UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Reliable, Scalable, and Deployable Millimeter-Wave Networks for Communication, Sensing, and Security

Permalink

https://escholarship.org/uc/item/9j89r1nm

Author

Jain, Ish Kumar

Publication Date

2024

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA SAN DIEGO

Reliable, Scalable, and Deployable Millimeter-Wave Networks for Communication, Sensing, and Security

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

in

Electrical Engineering (Communication Theory and Systems)

by

Ish Kumar Jain

Committee in charge:

Professor Dinesh Bharadia, Chair
Professor Pamela Cosman
Professor Ryan Kastner
Professor Shivendra Panwar
Professor Bhaskar Rao
Professor Xinyu Zhang

2024

Copyright

Ish Kumar Jain, 2024

All rights reserved.

The Dissertation of Ish Kumar Jain is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

Dedicated to my parents, Meena and Satish Jain, and my brother Atul Jain, for their continued support and encouragement in my endeavours of my research career.

TABLE OF CONTENTS

Dissertation Approval Page iii

Dedication iv

Table of Contents v

List of Figures ix

List of Tables xv

Preface xvi

Acknowledgements xvii

Vita xxi

Abstract of the Dissertation xxiii

Introduction 1
0.1 Systems Developed 3
0.1.1 Reliable and Scalable Communication for NextG Networks 4
0.1.2 Wireless Sensing for NextG Networks 6
0.1.3 Security in NextG Communication and Sensing 7
0.2 Contributions 8
0.3 Impact beyond this thesis 9
0.4 Organization 10

Chapter 1 mmReliable: Reliable and High Throughput mmWave Links 11
1.1 Background and Motivation 15
1.1.1 mmReliable: Making Multi-beam Practical 16
1.1.2 Beam Maintenance for Reliable mmWave 16
1.2 Design of Multi-beam System 18
1.2.1 Multi-beam Links are More Reliable 18
1.2.2 Multi-beam Throughput Model 20
1.2.3 Creating Constructive Multi-beam 26
1.2.4 Delay-Phased Array Beamforming for Wideband Channel 29
1.3 Proactive multi-beam tracking 33
1.3.1 Proactive Blockage Mitigation 34
1.3.2 Proactive User-Mobility Management 34
1.3.3 Superresolution: Per-Beam Tracking 36
1.3.4 Generalization to Multi-beam UE 38
1.4 mMobile: A Compact, Real-time mmWave Testbed 40
1.4.1 Requirements and Challenges 41

1.4.2 Related mmWave Testbeds in Literature 43
1.4.3 mMobile Testbed Architecture 45
1.5 mmReliable Implementation 47
1.5.1 Phased Array and Beamforming 48
1.5.2 Baseband Module and 5G-NR Compliance 49
1.6 Evaluation 50
1.6.1 Micro Benchmarks 50
1.6.2 End-to-end Results 55
1.7 Related Work 57
1.8 Discussion and Future Work 58
1.9 Acknowledgements 59

Chapter 2 mmFlexible: Flexible and Scalable Multi-user mmWave Networks 61
2.1 Background and Motivation 65
2.1.1 Primer on phased arrays and true-time delay arrays 66
2.1.2 Current mmWave Architectures are Incapable for Flexible Directional-Frequency Multiplexing 67
2.2 mmFlexible's DPA hardware design 69
2.2.1 Architecture of Delay-Phased Array (DPA) 69
2.2.2 Range of Delay Element in DPA 70
2.3 mmFlexible's DPA Software Design 71
2.3.1 Requirements for mmFlexible 71
2.3.2 Meeting mmFlexible's Requirements with DPA 71
2.3.3 FSDA Algorithm for Estimating Delays and Phases in DPA 73
2.4 Evaluation 76
2.4.1 Implementation and Emulation with 28 GHz Dataset 76
2.4.2 End-to-end Results 78
2.4.3 Benchmarks 80
2.5 Related Work 82
2.6 Discussion and Future Work 84
2.7 Conclusion 85
2.8 Acknowledgements 85

Chapter 3 FlexLink: Decoupling Control and Data Beams in mmWave Networks 87
3.1 Background and Motivation 91
3.1.1 Primer on 5G NR SSB Control Signals 91
3.1.2 Problem Statement: Coupled Control and Data Signaling 92
3.1.3 Motivation: Decoupling Control and Data with DPA 92
3.2 Design for FlexLink 94
3.2.1 Fast O(1) Computation of Delay and Phase in DPA 94
3.2.2 Generalization to an Arbitrary Number of Beams with O(1) Computation Complexity 104
3.2.3 Supporting Arbitrary Size Bandwidth Part with DPA 109
3.2.4 Design and Implementation of DPA Hardware 111

3.3 Evaluation . . . 114
3.3.1 Hardware Results . . . 114
3.3.2 HFSS Evaluations . . . 115
3.3.3 Scaling to Large Arrays . . . 116
3.4 Related Work . . . 120
3.5 Conclusion . . . 122
3.6 Acknowledgements . . . 122

Chapter 4 CommRad: Collaborative Bi-Directional Sensing-Driven Millimeter-Wave Networks . . . 123
4.1 Background and Related Work . . . 127
4.2 CommRad Design . . . 129
4.2.1 What is the Context that Radar Lacks? . . . 131
4.2.2 Acquiring Context Through Radar+Radio Collaboration . . . 133
4.2.3 Context-Aware Beam Management using Radar Sensor . . . 137
4.3 Implementation . . . 143
4.3.1 Radio Hardware and Software . . . 143
4.3.2 Radar Hardware and Software . . . 145
4.4 Evaluation . . . 145
4.4.1 Overall Throughput Improvement . . . 146
4.4.2 Overhead Reduction . . . 148
4.4.3 Angle Tracking Accuracy . . . 149
4.4.4 User/Reflector Tracking Accuracy . . . 150
4.4.5 Blockage Prediction and Mitigation . . . 151
4.5 Discussion and Future work . . . 151
4.6 Conclusion . . . 152
4.7 Acknowledgements . . . 152

Chapter 5 Security Challenges in Communication and Sensing Systems . . . 154
5.1 mmSpoof: Resilient Spoofing of Automotive Millimeter-wave Radars using Reflect Array . . . 154
5.2 BeamArmor: Seamless Anti-Jamming in 5G Cellular Networks with MIMO Null-steering . . . 158
5.2.1 MIMO null-steering model for jammer mitigation . . . 161
5.2.2 Practical, Real-time, and O-RAN-driven RAN Controller for BeamArmor 163
5.2.3 BeamArmor Application for Jammer Monitoring and Nulling . . . 165
5.2.4 Over-the-air Experiments . . . 166
5.2.5 Monitoring Jammer Interference with BeamArmor . . . 166
5.2.6 Mitigating Jammer with BeamArmor . . . 166
5.2.7 Platform Evaluation . . . 170
5.2.8 Related Work . . . 171
5.2.9 Limitations and Future Directions . . . 172
5.3 Conclusion and Outlook . . . 173
5.4 Acknowledgements . . . 173

Chapter 6 Conclusion 175
6.1 Concluding Remarks 175
6.2 Future Directions 178

Bibliography 181

LIST OF FIGURES

Figure 0.1. **Experimental testbeds:** Built to demonstrate our systems in real-world settings for (a) mmReliable [Sigcomm 2021], (b) mmFlexible [Infocom 2023], (c) CommRad [Preprint 2023], (d) mmSPoof [IEEE S&P 2023]. . . 3

Figure 0.2. **mmReliable:** Reliable constructive multi-beams. 4

Figure 0.3. **mmFlexible:** Flexible frequency-dependent multi-beams. 5

Figure 1.1. **mmReliable idea:** mmReliable utilizes beams with multiple lobes (*Multi Beam*) to provide high reliability compared to a single-lobe beams. 12

Figure 1.2. **Building blocks of mmReliable:** Beam management that leverages initial beam training for angle estimation and reference signaling for multi-beam establishment and maintenance. 17

Figure 1.3. **Constructive multi-beam:** Enhances the throughput & reliability of mmWave link. 19

Figure 1.4. **Results on Reflector Attenuation:** (a) CDF of the relative attenuation of the strongest reflected path compared to the direct path at over various locations. (b) A heatmap of strong paths in the environment while the UE moves, strong reflectors appear at different points in time. 25

Figure 1.5. **Channel probing procedure:** To obtain parameters (δ, σ) for multi-beam constructive combining that's robust to CFO/SFO impairments. 28

Figure 1.6. **Proposed architecture for delay phased array:** Delay-based multi-beamforming to handle wideband transmission. Note that this architecture can solve wideband issues due to multi-path channels. 30

Figure 1.7. **Constructive Combining for Wideband frequencies:** Illustration of the advantages of our delay compensation with delay phased array architecture. We achieve a flat frequency response despite the presence of multi-path delays by compensating for the additional delays caused by the multi-path. 32

Figure 1.8. **SNR with a range of frequencies:** The delay-optimized mmReliable response is flat across all frequencies irrespective of any channel delay spread. It also shows for two different delay spreads of 5 ns and 10 ns, the SNR suffers at certain frequencies. 33

Figure 1.9. **Overview:** mmReliable's proactive response against blockage and user mobility. 34

Figure 1.10. **User mobility:** Causes beams in multi-beam misalign with the user. Each beam may undergo a different angular deviation. 36

Figure 1.11. **Efficiency of superresolution algorithm:** (a) Simulation: Superresolution provides low MSE of per-beam power estimate even when the relative ToF is less than system resolution of 2.5 ns. (b) Hardware measurement: shows two sincs can be recovered from measured combined CIR. 37

Figure 1.12. **Multi-Beam UE:** Demonstrating beam misalignment for a multi-beam UE. 38

Figure 1.13. **Compact mobile setup** with 28 GHz phased array and a PlutoSDR streaming OFDM waveform using USB stick. 40

Figure 1.14. **Block diagram** mmReliable testbed 42

Figure 1.15. **Experimentation Testbed:** with FPGA and USRP baseband and 64 element phased array 48

Figure 1.16. **Sensitivity analysis (simulation):** with varying amplitude and phase per beam 51

Figure 1.17. **Benchmark Results on mmReliable:** with different phase/amplitude per beam, frequency, and beamforming type 51

Figure 1.18. **Blockage-Resilience Result:** Multi-beam is more reliable and blockage resilient compared to a single beam. 52

Figure 1.19. **Mobility Tracking Results:** with UE rotation and translation 52

Figure 1.20. **End-to-End Results:** for throughput, reliability and system overhead 52

Figure 2.1. **mmFlexible Idea:** It enables efficient use of the available mmWave spectrum resources through flexible directional-frequency multiplexing, allowing multiple users to be served simultaneously with low latency and high spectrum utilization. 62

Figure 2.2. **Desired frequency-space beam response:** 3 users in 3 different angular directions. 65

Figure 2.3. **Comparison of different mmWave front-end architectures:** Corresponding frequency-space patterns they created. 66

Figure 2.4. **Explanation on DPA:** How DPA creates a desired frequency-space image with two beams at -20° and 20°. 68

Figure 2.5. **FSDA Algorithm:** Three steps of our FSDA algorithm to estimate per-antenna delays and phases that generate a desired frequency-space image. The core step is our novel 2D FSDA transform with some pre-processing and post-processing of inputs and outputs respectively. 73

Figure 2.6. **Implementation overview for mmFlexible:** Four main components: Data traffic generator, Fast beam scan angle estimation, a scheduler, and our FSDA algorithm. 76

Figure 2.7. **End-to-End Results:** Performance gain with DPA compared to baselines under different scenarios 76

Figure 2.8. **Benchmark Results:** Performance gain of mmFlexible compared to baselines under different scenarios: (a) Angle separation between two users (b) (c) Subcarrier allocation between users (d) How much minimum frequency resources can be allocated without SNR degradation. 78

Figure 2.9. **Large-scale Simulation:** Scaling the mmFlexible's performance with (a) Increasing number of user directions and (2) Increasing number of antennas. 79

Figure 3.1. **FlexLink** designs and implements a novel radio architecture that can decouple control and data signaling through flexible frequency resource distribution of control and data to different directions with narrow pencil beams, much like a flexible configurable prism. 88

Figure 3.2. **Delay-phased array (DPA)** is an analog front-end architecture with programmable delay and phase elements. 91

Figure 3.3. **Proof of estimating closed-form delay and phase** values for a given two-beam response at $\pm\theta_0$. 96

Figure 3.4. **Desired frequency-space beam response** for 2 users at directions $-\theta_0$ and θ_0. 97

Figure 3.5. **Sketch of proof**: We show the phase response for each antenna is a step function with a variable step size that depends on the antenna index 97

Figure 3.6. **Generalized setting for multi-beams** with an arbitrary number of beams D, beam angles, and bandwidth parts. The beam angles contributes to ϕ_d, and the bandwidth part is defined as $\alpha_d B$ for system bandwidth B and fraction $\alpha_d < 1$ for beam index $d \leq D$. 104

Figure 3.7. **Effect of delay and phase quantization** in overall array gain of 16 antennae DPA. 110

Figure 3.8. Effect of delay and phase quantization in overall array gain of 16 antennae DPA. 111

Figure 3.9. **Implementing delay-phased array** at 4-6 GHz band. 111

Figure 3.10. **Hardware components** of TTD unit and phase shift unit. 111

Figure 3.11. **Beam pattern measurement** in an anechoic chamber. 111

Figure 3.12. **Experimental results** of beamforming gain for two beams at 30^o and -30^o with different bandwidth parts per beam. 115

Figure 3.13. **Evaluation Results:** FlexLink improves spectral efficiency and increases spectrum utilization. 116

Figure 3.14. **HFSS hardware simulation** of DPA. 117

Figure 3.15. **Simulation for large antenna arrays** and 3+ number of simultaneous beams. 118

Figure 4.1. **CommRad** is a collaborative learning framework for context- aware sensing- driven mmWave communication. 124

Figure 4.2. **CommRad's bi-directional radar-radio collaborative learning** framework to improve data communication efficiency by reducing radio beam scan overhead. 126

Figure 4.3. **Overview** of CommRad's end-end system implementation of collaborative bi-directional learning with radar+radio integration. 131

Figure 4.4. **Challenges** with lack of context for user identification 133

Figure 4.5. **Challenges** with lack of context for reflector identification. 133

Figure 4.6. **Flow Diagram:** Identifying users and reflectors in radar view using contextual information from radio. 134

Figure 4.7. **Reflector location estimation** using radio beam scan measurement. 135

Figure 4.8. **Reflector orientation and size** estimation. 136

Figure 4.9. **Clutter removal** and user localization 138

Figure 4.10. **Proposed user tracking** algorithm. 139

Figure 4.11. **Tracking** both direct path and reflected path. 140

Figure 4.12. **Blockage tracking** and proactive prevention. 141

Figure 4.13. **Setup:** Base station with synchronized Radio and Radar setup 142

Figure 4.14. **Experiment scenarios** with the static radar-radio platform with four mobile users. We consider diverse scenarios indoors and outdoors with various reflectors such as concrete walls, monitor screens, glass walls, etc. 142

Figure 4.15. **End-end results of CommRad:** Improving median throughput by 2.5x compared to non-collaborative baseline. Supporting both single-beam (SB) and Constructive multi-beam (CMB) schemes with blockage mitigation. Reducing beam management overhead in 5GNR. 144

Figure 4.16. **Scenarios:** Four users crossing each other. We remove the clutter and then show angle-tracking accuracy with collaborative learning. 147

Figure 4.17. **Results:** CommRad's end-end angle tracking accuracy. 148

Figure 4.18. **Radar Prediction Accuracy:** CommRad acquires ground truth by performing continuous beam training at all times. We show that the angle acquired by this beam scan (Oracle) matches the radar prediction. 150

Figure 4.19. **Blockage adaptation:** CommRad adaptation to dynamic blockages. 151

Figure 5.1. **mmSpoof:** spoofing technique is based on a mmWave reflect array, and it does not require prior knowledge about victim radar, to spoof arbitrary distance and velocity values. 155

Figure 5.2. **The experimental setup** consists of the mmReliable device mounted on a car, with the radar on a tripod. The top picture shows an overview of the setup, while the bottom figure shows a zoomed-in view. 156

Figure 5.3. **BeamArmor:** a controller platform for jammer monitoring and MIMO null- steering anti- jamming solutions implemented with open-source srsRAN 5G cellular stack. 159

Figure 5.4. **Core functional blocks** of our RAN controller for BeamArmor in interaction with the srsenb application and the USRP radio peripheral. 164

Figure 5.5. **Experiment set up** with the active elements (srsenb, srsUE, and jammer) in a lab environment. We tested with 3 different jammer locations. 166

Figure 5.6. **BeamArmor's cancellation abilities:** We present the mean received signal strength computed from received signal y in the form of $|y|^2$ as jammer power increases from 0 to 30 dBm. The red curve demonstrates how BeamArmor's null-steering decreases received jammer signal power. 167

Figure 5.7. Metrics SINR, throughput, and BLER for **Single Tone jammer.** with and without BeamArmor's anti-jamming solution (BA on/off). The green dotted line shows the ground truth scenario without (w/o) any jammer present. The red line represents the median value of one boxplot. 168

Figure 5.8. Metrics SINR, throughput, and BLER for **5 MHz Wideband jammer** with and without BeamArmor's anti-jamming solution (BA on/off). The green dotted line shows the ground truth scenario without (w/o) any jammer present. The red line represents the median value of one boxplot. 168

Figure 5.9. **UL SINR during active BeamArmor**controller communication at various down-sampling rates and at a fixed periodicity of every 10 TTI. UL SINR during an inactive BeamArmor controller serves as a benchmark performance. 170

LIST OF TABLES

Table 1.1. Comparison of mmWave testbeds. 45

Table 2.1. Delay-phased array & other baselines for 10 4k VR users (*We mention over-the-air worst case latency). 68

Table 5.1. UL SINR vs. increasing jammer TX power. 167

PREFACE

The next generation (NextG) of wireless networks would connect billions of devices using existing and limited natural spectrum resources, all the while making it secure and sustainable with limited power and antennas. This thesis takes a holistic approach to optimizing these resources to deliver not only high data rates but also enhance reliability and scalability and facilitate practical deployments. Moreover, as the wireless infrastructure is deployed ubiquitously in both indoor (e.g., WiFi) and outdoor (e.g., 5G cellular) environments, there is an opportunity to provide sensing services alongside data communication, eliminating the need for additional sensor deployment. By leveraging existing wireless infrastructure, we develop radio-frequency-based sensing solutions such as localization and tracking of mobile devices. Furthermore, this thesis recognizes that as communication and sensing systems become more accessible, they are vulnerable to security threats such as spoofing and jamming attacks, so we designed and implemented appropriate measures to counter these attacks.

ACKNOWLEDGEMENTS

I am deeply grateful to my mentor and advisor, Dinesh Bharadia, for his belief in me, his patient teaching, and his nurturing guidance, which have made me a confident and knowledgeable professional in my field. When I joined the WCSNG lab in 2018 with a theoretical background in wireless communication, I was unfamiliar with systems research. During those early days, Dinesh not only encouraged us to work on system implementations but also spent hours with us going through component datasheets and debugging hardware and systems. His cheerful and positive spirit always helped lift our mood during challenging times, providing new ideas and directions. Dinesh's mentorship instilled in me a passion for trying new ideas quickly and implementing them repeatedly until our system worked successfully. For example, when we faced the issue of lower throughput in the mmReliable project, it was Dinesh's consistent push, creative thinking, and out-of-the-box ideas that enabled us to not only improve but exceed traditional systems' throughput. Through his guidance, I learned to maintain a positive mindset toward research, pursue broad and innovative research directions, and adopt a persistent approach to trying, failing, and re-trying until achieving our goals. Thank you, Dinesh, for instilling these virtues in me.

I highly appreciate the positive and collaborative lab culture that we cultivated among all labmates— "Help will always be given at WCSNG to those who ask for it."—truly helping out each other voluntarily and passionately during deadlines and tough times. I am very thankful to my labmates for giving me support and confidence throughout the PhD journey. Particularly, I am fortunate to have my colleagues and collaborators Raghav Subbaraman, Rohith Reddy Vennam, and Frederik Zumegen, who immensely helped with my papers from refining the ideas to insightful discussions and helping with experiments. A special thanks to Raghav Subbaraman for my first collaboration on mmReliable and mmFlexible projects. Raghav was my go-to person for any hardware bugs and logistics issues. Through him, I learned the process of automating the experiments, conducting large-scale evaluations, and maintaining a sharable GitHub codebase for everyone to benefit from our learnings. My other collaboration with Rohith was incredible too;

While mentoring Rohith on mmFlexible and mmSpoof projects, I learned from his impeccable positive energy, hard work, and his style of questioning every fundamental assumption and building systems at bigger, larger scales. Frederik came from a different background and it was fun to mentor him for a wireless project on BeamArmor. I also get to learn from his amazing skills to implement and debug large-scale open-source code bases such as srsRAN. In addition, many Masters and undergraduate students helped a lot with my projects and it was a mutual learning experience while mentoring them on our projects- I am thankful to Tian Qiu, Tejas Sadarahalli, Raini Wu, Rohit Kumar, Aman Goyal, Navin Adarsh, Yeswanth Guntupalli, Indrajot Singh, Hou-Wei Lin, Weginbara (Michael) Youpele, Xiangwei Shao, YungYi Sun, Junyi Xu, Sonny Cao, Suriyaa MM, and Sesha Sai Rakesh for their time and efforts in building some of the amazing hardware prototypes for mmWave and O-RAN testbed, and helping out with PCB design and FPGA programming. Truly, my PhD would not have been possible without your help.

Thanks to the collaborative culture in our lab, I get to experience and contribute to many projects of my colleagues Roshan Ayyalasomayajula, Aditya Arun, Agrim Gupta, Kshitiz Bansal, Manideep Dunna, Adel Heidari, Ushasi Ghosh, and Sushila Seshasayee. I learned a lot about wireless sensing, Robotics, and Machine Learning from Roshan, Aditya, and Kshitiz. Our endless discussions with Agrim were deep and insightful all the way from criticizing and appreciating papers in our community, new crazy research ideas, and serious discussions on where we are moving as a community. Manideep has been always available for any discussion on any topic. Thanks to him, I could revise both my deep theoretical concepts and improve PCB designing and hardware knowledge. In addition, this thesis is incomplete if I don't mention my other labmates, Sreevatsank Kadaveru, Nagarjun Bhat, Hadi Givehchian, Richard Bell, Kevin Mills, Wei Sun, Archana Bura, and Srivatsan Rajagopal– all our group hangouts, coffee chats, weekly sports, hikes, and above all, the positive energy and compassion was needed for my mental and physical support to write this thesis. I am also thankful to Prof. Xinyu Zhang for his continuous encouragement and discussions and to my colleagues in Prof. Xinyu's lab– Song Wang, Renjie Zhao, Tim Woodford, Ke Sun, Kai Zheng to having insightful discussions and for

helping me borrow any electronics and equipment from their lab during deadline times.

I am incredibly lucky to have spent my PhD journey with my flatmates Manideep Dunna, Agrim Gupta, and Ashwin Verma at 8530 Costa Verde Blvd, a place I can call my second home outside India. They have been like a family to me, sharing so many moments from having dinner together, cooking together, watching movies together to celebrations during festivals. The beautiful beaches, hikes, and sunsets in San Diego also deserve credit and all my amazing friends who accompanied me to these social activities. I cannot imagine my life without their presence.

Chapter 1, in full, is a reprint of the material as it appears in **IK Jain**, R Subbaraman, D Bharadia, "Two beams are better than one: Towards Reliable and High Throughput mmWave Links", *ACM SIGCOMM 2021*. The dissertation author was the primary investigator and author of this paper.

Chapter 1, in part, is a reprint of the material as it appears in **IK Jain**, R Subbaraman, TH Sadarahalli, X Shao, H Lin, D Bharadia, "mMobile: Building a mmWave Testbed to Evaluate and Address Mobility Effects", *4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems (Mobicom Workshop - mmNets)*, 2020. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in full, is a reprint of the material as it appears in **IK Jain**, RR Vennam, R Subbaraman, D Bharadia, " mmFlexible: Flexible Directional Frequency Multiplexing for Multi-user mmWave Networks", *IEEE Infocom 2023*. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is a reprint of the material as it appears in **IK Jain**, D Bharadia, "FlexLink: Decoupling Control and Data Communication in mmWave Networks" *Preprint - Under submission to Tier-1 systems conference*. The dissertation author was the primary investigator and author of this paper.

Chapter 4, in full, is a reprint of the material as it appears in **IK Jain**, Suriyaa MM, D Bharadia, "CommRad: Collaborative Learning for Sensing-Driven mmWave Network" *Preprint - Under submission to Tier-1 systems conference*. The dissertation author was the primary

investigator and author of this paper.

Chapter 5, in full, is a reprint of the material as it appears in F Zumegen, **IK Jain**, D Bharadia , "BeamArmor: Anti-Jamming in 5G Cellular Networks with MIMO Null-steering" accepted for publication in *ACM Hotmobile* 2024. The dissertation author was the primary investigator and author of this paper.

VITA

2018 - 2024 PhD, University of California San Diego, Electrical and Computer Engineering

2016 - 2018 MS, New York University, Tandon School of Engineering, Electrical and Computer Engineering

2012 - 2016 B.Tech., Indian Institute of Technology (IIT Kanpur), India, Electrical Engineering

June–Dec 2022 Research Intern, **VMware,** Palo Alto, CA, USA

June–Aug 2017 Research Intern, **Nokia Bell Labs,** Murray Hill, NJ, USA

PUBLICATIONS

Preprint 2024	**IK Jain**, D Bharadia, "FlexLink: Decoupling Control and Data Communication in mmWave Networks" *Preprint - Under submission to Tier-1 systems conference*.
Preprint 2024	**IK Jain**, Suriyaa MM, D Bharadia, "CommRad: Collaborative Learning for Sensing-Driven mmWave Network" *Preprint - Under submission to Tier-1 systems conference*.
Preprint 2024	S Jonnavithula, **IK Jain**, D Bharadia, "MIMO-RIC: RAN Intelligent Controller for MIMO xApps" *Preprint - Under submission to Tier-1 systems Workshop*.
Hotmobile 2024	F Zumegen, **IK Jain**, D Bharadia , "BeamArmor: Anti-Jamming in 5G Cellular Networks with MIMO Null-steering" accepted for publication in *ACM Hotmobile* 2024.
Infocom 2023	**IK Jain**, RR Vennam, R Subbaraman, D Bharadia, "Flexible and resource efficient multi-user mmWave system", *IEEE Infocom 2023*. *Acceptance 252/ 1312 (19%)*
IEEE S&P 2023	RR Vennam, **IK Jain**, K Bansal, J Orozco, P Shukla, A Ranganathan, D Bharadia, "mmSpoof: Resilient Spoofing of Automotive Millimeter-wave Radars using Reflect Array", *IEEE Security and Privacy 2023*. *Acceptance 115/ 952 (12%)*
HotCarbon 2022	A Gupta, **IK Jain**, D Bharadia, "Multiple smaller base stations are greener than a single powerful one: Densification of Wireless Cellular Networks", *ACM HotCarbon Workshop 2022*
WPMC 2022	T Qiu, **IK Jain**, R Wu, P Cosman, D Bharadia, "Streaming 360-degree video with Viewport-adaptive Truncation", *nternational Symposium on Wireless Personal Multimedia Communications (WPMC) 2022*.

Sigcomm 2021	**IK Jain**, R Subbaraman, D Bharadia, "Two beams are better than one: Towards Reliable and High Throughput mmWave Links", *ACM SIGCOMM 2021*. *Acceptance 55/ 241 (22%)*
NSDI 2020	R Ayyalasomayajula, A Arun, C Wu, S Rajagopalan, S Ganesaraman, A Seetharaman, **IK Jain** and D Bharadia, "LocAP: Autonomous Millimeter Accurate Mapping of WiFi Infrastructure.", *NSDI 2020* *Acceptance 48/ 275 (17%)*
Mobicom—mmNets 2020	**IK Jain**, R Subbaraman, TH Sadarahalli, X Shao, H Lin, D Bharadia, "mMobile: Building a mmWave Testbed to Evaluate and Address Mobility Effects", *4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems (Mobicom Workshop - mmNets)*, 2020.
MDPI Journal 2019	A Choromanska, **IK Jain**, "Extreme Multiclass Classification Criteria", *MDPI Computation Journal*, 2019.
JSAC 2018	**IK Jain**, R Kumar, S Panwar, "The Impact of Mobile Blockers on Millimeter Wave Cellular Systems", *IEEE Journal on selected areas in communications (JSAC)*, 2018.
IEEE ITC 2018	**IK Jain**, R Kumar, S Panwar, "Driven by Capacity or Blockage? A Millimeter-wave Blockage Analysis" *IEEE International Teletraffic Congress (ITC)*, 2018.

FIELDS OF STUDY

Major Field: Electrical Engineering (Communication Theory and Systems)

Studies in Wireless Systems
Professor Dinesh Bharadia

ABSTRACT OF THE DISSERTATION

Reliable, Scalable, and Deployable Millimeter-Wave Networks for Communication, Sensing, and Security

by

Ish Kumar Jain

Doctor of Philosophy in Electrical Engineering (Communication Theory and Systems)

University of California San Diego, 2024

Professor Dinesh Bharadia, Chair

The next generation (NextG) of wireless networks would connect billions of devices using existing and limited natural spectrum resources, all the while making it secure and sustainable with limited power and antennas. In this thesis, I take a holistic approach to optimizing these resources to deliver not only high data rates but also enhance reliability and scalability and facilitate practical deployments. Moreover, as the wireless infrastructure is deployed ubiquitously in both indoor (e.g., WiFi) and outdoor (e.g., 5G cellular) environments, there is an opportunity to provide sensing services alongside data communication, eliminating the need for additional sensor deployment. By leveraging existing wireless infrastructure, my research focuses on

developing radio-frequency-based sensing solutions such as localization and tracking of mobile devices. Furthermore, my research recognizes that as communication and sensing systems become more accessible, they are vulnerable to security threats such as spoofing and jamming attacks, so I designed and implemented appropriate measures to counter these attacks.

Introduction

Over the last decade, we have seen many applications such as autonomous vehicles, drones, and extended reality (XR) devices that require high-speed wireless connectivity with the network infrastructure. At the same time, the radio frequency signals can be used as "Signals of opportunity" to provide sensing services such as localization and tracking without requiring a dedicated sensor deployment. Consequently, as we see heavy deployment of network infrastructure, it is vulnerable to security threats such as spoofing and jamming attacks. So, can we re-imagine and build the next-generation networks while considering the requirements for communication, sensing, and security? This thesis develops advanced technologies for these three pillars of NextG networks:

- **Communication:** Millimeter-wave links are essential for NextG networks due to abundant bandwidth, but they are inherently unreliable due to blockage and struggle to scale for a large number of devices. This thesis identified a fundamental problem in mmwave links, which rely on a "single directional beam" that creates a single point of failure and limits service to one user direction at a time. In mmReliable [Sigcomm 2021], we challenged the conventional single-beam assumption and introduced a constructive multi-beam system that improves both the reliability and throughput of mmWave links. Furthermore, in mmFlexible [Infocom 2023], we developed a new class of frequency-dependent multi-beams that serve multiple devices in a single timeslot, establishing a scalable and efficient multi-user system. We further explored new applications of multi-beams in decoupling control and data beams in 5G NR systems in Flexlink [Preprint 2024].

- **Sensing:** It is extremely hard to deliver on the above promises in a realistic world with hundreds of devices with complicated mobility patterns and complex propagation environments. My approach is to integrate communication radio with novel radio frequency-based sensing (CommRad [Preprint 2024] and LocAP [NSDI 2020]) to bring our technology out of the lab environment to challenging commercial deployments, with support from Qualcomm.

- **Security:** We demonstrated a spoofing attack on automotive radar sensors by employing a mmWave reflect array as a *spoofing device* in mmSpoof [IEEE S&P 2023]. In another initiative, we built an Open-RAN-driven jammer monitoring and anti-jamming system in BeamArmor [ACM Hotmobile 2024, IEEE Milcom 2023]. Moreover, we effectively implemented some of these techniques in a commercial setting at VMware.

Research Methodology: The research methodology in this thesis bridges the gap between theory and systems. As an MS student at New York University, I established a strong theoretical foundation through projects such as stochastic geometry-based network optimization [JSAC 2018] and theoretical machine learning [MDPI 2019]. Moving forward, during my PhD at UC San Diego, I started building large-scale systems to demonstrate our ideas in practice and achieve real-world impact. Specifically, we built a 28 GHz mmWave testbed called M-Mobile [mmNets 2020], [Sigcomm 2021] based on phased arrays and software-defined radios such as USRPs and FPGAs as shown in Figure 0.1(a).

These experiences helped me follow a threefold approach to addressing fundamental problems in wireless communication: Start with understanding the theory behind the problem, followed by creating a simulation and emulation model to verify the theory, and finally, develop algorithms that actually work on hardware systems.

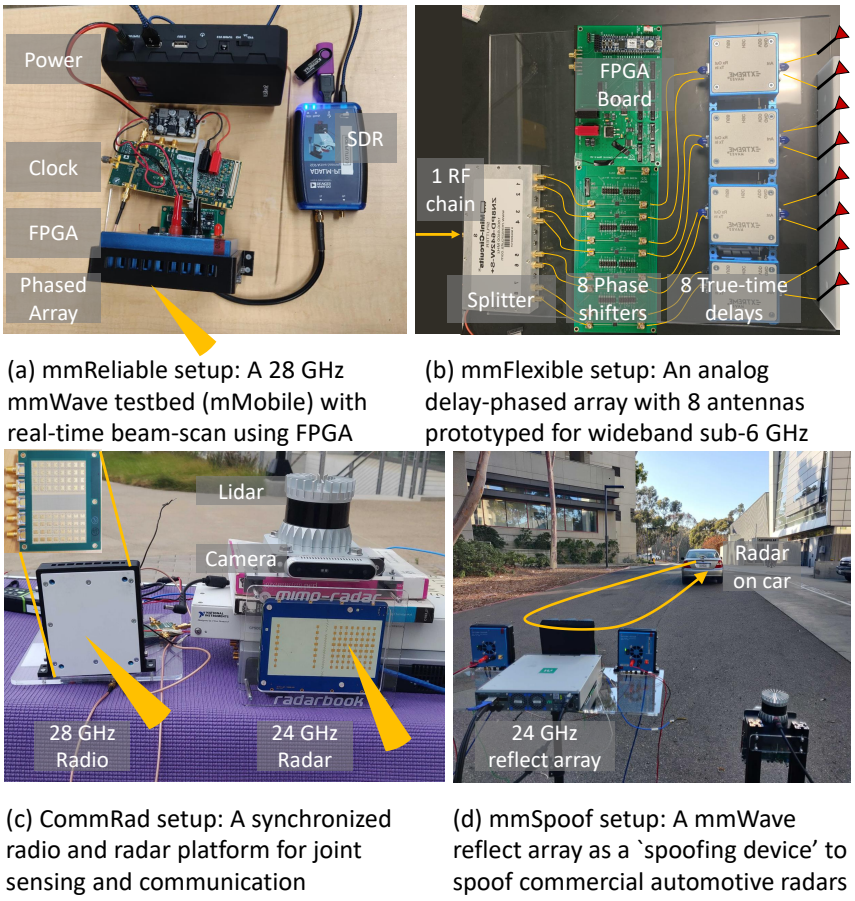

(a) mmReliable setup: A 28 GHz mmWave testbed (mMobile) with real-time beam-scan using FPGA

(b) mmFlexible setup: An analog delay-phased array with 8 antennas prototyped for wideband sub-6 GHz

(c) CommRad setup: A synchronized radio and radar platform for joint sensing and communication

(d) mmSpoof setup: A mmWave reflect array as a `spoofing device' to spoof commercial automotive radars

Figure 0.1. Experimental testbeds: Built to demonstrate our systems in real-world settings for (a) mmReliable [Sigcomm 2021], (b) mmFlexible [Infocom 2023], (c) CommRad [Preprint 2023], (d) mmSPoof [IEEE S&P 2023].

0.1 Systems Developed

We developed multiple platforms for nextG networks, including a compact and real-time mmWave testbed using phased arrays, FPGAs, and software-defined radios shown in Figure 0.1(a) and a new antenna array architecture called delay-phased array for scalable and low-latency mmWave communication in Figure 0.1(b). We further build experimental platforms for sensing and security including a sync radar and radio platform for sensing-driven communication in Figure 0.1(c) an a mmWave reflect array that can carry out an attack on commercial automotive radars in Figure 0.1(d). These testbeds span across the three areas of research in Communication, Sensing, and Security.

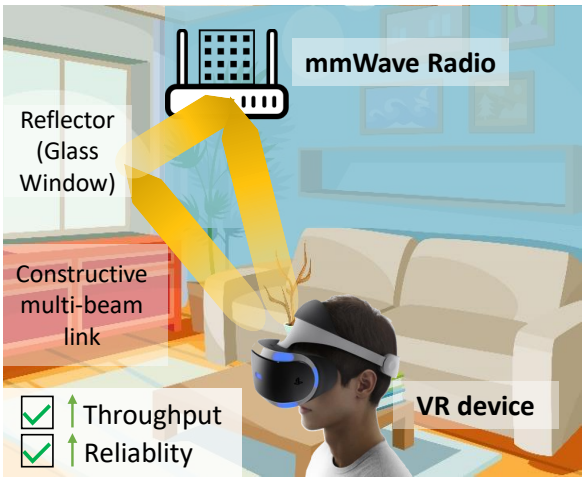

Figure 0.2. mmReliable: Reliable constructive multi-beams.

0.1.1 Reliable and Scalable Communication for NextG Networks

Reliable communication through constructive multi-beamforming:

Millimeter-wave technology has proven highly effective for high-throughput applications such as vehicular connectivity and video streaming. However, its deployment has predominantly been limited to fixed wireless settings and has yet to make substantial inroads into mobile access deployments. A critical challenge stems from mmWave's inability to offer a reliable link similar to sub-6 frequencies, where signals often degrade due to mobile users and obstacles. Traditional single-beam solutions are particularly susceptible to blockages because they create a single point of failure. We built a system called mmReliable [Sigcomm 2021] that deploys two or more beams at the same time to the same user to improve reliability. Intuitively, if one path to the user is blocked, an alternate path via some reflector is immediately available to maintain a reliable link as shown in Figure 0.2.

One significant challenge with multi-beams is the potential reduction in per-beam power, which could reduce overall throughput. Remarkably, we demonstrate that mmReliable not only maintains high throughput alongside high reliability but can also effectively double throughput under favorable channel conditions. The key idea was to create *constructive multi-beams* wherein

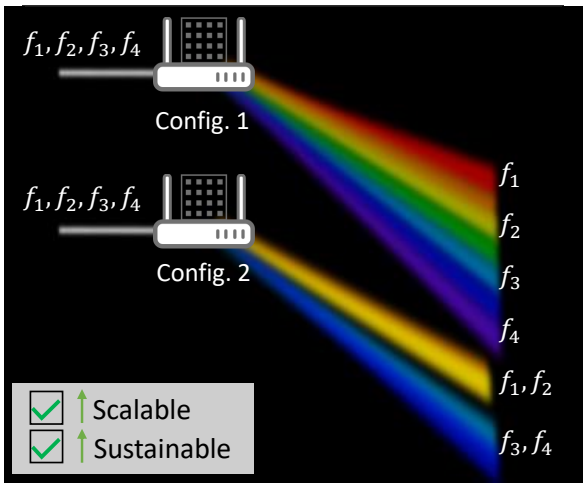

Figure 0.3. mmFlexible: Flexible frequency-dependent multi-beams.

the phase and power of each beam are optimized to ensure per-beam signals add constructively at the receiver. We implemented mmReliable with a single RF-chain-based testbed with 5G NR waveforms and demonstrated close to 100% reliability in mobile environments.

Scalable communication through frequency-dependent multi-beamforming:

Millimeter-wave bands support 100× more bandwidth than sub-6 systems, which could have easily been used to scale mmWave connectivity to hundreds or thousands of devices in a network, but it is limited to only a few devices. The underlying problem is not the availability of resources but the inefficient distribution of those resources to all users. The resource distribution in mmWave is limited to time-based scheduling, which serves one user at a time using a single directional beam. So, one user gets the entire bandwidth in that timeslot independent of the user's demand, leading to a wastage of spectrum resources.

I built mmFlexible [Infocom 2023] with the goal of distributing both time and frequency resources to multiple users in a flexible manner that scales with user demand. We introduced a special class of frequency-dependent multi-beams that can distribute different frequency bands to different users. A key feature of this system is the ability to assign an arbitrary set of frequencies to an arbitrary set of directions, much like a configurable prism as shown in

Figure 0.3. We achieved this flexibility by introducing a true-time delay element in the phased array and jointly programming the delays and phases at each antenna. In fact, for the first time, we have mathematically derived a closed-form formula for this problem and built the founding stones for a new area of research in flexible frequency-dependent beamforming.

As an extension to mmFlexible, we built a hardware prototype for the new antenna system as shown in Figure 0.1(b) and explored new applications in concurrent sensing and communication. My work FlexLink [Preprint 2024] achieves this by decoupling the data and control beams in mmWave 5G NR networks using our delay phased array architecture and demonstrating it on a hardware platform.

0.1.2 Wireless Sensing for NextG Networks

Radio-Frequency-based sensing: As Wi-Fi radios and NextG cellular networks are ubiquitously deployed in indoor and outdoor environments, these radios can also provide accurate sensing (e.g., positioning and tracking) of mobile devices in the event that traditional sensors such as GPS fail, e.g., under obfuscation by buildings and foliage. As part of a collaborative project, we demonstrated LocAP [NSDI 2020], a radio-frequency-based navigation system that is accurate and does not require prior knowledge about the WiFi infrastructure. To this end, we deployed an autonomous bot, loaded with a WiFi sensor, to estimate the location, orientation, and antenna spacing of all WiFi access points in that environment. Here, I contributed to conceptualizing and simulating the core algorithm that enabled millimeter-accurate localization of access point antennas.

Radar sensing-driven communication: The millimeter-wave network for mobile access requires tracking of mobile users to maintain directional links with them. Notably, 5G New Radio employs periodic beam scanning to estimate beam directions with mobile devices. However, this approach suffers from a trade-off between accuracy and time overhead. Achieving precise beams comes at the cost of increased overhead, manifested in prolonged beam scanning

duration and higher periodicity of the 5G beam scan procedure, which can reach as high as 25% overhead (5 ms beam scan duration at a 20 ms interval). The more time allocated to beam scanning, the less time remains (only 75%) for conveying valuable information bits to user devices.

In CommRad [Preprint 2023], we integrated a mono-static radar sensor with mmWave radio to sense and track mobile users as shown in Figure 0.1(c). A key challenge with mono-static radar sensing is that radar sees many objects in its view, but it lacks context, i.e., it cannot distinguish whether the objects represent communicating users or non-communicating entities. Similarly, it cannot differentiate between reflectors that can assist or obstruct communication links. My solution in CommRad is inspired by the realization that this contextual information can be obtained by the communicating radio itself. We used the communication waveform to obtain the context in the form of the user and reflector identification in a cluttered environment and integrated it with the radar tracking system. This combined radar-radio framework, dubbed collaborative learning, helps provide the context to the radar sensor in terms of initial user and reflector location, and in return, the radar aids in user tracking and sustains mobile links even in obstructed scenarios, resulting in robust directional links for all mobile users while providing a ten-fold reduction in the radio beam scanning overhead.

0.1.3 Security in NextG Communication and Sensing

Indistinguishable spoofing attacks on automotive sensing radars: Modern automotive vehicles are equipped with radar, which provides accurate all-weather perception to detect surrounding objects for safer driving. But are these radars safe? We built mmSpoof [S&P 2023], a spoofing system that exploits the radar FMCW waveform to create an 'illusion' of fake objects in the radar view which looks completely realistic to the radar. We achieve this by developing a reflect array as a spoofing device, which receives the radar signal and reflects back with a configurable frequency shift that changes the delay and Doppler profile at the radar, thus creating a fake object in the radar view. We demonstrated a 99% spoofing accuracy

through a real-time outdoor testbed as shown in Figure 0.1(d) with commercial radars. Our system is capable of working with any commercial radar without any prior knowledge of FMCW parameters on the radar, such as chirp duration and slope. These parameters are estimated in real-time using novel estimation algorithms we developed and deployed at the reflect array. We also discussed several anti-spoofing measures, such as multi-radar fusion and designing advanced radar waveforms with frequency-hopping. This is a joint work with Rohith Reddy Vennam as the lead author. My main contributions were formulating the mathematical framework of distance and velocity spoofing, building an mmWave array and USRP-based testing framework, and proofreading the paper while supervising Rohith.

