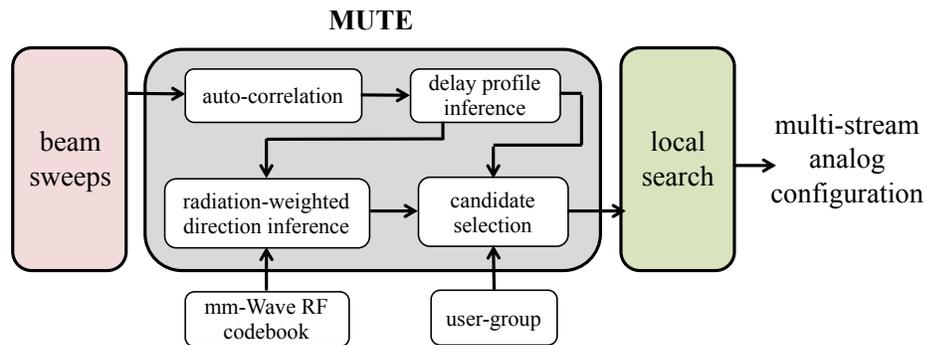


Figure 4.8: MUTE system architecture.

that each path may fall into. For example, in the beam pattern examples of Fig. 4.7, by receiving the same path (inferred via time delays in the PDP) with both beams, we can correspondingly weight the likelihood that the path direction is within the overlapping area of their main lobes.

MUTE leverages this information to construct candidate sets of transmit and receive beams at the AP and the target user group, respectively. This candidate set selection is obtained with *zero* additional overhead, only by repurposing the link establishment beam sweeps. As shown in Fig. 4.8, the final analog configuration is discovered by a local search over the candidate set of beams at the AP and clients.

Lastly, we note that MUTE requires directional beam patterns so that if beam patterns are perfectly omni-directional (or are otherwise identical), we cannot infer the path angles (directions). Furthermore, if the beam patterns are “perfect” and divide 360 degrees into non-overlapping regions, then the solution is quite trivial: one only needs to select distinct high SNR beams for multi-stream transmission, as distinct non-overlapping beams cannot share a common path. Thus, our approach is applicable not only to irregular beams (e.g., Fig. 4.7), but also symmetric side lobes and any other deviation from strictly non-overlapping beams.

4.2.2 Primer on Initial Beam Acquisition

Commercial products [16] and WLAN standards such as IEEE 802.11ad [13] and IEEE 802.11ay [6] establish directional links through a training mechanism, in which one node sends training frames sequentially across all beams in the predetermined RF codebook while the other node employs a quasi-omni antenna pattern to find the beam providing the highest signal strength. Repeating this procedure at both ends achieves a beam-pair configuration that can support single-stream (i.e., SISO) directional communication. Although prior efforts have attempted to reduce the frequency of such beam training procedures in mobile 60 GHz WLANs via in-band and out-of-band solutions [36, 46, 48, 79], the AP is bound to periodically repeat these beam sweeps in order to retain directional link connectivity. For example, link failure due to blockage or mobility or the presence of a new user triggers the beam acquisition sweeps.

In MUTE, we assume that the channel between the AP and each client is reciprocal, i.e., under a fixed analog configuration at the AP and the client, the uplink channel is the same as the downlink channel. Thus, the AP sending a training frame via a directional beam to a client in quasi-omni reception, is equivalent to the client sending the training frame via its quasi-omni pattern and the AP receiving it via the same directional beam. Hence, we run MUTE at the AP and process the received training frames in two stages: *(i)* the AP sweeps across its directional beams while the quasi-omni client sends training frames, *(ii)* the client sweeps and sends training frames while the AP is in quasi-omni reception. The underlining rationale is that the AP usually has more computation resources than mobile clients and thus MUTE can easily process the received training frames as discussed below.

4.2.3 Stream Separability Inference

4.2.3.1 Beam-specific Delay Profile

MUTE specifies the dominant paths by repurposing the received training frames during the initial beam acquisition phase. The GHz-scale sampling rate and sparsity of 60 GHz channels provides the unique opportunity to obtain the high-resolution PDP in contrast to sub-6 GHz bands [34]. PDP gives the distribution of signal power received over a multipath channel as a function of propagation delays and is specified as the spatial average of the complex baseband channel impulse. PDP is typically computed via transmission of a known pilot block [80, 81]. The auto-correlation of the received signal with a local copy of the ideal pilot block provides an estimation of the channel impulse response and leads to PDP estimation.

Unlike sub-6 GHz bands, we cannot represent a 60 GHz channel via a single power delay profile. The reason is behind the fact that omnidirectional transmission is not feasible with an omnidirectional reception in mmWave networks [13]. In other words, due to higher path loss, at least one side needs to be directional. Employing directional beam biases a number of paths by amplifying their strength over others and thus makes the PDP change based on the analog beams in use.

Let $\mathbf{H} \in \mathbb{C}^{N_{cl} \times N_{AP}}$ be the wireless channel between a client (with N_{cl} antennas) and the AP (with N_{AP} antennas). The uplink signal received by the AP can be written as

$$\mathbf{x}[m] = \mathbf{H}[m]\mathbf{w}_{cl}s[m] + \mathbf{n}[m], \quad m = 0, 1, \dots, M - 1, \quad (4.1)$$

where $s[m]$ is the m th symbol of the transmitted training frame (M is total number of symbols), \mathbf{w}_{cl} represents the client's analog beamforming vector and \mathbf{n} is additive white Gaussian noise. With the phased-array antenna, we can only access the signal

at the RF chain. The received signal at the RF chain, denoted by $\bar{\mathbf{y}}$, is expressed as:

$$y[m] = \mathbf{w}_{\text{AP}}^H \mathbf{x}[m], \quad m = 0, 1, \dots, M - 1, \quad (4.2)$$

where \mathbf{w}_{AP} is the analog beamforming vector at the AP. Both \mathbf{w}_{AP} and \mathbf{w}_{cl} are implemented using analog phase shifters; hence; the modulus of all the elements is 1. Furthermore, for an omnidirectional case, the phase delays are equal to zero for all antenna elements (i.e., $\mathbf{w}_{\text{AP}} = \mathbf{1}_{N_{\text{AP}} \times 1}$ or $\mathbf{w}_{\text{cl}} = \mathbf{1}_{N_{\text{cl}} \times 1}$). From Equation (4.1) and (4.2), we can write:

$$y[m] = h_{\text{eff}}[m]s[m] + \mathbf{w}_{\text{AP}}^H \mathbf{n}[m], \quad m = 0, 1, \dots, M - 1, \quad (4.3)$$

where \mathbf{h}_{eff} is the effective channel between the client and the AP when they employ \mathbf{w}_{cl} and \mathbf{w}_{AP} , respectively and can be written as follows:

$$h_{\text{eff}}[m] = \mathbf{w}_{\text{AP}}^H \mathbf{H}[m] \mathbf{w}_{\text{cl}} \quad (4.4)$$

From Equation (4.4), the effective channel clearly depends on the choice of analog beams. During the initial beam acquisition phase, the AP and each client exchange a series of training frames while sweeping through their codebook entries while the other side is in omnidirectional mode. These training frames are typically fixed across codebook entries and are known at both the AP and client. Hence, the autocorrelation of the received training signal with the known training frame obtains the PDP as follows :

$$(\bar{\mathbf{y}} \star \bar{\mathbf{s}})[n] = \tilde{h}_{\text{eff}}[n] \quad (4.5)$$

where $\bar{\mathbf{y}} = [y[0], y[1], \dots, y[M - 1]]^T$ is the received signal and $\bar{\mathbf{s}} = [s[0], s[1], \dots, s[M - 1]]^T$ is the ideal known training frame.

Peak Detection. A physical path between the AP and the client is reflected as a peak in $|\tilde{h}_{\text{eff}}|^2$. MUTE accounts for noise and hardware impairments by setting thresholds for peak detection (e.g., the peak value should be at least 10 times greater than the noise level). To emphasize that the identified multi-path components are beam dependent, we denote $P_b(\tau)$ the delay profile when the AP employs beam b (with quasi-omni client). $P_b(\tau)$ includes the intensity of detected peaks in $|\tilde{h}_{\text{eff}}|^2$ as a function of time delay τ .

Hardware Imperfections. The estimated power delay profile is subject to error due to carrier frequency offset or packet detection delay. Fortunately, prior work presented several solutions to deal with such hardware imperfections in practice [82, 83]. In this work, we obtain the PDP measurements directly from our platform (see Sec. 4.3) which deals with carrier frequency offset and synchronization issues. Today’s WiFi drivers already provide channel impulse response (CIR), RSSI, and SNR information and we expect future drivers for the 60 GHz devices to continue this trend and allow access to PDP.

Next, without loss of generality, we explain MUTE’s design for selecting transmit beams (at the AP) to enable downlink MIMO transmission. At the end, we briefly explain how candidate beam selection is performed for the MIMO reception at clients in a similar manner.

4.2.3.2 Delay Profile Aggregation

MUTE obtains the PDP corresponding to all codebook entries. Fig. 4.9 depicts an example scenario in which a LOS path and two reflected paths exist between the AP and the client. Two analog beam patterns are shown. Beam m (blue) has high directivity along LOS path 1 and NLOS path 3 while having low directivity gain along NLOS path 2. Hence, $P_m(\tau)$ includes two peaks at τ_1 and τ_3 , corresponding to path

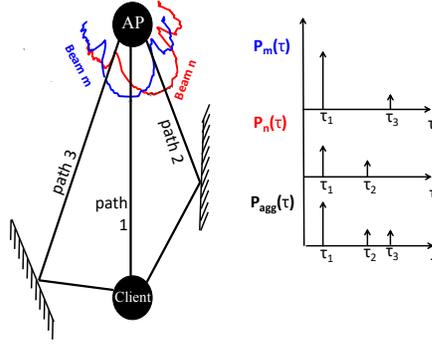


Figure 4.9: An example scenario with 3 dominant physical paths between the AP and client.

1 and path 3, respectively. Likewise, beam n (red) captures two dominate paths 1 and 2. The common LOS path between is reflected in both profiles by a peak at time τ_1 . MUTE aggregates all beam-specific delay profiles to identify all dominant paths between the AP and the client (denoted by $P_{agg}(\tau)$ in Fig. 4.9). We avoid double-counting the peaks if their time delay difference is below a configurable threshold to account for noise and hardware impairments that can cause a slight shift in the delay profile. Note that such aggregation over beam-specific PDPs do not provide any information about the relative strength of different paths as the strength of a path is always biased by the analog beams in use. To bold this fact, we depict same-length peaks for all paths in $P_{agg}(\tau)$ in Fig. 4.9.

4.2.3.3 Ambiguity Across Different Clients

By comparing the delay profile of any two beams, we can observe whether they capture similar paths, diverse paths, or completely orthogonal paths. Hence, beam-specific delay profiles can be directly leveraged to infer stream separability in the analog domain for *single user* MIMO transmissions. However, there is ambiguity in stream separability inference for multi-user transmissions. The reason is that $P_{agg}(\tau)$ represents the relative delay of dominant paths with respect to the shortest path (with

smallest time delay); comparing the delay profile of different users is problematic as there is no common time reference point. Even if there exists a global time reference point in PDP, we would still not be able to perfectly identify orthogonal paths to multiple users without having a sense of path direction. For instance, two users each having one path with the same path angle (see U_1 and U_2 in Fig. 4.11a) may experience different time delays as the time of flight or distance to the AP might be different. Hence, solely adding PDP information does not solve the multi-stream beam selection problem as it does not contain any direction information. MUTE tackles this issue by coupling PDP estimates with the *known* radiation patterns.

4.2.4 Radiation-Weighted Direction Inference

MUTE integrates the obtained beam-specific PDP with the knowledge of beam patterns to estimate the direction of all dominant paths from the AP to each client. Since the AP is the common reference point in space, MUTE is able to infer stream separability for any choice of analog beams in both multi-user and single-user multi-stream transmissions. We denote C_{AP} the AP's RF codebook. In particular, $C_{AP} = [c_1, c_2, \dots, c_N]$ where c_b denotes the radiation pattern of codebook entry b and $c_b(\theta)$ denotes its directivity gain along azimuth angle θ . Furthermore, we define \bar{c}_b as the average directivity of beam b and compute it as $\bar{c}_b = \frac{1}{|c_b|} \sum_{\phi=0}^{2\pi} c_b(\phi)$. Fig. 4.10a depicts an example beam pattern and its average directivity. For a pencil-shaped beam pattern, only the directivity of the main lobe is above the average directivity; however, for an irregular beam pattern (shown in Fig. 4.10a), multiple lobes might provide higher than average directivity.

Our key idea is based on the following two insights:

- The PDP of beam b (i.e., $P_b(\tau)$) **containing** peak p (corresponding to a path) implies that the angle of that path is *likely* to be among those directions satis-

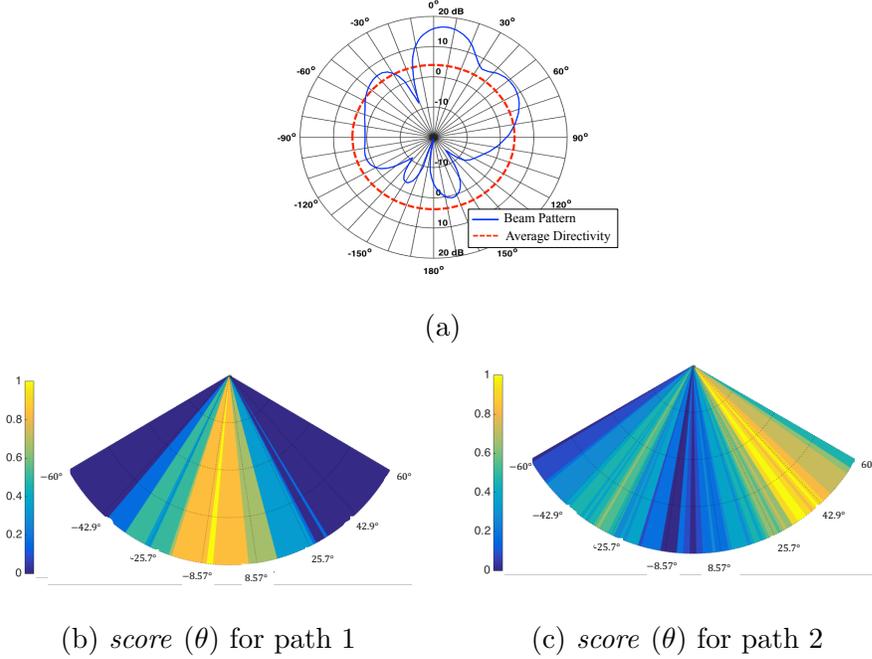


Figure 4.10: (a) An example irregular beam pattern and its average directivity, (b) $score(\theta)$ for path 1, and (c) $score(\theta)$ for path 2 in Fig. 4.9.

fying $c_b(\theta) > \bar{c}_b$.

- The PDP of beam b (i.e., $P_b(\tau)$) **not containing** peak p (corresponding to a path) implies that that the angle of that path is *unlikely* to be among those directions satisfying $c_b(\theta) > \bar{c}_b$.

Therefore, by weighting each PDP by the known directional gain for that radiation pattern, we can narrow the direction interval that each path may fall into. Algorithm 1 presents the details of MUTE's direction inference mechanism. This algorithm takes beam-specific delay profiles $P_b(\tau)$, the aggregated delay profile $P_{agg}(\tau)$, and the pre-known AP's beam patterns C_{AP} . MUTE examines peaks in $P_{agg}(\tau)$ (corresponding to physical paths) one by one. The angle of each path is estimated after examining all beam patterns and their collected delay profiles. For each beam pattern, the algorithm first computes the average directivity gain (line 3) and then indicates whether the delay profile of that beam contains the under-examined path or not. If beam b

Algorithm 1 MUTE's Radiation-Weighted Direction Inference

Input: $\{P_b(\tau)\}, P_{agg}(\tau), C_{AP}$
Output: $A(\Theta)$

```

1: for each peak  $p$  in  $P_{agg}(\tau)$  do
2:   for  $c_b \in C_{AP}$  do
3:      $\bar{c}_b = \frac{1}{|c_b|} \sum_{\phi=0}^{2\pi} c_b(\phi)$ 
4:     if  $P_b(\tau)$  contains peak  $p$  then
5:        $I(b, p) = 1$ 
6:     else
7:        $I(b, p) = -1$ 
8:     end if
9:   end for
10:   $score(\theta) = \sum_{c_b \in C_{AP}} I(b, p) \left( \frac{|c_b(\theta) - \bar{c}_b| + c_b(\theta) - \bar{c}_b}{2} \right)$ 
11:   $\theta_p^* = \arg \max_{\theta} score(\theta)$ 
12: end for

```

captures peak p , the algorithm sets the indicator $I(b, p)$ to 1 and -1 otherwise (line 4-7). Then, we weight the likelihood of directions with higher than average directivity via a score-based mechanism (line 10):

$$score(\theta) = \sum_{c_b \in C_{AP}} I(b, p) \left(\frac{|c_b(\theta) - \bar{c}_b| + c_b(\theta) - \bar{c}_b}{2} \right) \quad (4.6)$$

If beam b captures the peak (i.e., $I(b, p) = 1$), we increase $score(\theta)$ by $c_b(\theta) - \bar{c}_b$ for those angles with higher than average directivity gain (note that $\frac{|x|+x}{2} = 1$ for $x > 0$ and is zero otherwise). Otherwise (i.e., $I(b, p) = -1$), we decrease $score(\theta)$ for those angles with higher than average directivity gain to show the likelihood reduction. Finally, the algorithm returns the direction with the highest score for each path and creates an angular profile $A(\Theta)$ containing the direction of all paths for each user (line 11).

Fig. 4.10b and Fig. 4.10c demonstrate $score(\theta)$ for path 1 and path 2 in the example scenario of Fig. 4.9. We observe that the high score (yellow region) matches the true physical angle. These directions were inferred using imperfect beam patterns

generated by practical phased array antennas (see Sec. 4.3 for details). In general, the accuracy of MUTE's direction inference algorithm depends on the shape of beam patterns, their beamwidth, and the overlap of different codebook entries' beam patterns.

4.2.5 Candidate Beam Selection

So far, we have described how MUTE applies initial direction acquisition beam sweeps to estimate the delay profile and angular profile of each user. Here, we discuss how MUTE leverages this information to select a subset of candidate beams. MUTE aims to select the analog beams with diverse paths or ideally orthogonal paths. The rationale is that two analog beams sharing a dominant path would have high channel correlation which hinders the multiplexing gain of multi-stream transmission. In particular, [84] showed that in 2×2 single-user MIMO, both streams cannot operate under LOS condition unless they are separated with orthogonal polarizations.

Let G be the target user group. For any user u in G , MUTE attempts to maximize the received signal strength at u while minimizing the interference at all unintended users in G . To this end, for any dominant path i between the AP and user u , MUTE includes a beam in the candidate set that provides maximum directivity along path angle $\theta_{u,i}$ (to increase signal strength at the intended user), while inducing minimum directivity along the path angles of all other unintended users in G (to reduce interference). The general optimization for candidate beam selection is therefore:

$$B_u(G) = \left\{ \arg \max_b \frac{c_b(\theta_{u,i})}{\sum_{\substack{v \in G \\ v \neq u}} \sum_{\substack{\theta_{v,x} \in A_v \\ \theta_{v,x} \neq \theta_{u,i}}} c_b(\theta_{v,x})}, \forall \theta_{u,i} \in A_u(\Theta) \right\} \quad (4.7)$$

where $c_b(\theta)$ is the directivity of b^{th} entry in the AP's codebook along azimuth angle θ and $A_u(\Theta)$ represents the angular profile of user u . We further elaborate on MUTE's

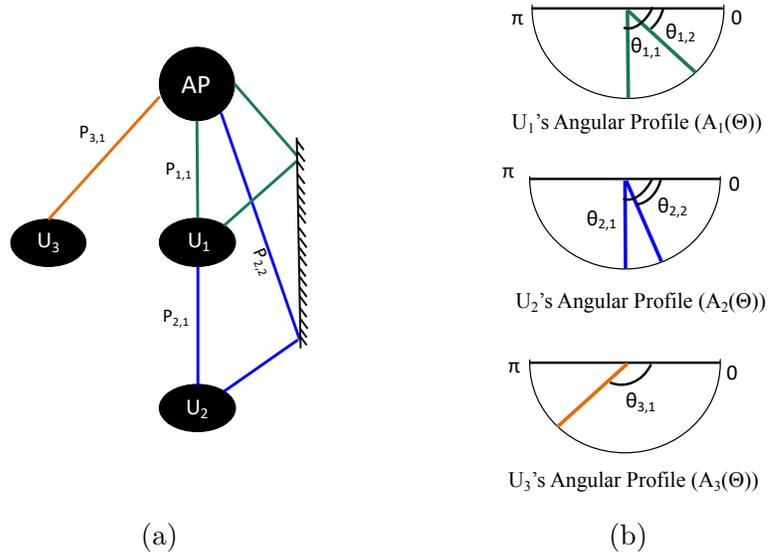


Figure 4.11: Candidate selection example.

candidate selection mechanism via a simple scenario with three users U_1 , U_2 and U_3 as depicted in Fig. 4.11a. Fig. 4.11b shows the angular profiles of these users; U_1 and U_2 hold two dominant physical paths with the AP, i.e., $A_1(\Theta) = \{\theta_{1,1}, \theta_{1,2}\}$, and $A_2(\Theta) = \{\theta_{2,1}, \theta_{2,2}\}$ while $\theta_{1,1} = \theta_{2,1}$ (they share a LOS path). Note that in practice, we flag such shared path if the path angle distance (here, $|\theta_{1,1} - \theta_{2,1}|$) is below a configurable threshold. Furthermore, U_3 has only a LOS path with the AP, i.e., $A_3(\Theta) = \{\theta_{3,1}\}$. We refer to U_1 and U_2 as *path-sharing users* because their angular profiles share a common path. We also call U_1 and U_3 *distinct users*. Next, we discuss two examples of candidate beam selection:

4.2.5.1 Example 1: Distinct Users

First, we target selection of candidate transmit beams for a two-stream MIMO transmission to U_1 and U_3 . To maximize the stream separability in the analog domain, we choose analog beams such that they provide high directivity gain along path angles of U_1 while inducing minimum interference at U_3 , and vice versa. According to Equation (4.7), MUTE constructs the candidate set of beams, denoted by $B(U_1, U_3)$,

as follows:

$$B(U_1, U_3) = \left\{ \arg \max_b \frac{c_b(\theta_{1,1})}{c_b(\theta_{3,1})}, \arg \max_b \frac{c_b(\theta_{1,2})}{c_b(\theta_{3,1})}, \arg \max_b \frac{c_b(\theta_{3,1})}{c_b(\theta_{1,1}) + c_b(\theta_{1,2})} \right\} \quad (4.8)$$

This candidate set contains: (i) The beam that has maximum directivity along $\theta_{1,1}$ while having minimum directivity along $\theta_{3,1}$; (ii) The beam with maximum directivity along $\theta_{1,2}$ and minimum directivity along $\theta_{3,1}$; and (iii) The beam with maximum directivity along $\theta_{3,1}$ while having minimum directivity along $\theta_{1,1}$ and $\theta_{1,2}$. The candidate set includes three beams in which two of them are eventually selected for 2×2 MU-MIMO transmission. In this toy example, it is not hard to predict that the third candidate beam in the $B(U_1, U_3)$ should be one of the final selected beams as it is the only candidate beam that provides high directivity gain (and SNR) at U_3 .