Jammer monitoring and anti-jamming in NextG networks: Jamming of NextG cellular networks is a big concern because of interference by nearby base stations in dense small-cell settings and intentional jamming using cheap SDRs like HackRF. Yet, there is a lack of research platforms to understand and provide real-time insight into jammer monitoring in operational RANs. Therefore, we built BeamArmor [Hotmobile 2024, Milcom 2023 Demo], a RAN controller platform built with the srsRAN software stack to monitor, locate, and mitigate jammer activities using MIMO null-steering techniques, demonstrating bit error rate and SNR improvements by over 15 dB. Broadly, this controller platform can demonstrate many other MIMO-based applications in cellular networks, such as positioning and channel prediction in real time.

0.2 Contributions

While building these systems, this thesis makes the following important contributions:

- mmReliable establishes the concept of constructive multi-beamforming in mmWave networks. It is the first demonstration of a counter-intuitive principle that constructive multi-beams can achieve higher throughput than single-beam systems while using standard-compliant phased arrays.

- mmFlexible introduces a new kind of beamforming called frequency-dependent multi-beams. It unlocked a new fundamental dimension of frequency into analog beamforming and demonstrated them via a novel delay-phased array antenna architecture.
- FlexLink is a novel application of frequency-dependent multi-beams in decoupling control and data beams to achieve high throughput and increase spectrum utilization. This is the first work to develop a solid mathematical foundation of frequency-dependent multi-beams with closed-form expressions and the first system prototype demonstration over the air.
- CommRad introduces a novel bi-directional learning framework where the mono-static capabilities of a radar sensor and bi-static capabilities of a radio sensor are combined to deliver context-aware accurate sensing and localization.
- BeamArmor and mmSpoof expose security vulnerabilities in communication and sensing systems and build end-end testbeds to demonstrate and counter them.

0.3 Impact beyond this thesis

The interdisciplinary approach of connecting theory and systems is reflected in various publications across prominent venues, including IEEE/MDPI Journals and Top-Tier ACM systems conferences, and resulted in six patent submissions, with three already granted. Besides this, many systems in this thesis have also caught attention in the popular press, including Interesting Engineering [73], TechXplore [74], InceptiveMind [148], and Science Daily [65], among others. The impact extends to various platforms and awards. I have been selected for the Marconi Society Scholar in Residence program, contributed to the Escribamos Ciencia program for K12 students, earned the best poster runner-up award at Hotmobile'23, and received the best 3-minute thesis presentation award at ACM Mobisys'20, Mobicom'21, and Mobicom'22. My open-sourced datasets and code have been used by multiple universities, including UCLA, NYU, and Texas A&M University.

Furthermore, some of the work in this thesis has achieved commercial impact. For instance, our work on Open-RAN-driven anti-jamming and interference management is deployed at VMware and has secured a grant of $240,000 to further advance this research. Additionally, our work on sensing-driven communication has received support through the Qualcomm Innovation Fellowship ($100,000).

0.4 Organization

The thesis is organized into six chapters with detailed mathematical, algorithmic, and hardware details of each system described above. Chapter 1 discusses how we achieve reliable and high throughput mmWave links through constructive combining using traditional phased arrays. It also delves into a new delay-based architecture called delay-phased array that achieves wideband constructive combining capabilities in strong multi-path environments. We then explore new use cases for the delay-phased array architecture in Chapter 2 and Chapter 3 with a focus on two different applications: (i) multi-user frequency multiplexing to improve latency and scalability and (ii) decoupling control and data beams to improve overall system throughput, respectively. Chapter 4 discusses novel applications of sensing-driven communication by using a radar sensor to aid a mmWave radio perform efficient beam management and blockage mitigation using a bi-directional learning framework. It is followed by Chapter 5 which discusses the security aspects of communication and sensing systems such as spoofing and jamming attacks and the building of an end-end O-RAN based real-time cellular platform to address these attacks. We finally conclude the thesis with some concluding remarks and discuss future research directions in Chapter 6.

Chapter 1

mmReliable: Reliable and High Throughput mmWave Links

5G New Radio (NR) is expected to support cutting-edge applications such as vehicular (V2X), factory automation, autonomous driving, and remote surgery [27, 226]. A key requirement for such applications beyond the high data rate is exceptionally high reliability, defined as the fraction of time during which the link does not suffer an outage [115, 107]. As such, 5G NR utilizes millimeter-wave (mmWave, FR2) frequencies due to abundant bandwidth (400MHz-2GHz) that can provide a high data rate compared to traditional sub-6 GHz FR1 (20-100MHz bandwidth) bands [180]. However, unlike its sub-6 GHz counterparts, mmWave uses directional beams to reduce the impact of higher path loss at mmWave frequencies and is highly susceptible to blockage and user movement. For instance, a measurement study of 28 GHz network deployment in Chicago [154] shows that blockage by humans always causes link outages, which is detrimental to applications requiring high-reliability [115].

A large fraction of existing work focuses on improving mmWave link establishment [114, 44, 207, 30, 98], throughput [208, 255, 165, 227] and coverage [252, 230]; while the link reliability requirement takes the backseat. The prevailing methods aim to reduce link outages by performing *beam-training*, a process where the mmWave base station and client tediously scan multiple beams to establish a directional link. Reducing beam-training time does not prevent outages as the process is usually reactive [259, 114, 43, 256, 44, 207, 30, 98] and kicks in only

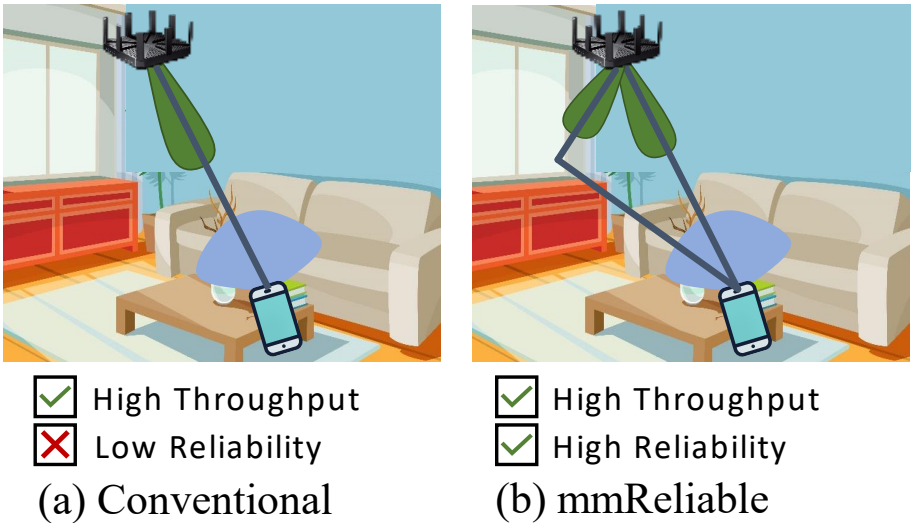

Figure 1.1. mmReliable idea: mmReliable utilizes beams with multiple lobes (*Multi Beam*) to provide high reliability compared to a single-lobe beams.

after a significant degradation in link quality. To improve reliability, some authors propose proactive approaches that constantly track the client using side-channel information from out-of-band LTE/WiFi [205, 159] or external sensors like GPS, radar or lidar [95, 229, 218, 231]. These solutions are not self-reliant, have limited accuracy, and possibly require sensitive information like user location [229, 218]; making the system dependent on external information and difficult to deploy with existing standards. In contrast, we aim to create the ideal mmWave link: one that can be created and maintained with high reliability (no link outage), supports high data rates, and is easy to integrate into current standards.

In this chapter, we introduce mmReliable, a system that achieves both reliable and high-data-rate mmWave links while being protocol-compliant. Inspired by sub-6 GHz communication, mmReliable delivers on reliability and throughput by exploiting the multipath diversity in environmental reflections. These reflections are typically strong for mmWave and can sustain the link even if the direct path is occluded or unavailable to establish an independent link. We observe that conventional single-beam links that utilize the reflectors are still unreliable because they are prone to outage due to blockage or misalignment. mmReliable uses a single radio frequency (RF) chain and phased array to create custom beam patterns with multiple directional beams

(*multi-beam*) that are each aligned towards direct and reflected paths as shown in Figure 4.1(b). The idea here is that the probability of multiple beams simultaneously facing a blockage is low, reducing the possibility of link outage and increasing reliability.

An obvious implication of splitting a single beam into a multi-beam is reduced power radiated per beam to conserve the Total Radiated Power (TRP) for FCC compliance [75]. The reduction in per-beam power could reduce the received signal strength and lead to throughput loss at the receiver, i.e., trading off throughput for reliability. It turns out we can break this trade-off: we observe that the received signal is the sum of multiple copies of the transmitted signal, one from every beam in a multi-beam as shown in Figure 4.1(b). If we ensure constructive addition of these copies at the receiver (Section 1.2), the received power and throughput can be increased to a value higher than that of a single beam while keeping TRP constant.

Let's take an intuitive example of how multi-beam improves throughput. Consider a mmWave channel consisting of two paths, and two beams are aligned along those two-channel paths. For simplicity, assume the path loss is unity for each path, and the signal traveled along each path incurs the same phase. mmReliable's transmitter splits the total power of a^2 equally into two beams with power $\frac{a^2}{2}$ on each beam, i.e., amplitude $\frac{a}{\sqrt{2}}$ on each beam, which propagates through the channel and is received at the Omni-receiver. Since the phase aligns for both paths, the received signal adds in the signal domain, i.e. total amplitude of the received signal $\frac{a}{\sqrt{2}} + \frac{a}{\sqrt{2}} = \sqrt{2}a$. The received signal upon converting to SNR (signal-to-noise ratio) is proportional to $2a^2$. In contrast, a traditional mmWave single-beam link does not exploit the second path and would provide an SNR proportional to a^2, which is 3 dB lower compared to the multi-beam case. Higher SNR results in lower bit error rates and allows higher modulation and coding schemes, leading to higher throughput for multi-beam. In Section 1.2, we go beyond this simple example to consider practical channels and show that, if the beam directions, phase, and power per beam of multi-beam are chosen optimally, then multi-beam will always achieve SNR higher than single beam *even in wireless channels with weak multipath while using the exact total radiated power*. A multi-beam created using optimal parameters is hereafter referred to as

constructive multi-beam.

As a first, we develop an algorithm to establish a constructive multi-beam link with low-overhead and protocol compliance. To achieve constructive multi-beam, one needs to estimate the optimal beam directions, beam-phase, and power per beam. Our key insight here is to decouple this problem into two: finding beam directions and estimating the per-beam power and phase. mmReliable first learns the directions of strong paths in the environment (i.e., paths viable for communication) during the mandatory beam-training phase. This could be done using exhaustive beam-scanning or any other improved algorithm [114, 44, 87, 118]. Then, with just two additional channel probes per viable beam, it learns the optimal power and phase to be applied to each beam in the multi-beam. Only two to three viable beams exist in typical environments due to the sparse nature of reflection clusters at mmWave [259, 213]; therefore, mmReliable's algorithmic overhead remain fixed and independent of the number of elements in the antenna array.

Once the constructive multi-beam is established, the challenge is in continuously maintaining it, as the wireless channel changes due to blockages and mobility. Link blockages could completely occlude one or more beams of a multi-beam. Mobility adversely affects the established links by causing misalignment between transmit and receive beams, degrading both throughput and reliability. Multi-beam management design has to address some uniquely challenging aspects. In contrast to a single-beam system where only one beam is misaligned, in a multi-beam system, all the beams get misaligned simultaneously due to mobility. In addition, using a phased array system with a single RF chain means that signals from multiple beam directions are always superposed into one. The superposition makes it difficult to assess the effects of blockage and misalignment.

We develop a beam-maintenance algorithm, which proactively re-aligns each of the beam parameters (direction, phase, and amplitude) in a multi-beam and continuously maintains its constructive nature. First, mmReliable uses a super-resolution algorithm to tease apart the properties of individual beams from their sum. Next, it uses the estimated per-beam properties to

measure power loss and infer the underlying cause, i.e., blockage or mobility. Instead of waiting for beams to degrade, mmReliable proactively makes this inference and optimizes per-beam properties, and maintains constructive multi-beam. mmReliable uses low-overhead standard-compliant channel probes that are embedded in the communication waveform to achieve the elusive goal of creating and maintaining reliable high throughput mmWave links, free from sporadic outages.

We build a first-of-its-kind software-defined 5G NR testbed on 28 GHz mmWave bands with 64 elements phased array that supports channels of up to 400 MHz [113]. We implement mmReliable using 5G NR-compliant reference signals. We perform experiments at 28 GHz in our university area by deploying mmReliable in various indoor and outdoor scenarios with link distances of up to 80 m. Through empirical measurements, we find that indoor walls and large outdoor buildings are powerful reflectors; the median attenuation of the reflected path with respect to the direct path in outdoor scenarios is only 5 dB. Our evaluation shows that in scenarios where user movement and blockage co-occur, mmReliable has reliability close to 100% while maintaining an average throughput of 1.5 bits/sec/Hz as compared to the single-beam reactive baseline, which has a reliability of 65% and on-average throughput of 1 bits/sec/Hz. The throughput-reliability product is improved by 2.3$\times$ compared to the best reactive baseline. The artifacts for mmReliable are available online[1].

1.1 Background and Motivation

Directional beamforming in mmWave creates the need for beam-management schemes for link creation and maintenance. Beam-management can be divided into two broad functions: the first is **beam-training**, and involves the search and creation of a beam [259, 114, 43, 256, 44, 207, 30, 98]. The second is **beam-maintenance** and involves the upkeep of the beam despite environmental factors such as blockages, mobility, and fading. Even if multi-beam is resilient to blockages, it is not practical unless it can be efficiently created and maintained

[1]Artifacts link: https://wcsng.ucsd.edu/mmreliable

under the constraints of existing standards. The efficacy of the beam-maintenance function is especially critical to enabling reliable mmWave links with low overhead. Protocols such as 5G-NR and IEEE 802.11ad contain exclusive signaling and control provisions for such beam-management [27, 2, 89].

1.1.1 mmReliable: Making Multi-beam Practical

While single-beams have been the default choice for communication, there are some recent interests in using multi-beams to enable auxiliary mmWave capabilities. Multi-beams have been shown to aid with the beam-training process [98, 39, 40, 76]. In [98], the authors employ multi-beams for fast beam-training, but use them to set up single-beam links for communication. In [39], the authors propose a beam-training method to create multi-beams for blockage resilience. However, they do not derive or discuss that a multi-beam link is better than a single-beam link for throughput. In addition, all of [98, 39, 40, 76] depend on repeated reactive beam-training to re-establish the link when it is blocked or misaligned, leading to intermittent outages. In contrast, we propose multi-beam as the new standard for beamforming in mmWave communication systems. We show that multi-beams are reliable while providing superior throughput to single-beams in all scenarios. As a first, we develop an efficient algorithm to estimate parameters and create multi-beams in state-of-the-art phased arrays. Our algorithm is protocol-compatible and can be deployed over any beam-training scheme.

1.1.2 Beam Maintenance for Reliable mmWave

Beam maintenance is important to reduce the frequency and overhead of repeated beam-training as shown in Figure 1.2. For instance, in the context of 5G-NR, a beam-training phase could take up to 5 ms to probe 64 beam directions with a default periodicity of 20 ms, leading to a 25% overhead [89]. Even if a more efficient training algorithm from [98, 39, 40, 76, 79] is used, the overhead will make beam-management solely based on beam-training intractable. The high overhead of beam-training can be avoided by carefully utilizing the intrinsic features of mmWave protocols. They use known reference signals interspersed with data communication to

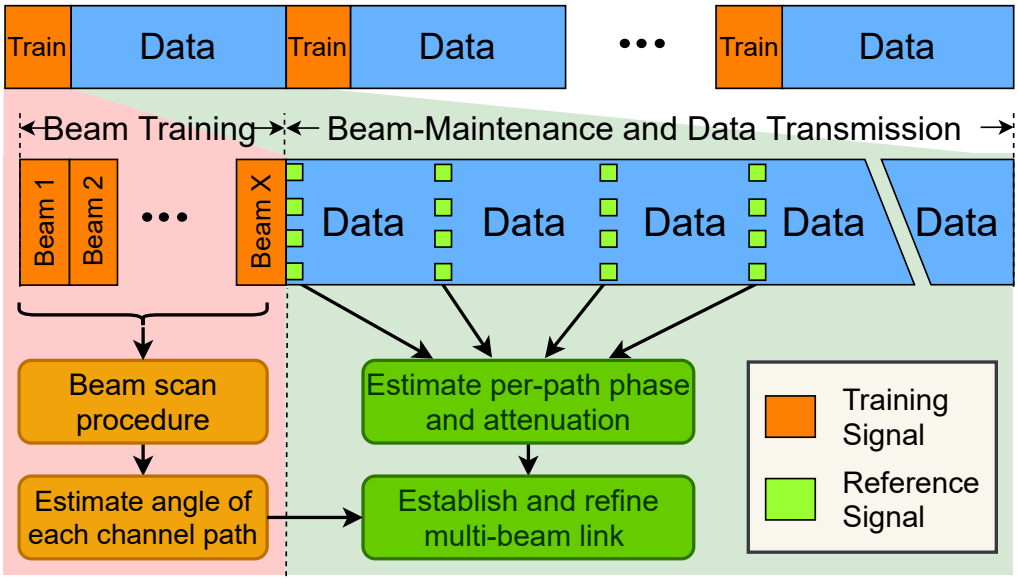

Figure 1.2. Building blocks of mmReliable: Beam management that leverages initial beam training for angle estimation and reference signaling for multi-beam establishment and maintenance.

perform beam-maintenance and refinement (See Figure 1.2). These reference signals provide channel measurements with an option to set any desired beam. Due to the sparsity of these reference signals, they cannot be used for a full-fledged beam-training, and only minor changes to the beamforming vector can be made using a few channel estimates. A low-overhead beam-maintenance scheme is, therefore, a must for reliable mmWave links.

While beam-maintenance for single beam links has been studied by some authors [38, 147, 113], there is no work that addresses multi-beam links. Due to this gap, the utility of multi-beam for communication is lessened despite its benefits. We are the first to propose a low overhead beam-maintenance scheme using multi-beam. Our approach periodically updates the per-beam angle, phase, and amplitude with negligible overhead, requiring only three channel estimates for a 2-beam multi-beam and five estimates for a 3-beam multi-beam. The entire process is completed within 0.6 ms for the latter, making multi-beam maintenance possible with low overhead. With this, our work makes a compelling case for multi-beam and addresses the end-to-end challenges in their creation and maintenance.

1.2 Design of Multi-beam System

This section describes how mmReliable creates and leverages multi-beam for reliable, high-throughput mmWave communication. We first show how links can be made reliable to blockage by using multiple beams at once. Then, we derive mathematically optimal beamforming weights to maximize the signal-to-noise ratio (SNR) at the receiver. We show that optimal beamforming is nothing but constructive multi-beam for most typical mmWave channels.

We make an interesting observation that constructive multi-beam is channel-dependent and depends on the number of paths, their angles, relative phase, and attenuation. Naturally, to implement constructive multi-beam, we need to estimate these parameters. We propose a two-step procedure to estimate these parameters using the beam-management framework shown in Figure 1.2. We leverage any standard beam-training procedure to calculate the angle of each path [87, 118]. Then, we design a low-overhead method using reference signals to estimate the per-path phase and attenuation. Finally, we use these estimates to set up and maintain the optimal constructive multi-beam despite the adversity of blockage and user mobility.

1.2.1 Multi-beam Links are More Reliable

To compare various techniques on equal footing, we formally define link reliability as the fraction of time when the link is available for communication within a large observation interval. Link outage (due to natural effects) or procedures like beam-training reduce the reliability as they temporarily render the link unavailable for communication. Therefore, we can express reliability as:

$$\text{Reliability} = 1 - \text{Prob}(\text{Outage}) \tag{1.1}$$

We empirically compute the probability of outage as the fraction of duration where the SNR is below a minimum threshold. The directional nature of single-beam mmWave links makes them susceptible to link outages since blockage & user mobility effects can reduce SNR by up to 30 dB [141] (Figure 1.18). Multi-beams prevent these outages by avoiding a single point of failure.

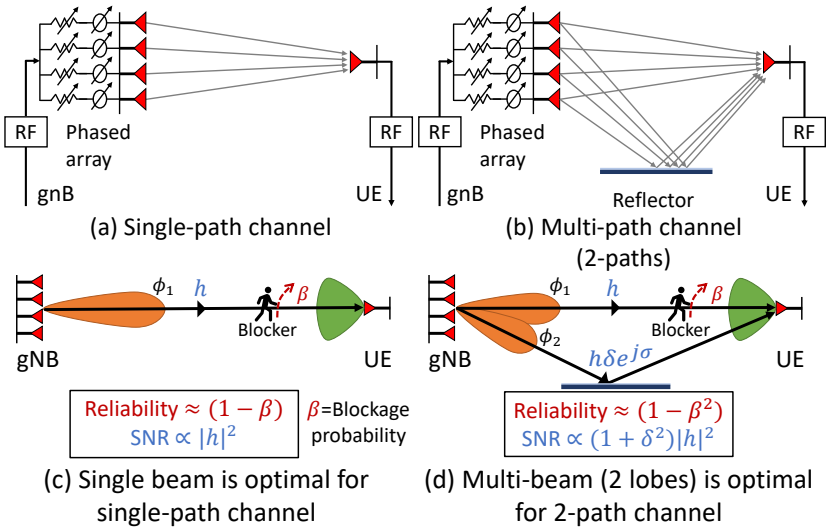

Figure 1.3. Constructive multi-beam: Enhances the throughput & reliability of mmWave link. (a) System with a single-path channel, (b) With multipath (2-path) channel. (c) The optimal beamforming technique for a single-path channel is a single beam with all the energy directed along that channel path. (d) A constructive multi-beam with two lobes is more reliable and provides higher SNR compared to a single beam.

Multi-beam links are reliable even under the impact of blockage events. Consider a blockage probability β $(0 \leq \beta \leq 1)$, which represents the fraction of time the link is in outage during the observation interval. For simplicity, we assume the beams in our multi-beam are blocked independently. Thus, the probability that k beams simultaneously experience an outage is β^k. By definition, the reliability would be $1-\beta$ for the single-beam case and $1-\beta^k, k \geq 2$ for the multi-beam case. Multi-beam provides higher reliability because it prevents the link from suffering an outage due to blockage, unlike its single-beam counterparts [208, 141]. In practice, blockages may be correlated, and scenarios could arise where two or more paths are blocked simultaneously. While no solution can prevent link outage if all paths are blocked (e.g., the user inadvertently blocks all antennas with their body), multi-beam links can be sustained as long as there is at least one unblocked path to the receiver. Simultaneous blockage of multiple paths is always less probable than blockage of a single path [208]; therefore, multi-beam links are *always* more reliable than single-beam links.

1.2.2 Multi-beam Throughput Model

While multi-beams are intuitively resilient to blockage by virtue of having multiple beams separated in angle, each beam's power is reduced since the total power is split amongst multiple beams. One may wonder whether this leads to reduced throughput since it could mean less power is incident on the receiver. But this is not the case, and it can be shown that a multi-beam link with optimal constructive combining and power-control can, in fact, provide higher throughput than a single-beam link. To prove this, we introduce a model for a multi-antenna base station transmitting to the user device using a mmWave link and derive the optimal beamforming solution that maximizes the received SNR. We make two important observations. First, we show that a single-beam system is optimal only when the wireless channel consists of a single path from the transmitter to the receiver. A single-beam system is sub-optimal for a multipath channel, which is generally the case even for mmWave links [203]. We then show that a multi-beam with two beams is optimal for a two-path channel and maximizes the received SNR. We conclude with a note on generalizing our formulation to an arbitrary number of beams in multi-beam.

■ **Primer on optimal beamforming:** We derive an expression for optimal beamforming weights at a multi-antenna base-station (gNB phased array) communicating with a single antenna receiver using mmWave (Multi-antenna receiver discussed in Section 1.3.4). The gNB uses a uniform linear phased array with N antenna elements. Beamforming is implemented by applying a $N \times 1$ beam weights vector $\mathbf{w}$ at each antenna using a combination of phase shifters and attenuators (Figure 1.3(a)). Our goal is to determine the optimal $\mathbf{w}$ that maximizes the received SNR.

Let s be the transmit signal from the gNB, the received signal y is expressed as:

$$y = \mathbf{h}^T \mathbf{w} s + \eta, \tag{1.2}$$

where $\mathbf{h}$ is $N \times 1$ channel from N transmit antennas to one receive antenna and η is white Gaussian noise at the receiver. The capacity or maximum throughput of a wireless link is

evaluated in terms of SNR (the ratio of signal power to noise power) as:

$$\text{SNR} = ||\mathbf{h}^T\mathbf{w}||^2 P_s/P_\eta, \tag{1.3}$$

where P_s is the average transmit power (without transmit array gain) and P_η is the noise power[2].

Our goal is to estimate beamforming weights $\mathbf{w}$, that maximize the SNR. Using Cauchy-Schwartz inequality, it follows that $||\mathbf{h}^T\mathbf{w}||$ is maximized when $\mathbf{h}^*$ (complex conjugate of $\mathbf{h}$) and $\mathbf{w}$ align in vector space [216]. Therefore, *the optimal weight vector is channel-dependent*. Intuitively, the weight vector cancels the phases in the channel and creates an inner-product form that maximizes the absolute value of their product. Additionally, the weights need to be unitary to keep the power constant, i.e., $\|\mathbf{w}\| = 1$. This way, we get the optimal weights $\mathbf{w}^{\text{opt}}$ as:

$$\mathbf{w}^{\text{opt}} = \mathbf{h}^*/||\mathbf{h}||, \tag{1.4}$$

The optimal beamforming provides the highest SNR of $\frac{||\mathbf{h}||^2 P_s}{P_\eta}$. The practical wireless channel—even mmWave channel—consists of multiple paths (direct or reflected paths). Since the optimal beamforming vector depends on the channel $\mathbf{h}$, it is affected by the nature of these paths.

■ **Single-beam is optimal for a single-path channel:** Extending our discussion, it is natural to wonder how to provide a structure to the general nature of optimal beamforming vectors ($\mathbf{w}^{\text{opt}}$). Consider the simplest of wireless channels with a single path from the transmitter to the receiver. It turns out that the single beam used by conventional mmWave systems is optimal for such a channel. To see this, we represent the single-path channel vector by $\mathbf{h}^{\text{single}}$ and derive the optimal weights. A single-path channel is entirely defined by two parameters: the direction of departure ϕ_1 and complex attenuation h:

$$\mathbf{h}^{\text{single}}[n] = he^{-j2\pi\frac{d}{\lambda}(n-1)\sin(\phi_1)}, \tag{1.5}$$

[2]To derive the expression for SNR: $E[yy^H] = E[(\mathbf{h}^H\mathbf{w}s+n)(\mathbf{h}^H\mathbf{w}s+n)^H]] = \|\mathbf{h}^H\mathbf{w}\|^2 E[ss^H] + E[nn^H]$ and we denote $P_s = E[ss^H]$ and $P_\eta = E[nn^H]$.

where n is the transmit-antenna index, d is the antenna spacing, and λ is the wavelength of carrier frequency ($d = \frac{\lambda}{2}$ in our phased array). Using (1.4) and (1.5), we obtain the optimal weights for this channel as $\mathbf{w}_{\phi_1} = \frac{\mathbf{h}^{\text{single}^*}}{\|\mathbf{h}^{\text{single}}\|}$ which can be simplified to:

$$\mathbf{w}_{\phi_1} = \frac{1}{\sqrt{N}}[1, e^{j2\pi\frac{d}{\lambda}\sin(\phi_1)}, \ldots, e^{j2\pi(N-1)\frac{d}{\lambda}\sin(\phi_1)}]^T, \tag{1.6}$$

which is the familiar single-beam weight vector [208, 255, 207, 251].

■ **Constructive multi-beam is optimal for multipath channel:** Since the optimal beam is channel dependent, let us observe how it changes when we introduce a reflector into the channel as shown in Figure 1.3(b). We continue to use h for the complex attenuation of the first path. The second path's attenuation is expressed as $h\delta e^{j\sigma}$, where $\delta \in \mathbb{R}^+$ and $\sigma \in [0, 2\pi]$ are respectively the relative attenuation and phase shift of the second path with respect to the first. Now, we can write the expression of the two-path channel as:

$$\mathbf{h}^{\text{multi}}[n] = he^{-j2\pi\frac{d}{\lambda}(n-1)\sin(\phi_1)} + h\delta e^{j\sigma}e^{-j2\pi\frac{d}{\lambda}(n-1)\sin(\phi_2)} \tag{1.7}$$

where ϕ_1 and ϕ_2 are the respective directions of departure of the two paths. If we naively use the single-beam weights from (1.6) here, it ignores the second path and severely attenuates power along ϕ_2 (Figure 1.3(c)) resulting in an approximate SNR of:

$$\text{SNR}^{\text{single}} = ||(\mathbf{h}^{\text{single}})^T\mathbf{w}_{\phi_1}||^2 P_s/P_\eta \approx |h|^2 P_s/P_\eta. \tag{1.8}$$

In contrast, the optimal weights for the multipath channel are obtained from (1.4) as $\mathbf{w}^{\text{multi}} = \mathbf{h}^{\text{multi}^*}/\|\mathbf{h}^{\text{multi}}\|$ which is visualized in Figure 1.3(d). From Eq. (1.3), the SNR of the optimal weights $\mathbf{w}^{\text{multi}}$ is:

$$\text{SNR}^{\text{multi}} = \|h^{\text{multi}}\|^2 P_s/P_\eta \approx (1+\delta^2)|h|^2 P_s/P_\eta. \tag{1.9}$$

In comparison to the naive single-beam approach, the optimal SNR is higher by a factor of

$1+\delta^2$, which provides 2x gain (3dB higher SNR) for two equally strong paths ($\delta = 1$). In general, $\mathbf{w}^{\text{multi}}$ is the optimal constructive multi-beam weight that splits the transmit signal along both paths such that they maximally combine at the receiver.

■ **Constructive Multi-beam is optimal for SNR and throughput for general k-path channel:** Here we show that the multi-beam patterns with appropriate power and phase control are optimal such that they maximize the SNR at the receiver. To demonstrate the multi-beam optimality, we first construct the maximum likelihood (ML) based solution and then show that the multi-beam approaches the ML estimate. Recall the optimal beamforming weights for a general multi-path channel $\mathbf{h}$ is given by:

$$\mathbf{w}^{\text{opt}} = \frac{\mathbf{h}^*}{||\mathbf{h}||}, \tag{1.10}$$

which provides the maximum SNR of $\frac{||\mathbf{h}||^2 P_s}{P_\eta}$. Note that the optimal beamforming weights is dependent on the multi-path channel taps.

Our goal is to show how it compares against the SNR of a single beam and our multi-beam system. We first write the channel in a geometric model representation [33, 191]:

$$\mathbf{h}(t) = \sum_{\ell=0}^{L-1} \mathbf{a}(\phi_\ell)\gamma_\ell e^{2\pi f_c \tau_\ell}\delta(t-\tau_\ell), \tag{1.11}$$

where we assume the channel is bounded by L paths and $\phi_\ell, \gamma_\ell, \tau_\ell$ represents the AoD, attenuation, and ToF of the ℓ^{th} path, respectively. The steering vector $\mathbf{a}(\phi_\ell)$ represents the relative phase shift introduced by each path at every antenna element (out of N elements) as

$$\mathbf{a}(\phi_\ell) = [1, e^{-j2\pi\frac{d}{\lambda}\sin(\phi_\ell)}, \ldots, e^{-j2\pi(N-1)\frac{d}{\lambda}\sin(\phi_\ell)}]^T$$

In the frequency domain, the channel is written at subcarrier k as follows:

$$\mathbf{h}(k) = \sum_{\ell=0}^{L-1} \mathbf{a}(\phi_\ell)\gamma_\ell e^{j2\pi(f_c+k\Delta f)\tau_\ell} \tag{1.12}$$

where f_c is center frequency and Δf is the subcarrier spacing.

The traditional single-beam pattern ignores the multi-path and creates a directional pattern along the strongest path (LOS or one of the reflectors). Assuming ϕ_0 is the direction of the strongest path, the single-beam antenna weights follows:

$$\begin{aligned} \mathbf{w}^{\text{single}} &= \mathbf{a}^*(\phi_0)/||\mathbf{a}(\phi_0)|| \\ &= \frac{1}{\sqrt{N}}[1, e^{j2\pi\frac{d}{\lambda}\sin(\phi_0)}, \ldots, e^{j2\pi(N-1)\frac{d}{\lambda})\sin(\phi_0)}]^T. \end{aligned} \tag{1.13}$$

Assuming a directional single-beam pattern severely attenuates other channel paths, the single-beam SNR consists only of the strongest path as:

$$\text{SNR}^{\text{single}} = \frac{||\gamma_0||^2 P_s}{P_\eta}. \tag{1.14}$$

Instead, we design phase-coherent multi-beam patterns which exploit B channel directions out of L paths as follows:

$$\mathbf{w}^{\text{multi}} = \frac{\sum_{b=0}^{B-1} \mathbf{a}^*(\phi_b)\gamma_b e^{j\sigma_b}}{||\sum_{b=0}^{B-1} \mathbf{a}^*(\phi_b)\gamma_b e^{j\sigma_b}||}, \tag{1.15}$$

where we orient all the B beams in the multi-beam along the B strongest paths in the channel. Note that our multi-beam patterns ensure that per-beam amplitude γ_b and phase σ_b (averaged over all frequency subcarriers) are aligned with per-path channel attenuation and phase respectively. We clearly see that when $B = L$, i.e., the multi-beam patterns are oriented along all the paths in a multi-path channel. In this case, the multi-beam weights are the same as the optimal weights in (1.10):

$$\mathbf{w}^{\text{multi}} = \mathbf{w}^{\text{opt}}, \quad \text{when } B = L. \tag{1.16}$$

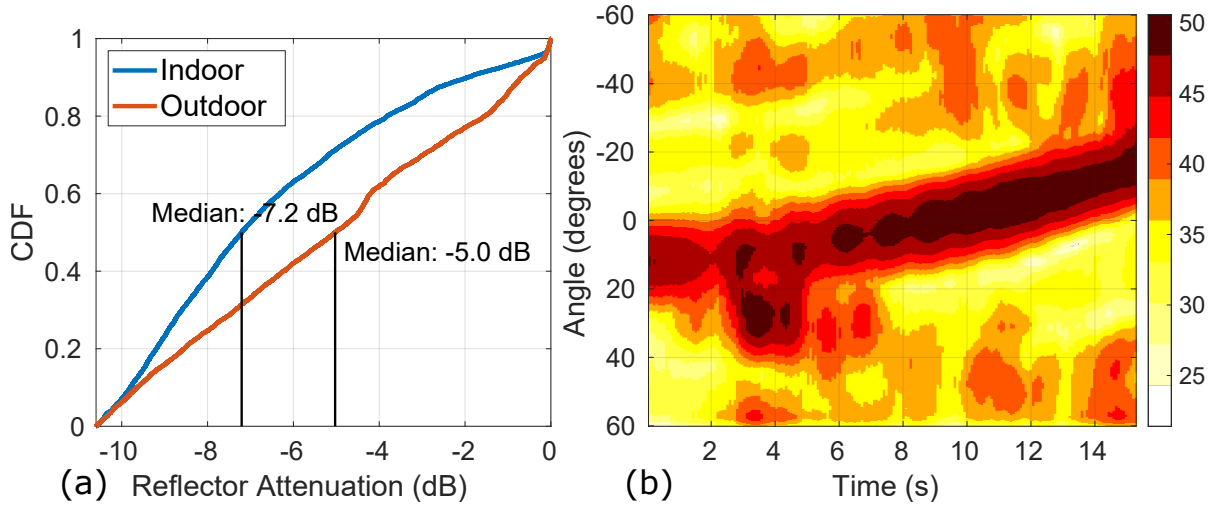

Figure 1.4. Results on Reflector Attenuation: (a) CDF of the relative attenuation of the strongest reflected path compared to the direct path at over various locations. (b) A heatmap of strong paths in the environment while the UE moves, strong reflectors appear at different points in time.

Thus, a multi-beam pattern utilizes B out of L channel paths to improve the SNR as follows:

$$\text{SNR}^{\text{multi}} = \frac{\sum_{b=0}^{B-1} ||\gamma_b||^2 P_s}{P_\eta}, \tag{1.17}$$

which converges to optimal SNR for $B = L$. We can write the capacity of multibeam-link as

$$C_{mb} = \log_2(1 + \text{SNR}^{\text{multi}}) \tag{1.18}$$

The optimal ML beamformer requires per-antenna channel estimate which has a high complexity in terms of beam probing overhead [165]. *Since mmWave channel is sparse and there are only 1 or 2 strong reflected paths in addition to the direct path, the multi-beam with 2-3 beams provides SNR gain comparable to the optimal beam with significantly lower overhead.*

■ **Strength of mmWave multipath:** At this point, one may ask if multiple paths exist in practical mmWave channels, and if so, how strong are they? There are numerous measurement studies that indicate the presence of strong multipath at both 28 GHz and 60 GHz [180, 259, 213, 158, 233, 45]. These studies show that in a typical deployment, reflectors like metals, concrete

walls, and tinted glass strongly reflect signals while attenuating them by about 5-6 dB ($\delta \approx 0.5$). In addition, we perform a measurement study of reflector strengths at various indoor (5m-10m link) and outdoor (10m-80m link) locations (Overall 10K data points). At each location, we perform a full 120^o scan to identify the distribution of signal strength across space. Shown in Figure 1.4(a), common reflectors cause 1-10 dB attenuation (relative to the direct path) with a median attenuation of 7.2 dB indoors and 5 dB outdoors. Figure 1.4(b) shows a particular data point where signal strengths along all angles are evaluated during user motion. Sometimes, reflectors as strong as the direct path exist and can be utilized to establish a multi-beam link.

1.2.3 Creating Constructive Multi-beam

So far, we have established that constructive multi-beam is the optimal beamforming technique for a typical mmWave channel. To achieve optimal beamforming, we need to estimate the channel $\mathbf{h}[n]$ at each antenna at the gNB, the complexity of which scales with the total number of antennas. In sub-6 GHz communication, the channel at each antenna is readily measured as each antenna could be connected to a separate RF chain. In contrast, mmWave hardware usually consists of a single RF chain attached to a large phased array with many antennas. Performing channel estimations for each antenna element is intractable in practical mmWave systems.

The intuition behind constructive multi-beam is that it optimally uses the available multipath. We make the key observation that in mmWave, the number of multipath reflections is usually sparse [259, 213]; therefore, there must be a simpler way to measure $h[n]$. Indeed, from (1.7), we see that the channel can be linearly decomposed into its constituent single-beam directions. Said differently: we can reconstruct $h[n]$ using channel measurements on individual multipath directions in the environment. Once we have $h[n]$, the optimal weights follow using (1.4).

Consider our example from the previous sub-section, where two paths were present in the environment at angles ϕ_1, ϕ_2. If the relative attenuation δ and the relative phase σ between

each path is measured, we can create the constructive multi-beam weights by simple addition:

$$\mathbf{w}(\phi_1, \phi_2, \delta, \sigma) = \frac{(\mathbf{w}_{\phi_1} + \delta e^{-j\sigma} \mathbf{w}_{\phi_2})}{\|(\mathbf{w}_{\phi_1} + \delta e^{-j\sigma} \mathbf{w}_{\phi_2})\|}, \tag{1.19}$$

The denominator ensures that the TRP is conserved by enforcing $\|\mathbf{w}\| = 1$. We then appropriately quantize the beamforming vector phases & amplitudes to be compatible with our phased array. In [39], a sub-optimal multi-beam is created by splitting the array into sub-arrays, each responsible for a particular component beam. In contrast, mmReliable utilizes phase and amplitude control to create the optimal channel-dependent multi-beam.

■ **Estimating parameters for constructive multi-beam:** Reliably estimating the parameters to establish a multi-beam is integral to maintaining its constructive nature. In the two-beam case, we need to estimate δ (relative attenuation, $\delta \in \mathbb{R}^+$) and σ (relative phase shift, $0 < \sigma \leq 2\pi$) of the reflected path w.r.t. the direct path. We already know the directions ϕ_1, ϕ_2 of the two paths from the beam-training phase. We denote the narrow-band complex wireless channel along each path as h_1 and h_2, respectively. If the gNB measures both h_1 and h_2, then it can estimate the relative amplitude and phase shift using their ratio. A simple method would be to measure each beam's channel one by one by setting the beam to ϕ_1 and then ϕ_2, as suggested in [39]. However, hardware offsets such as Carrier Frequency Offset (CFO) and Sampling Frequency Offset (SFO) cause time-varying and sometimes unpredictable channel phases, making these channel estimates unreliable [165, 70, 128]. The channel magnitude remains fixed despite these offsets. Therefore, we develop an estimation method using channel magnitude alone. First, we observe that only the *relative channel between the two beams* is of interest to us, which means that we can treat h_1 as a reference and assume $h_1 \in \mathbb{R}^+$ without loss of generality. Then, we observe that the magnitudes $p_1 = |h_1|^2$ and $p_2 = |h_2|^2$ are already available from the initial beam-training phase. Now, we develop a method to estimate their ratio using just two extra channel probes. Shown in Figure 1.5, the gNB uses two RS to probe two 2-beam patterns: first it sets the beam $w(\phi_1, \phi_2, 1, 0)$, and then

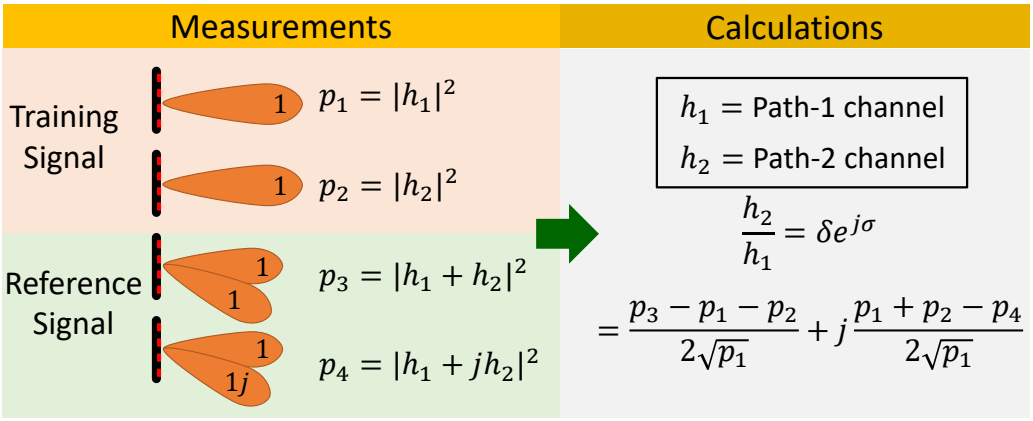

Figure 1.5. Channel probing procedure: To obtain parameters (δ, σ) for multi-beam constructive combining that's robust to CFO/SFO impairments.

$w(\phi_1, \phi_2, 1, \frac{\pi}{2})$. For each of the two probes, it estimates the channel magnitudes p_3 and p_4 as:

$$\begin{aligned} p_3 &= |h_1 + e^{j0} h_2|^2 = |h_1|^2 + |h_2|^2 + 2h_1\,\mathrm{Re}(h_2) \\ p_4 &= |h_1 + e^{j\frac{\pi}{2}} h_2|^2 = |h_1|^2 + |h_2|^2 - 2h_1\,\mathrm{Im}(h_2) \end{aligned} \tag{1.20}$$

where $j = \sqrt{-1}$. Using Eq. (1.20), we can estimate $\mathbf{h}_1$ and $\mathbf{h}_2$ individually, and express their ratio as:

$$\frac{h_2}{h_1} = \hat{\delta} e^{j\hat{\sigma}} = \frac{p_3 - p_1 - p_2}{2\sqrt{p_1}} + j\frac{p_1 + p_2 - p_4}{2\sqrt{p_1}}, \tag{1.21}$$

where $\{\hat{\delta}, \hat{\sigma}\}$ are the estimates of the relative amplitude and phase respectively between the two multi-beam paths. Thus, our algorithm estimates the required parameters using only two consecutive RS probes for the two-beam case. The algorithm can be generalized to any K multi-beam by performing two probes for each additional beam and solving Eq. (1.21) to yield $K-1$ relative channels w.r.t h_1. In [165], the authors develop a similar technique, albeit for measuring the per-element channel in a phased array. The channel probing overhead in [165] is proportional to the number of antennas N ($\approx 5N$). In contrast, our solution to measure the per-beam channel requires $2(K-1)+K$ probes for K beams (including the K estimates from the beam-training phase) and is independent of the number of antenna elements in the phased

array; making it tractable for large arrays.