4.2.5.2 Example 2: Path-Sharing Users

Second, we target a two-stream MIMO transmission to U_1 and U_2 in Fig. 4.11. Here U_1 and U_2 share a common LOS path (i.e., $\theta_{1,1} = \theta_{2,1} = \theta'$) and the transmission of two streams under LOS condition would result in significant channel correlation. Therefore, the analog beam along LOS path should be employed for either user 1 or user 2, and not both. In the former case, its directivity gain along $\theta_{2,2}$ should be minimized while in the later case, its directivity gain along $\theta_{1,2}$ should be minimum. Thus, the set of candidate beams for MU-MIMO to U_1 and U_2 can be directly derived from Equation (4.7) as follows:

$$B(U_1, U_2) = \left\{ \arg \max_b \frac{c_b(\theta')}{c_b(\theta_{2,2})}, \arg \max_b \frac{c_b(\theta')}{c_b(\theta_{1,2})}, \arg \max_b \frac{c_b(\theta_{1,2})}{c_b(\theta_{2,2}) + c_b(\theta')}, \arg \max_b \frac{c_b(\theta_{2,2})}{c_b(\theta_{1,2}) + c_b(\theta')} \right\} \quad (4.9)$$

Note that single-user MIMO beam steering is a special case in Equation (4.7) in which the angular profiles of all users in G are identical.

Discussion. We have described how MUTE selects a set of candidate transmit beams by processing the received training frames at the AP. There are two possible ways to replicate the same procedure for client-side beam selection: *(i)* Each client processes the received training frames during the initial beam acquisition sweeps, taking into account its known RF beam patterns; or *(ii)* clients send beam training frames (while sweeping over their directional beams) and the AP receiving the training frame takes care of processing. In the latter case, the AP needs to know the clients' RF beam patterns and announce the candidate beams to the target user group before the final local search.

4.2.6 Local Search

Once the sets of candidate transmit and receive beams are constructed, local training is triggered to discover the optimum analog configuration. This training involves testing all possible beam combinations in the candidate sets for a multi-stream transmission to the target user group and estimating the achievable aggregate rate under each configuration. While MUTE's candidate beam selection does not entail any additional overhead, this local search requires active signaling and thus incurs overhead that is proportional to the number of candidate beams at the AP and clients. We evaluate the candidate set size in Sec. 4.5.

4.3 MUTE Implementation

We conduct over-the-air experiments utilizing X60, a programmable testbed for wide-band 60 GHz WLANs [1]. X60 provides signal level accessibility, and is engineered to

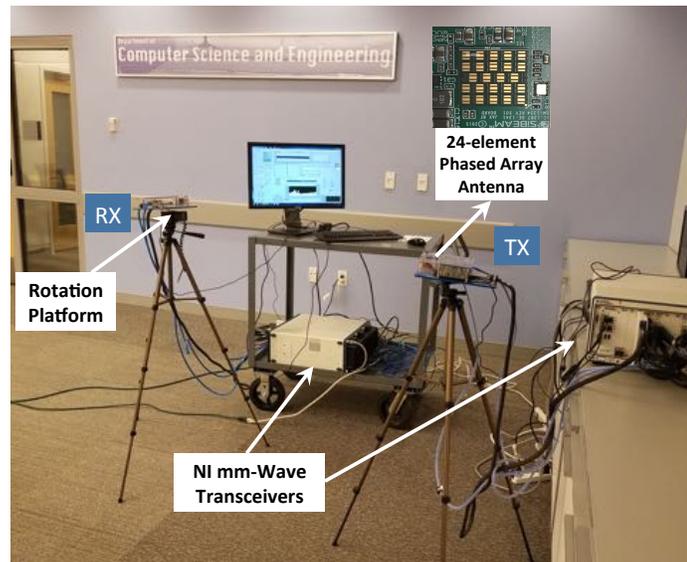


Figure 4.12: The X60 platform for 60 GHz band.

provide CSI, SNR, and PDP. Fig. 4.12 depicts X60 in which each node is equipped with National Instruments' mm-Wave transceiver systems and a *user-configurable* 24-element phased array antenna from SiBeam. It enables communication over 2 GHz channels via fixed codebook based beam patterns that can be steered in real-time (electronic switching in $< 1\mu s$). X60 enables fully programmable PHY, MAC, and Network layers. The reference PHY implementation allows for modulation and coding combinations from 1/5 BPSK to 7/8 16 QAM, resulting in bit rates from 300 Mbps to 4.75 Gbps. Data transmission takes place in 10 *ms* frames, which are divided into 100 slots of 100 μs each.

The built-in phased array has 24 antenna elements; 12 for transmission and 12 for reception. The phase of each antenna element can take one of the four values: 0, $\pi/2$, π , and $3\pi/2$. SiBeam's reference codebook consists of 25 3D beam patterns spaced roughly 5° apart (in their main lobe direction) and covering a sector of -60° (corresponding to beam index -12) to 60° (corresponding to beam index +12) around the antenna's broadside direction. The half power beamwidth of beams are $25^\circ - 35^\circ$. X60's beam patterns (two of them are shown in Fig. 4.7) are similar to the patterns of

commodity off-the-shelf 802.11ad devices which have main lobe overlaps and strong side-lobes [16]. Thus, *X60* allows us to evaluate realistic, imperfect and irregular beam patterns, and their impact on multi-stream beam steering.

Due to hardware limitations (availability of only one RF chain at each node), over-the-air MIMO transmission is not feasible. Our key experimental methodology is to collect channel samples (in time and frequency domain) from over-the-air measurements and subsequently perform trace-driven emulation to study MIMO beam steering in millimeter-wave networks. Sec. 4.5 elaborates more on our measurement setup.

4.4 Benchmarking Algorithms

In this section, we describe two benchmarking algorithms for evaluation purposes.

4.4.1 SNR-based Beam Selection (Baseline)

We introduce a baseline scheme that down selects a subset of beams from the fixed RF codebook based on their achievable SNR in SISO communication. The underlying rationale is that higher received signal strength provides greater margin for interference tolerance. Therefore, this scheme selects a subset of beams (for multi-stream transmission) entirely based on their SNR values and exploits zero-forcing to mitigate or ideally cancel any residual interference. In particular, the baseline scheme picks the top n beams at the AP and the top m beams at each client, according to their achievable SNR in the initial AP-side and client-side beam sweeps. The final analog configuration is realized via a local search among all possible combinations of candidate beams.

Similar to MUTE, this baseline approach does not introduce additional overhead

for candidate selection as the received SNR associated with each beam is already available after initial beam sweeps as discussed in Sec. 4.2.2. However, unlike MUTE which infers stream separability by building an angular profile, the baseline scheme select beams entirely based on SNR. The final analog setting is found by further training whose overhead depends on the number of transmit and receive beams in candidates sets (i.e., n and m). Throughout the chapter, we refer to this method as baseline scheme or SNR-based beam selection strategy.

4.4.2 Exhaustive Search

Exhaustive search assesses all beam combinations to find the optimal configuration. For simultaneous transmission of m data streams, exhaustive search tests total $O(|C_{AP}|^m \times |C_{cl}|^m)$ analog combinations where $|C_{AP}|$ is the transmit codebook size at the AP and $|C_{cl}|$ is clients' codebook size. Implementation of exhaustive search may not be practical in real scenarios due to the prohibitively large beam training overhead. Nonetheless, we study this algorithm for comparison purposes as it provides an upper-bound for the achievable sum-rate of multi-stream transmission.

4.5 Experimental Results

In this section, we conduct over-the air measurements to evaluate and compare the performance of MUTE against the benchmarking algorithms.

Setup. We deploy X60 nodes and conduct an extensive set of experiments in multiple indoor environments and many AP-client settings. In this work, as the first attempt to explore MIMO beam steering in 60 GHz WLANs with over-the-air channel traces, we zoom into one experimental setup depicted in Fig. 4.13 including the AP and 12 client locations (represented by square boxes). The AP is mounted

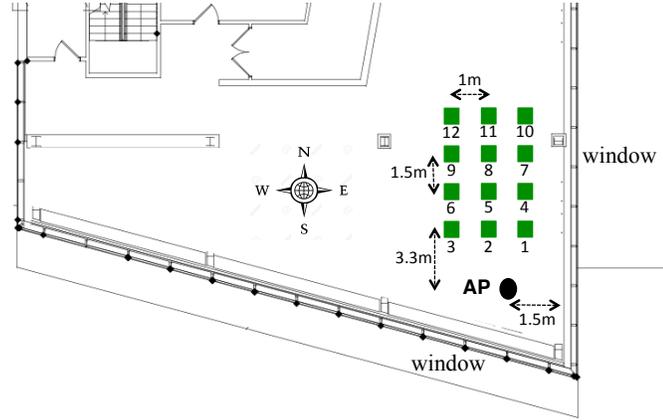


Figure 4.13: Experimental floorplan. Square boxes represent client positions.

on a tripod at 0.9 m height from the floor and pointing North. The client is at the same height but pointing South at all locations. The presence of windows and metal coating beneath them (not shown) create reflections.

Methodology. For each AP-client setting, we collect channel statistics for all possible 625 (25×25) beam-pair combinations. For each beam-pair, 100 frames are transmitted at Modulation and Coding Scheme (MCS) 0 and SNR, channel magnitude and phase, and PDP are logged every four frames (every 40ms). We assume that the multiple virtual RF chains are co-located at the AP. This is because we find out that changing the location of the AP by $\lambda/2$ or $2.5mm$ does not change the received PDP, composite channel, and SNR. In other words, due to wide ($\approx 25^\circ$) beamwidth of codebook entries, the physical paths being captured by a beam are not sensitive to small movement of the AP or client. To emulate hybrid analog/digital beamforming, we process the channel traces under a fixed choice of analog beams at the AP and clients and compute zero-forcing weights. We then map the SINR to the corresponding data rate using the protocol-specific minimum SNR tables [85].

4.5.1 Performance Analysis of MUTE

First, we explore the performance of MUTE in selecting the best analog beams at the AP and client(s) to be used for multi-stream simultaneous transmission and reception, respectively. For simplicity, we focus on a two-stream case (2×2 multi-user and single-user MIMO configurations) and later, in Sec. 4.5.4, we increase the number of spatial streams.

4.5.1.1 Multi-User MIMO

We consider a downlink two-user MIMO case in which one client, R_1 , is placed at position 1, while we consider all other 11 positions for the second client (R_2) by sequentially repeating experiments at these positions. For each position of R_2 , MUTE generates a subset of candidate beams at the AP and each client by post processing the measured beam-specific PDPs. Then, to discover the final multi-stream configuration, MUTE performs a local search and computes the zero-forcing weights for any combination of analog beams in the candidate sets. Applying the zero-forcing weights, we compute the expected SINR at R_1 and R_2 and infer the per-user data rate by employing the protocol-specific minimum SNR tables [85]. In particular, for each user and potential multi-stream analog configuration, MUTE selects the MCS index whose corresponding SNR is less than or equal to the calculated SINR. The corresponding number of data bits per symbol is the per-user data rate (each stream can use a different MCS but the coding rate for two streams are the same). The final MIMO beam configuration is the one providing highest aggregate data rate. For comparison, we also implement Exhaustive Search and the baseline SNR-based Beam Selection. To have a fair comparison, we ensure that the candidate set size for the baseline scheme is the same as MUTE, e.g., if MUTE provides k candidate beams, the baseline scheme would down select the top k beams (based on SNR).

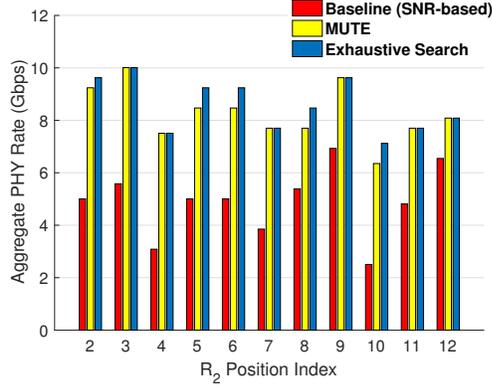


Figure 4.14: Aggregate PHY rate of a two-user MIMO transmission to R_1 (fixed at position index 1) and R_2 when placed at other 11 positions.

Fig. 4.14 shows the aggregate PHY rate of the baseline, MUTE, and Exhaustive Search for a two-user simultaneous transmission to R_1 and R_2 . First, we observe that MUTE is able to achieve more than 90% of Exhaustive Search’s rate across all locations. This implies that MUTE’s multi-path inference successfully discovers candidate beams with maximum stream separability. Second, Fig. 4.14 shows that the baseline strategy provides around 60% of the Exhaustive Search’s aggregate rate, except for when R_2 is placed at position indices 4, 7 and 10. In those cases, R_1 and R_2 are approximately along the same LOS path with the AP. Hence, the baseline scheme selects analog beams along the shared LOS path and induces inter-user interference. In other words, a candidate beam (at the AP) that is intended for R_1 has high directivity along the LOS path and will incur high interference at R_2 . Transmission of two streams along the LOS path reduces the aggregate PHY rate. In contrast, MUTE is able to create “separate beam” transmissions via an NLOS path, even if users are along the same LOS path, and still obtain 90% of maximum aggregate rate.

Finding: MUTE achieves more than 90% of Exhaustive Search’s aggregate rate in two-user MIMO case: If the LOS paths from the AP to two users have enough spatial separation, MUTE selects beams along the LOS path for both users; otherwise, MUTE is able to create “separate beam” transmissions via NLOS paths.

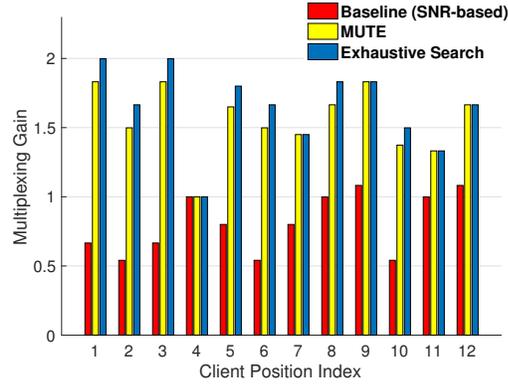


Figure 4.15: Multiplexing gain of 2×2 single-user MIMO as a function of client position.

4.5.1.2 Single-User MIMO

Next, we perform a similar experiment for 2×2 single-user MIMO. In contrast to downlink multi-user MIMO in which users are spatially separated, single-user MIMO requires multiple independent data streams to be successfully decoded at one spatial location. This makes the problem of beam selection even more challenging. To study single-user MIMO, we employ the same node deployment as in Fig. 4.13, and explore the performance of MUTE across all client positions and compare it against Exhaustive Search and the baseline scheme. Note that again, to ensure a fair comparison, the number of candidate beams in the baseline is the same as in MUTE.

Fig. 4.15 depicts the multiplexing gain of 2×2 single-user MIMO across different client positions. The multiplexing gain is computed by dividing the aggregate PHY rate of two-stream transmission over the maximum data rate of the corresponding SISO transmission being realized by employing the (TX, RX) beam-pair with the highest SNR. Theoretically, this setup should achieve close to $2\times$ gain; however, this does not hold true for every client position even under Exhaustive Search as shown in Fig. 4.15. This is because we find in measurements that the signal strength along the NLOS path is typically lower than the LOS one. Hence, even if the inter-stream interference is negligible, the aggregate PHY rate of a two-stream transmission (via a

LOS and an NLOS path or two distinct NLOS paths) might not obtain $2\times$ gain over the SISO transmission.

Although single-user MIMO beam training is inherently more challenging, MUTE is still able to obtain 90% of the Exhaustive Search’s multiplexing gain across all receive positions. The reason is that MUTE, by design, identifies all dominant paths and includes analog beams capturing diverse or ideally orthogonal paths.

In contrast, Fig. 4.15 reveals that the baseline scheme does not support transmission of two streams as the multiplexing gain is lower than unity for several client positions. This is because the baseline scheme selects analog beams that capture the LOS path as they provide higher SNR. Hence, to due lack of rich scattering, the vector channel of first stream is highly correlated with the second stream and the client cannot successfully decode both streams. Sec. 4.5.2 elaborates more on the limitations of the baseline (SNR-based Beam Selection) scheme.

Finding: In single-user MIMO, transmission of multiple streams across the LOS path hinders multiplexing gain; however, MUTE successfully creates separate beam transmissions via diverse paths and achieves 90% of optimal performance.

4.5.2 Limitations of SNR-based Beam Selection

4.5.2.1 Sparsity of Channel, Richness of Strong Beams

We hypothesize that only a few beams can provide sufficient SNR for multi-Gbps communications. This implies that there might be a few beams that can potentially support MIMO as most beams would not even provide the sufficient link budget for a SISO transmission. To investigate this premise, we conduct over-the air experiments and explore the distribution of “strong” beams (i.e., beams that support at least 1 Gbps data rate in the SISO configuration). We use the same node placement as depicted in Fig. 4.13 and measure the received SNR for all 25×25 beam-pair

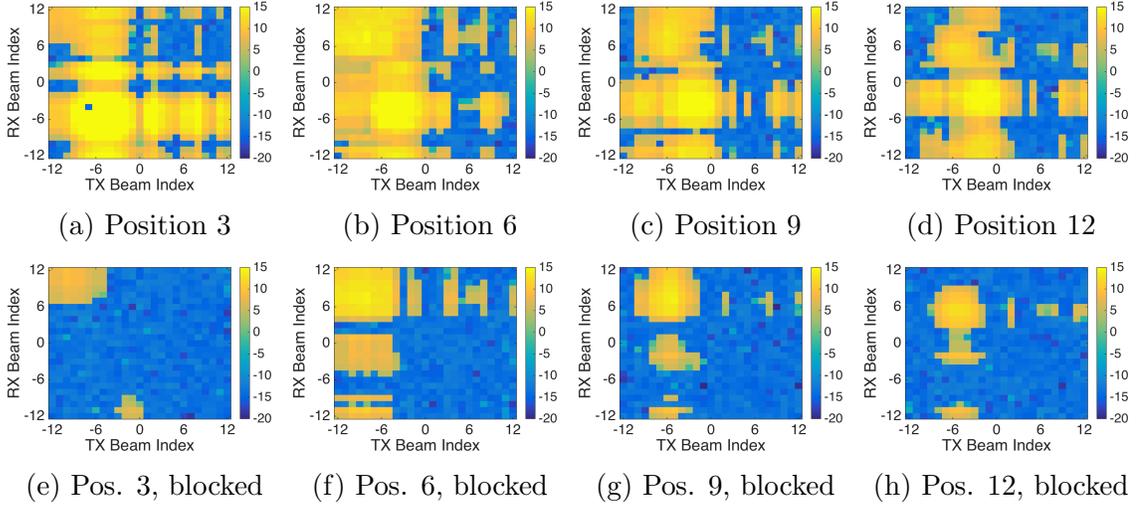


Figure 4.16: The SNR heatmaps for all 25×25 beam-pairs for client position indices 3, 6, 9 and 12 under LOS connectivity (top row) and blockage (bottom row).

combinations under two scenarios: *(i)* the AP has a LOS path to the client; *(ii)* the LOS path is blocked with a wooden table as shown in Fig. 4.17. We represent each beam sweep as a heatmap of corresponding SNR values with TX beam indices along the x-axis and RX beam indices along the y-axis.

Fig. 4.16 presents the SNR heatmaps for client position indices 3, 6, 9 and 12 under LOS connectivity (top row) and blockage (bottom row) (due to space limit, we do not show the SNR heatmap for all positions). The SNR range in the heatmaps is between -20 dB to 15 dB with yellow colored regions indicating beam-pairs with SNR above 10 dB whereas blue regions indicate beam-pairs with negative SNR.

Surprisingly, we observe several beam-pairs providing above 10 dB SNR that corresponds to 1 Gbps data rate in our platform. The received SNR corresponding to an analog configuration is a function of its captured physical paths and the directivity gain along them. Imperfect beam patterns cause a physical LOS/NLOS path to be captured by multiple beams, albeit with different directivity gains. We confirm that the beam-pair with the highest SNR corresponds to the physical LOS path in Fig. 4.16(a)-(d). For instance, the direction of the geometrical LOS path between



Figure 4.17: LOS blockage with a wooden table.