■ **Handling wideband channels:** Typically, the Channel State Information (CSI) is measured across multiple frequency subcarriers over a wide bandwidth. The CSI is readily available for the 5G NR system from the reference signal [27]. Even for 802.11ad commercial routers, the CSI can be extracted in the firmware [178, 226, 165]. The prior analysis can easily be extended to multiple subcarriers by simply treating each subcarrier's channel independently. We denote $\mathbf{h}_1(f)$ and $\mathbf{h}_2(f)$ as the wideband CSI across frequency index f. We first estimate $\mathbf{h}_1(f), \mathbf{h}_2(f)$ using Eq.(1.21). Then, we formulate an optimization problem that maximizes the average received signal strength overall frequencies to jointly estimate $\{\hat{\delta}, \hat{\sigma}\}$ as:

$$\{\hat{\delta}, \hat{\sigma}\} = \underset{\delta,\sigma}{\arg\max} \ \left\| \mathbf{h}_2(f) - \delta e^{j\sigma} \mathbf{h}_1(f) \right\|^2. \tag{1.22}$$

We solve the above to obtain a closed-form solution ($\langle \cdot, \cdot \rangle$ is the inner product over the frequency dimension):

$$\hat{\delta} e^{j\hat{\sigma}} = \langle \mathbf{h}_1(f), \mathbf{h}_2(f) \rangle / \left\| \mathbf{h}_1(f) \right\|^2 \tag{1.23}$$

which reduces, as expected, to $\hat{\delta} e^{j\hat{\sigma}} = h_2/h_1$ for a narrowband channel (e.g. in IEEE 802.11ad 60 GHz standard [2]). The optimal phase, amplitude, & angle will change over time as the user moves. We periodically estimate these parameters using our tracking algorithm discussed in the next section.

1.2.4 Delay-Phased Array Beamforming for Wideband Channel

So far, we have discussed how to create constructive multi-beams using a phased array. However, a phased array-based multi-beamforming has a limitation in maintaining constructive combining for all frequency components in the signal. Notably, the phased array can only maintain constructive combining for one frequency and fails to maintain for other frequencies. So, this approach would only work for narrow-band channels (when multipath delays are small) and fail for wideband channels. We ask if there are new antenna architectures that enable

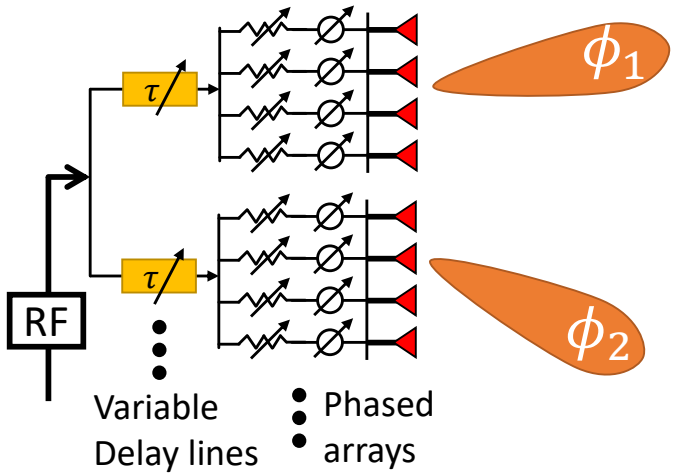

Figure 1.6. Proposed architecture for delay phased array: Delay-based multi-beamforming to handle wideband transmission. Note that this architecture can solve wideband issues due to multi-path channels.

wideband constructive combining for all frequencies in the signal.

Here we explore a new architecture called delay phased array, which uses two phased arrays connected with a network of delay lines and then to a single RF chain. Each phased array creates a single beam in a different direction such that they coherently add to a single receiver. Such architecture can handle multi-path channel bandwidth, providing flat frequency response.

We first show that a single beam system does not require the new architecture because there is only one path that causes the same delay at each antenna.

Let us reconsider the single path channel from N antenna Base station to a single antenna user. Assume the signal direction of departure is ϕ_1. At the far away location of the receiver, the signal travels a different distance from each of the transmit antenna which causes a delay in signal arrival given by τ_1 Considering this delay, we rewrite the expression for single path channel from (1.5) to the following expression of $\mathbf{h}^{\text{single}}(t,n)$:

$$\mathbf{h}^{\text{single}}(t,n) = h_1 e^{j2\pi(n-1)\frac{d}{\lambda}\sin(\phi_1)}\delta(t-\tau_1) \tag{1.24}$$

Here the delay τ_1 is constant for all antennas and can be handled by baseband processing.

Therefore, the optimal beamformer for such a single path channel remains a single beam pattern given by (1.6).

Multi-beamforming with new delay phased array architecture Next, we show that the traditional phased arrays optimal for a single-beam system don't work well for a wideband multi-beam system. Then we propose a new architecture called delay phased array based on delays and phase shifters that achieves optimal multi-beam performance despite a large delay spread of multi-paths.

The multi-path channel with two paths consists of delays τ_1 and τ_2 as:

$$h(t,n) = h_1 e^{j2\pi(n-1)\frac{d}{\lambda}\sin(\phi_1)}\delta(t-\tau_1) + h_2 e^{j2\pi(n-1)\frac{d}{\lambda}\sin(\phi_2)}\delta(t-\tau_2) \tag{1.25}$$

Here due to different delays of the two paths, the multi-beam beamformer will suffer from wideband issues. That's when two copies of the same signal are delayed by a different amount when added up at the receiver, creating an interference pattern at different frequencies. Some frequencies have constructive addition while others may have a destructive interference over a wideband scenario. To solve this wideband issue due to multi-path delays, we propose a new delay phased array architecture shown in Figure 1.6. Here we divide the phased array into two or more groups, each for one beam in multibeam. We then create additional delays for each sub-phased array that cancels the difference in delays caused by the channel multi-path. To create a two-beam beamformer in this case, we divide the antennas into two groups of $N/2$ antenna each and connect them with a variable delay line. The expression of a two-beam beamformer is:

$$\mathbf{w}_{\phi_1,\phi_2}(n) \begin{cases} e^{-j2\pi(n-1)\frac{d}{\lambda}\sin(\phi_1)}\delta(t-\Delta\tau_1), & \text{if } n \le N/2 \\ e^{-j2\pi(n-1)\frac{d}{\lambda}\sin(\phi_2)}\delta(t), & \text{if } N/2 < n \le N \end{cases} \tag{1.26}$$

where we apply a delay shift of $\Delta\tau_1$ at the first sub-phased-array to compensate for the multi-path delay difference in two paths.

Such a two-beam beamformer is better than the one in (1.19) because it will work for a

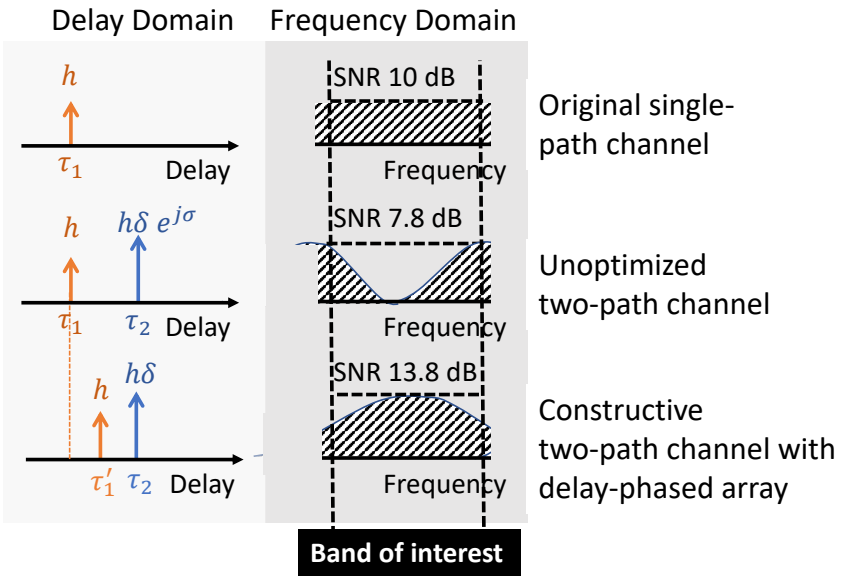

Figure 1.7. Constructive Combining for Wideband frequencies: Illustration of the advantages of our delay compensation with delay phased array architecture. We achieve a flat frequency response despite the presence of multi-path delays by compensating for the additional delays caused by the multi-path.

wide range of frequencies as shown in Figure 1.7. Here, we show that a single-path channel has a flat frequency response, but a two-path channel suffers from an imperfect frequency response where some frequency bands are severely attenuated. We also show that by compensating for the multi-path delays, we achieve frequency flat response, while providing 2x gain compared to a single beam system.

In Figure 1.8, we show through simulation that the optimal mmReliable beamformer created using our delay phased array architecture has a flat response across the entire frequency band, while a multi-beam without the proposed delay phased array architecture suffers from low SNR for certain frequencies. We plot two baselines with a natural multipath delay spread of 5 ns and 10 ns respectively. While non-optimized mmReliable suffers from signal degradation, delay-optimized mmReliable can compensate for this delay spread and achieve optimal flat performance.

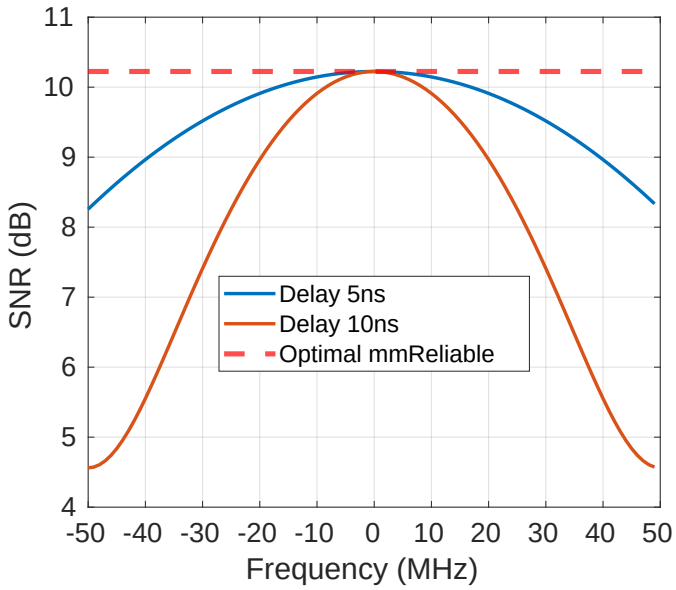

Figure 1.8. SNR with a range of frequencies: The delay-optimized mmReliable response is flat across all frequencies irrespective of any channel delay spread. It also shows for two different delay spreads of 5 ns and 10 ns, the SNR suffers at certain frequencies.

1.3 Proactive multi-beam tracking

Once a constructive multi-beam is established, mmReliable maintains it over time to honor the goal of reliable mmWave communication. The purpose of beam maintenance is to keep a high-throughput link even when the user is mobile or sees a random blockage. One choice for beam maintenance could be to periodically repeat the set of procedures in Section 1.2.3 to re-establish constructive multi-beam with a mobile user. However, it will require repeating the beam-training phase, which incurs significant delays and probing overhead (Section 1.1.2). We observe that the effect of blockage and mobility is embedded naturally in the form of variations in the wireless channel. Thus, our approach is to leverage OFDM channel estimate and reference signaling to maintain a multi-beam link and avoid impending blockage and mobility events while the communication link is active. We show the overall functioning of mmReliable using a flow-chart in Figure 1.9. mmReliable continuously monitors the OFDM channel to identify blockage or mobility on a per-beam basis. User mobility is tracked in the background, and the multi-beam is refined periodically with a low-overhead. When the beam is no longer recoverable

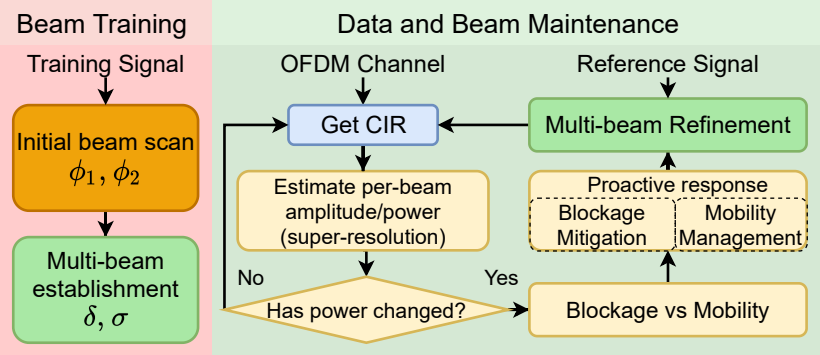

Figure 1.9. Overview: mmReliable's proactive response against blockage and user mobility.

by tracking alone (due to accumulated errors in tracking over time), mmReliable re-calibrates the system using the beam training phase.

1.3.1 Proactive Blockage Mitigation

Mobile blockers interact with the multi-beam system by suddenly occluding one or more active beams. mmReliable analyzes the rate of change of amplitude per beam to detect such events. Empirical evaluation shows that typical blockage events cause per-beam amplitude to degrade by 10 dB in just 10 OFDM symbols. Once a blockage event is detected, mmReliable responds by re-purposing the power on the blocked beam to other unblocked ones. Since there are always multiple beams active, there is no significant impact on link reliability, even if one is blocked. Said differently, the transceiver reduces the number of beams whenever blockage is detected along some paths. It is rare that all paths are blocked simultaneously; nonetheless, in case of a complete outage, the radio can initiate a new beam training phase to search for other alternate paths or perform a handover [229].

1.3.2 Proactive User-Mobility Management

When the user is mobile, the initial multi-beam may no longer be supported. A mere angular movement of 14^o would cause a 20dB loss in signal strength leading to an outage. A natural solution is to request a new beam training to locate the new user position, but beam training causes high overhead in tedious scanning. We propose to proactively track the user

mobility using channel measurements and refine the beam periodically. Traditional in-band tracking based on single-beam is not applicable to mmReliable since each beam in multi-beam may undergo different angular deviation making it challenging for tracking as shown in Figure 1.10.

Our insight is that we can track each beam by observing the gradual changes in per-beam power. We propose a model-driven approach to estimate per-beam angular deviation from the per-beam power measurements, $P_k(t)$ given by

$$P_k(t) = G_T(\phi_k + \varphi_k(t)) + G_R + P_T - P_h \text{ (in dB)}. \tag{1.27}$$

where $G_T(\phi_k + \varphi_k(t))$ and G_R are transmit and receive gain respectively, where initial angle ϕ_k changes by $\varphi_k(t)$ over time, P_T is transmit power, and P_h is power decay due to channel impairments. We take the difference $P_k(t_0) - P_k(0)$ to find the relative change in the per-beam power as:

$$P_k(t_0) - P_k(0) = G_T(\phi_k + \varphi_k(t_0)) - G_T(\phi_k) \tag{1.28}$$

where we assume that the channel loss P_h (due to path loss or reflection loss) and receiver gain G_R of omni-user is static for the small duration of user motion. Now, to estimate $\varphi_k(t_0)$, our insight is that the direct path power correlates with the beam pattern at the gNB, which is a function of spatial angle:

$$G_T(\theta) = \frac{\sin(N\theta/2)}{N\sin(\theta)} \tag{1.29}$$

where N is the number of antennas in a uniform linear array and θ is spatial angle. Therefore, we can use an inverse function to estimate the angle from the measured per-beam power. However, the beam pattern is usually symmetric, and two possible values $\varphi(t_0)$ and $-\varphi(t_0)$ could have caused the observed change in G_T. To deal with the ambiguity of the direction of motion, mmReliable tries one of the two possibilities using reference signal probing in the hope that it improves the SNR. If the probe doesn't improve SNR, possibly the other angle is correct, and

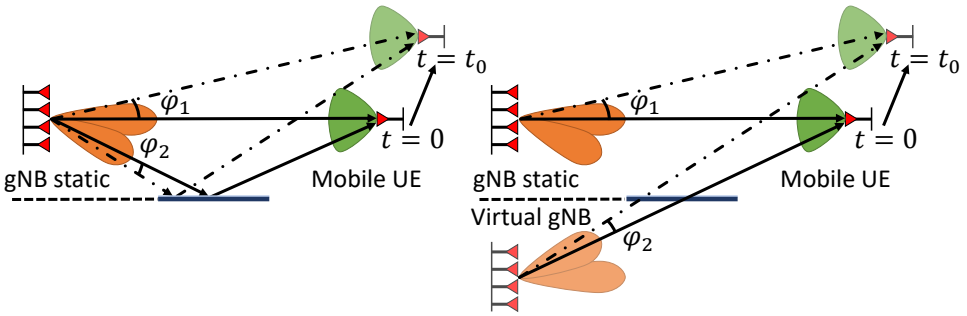

Figure 1.10. User mobility: Causes beams in multi-beam misalign with the user. Each beam may undergo a different angular deviation.

thus, mmReliable refines the beam to that angle. The refinement adds an overhead of only one additional probe in addition to $2(K-1)$ probes required for constructive multi-beam for the K-beam case. mmReliable periodically estimates the value of the per-beam angle, phase, and amplitude and updates the beam pattern at the gNB to realign towards the user.

1.3.3 Superresolution: Per-Beam Tracking

Here we describe how mmReliable estimates the per-beam power required for the tracking algorithm. A simple solution would be to scan single-beams in all K directions and observe the corresponding signal strengths. But, such scanning is intractable due to a high overhead in acquiring fine-grained power estimates for efficient tracking. Our idea is to estimate per-beam power using the channel impulse response (CIR) of the current multi-beam in the background without any tedious scanning. When receiving a multi-beam transmission, it can be shown that the signal at the receiver consists of a superposition of a delayed and attenuated version of each individual transmit beam. If we consider the delay and attenuation experienced by beam-index k as δ_k and α_k respectively, then the effective multi-beam channel can be expressed as:

$$h_{\text{eff}} = \sum_k \alpha_k \delta(t - \tau_k), \tag{1.30}$$

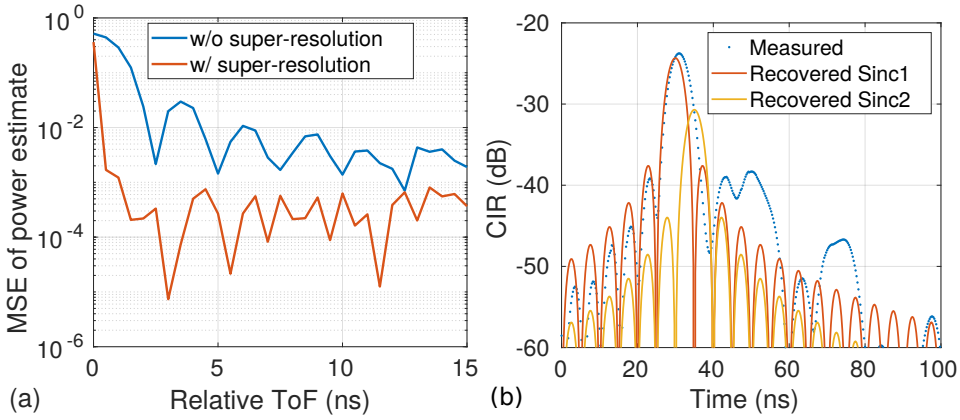

Figure 1.11. Efficiency of superresolution algorithm: (a) Simulation: Superresolution provides low MSE of per-beam power estimate even when the relative ToF is less than system resolution of 2.5 ns. (b) Hardware measurement: shows two sincs can be recovered from measured combined CIR.

where α_k is a function of the transmitter and receiver beam pattern, while δ_k relates to the Time of Flight (ToF) of each multi-beam channel path. Recall that our goal is to estimate the individual α_k for each beam ($k = 1, \ldots, K$).

For a frequency-selective wideband system the received signal is sampled at a sampling rate and sinc interpolated due to limited bandwidth as follows:

$$h_{\text{eff}}[n] = \sum_k \alpha_k \text{sinc}\left(B(nT_s - \tau_k)\right) \tag{1.31}$$

where B is the bandwidth, T_s is the sampling rate of the receiver, and α_k is total signal attenuation along path k.

We can re-write the above formulation as an optimization problem where our objective is to find α_k such that it fits the channel best. Let say the collected CIR is represented by column vector $\mathbf{h}_{\text{CIR}}$. We solve the following optimization to extract $\alpha = [\alpha_1 \alpha_2 \ldots \alpha_K]^T$, which amplitude per beam.

$$\alpha_{est} = \arg\min_{\alpha} \|\mathbf{h}_{\text{CIR}} - \mathbf{S}\alpha\|^2 + \lambda \|\alpha\|^2 \tag{1.32}$$

where the matrix $\mathbf{S}$ consists of all the ToF. We are essentially solving a super-resolution problem by fitting a sinc model over the entire CIR response. The optimization is a convex formulation,

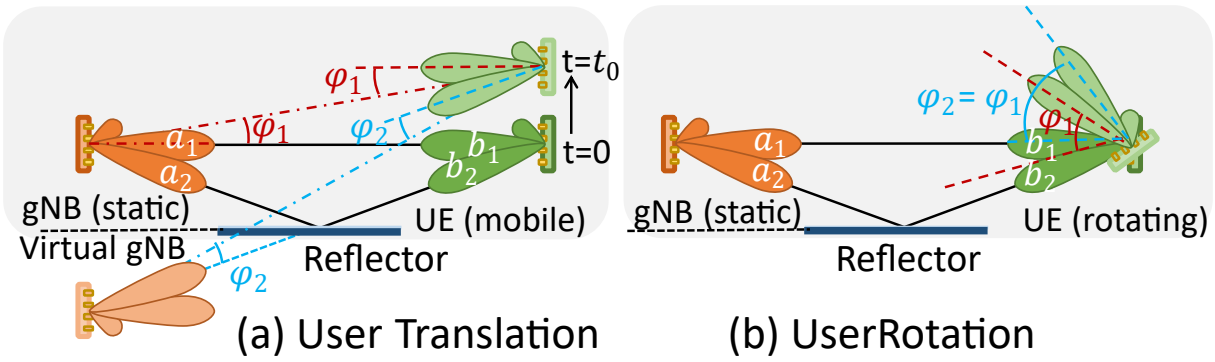

Figure 1.12. Multi-Beam UE: Demonstrating beam misalignment for a multi-beam UE.

with L2 norm regularization. We use standard techniques [51, 92] to solve this problem in 100μs. Here, we make a key observation that, after training, the absolute ToF may have changed, but the relative ToF changes slowly. Our key idea is that we shift the $\mathbf{h}_{\text{CIR}}$ first so that the strongest path is shifted to zero delays and since we know relative to the first path, the delay of the second and third path. We can populate the $\mathbf{S}$ matrix with only a few columns, thereby achieving accurate and reliable solutions to α_k. We account for small variations in relative-ToF by trying few values around the initial value that best fits our model. In summary, we have built an accurate super-resolution algorithm by leveraging the initial relative-ToF information between the multi-beams.

■ **Efficiency of super-resolution algorithm:** In Figure 1.11(a), we show that our super-resolution algorithm can indeed achieve high resolution in estimating per-beam amplitude even when the relative ToF is lower than what corresponds to the resolution (2.5 ns for 400 MHz bandwidth). We also measure the channel impulse response through our testbed (6m link with a reflector placed at 30^o) in Figure 1.11(b) and show that our super-resolution can efficiently extract the two sinc which are superimposed in the received CIR.

1.3.4 Generalization to Multi-beam UE

So far we developed mmReliable for a quasi-omni beam pattern at the user equipment (UE). Directional beams may be required at the UE side when the SNR is low, e.g., for longer

outdoor links. mmReliable extends naturally to the scenarios where the UEs also have directional beams. In order to do so, mmReliable utilizes the design concepts discussed before, like super-resolution, proactive tracking to mitigate the multi beam link against outages due to mobility, thus addressing the new system level challenges posed by directional UEs.

The main challenge posed by a directional UE is misalignments on both the UE and gNB side due to UE mobility. Consider the scenario in Figure 1.12 (a), in which because of UE moving from t=0 to t=t_0, the two beams at the gNB are misaligned by φ_1, φ_2 angle differences respectively. We see a similar misalignment at the UE side as well. If the UE and gNB were able to estimate these misalignment angles, they could compensate for the UE mobility by realigning the beams a_1, b_1 and moving them $\varphi_1, -\varphi_1$ respectively, and a_2, b_2 will be realigned by $-\varphi_2, \varphi_2$ respectively. Along with estimation of φ_1, φ_2, this step also requires association of beams a_1, b_1 and a_2, b_2, since there could be fallacious solutions by associating a_1, b_2 or a_2, b_1 together. Hence, handling UE mobility involves solving the two step problem, by first correctly associating one multi beam to the other, and then estimating the misalignment angles. The first problem follows as a corollary of our super-resolution algorithm, which allows to resolve the per path CIR, and by utilizing the unicity of ToF, we can associate two per path CIRs and hence associate the two beams which are along these respective paths.

In order to mitigate against the misalignments due to user mobility, mmReliable needs to repeatedly estimate and compensate for the misalignment angles. For estimation of the misalignment angles, we develop individual models for tracking translation and rotation similar to our approach in Eq. (1.28). The rotation of UE beam causes changes in UE antenna gain from which we can estimate the angle of rotation using a inverse function (Figure 1.12). The translation case causes changes in both the gNB's and UE's antenna gain. The received power in this case will be the sum of the two antenna gains. To estimate the angle in this case, we leverage our observation that the beams at gNB and the UE gets misaligned by the same angle (Figure 1.12). We estimate this angle by inverting the sum of the beam pattern at gNB and UE. In this way, we track the user motion and proactively refine the beams.

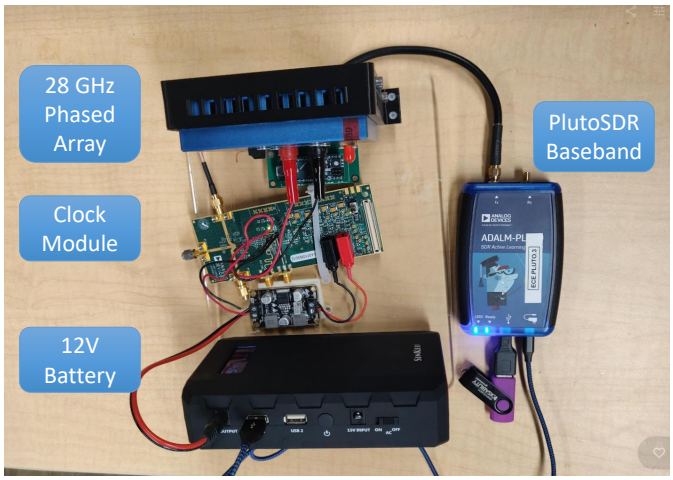

Figure 1.13. Compact mobile setup with 28 GHz phased array and a PlutoSDR streaming OFDM waveform using USB stick.

1.4 mMobile: A Compact, Real-time mmWave Testbed

Beamforming plays a critical role in millimeter-wave (mmWave) radios, leveraging phased arrays to emit highly directional waves at these high frequencies. The complexity of selecting optimal beams to maintain connections while accommodating dynamic factors like user mobility escalates with the number of users. Additionally, there are multiple environmental factors such as line-of-sight blockers, which can cause instant link disruptions, and static or dynamic reflectors, which enhance link reliability [198, 111, 137]. Consequently, beam selection poses a non-trivial problem that necessitates experimental investigations to develop algorithms addressing various challenges in mmWave communications. However, conducting experiments involving truly mobile users, operating in real-time on existing cellular protocols presents considerable difficulties. The lack of platforms capable of real-time processing of complex networking protocols, simultaneous beam selection, and compact form factor compounds these challenges. Existing mmWave testbeds [243, 188, 173, 248, 131, 113, 29, 110, 130, 189] are typically either cumbersome or lack fast microsecond-level beamforming capabilities

We showcase a truly mobile platform that is compact enough so that a person can carry it around easily as they would a normal smartphone while being powerful and fast enough to

allow for complex signal processing pipelines to be implemented. We use off-shelf Software-Defined Radios (SDRs) that are easily available and compact for baseband processing and signal generation. This allows experimenters to use vast amounts of the existing resources to create the required pipelines without a new learning curve. We also use state-of-the-art phased arrays that are capable of performing microsecond-level beam-switching that allows the development of algorithms that comply with 5G NR protocols. We present a compact interface with the SDR and phased array using FPGA, enabling real-time message passing between the network and the radio. This setup also ensures that all the systems are time-synchronized to the sample level so that we can have complete freedom and control over the algorithm researchers might want to deploy.

1.4.1 Requirements and Challenges

This section discusses the requirements and challenges in building a testbed oriented toward evaluating the effects of mobility in mmWave systems.

(i) Small form factor for mobile node: A form factor would allow the system to be mounted on vehicles for evaluating V2X scenarios. It requires a small and modular design of phased arrays integrated with mixers and amplifiers. Moreover, the mobile node's baseband processing should support 100 MHz bandwidth requisite by 5G NR standards. Finally, all these components have different power requirements that need to be supplied through a single battery (e.g., an available 12 V supply from a car).

(ii) High processing for gNB: The gNB needs to support up to 400 MHz of bandwidth, which cannot be met by commercial SDRs. It requires high-speed ADCs controlled by a powerful processor and FPGA. The gNB would be static and can be powered by an electric socket.

(iii) Fast control and calibration: The phased arrays need to interface with a microcontroller for fast switching between beam patterns. Moreover, some phased arrays [174] require calibration of antenna phases to generate desired beam patterns.

(iv) Clock and synchronization: The baseband sampler and up/down converter require a stable

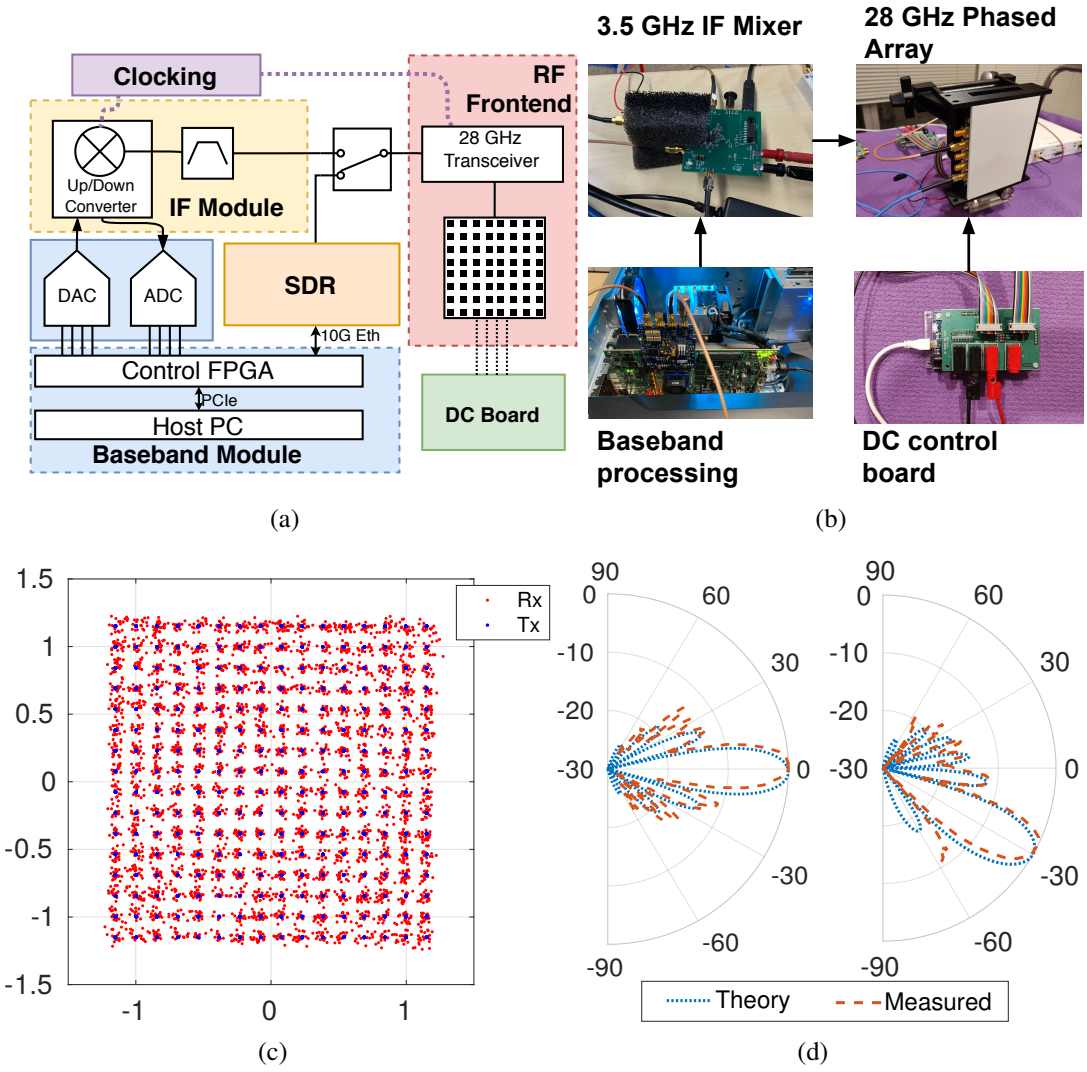

Figure 1.14. (a) Block diagram mmReliable testbed that has 2 ADC's and 2 DAC's, each sampling at 1 Gsps, which is converted to the IF stage of 3.5 GHz, and then the signal is provided to RF frontend which upconverts to 28 GHz and transmits through a phased array with 64 antennas. The node can configured to be transmit/receive mode with a T/R switch (b) Setup of mmReliable testbed for gNB (c) 256-QAM constellation. (d) Beam patterns measured in an anechoic chamber compared with theoretical patterns.

clock to reduce system phase noise. Moreover, 5G NR uses OFDMA-based waveform, which is susceptible to frequency offsets between gNB and UE and requires tight synchronization.

(v) Experimentation and mobility dataset: While the testbed architecture itself may be useful to members of the community to build one of their own, and it would be beneficial to be able to test beam-management algorithms without the need to invest and set up hardware. Allowing

remote experimentation and access is one straightforward way, but it involves a significant investment in the testbed's maintenance year-long. Added to that, limiting the testbed to remote experimentation would force the experiments to be confined to a few pre-decided scenarios, and a few users. Therefore, we believe that a well-designed dataset from multiple measurement campaigns [13] would be far more tractable and useful.

1.4.2 Related mmWave Testbeds in Literature

Hardware testbed: To evaluate the performance of cutting-edge applications and verify novel protocols, we need a wireless testbed with flexible configurability. Most testbeds [243, 188, 173, 248, 131] are designed for 60 GHz research based on IEEE 802.11ad which is not suitable for our design at 28 GHz for 5G New Radio. Moreover, OpenMilli [243] uses a mechanically steerable horn antenna not suitable for mobile networks. X60 [188] uses expensive hardware equipment from National Instruments (over $200,000) and is not affordable for most research groups. M-cube [248] uses cheap commercial phased arrays and is designated for Massive MIMO research. A similar 60 GHz testbed [228, 103] is used to demonstrate the high overhead of beam management in V2X scenarios. However, 60 GHz testbeds are not suitable for 28 GHz research because of vast differences in the phased array and circuit design. The phased arrays at 60 GHz suffer from higher feed loss and often produce irregular patterns [87]. Also, 60 GHz propagation suffers from higher path loss and atmospheric oxygen absorption at that frequency. In contrast, our testbed at 28 GHz provides realistic channel measurements for evaluating outdoor mobility for V2X scenarios.

There are few testbeds at 24 and 28 GHz as well. MIRA [29, 98, 118] testbed uses 24 GHz phased array but is limited to only 8 antennas. IBM phased array [186] uses 64 dual-polarized antennas for beamforming, but the baseband bandwidth is limited to that provided by commercial SDRs. In contrast, our mmReliable testbed consists of state-of-the-art phased arrays [124] with 64 antennas that work very well over 4GHz bandwidth (27-31 GHz). mmReliable's baseband processing is done on an FPGA platform that supports up to 1 GHz of baseband

bandwidth.

While commercial 5G deployment is on its early stage, many organizations are building a large-scale testbed that closely resembles a real-world 5G deployment [12, 183]. For instance, the COSMOS project [183] is being developed in New York City in collaboration with multiple universities and industry partners. However, such a large-scale design comes with additional cost, high complexity, and is hard to replicate elsewhere for mobile experiments. The COSMOS testbed is used for 28 GHz channel measurements [59] in NYC streets, focusing on large-area coverage. However, it lacks insights into improving the mmWave connectivity and throughput with proper beam management in mobile scenarios. In contrast, we have demonstrated that efficient beam tracking for a mobile user can significantly improve the link throughput for 5G NR. A comparison of mmWave testbeds is presented in Table 1.1.

Mobility management: There is a decade of work that targets initial beam-acquisition, which incurs high overhead proportional to the number of beams [207, 98]. In contrast, beam tracking methods can predict beam direction under user mobility without scanning a large set of beam patterns. Most tracking algorithms proposed in the literature are evaluated in simulation with unrealistic channel models [258, 194, and references therein]. There are few approaches that are tested on a testbed with limited settings. For instance, [96] proposes the use of a wide beam to combat high mobility, which is not viable for long communication links in 5G NR. Beam-forecast [255] uses a narrowband channel model for beam tracking, which is also not suitable for wideband 5G NR. Most importantly, both techniques [96, 255] are evaluated on a 60 GHz testbed with trace-driven evaluation using a mechanically steerable horn antenna that does not emulate phased array beamforming well. In contrast, we demonstrate our tracking algorithm on our 28 GHz testbed with phased arrays using 5G NR waveforms. Finally, location assistant devices such as GPS or external sensors can be used for finding the user's direction from gNB [251, 229, 205]. However, external sensors are not ubiquitous and suffers from low accuracy. On the other hand, our tracking approach completely relies on mmWave radios without requiring external sensors.

Table 1.1. Comparison of mmWave testbeds.

Testbed	Cent.Freq	Ant.Type	#Ant.	Bandwidth
OpenMilli [243]	60 Ghz	Horn	1	1GHz
X60 [188]	60 GHz	Array	12	2Ghz
mm-Flex [131]	60 Ghz	Array	16	2Ghz
Millimetera [173]	60 GHz	Array	8	1GHz
MIRA [29]	24 GHz	Array	8	NA
IBM [186]	28 Ghz	Array	128	100 MHz
mmReliable	**28 GHz**	**Array**	**64**	**1 GHz**

1.4.3 mMobile Testbed Architecture

The need for 5G NR compliance sets up some requirements for mMobile. An RF bandwidth of at least 400 MHz is required for gNB class performance. Further, low phase noise and synchronization have to be ensured for OFDM performance. With a large RF bandwidth, it follows the requirement of fast baseband processing to support it. Typical mmWave systems further require an IF bridge between the baseband and RF. Finally, to allow for larger flexibility, especially for mobile experiments, there is the requirement of a modular design.

Overview: mMobile uses a superhet architecture with an IF frequency of 3.5 GHz, typical to other similar systems. The baseband processing is implemented on an FPGA with a commercial Gigasample ADCs/DACs module, capable of supporting well over the required bandwidth. The wideband processing is matched at the IF and RF using wideband mixers and phased arrays. A clock distribution system supplies reference frequencies for the IF and RF mixing stages. Each module in the system may be replaced with a COTS alternative of similar specifications. The architecture is illustrated in Figure 1.14(a) and our setup is shown in Figure 1.14(b).

Configurable Baseband module: Software-defined radios (SDRs) such as USRPs are commonly used as baseband platform since they provide easy integration, experimentation, and debugging [186]. Most SDR works for a wide range of frequencies up to 6 GHz, and therefore they are suitable for directly providing the IF signal required for the 28 GHz link. However, the downside of SDRs is that they are limited by bandwidth ($<$ 150 MHz for USRP X310). 5G NR

specifications allow for up to 400 MHz of bandwidth in FR2. Such high bandwidth baseband processing requires high-performance FPGAs and ADC/DACs, but the FPGA setup may be bulky and power-hungry. To allow for both low and high-bandwidth applications, we use a configurable baseband module that can be switched between a USRP or an FPGA platform. For the high bandwidth node, a Xilinx KCU105 is used in conjunction with the Analog Devices FMCDAQ2. An external mixer (Qorvo RFFC5071A) is used to upconvert the FMCDAQ2 I/Os to the required IF frequency of 3.5 GHz. The digital processing pipeline uses PCIe to link to the host computer using a reference implementation like [243]. The low bandwidth SDR-based baseband module consists solely of a USRP X300 SDR with SBX-120 daughtercard to support 100 MHz bandwidth requirements at the UE.

Phased array and Calibration: Most phased arrays pose a random phase offset due to hardware imperfections at each antenna element corrupting the desired beam pattern. We use a simple scheme to calibrate the array as follows. We fixed one of the antennas as a reference and set its phase to zero. We then take one other antenna and vary its phase for all possible values. We measure the received signal strength (RSS) using a horn antenna in an anechoic chamber for all these phase values. Mathematically, the RSS is $|1+e^{j\phi}e^{j\psi}|^2 = 4\cos^2(\phi/2+\psi/2)$ where ϕ is the configured phase shift and ψ is the random phase offset between the two antennas. We observe that we can estimate ψ by simply looking for which configured phase gives the highest RSS. However, looking for the highest RSS is prone to errors because of a broad maximum. We instead look for a sharp minimum and add 180^o offset to get the phase calibration factor ψ. We can repeat the process to calibrate all the antennas relative to the reference antenna. We tested and verified our calibration routine with phased arrays from pSemi [174]. We also worked with phased arrays [124, 140], which do not require calibration.

IF design and Image rejection: The IF design requires the conversion of the signal from baseband to the IF frequency. We explore two choices of IF mixer design. An IQ mixer takes the I and Q signals from the DAC and combines them using two 90^o phase-shifted LO. However, we found that most commercial IQ mixers in our desired frequency range are either passive

mixers requiring high-power LO and amplifiers or suffer from low bandwidth. In our current implementation, we choose another IF design that implements IQ mixing in the digital domain and uses a single channel mixer. We handle the image rejection at the receiver using a requisite filter.

Clock distribution module: The clock distribution is an essential aspect of a 5G testbed. In a heterodyne architecture, we require three clocks for the baseband module, IF mixer, and the RF mixer. We use an ADF5355-based clock distribution system to generate synchronized 12.5 GHz and 3.5 GHz clocks for the IF and RF. We also connected a GPS disciplined oscillator (GPSDO) that uses GPS signals and a high-quality oscillator to provide a stable clock and remove any sampling frequency offsets.

We use a custom transceiver module to up/downconvert between the IF and 28 GHz. The transceiver module drives the 8x4 dual-polarized phased array (64 antennas) with 6 bits phase depth and 27 dB amplitude range [140]. The transceiver module contains a 4x clock multiplier and is clocked by a 6 GHz signal from a ADF5355 PLL. The transceiver uses this clock to upconvert signals from an IF of 4 GHz. We use an Artix-7 based FPGA to control the time-critical phased array beamforming weights through SPI. We also designed a power regulator board (DC Board) that takes a single 12 V supply and powers the array, up/down converters, and fan. We have tested our architecture with three different phased arrays [124, 174, 140].

1.5 mmReliable Implementation

We use a custom 28 GHz mmWave testbed to evaluate the performance of mmReliable. Our testbed consists of two major components as shown in Figure 1.15(a): (i) an 8×8 mmWave phased array (PA) with real-time beamforming, and (ii) a baseband module that generates and processes standard-compliant OFDM waveforms. For more details on the architectural choices and performance of our testbed, we refer the reader to [113]. In this section, we describe components relevant to this work in greater detail.

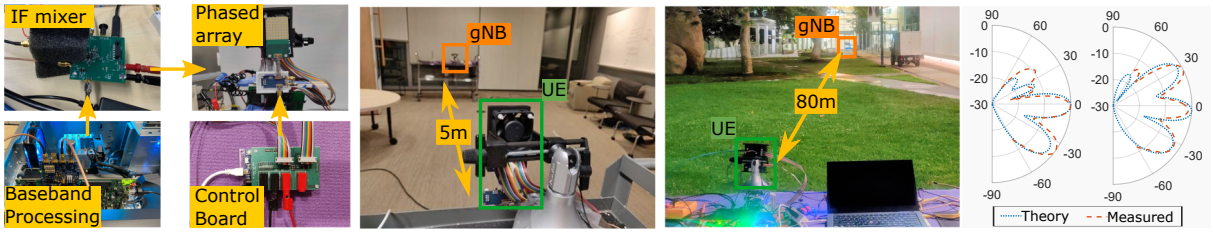

(a) Experiment testbed (b) Indoor scenario (c) Outdoor scenario (d) Multi-beam patterns

Figure 1.15. Experimentation Testbed: (a) Our experimental testbed with FPGA baseband and 64 element phased array [113], (b) indoor 5m and (c) outdoor 80m test scenarios. Our phased array is mounted on a gantry system which provides precise translation and rotation measurements used as ground truth. (d) Comparison of multi-beam patterns from theoretical analysis and our measurements.

1.5.1 Phased Array and Beamforming

Our 8×8 Phased Array has only one RF chain and is designed in-house [113, 124, 140]. It features an integrated up-down converter that operates with an IF frequency of 3.5 GHz. The array is capable of 2 GHz bandwidth centered at 28 GHz and includes amplifiers, image-reject filters in both transmit and receive directions. We use only the azimuth for beamforming and set all the weights along the elevation to the same value. We use a Cinetics Axis-360 [7] gantry system to precisely move the array for our evaluations.

■ **Beamforming Control:** Inspired by state-of-the-art design [186], we use an Artix-7 FPGA [232] with custom Verilog modules to program the phased array with high speed (100 us per beam) and accurate timing over an SPI bus. In practical systems, beam-weights that span a limited set of angular directions (typically numbering 64-1024) are programmed into the memory of the FPGA device [165]. Multi-beams require beam-specific amplitude and phase in addition to angular directions, and storing all possible combinations in memory could be intractable. Since multi-beams are linear sums of single-beam-weights, they can be generated using simple addition and multiplication operations on the FPGA. It would allow the fast creation of multi-beam weights while requiring storage of only single-beam weights.

■ **Beam Weight Quantization:** Phase-coherent multi-beam patterns can be generated using minimum 2-bit phase shifters and 1-bit amplitude control (turning antenna on or off) [164, 138, 165] which is available in commercial 802.11ad devices [178, 9, 143] as well as 28 GHz 5G

NR [190, 183]. Having higher resolution in phase and amplitude control helps create accurate single-beam and multi-beam patterns that reflect mathematical models. Our array provides 6 bits of phase control and 27 dB of gain control per antenna element similar to commercial products [190, 183, 186][3]. Figure 1.15(d) shows that our phased arrays generate accurate multi-beam patterns.

1.5.2 Baseband Module and 5G-NR Compliance

Figure 1.15(d) Shown in Figure 1.15(a), we use a 1 GSPs ADC/DAC in conjunction with a KCU-105 FPGA [6, 5] for creating 400 MHz baseband OFDM waveforms. We use an off-the-shelf IF mixer [8] to up/down-convert to the IF frequency of 3.5 GHz. For outdoor experiments (Figure 1.15(c)), we use a compact setup with a USRP X300 (100 MHz bandwidth) [113]. A host PC running MATLAB controls waveform generation, post-processing, and beamforming. We synchronize the array's clock with that of the baseband module to switch beams at intervals with an accuracy of $< 0.1\mu s$.

■ **5G-NR Compliance:** To ensure that mmReliable is compatible with 5G NR standards, we use OFDM waveforms with an effective bandwidth of 400 MHz or 100 MHz and numerology that yields 120 kHz sub-carrier spacing [10, 27]. The beam-training phase in 5G-NR is performed using channel estimates from Synchronization Signal Block (SSB), which repeats with a default period of 20 ms [10, TS 38.215]. mmReliable leverages channel estimates from the Channel State Information Reference Signal (CSI-RS) [10, TS 38.215] for beam-maintenance. CSI-RS can be sent arbitrarily with a spacing ranging from 0.5 ms to 80 ms. A CSI-RS occupies only one symbol (8.93μs@120 kHz) in the slot and only a few frequency subcarriers [10, TS 38.211], contributing much less overhead than an SSB. For example, by using only one CSI-RS symbol every 20 ms for beam-maintenance, we incur an overhead of less than 0.04%. While beam-training cannot be avoided altogether, the maintenance scheme can reduce its periodicity. If the

[3]Verizon, Ericsson, & other commercial vendors deploy high-resolution 28 GHz phased arrays developed by Samsung, IBM & others [11, 3]

periodicity of a 5 ms long SSB can be extended to 1 second, then the overhead of the entire beam-maintenance process is 0.5% of channel air time. During post-processing, we use channel estimates only at standard specified intervals to evaluate all algorithms.

1.6 Evaluation

We evaluate the performance of mmReliable for various indoor and outdoor settings: a $7m \times 10m$ conference room occupied with wooden furniture, whiteboard, and reflective glass walls (Figure 1.15(b)) and an outdoor $30m - 80m$ link next to a large building with glass walls (Figure 1.15(c)). The Tx module of the testbed is used as a fixed gNB, while the Rx module is movable as a mobile user placed on a cart. The phased array antenna of the Tx is kept vertical, facing the user, and parallel to the line of user translation. The precision gantry setup allows 70 cm linear motion and 360-degree rotation about its axis accurate to 1 cm and 0.1^{o}, respectively, for controlled experiments. The Rx phased array mounted on the gantry (see Figure 1.15 (a)) moves as a mobile user with 24^{o}/s rotation (typical speed for VR headset motion [175]) and 1.5m/s translation speed in indoor scenarios. For outdoor experiments, we manually move the cart to a predefined trajectory for accurate ground truth information. We also experiment with natural motion for end-end experiments.

We perform multiple experiments of 1-sec duration with mobility and blockage. At the beginning of each experiment, we perform beam training to estimate the angle to the user and establish directional communication. We then run our constructive multi-beam and user tracking algorithm to realign the beam patterns in real-time as the user moves. Our testbed reports the complex channel per sub-carrier, Signal to Noise Ratio (SNR), and throughput of the wireless link as performance metrics similar to others [257, 207, 208, 248, 183].

1.6.1 Micro Benchmarks

Here we present benchmarks for both constructive multi-beam and proactive user tracking to characterize our system.

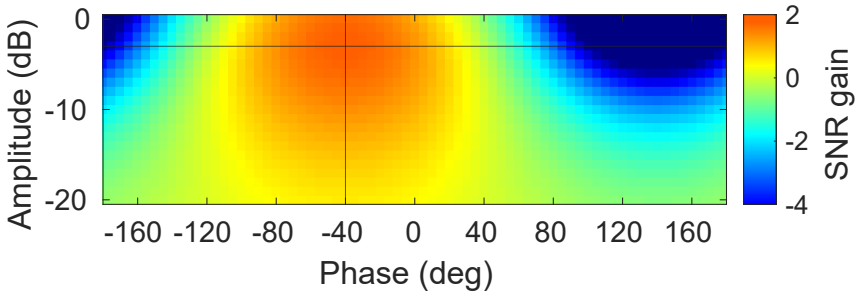

Figure 1.16. Sensitivity analysis (simulation): SNR gain (dB) of 2-beam w.r.t. single-beam for varying 2nd beam's phase $\hat{\sigma}$ and amplitude $\hat{\delta}$ in multi-beam. We show for a -3 dB weaker multipath channel, the multi-beam's highest SNR gain is 1.76 dB which is not very sensitive even for a high phase error of $\pm 75^{o}$. However, if the phase error is highest (180^{o}), the SNR drops significantly.

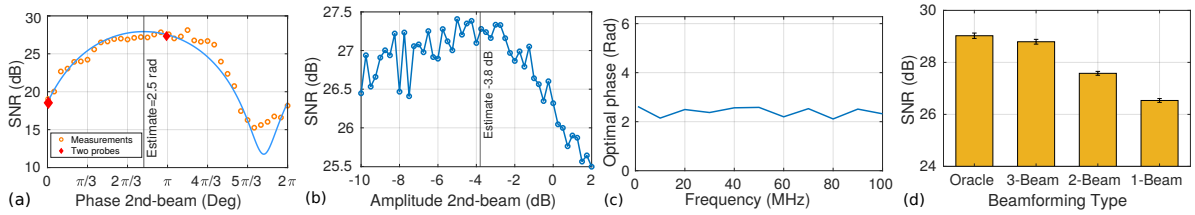

Figure 1.17. Benchmark Results on mmReliable: (a) SNR measurements with different phases of the 2nd beam in our 2-beam pattern, (b) SNR measurements with various amplitudes of the 2nd beam through brute-force probing. We also show that our constructive multi-beam requires only two probes and provides accurate phase and amplitude estimates without brute-force scanning. (c) The optimal phase is stable over a frequency range of 100 MHz. (d) SNR gain of multi-beam against the single beam and oracle (knows the best channel-dependent beam) for a static unblocked link.

■ **Sensitivity Analysis of Multi-Beam:** Accurate estimation of phase and amplitude provides the highest SNR; however, it is not very sensitive to estimation errors in these parameters. To quantify the sensitivity, we simulate a two-path channel and set the relative phase and attenuation to -40^{o} and -3 dB, respectively. We vary the relative phase and amplitude of the 2nd beam in a 2-beam pattern (w.r.t. the first beam) and show the 2D plot of SNR gain in Figure 1.16. The highest SNR gain is 1.76 dB for perfect estimation, which reduces gradually with an increasing mismatch between channel parameters and their estimations. The multi-beam SNR exceeds the single-beam SNR and can tolerate errors of $\pm 75^{o}$ in phase estimation and amplitude errors of -20 dB.

■ **Constructive combining accuracy:** We first show the need for constructive combining in a multi-beam system and then evaluate the performance of mmReliable's algorithm for estimating

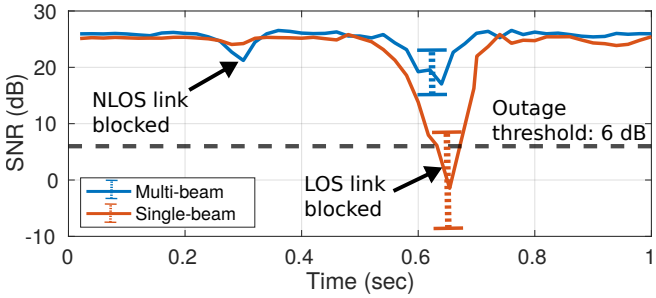

Figure 1.18. Blockage-Resilience Result: Multi-beam is more reliable and blockage resilient compared to a single beam.

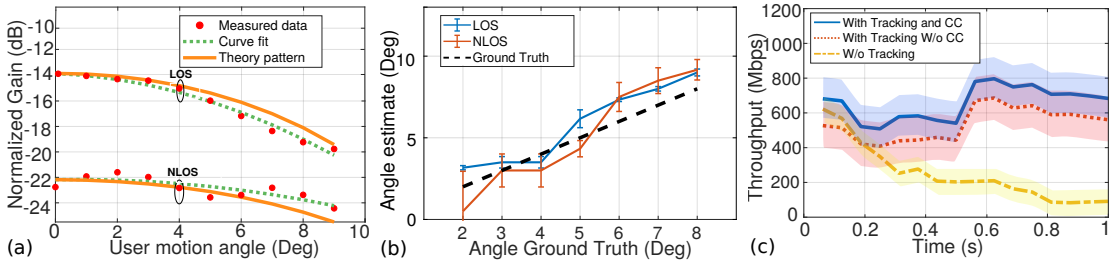

Figure 1.19. Mobility Tracking Results: (a) Effect of noise and channel impairments on user tracking for UE rotation. We show that even NLOS paths with low signal strength can be tracked effectively. (b) Tracking angle accuracy: mmReliable can estimate the angle within 1^{o} error on average. (c) Tracking time series showing the advantage of constructive combining (CC) & tracking; the shaded region shows standard deviation across multiple experiments, and lines show mean throughput.

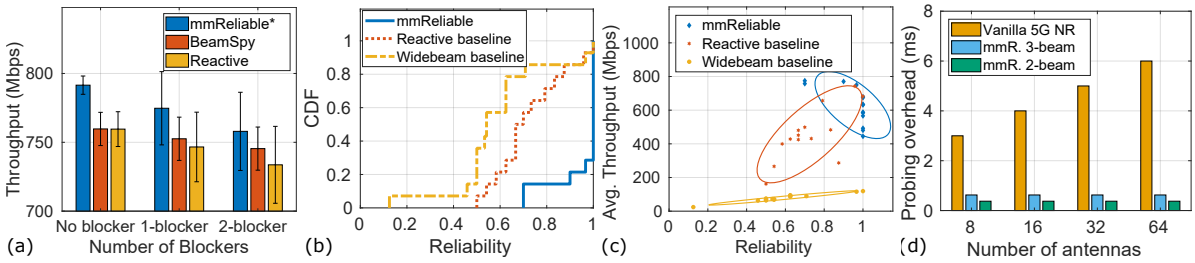

Figure 1.20. End-to-End Results: Performance gain of mmReliable compared to baselines under different scenarios: (a) Static link with blockages (*evaluates mmReliable without proactive user tracking) (b) Reliability for a mobile link (c) Overall Throughput-Reliability Trade-off; ellipses represent standard deviation in two dimensions. (d) mmReliable (mmR.) has a low probing overhead than vanilla 5G NR.

the phase σ and amplitude δ needed for constructive combining. We establish an indoor multi-beam link with two beams towards a static UE at 7 m along a LOS angle of 0° and NLOS angle of 30^{o}, respectively. We set the phase and amplitude of the first beam to be 0 radians, 0 dB, respectively. We then perform an exhaustive scan of the second beam's phase in the range $[0, 2\pi]$

and observe the SNR variations in Figure 1.17(a). We see that the SNR is a maximum of 27 dB with a minor variation of around 1 dB in the range of $\pm 70^{\circ}$. Any estimate in this range would give a high SNR. But, if the phase estimate is off by 180°, the SNR degradation could be as significant as 13 dB, which motivates the need for constructive multi-beam. It also indicates that mmReliable's two-probe method accurately estimates the phase of 2.5 radians, which maximizes the SNR. Similarly, we perform an exhaustive scan of the second beam's relative amplitude in the range of $[-10, 2]$ dB and show the SNR variations in Figure 1.17(b). It indicates that the SNR is highest for a range of amplitude around -5 dB to -3 dB. Therefore our estimate of -3.8 dB using the two-probe method is reasonably accurate. Finally, we show the per-beam phase's stability over the 100 MHz frequency range in Figure 1.17(c). We observe that phase variation is less than 1 rad, which doesn't impact the SNR gain of our multi-beam system.

■ **SNR gain due to constructive multi-beam:** We now verify that constructive multi-beam leads to higher SNR than a single beam system, even for a static unblocked link. We reuse our previous set-up and measure that our two-beam multi-beam link provides 1.04 dB SNR improvement on average compared to a single beam oriented towards the LOS path (Figure 1.17(d)). The SNR gain is achieved by constructively combining the signal along the two-channel paths. To understand how many beams are required in a multi-beam, we establish an oracle channel-dependent beam that uses the channel at each antenna element of the phased array. Obtaining this oracle beam requires an extensive measurement process with a high overhead [165]. This oracle beam provides an SNR gain of 2.5 dB w.r.t. single beam. We show that our multi-beam pattern with only three beams (utilizing the three most strong channel paths) provides a 2.27 dB gain w.r.t. single beam, which performs similarly to the oracle beam. *This analysis indicates that 3-beams are enough to achieve within 92% of the optimal beam with significantly less overhead.*

■ **Static link with blockages:** To study the impact of blockage on a multi-beam link, one of the authors walks across our established link multiple times with a consistent walking speed. From Figure 1.18, we observe that our multi-beam link with two beams is blocked twice as the blocker passes both the NLOS and LOS link. However, even for the LOS link blockage, the

SNR reduction is only 7 dB, which does not cause a link outage. In comparison, a single beam link SNR degrades by 26 dB, below the outage threshold of 6 dB SNR (required for decoding 5G-NR OFDM signals). *Therefore, a multi-beam link is resilient to outages because it avoids a single point of failure.*

■ **Accurate per-beam power estimation:** We now test our multi-beam link under mobility. We set up a multi-beam with two beams at the transmitter and rotate the transmit phased array using our gantry. We then collect the channel estimates with our Omni-directional receiver and estimate per-beam power using our super-resolution algorithm. We plot the per-beam power of the two beams in multi-beam with ground truth angle in Figure 1.19(a). We observe that per-beam power reduces with the angle of rotation and follows the beam pattern. We observe perturbations in power values due to measurement noise, which may affect tracking performance. To improve tracking accuracy, mmReliable takes time average of power values with a forgetting factor & fits a quadratic polynomial to smooth the data. With proper scaling, the curve fitting approximates the beam pattern within 1 dB mean error (Figure 1.19(a)).

■ **Accurate rotation angle estimation:** We next evaluate the performance of our tracking algorithm. We perform multiple experiments with a uniform rotation of the transmitter array. The frequency of rotation is varied over a range of 2^{o} to 8^{o}, and the average angle estimated by our algorithm is reported in Figure 1.19(a). The mean angle estimate is accurate to 1^{o} for tracking both LOS & the NLOS link when compared against ground truth.

■ **Tracking accuracy time series:** Accurate angle estimation by mmReliable's proactive tracking algorithm helps maintain a high throughput link under user mobility. Figure 1.19(c) shows the throughput of our multi-beam links for a uniform translation of 1 sec duration (at 1.5 m/s speed) under three situations. First, we show that the link throughput degrades from 600 Mbps to $\sim$ 100 Mbps below our outage threshold if no tracking is performed. In contrast, multi-beam tracking with constructive combining (CC) approach maintains a consistently high throughput 1-sec of user mobility. Finally, if we perform tracking alone, without doing constructive combining optimization, the throughput improvements are lower by an average of 100 Mbps compared to

the case of tracking plus CC. *Thus, mmReliable provides consistently high throughput due to its super-resolution, per-beam tracking, & constructive combining algorithms.*

1.6.2 End-to-end Results

We present end-to-end results for mmReliable and compare the performance against other baselines [208, 98, 96].

■ **Multi-beam performance for static link:** We first show that mmReliable with constructive multi-beam and proactive response to blockages can beat other single-beam-based baselines for static links. For this, we implement mmReliable without the user tracking algorithm. We also implement two baselines; (i) BeamSpy [208] which detects link blockage by exploiting spatial channel profiles to find the alternate unblocked link without extensive training. (ii) Reactive baseline [98], which implements a fast beam training in response to the blockage event. Both baselines are based on a single beam, and they suffer from a single point of failure to link blockage and act after an outage is detected. We see from Figure 1.20(a) that mmReliable outperforms both baselines and provides a high throughput mmWave link even under the impact of blockage. mmReliable's throughput drops only by 4% even when there are two blockers near the beams because of the proactive utilization of reflectors.

■ **Improvement in Reliability for mobile links:** We perform user translation and rotation experiments in multiple indoor and outdoor environments. As the user moves, a human blocker is introduced midway between gNB and UE, blocking the link for a duration chosen uniformly between 100 ms to 500 ms [180], over 1-sec experiments. We perform 100 such experiments and combine the results to form one point, and various user mobility and blockage patterns across two indoor environments are reported on the throughput and reliability curve. Figure 1.20(b) shows that mmReliable achieves close to 100% reliability (median value 1). The reactive baseline suffers from lower reliability of median value 0.65, while the widebeam baseline is at 0.5. Thus, mmReliable achieves its goal of high reliability.

■ **Throughput-Reliability Trade-off:** We introduce a new method to evaluate mmWave

systems by interpreting the data from the last sub-section using both reliability and throughput metrics. Using both metrics allows us to compare mmReliable with other baselines to assess the *throughput-reliability tradeoff* in Figure 1.20(c). Across all experiments, mmReliable is not significantly affected by throughput loss, while the other benchmarks show degradation due to blockage and user mobility. The results indicate that mmReliable delivers an average throughput improvement of 50% over the reactive baseline (200 Mbps improvement over 400 Mbps). More importantly, mmReliable offers a consistent throughput with low variations compared to the baselines, essential to many mission-critical applications. *The evaluation reveals that mmReliable is advantageous over throughput-based design on both throughput & reliability.*

■ **Beam probing overhead:** mmReliable reduces the beam probing overhead by maintaining a multi-beam link through tracking and beam refinements. The proposed beam refinement procedure requires three CSI-RS probes for two-beam case: two for estimating δ, σ and one probe for detecting the direction of motion. For the 3-beam case, it only increases to 5 probes. On the other hand, an entire beam scan using SSB in traditional 5G NR requires a large number of SSB probes proportional to the number of spatial directions. One CSI-RS occupies one slot, which is 0.125 ms at 120 kHz sub-carrier spacing, and one SSB takes four slots (0.5 ms). Let's consider the best scanning method, which requires the number of probes proportional to only a logarithmic number of antennas [98]. We show in Figure 1.20(d) that traditional 5G NR suffers from higher probing overhead compared to mmReliable. The 5G NR probing overhead for eight antennas base station is 3 ms, which increases to 6 ms for 64 antennas (because of the high directionality of beam patterns). In contrast, the overhead of mmReliable remains as low as 0.4 ms for 2-beam & 0.6 ms for 3-beam cases independent of the number of antennas. *Thus, mmReliable is an efficient solution for beam management, which improves mmWave link reliability & throughput while maintaining low system overhead.*

1.7 Related Work

mmReliable is closely related to extensive work in mmWave communication literature, therefore we have grouped the past work into following categories.

■ **Channel-dependent or multi-beamforming:** Multi-beam patterns belong to the family of channel-dependent beamforming [203, 165] where the beamforming weights are obtained to maximize the SNR for a given channel condition. For instance, [203] showed through channel-sounding measurements that combining multiple signals across four directions could achieve 28 dB improvement in path-loss at 28 GHz; however, they require multiple phased arrays. Traditionally these optimizations are proposed for a MIMO system that uses channel measurements at each antenna [72, 99, 71] for digital beamforming. In contrast, mmReliable focuses on analog beamforming where the signal from each antenna element is connected through a single RF chain. ACO [165] proposed a procedure to estimate per-antenna CSI for analog array by designing a specific beam scan procedure.

However, all work in this field incurs high overhead proportional to the number of antennas and cannot be repeated often as the channel changes. In contrast, we propose a multi-beam that approximates the optimal beam pattern and is easy to maintain with low overhead. A similar system [39] proposes constructive multi-beam for communication that focuses solely on beam-training and multi-beam creation. Specifically, it does not describe any beam-maintenance scheme and therefore would incur high overheads and be unreliable in practice. mmReliable is the first system that enables multi-beam constructive combining with low-complexity beam training and continued beam maintenance, mitigating blockage and mobility effects.

■ **Blockage and mobility:** BeamSpy [208] and Beam-forecast [255] provide a model-driven approach to handle blockage and mobility, respectively. However, they suffer from high run-time optimization and table-lookup complexity. To recover from a complete outage, UBig [206] performs efficient handovers, and mmChoir [241] proposes joint transmission from multiple base stations. These approaches may improve system reliability but add additional signaling overhead.

Location estimation using extensive beam training from single or multiple gNBs is proposed in [96, 181, 169, 82, 166]. In contrast, mmReliable performs tracking of the user without any training overhead. [138] developed a mechanism to detect the direction of user movement and adjust beam direction in small increments. Though they use a 2-lobe pattern to detect mobility, they switch back to a single-beam pattern for communication. The fundamental difference is they use a single beam and are generally reactive.

■ **Beam Training:** A vast literature focus on reducing the initial beam training delay through single-beam [114, 43, 256, 256, 44, 207] or multi-beams [30, 98, 215, 40, 39]. In contrast, mmReliable focuses on link maintenance and reliability and therefore is complementary to initial beam training approaches. High training overhead may be reduced by taking assistance from device positioning systems [229, 218] or light sensors [95], out-of-band WiFi [205, 159], and multi-user coordination techniques [195, 176]. However, in contrast to mmReliable, they suffer from relatively lower accuracy with added complexity while dependent on external devices. mmReliable high accuracy is achieved based on the super-resolution algorithm combined with the motion tracking algorithm.

■ **Reflecting surfaces:** Prior work on active mmWave relay [28, 212] requires additional active devices (antenna array) with beamforming capabilities to relay the signal to the user. These devices increase cost and power consumption. On the other hand, mmReliable relies entirely on natural reflectors present in the environment. The authors of [257, 84, 247, 123, 259, 230, 254] leverage reflective surfaces to improve indoor mmWave connectivity & coverage. Their model is suitable for one-time sensing for initial network deployment, which complements our goal of improving the link reliability in real time.

1.8 Discussion and Future Work

Since multi-beams are a new class of beamforming, designing a robust networking system while multiplexing users is an open challenge. We discuss some limitations of mmReliable and

identify possible future research directions.

Strong reflectors: The functioning of mmReliable depends on the availability of strong reflectors in the environment, otherwise, the multi-beam system falls back to a single-beam system. We envision future deployments where intelligent reflecting surfaces [70, 28, 225, 212] are deployed in the environment to engineer strong reflections that improve the throughput and reliability of mmWave links.

Tracking re-calibration: While we can track the angle of movement of both direct and reflected paths using super-resolution and tracking algorithms, the tracking of reflected paths would be affected when the path is blocked, or the reflector vanishes. mmReliable detects such cases by observing per-beam power and reallocates the power along other unaffected beams. If the tracking error accumulates over time, mmReliable needs to initiate a beam training to refresh the system or perform a handover.

Hybrid beamforming and mmReliable: Using multiple mmWave RF chains and phased arrays can provide degrees of freedom to multiplex users or improve throughput. Prior work [118, 87] has proposed interference-aware spatial multiplexing of beams in different directions for multiple users. These techniques fit with mmReliable in many ways. First, many multi-beams can be created, one from each RF chain, allowing mmReliable's improvements for multiple users. Second, in addition to per-beam phase, amplitude, and angle, multiple RF chains allow for control over group delay, enabling arbitrarily wideband operation. Finally, we can jointly use some spatial beams for enhancing reliability while others for improving multi-user coexistence.

1.9 Acknowledgements

We are grateful to the anonymous reviewers and our shepherd, Prof. Romit Roy Choudhury, for their insightful feedback. We thank Prof. Gabriel Rebeiz, UCSD, and his group for the mmWave 5G phased arrays and transceivers. The research is supported by NSF CCRI # 1925767.

Chapter 1, in part, is a reprint of the material as it appears in **IK Jain**, R Subbaraman, D Bharadia, "Two beams are better than one: Towards Reliable and High Throughput mmWave Links", *ACM SIGCOMM 2021*. The dissertation author was the primary investigator and author of this paper.

Chapter 1, in part, is a reprint of the material as it appears in **IK Jain**, R Subbaraman, TH Sadarahalli, X Shao, H Lin, D Bharadia, "mMobile: Building a mmWave Testbed to Evaluate and Address Mobility Effects", *4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems (Mobicom Workshop - mmNets)*, 2020. The dissertation author was the primary investigator and author of this paper.

Chapter 2

mmFlexible: Flexible and Scalable Multi-user mmWave Networks

Millimeter-wave (mmWave) networks have the potential to provide wireless connectivity to a growing number of users with their vast bandwidth resources. However, current mmWave systems have a significant limitation in that they are unable to simultaneously serve multiple users by distributing small chunks of frequency resources to different users who are in different directions. Unlike sub-6 systems, that use Omni-antennas to radiate signal in all directions, mmWave systems use pencil beams that illuminate a small region in space, meaning that all the frequency components are directed towards a fixed direction and cannot be distributed to other directions. This inflexibility leads to two main issues, as illustrated in Figure 2.1(a). Firstly, it leads to high latency, as the base station (gNB) must serve different user directions in a time-division manner, causing some users to experience long wait times, which is detrimental to latency-sensitive applications. Secondly, it leads to low effective spectrum usage. When a gNB serves one device at a time, each device gets a lot of instantaneous capacity which it may fail to utilize due to limited demand. But because the gNB cannot direct the remaining frequency resources in other directions, those resources are wasted, leading to low effective spectrum usage. Furthermore, other users in other directions could have used these wasted bands to improve overall spectrum usage.

In this chapter, we ask the question "*whether a mmWave base station can transmit or*

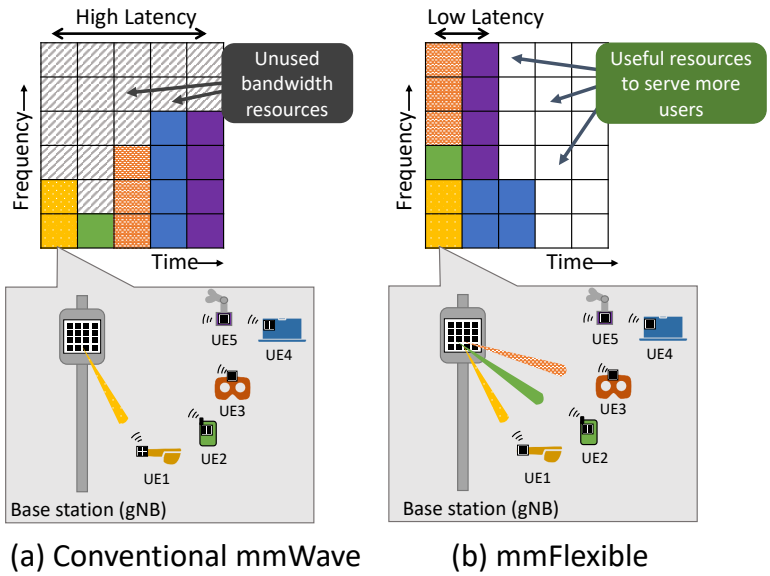

Figure 2.1. mmFlexible Idea: It enables efficient use of the available mmWave spectrum resources through flexible directional-frequency multiplexing, allowing multiple users to be served simultaneously with low latency and high spectrum utilization.

receive to any set of arbitrary directions using any set of contiguous frequency bands, creating a flexible frequency-direction beamforming response". This is possible using massive antennas and digital beamforming, but traditional mmWave systems rely on analog phased arrays with a single RF chain for cost and power efficiency. However, analog phased arrays cannot create such a frequency-direction response, as they take a single input from the RF chain and radiate all frequencies in the signal in one fixed direction. One naive solution is to split the phased array into multiple sub-arrays and program them to radiate in different directions, but this reduces the directivity in each direction, reducing signal strength, range, reliability, and even data rate.

Recently, new mmWave front-end architectures such as true-time delay (TTD) array [133, 236, 50] and leaky-wave antenna [88] have been proposed for frequency-dependent beamforming that spread different frequencies in different directions. However, these beam patterns have limitations in that each user only receives a tiny fraction of the bandwidth in one timeslot, which may not meet their demand in a reasonable time frame. Additionally, these architectures do not provide control over the number of beams, beam directions, and ***beam-bandwidth***[1], resulting

[1]**Beam-bandwidth** is defined as a fraction of system bandwidth that has high beamforming gain in the desired

in large chunks of frequency resources being wasted in directions where there is no active user, leading to low spectrum utilization.

We propose a system called mmFlexible that performs *flexible directional-frequency multiplexing* by allocating a sub-set of contiguous frequency resources to each user, regardless of their direction, preventing spectrum wastage. This is achieved by creating multiple concurrent pencil beams (multi-beams) in different user directions, where each beam carries a separate frequency band according to the user's demand in that beam direction. A key feature of mmFlexible's multi-beam response is that it preserves beamforming gain across all beams while re-distributing power to only desired frequency-direction pairs with minimal leakage in other directions and frequency bands. This set of frequency-direction pairs can be chosen arbitrarily, providing flexibility in performing directional-frequency multiplexing. This enables efficient use of the entire frequency band (up to 800 MHz for 5G NR and 2.3 GHz for IEEE 802.11ax bands), reducing spectrum wastage and providing low latency network access, as shown in Figure 2.1(b).

To implement mmFlexible, we introduce a novel mmWave analog array architecture called delay phased array (DPA) that can generate multi-beams with a flexible number of beams, with arbitrary beam directions and bandwidth parts (Refer to Figure 2.2). Unlike fixed delay architectures, DPA uses variable delay and phase elements at each antenna to create any desired frequency-direction response. Our insight is to use delays and phases in a complementary manner, with variable delay providing *frequency selectivity* and variable phase providing *direction steerability*, allowing DPA to create multi-beams towards multiple arbitrarily chosen frequency-direction pairs.

We architect the design of DPA and provide insights on the hardware requirements for creating the antenna array, and perform an analysis on the range and resolution of variable delay values required to generate our desired multi-beams. One of the challenges in implementing large delays on circuits is the size and complexity of the transmission lines, and integrating them onto an IC becomes even more difficult at mmWave frequencies due to bandwidth and matching

beam direction and low elsewhere.

constraints [86]. Our design addresses this by significantly reducing the range of delay values required, compared to traditional TTD array designs, making it practical and easy to manufacture. For example, TTD array requires monotonically increasing delays at each antenna, requiring a delay of 20 ns even for a 16-element array, which is 13x more than state-of-the-art mmWave delay designs [86]. In contrast, our DPA design consists of increasing and decreasing delay values for consecutive antennas, resulting in lower delay range requirements than traditional TTD arrays. Furthermore, the delay range is independent of the number of antenna elements, making it scalable for large arrays.

The next challenge is the software programming of DPA to meet the beamforming requirements of mmFlexible. The software should determine the appropriate values for the variable delays and phases at each antenna of DPA. Solving for these discrete values is computationally difficult as it is a non-convex and NP-hard problem. A naive solution is to pre-compute and store the delay and phase values for every possible frequency-direction pair, but this is infeasible due to a large number of such combinations (10^{28}). To overcome this, we develop a novel FSDA (frequency-space to delay-antenna) algorithm which provides a single-shot solution for estimating the delays and phases in real-time. It does this by mapping the desired frequency-space response to the delay-antenna space using a 2D transform and then extracting the corresponding delays and phases for each antenna. FSDA algorithm can be implemented in real-time using fast and efficient 2D FFT techniques. Our algorithm only requires the angles for each user and corresponding frequency resources allocated to users in the current time slot. The angles can be obtained from any standard compliant initial access protocol and so mmFlexible can be easily integrated into the standard 5G mmWave protocols.

In summary, ***we make the following contributions***:

- We propose mmFlexible, the first system that enables flexible directional-frequency multiplexing in mmWave networks, achieving higher spectrum usage, low latency, and scalability to support a large number of users.

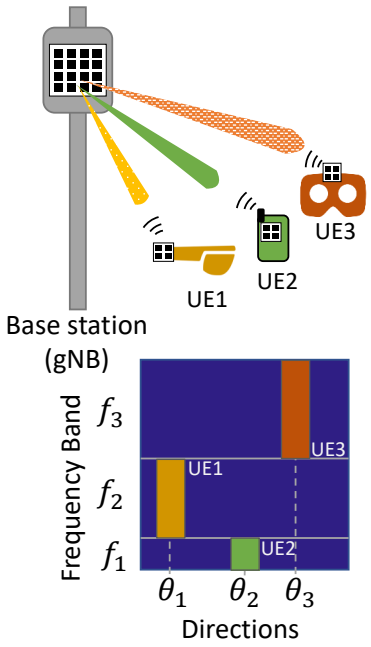

Figure 2.2. Desired frequency-space beam response: 3 users in 3 different angular directions.

- We design a novel mmWave front-end architecture called delay phased array that can generate multi-beams with a flexible number of beams, beam directions, and beam-bandwidths while maintaining high beamforming gains.
- We provide a new algorithm called FSDA (Frequency-space to delay-antenna) which estimates delays and phases in real-time using 2D FFT techniques and can be easily integrated with standard 5G mmWave protocols.
- We evaluate the performance of mmFlexible using real mmWave traces and show an improvement in latency by 60-150% compared to baselines. Furthermore, the multi-user sum-throughput is improved by 3.9x compared to true-time-delay array baseline [133].

2.1 Background and Motivation

In this section, we would provide a primer on different analog antenna array architectures, which have a single baseband radio frequency (RF) chain and mechanism to best achieve flexibility in resource allocation to multiple users. Next, we would show a realistic example

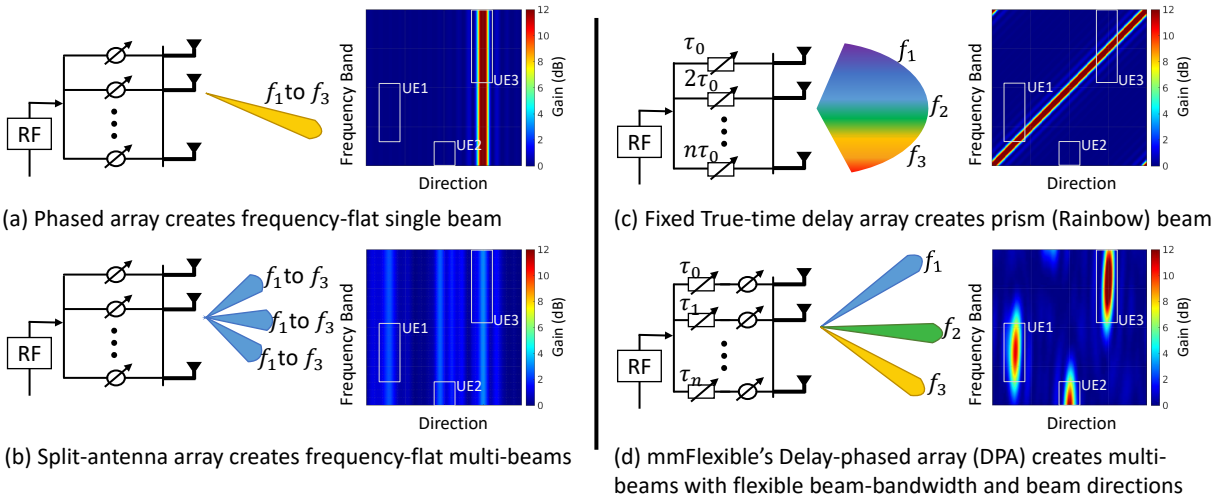

(a) Phased array creates frequency-flat single beam

(c) Fixed True-time delay array creates prism (Rainbow) beam

(b) Split-antenna array creates frequency-flat multi-beams

(d) mmFlexible's Delay-phased array (DPA) creates multi-beams with flexible beam-bandwidth and beam directions

Figure 2.3. Comparison of different mmWave front-end architectures: Corresponding frequency-space patterns they created. The traditional phased array and split-antenna array create a frequency-flat response, while TTD array creates a frequency-selective rainbow-like pattern radiating in all directions. In contrast, DPA provides a flexible frequency-direction response with a programmable number of beams, beam-bandwidth, and beam directions.

(Figure 2.2) that none of the architecture can meet the requirements for flexible directional-frequency multiplexing, and how our flexible DPA architecture can meet these requirements.

2.1.1 Primer on phased arrays and true-time delay arrays

■ **Phased array:** The phased array takes a single input from the digital chain, split into N copies, apply appropriate phase-shift, and radiates from N antennas (Figure 2.3(a)). The input signal with all of its constituent frequency bands is radiated in a direction specified by the phase setting. The set of phases at each antenna constitutes a weight vector w_{phase} as:

$$w_{\text{phase}}(n) = e^{j\Phi_n} \tag{2.1}$$

where Φ_n is the programmable phases for antenna index $n (n \in [0, N-1])$. Now, if we program the phases to create a directional beam at an angle say 30°, then this phase shift is applied to all the frequency components in the signal radiating them along that same angle 30°. Different frequency components in the signal cannot be radiated to other directions because of the frequency-

independent nature of weights in a phased array. Therefore, a phased array can serve users in only one direction at a time and cannot perform flexible directional-frequency multiplexing with many users. Split-antenna phased array (Figure 2.3(b)) uses the phased array architecture split into multiple sub-arrays to create multiple beams towards different users in one TTI, but suffers from lower antenna gain, range, and throughput.

■ **True-time delay array:** The true-time delay (TTD) array uses a delay element to replace the phase shift element as shown in Figure 2.3(c). The antenna weights with delay elements is given by:

$$w_{\text{delay}}(t,n) = \delta(t - \tau_n) \tag{2.2}$$

Past work on TTD array [236, 49, 50, 133] have shown that this architecture creates a prism-type response by radiating all the frequencies in all the directions in a linear fashion (Figure 2.3(c)) by using a fixed set of delays at each antenna given by $\tau_n = \frac{n}{\text{B}}$, for bandwidth B. This response has two major limitations: 1) most frequency bands would be wasted in space if there are no user in that direction and 2) each user gets a small frequency band and cannot scale it up arbitrarily. Due to these limitations, TTD array is not suitable for flexible directional-frequency multiplexing.

2.1.2 Current mmWave Architectures are Incapable for Flexible Directional-Frequency Multiplexing

Let us take a simple network scenario to understand the performance tradeoffs for the above systems. Typically for 4K VR applications, end-end latency should be $<$ 100 ms, Where transmission goes from the public cloud, network provider, base station (edge), and device. Over-the-air (base station to device) transmission latency comes down to a strict latency requirement of $<$ 1 ms [179]. Consider 10 users in the network with similar traffic demand: each user requires 60 Mbps (4K VR) throughput and $<$ 1 ms over-the-air latency due to the interactive nature of VR [170, 14, 179, 144]. The aggregate throughput provided by the users is 600 Mbps, only 30% of the max physical layer throughput of a 400 MHz FR2 gNB in downlink (2.2 Gbps). We assume all users are distinctly located (in different angular directions) and RF conditions

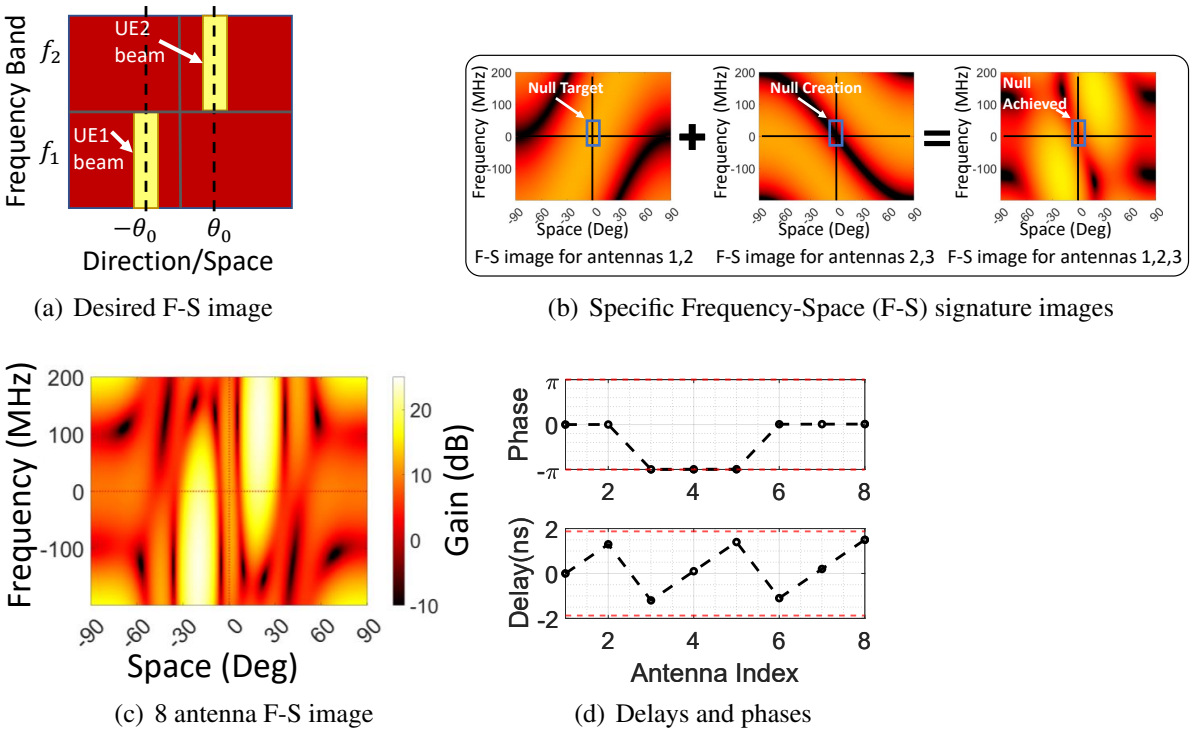

(a) Desired F-S image

(b) Specific Frequency-Space (F-S) signature images

(c) 8 antenna F-S image

(d) Delays and phases

Figure 2.4. Explanation on DPA: How DPA creates a desired frequency-space image with two beams at $-20°$ and $20°$. The first beam occupies a frequency band in [-200,0] MHz and the second beam occupies a frequency band of (0,200] MHz. DPA can create the desired image in (c) using delay and phase values in (d).

are good (e.g. LOS and no blockage). This allows us to control common external factors and exclusively evaluate architectural capabilities. We assume our DPA and split array can create up to 8 concurrent beams and TTD array creates a prism beam pattern. The comparison is shown in Table 2.1. From our end-end evaluations (Figure 2.7(a)), we can see that in edge scenarios all the baselines failed to meet the 1 ms over-the-air latency requirement. It is evident that mmFlexible is the best to support latency-critical applications while satisfying higher throughput demands.

Table 2.1. Delay-phased array & other baselines for 10 4k VR users (*We mention over-the-air worst case latency).