TX and RX at position 3 is $\tan^{-1}\left(\frac{-1.5}{3.3}\right) \approx -25^\circ$. Since the main lobes of $X60$ beam patterns are spaced roughly 5° apart, beam index -5 at TX and -5 at RX should provide highest directivity gain along the LOS path. Fig. 4.16a confirms that beam-pair (-5,-5) is within the high SNR region; however, due to overlap between neighboring beams, multiple beams include the LOS path and we see a cluster of high SNR beam-pairs around (-5,-5).

Under LOS blockage, the yellow region corresponding to the LOS component disappears in the bottom row plots of Fig. 4.16. This confirms that the signal strengths of neighboring beams are highly correlated as they capture one common path. Surprisingly, we observe that other high SNR beam-pairs with *large codebook distance* from the LOS region also experience significant SNR reduction after blockage. This implies that whether two beams capture the same path cannot simply be inferred from their RF codebook distance due to the irregularity and imperfections of beam patterns. Lastly, the highest SNR region after blockage achieves similar SNR under LOS conditions, i.e., LOS blockage has not degraded their SNR; thus, these beam-pairs must be capturing a reflected path. While MUTE discovers the LOS path as well as reflected paths, the baseline scheme likely selects LOS beams as they provide higher SNR.

Finding: While there exist a few physical paths between any two nodes, several beams may capture at least one path and thus provide high SNR. Whether two beams capture the same path or not cannot simply be inferred from their codebook distance due to the irregularity and imperfections of phased array generated beam patterns.

4.5.2.2 Zero-Forcing to the Rescue?

Traditionally, digital precoding schemes such as zero-forcing are employed to mitigate or ideally cancel inter-stream interference. Here, we explore the impact of such digital precoding techniques on the performance of SNR-based beam selection. In particular, we explore whether digital precoding can compensate for a *bad* choice of analog beams. To this end, we repeat the same experiment as in Sec. 4.5.1.2 and consider a special case that the candidate set size at the AP and client are equal (i.e., $m = n$). In particular, we vary m from 2 to 25, which is the total number of available beams in X60 (note that for $m = 25$, the baseline scheme turns into the Exhaustive Search).

Fig. 4.18 shows the normalized aggregate PHY rate with and without zero-forcing for the baseline scheme as the candidate set size varies between 2 to 25. The normalization is computed based on the achievable aggregate rate of the Exhaustive Search (i.e., $m = n = 25$). As expected, applying zero-forcing mitigates interference and boosts the achievable aggregate rate. Interestingly, to provide 90% of Exhaustive Search's aggregate rate, the baseline scheme requires the candidate set size to be 16 (which is more than half of the RF codebook size) for a two-stream analog-only beam steering. Applying zero-forcing mitigates interference and enables the baseline approach to achieve 90% of Exhaustive Search's aggregate rate with 10 candidate beams. This result implies that although zero-forcing reduces the inter-stream interference, it cannot completely compensate for a *bad* analog beam selection as the performance gap with the optimal solution remains significant.

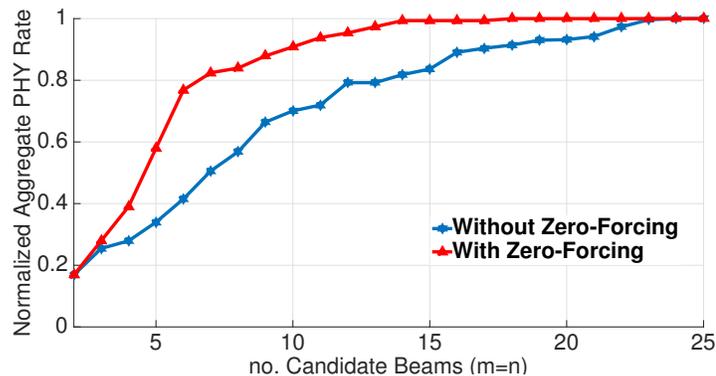


Figure 4.18: Aggregate PHY rate of baseline scheme as a function of candidate set size in 2×2 single-user MIMO.

The reason is behind how zero-forcing performs: Zero-forcing cancels inter-user interference by projecting the channel vector of a user on a precoding vector that is orthogonal to the channel vector of the other user so that the precoded (i.e., projected) channels become orthogonal to each other. The penalty one pays for such interference cancelation depends on the mutual channel correlation between users, i.e., channel projection incurs signal energy loss if the original channel vectors are not orthogonal. In sub-6 GHz bands, the rich scattering propagation causes semi-orthogonal channels and thus zero-forcing can successfully cancel interference (with low penalty). In contrast, 60 GHz channels are sparse and, more importantly, the effective channel vector of each user depends on the choice of analog beams that amplifies certain paths and weakens others. We have demonstrated that the high SNR beams typically share a common LOS path and thus incur high channel correlation. Hence, exploiting zero-forcing in the digital domain cannot compensate for low stream separability in the analog domain.

Finding: Although zero-forcing mitigates inter-stream interference, it cannot compensate for a bad choice of beams with high channel correlation and low analog domain separability.

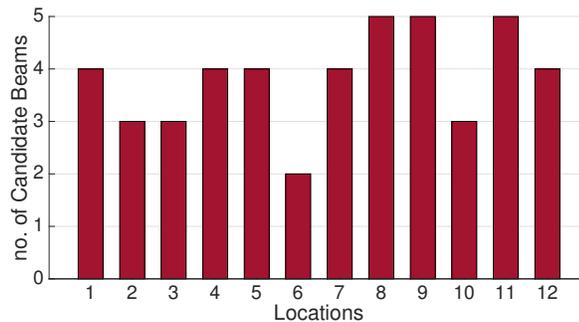


Figure 4.19: The number of candidate beams by MUTE in each client’s location.

4.5.3 Training Overhead

As discussed in Sec. 4.2, MUTE constructs a candidate set of beams and the final analog configuration is found by a local search among all combinations of beams in the candidate sets. Although MUTE obtains the candidate beams by passively overhearing training frames from initial beam acquisition phase with zero additional overhead, local search incurs additional time overhead that is proportional to the candidate beam set size.

Fig. 4.19 presents the size of candidate sets selected by MUTE for single-user MIMO configuration (the number of candidate beams for a multi-user transmission depends on the choice of target user group). We observe the candidate set size can be as low as 2 and is at most 5. MUTE, by design, adds exactly one candidate beam corresponding to every physical path in the aggregate delay profile; hence, the sparsity of 60 GHz channels results in selection of only a few beams (on average four).

Combining this result with the one in Sec. 4.5.1, we conclude that MUTE is able to achieve 90% of optimal aggregate rate by performing local search over only a few (on average four) candidate beams at the AP and client. In contrast, Exhaustive Search has to test all $\binom{25}{2} \times \binom{25}{2}$ beam combinations. Therefore, MUTE can achieve 90% of optimal aggregate rate while incurring only $\frac{\binom{4}{2}\binom{4}{2}}{\binom{25}{2}\binom{25}{2}} = 0.04\%$ of Exhaustive Search’s training overhead. In other words, MUTE reduces the training overhead by

99.6% with only 10% throughput loss. Furthermore, we observed in Sec. 4.5.2.2 that the baseline scheme provides 90% of the optimal aggregate rate when the candidate set size is 10. Consequently, under similar aggregate rate, MUTE requires only 1.8% of the baseline schemes' training overhead ($\frac{\binom{4}{2}\binom{4}{2}}{\binom{10}{2}\binom{10}{2}} \approx 1.8\%$).

Finding: MUTE achieves 90% of optimal aggregate rate while inducing only 0.04% of Exhaustive Search's training overhead or 1.8% of baseline scheme's overhead.

4.5.4 Scaling the Number of Spatial Streams

So far, we have evaluated the performance of MUTE for two-stream transmission. Here, we increase the number of spatial streams in a multi-user MIMO setting. To this end, we use the same node deployment as in Fig. 4.13 and for a k stream transmission, we consider all possible user groups consisting of k users out of 12 (i.e., total of $\binom{12}{k}$ different user groups). For each user group, we find the beam steering configuration under baseline, MUTE, and Exhaustive Search and report the achievable aggregate PHY rate after applying zero-forcing.

Fig. 4.20 depicts the achievable aggregate PHY rate as a function of the number of spatial streams. First, we observe that exhaustive beam steering is able to achieve on average about $2\times$ and $3\times$ throughput gain via simultaneous transmission of two, and three data streams, respectively. However, by further increasing the number of spatial streams from 4 to 6, the system's aggregate rate deviates from the *ideal* case which would linearly scale with number of spatial streams. This implies that the MIMO multiplexing gain does not endlessly increase proportionally with the number of streams because of undesired channel correlations. However, the saturation point depends on the AP-clients setting.

Second, while the baseline strategy provides 69% of the maximum sum-rate for two-stream transmission, its performance gap with Exhaustive Search increases with

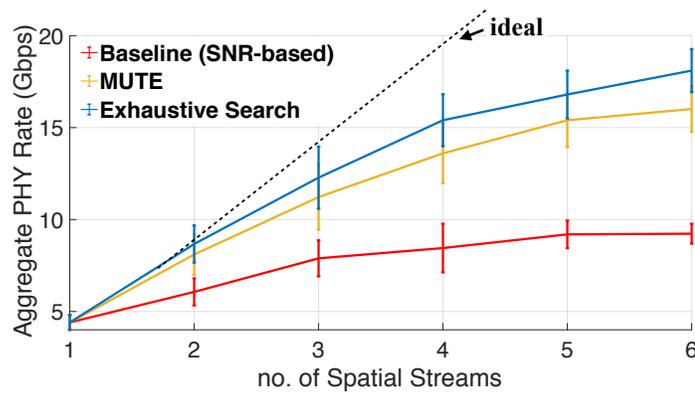


Figure 4.20: The aggregate PHY rate as a function of the number of spatial streams.

more streams such that it provides only 50% of optimal rate with 6 spatial streams. This is because the baseline scheme attempts to choose beams based on their achievable signal strength at the intended client and relies on zero-forcing to cancel inter-stream interference. However, the undesired channel correlation among a larger set of users cause inefficient interference cancellation and degradation in the relative performance of the baseline scheme. In contrast, MUTE accounts for undesired channel correlations by selecting beams over diverse or orthogonal paths. As a result, with increasing the number of spatial streams, MUTE’s gain remains close to the performance of Exhaustive Search (i.e., with marginal loss).

Finding: While multi-stream beam streaming becomes more challenging with an increasing number of spatial streams, MUTE is able to provide 87% of the maximum aggregate rate (realized by Exhaustive Search) with six streams.

4.6 Related Work

Prior work on 60 GHz WLANs mainly focuses on efficient beam training and tracking for *single-stream* transmission. To the best of our knowledge, MUTE is the first work on multi-stream beam steering.

Single-Stream Beam Training. Prior efforts reduced the frequency of single-stream beam training procedures in mobile 60 GHz WLANs via a variety of in-band and out-of-band solutions [36, 46, 48, 79, 86]. Such work is complimentary to MUTE as our method can be employed whenever the AP invokes beam sweeps, even if their frequency has been optimized. For example, link failure under blockage and mobility and the presence of a new user triggers the direction acquisition beam sweeps. MUTE passively overhears these beam sweeps and infers the beam separability by estimating PDPs.

Channel Profiling. Reverse-engineering 60 GHz channels has been explored in prior work with a different purpose of improving network connectivity in mobility and blockage [42–44]. Although such mechanisms obtain an aggregate channel profile (including LOS direction and location of reflectors), they do not provide stream separability inference, primarily due to lack of beam-specific multi-path profile knowledge, especially with imperfect beam patterns generated by practical phased arrays.

Direction Estimation. Direction inference techniques in sub-6 GHz bands have been studied in prior work [87–89]. These techniques employ the phase difference at multiple antennas for AoA estimation. However, due to a different node architecture at 60 GHz band (lacking one RF chain per antenna), we can only acquire a composite channel at the RF chain, where signals from multiple antenna elements are mixed, thereby thwarting AoA estimation. Likewise, direction estimation for sub 6 GHz MU-MIMO with hybrid beamforming was explored in [75]. However, their approach is limited to LOS detection and requires the analog weight vectors to be orthogonal, which does not hold true for practical phased arrays with limited phased levels. In contrast, MUTE takes advantage of GHz-scale sampling rate and sparsity of 60 GHz channel to obtain high resolution power delay profiles yielding to stream separability inference.

60 GHz MIMO. Prior work studied the potential of beam steering and spatial multiplexing in 60 GHz WLANs and showed that its directional nature motivates spatial reuse [41, 72, 74, 90]. Other works have explored hybrid beamforming [57] and user selection for multi-user 60 GHz WLANs [2]. In contrast, we aim to find the best analog configuration for multi-stream beam steering to a given target user-group.

Single-Shot Single-Antenna Motion Sensing in THz Networks

In this chapter, we present OPERA, One-shot Path discovEry with a THz RAinbow, a novel system that senses and tracks user motion and the environment via a single-shot THz pulse transmission to adapt THz wireless network to mobility and blockage. In particular, for efficient mobility management, OPERA tracks the LOS path between the AP and all clients as well as the receiver rotation so that nodes can proactively realign their beams. To recover from blockage, OPERA also identifies a small number of NLOS paths (that exist as a result of signal reflecting off of existing objects in the environment) that can provide high-rate backup links.

Our key idea is to transmit distinct signals with unique signatures across different angles such that each physical path has its own signature. To create a unique signature in each direction *simultaneously*, we introduce a novel THz node architecture in which the AP and clients are equipped with a Leaky-Wave Antenna (LWA).¹ A LWA allows traveling wave to leak from a waveguide into free-space with the unique property that the emission angle from the waveguide is coupled to the frequency of the input signal

¹We interchangeably use leaky waveguide and leaky-wave antenna in this chapter.

via a simple closed form and monotonic relationship [91]. We exploit this property in new ways and develop the *first* LWA-based path discovery system: First, we create a “THz Rainbow” by exciting the LWA with a time-domain THz pulse (equivalently, a broadband signal in frequency-domain) such that the signal emitted at each angle has a unique signature, thus filling the entire angular space with signal. Namely, each direction has a unique spectral peak or “color” due to the LWA’s frequency-angle coupling. Second, in addition to its spectral peak, known from [91], we discover that each direction also has a unique spectral *signature*. Consequently, we develop a model for predicting the angle-dependent spectral signature, based on the physics of EM waves propagating through and emitting from LWAs. The model is a deterministic function of the LWA geometric properties and thus can be known *a priori* by the AP and clients. With this design, the receiver can correlate its received portion of the THz Rainbow with the known transmission signatures, thereby identifying the transmission direction of each of the sender’s paths and tracking the amount of receiver rotation between two consecutive THz rainbow transmissions.

Our design is a non-coherent path discovery scheme in that it only requires power measurement at the receiver, and not phase information. This relaxation simplifies the THz node architecture, eliminates the need to keep tight synchronization between the transmitter and receiver, and is robust to small-scale channel variation (as opposed to phase with substantial sensitivity to sub-millimeter scale movement in the THz regime). Moreover, path discovery in OPERA is scalable, such that *all* clients can simultaneously and independently identify their LOS and NLOS paths via analyzing their received power spectrum from a single-shot transmission of the THz pulse shaped by a LWA.

Leveraging the design of OPERA, we present the *first single shot single antenna path discovery system in THz networks*. We implement the key components of

OPERA on a THz testbed and perform extensive over-the-air experiments in various settings. Our testbed consists of a THz pulse generator, broadband receivers, and custom LWAs. Our key findings are as follows:

(i) We experimentally demonstrate a THz Rainbow for the first time and measure its complex spatial-spectral profile. We find that any frequency component within the bandwidth of source emerges from the LWA in a range of emission angles, including but not limited to the angle known from literature [91]. In addition, we observe that the angular span of such emissions is itself frequency-dependent causing a unique spatial-spectral signature at each direction.

(ii) Even though the spectral peak has a one-to-one relationship with emission angle, our experiments reveal that discovering paths solely via this relationship incurs critical limitations: (1) estimation accuracy is not robust across the entire angular range and (2) reflected paths non-uniformly impact power over different frequencies which can lead to a spectral peak shift and subsequent estimation inaccuracy.

(iii) Instead, OPERA takes advantage of the unique spatial-spectral signatures of the THz Rainbow and achieves average estimation error of $< 1.1^\circ$ for the first path and $< 1.9^\circ$ for the second path. More interestingly, over the range of accessible rotation angles, an average estimation error of less than 2° is obtained.

(iv) Moreover, we exploit the path directions provided by OPERA to adapt directional THz beams under client mobility. We find that the non-zero error in path estimates causes slightly misaligned beam steering and subsequent SNR penalty of 0.29 dB on average. Yet, as the time interval between two path discovery measurements increases, the corresponding SNR loss becomes larger. Interestingly, our analysis shows that the time overhead of a THz Rainbow is orders of magnitude smaller than regular client mobility time-scales. Hence, the overhead of OPERA, even if executed before every data transmission, is relatively small.

The rest of this chapter is organized as follows: Sec. 5.1 presents a primer on leaky-wave antenna. Sec. 5.2 illustrates our system architecture. In Sec. 5.3, we introduce the idea of THz Rainbow and model its unique characteristics in Sec. 5.4. Sec. 5.5 and Sec. 5.6 discuss the design details of OPERA’s path sensing and rotation sensing modules, respectively. Sec. 5.7 introduces our experimental platform. Sec. 5.8 and Sec 5.9 describe the experimental results on path discovery and rotation tracking, respectively. Finally, Sec. 5.10 reviews the related work.

5.1 A Primer on Leaky-Wave Antenna

A LWA consists of a pair of parallel metal plates with open sides and a slot in one of the metal plates. Fig. 5.1 depicts a LWA device with separation b between its two metal plates. By opening a single slot in one of the two plates, we permit some of the radiation to “leak” out into free-space.

Angle-Frequency Relation. The energy emitted at a particular angle is a function of input frequency with the overall emission pattern known to be very complex. With some simplifications, Maxwell’s equations with boundary conditions between the waveguide mode and free-space mode yield a direct relation between the maximum emission angle and the frequency of the input signal [92]. For the lowest transverse-electric mode, we can derive [93]:

$$\phi_f = \sin^{-1} \left(\frac{c}{2bf} \right), \quad (5.1)$$

where f is the frequency of the input signal, c is the free-space light speed, and b represents the distance between the two metal plates. Other geometrical parameters, such as the width and length of the leaky-wave aperture, can impact the efficiency of energy transfer between the guided mode and free-space, but not the maximum-

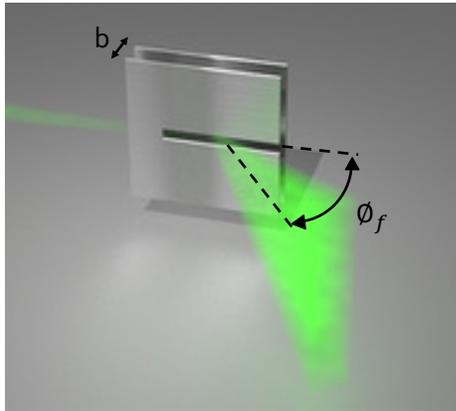


Figure 5.1: A Leaky-Wave Antenna in which the input signal emerges from the slot such that the maximum emission angle ϕ_f is related to input frequency via Eq. (5.1). The larger the input frequency, the lower the emission angle relative to the plate's axis.

energy emission angle. The subscript f in ϕ_f emphasizes that, given a fixed LWA structure, this angle is solely a function of input frequency. Moreover, Eq. (5.1), which we also refer to as a first principle model, suggests that the larger the input frequency, the lower the emission angle from the slot.

Conversely, when operated as a receiver, free-space waves with the frequency f would couple best into the waveguide if they are arriving at the slot with the correct angle ϕ_f . Thus, one can expect a simple and monotonic relationship between frequency and maximum-energy angle, for both the case of a guided wave radiating into free-space and the case of a free-space wave impinging on the device and coupling to a guided wave (i.e., for both transmission and reception of signals).

5.2 System Architecture

We exploit LWAs in the control plane to discover dominant paths in highly directional THz channels. We equip both the AP and client with LWAs. Fig. 5.2 depicts the core of our proposed node architecture which consists of a LWA antenna, a broadband receiver (co-located with the LWA), and a broadband transceiver that can generate

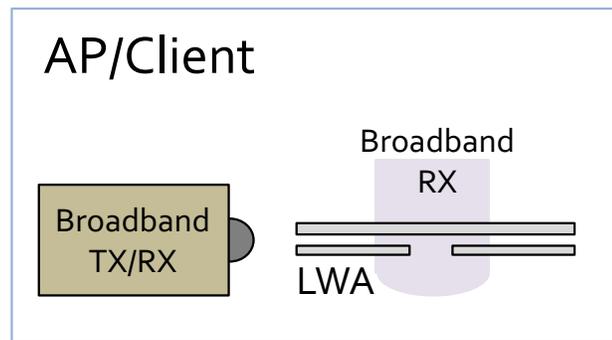


Figure 5.2: The control plane architecture consisting of a LWA, a broadband transceiver, and a broadband RX.

THz pulses or detect broadband (up to 1 THz) signals. In transmission mode, the injected signal to the LWA leaks out with an emission that falls within a sector of 90° , depending on the frequency of the signal. In reception mode, the waves are coupled into the waveguide and then guided toward the broadband receiver located at end of the metal plates.

Note that Fig. 5.2 shows the control plane architecture only. For data transmission, nodes can use another beam steering architecture (such as phased-array architecture in [94] or pixelated metasurface structures in [95]). We consider that such data plane components are either co-located with the LWA or their geometrical properties are known so that we can directly map the discovered paths in the control plane to the data plane.

Deployment Costs. LWA structures are passive, light (< 3 oz), and cheap ($< \$10$) and can easily meet the power, size, weight, and cost considerations of future handheld devices. Further, THz broadband transceivers have been designed and implemented using CMOS-technology in prior work [96]. With CMOS-technology, the cost of a broadband transceiver in the THz regime would not be significantly higher than today's commercialized transceivers in lower bands. Hence, we anticipate that the large-scale deployment of OPERA would not raise any significant cost concerns.