Architecture	Latency* (ms)	Packet Loss	Throughput
phased array (TDMA)	$>$ 1.25 ms	24.0%	54.9 Mbps
split-antenna array	$>$ 2 ms	33.3%	47.3 Mbps
TTD array	$>$ 3.5 ms	76.4%	18.3 Mbps
DPA	$<$ 0.5 ms	0.0%	71.3 Mbps

2.2 mmFlexible's DPA hardware design

mmFlexible introduces a new mmWave front-end architecture, delay-phased array (DPA), to enable flexible directional-frequency multiplexing for mmWave networks. The DPA design addresses the limitations of existing mmWave systems with phased arrays and TTD arrays by creating a multi-beam response with flexible beam directions and bandwidth parts. In this section, we discuss the practical implementation of DPA and show how we can build it with shorter delays.

2.2.1 Architecture of Delay-Phased Array (DPA)

DPA architecture consists of programmable delay and programmable phase element per antenna with a single-RF chain as shown in Figure 2.3(d). These elements can be programmed together to create flexible beam responses that are not possible by either of the two elements alone. Our insight is to control two knobs: delays τ_n and phase Φ_n to get the desired response. We define beam weights for DPA as:

$$w_{\text{dpa}}(t,n) = w_{\text{phase}}(n) w_{\text{delay}}(t,n) = e^{j\Phi_n} \delta(t-\tau_n) \tag{2.3}$$

Notice the dependence of DPA weights with time, which leads to a dependence on frequency. Upon taking FFT, the exponential term in the weights becomes $\Phi_n - 2\pi f \tau_n$, which is a function of frequency f at antenna index n. Therefore, the beamforming response of this weight vector will be a function of both frequency and direction. The beamforming gain of an antenna array represents the power radiated by the antenna array in different directions. The expression for beamforming gain $G(f,\theta)$ for DPA at a frequency f and direction θ is:

$$G(f,\theta) = \sum_{n=0}^{N-1} \mathscr{F}(w_{\text{dpa}}(t,n)) e^{-jn\pi \sin(\theta)} \tag{2.4}$$

where $\mathscr{F}$ is Fourier transform and $e^{-jn\pi\sin(\theta)}$ is the standard steering vector transformation from the antenna to space ($\sin(\theta)$)-domain[2]. Essentially, the response is the sum of the individual contribution from all antennas. By equating exponential terms of DPA weights to that of array response, we get: $\Phi_n - 2\pi f\tau_n = n\pi\sin(\theta)$. The variable delay causes a slope in frequency (f) - space ($\sin\theta$) plot, while the variable phase causes a constant shift along the space axis. Together, they can create an arbitrary line with a configurable slope and intercept in the frequency-space domain. With this insight in place, we will discuss how to program the phase and delay values at each antenna to get the desired beam response in Section 2.3.

2.2.2 Range of Delay Element in DPA

Before discussing DPA software programming, we emphasize that it is important to analyze the set of possible delay values that the hardware can support practically. Here we describe the requirements for the delay range and how it helps create a practical circuit board. Delay elements are implemented with variable-length transmission lines on a circuit board. Building large delay lines in IC at mmWave frequencies is prohibitive because of large size, bandwidth, and matching constraints [86]. Therefore, our design ensures that the delay range is not too large. While traditional true-time delay (TTD) array requires a delay range proportional to the number of antenna N, which is large for large antenna array (18.75 ns for 16 antenna array) [133]. In contrast, the delay range for mmFlexible is independent of the number of antennas. For the two-beam case, the delay range for mmFlexible is $\frac{3}{2B}$ (shown later in (2.5)), which is $3.7ns$ for 400 MHz bandwidth for 5G NR; significantly less than that required by TTD arrays. The delay range increases with the number of concurrent beams, but is independent of the number of antennas, making it scalable to large arrays.

Delay control with sub-ns accuracy has been demonstrated in full-duplex circuits for interference cancellation [153]. Recently, authors of [85] have shown accurate delay control with

[2]Note that the steering vector for a linear antenna array is given by $e^{-jn2\pi\frac{d}{\lambda}\sin(\theta)}$, which depends on the array geometry (antenna spacing d) and signal wavelength λ [47]. We assume $d = \frac{\lambda}{2}$ and approximate the steering vector as $e^{-jn\pi\sin(\theta)}$.

0.1 ns resolution and with 6-bit control (64 values until 6.4 ns), which satisfies the requirement of mmFlexible.

2.3 mmFlexible's DPA Software Design

2.3.1 Requirements for mmFlexible

Our goal is to construct arbitrary frequency-direction response $G(f, \theta)$ via DPA architecture, which would be energy efficient and enable efficient resource utilization with low latency. If we carefully notice the example in section 2, we observe that the split antenna achieves the frequency-direction mapping but with a loss of 6 dB SNR, which is a corollary of the entire 400 MHz radiated in each of the four directions. This leads to our first requirement: **Req.1:** *The system must be able to transmit/receive signals in the specific frequency-direction pairs associated with each user, with minimal energy leakage in other directions and frequencies.* Furthermore, we should be able to control the amount of bandwidth assigned to each user, which leads to: **Req.2:** *Flexibility in allocating bandwidth to each user, allowing for narrow beams in space for higher antenna gain and wide beams in frequency to support high-demand users.*

2.3.2 Meeting mmFlexible's Requirements with DPA

With this intuition in place, we revisit and explain how can mmFlexible achieve these requirements through a simple two-user example in Figure 2.4. Let us consider the two users are located at $-\theta_0$ and θ_0 respectively, and the base station wishes to serve these two users with equal bandwidth parts of $B/2$ each, where B is the total system bandwidth. To support such flexible directional-frequency multiplexing, the base station must create a frequency-direction beam response shown in Figure 2.4(a). We call such 2D beam patterns as *frequency-space (F-S) images* for simplicity. So how does DPA create these images and meet the above requirements?

We provide a closed-form expression for the set of delays τ_n and phases Φ_n for each

antenna that would generate the above beamforming response as follows:

$$\tau_n = \left(\frac{3}{2B} n \sin(\theta_0) + \frac{3}{4B}\right) \mod \frac{3}{2B} \tag{2.5}$$

$$\Phi_n = \text{round}(n \sin(\theta_0))\pi \mod 2\pi \tag{2.6}$$

A more generalized expression for an arbitrary number of beams, beam directions, and bandwidth part along with their proof is presented in the next FlexLink chapter 3.

Now, we achieve the requirements for mmFlexible by assigning a complementary F-S images to a subset of antennas. For instance, we create a positive slope in F-S image using antennas 1 and 2, and then create a complementary negative slope with antenna 2 and 3 as shown in Figure 2.4(b). When the two responses are combined together, we observe a frequency-space image where they combine constructively at desired user locations while creating a null (low gain) at other locations (Meeting **Req-1**).

For the second requirement, we create such beams by choosing the number of antennas for creating a positive or negative slope. The intuition is that a higher number of antennas makes the corresponding signature image narrow in space. For instance, it is clear from Figure 2.4(d) that three consecutive antennas (e.g. antenna 3,4,5) have increasing delays, while only two antennae (e.g. 2,3) have decreasing delays. This helps in making the positive slope in the F-S image narrow (3 antenna contribution), while the negative slope remains wide (2 antenna contribution). This effect causes beams that are narrow in space, but arbitrarily wide in frequency as shown in Figure 2.4(c) (Meeting **Req-2**).

This intuition helps in understanding how a simple 2-beam frequency-space image is created. We use this insight bandwidth partsvel FSDA algorithm that estimates the delay and phase values for any frequency-space image with an arbitrary number of beams, beam directions, and beam-bandwidths.

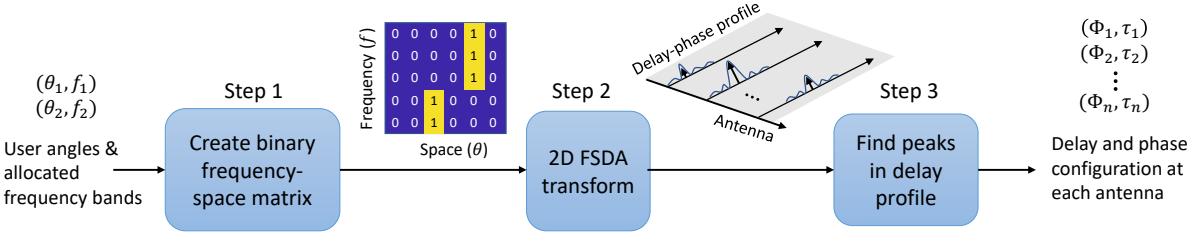

Figure 2.5. FSDA Algorithm: Three steps of our FSDA algorithm to estimate per-antenna delays and phases that generate a desired frequency-space image. The core step is our novel 2D FSDA transform with some pre-processing and post-processing of inputs and outputs respectively.

2.3.3 FSDA Algorithm for Estimating Delays and Phases in DPA

To create a desired frequency-direction beam response, the base station needs to estimate the corresponding delays and phases per-antenna in DPA. One naive solution is to try different discrete values in a brute-force way using a look-up table to get the desired beams. But, this solution is computationally hard and memory intensive since there is a large set of possibilities for the delays and phases at each antenna. For instance, with 64 delays and 64 phases per antenna (assume both are 6-bit, so $2^6 = 64$), a brute-force look-up table search would require $64^N \times 64^N \approx 10^{28}$ probes to try each combination and then storing it all in memory, which is impossible to solve with even high memory and high computing machines. Our insight is that we can pose this problem as an optimization framework and solve them in a computationally efficient way.

We now formulate the optimization problem with insights we have obtained from the previous subsection and from the fundamentals of digital signal processing. The goal is to relate weights (delays and phases) to the antenna gain pattern in (2.7) in a way that simplifies our estimation problem.

Our insight is that similar to how frequency and time are related by a Fourier transform, there is a similar transform that relates space and antenna using steering matrices. So, there are two transforms that bridges the world of antenna weights to the desired gain pattern: time to frequency transform and antenna to space transform. Mathematically, we re-write the gain

pattern of DPA to emphasize this 2D transform:

$$G(f,\theta) = \sum_{k=0}^{K-1} \sum_{n=0}^{N-1} \mathbf{U}(f,k) w_{\text{dpa}}(k,n) V(n,\theta) \tag{2.7}$$

where $\mathbf{U}(f,k)$ is a discrete domain Fourier transform (DFT) and the steering matrix V is defined per-element as $V(n,\theta) = e^{-jn\pi\sin(\theta)}$. Now since the signal is actually sampled only discretely with a sampling time of T_s, our original delay weight element $w_{\text{delay}}(t,n)$ would reduce to $w_{\text{delay}}(k,n) = \delta(kT_s - \tau_n)$ as described in (2.7).

We then represent the gain pattern by a discrete frequency-space matrix $\mathbf{G}$ and the weights as discrete time-antenna matrix $\mathbf{W}$ and relate them with the following 2D transform:

$$\mathbf{G} = \mathbf{UWV} \tag{2.8}$$

where $\mathbf{U}$ is time to frequency transform matrix and $\mathbf{V}$ is antenna to space transform matrix. Here we formulate $\mathbf{W}$ as $K \times N$ matrix, where K is the number of discrete time values and N number of antennas.

We follow a three-step process to estimate the weight matrix $\mathbf{W}$ that creates our desired frequency-space image intuitively explained in Figure 2.5. There are two inputs to our algorithm: Angles and desired frequency bands for each user. These two inputs are enough to represent the given frequency-space image. As a first, we create a binary frequency-space matrix that consists of 1s at desired frequency-space locations and 0s otherwise, we denote it by $\mathbf{G}_{\text{desired}}$. We then formulate the following optimization problem:

$$\begin{aligned} \hat{\Phi}_n, \hat{\tau}_n &= \min ||\mathbf{G}_{\text{desired}} - \mathbf{UWV}||^2 \\ \text{s.t.} \quad \mathbf{W}(k,n) &= e^{j\Phi_n}\delta(kT - \tau_n) \end{aligned} \tag{2.9}$$

This optimization is a non-convex due to the non-linear terms such as exponential in phase and delta in delay. Moreover, the constraint of having a discrete set of values for delays and phases

makes it NP-hard. We make an approximation by relaxing the delta constraint and letting the weights at each antenna take any variation over time. It means that we allow weights to take the form of a continuous profile over time at each antenna rather than a delta function which is non-zero at only one value and zero otherwise. We call it the delay-phase profile at each antenna. How do we estimate this delay-phase profile?

Our insight is that we can write an inverse transform of $\mathbf{U}$ and $\mathbf{V}$ to go from the frequency-space domain to the time-antenna domain. The logic behind such formulation is that using an appropriate discrete grid along the time and space axis, we can formulate $\mathbf{U}$ and $\mathbf{V}$ as linear transforms, i.e., $\mathbf{U}^{\dagger}\mathbf{U} = \mathbf{I}$ and $\mathbf{V}\mathbf{V}^{\dagger} = \mathbf{I}$ for identity matrix $\mathbf{I}$ (Note $(.)^{\dagger}$ is pseudo-inverse of a matrix). Therefore, it is easy to write their inverse by simply taking the pseudo-inverse. We estimate $\hat{\mathbf{W}}$ as:

$$\hat{\mathbf{W}} = \mathbf{U}^{\dagger}\mathbf{G}_{\text{desired}}\mathbf{V}^{\dagger} \tag{2.10}$$

The final step of our algorithm is to extract delays and phases from $\hat{\mathbf{W}}$. Note that each column in $\hat{\mathbf{W}}$ contains the delay-phase profile. We find the maximum peak in this profile and the index corresponding to this peak gives the delay and the max value at this peak gives the phase term. Note that since we did not put any restriction on the number of non-zero delay taps, we could get more than one delay tap per antenna. We empirically found that the estimated delay profile has only one significant peak with high magnitude than other local peaks (See Section 2.4). Also, the intuition comes from our insights from the previous section that usually one delay per antenna suffices in creating the desired response.

■ **Weights Quantization:** The delay and phase values obtained from FSDA algorithm are still continuous in nature and must be discretized to be fed into the DPA hardware. We quantize both the phase and delay values with a 6-bit quantizer in software before feeding to the array. The quantized phase takes one of the 64 values in $[0^{\circ}, 360^{\circ})$ and the quantized delay varies in the range [0, 6.4ns) with an increment of 0.1 ns.

■ **Computation complexity of FSDA:** The run-time complexity of FSDA is dominated by

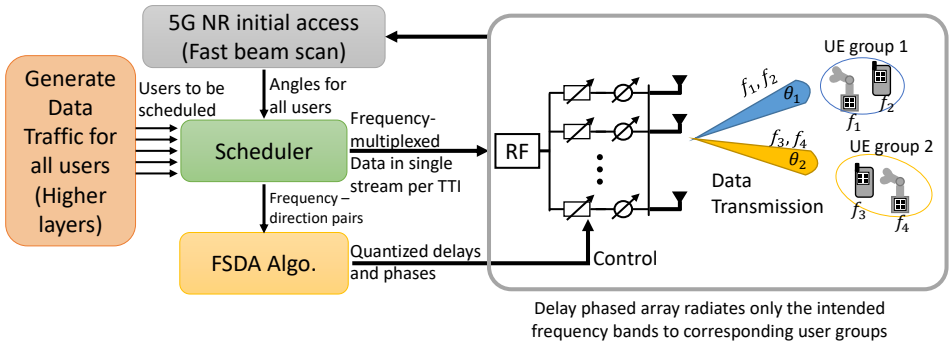

Figure 2.6. Implementation overview for mmFlexible: Four main components: Data traffic generator, Fast beam scan angle estimation, a scheduler, and our FSDA algorithm.

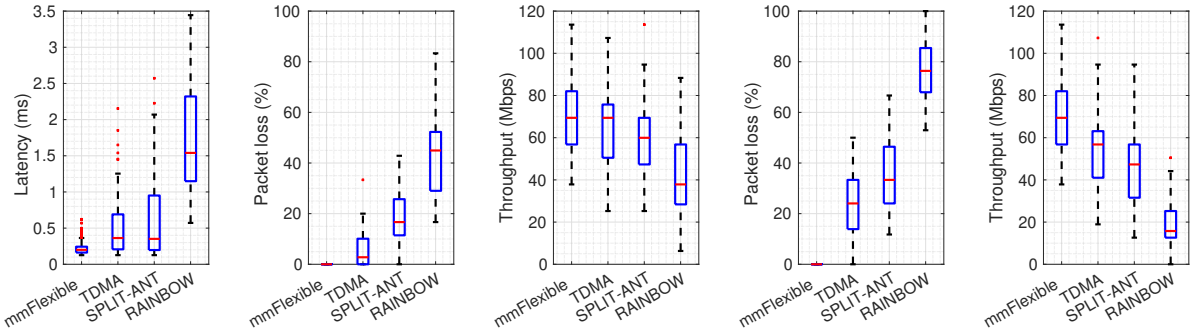

(a) Latency (b) Loss for <1ms latency (c) Thput for <1ms latency (d) Loss for <0.5ms latency (e) Thput for <0.5ms latency

Figure 2.7. End-to-End Results: Performance gain with DPA compared to baselines under different scenarios: (a) mmFlexible is able to meet strict latency constraints ($<$ 1 ms), but the baselines are not; (b) to (d) Percentage of packet loss and Throughput performance for latency constraint $<$ 1 ms and $<$ 0.5 ms respectively.

the 2D FFT transform on a given frequency-space image. Given the frequency, the axis is divided into M subcarriers and the space axis into D directions, the run-time complexity is $\mathcal{O}(MD(\log(M)+\log(D))$.

2.4 Evaluation

2.4.1 Implementation and Emulation with 28 GHz Dataset

We implement an end-end system of mmFlexible with four major components, as illustrated in Figure 2.6. The components of the systems are as follows:

1. ***Data traffic generator:*** We generate MAC layer packets with the throughput and latency

constraints mentioned in the example in Sec. 3.1 based on [144, 14, 179]. We test with two different latency constraints: 1 ms and 0.5 ms. We use the same traffic generator for our system and all baselines.

2. ***Users angle estimation (initial access):*** We use the *channel collected from mobile 28 GHz testbed [113]*, using switched beamforming techniques [52, 165] for user's angle estimation. We leverage the existing 5G NR SSB Beam scan [89] using an exhaustive search to estimate angles. The gNB scans 64 beams in the codebook, and each UE reports the best beam index that maximizes the received signal strength from which the gNB determines the user's angle.

3. ***Data Scheduling:*** We implement a Proportional Fair (PF) scheduler [35] to allocate spectrum resources to users on a per-TTI basis. The available 120° field of view is mapped into 10 groups with each 12° half-power beam width. The scheduler uses user grouping, demand generation information, and SNRs (mapped to CQIs) to determine which user group to support and how many subcarriers to allocate for each user. Throughput is then calculated as a function of the allocated resources and channel to each user.

4. ***FSDA Algorithm:*** Our system's front-end uses DPA, requiring delays and phases as input. The FSDA algorithm provides quantized delays and phases, which are applied to the DPA to generate beams in desired directions and frequency bands. Array gain from the FSDA algorithm or respective baselines is then fed back to the scheduler for SNR computation.

We evaluate the performance of the baselines and mmFlexible with a PF scheduler by emulating 2000 TTIs of $125\mu s$ each, for a total duration of 0.25 seconds, across all users. The system undergoes the above four-step procedure for each TTI.

■ **Baselines:** We compare the performance of mmFlexible with three baselines: TDMA, Split antennas [107], and Rainbow-link [133]. The TDMA approach has the scheduler assign one direction per TTI and beam in that direction over all frequencies. The Split antennas baseline has the scheduler assign one or multiple directions based on SNR & demand requirements

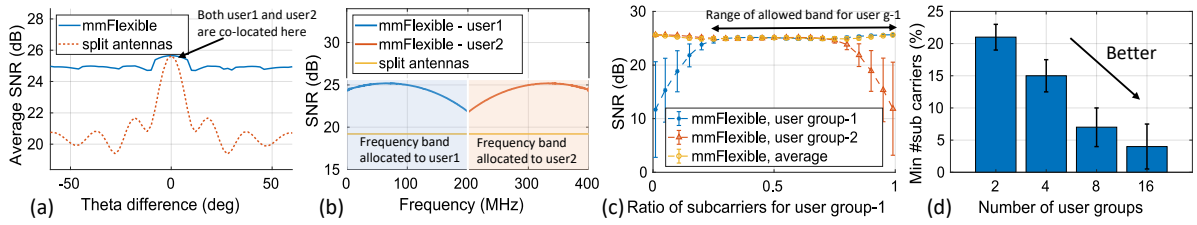

Figure 2.8. Benchmark Results: Performance gain of mmFlexible compared to baselines under different scenarios: (a) Angle separation between two users (b) (c) Subcarrier allocation between users (d) How much minimum frequency resources can be allocated without SNR degradation.

support the given directions in all subcarriers. The Rainbow-link [133] transmits in all directions regardless of the number of user directions, using only subcarriers in user directions and wasting the rest.

2.4.2 End-to-end Results

1) **Latency**: Latency is the time for a packet to travel from source (gNB) to destination (UE) over the air. We evaluate the latency distribution across the baselines and present the results in Figure 2.7(a). We see that mmFlexible has a median latency of 0.2 ms, while TDMA and Split-antenna baselines have a higher median latency of 0.32 ms and 0.26 ms respectively. Our implementation features equal offered throughput for every user direction, which is the best case scenario for Rainbow-link operation; despite this, the median latency for the Rainbow-link is 1.5 ms because of its inability to assign bandwidth to a user that is proportional to its demand. The worst-case latency of mmFlexible is well below 1 ms. Notably, all baselines have weighted right tail distributions of latency, making their worst case much worse than 1 ms. The inability of the baseline methods to honor the latency constraint leads to dropped packets and ultimately lower link throughput and reliability.

2) **Packet loss**: If a packet's latency constraint is not met, then the packet is considered undelivered and lost. In Figure 2.7(b) and Figure 2.7(d), we see that mmFlexible is able to function without any packet loss. However, due to their inability to honor the latency constraints, TDMA, Split-antenna, and Rainbow-link baselines result in a median packet loss of

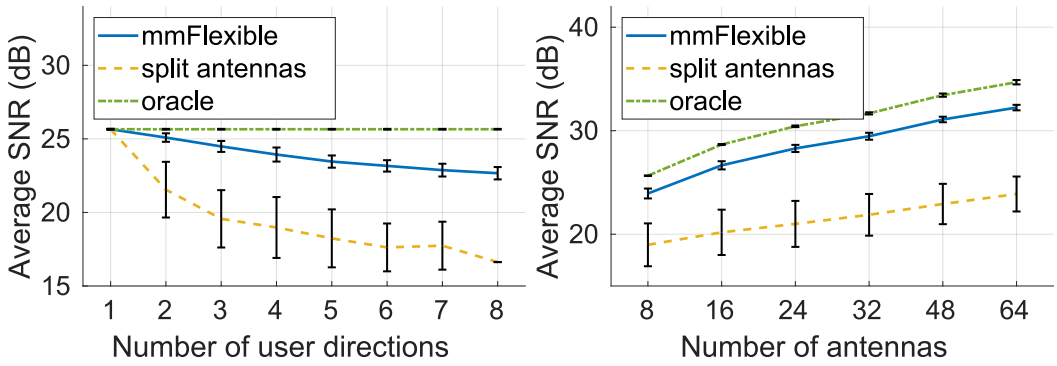

Figure 2.9. Large-scale Simulation: Scaling the mmFlexible's performance with (a) Increasing number of user directions and (2) Increasing number of antennas.

24.0%, 33.3%, and 76.4% respectively. mmFlexible's ability to serve multiple users in different directions enables it to optimally allocate resources without any power degradation and meet both throughput and latency constraints. TDMA is forced to serve one use direction at a time, resulting in a violation of latency constraints. Split-antenna baseline attempts to serve users in multiple directions simultaneously, but suffers from reduced throughput due to SNR degradation, resulting in high latency and packet loss compared to mmFlexible and TDMA. The Rainbow-link baseline allocates too few resources to each user and is the slowest and most unreliable in delivering packets.

3) **Per-user Throughput**: As shown in Figure2.7(c) and Figure2.7(e), mmFlexible outperforms all three baselines TDMA, Split antennas, and Rainbow-link. TDMA can only support one user group direction out of all the presented user demand directions, reducing its efficiency. However, in scenarios with heavy throughput demand, it performs similarly to mmFlexible. The mean throughput of mmFlexible is 1.3× that of TDMA throughput in the case of 0.5 ms latency requirement. The Split-antenna approach serves users in different directions by splitting its antennas, resulting in reduced overall throughput. mmFlexible provides 1.5× more throughput than the split scenario. Rainbow-link performs better only in scenarios with low throughput demand and users in all directions, but performance degrades in all other cases. mmFlexible provides 3.9× more throughput than the Rainbow-link baseline.

2.4.3 Benchmarks

We benchmark various theoretical and systems aspects of mmFlexible and compare the results with a split-antenna baseline. We chose the split-antenna baseline as it is closest to mmFlexible in creating multiple simultaneous beams in different directions. We present our results by calculating SNR for various scenarios. We use a channel dataset (28 GHz testbed [113]) and compute SNR per user for all subcarriers by evaluating over LOS channel model [4] with no blockage and equal distance consideration for all users.

1) **Effect of angular separation between users on SNR:** We evaluate the performance of mmFlexible and a split-antenna approach in serving two users in different locations (θ_1, θ_2) with a single RF chain and an equal number of subcarriers. Figure 2.8(a) shows the mean SNR of the two users for different angular separations $(\theta_1 - \theta_2)$. We observe that mmFlexible performs optimally in all cases compared to the baseline. When the angular separation is close to zero, both users are in the same direction and can be served by a single beam. In these scenarios, both approaches converge to a single beam and perform similarly. The baseline becomes inefficient as the angular separation increases, while mmFlexible provides a stable response after 12^o separation with a 3-5 dB higher gain than the baseline approach.

2) **Does mmFlexible have interference due to multiple user transmissions from different directions?** No, mmFlexible is designed to support frequency multiplexing of users in different directions without interference. We evaluated this by comparing mmFlexible to a split-antenna baseline for two users separated by 30^o angle and equal bandwidth, as shown in Figure 2.8(b). Both approaches receive data from user-1 in the 0-200 MHz bandwidth and from user-2 in the remaining 200-400 MHz bandwidth, without sharing any subcarriers between users. This results in no interference from simultaneous multiple-user transmissions for both approaches. Additionally, mmFlexible provides a higher gain to users by radiating only in allocated subcarriers as illustrated in Figure 2.3(d), resulting in higher SNR than the baseline. The average SNR over all subcarriers shows that mmFlexible has a gain of 5dB compared to the

split-antenna approach, even a 3dB higher gain at edge subcarriers.

3) **Impact of subcarrier allocation on SNR**: mmFlexible can reliably support both low-bandwidth *Ultra-Reliable Low-Latency Communication (URLLC)* and high-bandwidth *AR/VR* devices in different directions simultaneously. Here we evaluate the effect of subcarrier allocation on SNR for two user groups. Figure2.8(c) shows the SNR achieved with subcarrier allocation to user group-1 from 1% to 99% (remaining subcarriers are allocated to user group-2). The average SNR over all users remains at 25 dB and does not depend on the proportion of subcarriers allocated to individual user groups. It is observed that the SNR of user group-1 converges to the average SNR when it is allocated more than 20% of the total subcarriers. Therefore, to achieve the benefits of mmFlexible in a URLLC link, at least 20% of the total subcarriers should be allocated to a user group. Additionally, as presented in Figure2.8(d), as the number of user groups or user directions increases, the convergence point occurs at a lower percentage of subcarriers allocated. For instance, when there are 4 users in the system, then each user can be allocated with a minimum 15% of subcarriers without degrading the SNR performance. This convergence point drops down to 4% for 16 users which is favorable for URLLC applications which requires low-bandwidth and low latency.

Theoretical Baseline (Oracle): Oracle is a theoretical entity that can perfectly transmit only in desired subcarriers without any degradation in power at edge subcarriers (as shown in Figure 2.2 desired frequency-space response).

4) **Impact of the number of user directions**: The performance of the mmFlexible improves as the number of supported user directions increases. We tested the system using a base station antenna array with 8 transmit/receive antennas, varying the number of user directions from 1 to 8. Figure 2.9(a) illustrates that the relative gain (difference between the average SNR of mmFlexible and split antennas) increases with an increase in user directions. As the number of user directions increases, the split antenna approach divides the antennas per beam, resulting in reduced gain. Conversely, mmFlexible transmits power only in desired frequency bands and angular directions resulting in a higher gain. The Oracle creates a digital frequency filter

at each antenna, resulting in a perfect frequency-space slicing which ensures that the average SNR remains constant regardless of whether it serves in one direction or multiple directions simultaneously. In contrast, mmFlexible has one delay per antenna (for hardware feasibility), which makes it difficult to create an ideal frequency-space slicing, leading to power degradation at the edge subcarriers. Despite this, even after eight splits, the degradation is less than 2.5 dB with the Oracle, and the gain is more than 6 dB higher compared to the split antenna baseline. Error bar in Figure 2.9(a) indicates average SNR variations with users in different angle separations.

5) **Impact of the number of antennas**: We show that mmFlexible performs better with the increase in the number of antennas. We evaluated this hypothesis by serving four users and varying antennas from 8 to 64. Figure 3.15(d) shows gain variations from 8 antennas to 64 antennas for 4 users; it is clearly evident that mmFlexible outperforms with the increase in antennas over the split antennas baseline. The error bar in the figure shows the variations in the average SNR when serving four users at different angle separations. Similar to the Oracle, mmFlexible's performance remains constant even as the number of antennas increases because the number of frequency splits is determined solely by the number of user directions, which are the same in all cases.

2.5 Related Work

mmFlexible builds upon previous work in mmWave and THz communications, but sets itself apart by introducing a system that can perform flexible directional-frequency multiplexing, while maintaining energy efficiency and high performance in terms of range, throughput, and link reliability. To the best of our knowledge, no existing literature has achieved this level of frequency slicing without compromising on performance. mmFlexible's unique approach enables efficient use of the entire frequency band, reducing spectrum wastage and providing low latency network access.

■ **Split antenna phased array:** Traditional phased array beamforming does not support flex-

ible directional-frequency multiplexing because of a single narrow pencil beam. In the past, split-antenna arrays have been used to create concurrent multi-beams across multiple directions [112, 39, 242, 98, 258]. However, this approach often results in lower beamforming gain and throughput for each user. The array gain reduces proportionally to the number of beams as the total available power is distributed along multiple directions and across the entire bandwidth. In contrast, mmFlexible uses a unique split-beam mechanism with frequency selectivity that radiates only in the desired frequency band, preserving high directivity, signal strength, and throughput.

■ **True-time delay array architecture:** Previous work on True-time delay arrays (TTD) has primarily focused on beam steering for ultra-wideband signals [185, 81] and, more recently, single-shot beam training [236, 222, 50], compressive channel estimation [48], wideband tracking [211] and THz communication [210]. However, none of these works address the problem of flexible low-latency multi-user communication. Rainbow-link [133] uses TTD arrays for multi-user communication, but it is limited to fixed low-throughput IoT applications (limited to 7.8 MHz per user [133]) and cannot flexibly allocate a large number of subcarriers to a single direction for broadband users. mmFlexible addresses this limitation by using variable delay elements to create arbitrary frequency slicing and the ability to radiate those frequencies in any desired direction. mmFlexible's beamforming is orthogonal to previous work in this area, but it can leverage the fast beam training capability of TTD arrays.

Recently, a new approach called Joint Phase-Time Array (JPTA) has been proposed for creating frequency-dependent beams [182] that is similar to mmFlexible in its goals. However, mmFlexible and JPTA have several key differences in their implementation and performance. Firstly, JPTA does not provide any details on the specific hardware architecture and the range of delay values required to support their beamforming patterns. mmFlexible is the first work that provides a clear explanation of the DPA hardware architecture and the range of delay values required to support our beamforming patterns. We have shown that our DPA patterns can be realized with shorter delays and independent of the number of antennas making it practical and

realizable with current technology, and scalable to large antenna arrays. Secondly, JPTA requires digital processing over each frequency subcarrier, while DPA programming is fully analog and does not require any digital programming over the frequency subcarriers. Finally, JPTA has a high run-time complexity for estimating the delay values which makes it inadequate for real-time operation, while mmFlexible provides a closed-form mathematical expression for delays that can be used as plug-and-play with a constant run-time complexity.

■ **Other front-end mmWave architectures:** In [83], a new mmWave receiver architecture for frequency multiplexing was proposed, which utilizes a network of mixers at each antenna to receive different frequency components from different directions. However, this approach has a fixed hardware structure that is not scalable to support a large number of users, requiring different hardware for different numbers of users. In contrast, mmFlexible's delay-phased array (DPA) is flexible and programmable, providing a scalable solution for supporting a large number of users.

2.6 Discussion and Future Work

We discuss potential future work ideas:

■ **Circuit for DPA:** Implementing circuit delays in mmWave frequencies is challenging due to non-linearity, bandwidth and matching constraints [61, 102]. In contrast, delay elements can be more accurately implemented at intermediate frequencies (IF) (sub-6 GHz) using techniques such as voltage-time converters [85] and switched-capacitor arrays [86]. Recent work [86] has shown that efficient mixers, phase-shifters, and IF true-time-delays can be used to make a DPA that meets the requirements of mmFlexible.

■ **Single RF vs Multi-RF systems:** mmFlexible works with a single-RF chain and radiating each frequency component in different directions, while past work on the single-RF system stream all frequencies in one direction. Multi-RF systems (Hybrid arrays) [197, 245] offer freedom to create multi-stream to multi-directions, but each stream carries the entire bandwidth. Thus they

also cannot perform flexible directional-frequency multiplexing. Our work can be extended to the multi-RF system where each RF is connected to a DPA that creates large sectors where each DPA can serve in a sector with our flexible directional-frequency multiplexing technology.

■ **Applications beyond flexible frequency multiplexing:** Delay-phased arrays have the potential to enable a plethora of applications in communication and sensing beyond flexible directional-frequency multiplexing. For instance, the ability to create arbitrary and controllable frequency-space beams can help faster localization and tracking of multiple targets. Delay-phased arrays also enable simultaneous communication and sensing paradigms where some frequency bands are used for communication while other bands can be used for sensing. All these applications can be enabled with simple software or firmware updates on the same underlying hardware. We leave these applications for our future work.

2.7 Conclusion

mmFlexible presents a new delay-phased array architecture, which can enable frequency-dependent directional beamforming in millimeter-wave networks. It can create multiple frequency slicing, with each slice can be radiated in the different spatial direction with high beamforming gain. mmFlexible achieves this by developing the DPA architecture to create arbitrary frequency-space radiation patterns, and only radiating in the desired directions. mmFlexible leverages the frequency-space patterns to improve frequency-time resource scheduling for mmwave communication and significantly outperform the base-lines, and pave the way to make 5G mmwave scalable to mutli-user scenarios and enable new low-latency applications.

2.8 Acknowledgements

We are grateful to the anonymous reviewers their valuable feedback, as well as to the WCSNG group at UC San Diego for their input. The research was supported by NSF #2211805.

Chapter 2, in part, is a reprint of the material as it appears in **IK Jain**, RR Vennam,

R Subbaraman, D Bharadia, “ mmFlexible: Flexible Directional Frequency Multiplexing for Multi-user mmWave Networks”, *IEEE Infocom 2023*. The dissertation author was the primary investigator and author of this paper.

Chapter 3

FlexLink: Decoupling Control and Data Beams in mmWave Networks

Millimeter-wave bands in Frequency Range 2 (FR2) are vital for next-generation 5G and beyond networks to support high throughput and low latency applications such as autonomous vehicles, industrial IoT, XR streaming, etc. To facilitate communication at these bands, 3GPP adopts OFDMA to optimize time and frequency resources (in the form of resource blocks or RB) with various control and data signals. However, current mmWave systems suffer from low spectrum utilization and cannot fill all orthogonal RBs with control and data signals due to a constraint posed by directional links. Millimeter-wave systems deploy directional beams via a phased array to combat high path loss at these frequencies, but they radiate all RBs in a single direction, meaning all the control and data is radiated in one single direction. The problem is that the SSB control beam has to cover all 360 space since a new user may appear at any angle around the base station. However, the data beam has to be focussed towards the active user direction, which may not be located along the SSB control beam direction. For instance, when the control beam is at 30 degrees, and there are no active users at 30 degrees, then the base station can only send control and not data signal in that time slot. Additionally, since these control signals occupy a small number of resource blocks, the remaining RBs could go unused/wasted resources. For instance, SSB control signal requires only 7% of the entire 400 MHz band; the remaining 93% RBs are unused and get wasted as they are unable to serve any user with data

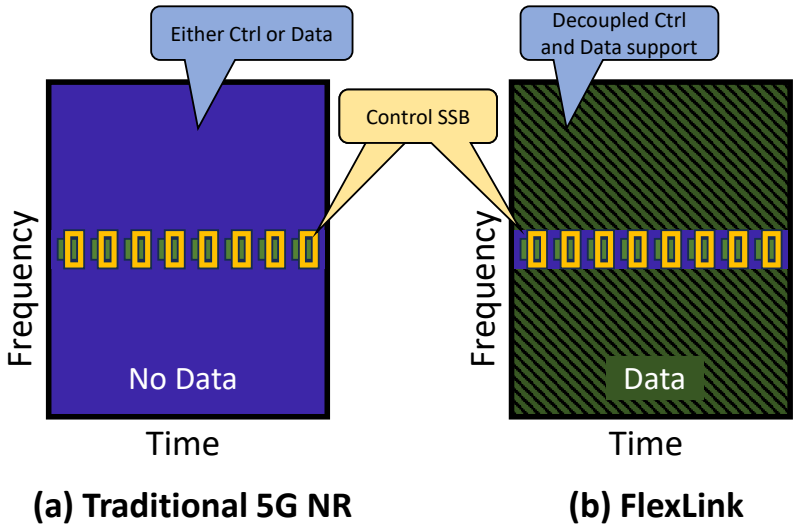

Figure 3.1. FlexLink designs and implements a novel radio architecture that can decouple control and data signaling through flexible frequency resource distribution of control and data to different directions with narrow pencil beams, much like a flexible configurable prism.

communication [27]. This frequency-direction constraint leads to the coupling of control and data signals.

To decouple control and data signals, multi-beams architectures are proposed, that can support multiple directions simultaneously, i.e., one direction for control signals and the remaining directions for data communication in a single time slot. Traditional phased arrays can be used to create multi-beams, but since they utilize an antenna-splitting mechanism [112, 39, 40], they compromise performance in terms of lower throughput and coverage. Splitting a beam into two reduces beamforming gain by half for each split beam to preserve the total radiated power. TTD-based architecture is proposed to create infinitely many beams at the same time by spreading each frequency subcarrier to different directions like a prism [133, 236, 50, 109, 182]. But this extreme of infinitely many beams radiates in all the directions even those which may not have an active user, thus wasting spectrum resources. Moreover, each direction receives signals from a tiny fraction of bandwidth, which is not enough to schedule a wideband signal transmission for data signals. Therefore, there is a gap in the literature to provide support for decoupling control and data signals without wasting precious spectrum resources.

In this chapter, we propose FlexLink, a novel multi-beamforming system to decouple the control and data beams in a flexible, 5G-compliant manner, that achieves high spectrum utilization. FlexLink uses an antenna array architecture called delay-phased array (DPA) [109, 182]. The DPA architecture consists of both delays and phases, unlike traditional phased arrays which only have phase elements, or TTD arrays which only have delay elements. Recent work on DPA shows the ability to create flexible frequency-dependent multi-beamforming, with the ability to stream a subset of frequencies in one direction and another subset in another direction until all frequency resources are utilized without wastage. While this architecture has been previously used for multi-user communication with frequency multiplexing, we propose a new application for decoupling control and data signals, while meeting more stringent requirements for this application. We particularly make three important contributions:

1) Fast configuration of multi-beams: First, we develop a new optimization framework to estimate delays and phases per antenna to create a desired multi-beam response, for instance, one beam for control and another for data at the same time using orthogonal frequency RBs. Unlike previous iterative methods, which suffer from high computation complexity ($O(MN\log(MN))$ for M directions and N frequency subcarriers), our formulation delivers the same result in one shot with O(1) complexity. We achieve this fast estimation by approximating an NP-hard optimization problem as an L2 norm minimization problem with non-linear constraint and solving it to a closed-form expression. To the best of our knowledge, this is the first work to achieve a closed-form formula for delay and phase values to program DPA.

2) Supporting a wide range of bandwidth parts: Second, the multi-beams should support both control and data signals with a wide range of bandwidth parts. The control signal may require as small as 7% of bandwidth, while the data may require a large 93% of bandwidth. So the hardware should support such diverse bandwidth parts. We show that the traditional solution is limited to a minimum of 20% bandwidth part support. Lower than 20% bandwidth part, the beamforming gain degrades exponentially and cannot be supported. In contrast, we optimize the framework to provide support for this wide range of bandwidth-parts while supporting less

than 10% bandwidth-part for the control beam, while the remaining 90% goes to the data beam.

3) Designing hardware prototype for FlexLink: Finally, we are the first to build a hardware prototype for DPA and demonstrate the performance of frequency-dependent multi-beams in real-world settings with over-the-air experiments. We discuss various design choices in building this hardware prototype, including the impact of quantization in delay and phase values, the impact of range and resolution of delay and phase values per antenna, the methodology to achieve wideband delay element using a switched array architecture, and system integration. We provide details on our PCB for control and power distribution and FPGA implementation for fast beamforming, which will be open sources to the community.

FlexLink Evaluation Overview: We verify the performance of FlexLink through both system and circuit-level simulations and over-the-air hardware measurements. We compare antenna gain and throughput performance with two baselines: One DPA-based heuristic algorithm which suffers from high computational complexity and another phased array multi-beam baseline which suffers from low SNR and throughput. We show that FlexLink is able to create any arbitrary multi-beam response with a configurable bandwidth part with O(1) complexity, improving the spectral efficiency of both control and data beams simultaneously.

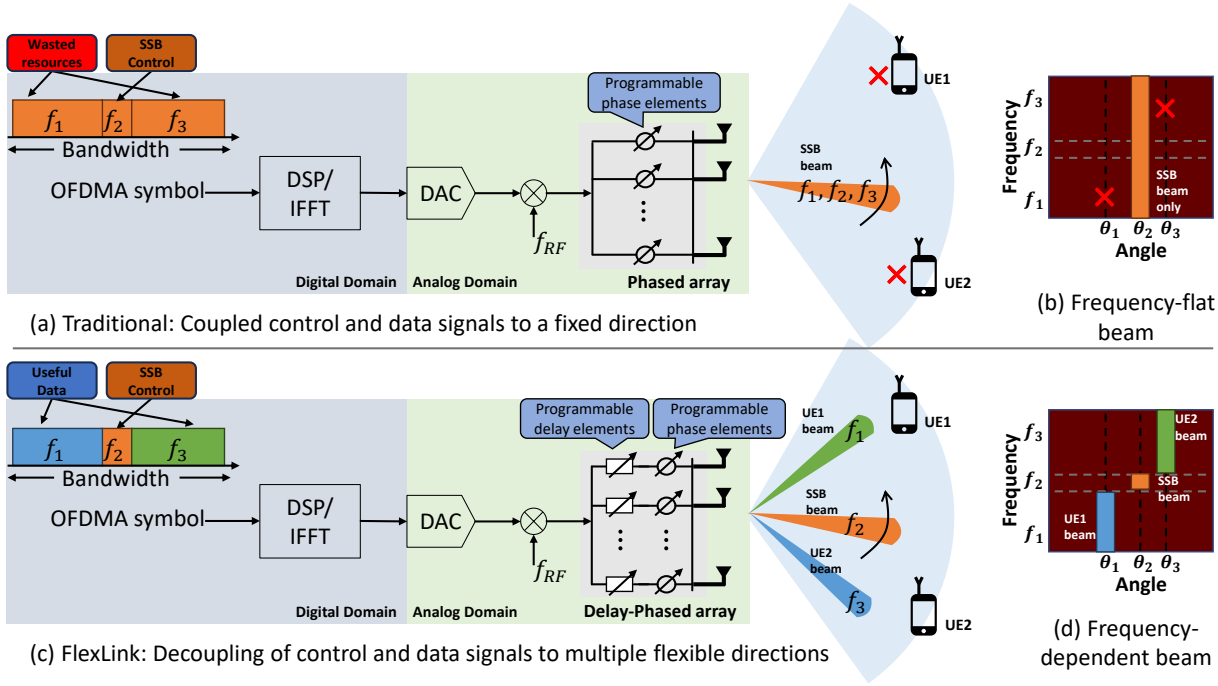

Figure 3.2. Delay-phased array (DPA) is an analog front-end architecture with programmable delay and phase elements.