Towards 360° Coverage. Our system uses a single LWA to transmit within a sector of 90 degrees. The angular range can be increased to 360° by extending the node architecture. In particular, employing a multi-face LWA structure increases the angular coverage. Yet, simultaneous transmissions from multiple LWAs raises an ambiguity concern for the receiving node, who may not be able to tell from which LWA a detected signal originated. One way to tackle this issue is to exploit an additional degree of freedom in LWAs (e.g., the shape of the leaky-wave slot, which need not have a constant width along its length) to generate transmission patterns with slightly different spectral-spatial signatures for each segment. Building and testing a multi-face LWA structure would be an interesting engineering problem that we leave for future work. Here, for ease of exposition, we limit scope to one 90° sector.

5.3 THz Rainbow: A Unique Pattern

Leaky-wave antennas have a long history, having been used in the RF region since at least the 1940's [97] and more recently, in the THz regime [93]. However, this work is the first to employ LWAs as a control plane for the THz WLANs. We propose, for the first time, to excite a LWA with a THz pulse. THz pulse generation and its conventional applications in spectroscopy has been studied in prior work [98]. Here we excite the LWA with a THz pulse source (or equivalently, an ultra-broadband input signal in frequency-domain).¹ Due to the angle-frequency coupling in a LWA, different frequency components are decomposed and radiated into the free-space at different angles forming what we call a “THz Rainbow”.

In this scheme, the entire space is filled with THz signals such that different frequencies (within the bandwidth of the source) simultaneously directed to different angles across the entire angular range. Fig. 5.3 depicts the THz Rainbow with angular

¹In our platform, the 3 dB bandwidth of the broadband source is 1 THz.

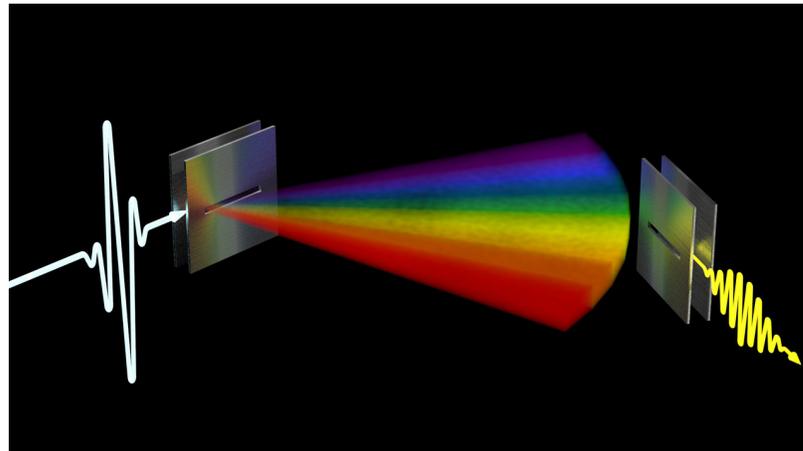


Figure 5.3: Creation of a THz Rainbow by exciting a LWA device with a THz pulse.

dependency in the spectral signature of radiating signals. In particular, in lower angles (relative to the axis of the LWA), the radiating waves contain higher frequency components and are fast decaying in the time-domain. In contrast, waves emitting at larger angles are wider in the time-domain and have relatively smaller frequency components. A broadband receiver in the far field of the LWA captures only a portion of the THz Rainbow. Thus, OPERA takes hints from the received signal to estimate the angles of departure and arrival corresponding to the LOS and all reflected paths between a TX and RX.

As a result of diffraction from the finite LWA aperture, emission of waves from LWA to free-space spans in 3D (see Fig. 5.1 for the emission of a single-frequency signal). Although not shown in Fig. 5.3, the THz Rainbow has a 3D transmission pattern. Hence, in principle, we can extend the spatial-spectral signatures and mathematical formulation to extract both azimuth and elevation angles from the received signal. We leave exploration of single-shot 3D path discovery with THz Rainbow for future endeavors and focus on 2D path discovery in this thesis.

5.4 Spatial-Spectral Signatures of THz Rainbow

In order to relate the received signal to the properties of the underlying physical paths (e.g., AoA and AoD), we need to characterize the mediums in which waves propagate to reach the RX. Fig. 5.4 depicts different components of an end-to-end channel between the source and RX: An input signal is guided inside the LWA until it radiates out of the slot. We assume propagation inside the LWA is ideal. While the waveguide introduces some dispersion (since this is a transverse electric waveguide mode) and loss (since the plates are made of real metals with non-infinite conductivity), both of these effects are quite small [99]. However, radiation from the LWA to the air imposes spatially-biased attenuation. The signal subsequently traverses the available paths of the wireless medium. Finally, impinging wave are coupled into the LWA and guided toward the RX.

For radiation out of the waveguide into free-space, we know the relation between spectral peak and emission angle from Eq. (5.1). However, our measurements show that any specific frequency component is not a spatial delta function. Instead, the power spectrum of signals leaking from the waveguide can appear at a range of angles, albeit encountering different coupling losses. More importantly, our measurements indicate that the angular width of the signal at any particular frequency is frequency dependent.

To understand this broader angular width, we analyze the underlying physical properties of leaky waves. We treat the leaky waveguide slot as a finite-length aperture, which produces a diffraction pattern in the far field. In this case, the angular distribution of the diffracted field (in the plane of the slot) can be derived for a given frequency f by treating the slot as a diffracting aperture of length L , such that the far-field pattern is the Fourier transform of the aperture function [92]:

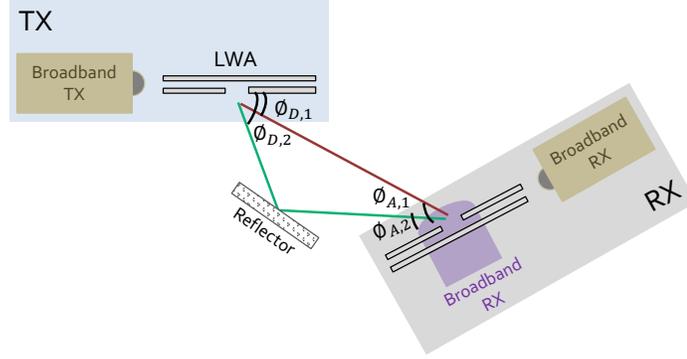


Figure 5.4: The end-to-end channel between a broadband source and a broadband RX.

$$G(\phi, f) = \text{sinc}\left[(\beta - j\alpha - k_0 \cos\phi) \frac{L}{2}\right] \quad (5.2)$$

where k_0 is the free-space wave vector ($k_0 = \frac{2\pi f}{c}$), β is the wave vector of the TE_1 guided wave ($\beta = k_0 \sqrt{1 - (\frac{c}{2bf})^2}$), and α is a parameter that describes the loss of energy in the guided mode due to leakage out of the slot. Note that for $\alpha \ll \beta$, the peak value of $G(\phi, f)$ happens at $\cos\phi = \frac{\beta}{k_0}$ which is the same as Eq. (5.1) except for a change in the angle definition. Note that Eq. (5.2) suggests that the diffraction pattern is not a function of transmitter-receiver distance.

While the rigorous diffraction formulation in Eq. (5.2) accurately describes the far-field radiation pattern, it is cumbersome to use as a fitting function, in particular because the side lobes associated with a sinc function fall below the dynamic range of our measurement system, and are therefore not observed in our data. Further, for the purpose of path discovery, the high-amplitude portion of the signature plays the key role. Hence, as an alternative, we introduce a second-order phenomenological model which is meant to characterize the high-amplitude spectral characteristics incorporating the observations from our measurements. For a given frequency f , we model LWA emission as a spatial filter or a directional radiation pattern with highest directivity along the previously predicted peak angle, and fall off on both sides. We

model this behavior as a filter that is a symmetric parabolic function of angle ϕ :

$$\tilde{G}(\phi, f) = \max(1 - \gamma(f)(\phi - \phi_f)^2, 0), \quad (5.3)$$

where ϕ_f is the maximum-energy angle corresponding to frequency f as defined in Eq. (5.1). The filter function $\tilde{G}(\phi, f)$ models the fraction of amplitude at frequency f that is emitted at angle ϕ w.r.t. the LWA's metal plate. At the center of the filter, i.e., at the optimum angle ϕ_f , the power is maximum and $\tilde{G}(\phi_f, f) = 1$. As we deviate from the optimum angle, power is attenuated with a frequency-dependent rate denoted as $\gamma(f)$. Specifically, our data indicates that the fall-off rate $\gamma(f)$ can be represented as a linear function of frequency such that the angular width becomes narrower as frequency increases. With a linear $\gamma(f)$ function, namely, $\gamma(f) = \gamma_0 + \gamma_1 f$, we have only two parameters to fit a family of quadratics. There is no closed-form expression available for $\gamma(f)$; yet, for a fixed node architecture (i.e., a given LWA geometry), it is a deterministic function of frequency which can be measured and known *a priori*.¹

For angles ϕ far from ϕ_f , $1 - \gamma(f)(\phi - \phi_f)^2$ becomes negative. For these values of ϕ , we assume that there is zero power detected (i.e., the signal simply misses the detector), and so we set $G = 0$. In Sec. 5.8, we validate this LWA-to-air emission model. Now that we have modeled the transmission pattern of a THz Rainbow, we can proceed to path discovery by exploiting the known and fixed THz Rainbow pattern as angular-spectral signatures.

¹For our particular LWA structure, we find $\gamma(f) = 0.032f + 0.01$, when frequency is in THz unit and angle is in degrees.

5.5 Path Sensing Module

In this section, we describe the design of OPERA that aims to discover dominant physical paths in THz channels in a single shot and using a single antenna.

5.5.1 Design Overview

The higher path loss in THz frequency range necessitates the use of directional beams at both sender and receiver. The best directions to transmit are typically unknown in advance and thus different directions need to be tested for their achievable signal strength. In most systems, the receiver needs to be explicitly informed of each direction the transmitter is testing so that it can notify the sender which direction was best. For instance, in directional 60 GHz WLANs, IEEE 802.11ad specifies a beam sweeping mechanism in which beacons are transmitted with an explicit and unique header in each direction [14]. Such a scheme involves several consecutive beacons to cover all possible directions and hence the time overhead increases with the total number of directions.

We propose OPERA, a single-shot single-antenna path discovery platform that identifies high-rate physical paths between the AP and clients. THz channel sparsity suggests that only a few dominant high-rate paths (LOS and possibly first-order reflected paths) exist between any two nodes [100]. Our key idea for path discovery is to embed path information into the inherent characteristics of the signal traveling along that path. In particular, we enable the receiver to correlate the sender's path direction with the spectral signature of the signal transmitted in that direction. For instance, one naive LWA implementation would be to deploy a beam sweeping mechanism analogous to that of IEEE 802.11ad, but vary the carrier frequency of beacons transmitted in different directions. Even though such an approach would assign unique spectral signatures (i.e., carrier frequency) to different directions, it would not

be a single-shot scheme as several consecutive transmissions would be involved.

Instead, we show how to use a single leaky waveguide antenna for efficient single-shot path discovery. In particular, we deploy a LWA device and introduce a *THz Rainbow* to simultaneously transmit waves whose spectral characteristics are a function of emission angle. While it may seem that obtaining the peak frequency as the direction dependent signature (i.e., angle-frequency coupling in Eq. (5.1)) would be sufficient for path discovery, our experimental analysis shows that this is not the case. Hence, we introduce a phenomenological model to characterize the angular-dependent spectral signature of waves composing THz Rainbow. Leveraging this model, we devise an optimization framework to extract the unique features of multiple distinct paths from the received power spectrum.

Finally, OPERA is an incoherent scheme, i.e., it solely relies on the power or signal strength measurement (and not phase information) to estimate the path angles. To exploit phase, the transmitter must first send a pre-defined preamble in order for the receiver to lock on to the phase; in contrast, OPERA eliminates the need to keep tight synchronization between TX and RX. Moreover, it is robust to small-scale channel variation as opposed to phase with substantial sensitivity to sub-millimeter scale movement in THz regime. Next, we illustrate the key components of OPERA.

5.5.2 One Shot Path Discovery

We envision periodic THz Rainbow excitation so that receivers are able to update their angular profile by assessing the received power spectrum. Next, we illustrate OPERA's AoD and AoA estimation techniques.

5.5.2.1 AoD Estimation

The THz Rainbow is composed of signals in all directions within a sector of 90° . However, only signals along the dominant paths of the over-the-air wireless channel can reach the receiver. Each path contributes to the power spectrum depending on its angular-dependent spectral signature. Hence, we can model the total power spectrum based on the aggregated effect of signals from all paths reshaped by the LWA spatial filter gain:

$$S(f) = \left| \sum_{p=1}^P \tilde{G}(\phi_{D,p}, f) \alpha_p(f) \right|^2 \quad (5.4)$$

where P is the number of multipath components, $\phi_{D,p}$ denotes the AoD corresponding to path p , and $\alpha_p(f)$ represents the complex wireless channel gain over path p .

While Eq. (5.4) models the expected power spectrum based on path directions, we can also measure the actual received power spectrum at our broadband detector. In practice, power measurements are available over a discrete set of frequencies depending on the clock sampling rate. We denote \mathbf{S}^{ms} as the vector of power measurements across the available set of frequencies. Similarly, \mathbf{S} captures the expected power spectrum based on Eq. (5.4) computed for the same set of frequencies. Note that \mathbf{S} , by definition, is a sparse tall vector as $\tilde{G}(\phi_{D,p}, f_i)$ is zero for frequencies far from the peak frequency corresponding to $\phi_{D,p}$. Hence, we estimate the AoD that best describes the power measurements across the entire spectrum:

$$\phi_{D,1}^* = \arg \min_{\phi} \|\mathbf{S} - \mathbf{S}^{\text{ms}}\|^2 \quad (5.5)$$

where $\phi_{D,1}^*$ is the estimate angle of the strongest path. OPERA assumes that other paths, if present, are significantly weaker such that the received signal signature is dominated by the spectral signature of the strongest path. Once the strongest path is identified, we remove its spectral signature from the measured power spectrum \mathbf{S}^{ms}

and repeat the optimization in Eq. (5.5) to discover the second path. OPERA stops this iterative process when the power to noise ratio of the residual signal drops below a configurable threshold. We emphasize that no phase information is used in this methodology, hence the detector can be non-coherent.

5.5.2.2 Opportunistic AoA Estimation

In principle, assuming that the THz wireless channel is reciprocal, AoA can be estimated similarly by generating a THz Rainbow at the client and assessing the received spectrum at the AP. Potentially, we can take advantage of the receiver's node architecture to opportunistically estimate AoA, at least for a range of angles, without initiating a separate THz Rainbow. Namely, comparing the spectral characteristic of impinging waves at the LWA before and after coupling into the waveguide provides insights on the AoA. Intuitively, if the client's LWA is parallel to the AP's LWA (i.e., $\text{AoD}=\text{AoA}$), then the air-to-waveguide coupling loss is negligible and we expect a similar spectral characteristics for coupled and impinging waves. However, when the AoA is different from AoD, then the coupled waves would experience a frequency-dependent power loss that can be exploited to extract the AoA. Of course if the AoA is far from AoD, one would expect that the client would receive no signal. However, our preliminary analysis shows that the spectrally broader emission at a specific angle enables a fairly large $|\text{AoA}-\text{AoD}|$ without complete loss of signal. In any case, the client can always generate a THz Rainbow and repeat the mechanism described above to explicitly estimate the AoA of the dominant paths. In our experiments, we focus on evaluation of AoD/AoA with the explicit THz Rainbow generation and we leave the exploration of opportunistic AoA estimation and its limitations for future work.

5.5.2.3 Overhead and Interference Analysis

The AP periodically generates the THz Rainbow while all clients measure the received signal with their broadband detectors to update their path profiles. We denote T_{rb} as the cycle duration of THz Rainbow transmissions. A smaller T_{rb} indicates that path information updates are available more frequently with the cost of higher time overhead. However, we argue that even if we update path information before every data packet transmission, the interference and overhead costs are insignificant:

First, THz Rainbow transmission incurs negligible overhead, in the range of few tens of nano second depending on the THz pulse generation method [96] whereas mobility is in millisecond scale. This means that to manage regular client mobility (e.g., walking or running speed), path information needs to be updated, at most, in sub-millisecond scale while OPERA is orders of magnitude faster, causing the ratio of control plane overhead to data transmission to be negligible. Moreover, OPERA has a scalable design such that all clients can simultaneously and independently update their AoD information via a single sender-side THz Rainbow transmission.

Second, it may seem that OPERA is a resource-intensive solution since it fills the entire space with THz signals composed of a broad range of frequencies that may cause interference with any on-going transmissions in neighboring cells. However, the THz Rainbow transmission only takes several nano seconds and not every frequency component exists in every direction. In fact, interference can take place only if a node in a neighboring cell is located at a specific angle and operates in a frequency band that matches the LWA angle-frequency relation. In other words, the footprints of the THz Rainbow are narrow in space and short in time. Further, as this transmission is broadband, the power of different sub-bands is limited and thus the signal is unlikely to be received at longer distances of neighboring cells. However, in an extreme case, the APs could coordinate to minimize the inter-cell interference.

5.6 Rotation Sensing Module

Rotation tracking is critical for the client so that it can proactively and locally realign its beam. Here, we show how to leverage the THz Rainbow for rotation estimation.

5.6.1 Design Overview

To explore the impact of rotation, we need to consider the reception model of a LWA. Interestingly, the reception model is reciprocal to the transmission model; hence, there is a coupling between the optimum incident angle and the frequency of impinging waves. If the client's waveguide is parallel to the transmitter's waveguide, then it is clear that a signal at a particular frequency will couple into the waveguide. However, if the client is rotated, then the two angles do not match and thus waves would encounter a coupling mismatch loss resulting in power reduction.

We know from the THz rainbow propagation model that there exists a spectral band at each angle. We show here that this broader spectrum enables a new method for rotation estimation of the mobile client. In particular, we track changes within the spectral band of rotated RX to estimate the amount of rotation between two consecutive transmissions of THz Rainbow. Next, we first introduce a ray optics model that can capture the spectral width and then illustrate our rotation estimation technique.

5.6.2 Ray optics model

An alternative description of spectral width at each angle relies on ray optics. For a LWA with infinitely thin metal plates, the energy leakage is determined only by phase matching. However, for a plate of finite thickness, the slot itself acts as a waveguide, which presents an impedance boundary between the TE_1 fast wave and

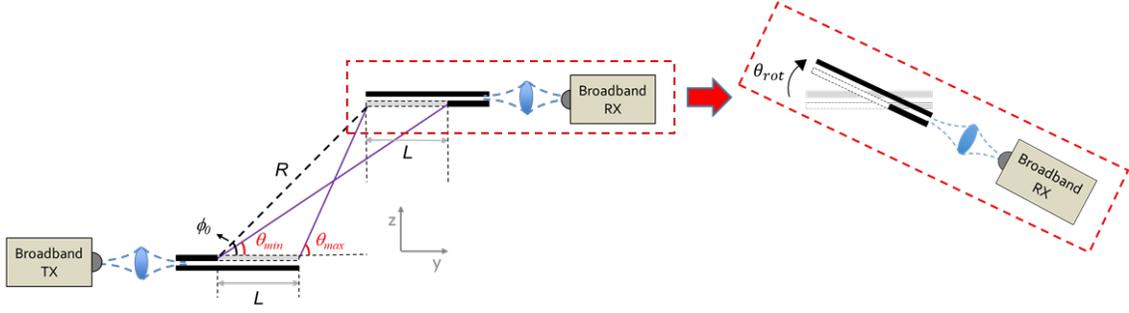


Figure 5.5: A schematic that illustrates the rotation of the RX waveguide and defines various parameters used in the discussion.

free-space. Rays can reflect from this boundary, and remain in the waveguide for a longer propagation distance before leaking out. As illustrated in Fig. 5.5, this results in a larger effective length for the emission region. In this figure, we denote the distance between the two LWGs along the z axis as Z_0 and along the y axis as Y_0 , such that $R^2 = Z_0^2 + Y_0^2$ (where R is defined in the figure). We can write two expressions for $\cos\phi_0$: (i) from geometry, we have $\cos\phi_0 = Y_0/R$, and (ii) from wave propagation, we have $\cos\phi_0 = k_y/k_0$. Here k_0 is the free-space wave vector and k_y is its y component. Combining these two expressions, we can write

$$Y_0 = \frac{k_y}{k_0} R \quad (5.6)$$

Similarly, we can describe Z_0 by

$$Z_0 = \frac{k_z}{k_0} R \quad (5.7)$$

Simple geometric considerations allow us to find:

$$\tan(\theta_{min}) = \frac{Z_0}{Y_0 + L} \quad \tan(\theta_{max}) = \frac{Z_0}{Y_0 - L} \quad (5.8)$$

After substituting Y_0 and Z_0 and some simplifications, we derive the minimum and maximum angles at which a light ray could be received, as:

$$\theta_{min} = \tan^{-1}\left(\frac{k_z R}{k_y R + k_0 L}\right) \quad \theta_{max} = \tan^{-1}\left(\frac{k_z R}{k_y R - k_0 L}\right) \quad (5.9)$$

Here, $k_z = \pi/b$ and $k_y = \sqrt{k_0^2 - k_z^2}$ are the z and y components of the free-space wave vector, respectively. We assume an effective slot length L which is identical for both transmitter and receiver. We note that this ray optics approach makes sense only in the limit where the rate of emission is large, such that the loss parameter α satisfies $\alpha L > 1$.

The above derivation is a simple ray-optics description which relies on knowledge of the (finite) length of the slot apertures in the two leaky-wave devices. In contrast, the underlying description of Eq. (5.2) involves diffraction through an aperture of finite length, and thus relies only on the length of the TX aperture, not the RX aperture.