3.1 Background and Motivation

3.1.1 Primer on 5G NR SSB Control Signals

Synchronization Signal Blocks (SSBs) are essential in 5G NR systems for synchronization and cell search, especially in mmWave frequencies [27]. They help user equipment (UE) synchronize with the cell, providing initial timing and frequency information. Multiple SSBs are transmitted in specific beam directions to cover the entire 360 degree search area, allowing the UE to find the strongest signal and report its quality to the base station. The base station then adjusts its beamforming parameters to optimize communication. This is crucial in mmWave communication due to narrow beamwidth and factors like user mobility and signal blockage.

However, using SSB introduces significant control signaling overhead, consuming both time and frequency resources. In the time domain, SSBs are transmitted periodically, ranging from 20ms to 120ms, and up to 64 SSBs are used in a burst of 5ms window. In the frequency

domain, SSBs occupy a dedicated bandwidth of 12 resource blocks (RBs), translating to a 17.3 MHz bandwidth out of a 400 MHz system bandwidth, leading to only 7% resource utilization [27]. The remaining 93% RBs could have been used for data signals, but they often go unused/wasted due to strong coupling between control and data signals posed by traditional analog front-end hardware.

3.1.2 Problem Statement: Coupled Control and Data Signaling

Despite having distinct and non-overlapping resource blocks allocated to control and data signals, they are coupled together to radiate in a single beam direction due to the constraints posed by traditional phased arrays. For example, as shown in Figure 3.2(a), three non-overlapping bands (sequence of contiguous non-overlapping subcarriers), namely f_1, f_2, f_3 are all radiated to a single SSB beam angle θ_2, through a single RF-chain based analog phased array, where there is no active user. The active users at angle θ_1 and θ_3 are unable to get served in this TTI. Thus a phased array-based RF front-end naturally suffers from coupling of control and data signals to a single direction.

3.1.3 Motivation: Decoupling Control and Data with DPA

The control and data signals are coupled to the same angular beam due to the hardware constraint imposed by traditional phased arrays. To decouple them, we want the capability to stream the SSB control beam towards angle θ_2, while the two UE's data beam towards angle θ_1 and θ_3 as shown in Figure 3.2(c,d). A naive way to decouple control and data signals is to deploy a wider beam or multiple split beams, hoping to find an active user over a wide angular region. However, this approach would lead to a loss in beamforming gain and ultimately lower throughput and coverage. Every split in beam reduces the gain by half [112]. In FlexLink, we ask the question: can we split the narrow pencil beam into different arbitrary directions for control and data signals without compromising on the beamforming gain in any of the directions? At first glance, it sounds counter-intuitive since there is no alternative source of power to radiate in

different directions when the total radiated power remains the same. We answer this dilemma with a novel antenna array architecture called delay-phased array (DPA) [109].

Unlike traditional phased arrays which consist of only a phase shifter element, the delay-phased array consists of both programmable phase elements and programmable delay elements [109]. The key idea behind this architecture is that the additional delay element, say τ_n at antenna index n generates a frequency-dependent phase of $2\pi f \tau_n$, which is a function of frequency f and antenna n. Thus, different frequency components in the signal are applied with different beamforming vectors that will make them radiate in different directions. Therefore, DPA is the most suitable architecture for FlexLink's application of decoupling control and data beams, allowing efficient resource utilization that leads to higher throughput, and lower latency, without breaking compliance with 5G NR protocols.

3.2 Design for FlexLink

FlexLink utilizes a novel DPA antenna array to decouple the control and data beams to different desired directions. A key component of FlexLink is to program the delay and phase values in DPA to generate desired multi-beam patterns that maximize the radiated energy towards the SSB control beam direction as well as the data beam directions, simultaneously. Past work formulates this as an estimation problem for the per-antenna delay and phase values [109]. Their iterative methods suffer from high computational complexity that hinders the practicality of DPA in a real-world system. In contrast, FlexLink demonstrates a novel closed-form estimation for delay and phase values with O(1) complexity, making them as easy to program as traditional phased arrays. We then discuss our innovative solution to support an arbitrary bandwidth splitting per beam and provide valuable insights that helped us create the first 8 antenna prototypes of DPA, to move from theory to practice.

3.2.1 Fast O(1) Computation of Delay and Phase in DPA

Our goal is to program the DPA delay and phase values to generate a desired frequency-dependent multi-beam response. The number of beams in our multi-beams can be arbitrary but finite. The direction of each beam can be anywhere in the field-of-view of DPA, ideally at different angles. Importantly, each beam should carry a fraction of bandwidth (as a set of contiguous subcarriers or bandwidth parts), such that the total bandwidth parts sum to the system bandwidth. These flexible multi-beams are achieved via DPA hardware by programming the delay and phase elements appropriately. Mathematically, we model DPA as a uniform linear antenna array with the phase value Φ_n and delay value τ_n at antenna index n. The antenna weight vector $w_{\text{dpa}}(n, f)$ is then given by:

$$w_{\text{dpa}}(n, f) = e^{j\Phi_n + j2\pi f \tau_n} \tag{3.1}$$

which is a function of antenna index n and frequency f. The beamforming gain $G(f,\theta)$ of DPA as a function of the weight vector is given by:

$$G(f,\theta) = \sum_{n=0}^{N-1} w_{\text{dpa}}(n,f) e^{-jn\pi\sin(\theta)} \tag{3.2}$$

which is a function of the beamforming angle θ and frequency f, much like a 2D energy map shown in Figure 3.2(d).

Our objective is to maximize the beamforming gain along the desired multibeam directions as:

$$\max_{\tau_n,\Phi_n} \left\|G(f,\theta)\right\|^2, \; s.t.(\theta,f) \in \{(\theta_1,f_1),\ldots,(\theta_K,f_K)\} \tag{3.3}$$

where the desired K multi-beam directions are represented by set $\{(\theta_1,f_1),\ldots,(\theta_K,f_K)\}$. Ideally, we do not want to radiate any energy in the undesired bands/angles, but it is inevitable due to the undesired side-lobes that appear due to a finite number of antennas.

Now, we ask what set of delays and phases per antenna would give us the beamforming gain pattern with the desired bandwidth part and beam direction. We first consider a simple case of two beams with equal bandwidth parts of $B/2$ each, where B is the total system bandwidth. We assume the two beams are directed along $(-\theta_0,\theta_0)$ respectively as shown in Figure 3.3(a). We will later discuss a general case with arbitrary bandwidth part and beam direction.

We formulated it as an optimization problem and solved it to a closed-form expression for the set of delays τ_n and phases Φ_n for each antenna n ($n = 0,1,\ldots,N-1$) that would generate the given beamforming response as follows:

■ **Theorem 1.** ***(2-beam symmetric case)*** *The closed-form expression for the set of delays τ_n and phases Φ_n for each antenna n ($n = 0,1,\ldots,N-1$) that would generate a given two-beam*

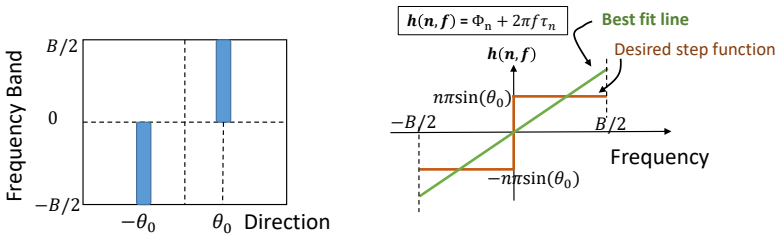

(a) Desired multi-beam response with given beam-bandwidth

(b) Solution for delay and phase is obtained from the best fit line

Figure 3.3. Proof of estimating closed-form delay and phase values for a given two-beam response at $\pm\theta_0$.

response with equal bandwidth part $B/2$ and angles $\pm\theta_0$ respectively is as follows:

$$\tau_n = \left(\frac{3}{2B} n \sin(\theta_0) + \frac{3}{4B}\right) \quad mod \ \frac{3}{2B} \tag{3.4}$$

$$\Phi_n = round(n \sin(\theta_0))\pi \quad mod \ 2\pi \tag{3.5}$$

Before proving the theorem, we want to draw two important insights from this closed-form formula. First, the computation of delays and phases is like a plug-and-play model by putting the desired angle-frequency pairs into the formula. Unlike past methods, this does not require iterative optimization, thus reducing the computational complexity to O(1), allowing fast computation in FPGAs. Second, similar to how phase is bounded by 2π, the delays are also bounded by a max value of $3/2B$, independent of the number of antennas. This shows the promise of scaling the DPA architecture to arbitrarily large antenna arrays without requiring longer delay lines that are very hard to achieve in practice (discussed in Section 3.2.4).

Proof of Theorem 1: Two-beam case We drive the expression of delays and phases per antenna that would give us the beamforming gain pattern with the desired bandwidth part and beam direction for the two-beam case.

We reformulate the optimization in (3) to include two-beam constraints as:

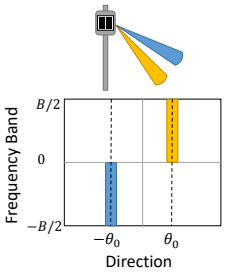

Figure 3.4. Desired frequency-space beam response for 2 users at directions $-\theta_0$ and θ_0.

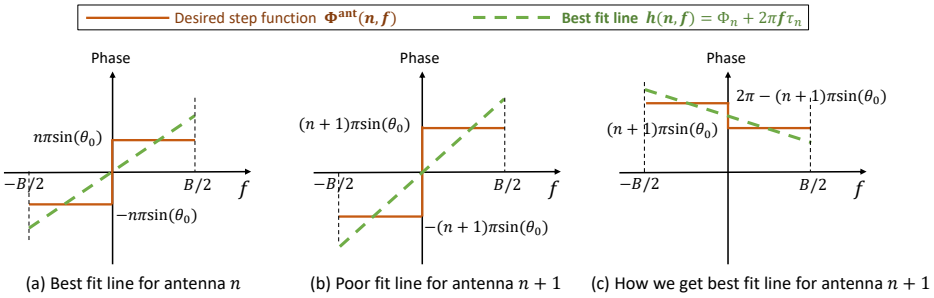

Figure 3.5. Sketch of proof: We show the phase response for each antenna is a step function with a variable step size that depends on the antenna index, n. (a) The objective is to fit a line to this step function in such a way that the error in the line fitting is minimized. (b) we show that as the antenna index increases to say, $n+1$, the step size of the step function also increases linearly with the antenna index. An increase in step size leads to higher errors in line fitting. (c) We address this issue by reducing the step size by 2π, which will not change the desired phase response but will help to reduce the error.

$$\begin{aligned} &\max_{\tau_n,\Phi_n} \left\| \sum_{n=0}^{N-1} e^{jh(n,f)} e^{-j\Phi^{\text{ant}}(n,f)} \right\|^2 \\ &s.t. \quad h(n,f) = \Phi_n + 2\pi f \tau_n \\ &\text{and} \quad \Phi^{\text{ant}}(n,f) = \begin{cases} -n\pi \sin(\theta_0) & f \in [\frac{-B}{2}, 0] \\ n\pi \sin(\theta_0) & f \in (0, \frac{B}{2}] \end{cases} \end{aligned} \tag{3.6}$$

where $h(n,f)$ is a function of variable phase Φ_n and delay τ_n at antenna n and the function $\Phi^{\text{ant}}(n,f)$ represents the constraints from the desired frequency-direction response. Let's see how these functions help us to simplify the optimization problem. Specifically, we apply triangle inequality to find an upper bound on the optimization variable and then maximize this upper bound. Triangle inequality states that the 'norm-of-sum is upper bounded by sum-of-norms', which we can apply to our optimization as:

$$\left\| \sum_{n=0}^{N-1} e^{jh(n,f)} e^{-j\Phi^{\text{ant}}(n,f)} \right\| \leq \sum_{n=0}^{N-1} \left\| e^{jh(n,f)} e^{-j\Phi^{\text{ant}}(n,f)} \right\| = N \tag{3.7}$$

So, the expression is maximized if each term in the sum is unity, i.e., $e^{jh(n,f)}e^{-j\Phi^{\text{ant}}(n,f)} = 1$, or, in other words, the two exponential terms are equal, i.e., $h(n,f) = \Phi^{\text{ant}}(n,f)$ for each antenna and for each frequency. It is impossible to achieve this solution for all frequencies because the two functions h and Φ^{ant} vary differently with frequency; h is linear, while Φ_{ant} is a step function as shown in Figure 3.5(a). The only case when the optimal results are possible is when the step size of the step function Φ^{ant} is zero, i.e., the two beams align to the same angle. In this case, the step function is reduced to a line, and optimal h can be obtained. However, this case only produces a single beam without any dependence on frequency. A natural question is how we can obtain general frequency-dependent multi-beams. We propose an optimization framework that can help to find a closed-form expression for delays and phases. Our optimization problem is formulated in a way that finds the line h that best fits the given step function Φ^{ant}. We achieve this by approximating our optimization problem in (7) to the best line-fitting on a per-antenna basis:

$$\min_{\tau_n,\Phi_n} ||h(n,f) - \Phi^{\text{ant}}(n,f)||^2 \tag{3.8}$$

We can visualize this optimization in Figure 3.5(a), where the line $h(n,f)$ is fit over the step function $\Phi^{\text{ant}}(n,f)$. The slope of the best-fit line gives the delay value, and the y-intercept gives the phase value. In this way, we can estimate both delay and phase values by solving for the best-fit line.

■ **Impact of linear approximation:** However, as antenna index n increases, the error in line fitting also increases due to the linear increase in the step size with n, as shown in Figure 3.5(b). This could lead to high error for large antenna arrays and limit our solution to scale with antennas.

■ **Minimizing linear approximation error:** We have an innovative and simple solution to address this issue. To address this issue, we utilize the concept of wrapping the phase of a signal by 2π, i.e., adding an integer multiple of 2π to the phase does not change the signal. We use this idea to strategically add a phase of multiple of 2π to a specific set of frequencies in order to

minimize the error in line fitting as shown in Figure 3.5(c). With this insight, we redefine the step function Φ^{ant} as:

$$\Phi^{\text{ant}}(n,f) = \begin{cases} k2\pi - n\pi\sin(\theta_0) & f \in [\frac{-B}{2}, 0] \\ n\pi\sin(\theta_0) & f \in (0, \frac{B}{2}] \end{cases} \tag{3.9}$$

where k is a constant integer. A natural question is how do we estimate this integer to minimize the error in line fitting? Our solution is a two-step process: we solve for the delays and phases as a function of k and then find the optimal value of k to minimize the error.

To solve for per-antenna delays and phases, we form a system of linear equations. We discretize the frequency as $f = m\Delta f$ for $m \in [-M/2, M/2]$, where the bandwidth is $B = (M+1)\Delta f$. Note that there are M frequency bins that can be a large number, i.e., $M \to \infty$ for creating a continuous frequency axis. We then formulate a set of linear equations for each frequency term to solve for the variable delay τ_n and phase Φ_n for each antenna n, given by:

$$\Phi_n + 2\pi m\Delta f\tau_n = \Phi^{\text{ant}}(n, m\Delta f)\ \forall m \in [-M/2, M/2] \tag{3.10}$$

which is in the form of a system of linear equations ($Ax = b$), with unknown $x = [\tau_n, \Phi_n]$. Solving this system of linear equations gives the desired estimate of delay and phase values.

We re-write the equations in a matrix form as follows:

$$Ax = b \tag{3.11}$$

where x is a 2×1 vector of variable phase and delay given by:

$$x = \begin{bmatrix} \Phi_n \\ 2\pi\Delta f\tau_n \end{bmatrix} \tag{3.12}$$

and the matrix A and vector b are constants given by:

$$A = \begin{bmatrix} 1 & \frac{-M}{2} \\ 1 & \frac{-M}{2}+1 \\ \vdots & \vdots \\ 1 & \frac{M}{2} \end{bmatrix} \tag{3.13}$$

$$b = \begin{bmatrix} \underbrace{k2\pi-\phi \ \ldots \ k2\pi-\phi}_{M/2} & \underbrace{\phi \ldots \phi}_{M/2+1} \end{bmatrix}^T \tag{3.14}$$

where $\phi = n\pi\sin(\theta_0)$ is introduced for simplicity. Specifically, we add $k2\pi$ to the first half of the frequency subcarriers, and this is reflected in the value of b.

The solution can be obtained by solving a system of linear equations as follow:

$$x = (A^T A)^{-1} A^T b \tag{3.15}$$

We first solve for $(A^T A)^{-1}$ and $A^T b$ separately and then multiply them together to get x:

$$\begin{aligned} A^T A &= \begin{bmatrix} 1 & 1 & \ldots & 1 \\ \frac{-M}{2} & \frac{-M}{2}+1 & \ldots & \frac{M}{2} \end{bmatrix} \begin{bmatrix} 1 & \frac{-M}{2} \\ 1 & \frac{-M}{2}+1 \\ \vdots & \vdots \\ 1 & \frac{M}{2} \end{bmatrix} \\ &= \begin{bmatrix} M+1 & 0 \\ 0 & 2\sum_0^{\frac{M}{2}} k^2 \end{bmatrix} = \begin{bmatrix} M+1 & 0 \\ 0 & \frac{M(M+1)(M+2)}{12} \end{bmatrix} \end{aligned} \tag{3.16}$$

Taking the inverse of the above 2x2 matrix, we get

$$(A^T A)^{-1} = \begin{bmatrix} \frac{1}{M+1} & 0 \\ 0 & \frac{12}{M(M+1)(M+2)} \end{bmatrix} \tag{3.17}$$

This solves for $A^T A$. Next, we obtain $A^T b$ as:

$$\begin{aligned} A^T b &= \begin{bmatrix} 1 & 1 & \dots & 1 \\ \frac{-M}{2} & \frac{-M}{2}+1 & \dots & \frac{M}{2} \end{bmatrix} \\ &\times \begin{bmatrix} \underbrace{k2\pi-\phi \ \dots \ k2\pi-\phi}_{M/2} & \underbrace{\phi \dots \phi}_{M/2+1} \end{bmatrix}^T \\ &= \begin{bmatrix} \phi + kM\pi \\ (2\phi - k2\pi)\sum_1^{M/2} m \end{bmatrix} \\ &= \begin{bmatrix} \phi + kM\pi \\ (\phi - k\pi)\frac{M(M+2)}{4} \end{bmatrix} \end{aligned} \tag{3.18}$$

This solves for $A^T b$. We now obtain the solution for unknown x as follows:

$$\begin{aligned} x &= (A^T A)^{-1} A^T b \\ &= \begin{bmatrix} \frac{1}{M+1} & 0 \\ 0 & \frac{12}{M(M+1)(M+2)} \end{bmatrix} A^T b \\ &= \begin{bmatrix} A^T b[1]\frac{1}{M+1} \\ A^T b[2]\frac{12}{(M(M+1)(M+2))} \end{bmatrix} \end{aligned} \tag{3.19}$$

Finally, using the definition of x from (3.12), we get the solution for the per antenna phase and

delay. The phase is given by $x[1]$ and delay is given by x[2] as follows:

$$\begin{aligned}\Phi_n &= \lim_{M\to\infty} x[1] = \lim_{M\to\infty} A^T b[1]\frac{1}{M+1} \\ &= \lim_{M\to\infty} \frac{\phi + kM\pi}{M+1} \\ &= k\pi\end{aligned} \tag{3.20}$$

where the approximation is taken for a large number of frequency bins, i.e., $M \to \infty$. Similarly, we get the expression for the delay:

$$\begin{aligned}\tau_n' &= \frac{x[2]}{2\pi\Delta f} = A^T b[2]\frac{12}{M(M+1)(M+2)}\frac{1}{2\pi\Delta f} \\ &= (\phi - k\pi)\frac{M(M+2)}{4}\frac{12}{M(M+1)(M+2)}\frac{1}{2\pi\Delta f} \\ &= \frac{3(\phi - k\pi)}{2\pi\Delta f(M+1)} \\ &= \frac{3(n\sin(\theta_0)}{2B} - \frac{3k}{2B} \\ &= \frac{3}{2B}(n\sin(\theta_0) - k)\end{aligned} \tag{3.21}$$

This gives a closed-form expression of delay. Note that the expression for delay and phase both depend on the unknown constant integer k. So, how do we solve for k to obtain a generalized formula for delay and phases? Out of many possible solutions, because of the presence of a random integer k, we need to choose the value of k that minimizes the error in line fitting. To solve for k, we make an observation that the step size of the step function is part of the exponential and therefore is bounded by 2π. Let's find the condition when the step size is

bounded between $-\pi$ and π as follows:

$$\begin{aligned}
& -\pi < \phi - (-\phi + k2\pi) \leq \pi \\
\Rightarrow\ & -\pi < 2\phi - k2\pi \leq \pi \\
\Rightarrow\ & -\pi < 2n\pi\sin(\theta_0) - k2\pi \leq \pi \\
\Rightarrow\ & \frac{-1}{2} < n\sin(\theta_0) - k \leq \frac{1}{2} \\
\Rightarrow\ & n\sin(\theta_0) - \frac{1}{2} < k \leq n\sin(\theta_0) + \frac{1}{2} \\
\Rightarrow\ & k = \text{round}(n\sin(\theta_0))
\end{aligned} \tag{3.22}$$

Since k is an integer, the only possible value of k is given by $\text{round}(n\sin(\theta_0))$. This gives us a unique solution for the integer constant k. Putting the in the expression of phase and delay, we get the required phase:

$$\Phi_n = \text{round}(n\sin(\theta_0))\pi \tag{3.23}$$

Similarly, we obtain the final expression of delay as:

$$\tau_n' = \frac{3}{2B}(n\sin(\theta_0) - \text{round}(n\sin(\theta_0))) \tag{3.24}$$

We notice the above expression of delay varies in the range of $\frac{-3}{4B}$ to $\frac{3}{4B}$, resulting in a total range of $\frac{3}{2B}$. Since negative delays are not possible to generate in a causal system, we can add a constant delay to all antennas to make delays positive. The minimum constant delay factor that we can add without compromising on the performance is $\frac{3}{4B}$. After adding this factor and simplifying the delay expression, we get the following:

$$\begin{aligned}
\tau_n &= \tau_n' + \frac{3}{4B} \\
&= \frac{3}{2B}(n\sin(\theta_0) - \text{round}(n\sin(\theta_0))) + \frac{3}{4B} \\
&= \left(\frac{3}{2B}n\sin(\theta_0) + \frac{3}{4B}\right) \mod \frac{3}{2B}
\end{aligned} \tag{3.25}$$

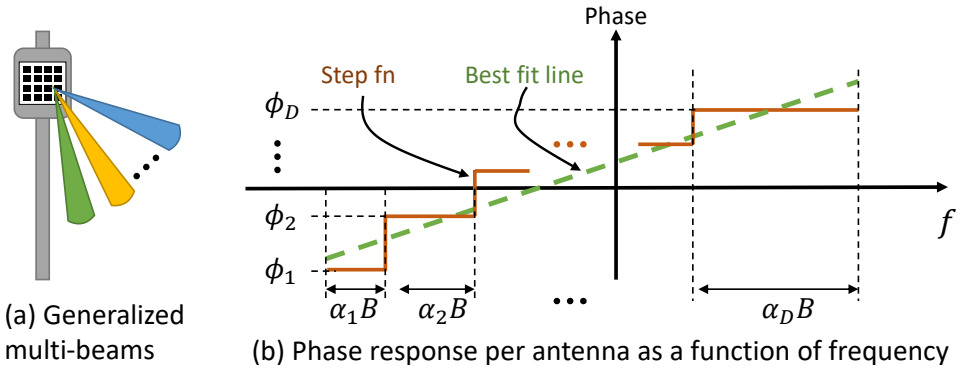

Figure 3.6. Generalized setting for multi-beams with an arbitrary number of beams D, beam angles, and bandwidth parts. The beam angles contributes to ϕ_d, and the bandwidth part is defined as $\alpha_d B$ for system bandwidth B and fraction $\alpha_d < 1$ for beam index $d \leq D$.

where the last step is a simplification of the round function into a modulo function for ease of understanding and emphasizing that the range of values of delays is independent of the number of antennas. The delay range is $\frac{3}{2B}$ for two beam case which depends inverse This is how we estimate the final expression of optimal phase Φ_n and delay τ_n in (3.5) and (3.4), respectively.

3.2.2 Generalization to an Arbitrary Number of Beams with O(1) Computation Complexity

We also show generalized beamforming response to an arbitrary number of beams with arbitrary beam directions and arbitrary bandwidth parts as shown in Figure 3.6.

■ **Theorem 2.** ***(Generalized case):*** *Let there are D beam directions with beam angles θ_d and bandwidth part $\alpha_d B$ for $\sum_d \alpha_d = 1$. We define $\phi_d = n\pi \sin(\theta_d)$ for simplicity. The per-antenna phases and delays in realizing such a generalized beamforming response is given by:*

$$\Phi_n = \sum_{d=1}^{D} \alpha_d (\phi_d + 2k_d\pi) \tag{3.26}$$

$$\tau_n = \sum_{d=1}^{D} \frac{3}{\pi B} (\phi_d + 2\pi k_d) \alpha_d \Big(\sum_{\ell=1}^{d} 2\alpha_\ell - \alpha_d - 1 \Big) \tag{3.27}$$

where,

$$\phi_d = n\pi \sin(\theta_d) \tag{3.28}$$

and the constant integer k_d for beam d and antenna n is:

$$k_d = \begin{cases} 0 \quad d = 1 \\ k_{d-1} + round(\frac{n\sin(\theta_{d-1}) - n\sin(\theta_d)}{2}) \;\; d \geq 2 \end{cases} \tag{3.29}$$

Before we provide a derivation for the generalized delay and phase expression, we first verify the expression for the two-beam case by putting $D = 2$ in the above expression. We would obtain a generic expression with an arbitrary bandwidth part and arbitrary beam directions and then simplify it to the case of an equal bandwidth part and symmetric direction similar to the previous analysis.

■ **Corollary 1.** *(Generalized two-beam case) For two beams with beam directions θ_1 and θ_2 and the bandwidth parts αB and $(1-\alpha)B$ respectively, such that the bandwidth fraction $0 \leq \alpha \leq 1$, the corresponding delays and phases that generate this two-beam response is given by:*

$$\Phi_n = \alpha\phi_1 + (1-\alpha)(\phi_2 + 2k_2\pi) \tag{3.30}$$

$$\tau_n = \frac{3}{\pi B}[(-\phi_1 + \phi_2 + 2\pi k_2)\alpha(1-\alpha)] \tag{3.31}$$

for $\phi_1 = n\pi\sin(\theta_1)$ and $\phi_2 = n\pi\sin(\theta_2)$, and the constant integer k_2 defined as

$$k_2 = round(\frac{n\sin(\theta_1) - n\sin(\theta_2)}{2}) \tag{3.32}$$

Proof of Corollary 1. We prove the corollary for the two-beam case by putting the number of beams $D = 2$ in the generalized expressions and further simplify them with $\alpha_1 = \alpha$ and $\alpha_2 = 1-\alpha$. Also, note that $k_1 = 0$, by definition. We now simplify the phase as follows:

$$\begin{aligned} \Phi_n &= \alpha_1(\phi_1 + 2k_1\pi) + \alpha_2(\phi_2 + 2k_2\pi) \\ &= \alpha\phi_1 + (1-\alpha)(\phi_2 + 2k_2\pi) \end{aligned} \tag{3.33}$$

and delay

$$\begin{aligned}\tau_n &= \frac{3}{\pi B}[(\phi_1 + 2\pi k_1)\alpha_1(\alpha_1 - 1) \\ &\quad + (\phi_2 + 2\pi k_2)\alpha_2(2\alpha_1 + \alpha_2 - 1)] \\ &= \frac{3}{\pi B}[(\phi_1)\alpha(\alpha - 1) + (\phi_2 + 2\pi k_2)(1-\alpha)(\alpha)] \\ &= \frac{3}{\pi B}[(-\phi_1 + \phi_2 + 2\pi k_2)\alpha(1-\alpha)]\end{aligned} \tag{3.34}$$

This proves the corollary. □

We will now verify the delays and phases for the special case of equal bandwidth part, i.e., $\alpha = 1/2$ and symmetric angles, i.e., $\phi_1 = -n\pi\sin(\theta_0)$ and $\phi_2 = n\pi\sin(\theta_0)$. In this case, we get $k_2 = -\text{round}(n\sin(\theta_0))$ and a simplified phase as:

$$\begin{aligned}\Phi_n &= \alpha\phi_1 + (1-\alpha)(\phi_2 + 2k_2\pi) \\ &= \phi_1/2 + \phi_2/2 + k_2\pi \\ &= k_2\pi \\ &= -\text{round}(n\sin(\theta_0))\pi \\ &= \text{round}(n\sin(\theta_0))\pi\end{aligned} \tag{3.35}$$

which is an integer multiple of π, where the integer multiple is given by $\text{round}(n\sin(\theta_0))$. This expression is the same as what we derived earlier in (3.23). We now get delay expression as:

$$\begin{aligned}\tau_n &= \frac{3}{\pi B}[(-\phi_1 + \phi_2 + 2\pi k_2)\alpha(1-\alpha)] \\ &= \frac{3}{\pi B}[(n\pi\sin(\theta_0) + n\pi\sin(\theta_0) + 2\pi k_2)\alpha(1-\alpha)] \\ &= \frac{3}{B}[(2n\sin(\theta_0) + 2k_2)\alpha(1-\alpha)] \\ &= \frac{3}{2B}(n\sin(\theta_0) + k_2) \\ &= \frac{3}{2B}(n\sin(\theta_0) - \text{round}(n\sin(\theta_0)))\end{aligned} \tag{3.36}$$

This delay expression is the same as what we have derived earlier for the simple 2-beam case in (3.24). This validates the generalized multi-beam expressions for two beams. We now prove the generalized expression for an arbitrary number of beams, beam directions, and bandwidth parts.

Proof of Theorem 2. **(Generalized case)** We follow the same formulation as before as a set of linear equations with a new set of constraints for the generalized case. We notice that the constant matrix A doesn't change for the generalized case, and only the constant vector b is modified. We would solve for delays and phases similar to the simple 2-beam case using pseudo-inverse (Recall $x = (A^T A)^{-1} A^T b$).

We start with a generic expression for vector b as:

$$b = \left[\underbrace{\phi_1 + k_1 2\pi \ldots}_{\alpha_1 M}, \underbrace{\ldots}_{\ldots}, \underbrace{\ldots \phi_D + k_D 2\pi}_{(\alpha_D) M} \right]^T \tag{3.37}$$

Therefore, solution for $A^T b$ is now:

$$A^T b = \begin{bmatrix} 1 & 1 & \ldots & 1 \\ \frac{-M}{2} & \frac{-M}{2}+1 & \ldots & \frac{M}{2} \end{bmatrix} b \tag{3.38}$$

We first solve for phase. Since the matrix A is unchanged, we directly apply $(A^T A)^{-1}$ from the two-beam case (3.20) and solve for the per-antenna phase as follows:

$$\begin{aligned} \Phi_n &= \lim_{M \to \infty} A^T b[1] \frac{1}{M+1} \\ &= \lim_{M \to \infty} \frac{\sum_{d=1}^{D} \alpha_d M (\phi_d + 2k_d \pi)}{M+1} \\ &= \sum_{d=1}^{D} \alpha_d (\phi_d + 2k_d \pi) \end{aligned} \tag{3.39}$$

This proves the expression of phase. We now solve for the delay in a similar way as we solved

for the two-beam case in (3.21) as follows:

$$\begin{aligned}
\tau_n &= \lim_{M\to\infty} A^T b[2] \frac{12}{M(M+1)(M+2)} \frac{1}{2\pi\Delta f} \\
&= \lim_{M\to\infty} A^T b[2] \frac{12}{M(M+2)} \frac{1}{2\pi B} \\
&= \lim_{M\to\infty} \Big[(\phi_1 + k_1 2\pi)(-\sum_{m=0}^{M/2} m + \sum_{m=0}^{\alpha_1 M - M/2} m) \\
&\quad + (\phi_2 + k_2 2\pi)(-\sum_{m=0}^{\alpha_1 M - M/2} m + \sum_{m=0}^{(\alpha_1+\alpha_2)M - M/2} m) \\
&\quad \cdots + (\phi_D + k_D 2\pi)(\sum_{m=0}^{\sum_{\ell=1}^{D} \alpha_\ell M - M/2} m - \sum_{m=0}^{\sum_{\ell=1}^{D-1} \alpha_\ell M - M/2} m) \Big] \\
&\quad \times \frac{12}{M(M+2)} \frac{1}{2\pi B}
\end{aligned} \tag{3.40}$$

where we put the expression for $B = (M+1)\Delta f$ for simplification. Still, the above expression is complex and messy, which we intend to write in a simplified closed-form solution. Our idea to simplify the above expression is that since there is a M^2 term in the denominator, we only collect the M^2 terms from the numerator and ignore other constant or linear terms. This is because we are interested in the case when the number of frequency bins is very high, i.e., $\lim_{M\to\infty}$. This leads to a simplified expression as follows:

$$\begin{aligned}
\tau_n &= \frac{3}{\pi B} \Big[(\phi_1 + k_1 2\pi)(-(1/2)^2 + (\alpha_1 - 1/2)^2) \\
&\quad + (\phi_2 + k_2 2\pi)(-(\alpha_1 - 1/2)^2 + (\alpha_1 + \alpha_2 - 1/2)^2) \\
&\quad \cdots (\phi_D + k_D 2\pi)((\sum_{\ell=1}^{D} \alpha_\ell - 1/2)^2 - (\sum_{\ell=1}^{D-1} \alpha_\ell - 1/2)^2) \Big] \\
&= \frac{3}{\pi B} \Big[(\phi_1 + k_1 2\pi)(\alpha_1(\alpha_1 - 1)) \\
&\quad + (\phi_2 + k_2 2\pi)(\alpha_2(2\alpha_1 + \alpha_2 - 1)) \\
&\quad \cdots (\phi_D + k_D 2\pi)(\alpha_D(\sum_{\ell=1}^{D-1} 2\alpha_\ell + \alpha_D - 1)) \Big]
\end{aligned} \tag{3.41}$$

This proves the generalized expression for phases and delays in (3.26) and (3.27), respectively, for an arbitrary number of beams, beam directions, and bandwidth part. □

We now discuss insights into how we obtain the expression for the unknown integer constant k_d to bind the error in line fitting. Recall in the two-beam system, we have seen that by strategically adding an integer multiple of 2π to the phase of one beam, we can reduce the error in line fitting without changing the actual value of the phase. This is due to the concept of phase wrapping, where adding 2π to a phase value results in the same point on the complex plane.

To generalize this insight for an arbitrary number of beams, we need to ensure that the phase difference between any two consecutive beams is less than π. Because if the phase difference goes more than π, then we can always add 2π to the phase of one of the two consecutive beams to make the phase difference less than π. By doing this, we can ensure that the step size doesn't grow high for a large antenna index, thus minimizing the error in line fitting. To obtain the expression for the integer constant k_d for each beam d, we start by fixing the phase of the first beam and then adjust the phase of the second beam with respect to the first beam and follow this process to all consecutive antennas: adjust the phase of dth beam with respect to the $d-1$th beam. This ensures that no two consecutive beams have a phase difference greater than π. This results in the final expression for k_d in (3.29), which expresses the integer constant multiple of 2π that is added to the phase of each beam to achieve the desired bound on the phase difference. This bound on phase difference leads to a bound on the error in line fitting, resulting in an accurate line fitting for generalized multi-beam systems.

3.2.3 Supporting Arbitrary Size Bandwidth Part with DPA

One of the main requirements of FlexLink is the ability to support an arbitrary bandwidth part. However, in practice, there is an error due to line fitting as discussed in the previous sub-section, that limits the performance of DPA multi-beamforming. To understand it better, we plot the beamforming gain of DPA as a function of the bandwidth part in two beam cases as

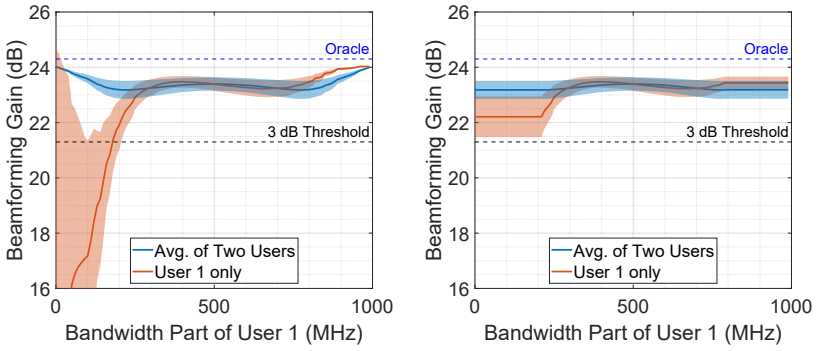

(a) Without optimization (b) With FlexLink optimization

Figure 3.7. Effect of delay and phase quantization in overall array gain of 16 antennae DPA.

shown in Figure 3.7(a). We observe a consistently high performance of overall beamforming gain of two users. However, we see a drop in performance when the bandwidth part of a user is small (typically lower than 20% of system bandwidth). In such a case, the user suffers from low gain despite having a higher gain for the other beam. The reason is that at smaller bandwidth parts, the step function is highly skewed and the error in line fitting is high for the user with smaller bandwidth parts, which results in more than a 10 dB drop in beamforming gain.

We develop a threshold mechanism to avoid this drop. Our idea is to not allow the beamforming gain of any beam to drop by a certain 3 dB below the optimal response. We can achieve this by setting a minimum threshold on beam bandwidth at the 3 dB mark, which corresponds to a 20% beam bandwidth as shown in Figure 3.7(a). If a beam requests a bandwidth part lower than 20% (and other beam request more than 20%), then to avoid lower gain, we compute the delay and phase values corresponding to the 20%-80% split case. This will ensure the beamforming gain of the lower bandwidth part beam freezes to a value only 3 dB lower than Oracle and does not go further down as shown in Figure 3.7(b). Thus, this thresholding scheme is effective in maintaining high gain even for low bandwidth part beams. This, however, comes at the expense of a drop in beaming gain of the other beam with high bandwidth part. It is not a problem because this drop is less than 1 dB which is acceptable given that it brings a lot of benefits to the other low bandwidth part beam. Using this heuristic approach, we can support an

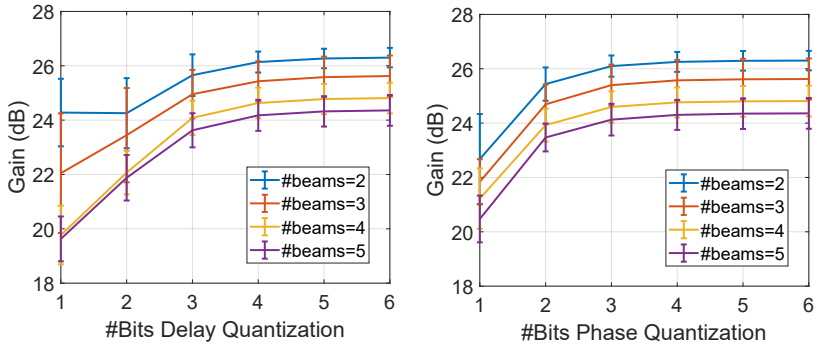

(a) Delay Quantization (b) Phase Quantization

Figure 3.8. Effect of delay and phase quantization in overall array gain of 16 antennae DPA.

arbitrarily small bandwidth part beams.

3.2.4 Design and Implementation of DPA Hardware

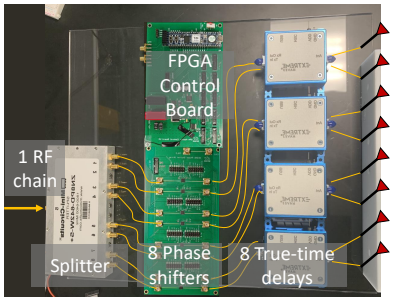

Figure 3.9. Implementing delay-phased array at 4-6 GHz band. There are 8 phase shifters on a custom PCB and 8 TTDs (4 on the bottom and 4 on top) connected to 8 antennas. The delay-phased array is driven by a single RF chain through USRP.

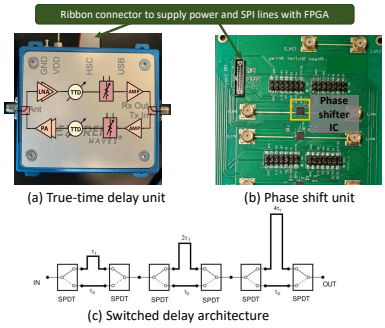

Figure 3.10. Hardware components of TTD unit and phase shift unit.

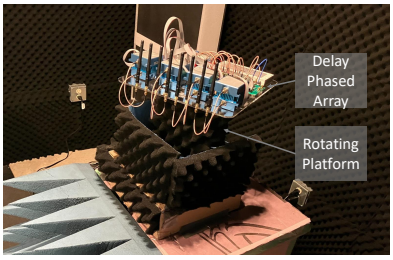

Figure 3.11. Beam pattern measurement in an anechoic chamber.

To demonstrate the practical performance of FlexLink, we designed a custom prototype of DPA hardware in the lab using commercial true-time delay chips and phase shifter chips operating at 4 GHz–6 GHz frequencies as shown in Figure 3.9. Although, generally, there are differences between sub-6 GHz and mmWave systems, the generation of desired beam patterns is agnostic to center frequency and can be validated with our sub-6 setup. Here, we discuss the

DPA hardware requirements in terms of range, resolution, wideband design of delay and phase circuits, and implementation setup.

■ **Range of delay and phase elements in DPA:** The range of phase values is a constant 2π that covers all scenarios because the exponential phase always wraps around 2π. However, the range of delay values is not straightforward because unlike phase, delays don't generally wrap around a certain value. In fact, prior works on using delay element for multi-beamforming showed an unbounded delay that increases linearly with a number of antennas [133], thus hard to build and scale in a practical circuit design because of large size, bandwidth, and matching constraints [86]. In contrast, we designed FlexLink to have shorter and bounded delay lines that do not scale with the number of antennas. For the two-beam case, the delay range for FlexLink is $\frac{3}{2B}$, which is $1.5ns$ for 1 GHz bandwidth. the range of delay can be further reduced with a smaller field of view of the array. A 120^o angular field-of-view requires only $1ns$ delay range.

■ **Resolution of delay and phase elements in DPA:** The resolution is measured in terms of quantization level in the values of delays and phases. The more quantization bits used, the better the array's performance will be, but this also increases the complexity of the hardware design. This creates a trade-off between the number of quantization bits and the gain performance of the array. We show in Figure 3.8 that the array gain improves as the number of quantization bits is increased from 1 to 6 for both delay and phase quantization. However, the improvement in the array gain becomes small for higher numbers of bits and reaches a saturation point. Specifically, the array gain improves by less than 1 dB when the quantization is increased from 4 bits. Therefore, practical FlexLink hardware can be designed with a 4-bit quantization in phase/delay while maintaining high performance.

■ **Wideband Delay and Phase Units:** To meet the ultra-high bandwidth of 5G and the next generation of cellular networks, we need to design a wideband circuit for delay and phase units. Delay lines can be implemented in an integrated circuit using techniques such as passive elements [132], all-pass filter [150], and Switched Capacitor design [136, 86] depending on applications and requirements. We implemented delays using a switched delay architecture [209]

as shown in Figure 3.10 because of its advantage over wideband operation. It deploys a series of SPDT switches connected to two transmission lines of variable length, switching to the longer transmission line would increase delay by a certain unit along the path. With m switches, we can generate m bit quantized delay with 2^m possible delay values. The phase shift can also be implemented using transmission lines, but transmission lines-based phase shift supports only narrowband frequencies. We choose a varactor-based phase shift design for wideband operation.

■ **Design of FPGA Control Board PCB:** We designed a PCB board that hosts a CMOD A7 FPGA [232] to program the delay and phase values in our hardware. The PCB also contains a power distribution module with appropriate LDOs that takes a single 12V input and distributes different voltage levels (e.g., 5V, -5V, and 3.3V) to the phase shifter and TTD chips via a ribbon cable.

■ **Beam Pattern Measurement Setup:** The whole setup was laid over a motor that was controlled via an automated software program as shown in Figure 3.11. To measure the DPA beam pattern, we rotate the motor 1° at a time in the $\pm 60^\circ$ range. We connected the Tx antenna to the VNA's port 1 and the output of the combiner to its port 2. For each angle, we saved the S21 parameters over the operating frequency range. Then, we compute the received power using the measured S21 parameters and plot the power versus frequency and angle. We verified the setup in an anechoic chamber, then we conducted over-the-air experiments in a lab environment. To reduce the channel effects, which may be misleading and affect the resulting image, we used radiation-absorbent materials to isolate the setup.