Both the diffraction formalism and the ray optics picture can be used to predict the spectral bandwidth of radiation emitted at any given angle from the leaky-wave slot, assuming that the waveguide is excited with a broadband input. Later, in Sec. 5.9, we show the agreement on spectral width from these two models, and with results from over-the-air measurements. Hence, for the interest of computational complexity, we rely on the ray optics approach for subsequent discussion on rotation estimation.

5.6.3 Rotation Tracking

To understand the changes in the received spectrum that result from a rotation of the client (i.e., a non-zero value of θ_{rot}), we use a similar ray-based approach to predict the frequency distribution at a specific angle, given the geometry of transmitting and

receiving waveguides, and their separation. We denote θ_{min} and θ_{max} as the minimum and maximum angles a light ray could be received by a LWA with slot length L and propagation distance R (see Fig. 1b of the main text). We consider a RX LWA which is parallel to, and at angle ϕ_0 with respect to, the transmit LWA. In such a case, the AoD of rays radiating out of the transmit LWA is the same as their AoA when impinging on the receiver LWA; thus, coupling into the waveguide with minimum coupling loss at all radiated frequencies. Hence, as discussed above, we measure a spectral band whose low-frequency and high-frequency edges are determined by ϕ_0 .

Fig. 5.6 illustrates the key components of our rotation estimation model. We denote the low-frequency and high-frequency edges of this initial spectral band (when there is zero rotation) as $f_{min,0}$ and $f_{max,0}$, respectively. Then, we measure the changes in this spectral region of interest in order to detect clockwise (CW) or counter clockwise (CCW) rotation. Specifically:

CW rotation: In this case, the effective AoA of impinging rays increases by the amount of rotation (θ_{rot}). Hence, the larger frequency components experience higher coupling loss causing a reduction in the high-frequency edge of the measured spectrum. For a given rotation angle θ_{rot} , of a receiver located in parallel and in the far field of TX LWA, the spectral change is described by

$$\Delta f(\theta_{rot}) = f_{max,\theta_{rot}} - f_{max,0} = \left[\frac{\partial f_{max}(\theta)}{\partial \theta} \Big|_{\phi_0} \right] \theta_{rot}, \quad (5.10)$$

where $f_{max}(\theta)$ is the max-frequency contour shown as a solid black curve in Fig. 5.6. Note that the f_{min} of the measured spectrum is expected to remain the same with a CW rotation. The reason is, as shown in Fig. 5.6, the initial low-frequency edge ($f_{min,0}$) falls into the spectral band of larger angles (i.e., $\phi_0 + \theta_{rot}$). Also, the receiver is blind to any spectral change beyond the initial spectral region of interest; hence, frequencies beyond the range $[f_{min,0} - f_{max,0}]$ are not received.

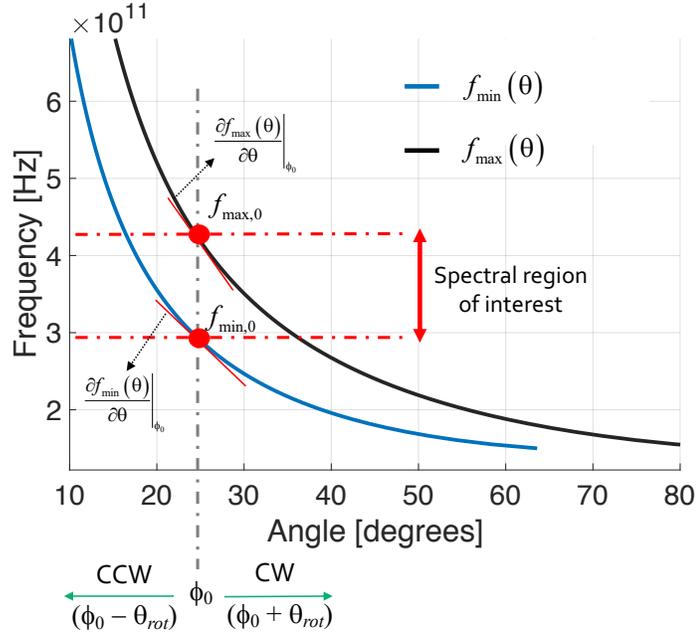


Figure 5.6: The key components of our rotation estimation model. We estimate the amount of rotation by tracking the changes in the spectral width of the received spectrum. A CW rotation causes a decrease in the high-frequency edge of received spectrum while a CCW rotation results in an increase of the low-frequency edge.

CCW rotation: Conversely, when rotated counterclockwise, the AoA of the impinging waves change to $\phi_0 - \theta_{rot}$ causing high coupling loss for lower frequency components in the spectral region of interest. Similarly, we relate the amount of rotation to the changes in the low-frequency edge (f_{min}) of the measured spectrum as follows:

$$\Delta f(\theta_{rot}) = f_{min,\theta_{rot}} - f_{min,0} = \left[\frac{\partial f_{min}(\theta)}{\partial \theta} \Big|_{\phi_0} \right] \theta_{rot}, \quad (5.11)$$

As above, the f_{min} of the measured spectrum after a CCW rotation is expected to remain the same and $f_{min}(\theta)$ is the low-frequency counter obtained via the ray optics model.

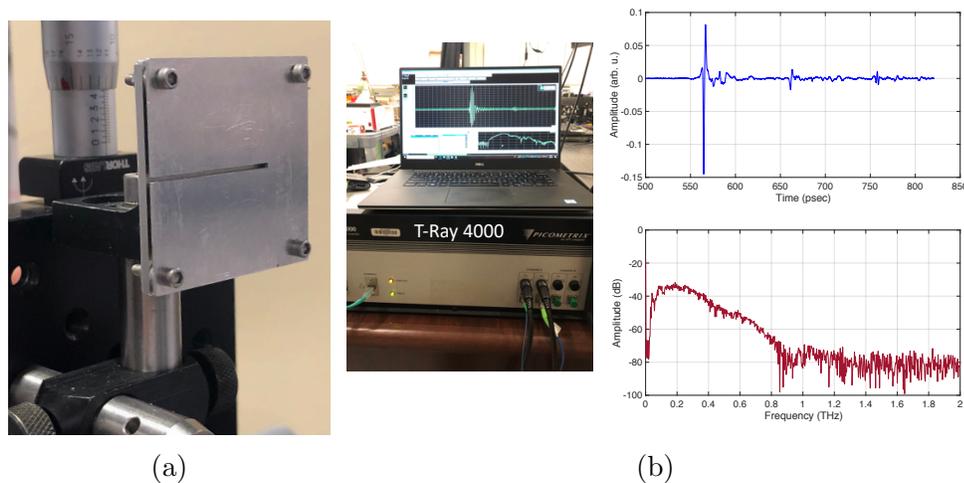


Figure 5.7: (a) Our custom $4 \times 4 \text{ cm}^2$ LWA; (b) The ultra-broadband transceiver system and its generated THz pulse in time and frequency domains.

5.7 Experimental Platform

We conduct extensive over-the-air experiments to characterize the THz Rainbow and evaluate the performance of OPERA. We exploit a custom LWA device, shown in Fig. 5.7a, which we build using $4 \times 4 \text{ cm}^2$ metal plates and 1.075 mm spacers between (i.e., plate separation $b = 1.075 \text{ mm}$). Other geometrical parameters include the plate thickness (1 mm), the slot length (3 cm), and the slot width (1 mm).

For THz pulse generation and detection, we use T-Ray 4000 TD-THz System from PICOMETRIX [101]. The interchangeable fiber-coupled sensor heads deliver a picosecond duration time-domain THz pulse. Fig. 5.7b depicts our ultra-broadband transceiver system and its generated THz pulse in time and frequency domains.

Fig. 5.8 shows our measurement setup. We first generate a THz pulse and focus it to the LWA via a lens with a focal distance of 6 cm. The signal leaking out of the waveguide generates a THz Rainbow. We exploit aluminum sheet as reflector to configure NLOS paths. To achieve different NLOS angles, we mount and move the reflector on a rail as shown in Fig. 5.8. Similarly, to vary the LOS angle, we mount the broadband detector on a rotating arm that rotates around the center of LWA.

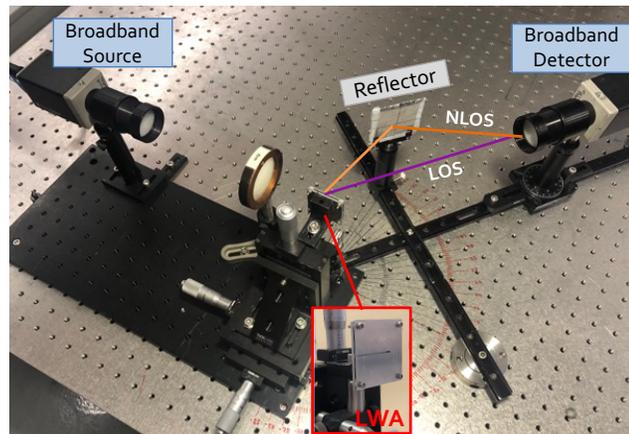


Figure 5.8: Our experimental setup.

The broadband detector provides raw measurements of the composite received signals in time and frequency domains. In our experiments, we limit the setup to one-path (LOS or NLOS) and two-path settings and leave exploration of more complex THz channels for future work.

Due to hardware limitations, i.e., low-power THz source, we are bound to conduct experiments in small scales (up to a meter). Scaling up the transmitter-receiver distance is feasible by increasing the power emitted by the transmitter. In particular, our THz source emits an average power of roughly -10 dBm, which is the time-averaged power, integrating over the entire (1 THz) spectral band. Hence, the power in a 1 kHz bandwidth is -100 dBm. Meanwhile, prior work reported a CMOS-based widely tunable source which achieves about -10 dBm in a 1 kHz bandwidth, over the entire range 100-300 GHz offering a factor of 10^9 increase in output power compared to our THz source [96]. Further, links at WLAN-scale distances (100+ meters) have already been demonstrated at many frequencies between 100 GHz and 557 GHz [102] and in several bands at higher frequencies up to 850 GHz [103]. Note that increasing the transmit power does not affect the spatial-spectral signature of THz Rainbow nor the design of OPERA.

5.8 Experimental Results on Path Discovery

In this section, we discuss our over-the-air experiments and evaluate the key components of OPERA.

5.8.1 Experimental Characterization of the THz Rainbow

We first experimentally characterize a THz Rainbow and validate the phenomenological model proposed in Sec. 5.4.

Setup. We deploy T-Ray 4000 THz system for THz pulse generation and broadband detection and integrate it with our LWA device. Fig. 5.9 depicts our experimental setup in which blue circles represent broadband detector locations. We conduct extensive experiments in multiple configurations covering different LWA-detector distances and relative angles. In each setting, we move the detector on a quadrant from 10° to 80° with the setup of 1° . We explore this angular range since we observe that below 10° and above 80° , the leaked waves are fundamentally very weak (close to noise floor), due to a decrease in the efficiency of coupling from the guided wave to free-space near the extrema of the angular range.

Methodology. In each setting, we identify the time window that the signal falls into through energy detection at our broadband detector. Then we lock the receiver to that time window and collect raw time-domain and frequency-domain data and apply signal processing techniques (such as smoothing, filtering, FFT, etc.).

5.8.1.1 Measuring The Angle-Frequency Relation

We first experimentally explore the angle-frequency coupling and compare it against the first principles model in Eq. (5.1). Fig. 5.10 plots the peak frequency of power spectrum vs. receiver's angle. The results show good agreement between the theory and experiment, at least for the angles above 20° . However, below 20° , the measure-

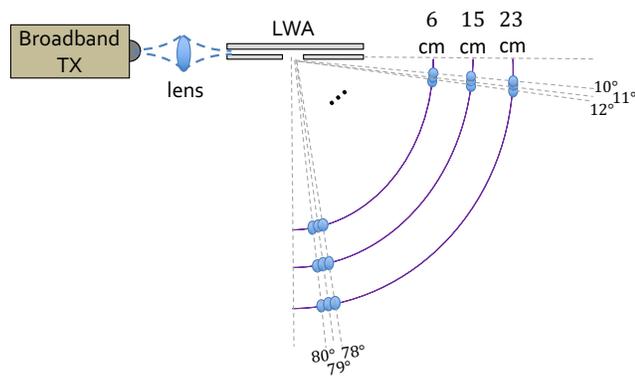


Figure 5.9: The schematic setup for THz Rainbow characterization.

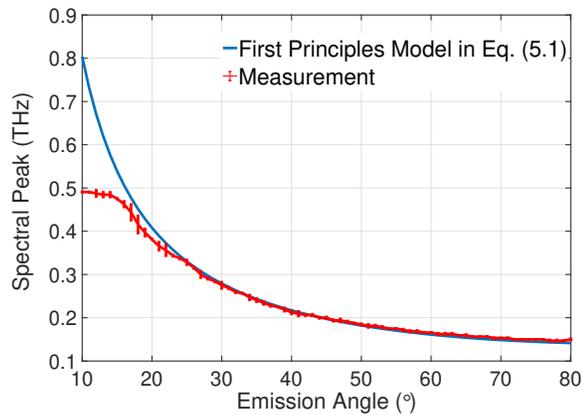


Figure 5.10: The LWA angle-frequency coupling, theory vs. experiment.

ment results deviate from theoretical behavior. The reason is that, for those lower angles, the corresponding peak frequencies are above 400 GHz and higher frequency components are fundamentally weaker, even in the spectrum of THz illumination source (see the spectrum of THz pulse in Fig. 5.7b). Hence, when exposed to path loss and coupling loss, the signal level degrades to the noise level increasing the measurement error. Also the model in Eq. (5.1) does not account for the realistic plate geometry. Namely, the non-idealities in the waveguide geometry such as the finite non-zero plate thickness contribute to a more complex spectrum, and these effects are likely to become more pronounced at small angles, where the effective propagation length inside the leaky-wave slot is larger.

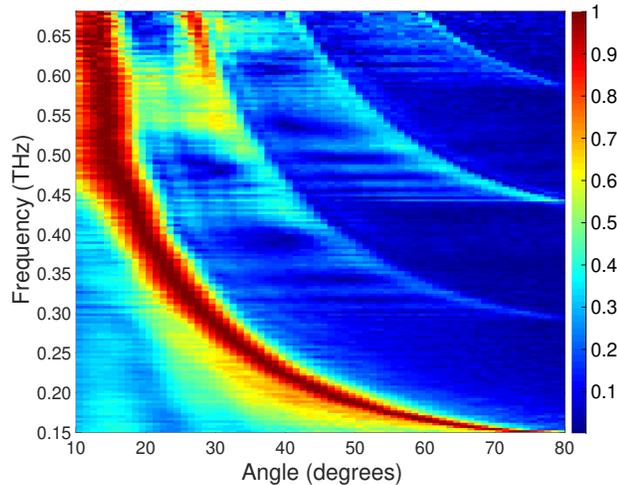


Figure 5.11: Spectrum-Angle heatmap of THz Rainbow.

Our results above show that the *peak frequency* (or the frequency that is maximally coupled) at a given angle matches the first principles model in Eq. (5.1). However, this model does not provide any information on other frequency components of the power spectrum. The overall spectral properties of waves composing a THz Rainbow is unexplored and no analytic solution exists even under idealistic geometry (i.e., zero plate thickness). In order to develop a phenomenological model, we directly measure the angular-dependent spectrum of a THz Rainbow, deploying the setup in Fig. 5.9.

Fig. 5.11 presents the heatmap of signal amplitude over different frequencies as a function of the LWA-receiver angle. For each frequency f , we normalize the corresponding amplitude over the entire angular range to the maximum value (i.e., normalizing every row of the heatmap matrix to the maximum value in that row). This heatmap includes the frequency range of 0.15 to 0.68 THz with the resolution of 3.1 GHz.

We observe that the high power region (with value one) follows Eq. (5.1). Interestingly, the heatmap shows high power associated with frequencies that are far from the frequency characterized by first principles model. In particular, at lower angles, the spectrum contains a wide range of high-power frequencies whereas at larger angles

mostly frequencies close to 0.15 THz hold relatively high power (i.e., the yellow bright region is much wider for lower angles). Further, Fig. 5.11 reveals similar second-order and third-order higher frequency echos that are caused by higher order transverse electric (TE) modes in the waveguide.

Finding: We experimentally demonstrated THz Rainbow and showed that the peak frequency is inversely related to the angle of emission via a non-linear function. However, the overall spectral profile is irregular such that lower frequency components tend to spread over a wider range of angles.

5.8.1.2 Validation of our Phenomenological Model

In Sec. 5.4, we modeled the THz Rainbow behavior as frequency-dependent spatial mask such that for a given frequency, the angular response is analogous to a spatial filter with a fall-off rate that itself depends on the frequency. To validate this model, we deploy the same setup as in Fig. 5.9 but this time we explore the change in the signal amplitude carried by frequency f as we deviate from the optimum angle. Even though our data set covers a wide range of frequencies, for the sake of space, we include results for few frequencies, namely, $f = 200, 250, 300,$ and 350 GHz.

Fig. 5.12 depicts normalized measured amplitude together with the physical model $G(\phi, f)$ in Eq. (5.2) and our phenomenological model $\tilde{G}(\phi, f)$ in Eq. (5.3). The normalization is computed such that the best angle (for each frequency) has a peak value of one. We observe good agreement between the experimental data and the developed phenomenological model, at least for an angular range of interest (angles whose corresponding measured power is within the 3 dB of maximum power). For larger angular offsets, irregularities can be seen in the measured power that is not captured by our phenomenological model. The rationale is that the signal strength is substantially weaker with larger angular offset; thus, even if a more complex model

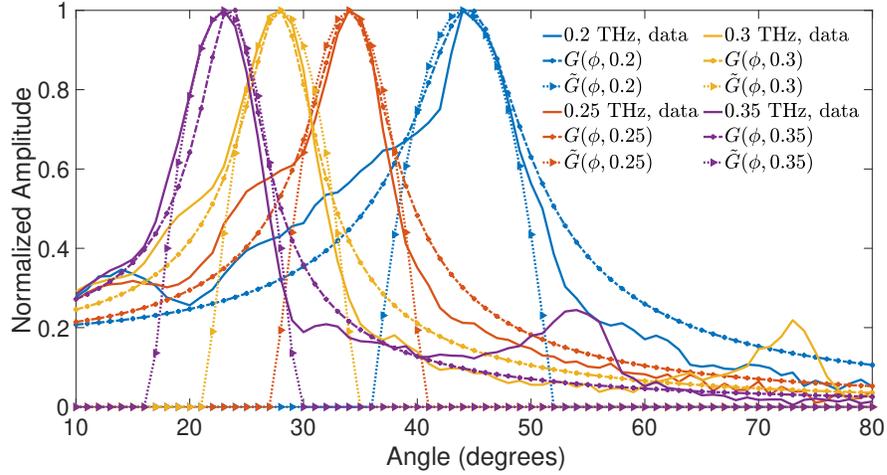


Figure 5.12: Validating the phenomenological model of spectral-spatial behavior in THz Rainbow.

is developed, it would not significantly improve the performance of OPERA.

Further, Fig. 5.12 confirms our approximation in Eq. (5.3) that the fall off rate (due to deviating from center angle on either side) is symmetric for the angular range of interest. Nonetheless, out of this range, we observe asymmetric behavior in a way that the fall off rate of the right tail is faster than the left tail. Finally, our experiments confirm that the fall-off rates are frequency-dependent such that the smaller the frequency, we observe less power drop given a fixed angular offset. We account for this dependency by including the frequency-dependent scalar $\gamma(f)$ in our expression in $\tilde{G}(\phi, f)$. Given a LWA device, $\gamma(f)$ is fixed and can be measured in advance.

Finding: We proposed and validated a model that characterizes the angular-dependent power spectral density of THz Rainbow. While the model does not capture the low-power irregularities at angles far from the optimum center angle, it accurately describes the power drop for a certain angular range of interest. OPERA takes advantage of such spatial-spectral signatures for path discovery.

5.8.2 OPERA Single-Path Discovery

We first evaluate the performance of OPERA in single-path settings (e.g., LOS or one reflected path only). For comparison purposes, we introduce a baseline scheme.

5.8.2.1 Baseline: Peak-Frequency Inversion

We introduce a baseline path discovery approach which we call Peak-Frequency Inversion. This scheme takes the spectral peak as signatures to find angle paths. The rationale is that the spectral peak has a one-to-one relation with the receiver's angle relative to the LWA. Besides, we expect paths to contribute to the overall spectrum based on their AoD, disregard of whether it is a LOS or a reflected path. In other words, we hypothesize that reflecting off a surface does not impact the spectral signature (in particular, the peak frequency) of the incident signal.

Fig. 5.13 describes the Peak-Frequency Inversion scheme consisting of a peak detection component and a frequency to angle mapping module. This scheme picks p frequencies, each corresponding to a path, from the received power spectrum and maps them to the corresponding angles. Our result in Sec. 5.8.1 revealed that the first principle model in Eq. (5.1) deviates from measurements for emission angles below 20 degrees. Hence, to eliminate this source of error, we exploit the measured angle-frequency curve in Peak-Frequency Inversion, as shown in Fig. 5.13.

The baseline scheme does not require any prior information other than the measured angle-frequency curve. Of course, one could even relax this constraint by substituting this measured curve with the Eq. (5.1), with the risk of higher estimation error for some path angles. Instead, OPERA assumes the spectral characteristics at different angles are known and exploits such signatures to find the path angles that best fits the entire power spectrum through a maximum likelihood optimization.