3.3 Evaluation

3.3.1 Hardware Results

We obtained over-the-air signal measurements with two beams at various angles and bandwidth parts. Figure 3.12 shows example beam patterns across frequency and angle with varying bandwidth parts. The two beams are pointing towards 30^o and -30^o as can be visualized through the shown beam patterns. We reduce the bandwidth part of one of the beams from 50% to 20% (consequently increasing it for the other beam from 50% to 80%). The gain across the frequency band for the smaller bandwidth part is capturing a smaller bandwidth, and vice versa. The gain is high in the desired frequency-angle bin shown by a black box and low elsewhere as expected. We compared the measured beam patterns with simulated ones to show the similarities. We ensured our simulation plots consider the same quantization levels and beam squint effect that we face in our hardware. We show in Figure 3.13(a) that the beamforming gain at the desired angle from hardware matches with the simulation. However, the hardware patterns suffer from additional degradation due to various factors such as multipath (we did over-the-air experiments in a conference room), interference with other devices such as WiFi at 5 GHz, and hardware artifacts (e.g., the matching effect of wideband transmission lines, PCB substrate, and delay/phase IC). Nonetheless, the general behavior suggests that the desired frequency-angle patterns are achievable with our hardware.

Spectral Efficiency Improvement: FlexLink improves spectral efficiency by providing high-gain beams for both control and data signals. The high gain will result in high spectral efficiency. Traditional phased array creates only a single directional beam to support the control signal, but a single beam cannot efficiently support data in the same TTI, even on different sets of bandwidth parts. This is shown in Figure 4.15(a) that the single beam baseline only supports high spectral efficiency for the control signal but low efficiency for the data signal. In FlexLink, we support both control and data with high gain. This results in a significant improvement (almost double in the worst 10% cases) of spectral efficiency for the data signal. The control signal however suffers

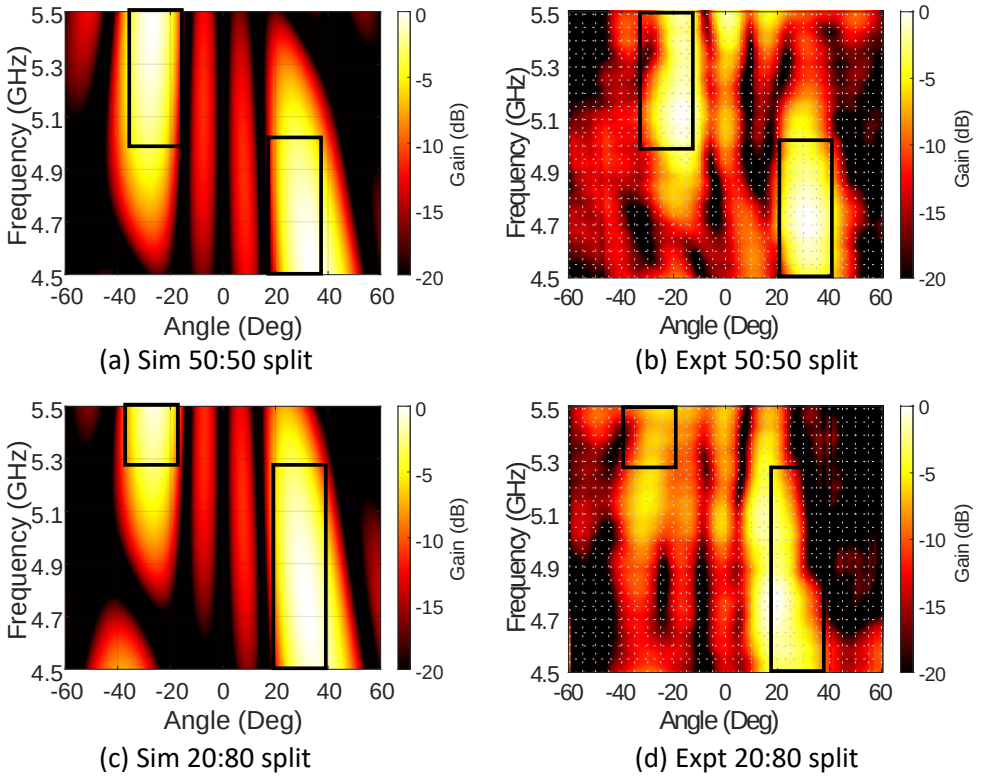

Figure 3.12. Experimental results of beamforming gain for two beams at 30^o and -30^o with different bandwidth parts per beam.

from a loss in efficiency. This is due to the linear approximation we make to obtain delay and phase values and the hardware impurities as explained before. Nonetheless, the overall efficiency of data and control is similar and it shows that it is possible to communicate both concurrently. We again emphasize that while the baseline suffers from only 7% spectrum utilization for control SSB signal, *FlexLink utilizes the entire 100% of the spectrum with control and data signal, both with spectral efficiency.*

3.3.2 HFSS Evaluations

We use the ANSYS HFSS (High-Frequency Structure Simulator) tool to simulate a FlexLink so that all the non-ideal effects involving substrate waves, material losses, coupling between adjacent elements, and individual antenna element patterns are modeled by providing the appropriate delay and phase values. The element is modeled using a simple inset-fed patch

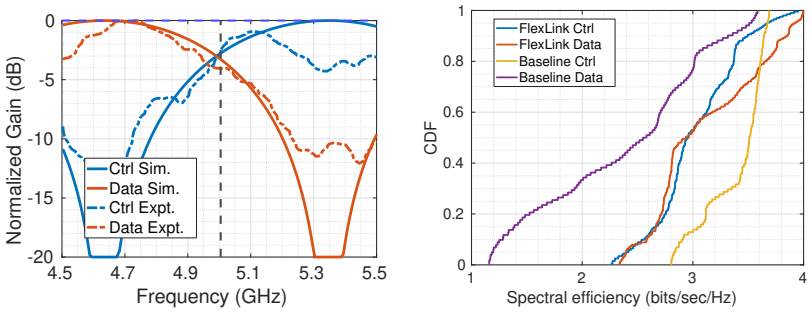

(a) Bandwidth part 50:50 (b) Gain with subcarrier allocation

Figure 3.13. Evaluation Results: FlexLink improves spectral efficiency and increases spectrum utilization.

antenna on a Rogers RO4350 substrate with a variable delay + phase applied at the input port (see Figure (3.14a)). The results of the simulation are shown in Figure (3.14b) with the main lobe along 0° and -30°.

3.3.3 Scaling to Large Arrays

Due to the limited 8 antenna array and 1ns max range of delay unit, we cannot perform hardware measurements for a large-scale system. We performed simulations to show the impact of the number of antennas (up to 64 antenans) and bandwidth parts.

Scaling to 3 or more beams: Here we provide an extensive evaluation for FlexLink under different scenarios. We compare the beamforming patterns that are produced by our O(1) mathematical expression vs high complexity heuristic-based FSDA algorithm [109]. We utilize a 32-element linear antenna array with programmable phase and delay elements per antenna in various scenarios with three or more beams (three-beam configuration shown in Figure 3.15(a)). FlexLink can reasonably achieve the desired multi-beam patterns with an arbitrary number of beams, beam directions, and bandwidth parts. We observe that the antenna gain is high in the desired frequency-direction pairs and low elsewhere, as expected. This verifies that the linear approximation of the non-linear per-antenna phase profile is fairly accurate. However, at the edges of each beam, there may be a gradual transition instead of a sharp one. This can be seen at the boundary of the frequency band supported by each beam, but it only occurs along the

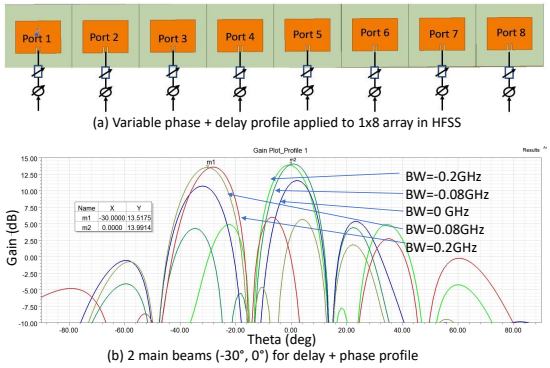

Figure 3.14. HFSS hardware simulation of DPA.

frequency band and not in the angular domain. The beams remain sharp in the intended directions and do not spread to other unwanted angular directions.

Low computation complexity: We show that FlexLink achieves the same high performance with low computational complexity compared to heuristic algorithm [109]. Figure 3.15(b) shows the corresponding delay and phase values obtained by FlexLink and baseline [109] optimization are the same values. These delay and phase values are used to create the beamforming pattern in Figure 3.15(a). FlexLink achieves this high performance with a simple O(1) complexity in estimating the delay and phase values, while the baseline requires high $O(NM\log(NM))$ complexity. For instance, $N = 64$ antennas and $M = 256$ directions, the baseline requires 69k multiplications, while FlexLink requires less than 10 multiplications, which makes it efficient to program in FPGAs.

Impact of bandwidth part allocation: Applications like concurrent communication and control, or use cases with users who require both low and high bandwidth, should have high gains for all beams to user directions to enable low latency and reliable links. Here we demonstrate that FlexLink provides high gains for all users regardless of their bandwidth allocation. To illustrate this, we considered a linear array with 16 antennas and $\frac{\lambda}{2}$ separation and two beams in different directions. Since there are only two beams, if we are giving $x\%$ of sub-carriers to beam 1 (for UE 1), then beam 2 (for UE 2) will get the remaining $(100-x)\%$ sub-carriers. We evaluated gain by varying beam 1 sub-carriers from 5% to 95%. The split antenna approach has low constant gain

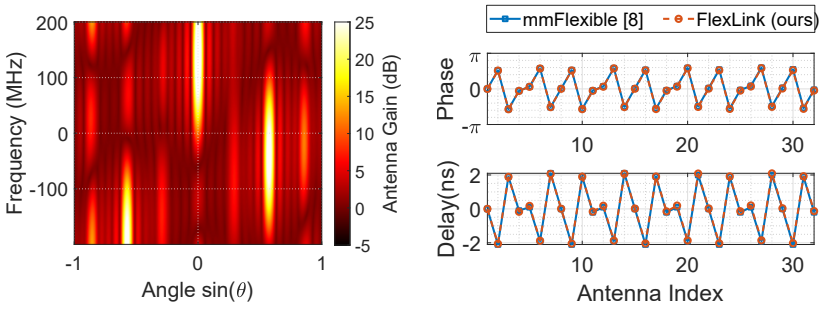

(a) 3-beam response (b) Phases and Delays

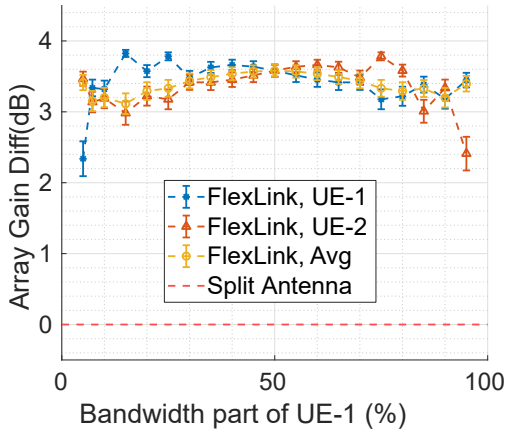

(c) Gain with subcarrier allocation

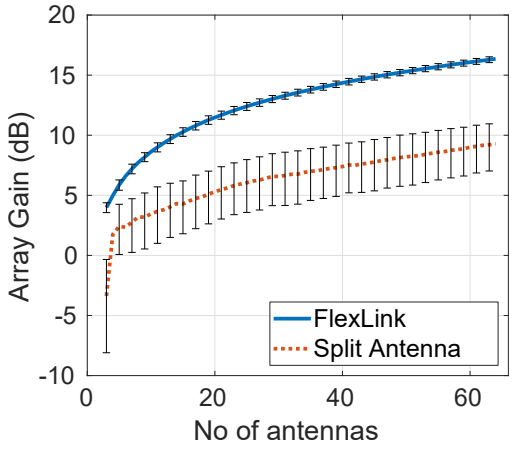

(d) Gain with number of antennas

Figure 3.15. Simulation for large antenna arrays and 3+ number of simultaneous beams.

irrespective of sub-carrier allocations, since radiating beams in user directions over the entire bandwidth. Array gain difference with respect to split antenna is shown in Figure (3.15(c)), demonstrating that even at edge scenarios with beam 1 at 5% sub-carriers or beam 1 at 95%,

Both users get similar high gains without prioritizing high bandwidth user. The error bars in the respective colors indicate gain variations due to different angle separations between the beams. Thus proving that FlexLink truly enables programmable bandwidth parts without compromising on individual user gain.

Impact of number of antennas: Here, we present gain evaluations with an increase in the number of antennas for both DPA and split antenna techniques. We evaluate both approaches in supporting 3 users in different directions and varying the number of antennas from 3 to 64. As illustrated in Figure (3.15(d)), with an increase in the number of antennas both DPA and split technique gain increases with a similar pattern. Note for the three antennas scenario (to support three user directions), the split antennas approach radiates the same as quasi-Omni, hence it has a lower gain than the rest. The error bars in black indicate the variations in gain accommodating different user directions (Monte-Carlo simulations varying user directions). Split antennas as high error bars indicate high gain variations compared to DPA.

3.4 Related Work

FlexLink is different from other state-of-the-art mmWave architectures. We review various architectures, and distinguish them with FlexLink.

Multi-beams with Phased arrays: Techniques like [112, 39, 40] can generate multiple concurrent beams from a single phased array but reduce EIRP per direction due to beam splitting. A wide Quasi-Omni beam [106] results in even lower SNR and throughput. These methods radiate the entire bandwidth in all directions, unlike FlexLink, which selectively targets applications. While they can't decouple data and control applications, they improve mmWave link safety [145], reliability, single receiver gain [112, 39, 40], joint communication and sensing [242], as well as fast initial access and mobility tracking [98, 30, 258]. In contrast, FlexLink creates frequency-dependent multi-beams that maintain high beamforming gain.

Fixed True-time delay (TTD): True-time delay (TTD) architecture [236, 50, 80, 222], frequency-scanning antennas [135, 134, 193], and Leaky-wave antennas [187, 88] create frequency dependent beam patterns but are different from FlexLink. They disperse wideband signals in all directions like a prism, causing each beam to carry a small fraction of bandwidth ($< 6\%$ [50, 222]), insufficient for user demand and leading to energy waste where no active users are present. In contrast, FlexLink generates flexible multi-beams programmable for specific directions and bandwidths per beam. Additionally, TTD methods require changes at both base station and user ends, conflicting with 5G NR protocols. FlexLink supports any 5G NR-compatible beam scan procedure at the base station without modifying mobile devices, incorporating both delay and phase components. If future cellular standards adopt TTD-based fast beam scan methods, FlexLink can adapt, ensuring compatibility with current standards and future-proofing against new TTD ideas.

Delay-phased array: Our previous work on the delay-phased array (DPA) architecture [109] and joint phase-time array [182] have been proposed for creating frequency-dependent multi-beams. FlexLink uses the same DPA architecture as [109], but is different in signal

processing and hardware implementation. None of these past works [109, 182] derive any closed-form expression for delays and phases nor provide insight into the delay range for practical implementation of the circuit board. [246, 156] focuses on mitigating beam squint, [240] on two-dimensional array, and [223] on codebook with this architecture. In contrast, we are the first to derive the closed-form expression, demonstrate a new application of control and data signal decoupling, and demonstrate the performance over the air through hardware implementation.

SSB beam scan for cell search: Extensive work on initial beam scan cell search is focussed on reducing the time overhead while ensuring a reasonable accuracy [114, 256, 44, 207, 98, 40]. [226]. The time overhead vs accuracy tradeoff is hard to achieve because reducing the number of beams over time to reduce time overhead results in a compromised accuracy of angle estimation. In contrast to these approaches, FlexLink focuses on reducing the frequency overhead of beam scan methods through novel frequency-dependent beamforming.

Hybrid multi-RF architectures: Hybrid architectures use multiple RF chains, each connected to a different phased array [244]. Although these arrays can create multiple streams through multiple RF chains, each stream is restricted to radiate in a single direction with the entire bandwidth, due to the phased array front-end. Thus, FlexLink complements hybrid arrays. Instead of the phased array, if we use DPA in a hybrid architecture, we can achieve both multi-stream spatial multiplexing and a programmable bandwidth part for each stream.

3.5 Conclusion

We presented FlexLink, a system for decoupling control and data signaling through a novel delay-phased array (DPA) antenna architecture. We presented a detailed mathematical analysis of the DPA and derived a closed-form mathematical expression for delay and phase values. We built a hardware prototype for DPA using wideband delay and phase units and showed the feasibility of multi-beam patterns in over-the-air testing and demonstrated its benefits in improving spectrum utilization for control and data signals.

3.6 Acknowledgements

We are thankful to members of WCSNG lab for their insightful feedback. Thanks to YungYi Sun and Sonny Cao for helping with designing the PCB and FPGA programming that was instrumental in demonstrating FlexLink on real hardware. Also, thanks to Mohammed Wael for helping with VNA measurements.

Chapter 3, in part, is a reprint of the material as it appears in **IK Jain**, D Bharadia, "FlexLink: Decoupling Control and Data Communication in mmWave Networks" *Preprint - Under submission to Tier-1 systems conference*. The dissertation author was the primary investigator and author of this paper.

Chapter 4

CommRad: Collaborative Bi-Directional Sensing-Driven Millimeter-Wave Networks

Millimeter-wave access is vital for the next-generation wireless network to connect the last mile devices with extremely high multi-Gbps data rates with ultra-low latency [10, 16]. With such high data rates, future autonomous vehicles, loaded with sensors such as Lidar, radars, or cameras, can share their data with the infrastructure or other vehicles for safer driving; the wireless AR/VR devices would enjoy high-quality XR streaming for an immersive experience [28]. However, a major challenge in supporting these use cases is user mobility – anything mobile is an issue for mmWave links. As millimeter-wave links deploy hundreds of antennas to create narrow beams to increase coverage range, they inadvertently fail to maintain these directional links in mobile scenarios. When the target users are mobile, the link requires tracking the users to maintain directional connectivity. Or when other objects in the environment are mobile, they can obstruct or block communication when they enter across an established link. While there are procedures for beam maintenance in 5G NR, such as beam training, where radio scans multiple beams to find the best direct or reflected-path beam for the users, these techniques suffer from high overhead manifested in their long duration and frequent repetitions. The more time a radio spends in beam training, the less time is available for communication, resulting in a high overhead of up to 25% (5 ms duration of beam training repeated every 20 ms, according

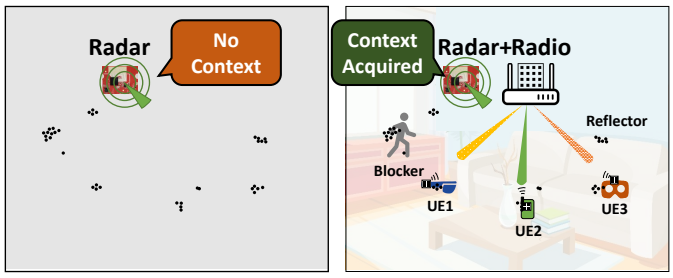

Figure 4.1. CommRad is a collaborative learning framework for context- aware sensing- driven mmWave communication.

to 5G NR protocol [27]). To make it worse, the overhead increases with the number of users, as each user demands separate time and frequency resources for feedback in the form of CSI reports, thus lacking scalability in multi-user scenarios.

To overcome the overhead, sensing-driven communication is envisioned as a pivotal technology where a dedicated sensor is deployed to aid in beam maintenance. The sensors can track users, identify good reflectors, proactively avoid environmental blockers, and instruct the radio about the best direct or reflected-path beam at a given time. Among many choices of sensors, a millimeter-wave radar sensor stands out for low-cost, privacy-preserving, and weather-independent sensing modality in both indoor and outdoor scenarios [249, 62, 42, 250, 69]. However, radar-sensor-driven communication suffers from a fundamental cold start problem as depicted in Figure 4.1. The problem is that the radar yields extensive reflections from many objects in the environment without having any context for those objects – which objects are mmWave users that need to be tracked, or which objects are reflectors/blockers that can assist or obstruct communication links? Due to this contextual deficiency, prior implementation of radar-aided communication is limited to a straightforward clutter-free environment with only one user and does not support reflector-assisted links [66, 139, 37, 58, 60]. These works rely on data-driven machine learning techniques to perform end-to-end training from the radar data to predict communication beam patterns directly, skipping the important step of acquiring context. The hope is that the model will somehow learn the context by feeding extensive data

while training. To the best of our knowledge, no prior work on radar-driven communication experimentally demonstrated context-aware multi-user tracking while supporting a reflected-path link dynamically and blockage adaptation.

CommRad: This chapter presents the design and implementation of CommRad, a sensing-driven communication framework driven by a fundamental insight: the communication radio can serve as a complementary sensing modality, providing the contextual awareness that radar sensing currently lacks. CommRad represents a ***collaborative bi-directional learning framework*** that seamlessly integrates wireless sensing (radar) and the communication radio as sensing modalities to collectively comprehend the environment and user mobility within a relevant context. In this framework, we first deploy the radio as a sensor to provide the required context to the radar through the mandatory beam training procedure. We then deploy the context-aware radar sensor to track any movement of objects and aid in maintaining the beams while the radio is in the communication phase, as shown in Figure 4.2. This bi-directional radar-radio collaborative learning reduces the overhead due to frequent beam training and ensures that radar has appropriate context to maintain the beams, ensuring reliable and high throughput mobile mmWave links.

A key observation within CommRad is the recognition that the radio and radar units provide *'complementary information'* essential for comprehending the environment and user mobility. **Radar as a Sensor:** The radar unit measures reflections from the environment and users through *mono-static sensing*, where the radar Rx is co-located and synchronized with the radar Tx. **Radio as a Sensor:** The radio unit conducts *bi-static sensing*, capturing signal transmissions toward users. The radio signal is reflected toward the user transceiver via reflectors instead of returning directly to the radio unit. Leveraging this observation, we develop a collaborative learning framework that combines wireless sensing (radar) and the communication link (radio) to gain a holistic understanding of the environment and user mobility, particularly in the context of beam management. We then develop beam tracking and blockage prediction mechanisms to ensure a reliable link.

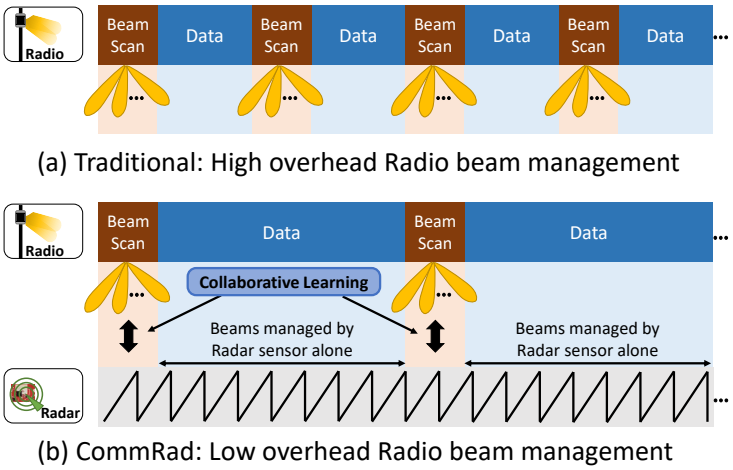

Figure 4.2. CommRad's bi-directional radar-radio collaborative learning framework to improve data communication efficiency by reducing radio beam scan overhead.

Acquiring Contextual Information: A key question is: what is the context that Radar lacks? We define context as the information about the users and environment that radar needs to aid in communication beam management. For active users, this context information includes their angular and distance parameters relative to the base station, while for reflectors, it is their precise location, orientation, and size (Sec 4.2.1). Having this context will ensure that radar can identify the user/reflectors in its view and distinguish them from other irrelevant objects in the environment. The next question is how to acquire this context from radio beam training. The base station radio emits a known preamble across multiple beams, received and utilized by active users to estimate Channel State Information (CSI) across angle and bandwidth. We developed algorithms that leverage this two-dimensional CSI to estimate user and reflector attributes and map them to the radar's view as contextual information as described in Sec 4.2.2.

Context-aware Beam Management: After acquiring context, the radar can distinguish active users and reflectors from other environmental objects. We then developed context-aware beam management by actively tracking the direct and reflected path angle towards all active users for the duration when the radio is busy with data communication. Our tracking solution is inspired by the spatial-temporal consistency of user motion, i.e., a certain maximum speed of motion typically bounds the user location, and further uses Kalman filtering to reduce the impact

of noise and provide accurate tracking (Sec 4.2.3). A key challenge, however, is when the direct path is blocked, the radio needs to find a reflected path beam to maintain a reliable connection. To ensure reliability, we track the reflected path angle via all reflectors identified in the previous radio beam training. We estimate the reflector parameters (location, orientation, and size) and develop a geometry transformation to derive the reflected path angle as a function of the user's direct path angle and known reflector parameters. This way, we track the direct and reflected path angles for all time. The base station uses these angle estimates to maintain a reliable beam towards the user via the direct or reflected paths, even in the presence of intermittent blockages. We developed proactive blockage management schemes that detect blockages in advance and instruct the radio to switch the communication beam to one of the unblocked paths to maintain a reliable connection. This is only possible with CommRad since we actively maintain both direct and reflected path angles using context-aware radar tracking.

Evaluation overview: We evaluate CommRad with an experimental setup that consists of a 28 GHz mmWave radio and 24 GHz COTS ISM band radar with four mobile users. Our setup supports mmWave 5G NR waveform for beam scan procedure. We synchronously collect radio beam scan measurements and radar frames for the duration of the experiment and evaluate CommRad and other baselines on this dataset. The results show that CommRad maintains a reliable and high throughput link while tracking mobile users and proactively handles blockages using radar measurements while periodically collaborating with the radio sensor. The median throughput is improved by 2.5x, while the worse-case 20th percentile throughput is improved by 8x while keeping the communication overhead as low as 0.5%.

4.1 Background and Related Work

Radar-aided mmWave systems: Radar provides high-resolution sensing of various object's location, distance, and velocity in the environment [168, 129]. They are low-cost and can operate in ISM bands such as 24 GHz, 60 GHz, or 77 GHz, which does not require

spectrum auctioning [26]. FMCW radars have been used for tracking mmWave radio beams using theoretical stochastic geometry models [152], and machine learning [66, 139, 37, 58, 60]. These works exploit an end-end machine learning approach where raw range-angle and range-Doppler images are fed into the ML model to get the beam index. Here, only the direct path angle is tracked using the radar, and a single-user system is assumed, which does not address the complexity of multi-user scenarios with reflectors. [37] relies on GPS at the user for beam prediction, but GPS is unreliable and unavailable on all mmWave devices. All these works are still primitive in handling multiple devices, lack beam tracking over time and integration with mmWave beam scan procedure, and do not focus on system integration. CommRad is the first end-to-end radar-driven mmWave communication system that leverages collaborative learning, user tracking, reflector modeling, and blockage mitigation to maintain a high throughput and reliable link.

Other sensors-aided mmWave systems: LiDARs [231, 125, 116] and cameras [34, 54, 104, 31, 53, 157] are demonstrated to track directional beams using signal processing and machine learning models. However, these models demand re-training for diverse environments, often involving labor-intensive data collection efforts with ground truth information, yet may still fall short of obtaining the requisite contextual awareness. In fact, the foundation of contextual awareness laid in CommRad can be applied to these sensing modalities to further improve their performance and generalize them to any indoor or outdoor scenarios. Other sensor-aided communication using device positioning systems [229, 218], light sensors [95], and out-of-band WiFi [205, 159] are less accurate and are mainly developed for indoor WiFi setting which is not common for NextG cellular deployment scenarios.

High-overhead beam training: There is extensive work in beam training methods to reduce the codebook size (scan fewer beams), but since they have to cover an entire angular space to detect a new user or reflector, —which can be anywhere in the angular space— they make a compromise of higher misdetection rates to reduce scanning overhead, especially in multi-user mobility scenarios [114, 256, 44, 207, 98, 40, 100]. To reduce this overhead, [226, 101] relies on

a pre-determined set of beams that are obtained for a site-specific location (e.g., road or highway) for V2X scenarios, but it does not generalize to arbitrary 6G cellular deployments.

Reflector Modeling: The authors of [257, 84, 247, 123, 259, 230, 254, 253, 255] leverage reflective surfaces to improve indoor mmWave connectivity & coverage using only the radio data at multiple static locations. Their model is suitable for one-time modeling during the initial network deployment, which complements our goal of improving the link throughput and reliability in real-time for mobile users using not only radio data but also both radar and radio collaborative learning.

Blockage Prediction: The authors of [67] developed radar-aided blockage prediction using an end-end machine learning model. BeamSpy [208], Terra [78] handle blockage by stationary objects without modeling mobile links and cannot adapt dynamically to blockage arrival patterns. Multi-connectivity [241] and handovers [206] are alternative ways of recovering from permanent blockages. In contrast to these works, we explore dynamic blockage events where a mobile blocker crosses the link or the mobile user crosses a static/mobile blocking object. Unlike [67], our end-end framework not only predicts blockages but also integrates with direct/reflect-path beam tracking with contextual awareness.

4.2 CommRad Design

In this section, we discuss the design and implementation of CommRad, a sensing-driven mmWave communication system to aid in mmWave beam management. Managing directional beams in mobile multi-user scenarios requires maintaining the beam direction towards the user either via a direct path or via a reflected path when the direct path is not available due to blockage events. The key information needed for this beam management at the radio is a time series of the best beam angles for all users at all times. The radio can obtain this information through the periodic beam training phase, but this information gets outdated when the radio is not doing beam training, i.e., it is communicating with active users. During this time, the radio relies

on the radar sensor to manage its beams. The radar sensor is great for detecting objects in the environment with high resolution in the distance and angle estimates thanks to the high bandwidth available (over GHz) in free ISM bands such as 24 GHz, 60 GHz, or 77 GHz center frequencies. Therefore, in CommRad, we explore integrating a radar sensor with the radio to aid beam management. To aid in radio beam management, the radar sensor must provide the time series of the best beam angles (direct or reflected path angles) for each user at all times. However, the radar sensor receives an extensive amount of reflections from many objects in the surroundings, and it cannot distinguish which objects are relevant for beam management, such as users, reflectors, and blockers.

In CommRad, we developed a three-stage process to acquire this crucial context and build a bi-directional collaborative framework for beam management, an overview of which is shown in Figure 4.3. **(Stage 1) What is context?** The first step is to define the context for the radar so that it can aid in beam management. The context should be relevant for identifying users, reflectors, and blockers in the environment and should be understandable by the radar, i.e., mapped into the radar's range-angle view. We discuss this context and how it is hard to obtain by the radar in Section 4.2.1. **(Stage 2) How to acquire context?** The second stage is to acquire this context for the radar. More importantly, the context of the users and reflectors should be updated periodically since the radar may lose this context in a dynamic environment where new users may appear or existing users may disappear. We discuss our algorithms for acquiring this context through the periodic radio beam training process in Section 4.2.2. **(Stage 3) How to do context-aware beam management?** Finally, as the radar is equipped with context, it has to continuously aid mmWave radio in maintaining the beams when the beam training is not immediately available. We discuss how radar aids beam management by tracking the user's direct and reflected path and proactively detecting blockers in Section 4.2.3. Actively tracking the reflected path along with the direct path enables us to implement advanced constructive multi-beam techniques [112] to improve link reliability. The overall system is implemented through a synchronized COTS radar and radio platform and evaluated in real-world settings, as

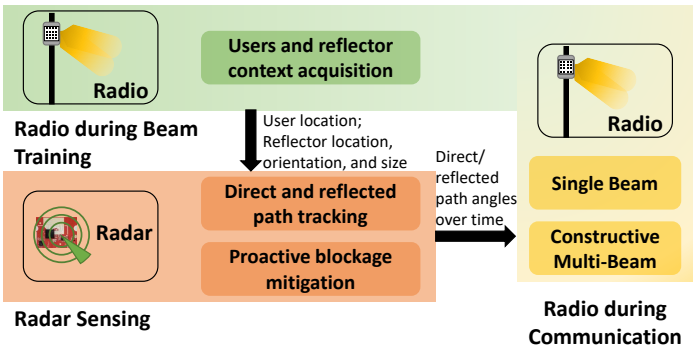

Figure 4.3. Overview of CommRad's end-end system implementation of collaborative bi-directional learning with radar+radio integration.

discussed in upcoming sections.

4.2.1 What is the Context that Radar Lacks?

The radar sensor detects different objects in the environment in the form of their distance and angle estimates. It creates a distance-angle view that shows multiple points of reflection from objects with varying strengths. We define the context for the radar as the set of points that belong to the relevant objects for the communication radio, such as users, reflectors, and blockers. Other objects in the environment that do not fall into this category are not relevant to the scope of this thesis. **(1) Context for users:** For the users, the context lies in the points in the distance-angle view that belong to an active user for which the radio has to maintain a directional beam for communication. Here, we define context as the approximate {distance, angle} of points that belong to an active user. **(2) Context for reflectors:** For the reflectors, the location is crucial too, but not enough. A reflector is modeled as a line segment, which is characterized by the location (any point on the reflector), orientation (slope of the line segment), and size (the start and end points on the reflectors). Learning these reflector parameters is essential to compute the reflected-path beam angle. **(3) Context for blockers:** Finally, we define the context for the blockers by their relative location, size, and speed, which is necessary to estimate the arrival time and duration of blockage events. Without this context information, the radar sensor cannot aid mmWave radio in direct/reflected path beam maintenance.

How Radar Lacks Context for Users?

The radar sensor detects multiple objects in the environment, but it does not know which objects belong to an active mmWave user. Without this crucial context, Radar cannot track user mobility and aid in beam maintenance. To demonstrate this limitation, consider a scenario with two mobile users in the environment crossing each other, i.e., one user is moving left to right, and the other is moving right to left in front of the radar. The radar can detect two users and track their angle over time. Without context information, the radar cannot distinguish between the two users and incorrectly assumes that both users continue along the same trajectory after the crossing, as shown in Figure 4.4(a). In reality, they move in opposite directions. This confusion can be resolved by periodically incorporating contextual information from the radio sensor, resulting in accurate mapping of user trajectories to user ID, as demonstrated in Figure 4.4(b).

How Radar Lacks Context for Reflectors?

The next challenge for radar is determining which objects in the radar's view serve as useful reflectors. Simply gauging signal strength on the radar to locate strong reflectors is not feasible, as objects that appear as a strong reflector for radar could be weak for communication. To illustrate this, we conducted an over-the-air experiment with our testbed involving varied reflector orientations. We plotted the Received Signal Strength (RSS) at both the radar and radio, as depicted in Figure 4.5(a). When the reflector is positioned at 0 degrees orientation, it optimally serves the user with the highest RSS, whereas for other reflector orientations, RSS diminishes as shown in Figure 4.5(b). In contrast, the radar obtains its highest RSS when the reflector is oriented at 60 degrees, a configuration that does not effectively serve the user through the reflected path, adhering to the law of reflection. We show in later sub-sections how radar-radio collaboration helps in obtaining the context for the reflectors and aiding in reflected beam maintenance even under user mobility.

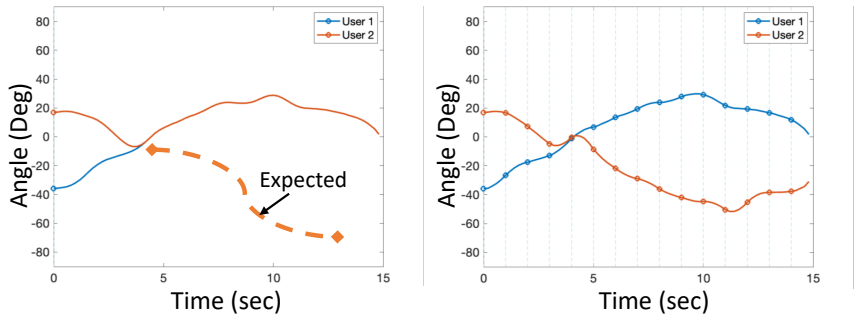

Figure 4.4. Challenges with lack of context for user identification

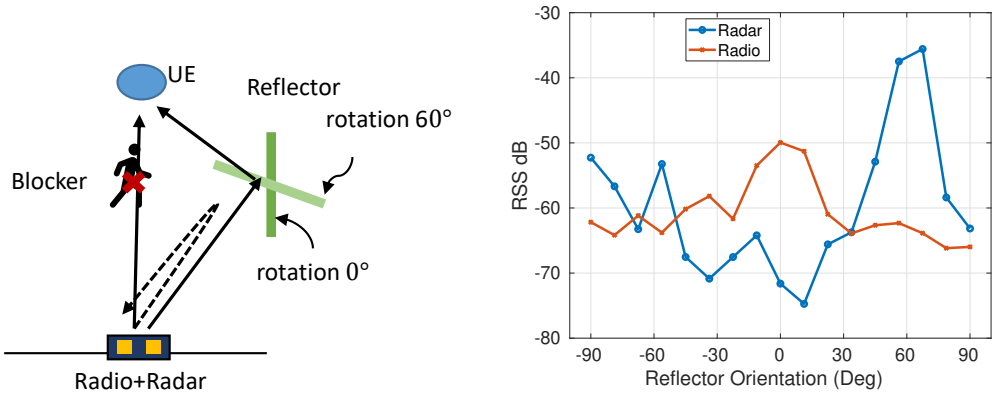

(a) Varying reflector orientation experiment setup

(b) RSS with varying reflector orientation

Figure 4.5. Challenges with lack of context for reflector identification.

4.2.2 Acquiring Context Through Radar+Radio Collaboration

So far, we have seen that the radar sensor alone cannot obtain contextual information about users and reflectors independently. This section describes how contextual information can be obtained using our radar+radio framework. This contextual information is then used to provide context to radar to map different objects as users or reflectors.

Acquiring Context for Users

Acquiring user context is a process of identifying which points in the radar's distance-angle view are active users. We use the radio beam scan procedure to estimate the user's angle and

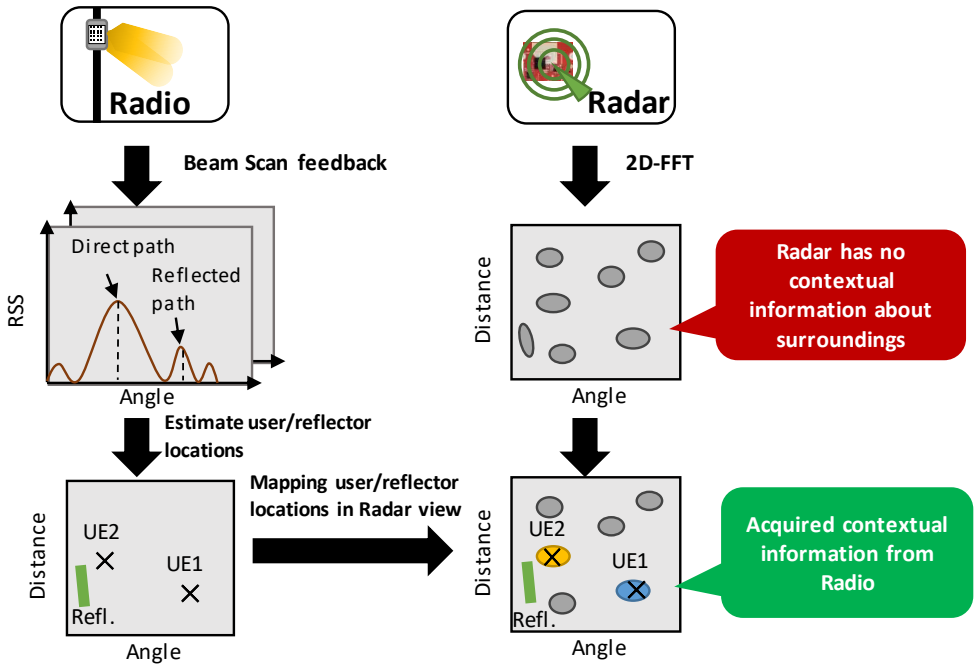

Figure 4.6. Flow Diagram: Identifying users and reflectors in radar view using contextual information from radio.

distance from the base station and then map it to the radar's view. First, consider a simple multi-user scenario without any reflectors and blockages and later analyze their impact. During the beam scan process, the base station transmits the preamble through multiple beams in different directions. Each user captures this signal and estimates the RSS per beam. The beam index that gives the highest RSS is identified as the best beam that serves the individual user. The direction at which this beam points is identified as the user's angle. We further estimate the user distance using the observation that the distance is inversely proportional to the RSS using Friss equation [196]. These estimates are used to map the user's location in radar view as shown in Figure 4.6.

Acquiring Context for Reflectors

Reflector identification is a crucial process in CommRad, as reflectors are key in maintaining reliable communication [112]. With its monostatic sensing capabilities, radar faces challenges in independently identifying reflectors, which contrasts with the bistatic nature of communication

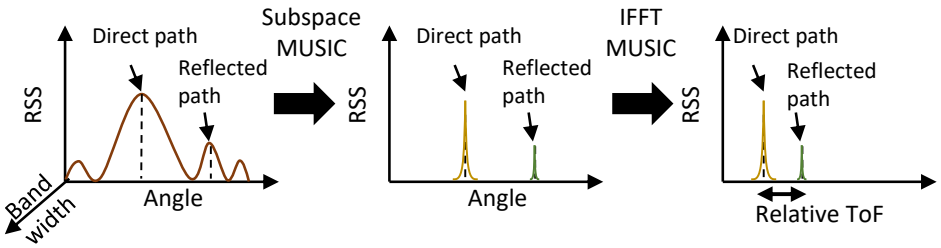

(a) Radio beam scan RSS with scan angle

(b) Subspace MUSIC spectrum

(c) Mapping user and reflector in radar view

Figure 4.7. Reflector location estimation using radio beam scan measurement.

systems. In CommRad, we leverage distinctive attributes from both radar and radio to acquire reflector identification.

Reflectors can be categorized as point reflectors (e.g., poles and lamp posts) or surface reflectors (e.g., buildings, glass walls, and monitor screens). Surface reflectors are particularly relevant for mobile links since their reflections are well-defined and predictable, unlike point reflectors, which are better suited for static users as they struggle to sustain mobile links. Therefore, our focus in this thesis is primarily focused on surface reflectors.

We represent surface reflectors as line segments in Cartesian coordinates, characterized by their location (any point on the reflector), orientation (slope of the line), and size (start and end points on the reflector). We estimate these parameters leveraging two-dimensional channel measurements across angle and bandwidth, following a three-step process as follows:

(i) Estimating reflected path angle and distance: As a first step, we must separate the direct path and reflected paths in the radio RSS measurements. We apply a variation of the MUSIC algorithm, "Subspace MUSIC," to separate the two paths across the angle domain, as we do not have per-antenna measurements in mmWave systems with analog beamforming arrays. We then apply MUSIC across bandwidth to obtain relative Time of Flight (ToF) information between the two paths, following the classic procedure in localization [128]. The direct path is identified as the one with the lower ToF, corresponding to the shorter signal travel distance. By utilizing the relative ToF and actual user location obtained through RSS measurements, we determine the angles and distances of both the direct and reflected paths. The procedure is

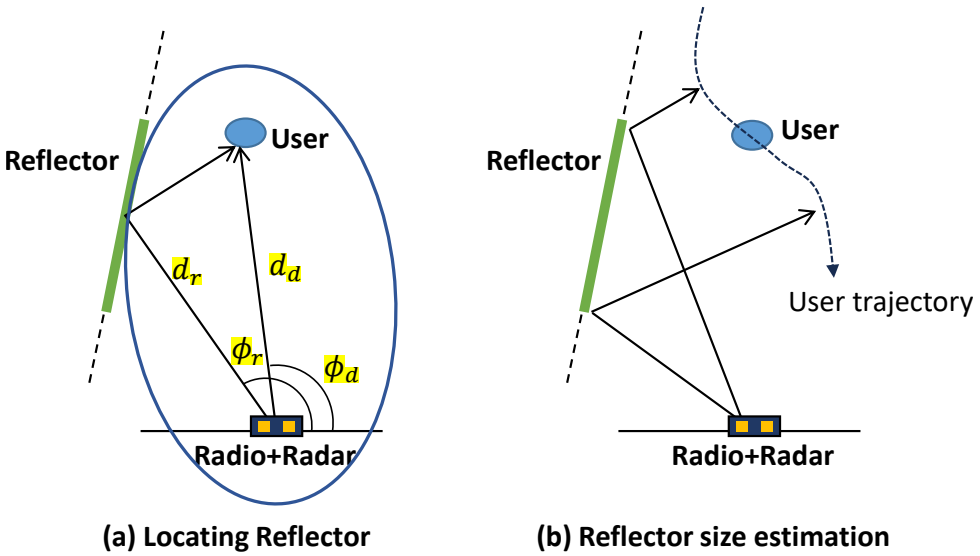

Figure 4.8. Reflector orientation and size estimation.

explained in Figure 4.7.