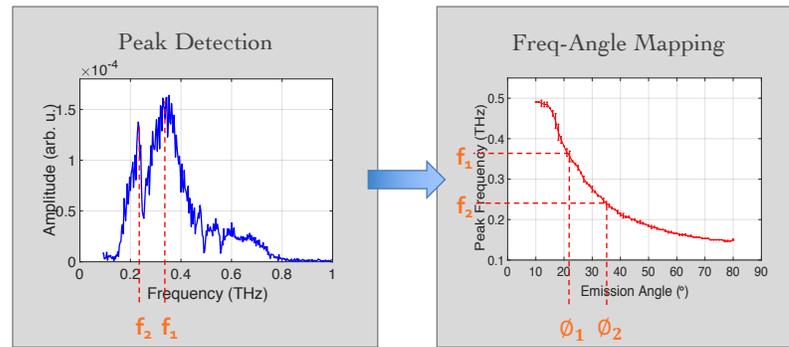
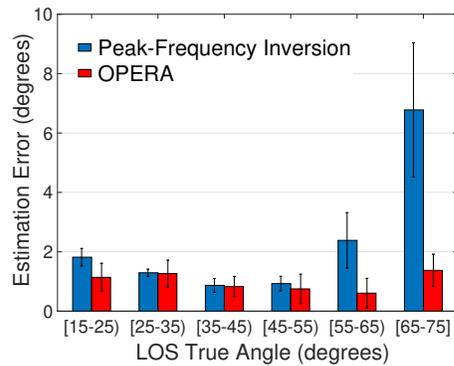


Figure 5.13: The baseline scheme that directly maps peaks in power spectrum to corresponding angles.

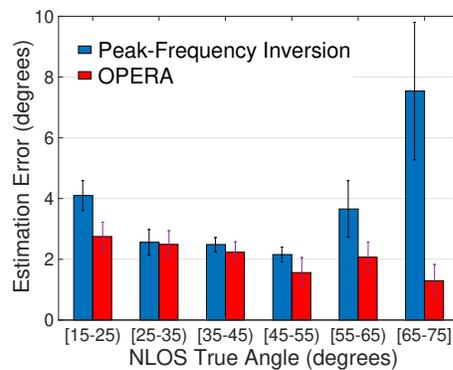
5.8.2.2 LOS vs. NLOS Detection

In this experiment, we compare the effectiveness of OPERA and Peak-Frequency Inversion in detection of a LOS path vs. a single reflected path. We consider various LOS settings with angles in the range of $15^\circ - 75^\circ$ and configure the setup carefully to ensure that no reflection exists. Each time, we excite the transmit LWA with a THz pulse and collect the signal received at the broadband detector. We repeat a similar experiment but this time we block the LOS path and carefully set up a metal reflector to add a reflected path. To compute the estimation error, we compare the predicted path direction with the geometric angle between the LWA and the center of the reflector.

Fig. 5.14 depicts the histograms of estimation error over different ground truth angular intervals for both LOS-only and NLOS-only experiments. In the LOS settings, i.e., Fig. 5.14a, we observe that the estimation error of Peak-Frequency Inversion is non-uniform across different angles and increases significantly for LOS angles above 55° . The reason is rooted in the non-linearity of frequency-angle relation in the LWAs. As shown in Fig. 5.10, the peak frequency is roughly a flat (or slowly varying) function of the emission angle for angles above 55 degrees; thus, a small measurement error in the spectral domain may cause significant error in angle estimates. Moreover,



(a) Single LOS path



(b) Single NLOS path

Figure 5.14: The histogram of estimation error in single-path settings.

the spectral resolution of any broadband receiver is finite and bound to its clock sampling rate. The frequency resolution in our setup is 3.1 GHz; thus, peak frequency may be mistakenly shifted by up to 3.1 GHz and yield to an estimation error of several degrees in this slow-varying regime. In contrast, OPERA takes into account the spectral-angular profile of THz Rainbow and the relative power across different frequencies for path detection. Thus, it is able to distinguish between path angles with similar corresponding spectral peaks and consistently achieves less than 2° error on average across the angular space.

Fig. 5.14b reveals two key differences in NLOS detection: First, for both OPERA and baseline, the average estimation error is slightly higher in all angular intervals.

Second, the Peak-Frequency Inversion approach obtains larger error at the two ends of the angular range, i.e., [15-25] and [65-75]. The reason is two-fold: given the same AoD, a reflected path endures higher loss due to the reflection loss and the increased path length. Hence, the NLOS path is weaker and prone to higher estimation error. Moreover, reflection imposes non-equal attenuation over different frequencies leading to spectral peak shift or generally adds noise to the spectral signature of the path. Higher frequency components are more exposed to these errors comparing to the lower frequencies because they endure higher reflection loss and path loss. Hence, from Fig. 5.14b, we observe that the baseline scheme obtains high errors for lower angles in the range of 15 – 25 that contain higher frequency components. OPERA is not immune to the non-perfect reflection loss either; however, by considering the relative power over the entire power spectrum, it can better diminish the impact of such errors.

Finding: Even though the baseline scheme takes advantage of the one-to-one relation between spectral peak and emission angle, we find that it encounters critical limitations in practice. First, for a range of angles, the peak frequency is a slow-varying function of the emission angle makes it susceptible to measurement error and the limited frequency resolution at the receiver. Second, reflection non-uniformly impacts power over different frequencies which can lead to a spectral peak shift and subsequent estimation inaccuracy.

5.8.2.3 Single-Path Estimation Accuracy

Next, we explore the overall performance of OPERA vs. Peak-Frequency Inversion in single-path extraction (LOS-only and NLOS-only combined). We choose 225 different configurations in which the ground truth AoD varies between 10 to 80 degrees. Fig. 5.15 presents the empirical distribution function of the estimation error, which is

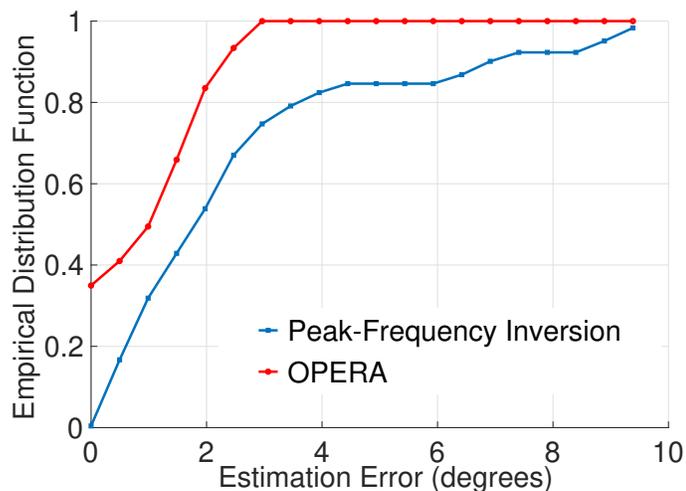


Figure 5.15: The estimation accuracy of OPERA in single-path settings.

the number of observations less than or equal to a specific error divided by the total 225 independent measurement readings.

The result reveals that OPERA successfully estimates the AoD within 2.2 degrees of the ground truth in 90% of cases. In addition, a simple frequency to angle mapping as in baseline achieves less than 3.8 degrees error in 80% of cases. Nonetheless, its curve slowly approaches one; in particular, in 18% of single-path settings the estimation error was more than 6 degrees and in 10% of instances the error was higher than 8 degrees. From the above analysis, we know that these high-error cases correspond to larger LOS/NLOS angles.

Finding: In single-path channels, Peak-Frequency Inversion achieves on average a comparable performance with OPERA. Yet, its accuracy is not robust across the entire angular domain.

5.8.3 OPERA Multipath Extraction

In multi-path channels, the receiver captures a superposition of signals traveling along different paths. In OPERA, by design, the spectral properties of such signals are

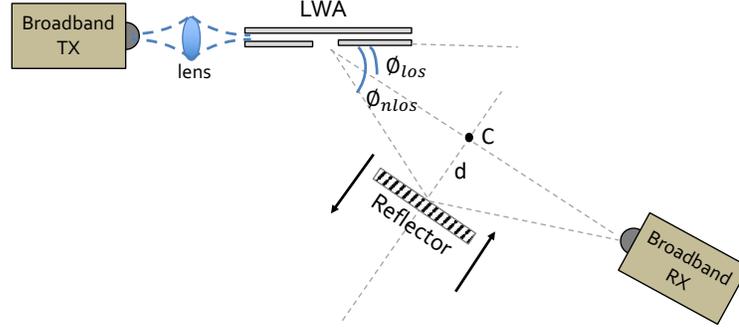


Figure 5.16: The schematic of our multipath experimental setup.

unique when emitted from the LWA. However, the key question is whether or not such spatial-spectral signatures are reliably extractable from the superposed signal in multipath settings.

Setup. We consider two-path configurations, with one LOS and one reflected path. The schematic of our multipath experimental setup is shown in Fig. 5.16. We set an aluminum sheet as reflector and slide it on a rail to vary the NLOS path angle. This rail is attached to another fixed rail at cross point C such that the distance from LWA to C is equal to RX to point C distance. Furthermore, the reflector is oriented in parallel to the rail that connects LWA to the RX. These considerations ensure that the portion of the emitted THz Rainbow that reflects off the reflector can be collected by the broadband RX. Moving the reflector from point C to the left increases the NLOS angle ($\phi_{nlos} > \phi_{los}$) and vice versa.

5.8.3.1 Minimum Inter-Path Angular Separation

OPERA identifies different paths via their angular-dependent frequency signature. Thus, when two paths have close AoD, their spectral profile might be indistinguishable. Hence, one key evaluation factor is the minimum required angular separation between paths so that we can successfully extract their signatures from the superposed received signal.

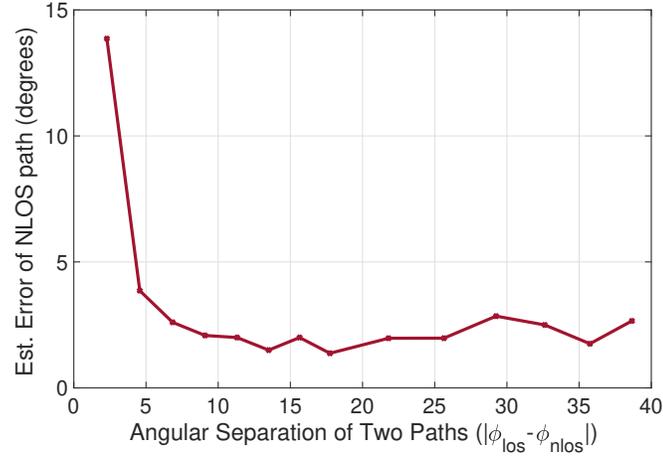


Figure 5.17: The estimation error of the reflected path as a function of inter-path angular separation.

To answer this question, we consider a fixed LOS configuration with path angle ϕ_{los} and move the reflector along the rail in the steps of 5 cm to vary the angular separation between the LOS and reflected path (i.e., $|\phi_{los} - \phi_{nlos}|$). We repeat this experiment for a discrete choice of LOS angles (i.e., $\phi_{los}=15$ to 75 in steps of 5 degrees). At each multipath setting, we collect raw data from the broadband receiver and implement OPERA path discovery.

We find that LOS detection is robust to inter-path angular separation. The reason is that the LOS path is stronger and its spectral signature dominates the characteristics of the received signal. However, Fig. 5.17 reveals this does not hold true for the reflected path. This figure shows the estimation error associated with detection of the NLOS path vs. the inter-path angular separation. We observe that when paths are not sufficiently separated (e.g., $|\phi_{los} - \phi_{nlos}| < 3^\circ$), the estimation error is significantly high. Since paths have similar spectral signatures in such cases, removing the contribution of first path from the superposed signal disrupts the detection of the second path.

Finding: When multiple paths are sufficiently angularly separated, i.e., $|\phi_{los} - \phi_{nlos}| > 3^\circ$, OPERA can successfully extract their spatial-spectral signatures.

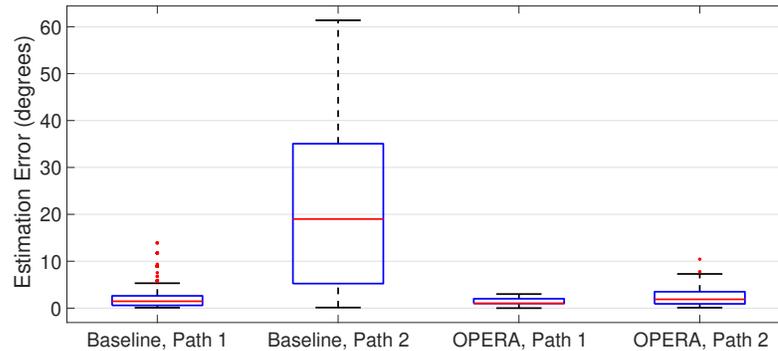


Figure 5.18: The box plot of estimation error for more than 300 independent multipath configurations.

5.8.3.2 Overall Multi-Path Estimation Accuracy

Finally, we analyze the overall performance of OPERA compared to the Peak-Frequency Inversion approach in multipath discovery. Hence, we place the receiver in different LOS angles relative to the transmit LWA (from 15 to 75 in the steps of 5 degrees) and configure the reflector in multiple locations along the rail (in the steps of 5 cm) to create a range of NLOS angles. Fig. 5.18 shows the box plot of the estimation error for more than 300 independent multipath measurements.

The Peak-Frequency Inversion scheme offers a relatively comparable performance on average in the discovery of the LOS path; nonetheless, its accuracy is non-uniform across the entire angular domain leading to higher variance. More importantly, Peak-Frequency Inversion demonstrates unreliable performance in the extraction of the second path. The reason is rooted in the irregularities in the spatial-spectral signatures of THz Rainbow. In other words, if such signatures were indeed a sharp delta function (i.e., each frequency was emitted from one and only one direction), then Peak-Frequency Inversion performance would have been enhanced as the spectrum of different paths would have been completely isolated and non-interfering. Instead, each path contributes to the overall spectrum with its wide and irregular spectral sig-

nature. Further, such contributions are not equally strong. In particular, the spectral peak of a secondary path might get dissolved in the irregularities of a stronger LOS spectral signature. In contrast, OPERA exploits the fact that THz Rainbow spectral profile is fixed and can be known a priori; therefore, the receiver correlates its signal with the known but irregular transmission signatures; thus, it is able to detect the second path direction with only 1.9° error on average.

Finding: Irregularities in the THz Rainbow profile together with the non-equal path strengths cause inaccurate path estimation by Peak-Frequency Inversion in multipath settings. Instead, OPERA exploits the known but irregular spectral signatures to successfully extract multiple paths.

5.8.4 OPERA for Beam Adaptation

Here, we evaluate the effectiveness of OPERA for beam adaptation. In particular, we explore two main factors: (i) The SNR penalty associated with non-ideal path estimates. This enables us to understand the performance penalty (ultimately in data rate) due to imperfect beam alignment. (ii) The rate at which path updates are provided, as stale path data can also degrade performance for mobile clients.

5.8.4.1 SNR Performance.

As discussed above, OPERA provide path directions within 2° of the ground truth. This non-zero error results in sub-optimal beam steering and an SNR penalty, consequently. We compare the SNR of OPERA beam adaptation with the maximum achievable SNR via an Oracle approach that has perfect knowledge of paths. Note that here, we assume the path tracking rate (the rate at which path information is updated) is exactly the same for Oracle and OPERA; hence, any difference in SNR is solely due to path estimation errors. Moreover, we expect the SNR difference to

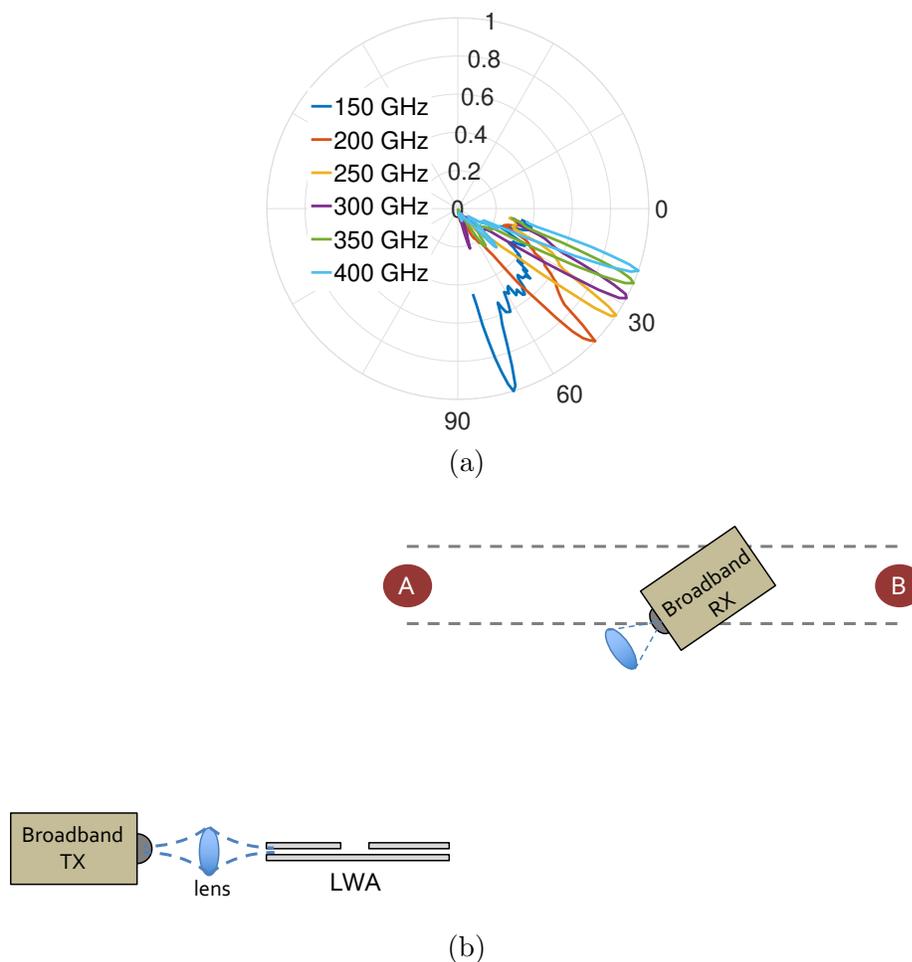


Figure 5.19: (a) Example directional radiation patterns generated by LWA; (b) Our mobility setup.

depend on the radiation patterns of directional beams in use, e.g., narrow beams are more sensitive to path estimation errors. Here, to have a fair comparison, we obtain a set of fixed beam patterns for both Oracle and OPERA schemes. In particular, we generate directional beams with a LWA by tuning the carrier frequency of the input signal. Example radiation patterns are shown in Fig. 5.19a.

Setup. We deploy the experimental setup in Fig. 5.19b, such that the transmitter is fixed and is equipped with a LWA. The broadband receiver is moving on a line from A to B. To emulate continuous motion, we collect raw data at the RX with high granularity of every 0.5 cm movement. Further, we repeat the same experiment in



Figure 5.20: The SNR performance of beam adaptation when path directions are provided by OPERA.

multiple parallel motion lanes with different distance to the TX. To isolate sender-side beam steering, we manually align the broadband RX as it moves along the translational line.

Fig. 5.20 plots the empirical distribution function of SNR loss in OPERA compared to the maximum achievable SNR via the Oracle Scheme. We observe the SNR loss of less than 2 dB in 90% of measurement instances. The non-zero error in AoD estimates yields to generation of a radiation pattern whose main-lobe is slightly misaligned from the ground truth angle and imposes an average SNR penalty of 0.3dB.

Finding: The non-zero error in path angle estimates provided by OPERA causes slightly misaligned beam steering, imposing an SNR penalty of 0.3 dB on average, comparing to the maximum achievable SNR via an Oracle scheme.

5.8.4.2 Path Tracking Rate.

A key factor affecting beam adaptation performance is the rate at which path angle estimates are available. In Sec. 5.5, we demonstrated that OPERA includes the periodic transmissions of THz Rainbow in time intervals of T_{rb} . Such broadband transmissions occupy the entire spectral band yet only for a nano-second scale period of time. Data transmissions in all bands need be stopped or will otherwise interfere

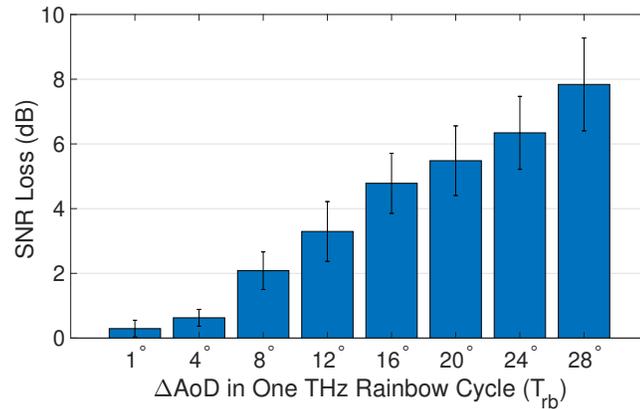


Figure 5.21: The impact of OPERA’s path tracking rate on the accuracy of beam adaptation.

with path discovery and vice versa. Moreover, the translational speed of the client also affects steering accuracy; the faster the speed, the harder the tracking since the client’s angular position may change more for the same estimation interval T_{rb} . Instead of evaluating these two factors separately, we consider the amount of change in client position relative to the transmitter during one cycle of OPERA, i.e., T_{rb} .

Fig. 5.21 presents SNR loss for different client movements in T_{rb} . The AoD of the LOS path varies as the client moves along the translational line. We use ΔAoD (between two consecutive THz Rainbows) as a metric for how fast the client moves since changes in the AoD indicate whether the transmit beam needs to be adapted or not. A larger number on the x-axis means that either T_{rb} is larger (for the same client speed) or the client is moving faster (for the same T_{rb}).

We observe that when path tracking is repeated most frequently (every 1° of client angular movement), the SNR loss (compared to the maximum achieved SNR via the Oracle scheme) is the lowest. Since the path angle estimates are being updated faster, we can adapt beams before client moves out of the AP’s steering direction and experiences a significant SNR loss. We observe that in such a case, the average SNR loss is only 0.29 dB. Furthermore, as the estimation interval increases, the average

SNR loss also becomes large, since there is a great change in AoD between two path discovery measurements.

As discussed in Sec. 5.5, THz Rainbow transmission incurs nanosecond-scale overhead whereas mobility is in millisecond time scales. For instance, consider a mobile client who is moving with speed of 10 m/s on a translation line with initial distance of 10 m from the transmit LWA. Then, the most frequent path tracking (i.e., every 1° of client movement) requires T_{rb} to be 17 msec meaning path directions need to be updated in millisecond-scale.

Finding: When the THz Rainbow is generated more frequently, the average SNR loss at a mobile client is smaller as beams can be better adapted based on the latest path directions. Yet, even under intensive path tracking (e.g., for every 1° change of client movement), the time overhead of updating paths is orders of magnitude smaller than the mobility scale.

5.9 Experimental Results on Rotation Tracking

In this section, we experimentally validate our rotation estimation technique.

5.9.1 Comparing Ray Optics vs. Diffraction Model

We introduced two models to understand the broader spectral band at each emission angle, one based on the diffraction and the other based on ray optics. In Fig. 5.22, we compare them against the measured heatmap of power spectrum.

Fig. 5.22a presents our measurement using a broadband detector staring directly at the emission point, without a second LWA. Each row of this image has been normalized to unity magnitude, in order to remove the frequency-dependence of the input signal from the THz-TDS transmitter, and emphasize the signals at higher

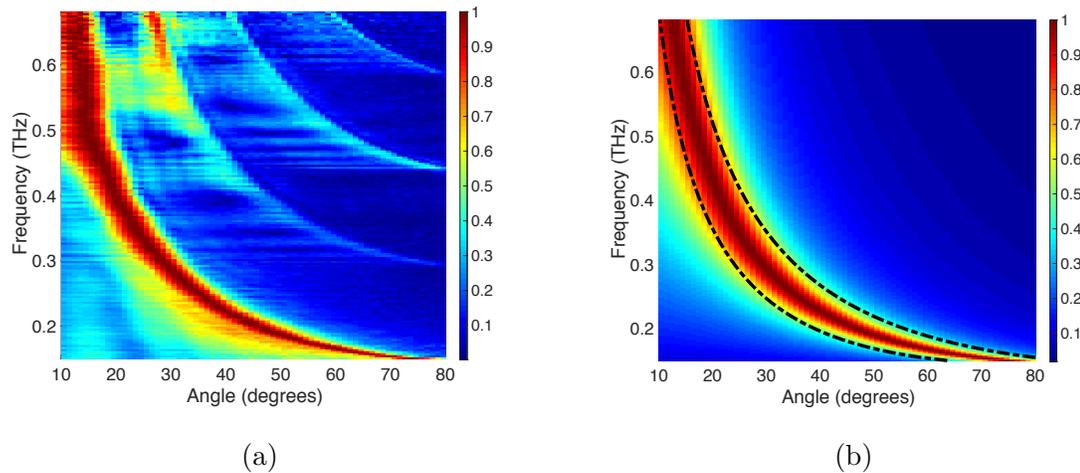


Figure 5.22: The spectrum of emitted radiation vs. emission angle. a) A plot of the measured spectrum of the radiation emitted by a LWG, after excitation with a broadband input; b) Two different predictions of the spectrum-angle relation by diffraction model and ray optics.

frequency. The prominent arc in the lower left region corresponds to the emission from the dominant TE_1 waveguide mode; the weaker arcs in the upper right arise from higher-order TE waveguide modes (TE_2 , TE_3 , and TE_4), which result from imperfect input coupling to the waveguide. The TE_1 mode signal represents about 90% of the total radiated energy.

In Fig. 5.22b, the false color plot corresponds to a diffraction model (Eq. (5.2)) using a value of $\alpha = 0.213 \text{ mm}^{-1}$ for the leakage parameter, while the two dotted black curves represent the results of a ray optics model (Eq. (5.9)).

The ray-optics description which relies on knowledge of the (finite) length of the slot apertures in the two leaky-wave devices (both the TX and the RX have the same slot length L). In contrast, the underlying description of Eq. (5.2) involves diffraction through an aperture of finite length, and thus relies only on the length of the TX aperture, not the RX aperture. In this sense, the two models are not equivalent in the general case; in order to render them equivalent, one would need to account for the finite aperture of the RX in the diffraction formalism. This could be accomplished,

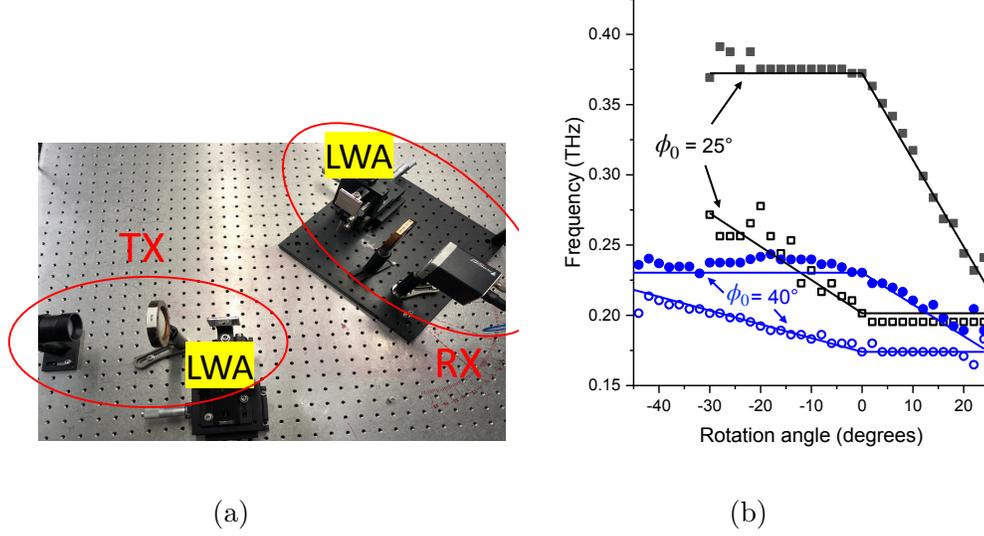


Figure 5.23: (a) Rotation experimental setup; (b) Characterization of client rotation and its impact on the received spectral edges for two different values of $\phi_0 = 25^\circ, 40^\circ$.

for example, by convolving the diffraction pattern of Eq. (5.2) with an aperture of finite size which subtends a finite angular range with respect to the TX waveguide; this would have the effect of broadening the spectrum that would be predicted at any given angle, relative to what is shown in Fig. 5.22b. However, in our case, the effects of this broadening would be minimal, since the TX-RX distance ($R = 15$ cm) is significantly larger than the slot length ($L = 3$ cm) (so therefore $\alpha L \approx 3.7 > 1$). This claim is substantiated by the good agreement between the measured spectra (Fig. 5.22b) and the spectra computed directly from the model.

5.9.2 Impact of Rotation on the Spectral Edges

To experimentally characterize the impact of rotation, we consider the setup in Fig. 5.23a in which the receiver is located at angle ϕ_0 with respect to the transmitter. We rotate the RX plane (including the LWA, the lens that focuses the waves into the broadband detector, and the broadband detector itself) by θ_{rot} degrees and measure the received power spectrum.

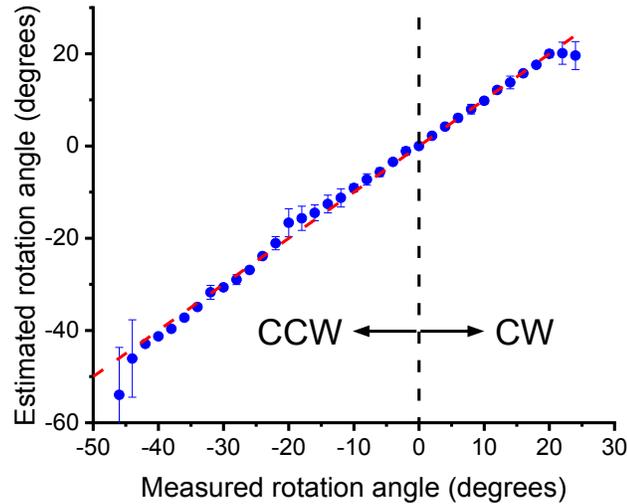


Figure 5.24: Evaluation of the accuracy of proposed rotation estimation technique over a range of CW and CCW measurements.

Fig. 5.23b presents the extracted values of f_{max} and f_{min} as a function of client rotation angle, for two different values of $\phi_0 = 25^\circ$ and $\phi_0 = 40^\circ$. The solid lines represent the predicted values based on ray optics. These results indicate the typical level of agreement between measurement and prediction, over the range of values of θ_{rot} where a non-zero spectral width is predicted by the theory.

5.9.3 Accuracy and Range of Rotation Tracking

By compiling all measurements at a given rotation angle (at each ϕ_0), we extract the measurement uncertainty in the rotation angle as a function of the degree of rotation. Fig. 5.24 depicts the estimated rotation angle against the measurement (ground truth rotation angle) for a range of CW and CCW rotation angles. We observe an overall good agreement; however, the uncertainty slightly increases for larger rotations, since a smaller signal is measured for larger rotations. Nevertheless, over the range of accessible rotation angles, an average estimation error of less than 2° is obtained

Another important evaluation question is the maximum range of rotation that

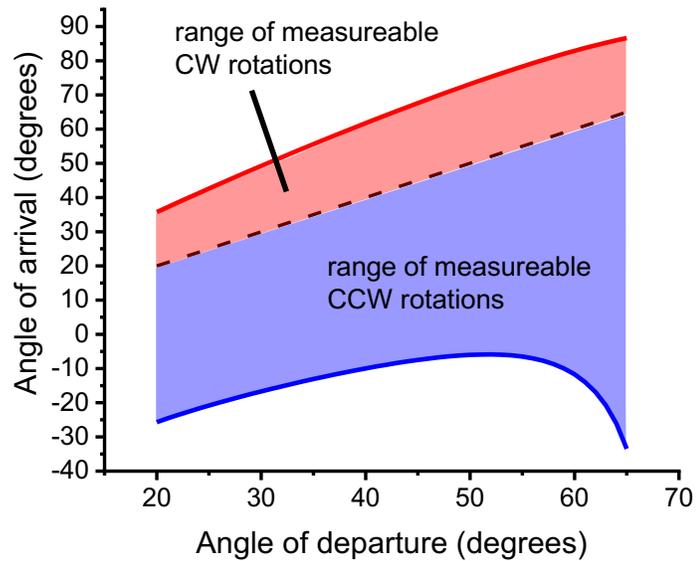


Figure 5.25: Maximum range of detectable rotation angles.

we can detect and track. Obviously, if the rotation angle θ_{rot} is too large, then no spectral information is received, and the rotation angle cannot be determined. However, Fig. 5.25 shows that for a surprisingly large range of angles, rotation can be accurately tracked using a single-shot measurement. This is mainly due to the wide spectral band emitted at each angle. Furthermore, Fig. 5.25 reveals that depending on the LOS angle ϕ_0 , the RX has different resilience toward CW and CCW rotation. Finally, the max range of measurable rotations is not symmetric with respect to the direction of rotation. This can be justified by our model in Eq. (5.10) and Eq. (5.11) by knowing the the rate of spectral change is higher for the high-frequency edge compared to the low-frequency edge for any given initial angle ϕ_0 (i.e., $\frac{\partial f_{max}}{\partial \theta} > \frac{\partial f_{min}}{\partial \theta}$).

5.10 Related Work

Leaky-Wave Antennas. LWAs have a long history, having been used in the RF region since at least the 1940's [97]. Most early examples were rectangular waveguides; however, in the THz range, the parallel-plate waveguide has been shown to be a very useful low-loss and low-dispersion platform for many purposes [93, 104, 105]. While LWAs facilitate beam steering toward a specific direction, discovering the right beam direction at the TX and RX remains a challenge in the literature. This work is the first-ever use of LWA devices as control plane to adapt THz WLANs and the first to use the LWA to transmit in all directions simultaneously.

Path Discovery. Prior work focuses on AoA estimation by exploiting large antenna arrays and taking advantage of advanced signal processing techniques [88, 89]. In particular, such approaches employ the phase difference at multiple antennas to infer the direction of wave incident. However, this method is not applicable in THz bands as even the state of the art THz node architecture cannot provide the amplitude and phase information corresponding to each antenna element in an antenna array.

More recent work proposed an alternative approach to direction finding using a time-modulated array. In this technique, by periodic on-off switching of the array elements, harmonic components are generated and analyzed to estimate the arrival direction. Compared with the aforementioned array-based AoA estimation techniques, the time-modulated array offers high precision, while its hardware complexity and calculation complexity are relatively low [106, 107]. However, the application of time-modulated arrays in high-frequency regimes is limited by the speed of RF switches. Further, this approach requires periodic transmission of the signal under different switching states and thus it is clearly not a single-shot nor a single-antenna approach.

Another body of work uses visible light sensing to estimate the AoA of the LOS path between the light source and a photodiode array at the client [108, 109]; how-

ever, such approaches are limited to the LOS path only. Authors in [110] study non-coherent path tracking in mmWave but their approach requires multiple beacon transmissions. Path angle estimation for THz communication has been studied recently in [111], in which the authors utilize a low frequency RF frontend to roughly estimate the AoA in a first step and refine it in THz regime in a second step. In contrast, our work proposes a single-shot and single antenna path discovery scheme that can identify all dominant paths between the AP and all clients.

Finally, a recent work in 60 GHz networks exploits the irregularity and side-lobes in 60 GHz beam patterns together with power delay profile measurements for path identification [4]. In contrast to practical antenna arrays in 60 GHz with wide and irregular beam patterns, THz beams are regularly shaped and quite narrow [15].

Directional Neighbor Discovery. Several papers addressed the problem of neighbor discovery for wireless networks that use directional antennas [112–114]. Their main objective is to discover the neighbors, i.e., nodes that are within communication range, around any particular node and store the neighborhood information. In particular, the challenge is to quickly steer the listening beam direction so that it can decode any of the multiple beacons transmitted directionally in a multi-shot discovery procedure. Instead, OPERA attempts to characterize all dominant paths between two nodes via one-shot transmission. Nonetheless, OPERA can be employed for directional neighbor discovery as well. To this end, the node aiming to identify its neighbors has to transmit a THz Rainbow. When OPERA identifies a path at a particular receiving node, one can conclude that those nodes are neighbors.

mmWave Beam Adaptation. There is extensive work to reduce overhead of beam acquisition and adaptation in mmWave networks. Beam acquisition solutions include IEEE 802.11ad/ay’s sector sweeping [6, 14], hierarchical codebook designs [115, 116], compressive sensing techniques to exploit channel sparsity [16, 117]

and model-driven beam steering via channel profiling [43, 44]. For beam adaptation under mobility prior work suggests a variety of in-band and out-of-band solutions including the use of legacy bands [46, 79] or visible light [108], and efficient beam searching [118, 119]. These efforts reduce training overhead and maintain alignment for mmWave links. However, given that THz links lack the “pseudo-omni” reception capability available at mmWave and are inherently more directional, these solutions are not applicable in THz bands. In contrast, this work provides path directions and enables efficient beam steering and beam adaptation in THz Networks.

Concluding Remarks

This thesis presents the design, implementation, and experimental evaluation of techniques for sensing-enhanced wireless directional networking in the emerging mmWave and THz spectrum. It introduces novel systems that leverage the wide available spectral band, large antenna arrays, unexplored high-frequency devices, unique propagation features of mmWave/THz waves to extend sensing with communication-dedicated hardware and signals. Specifically, I focus on sensing the features of nodes and the environment that enables efficient and intelligent coordination and adaptation of directional and mobile wireless networks. The contributions of this thesis include software-defined radio development, network control-plane design, physics of new high-frequency antennas, novel protocols, and algorithms.

This thesis demonstrates, for the first time, a sensing-assisted spatial multiplexing system for the efficient establishment of multi-stream directional links in multi-user mmWave MIMO networks. I showed how to fuse the imperfections of the phased-array-generated beams with sparse scattering to sense spatial paths across different users and analog beams. Subsequently, I proposed algorithms that exploit the discovered paths and pre-knowledge of radiation patterns in the RF codebook to mitigate inter-stream interference and undesirable channel correlations by intelligently choos-

ing beams that capture diverse or ideally orthogonal paths.

Then, I introduced the first single-shot single-antenna motion-sensing system in THz wireless networks. This system allows for ultra-fast detection and tracking of user moves (translational and rotation) with the transmission of wideband frequencies (i.e., a time-domain terahertz pulse) using an unexplored device named a leaky-wave antenna. I demonstrated how different frequencies of the terahertz pulse would simultaneously emerge at different angles, forming a pattern which I termed a THz Rainbow. Understanding the underlying properties of emitted waves as a function of spatial angle, I proposed techniques to extract user motions from the spectral characteristics of the received signal. I then leveraged this information to proactively re-align highly-directional beams under mobility or for fast blockage recovery. Our design requires a single-shot transmission of a terahertz pulse and a single leaky antenna at wireless node architecture. Combined, these innovations address the key challenges of directional networking, i.e., mobility, blockage, and scalability, in unprecedented ways.

This research builds a foundation for joint wireless communication and sensing in mmWave and THz spectrum yielding a paradigm shift in the design of next-generation wireless nodes, protocols, and standards. We envision that future wireless systems will exploit the wide and diverse spectral bands, novel materials, and the possibility of directionality in small form factors to achieve Tb/sec data rates jointly with ubiquitous high-resolution sensing.

References

- [1] S. K. Saha, Y. Ghasempour, M. K. Haider, T. Siddiqui, P. D. Melo, N. Somanchi, L. Zakrajsek, A. Singh, O. Torres, D. Uvaydov, J. M. Jornet, E. Knightly, D. Koutsonikolas, D. Pados, and Z. Sun, “X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays,” in *Proc. ACM WiNTECH*, 2017. (document), 1.1, 1.3, 3.4.1, 4, 4.2.1, 4.7, 4.3
- [2] Y. Ghasempour and E. W. Knightly, “Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs,” in *Proc. ACM MobiHoc*, 2017. (document), 1.3, 2.4, 4.1, 4.6
- [3] Y. Ghasempour, M. K. Haider, and E. W. Knightly, “Decoupling Beam Steering and User Selection for MU-MIMO 60 GHz WLANs,” *IEEE/ACM Transactions on Networking*, vol. 26, no. 5, pp. 2390–2403, 2018. (document), 1.3
- [4] Y. Ghasempour, M. K. Haider, C. Cordeiro, D. Koutsonikolas, and E. Knightly, “Multi-Stream Beam-Training for mmWave MIMO Networks,” in *Proc. ACM MobiCom*, 2018. (document), 1.3, 3.8, 5.10
- [5] Y. Ghasempour, M. K. Haider, C. Cordeiro, and E. W. Knightly, “Multi-User Multi-Stream mmWave WLANs With Efficient Path Discovery and Beam Steering,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 12, pp. 2744–2758, 2019. (document), 1.1, 1.3
- [6] Y. Ghasempour, C. da Silva, C. Cordeiro, and E. W. Knightly, “IEEE 802.11ay: Next-generation 60 GHz Communication for 100 Gbps Wi-Fi,” *IEEE Communications Magazine*, vol. 55, no. 12, pp. 186–192, 2017. (document), 1, 1.1, 3.1.1, 4.1, 4.2.2, 5.10
- [7] Y. Ghasempour, C.-Y. Yeh, R. Shrestha, D. Mittleman, and E. Knightly, “Single Shot Single Antenna Path Discovery in THz Networks,” in *Proc. ACM MobiCom*, 2020. (document), 1.3

-
- [8] Y. Ghasempour, R. Shrestha, A. Charous, E. Knightly, and D. M. Mittleman, "Single-Shot Link Discovery for Terahertz Wireless Networks," *Nature Communication*, vol. 11, no. 1, p. 2017, 2020. (document), 1.3
- [9] C. V. N. Index, "Global Mobile Data Traffic Forecast Update, 2015–2020," *Cisco white paper*, p. 9, 2016. 1
- [10] 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. 1
- [11] J. Zhang, P. Tang, L. Tian, Z. Hu, T. Wang, and H. Wang, "6–100 GHz Research Progress and Challenges from a Channel Perspective for Fifth Generation (5G) and Future Wireless Communication," *Science China Information Sciences*, vol. 60, no. 8, 2017. 1
- [12] M. Agiwal, A. Roy, and N. Saxena, "Next Generation 5G Wireless Networks: A Comprehensive Survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 3, pp. 1617–1655, 2016. 1
- [13] I. . W. Group, "IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," 2012. 1.1, 3, 3.1.1, 3.1.4, 3.2.2.1, 3.6.2, 4.2.2, 4.2.3.1
- [14] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. Widmer, "IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-second Wi-Fi," *IEEE Communications Magazine*, 2014. 1.1, 4.1, 4.1.2.1, 4.1.3, 5.5.1, 5.10
- [15] J. Federici and L. Moeller, "Review of Terahertz and Subterahertz Wireless Communications," *Journal of Applied Physics*, vol. 107, no. 11, p. 6, 2010. 1.1, 5.10
- [16] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices," in *Proc. ACM CoNEXT*, 2017. 1.1, 4, 4.2.1, 4.2.2, 4.3, 5.10
- [17] T. Bai and R. W. Heath, "Coverage Analysis for Millimeter Wave Cellular Networks with Blockage Effects," in *IEEE Global Conference on Signal and Information Processing*, 2013, pp. 727–730. 1.1
- [18] G. R. MacCartney, T. S. Rappaport, and S. Rangan, "Rapid Fading Due to Human Blockage in Pedestrian Crowds at 5G Millimeter-Wave Frequencies," in *Proc. IEEE GLOBECOM*, 2017, pp. 1–7. 1.1

-
- [19] V. Petrov, M. Komarov, D. Moltchanov, J. M. Jornet, and Y. Koucheryavy, “Interference and SINR in Millimeter Wave and Terahertz Communication Systems with Blocking and Directional Antennas,” *IEEE Transactions on Wireless Communications*, vol. 16, no. 3, pp. 1791–1808, 2017. 1.1
- [20] T. Bai and R. W. Heath, “Analysis of Self-body Blocking Effects in Millimeter Wave Cellular Networks,” in *Asilomar Conference on Signals, Systems and Computers*, 2014, pp. 1921–1925. 1.1
- [21] I. A. Ibraheem, N. Krumbholz, D. Mittleman, and M. Koch, “Low-Dispersive Dielectric Mirrors for Future Wireless Terahertz Communication Systems,” *IEEE Microwave and Wireless Components Letters*, vol. 18, no. 1, pp. 67–69, 2008. 1.1
- [22] S. Rajagopal, S. Abu-Surra, and M. Malmirchegini, “Channel Feasibility for Outdoor Non-Line-of-Sight mmWave Mobile Communication,” in *Proc. IEEE VTC*, 2012, pp. 1–6. 1.1
- [23] J. Ahmadi-Shokouh, S. Noghianian, E. Hossain, M. Ostadrahimi, and J. Dietrich, “Reflection Coefficient Measurement for House Flooring Materials at 57-64 GHz,” in *Proc. IEEE GLOBECOM*, 2009, pp. 1–6. 1.1
- [24] S. Jog, J. Wang, J. Guan, T. Moon, H. Hassanieh, and R. R. Choudhury, “Many-to-Many Beam Alignment in Millimeter Wave Networks,” in *Proc. USENIX NSDI*, 2019. 1.1
- [25] M. Cetin, I. Stojanovic, O. Onhon, K. Varshney, S. Samadi, W. C. Karl, and A. S. Willsky, “Sparsity-driven Synthetic Aperture Radar Imaging: Reconstruction, Autofocusing, Moving Targets, and Compressed Sensing,” *IEEE Signal Processing Magazine*, vol. 31, no. 4, pp. 27–40, 2014. 1.2
- [26] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller, “3D Tracking via Body Radio Reflections,” in *Proc. USENIX NSDI*, 2014. 1.2
- [27] W. Wang, A. X. Liu, and K. Sun, “Device-free Gesture Tracking using Acoustic Signals,” in *Proc. ACM MobiCom*, 2016, pp. 82–94. 1.2
- [28] D. Murray and A. Basu, “Motion Tracking with an Active Camera,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 5, pp. 449–459, 1994. 1.2
- [29] 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, pp. 77–88, 2019. 1.3

-
- [30] N. Instruments, “Introduction to the NI mmWave Transceiver System Hardware - National Instruments,” <http://www.ni.com/white-paper/53095/en/>, 2017, (Accessed on 06/25/2017). 2
- [31] “Comsol multiphysics.” [Online]. Available: <https://www.comsol.com> 2.1.3
- [32] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, “Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation,” in *Proc. ACM CoNEXT*, December 2015. 2.1.3, 2.4
- [33] “wil6210: Linux wireless.” [Online]. Available: <https://github.com/torvalds/linux/blob/master/drivers/net/wireless/ath/wil6210/wmi.h> 2.1.3
- [34] H. Xu, V. Kukshya, and T. S. Rappaport, “Spatial and Temporal Characteristics of 60-GHz Indoor Channels,” *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 3, pp. 620–630, 2002. 2.2.2, 2.3.2, 2.3.4, 2.4, 4, 4.2.1, 4.2.3.1
- [35] G. R. Maccartney, T. S. Rappaport, S. Sun, and S. Deng, “Indoor Office Wideband Millimeter-Wave Propagation Measurements and Channel Models at 28 and 73 GHz for Ultra-Dense 5G Wireless Networks,” *IEEE Access*, 2015. 2.2.2, 3.4.2, 3.8
- [36] M. K. Haider and E. W. Knightly, “Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation,” in *Proc. ACM MobiHoc*, 2016. 2.3.5, 2.4, 3.4.2, 4.2.2, 4.6
- [37] P. F. M. Smulders and L. M. Correia, “Characterisation of Propagation in 60 GHz Radio Channels,” *Electronics & Communication Engineering Journal*, vol. 9, no. 2, pp. 73–80, 1997. 2.4
- [38] C. R. Anderson and T. S. Rappaport, “In-building Wideband Partition Loss Measurements at 2.5 and 60 GHz,” *IEEE Transactions on Wireless Communications*, vol. 3, no. 3, pp. 922–928, 2004. 2.4
- [39] P. F. M. Smulders, “Statistical Characterization of 60-GHz Indoor Radio Channels,” *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 10, pp. 2820–2829, 2009. 2.4
- [40] A. Maltsev, R. Maslennikov, A. Sevastyanov, A. Khoryaev, and A. Lomayev, “Experimental Investigations of 60 GHz WLAN Systems in Office Environment,” *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 8, pp. 1488–1499, 2009. 2.4
- [41] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, “60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling,” in *Proc. ACM SIGMETRICS*, 2015. 2.4, 4.6

-
- [42] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, “BeamSpy: Enabling Robust 60 GHz Links Under Blockage,” in *Proc. USENIX NSDI*, 2016. 2.4, 4.6
- [43] A. Zhou, X. Zhang, and H. Ma, “Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering,” in *Proc. IEEE INFOCOM*, 2017. 2.4, 4.6, 5.10
- [44] T. Wei, A. Zhou, and X. Zhang, “Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors,” in *Proc. USENIX NSDI*, 2017. 2.4, 4.6, 5.10
- [45] S. Naribole and E. Knightly, “Scalable Multicast in Highly-Directional 60 GHz WLANs,” in *Proc. IEEE SECON*, 2016. 2.4, 3.4.2
- [46] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, “Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement,” in *Proc. IEEE INFOCOM*, 2015. 2.4, 4.2.2, 4.6, 5.10
- [47] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan, “OpenMili: A 60 GHz Software Radio Platform With a Reconfigurable Phased-Array Antenna,” in *Proc. ACM MobiCom*, 2016. 2.4
- [48] O. Abari, H. Hassanieh, M. Rodriguez, and D. Katabi, “Millimeter Wave Communications: From Point-to-Point Links to Agile Network Connections,” in *Proc. ACM HotNets*, 2016. 2.4, 4.2.2, 4.6
- [49] X. Tie, K. Ramachandran, and R. Mahindra, “On 60 GHz Wireless Link Performance in Indoor Environments,” in *Proc. PAM*, 2012. 2.4
- [50] 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. 2.4
- [51] S. K. Saha, A. Garg, and D. Koutsonikolas, “A First Look at TCP Performance in Indoor IEEE 802.11 ad WLANs,” in *Proc. INFOCOM Workshop*, 2015. 2.4
- [52] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, “A Feasibility Study of 60 GHz Indoor WLANs,” in *Proc. IEEE ICCCN*, 2016. 2.4
- [53] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, “Multi-Gigabit indoor WLANs: Looking beyond 2.4/5 GHz,” in *Proc. IEEE ICC*, 2016. 2.4
- [54] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, “Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks,” *IEEE Transactions on Communications*, vol. 61, no. 10, pp. 4391–4403, 2013. 3.1.1

-
- [55] A. M. Sayeed and V. Raghavan, “Maximizing MIMO Capacity in Sparse Multipath With Reconfigurable Antenna Arrays,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 156–166, 2007. 3.1.1
- [56] O. Bejarano, E. W. Knightly, and M. Park, “IEEE 802.11ac: from Channelization to Multi-User MIMO.” *IEEE Communications Magazine*, 2013. 3.1.1
- [57] A. Alkhateeb, G. Leus, and R. W. Heath, “Limited Feedback Hybrid Precoding for Multi-User Millimeter Wave Systems,” *IEEE Transactions on Wireless Communications*, vol. 14, no. 11, pp. 6481–6494, 2015. 3.1.1, 3.8, 4.6
- [58] K.-P. Ho, S. Cheng, and J. Liu, “MIMO Beamforming in Millimeter-Wave Directional Wi-Fi,” *ArXiv e-prints*, 2014. 3.1.1, 4.1
- [59] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, “Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs,” in *Proc. ACM MobiCom*, 2010. 3.1.4, 3.6.1
- [60] J. Wang, “Beam Codebook based Beamforming Protocol for Multi-Gbps Millimeter-Wave WPAN Systems,” *IEEE Journal on Selected Areas in Communications*, 2009. 3.2.2.2
- [61] N. Anand, E. Aryafar, E.W. Knightly, “WARPlab: a Flexible Framework for Rapid Physical Layer Design,” in *Proc. ACM workshop on Wireless of the students, by the students, for the students*, 2010. 3.4.2
- [62] S. Sun, G. R. MacCartney, and T. S. Rappaport, “Millimeter-Wave Distance-Dependent Large-Scale Propagation Measurements and Path Loss Models for Outdoor and Indoor 5G Systems,” in *Proc. IEEE EuCAP*, 2016, pp. 1–5. 3.4.2
- [63] M. K. Samimi, S. Sun, and T. S. Rappaport, “MIMO Channel Modeling and Capacity Analysis for 5G Millimeter-Wave Wireless Systems,” in *Proc. IEEE EuCAP*, 2016, pp. 1–5. 3.4.2
- [64] R. A. Stirling-Gallacher and M. S. Rahman, “Multi-User MIMO Strategies for a Millimeter Wave Communication System using Hybrid Beam-forming,” in *Proc. IEEE ICC*, 2015. 3.8
- [65] Y. Ghasempour, N. Prasad, M. Khojastepour, and S. Rangarajan, “Novel Combinatorial Results on Downlink MU-MIMO Scheduling with applications,” in *Proc. IEEE WONS*, 2017. 3.8
- [66] S. Sur, I. Pefkianakis, X. Zhang, and K. Kim, “Practical MU-MIMO User Selection on 802.11ac Commodity Networks,” in *Proc. ACM MobiCom*, 2016. 3.8

-
- [67] M. Esslaoui, F. Riera-Palou, and G. Femenias, "A fair MU-MIMO Scheme for IEEE 802.11ac," in *International Symposium on Wireless Communication Systems*, 2012. 3.8
- [68] T. Tandai, H. Mori, and M. Takagi, "Cross-layer-optimized User Grouping Strategy in Downlink Multiuser MIMO Systems," in *Proc. IEEE VTC*, 2009. 3.8
- [69] N. Anand, J. Lee, S.-J. Lee, and E. W. Knightly, "Mode and User Selection for Multi-User MIMO WLANs without CSI," in *Proc. IEEE INFOCOM*, 2015. 3.8
- [70] X. Xie and X. Zhang, "Scalable User Selection for MU-MIMO Networks," in *Proc. IEEE INFOCOM*, 2014. 3.8
- [71] J. Geng, Z. Wei, N. Li, C. Chen, and D. Yang, "An Efficient Multi-User Scheme for the 60 GHz Wireless System with Low Cost Arrays," in *Proc. IEEE VTC*, 2013. 3.8
- [72] Y. Ghasempour, N. Prasad, M. Khojastepour, and S. Rangarajan, "Link packing in mmWave networks," in *Proc. IEEE ICC*, 2017. 3.8, 4.6
- [73] A. Michaloliakos, W. C. Ao, and K. Psounis, "Joint User-Beam Selection for Hybrid Beamforming in Asynchronously Coordinated Multi-Cell Networks," in *Proc. IEEE ITA Workshop*, 2016. 3.8
- [74] Y. Ghasempour, N. Prasad, M. Khojastepour, and S. Rangarajan, "Managing Analog Beams in mmWave Networks," in *Proc. of Asilomar Conference on Signals, Systems and Computers*, 2017. 3.8, 4.6
- [75] Z. Chen, X. Zhang, S. Wang, Y. Xu, J. Xiong, and X. Wang, "BUSH: Empowering Large-Scale MU-MIMO in WLANs With Hybrid Beamforming," in *Proc. IEEE INFOCOM*, 2017. 3.8, 4.6
- [76] W. Roh, J. Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar, "Millimeter-Wave Beamforming as an Enabling Technology for 5G Cellular Communications: Theoretical Feasibility and Prototype Results," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 106–113, 2014. 4
- [77] T. Baykas, C.-S. Sum, Z. Lan, J. Wang, M. A. Rahman, H. Harada, and S. Kato, "IEEE 802.15. 3c: The First IEEE Wireless Standard for Data Rates over 1 Gb/s," *IEEE Communications Magazine*, vol. 49, no. 7, pp. 114–121, 2011. 4.1, 4.1.3.2
- [78] IEEE, "802.11ay task group," <http://www.ieee802.org/11/Reports/tgay-update.htm>. 4.1.1.3, 4.1.3
- [79] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Wireless Networks," in *Proc. of ACM MobiCom*, 2017. 4.2.2, 4.6, 5.10

-
- [80] K. C. Hung and D. W. Lin, "Pilot-Based LMMSE Channel Estimation for OFDM Systems With Power Delay Profile Approximation," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 1, pp. 150–159, 2010. 4.2.3.1
- [81] Y. J. Kim and G. H. Im, "Pilot-Symbol Assisted Power Delay Profile Estimation for MIMO-OFDM Systems," *IEEE Communications Letters*, vol. 16, no. 1, pp. 68–71, 2012. 4.2.3.1
- [82] J. Zhang, Y. R. Zheng, C. Xiao, and K. B. Letaief, "Channel Equalization and Symbol Detection for Single-Carrier MIMO Systems in the Presence of Multiple Carrier Frequency Offsets," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 4, pp. 2021–2030, 2010. 4.2.3.1
- [83] A. N. Mody and G. L. Stuber, "Synchronization for MIMO OFDM Systems," in *Proc. IEEE GLOBECOM*, 2001. 4.2.3.1
- [84] A. Maltsev, A. Sadri, C. Cordeiro, and A. Pudeyev, "Practical LOS MIMO Technique for Short-Range Millimeter-Wave Systems," in *Proc. IEEE ICUBW*, 2015. 4.2.5
- [85] C. R. C. M. da Silva, A. Lomayev, C. Chen, and C. Cordeiro, "Analysis and Simulation of the IEEE 802.11ay Single-Carrier PHY," in *Proc. IEEE ICC*, 2018. 4.5, 4.5.1.1
- [86] M. K. Haider, Y. Ghasempour, and E. W. Knightly, "Search Light: Tracking Device Mobility Using Indoor Luminaries to Adapt 60 GHz Beams," in *Proc. ACM MobiHoc*, 2018. 4.6
- [87] J. Xiong and K. Jamieson, "Towards Fine-grained Radio-based Indoor Location," in *Proc. ACM HotMobile*, 2012. 4.6
- [88] J. Xiong and K. Jamieson, "ArrayTrack: A Fine-grained Indoor Location System," in *Proc. of USENIX NSDI*, 2013. 4.6, 5.10
- [89] C. Wong, R. Klukas, and G. G. Messier, "Using WLAN Infrastructure for Angle-of-Arrival Indoor User Location," in *Proc. IEEE VTC*, 2008. 4.6, 5.10
- [90] S. Sun, T. S. Rappaport, R. W. Heath, A. Nix, and S. Rangan, "MIMO for Millimeter-Wave Wireless Communications: Beamforming, Spatial Multiplexing, or Both?" *IEEE Communications Magazine*, vol. 52, no. 12, pp. 110–121, 2014. 4.6
- [91] D. R. Jackson and A. A. Oliner, "Leaky-Wave Antennas," *Modern Antenna Handbook*, pp. 325–367, 2008. 5
- [92] F. B. Gross, *Frontiers in Antennas: Next Generation Design & Engineering*. McGraw-Hill Education, 2011. 5.1, 5.4

-
- [93] N. J. Karl, R. W. McKinney, Y. Monnai, R. Mendis, and D. M. Mittleman, "Frequency-division Multiplexing in the Terahertz Range using a Leaky-wave Antenna," *Nature Photonics*, vol. 9, no. 11, p. 717, 2015. 5.1, 5.3, 5.10
- [94] K. Sengupta and A. Hajimiri, "A 0.28 THz Power-Generation and Beam-Steering Array in CMOS based on Distributed Active Radiators," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 12, pp. 3013–3031, 2012. 5.2
- [95] N. Karl, M. S. Heimbeck, H. O. Everitt, H.-T. Chen, A. J. Taylor, I. Brener, A. Benz, J. L. Reno, R. Mendis, and D. M. Mittleman, "Characterization of an Active Metasurface using Terahertz Ellipsometry," *Applied Physics Letters*, vol. 111, no. 19, p. 191101, 2017. 5.2
- [96] T. Chi, M.-Y. Huang, S. Li, and H. Wang, "17.7 A Packaged 90-to-300GHz Transmitter and 115-to-325GHz Coherent Receiver in CMOS for Full-Band Continuous-wave mm-Wave Hyperspectral Imaging," in *Proc. of IEEE ISSCC*, 2017. 5.2, 5.5.2.3, 5.7
- [97] C. A. Balanis and C. R. Birtcher, "Antenna measurements," *Modern Antenna Handbook*, pp. 977–1033, 2008. 5.3, 5.10
- [98] M. C. Beard, G. M. Turner, and C. A. Schmuttenmaer, "Terahertz Spectroscopy," 2002. 5.3
- [99] M. Mbonye, R. Mendis, and D. M. Mittleman, "Measuring TE 1 Mode Losses in Terahertz Parallel-Plate Waveguides," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 34, no. 7-8, pp. 416–422, 2013. 5.4
- [100] A. Moldovan, M. A. Ruder, I. F. Akyildiz, and W. H. Gerstacker, "LOS and NLOS Channel Modeling for Terahertz Wireless Communication with Scattered Rays," in *IEEE Globecom Workshops*, 2014. 5.5.1
- [101] I. Duling and D. Zimdars, "Terahertz Imaging: Revealing Hidden Defects," *Nature Photonics*, vol. 3, no. 11, p. 630, 2009. 5.7
- [102] Y. Yang, M. Mandehgar, and D. R. Grischkowsky, "Understanding THz Pulse Propagation in the Atmosphere," *IEEE Transactions on Terahertz Science and Technology*, vol. 2, no. 4, pp. 406–415, 2012. 5.7
- [103] Y. Yang, M. Mandehgar, and D. R. Grischkowsky, "THz-TDS Characterization of the Digital Communication Channels of the Atmosphere and the Enabled Applications," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 36, no. 1, pp. 97–129, 2015. 5.7
- [104] R. Mendis, A. Nag, F. Chen, and D. M. Mittleman, "A Tunable Universal Terahertz Filter using Artificial Dielectrics based on Parallel-Plate Waveguides," *Applied physics letters*, vol. 97, no. 13, p. 131106, 2010. 5.10

-
- [105] S. Gupta, S. Abielmona, and C. Caloz, “Microwave Analog Real-Time Spectrum Analyzer (RTSA) Based on the Spectral-Spatial Decomposition Property of Leaky-Wave Structures,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 57, no. 12, pp. 2989–2999, 2009. 5.10
- [106] W.-Q. Wang, H. C. So, and A. Farina, “An Overview on Time/Frequency Modulated Array Processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 2, pp. 228–246, 2017. 5.10
- [107] P. Rocca, L. Poli, and A. Massa, “Instantaneous Directivity Optimisation in Time-Modulated Array Receivers,” *IET Microwaves, Antennas & Propagation*, vol. 6, no. 14, pp. 1590–1597, 2012. 5.10
- [108] M. K. Haider, Y. Ghasempour, D. Koutsonikolas, and E. W. Knightly, “LiS-teer: MmWave Beam Acquisition and Steering by Tracking Indicator LEDs on Wireless APs,” in *Proc. ACM MobiCom*, 2018. 5.10
- [109] Y.-S. Kuo, P. Pannuto, K.-J. Hsiao, and P. Dutta, “Luxapose: Indoor Positioning with Mobile Phones and Visible Light,” in *Proc. ACM MobiCom*, 2014. 5.10
- [110] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, “Noncoherent mmWave Path Tracking,” in *Proc. HotMobile*, 2017. 5.10
- [111] B. Peng, K. Guan, S. Rey, and T. Kurner, “Two-Step Angle-of-Arrival Estimation for Terahertz Communications based on Correlation of Power-Angular Spectra in Frequency,” in *In Proc. EuCAP*, 2018. 5.10
- [112] A. Zhou, T. Wei, X. Zhang, and H. Ma, “FastND: Accelerating Directional Neighbor Discovery for 60-GHz Millimeter-Wave Wireless Networks,” *IEEE/ACM Transactions on Networking.*, vol. 26, no. 5, pp. 2282–2295, 2018. 5.10
- [113] J. Ning, T.-S. Kim, S. V. Krishnamurthy, and C. Cordeiro, “Directional Neighbor Discovery in 60 GHz Indoor Wireless Networks,” in *Proc. ACM MSWiM*, 2009. 5.10
- [114] R. Murawski, E. Felemban, E. Ekici, S. Park, S. Yoo, K. Lee, J. Park, and Z. H. Mir, “Neighbor Discovery in Wireless Networks with Sectorized Antennas,” *Ad hoc networks*, vol. 10, no. 1, pp. 1–18, 2012. 5.10
- [115] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, “Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 5, pp. 831–846, 2014. 5.10
- [116] S. Hur, T. Kim, D. J. Love, J. V. Krogmeier, T. A. Thomas, and A. Ghosh, “Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell

-
- Networks,” *IEEE Transactions on Communications*, vol. 61, no. 10, pp. 4391–4403, 2013. 5.10
- [117] Z. Marzi, D. Ramasamy, and U. Madhow, “Compressive Channel Estimation and Tracking for Large Arrays in mm-Wave Picocells,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 3, pp. 514–527, 2016. 5.10
- [118] Z. Yang, P. H. Pathak, Z. Yunze, and P. Mohapatra, “Sensor-Assisted Codebook-Based Beamforming for Mobility Management in 60 GHz WLANs,” in *Proc. IEEE MASS*, 2015. 5.10
- [119] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, “Fast Millimeter Wave Beam Alignment,” in *Proc. ACM SIGCOMM*, 2018. 5.10