(ii) Estimating reflector orientation: After obtaining the angles and distances of each path, we construct an ellipse with the base station and user location as its foci. The reflector's position locus lies on this ellipse, as shown in Figure 4.8(a). To pinpoint the exact reflection point on the ellipse, we find the intersection between the ellipse and the line extending from the base station towards the reflector angle. This intersection point serves to identify the orientation of the reflector as the slope of the tangent line to the ellipse at that location.

(iii) Estimating reflector size: To complete the reflector characterization, we need to estimate its size, which is crucial for tracking the reflected path angle. We address this challenge through a learning process over time, as demonstrated in Figure 4.8(b). We collect multiple beam scan measurements from various user locations as they move around and repeat steps (i) and (ii) to identify all the points on the reflectors. In this way, we obtain the reflector size as a set of two endpoints on the line segment obtained in this process. Notably, this is a one-time process for static reflectors but can be repeated if changes occur in the reflector configuration.

Problem of unreliable RSS measurements: It is essential to note that the reflector identification process may accumulate errors if underlying parameters are inaccurate. A key source

of error is distance estimation from RSS measurements, especially in multipath environments. We have observed that radar distance measurements are more accurate, relying on time-of-flight measurements from a synchronized Tx and Rx architecture. Thus, our approach initially obtains course distance estimates from radio measurements and then refines them by identifying the nearest object in the radar view. This dual-step *radar-radio collaborative learning* process ensures accurate distance and angle measurements, enhancing reflector identification accuracy.

4.2.3 Context-Aware Beam Management using Radar Sensor

After identifying users and reflectors in the radar view through radar-radio collaborative learning, we next use the radar sensor to track the user motion through changes in both the direct path and reflected path over time. It is important for radar sensors to maintain and aid direction beams while the radio is occupied with data communication and not performing beam scans. Our approach to the direct path and reflected path tracking is described in this subsection.

Before discussing our tracking algorithm, we present a primer on the choices of range and resolution, followed by the techniques used to counter the challenges described in the previous sections.

Maximum Range and Range-Resolution:

The range resolution for FMCW radar is given by $\Delta R = \frac{c_0}{2B}$. CommRad uses a bandwidth of 200MHz, thus resulting in a range resolution of 0.75m. With a resolution under 1m, we are able to capture the movement of the user across time accurately. The maximum range of the FMCW radar is given by $R_{max} = \frac{c_0}{2B}N$, where N is the number of samples corresponding to the duration observed. For an N value of 1000, we get a maximum range of 7.5km.

Angle resolution: The angle resolution of a antenna array with N_{ant} is given by, $\theta_{res} = \frac{2}{N_{ant}}$. For a 16 virtual antenna system, angular resolution is 7.16^o, whereas for a 8 antenna system, angular resolution is 14.32^o. By using a virtual antenna array, we achieve better angular resolution, which helps improve the localization and tracking accuracy of the system.

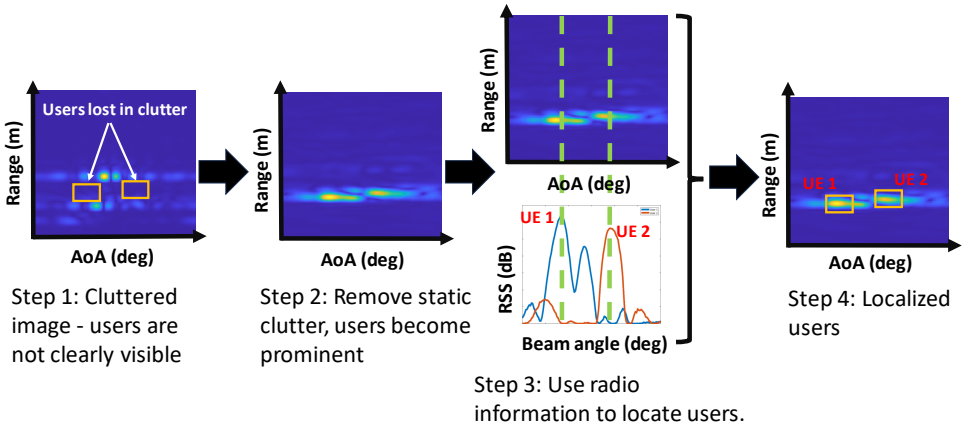

Figure 4.9. Clutter removal and user localization

Static Clutter Removal with Frame Subtraction

The Range-AoA representation from radar contains high-energy points corresponding to all environmental reflections, including users, reflectors, blockers, etc. While tracking the user, all static reflections coming from the environment can be considered as clutter. The presence of clutter causes ambiguity in distinguishing the users over time. To remove clutter, we use the static nature of those objects and cancel them across time. Each frame, corresponding to one chirp, is 1ms apart from the consecutive frame. Depending on the type of mobility (users in indoor scenarios have lesser mobility than in outdoor vehicular scenarios), we perform an average over multiple chirps and subtract the Range-AoA profile of the current frame from the averaged Range-AoA profile. This removes the static objects, and user tracking can be done accurately and reliably. We show the above-mentioned workflow in Figure 4.9. We first obtain the clutter-removed image using frame subtraction and then use the radio data to localize the users within this image.

Direct Path Beam Tracking

For direct path tracking using the radar sensor, we have devised a three-step process illustrated in Figure 4.10. The first step involves identifying the user's initial location within the radar view

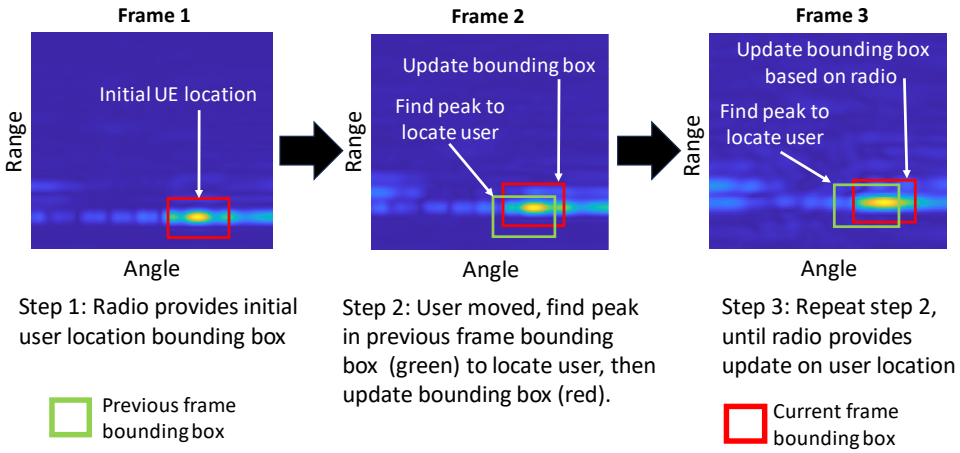

Figure 4.10. Proposed user tracking algorithm.

and creating a bounding box around each user. Creating a local bounding box around a user allows us to remove the dynamic clutter in the environment, like other users, blockers, etc., that may not be present near the user. In the second step, we examine the subsequent radar frame to locate the user with the highest signal strength while retaining the same bounding box as in the previous step. Finally, in step 3, we adjust the bounding box to align with the new user location identified in step 2. These last two steps are repeated for all consecutive frames to ensure continuous user tracking. By tracking the user with radar, we can still maintain directional beams toward the user, even when the radio resources are unavailable for beam training. Periodic updates of user location from the radio can then help maintain the tracking accuracy of the radar, especially when the radar loses the context of the users who might be too close.

Reflected Path Beam Tracking

Reflector tracking is more complicated than user tracking, as mobile users change the angle of the reflected path. We need to know the reflector's location, size, and orientation to track the reflector. Once this information is available, we can create a virtual representation of the base station relative to the reflector as an object behind a mirror. The line connecting this virtual

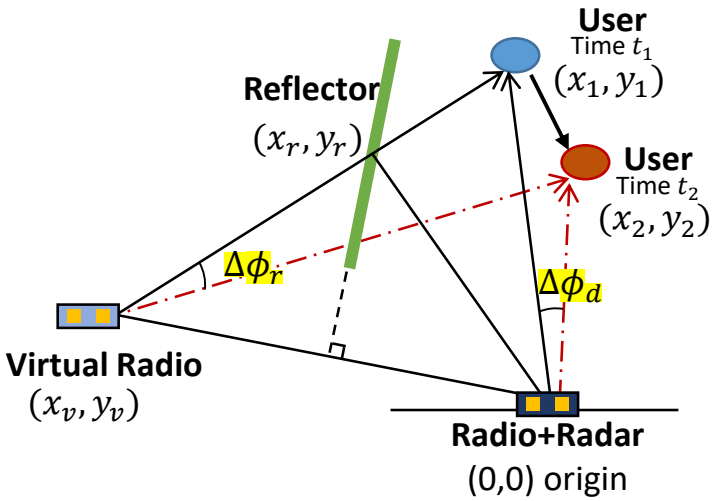

Figure 4.11. Tracking both direct path and reflected path.

base station to the user defines the angle of the reflected path. Let the user's location at time t_1 be denoted as (x_1, y_1), and at a subsequent time t_2, the user's location changes to (x_2, y_2) (See Figure 4.11). Assuming the radar and radio system is located at the origin $(0,0)$ and has a virtual location of (x_v, y_v), we can estimate changes in the reflected path angle, denoted as $\Delta\phi_r$, as follows:

$$\Delta\phi_r = \tan^{-1}\frac{y_2 - y_v}{x_2 - x_v} - \tan^{-1}\frac{y_1 - y_v}{x_1 - x_v} \tag{4.1}$$

Here, the virtual location of the radar/radio system $\times_v, y_v$ is unknown, which is dependent on the reflector's location and orientation, defined as:

$$(x_v, y_v) = (\tan(\phi), -1) \times \frac{2(y_r - \tan(\phi)x_r)}{1 + \tan(\phi)^2} \tag{4.2}$$

Where (x_r, y_r) represents an arbitrary point on the reflector, and ϕ denotes the reflector's orientation angle.

With the estimate of the reflector orientation, we can track the change in reflected path angle $\Delta\phi_r$, which can be used as an alternate beam path to aid communication when the direct path is blocked or unavailable for communication.

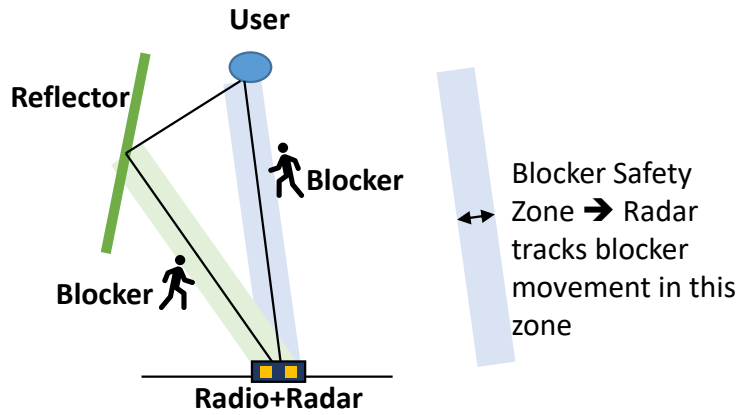

Figure 4.12. Blockage tracking and proactive prevention.

Context-Aware Blockage Mitigation using Radar

Blockage causes trouble to the mmWave link by reducing the signal strength and bringing the link to an outage. Therefore, to revive the link under blockage, alternate reflected paths that are not blocked are utilized. When all the paths are blocked, the link is unrecoverable. To recover the link for partial blockage, it is important to learn the blockage arrival and departure process promptly. Prior work on blockage mitigation is either reactive in nature, which acts when the link is already down, or they rely on radio alone, compromising accuracy in a large multi-user environment. Our proposed collaborative learning framework takes care of the blockage events. By leveraging unique features of radio and radar sensing, we create a signal processing pipeline to promptly detect blockage and take precautionary actions such as switching beams to a reflected path.

Defining a blockage-prone region: The first step in detecting blockage with the help of radar is to determine a blockage-prone region. We define a blockage-prone region as a rectangle patch between the user and the base station, as shown in Figure 4.12. Any blocker in this region can cause blockage. The radar tracks an object's velocity over frames to estimate the time-of-arrival of the blocker into the blockage-prone region. The blockage duration is estimated from blocker velocity and blocker size. A blocker of length l_B and velocity v_B would block the link for the duration of $l_B v_B$.

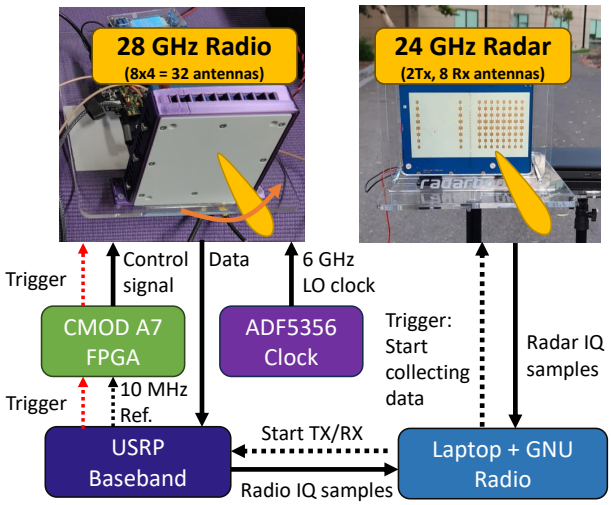

Figure 4.13. Setup: Base station with synchronized Radio and Radar setup

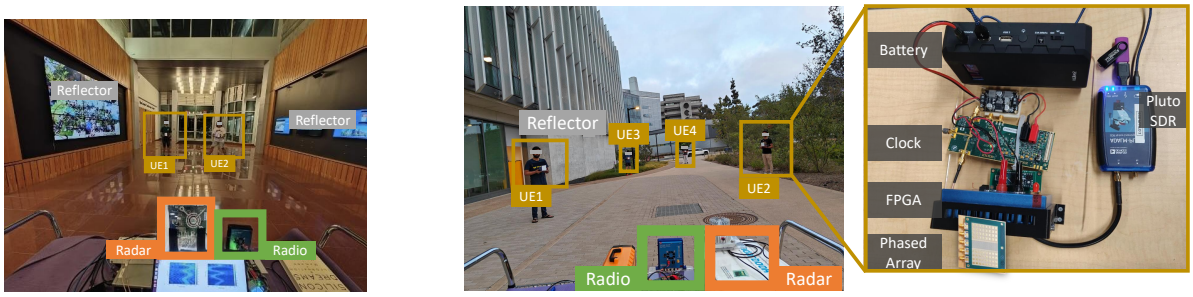

(a) Indoor scenario with reflectors

(b) Outdoor scenario with reflectors, four users, and detailed user setup

Figure 4.14. Experiment scenarios with the static radar-radio platform with four mobile users. We consider diverse scenarios indoors and outdoors with various reflectors such as concrete walls, monitor screens, glass walls, etc.

Proactive blockage mitigation: With the blocker velocity and duration known, the base station alerts the radio of the upcoming blockage event. The response to the blockage depends on the duration of the blockage event. If the blockage event is too slow and lasts longer, then the base station switches to one of the reflected beams. The reflected beam is served until the blockage event lasts. After the blockage is gone, the base station switches back to the original link.

4.3 Implementation

The implementation of the entire system is made in a way that enables experiments that can be conducted both indoors and outdoors. We describe the hardware and software used, the implementation of each subsystem, and the integration of these subsystems to create the joint synchronized platform as shown in Figure 4.13.

4.3.1 Radio Hardware and Software

The mmWave radio system is powered by a combination of commercial software-defined radios (SDRs) and custom-designed phased arrays. The SDRs are connected to the PC using SFP optical cables or USB cables. UHD and GNU-Radio are used to control the SDRs.

■ **Phase Array specifications and control hardware:** The phased arrays have 32 antennas in a 4x8 configuration [23]. Each antenna is dual-polarized, so we have 64 antennas. The antenna RF input and outputs are connected through SMA connectors to the SDRs and operate at 28GHz, which is up/down-converted from/to the 4GHz baseband signal with a 24GHz UDC. We use an Artix-7-based FPGA [21] for programming the phased array registers and activating beam training commands using the SPI protocol. The registers are modified to set the phase and gain values (supports 6-bit phase and 6-bit gain values) obtained from MATLAB corresponding to a particular beam. Look Up Tables (LUTs) present in the phased arrays allow for storing the phase and gain values in a continuous register location that can be loaded rapidly without microcontroller intervention every time. The phased array is capable of performing a beam switch under 6μs. The FPGA, SDRs, and phased arrays are fed with a common 10MHz reference signal generated from the Ettus OctoClock clock distribution module that ensures sample level synchronization within the radio. A power regulator board takes 12V input [15] and distributes power to phased arrays, clock modules [18] and fans.

■ **Base station setup:** For the base station, we use Ettus N310 USRP [22] that is connected to a PC through SFP optical cables. We operate the USRP at 4.05GHz center

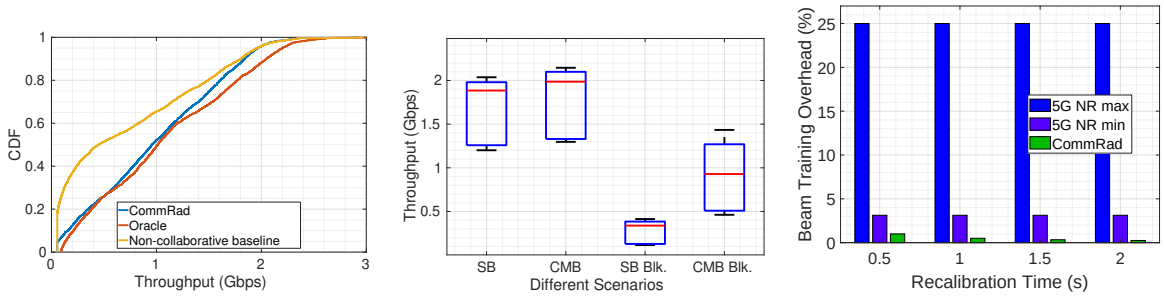

(a) Throughput CDF (b) Throughput in different scenarios (c) Overhead against recalibration time

Figure 4.15. End-end results of CommRad: Improving median throughput by 2.5x compared to non-collaborative baseline. Supporting both single-beam (SB) and Constructive multi-beam (CMB) schemes with blockage mitigation. Reducing beam management overhead in 5GNR.

frequency and a sampling rate of 30.72MHz using GNU-Radio software [24]. We program the pre-calculated phase and gain values for the required angles for an experiment into the LUTs present in the phased arrays. When the USRP starts transmitting, a trigger is sent to the FPGA, which then activates the beam training, and the phased array loads the register values from the LUTs sequentially.

■ **User setup:** We developed a compact mobile setup for the user, consisting of a mmWave phased array, clock, Arduino micro-controller, a battery, and PlutoSDR. We choose plutoSDR [17] at the user because it is compact and allows streaming data directly from a USB stick without needing a laptop, thus making it easier to scale to multi-UE experiments. The phased arrays are controlled through an Arduino microcontroller to set a static quasi-omni-beam at each user.

■ **Waveform design:** We use a modified 5G NR beam management protocol with a 12.5μs OFDM symbol with the preamble and data blocks. We choose this duration of OFDM symbol to get an integer number of samples at the sampling rate of 30.72MHz and allocate enough samples for the preamble. The preamble contains a typical 64-sample pseudorandom sequence, which is repeated twice and has a 32-sample cyclic prefix (CP). The preamble is followed by data that contains the repeated beam IDs that are also coded with a pseudorandom sequence to avoid the high Peak-to-Average-Power-Ratio (PAPR) prevalent in OFDM systems.

4.3.2 Radar Hardware and Software

We used an INRAS RadarBook FMCW radar, which operates at 24GHz [26]. The radar is connected to the same PC as the base station with a 1 Gbit Ethernet cable. The radar has 2 TX and 8 RX antennas that are combined in post-processing to give a 16-antenna virtual array. The radar uses a bandwidth of 200MHz, with one chirp containing 1000 samples, with a slope of 2 MHz/μs having a ramp-up time of 100μs and a total duration of 200 μs. Each data frame contains 200 chirps, and one frame lasts 200 ms. The chirps are split equally between the two TX antennas, i.e., 100 chirps for each antenna. For the above parameters, we have a range resolution of 0.75m, velocity resolution of 0.75 m/s (14 Hz frequency resolution), and angular resolution of 7.16 degrees.

■ **Synchronization and calibration:** radio and radar each save timestamps using the NTP protocol during data collection for synchronization purposes. The radio achieves sample-level synchronization via the USRP, which sends triggers to the FPGA precisely when transmission or reception starts. This, in turn, triggers the phased array to begin beam-sweeping. By chaining triggers from the USRP to the FPGA to the phased array, we align the beam training with the initial sample transmission or reception. Figure 4.13 illustrates the clock and trigger signals used for synchronization and control. Additionally, calibration is necessary because the radar and radio, positioned about 15 cm apart, have different local coordinate systems. We establish a global coordinate system at the radio's location and apply a linear transform to convert radar points into this global coordinate system.

4.4 Evaluation

We evaluate the tracking accuracy of CommRad and the resulting throughput and overhead reduction improvement in various indoor and outdoor scenarios as shown in Figure 4.14. We performed multiple experiments and collected data for different scenarios containing varying numbers of users and different mobility patterns, with each experiment lasting for 10 to 15

seconds.

■ **Oracle:** We consider Oracle as an entity that knows the best beam at all times for all users. The Oracle is obtained by performing a beam-sweep using the phased array at a very fast rate. The phased array scans 121 beams, with each beam lasting 12.5us. Thus, the total time required for a beam scan is 1.5ms. This procedure is executed for the entire duration of the experiment to capture the robust ground truth / Oracle data from the perspective of the radio. The Oracle perfectly captures the reflected beam when the direct path is blocked.

■ **Baseline:** We implement a non-collaborative baseline that does not utilize the periodic contextual information from the radio for aiding radar tracking. To ensure fairness, we provide a one-time context for radar measurement at the beginning of tracking and let the radar track the user without context for the remaining duration of the experiment. Here, the radar follows a similar approach for tracking by creating bounding boxes around the users after the radio gives the initial locations. Then, the radar tracks the user by updating the bounding box, but it does not recalibrate with the radio periodically. It essentially does the tracking on its own. We also compare the results with a reactive baseline, where the radio holds the beam until the beam is updated in the next beam training and has no input from the radar.

We first present end-end results with CommRad and then benchmark individual components of CommRad separately.

4.4.1 Overall Throughput Improvement

The goal of this evaluation is to show the throughput improvement of CommRad in diverse multi-user multipath environments. We obtain radio context information periodically, ensuring the radar can track and maintain beams in this duration. We obtain a CDF of throughput measurements of CommRad and compare it against the non-collaborative baseline and Oracle in Figure 4.15(a). Here, we combined all the measurements across different environments and mobility patterns. The median throughput of CommRad is 1000 Mbps, while the baseline offers only 400 Mbps median throughput, thus providing a 2.5x throughput gain over the baseline

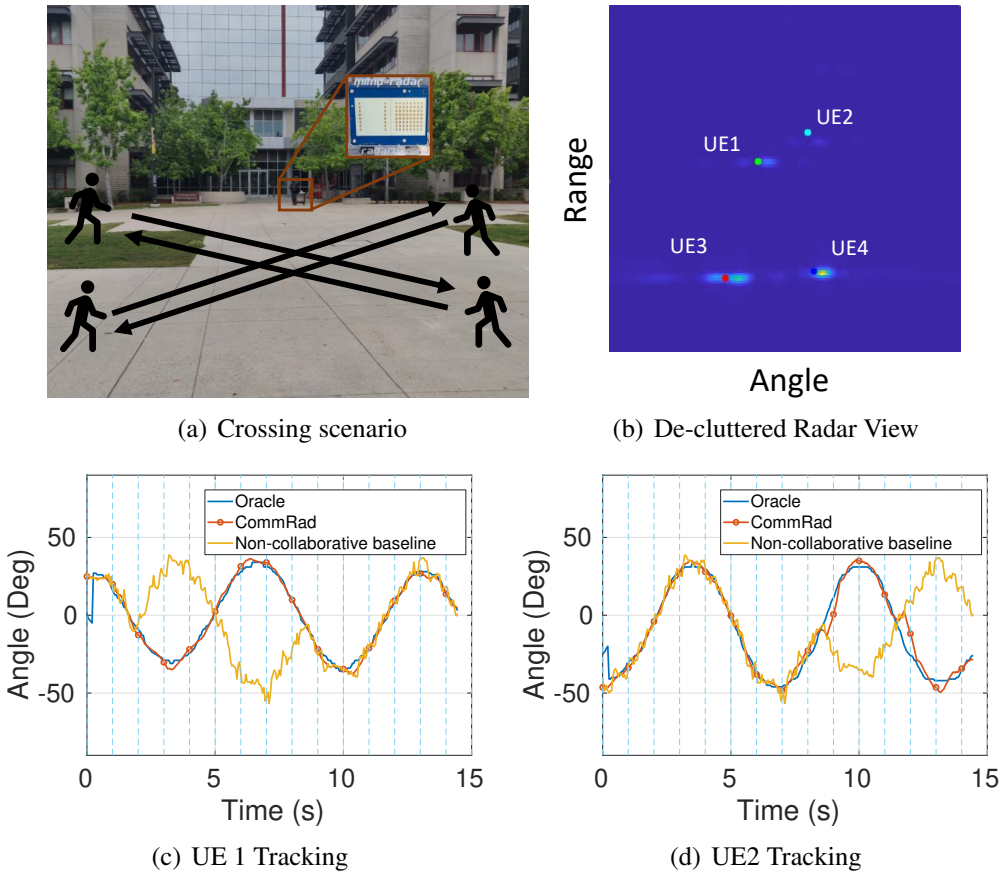

(a) Crossing scenario

(b) De-cluttered Radar View

(c) UE 1 Tracking

(d) UE2 Tracking

Figure 4.16. Scenarios: Four users crossing each other. We remove the clutter and then show angle-tracking accuracy with collaborative learning.

while performing almost similar to the Oracle system. The throughput improvement is more significant for the worst case, the 20th percentile of the cases, improving from 50 Mbps to 400 Mbps, providing a staggering 8x throughput gain. These results show that CommRad can help maintain a reliable link in difficult situations like device mobility and blockages. These improvements are due to robust radar tracking and blockage mitigation enabled by context awareness in CommRad's radar-radio collaborative framework.

Further, we evaluate the throughput improvement using traditional single-beam and constructive multi-beam techniques [112]. Constructive multi-beam uses two beams simultaneously, one in the direct path and the other in the reflected path direction, to provide reliable communication during blockages. The box plots in Figure 4.15(b) show that the mean throughput increases

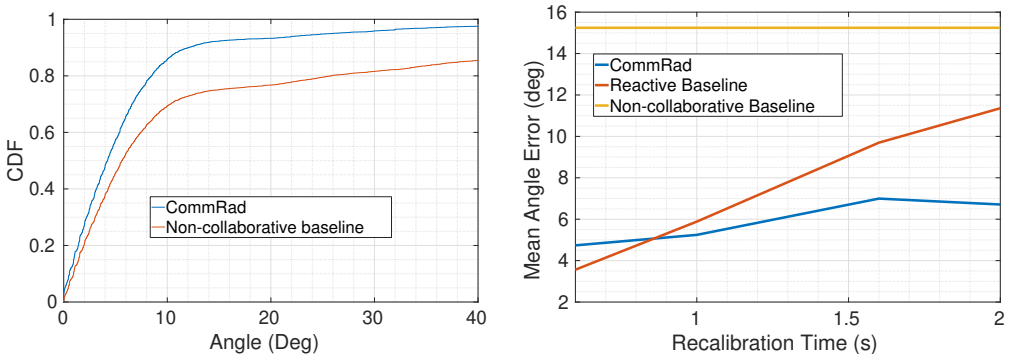

(a) CDF for angle error (b) Angle error over recalibration time

Figure 4.17. Results: CommRad's end-end angle tracking accuracy.

by about 100 Mbps without blockers. With blockage, the traditional single-beam experiences a significant drop in throughput due to the direct LOS path being blocked. In contrast, the constructive multi-beam, which uses the reflected path to support the communication even during blockage to direct LOS path, experiences a drop in throughput by half, which is less significant compared to almost a 10x drop in the single beam case.

4.4.2 Overhead Reduction

Next, we emphasize the advantage of using radar sensing to aid mmWave communication to reduce beam scan overhead. In cellular protocols like 5GNR, the base station (gNB) needs periodic beam scans and constant feedback for establishing and maintaining the beams towards the user, which creates a significant amount of overhead in terms of radio resources, affecting throughput and latency. Based on the periodicity of beam training in 5G NR, this overhead could be between 2.5% to 25%. CommRad can avoid this high overhead in many scenarios encountered in 5GNR. During initial access, the radar can indicate the arrival of a new object in the environment and its angle, thereby allowing gNB to scan only a small subset of angles. This can serve as a substitute for the beam refinement procedure suggested in 5G NR. During mobility, CommRad provides the optimal angle to the gNB, thereby reducing user feedback overhead. Finally, during blockages, while CommRad can provide alternate reflected paths, during the

absence of strong reflectors, radar can provide the user location immediately without needing to re-establish the link. We show that CommRad significantly reduces this overhead down to 0.5% when it relies on radar-based beam management with a periodicity of 0.5 seconds. At a higher periodicity of 2 seconds, this overhead is further reduced by 0.25%. This shows CommRad as a low-overhead mechanism for the radio to spend more time on communication than spending time for beam scanning.

4.4.3 Angle Tracking Accuracy

Accurate user tracking is the most important aspect of the system on which many features like reflector estimation, tracking, and blockage prediction impinge on. One of the main reasons for throughput gain for CommRad is because of its capability to track user angle accurately. We evaluate this by creating a link between the user and the base station radio and performing experiments where the user moves in various directions and speeds. We use radar to track the mobile user along with periodic updates from the radio.

An example is shown in Figure 4.16 with 4 users crossing each other diagonally. Here, we compare CommRad with a non-collaborative baseline comprising only radar tracking without collaboration with the radio. We observe that in complex scenarios where two users overlap, the non-collaborative baseline loses track of the user's identity and reports the same path for both users. With the periodic collaboration with radio, this effect is mitigated in CommRad.

To show the accuracy in user angle tracking, we plot the CDF of the error in angle prediction for CommRad compared to the non-collaborative baseline in Figure 4.17(a). The angle error is the absolute value of the difference in Oracle angle and CommRad/baseline angle. We show that the median angular error of CommRad is 4 degrees while the baseline suffers from 8 degrees of median error. The 90th percentile error is more significant, where the CommRad error of 10 degrees is significantly better than the baseline error of more than 40 degrees. Figure4.17(b) shows the trade-off between angle error and recalibration time. The higher the recalibration time, the higher the angular error, consistent with the results expected

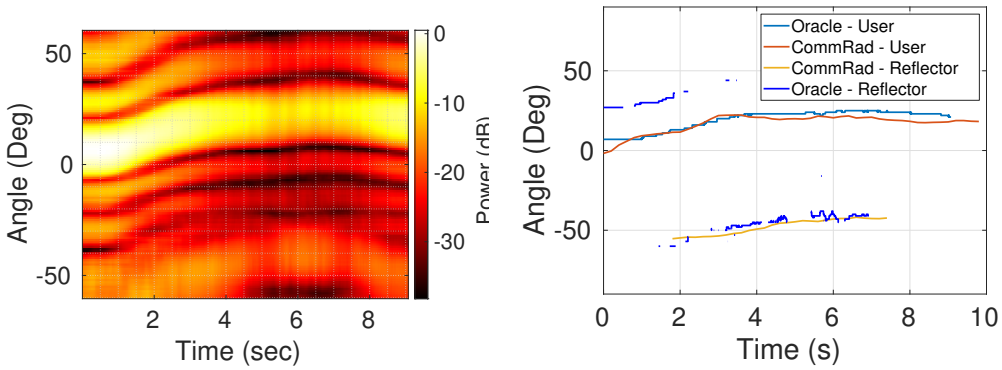

(a) RSS power over time

(b) Angle with time

Figure 4.18. Radar Prediction Accuracy: CommRad acquires ground truth by performing continuous beam training at all times. We show that the angle acquired by this beam scan (Oracle) matches the radar prediction.

from radar tracking objects without context for a long time. We also observe that around 2s, the angle error is around 7^o, within the beamwidth of 10^o of the radio with 8 antennas. Therefore, an ideal trade-off between overhead reduction and reliable communication can be achieved when we update the user location at 2s intervals, causing the overhead reduction up to 0.25% as estimated previously.

4.4.4 User/Reflector Tracking Accuracy

We benchmark the accuracy of our radar tracking algorithm for both direct and reflected paths. For comparison, we obtain the ground truth angle through a continuous beam training process described earlier. Figure 4.18(a) shows the RSS power across time and angle. This represents the ground truth (Oracle) information, where the high energy corresponds to the angle path traveled by the user over time. The second higher energy path corresponds to the reflected path seen by the radio. In Figure 4.18(b), we compare CommRad's angle tracking with the Oracle angle and show that they match with high accuracy. The reflected path tracking is remarkable as the radar could accurately predict when the reflector is available to the user at the 2-sec mark and when it vanishes at the 7-sec mark due to user mobility. This high accuracy is achieved thanks to accurate context acquired by radar in reflector location, orientation, and size.

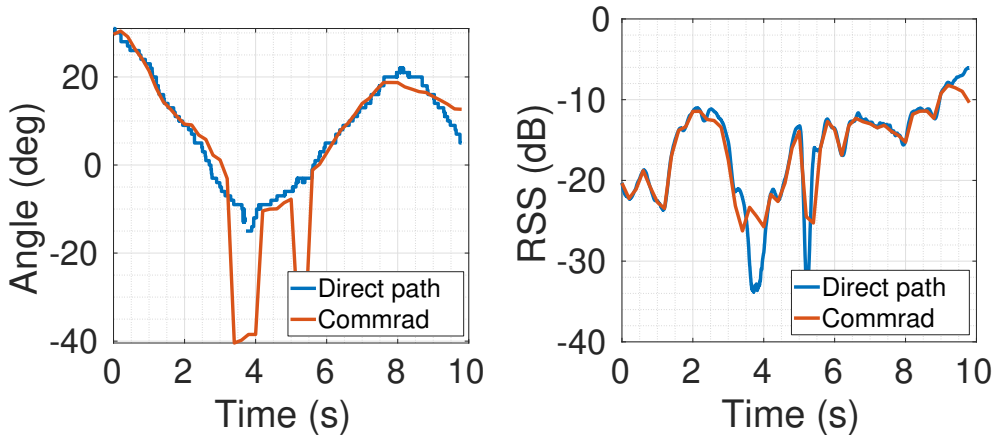

(a) Angle adaptation (b) RSS Gain

Figure 4.19. Blockage adaptation: CommRad adaptation to dynamic blockages.

4.4.5 Blockage Prediction and Mitigation

Blockage for mmwave links reduces SNR and throughput and causes link outages. Utilizing alternate paths to maintain the link while the direct path is blocked is one of the proven solutions to prevent link outages. Since the potential reflected path can be anywhere in the entire angular space, it becomes tedious to do a full beam scan. Instead, we rely on radar to detect the blockage event and identify potential reflected paths to serve the user during the blockage. To demonstrate this, we show a 10-sec trace of adapted angles and corresponding RSS in Figures 4.19(a) and 4.19(b), respectively. We compare CommRad's angle adaptation with a direct path baseline without any adaptation. As we can see, there is a sharp drop in angle twice at the 3-sec and 5-sec mark, which indicates a blockage event and beam switching to a reflected path. This adaptation in CommRad helped maintain high RSS, while the non-adaptive baseline suffered an RSS loss of more than 10 dB due to blockage. This shows the feasibility of radar-aided beam management and blockage adaptation in mmWave systems.

4.5 Discussion and Future work

We discuss potential limitations and future directions of CommRad:

■ **Sensor-fusion with contextual awareness:** CommRad deploys a 24 GHz Radar sensor in the ISM band to aid in beam management with contextual awareness. While Radars stand out as low-cost, privacy-preserving, and weather-resistant sensing modality, they can be fused with other sensors, such as radars at 60 GHz or 77 GHz ISM bands or light-based sensors such as cameras or lidars, to further improve beam management. Notably, CommRad's framework for contextual awareness can be generalized to other sensing modalities.

■ **User-side beamforming:** We provide directional beam tracking at the base station side, assuming a Quasi-Omni directional user-side beam. This represents a typical scenario where the user side has a lower number of antennas due to power and size constraints for hand-held devices. Alternative techniques such as IMU-based user beam adaptation [229] can be used along with CommRad.

4.6 Conclusion

In this chapter, we investigated a collaborative bi-directional learning framework through the collaboration of radar and radio to improve next-generation mobile mmwave links. We identified complementary sensing information from mono-static radar and bi-static radio and combine them under our framework to get the best of both worlds. As a result, we improved mmwave link quality for mobile users even under environmental vulnerabilities such as blockages and beam misalignment due to mobility.

4.7 Acknowledgements

We are thankful to WCSNG lab members at UC San Diego for insightful feedback and to Raghav Subbaraman for their contribution to the initial phase of this project. This work was generously funded by Qualcomm Innovation Fellowship to Ish Kumar Jain and Raghav Subbaraman.

Chapter 4, in part, is a reprint of the material as it appears in **IK Jain**, Suriyaa MM, D Bharadia, "CommRad: Collaborative Learning for Sensing-Driven mmWave Network" *Preprint - Under submission to Tier-1 systems conference*. The dissertation author was the primary investigator and author of this paper.

Chapter 5

Security Challenges in Communication and Sensing Systems

In this chapter, we discuss the security of wireless systems for both communication and sensing. For sensing systems, we explore the security vulnerabilities of mmWave automotive radar that is commonly deployed in autonomous vehicles and robots. We present mmSpoof, a spoofing attack that can mislead a radar sensor into detecting objects that are actually not present in the environment, thus posing security risks. For communication system security, we explored an O-RAN framework to carry out Jamming attacks on cellular networks and built a MIMO-based real-time anti-jamming with srsRAN software stack for end-end applications.

5.1 mmSpoof: Resilient Spoofing of Automotive Millimeter-wave Radars using Reflect Array

Many safety and security-critical applications now rely on radars for real-time situational awareness, especially with the rise of autonomous cyber-physical systems. mmWave frequency modulated continuous wave (FMCW) radars are particularly favored for their robust object detection capabilities in adverse weather conditions, such as fog and low light, where light-based sensors like cameras and Lidars fail [202, 97, 42, 217]. These radars are essential in modern automobiles' advanced driver assistance systems (ADAS) for tasks like pedestrian and blind-spot detection, adaptive cruise control, and emergency braking [204, 238, 120]. Additionally, they

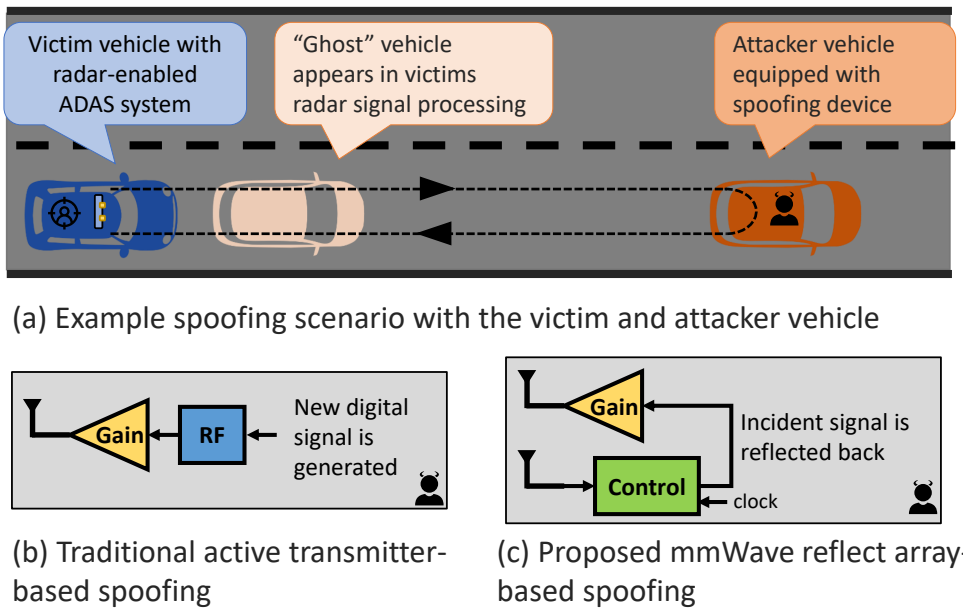

Figure 5.1. mmSpoof: spoofing technique is based on a mmWave reflect array, and it does not require prior knowledge about victim radar, to spoof arbitrary distance and velocity values.

enhance the autonomy of unmanned aerial vehicles [167, 199], support multilane traffic monitoring systems [105], and protect critical infrastructure by detecting unauthorized drones [122, 94]. Given these crucial applications, it becomes imperative to examine mmWave FMCW radars' security against adversarial threats.

FMCW radars transmit a known chirp signal that reflects off objects and is received back by the radar receiver for object detection. Attacks on radar systems fall into two main categories: Jamming and Spoofing [63, 239]. In a jamming attack, the attacker transmits a high-power noise signal in the same frequency band as the radar, which saturates it and hinders normal operation [57, 171, 235]. While this kind of malfunction is easily detected [64, 117, 224], spoofing attacks are more insidious. Spoofing attempts to manipulate radar measurements, either by introducing "ghost" objects or making the radar miss actual objects. For example, an attacker ahead of a radar-equipped vehicle could manipulate radar signals to cause the vehicle to falsely measure a different distance, potentially triggering sudden braking and risking an accident, as illustrated in Figure 5.1(a).

Several studies highlight the vulnerability of FMCW radars to spoofing attacks [121, 55,

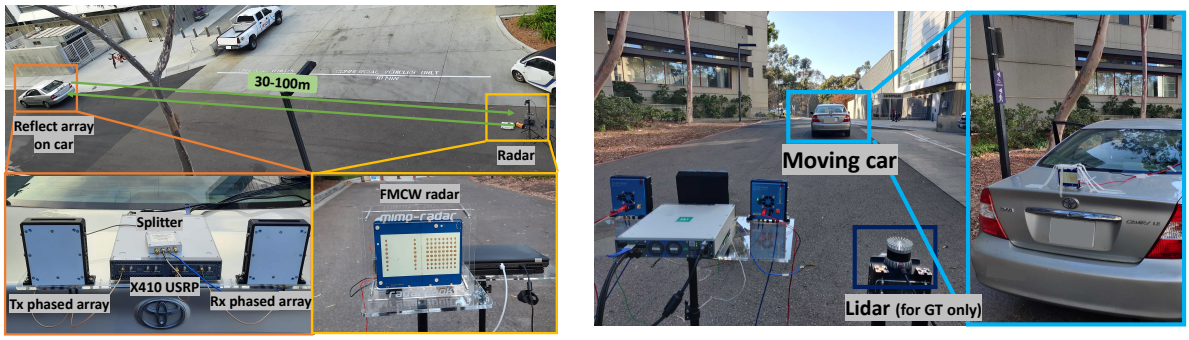

Figure 5.2. The experimental setup consists of the mmReliable device mounted on a car, with the radar on a tripod. The top picture shows an overview of the setup, while the bottom figure shows a zoomed-in view.

56, 149, 155, 127, 163]. In these scenarios, an attacker uses an active transmitter to generate a spoofing signal that the radar receives, leading to a spoofing attack (as shown in Figure 5.1(b)). However, these studies often fall short, requiring attackers to have complete knowledge of radar signal parameters, perfect synchronization with the victim's radar, and connection via wires. These methods also prove vulnerable to radar-specific security measures [151], such as introducing short transmission pauses or dynamically altering transmit signals' characteristics, like frequency sweep slope and chirp time period, to thwart spoofing attempts.

We present mmSpoof, a robust and flexible spoofing mechanism using a millimeter-wave (mmWave) reflect array. The reflect array captures incident radar signals and reflects them back to the radar, causing malicious detection of ghost objects. Unlike active transmitters, mmSpoof leverages reflection with modulation to perturb measured distance and velocity, as shown in Figure 5.1(c). This approach eliminates the need for synchronization between the victim radar and the attacker and preserves radar security encodings, making mmSpoof resilient to various security measures. Moreover, mmSpoof can spoof any arbitrary distance and velocity without requiring synchronization. We also develop a mechanism to learn critical radar parameters necessary for spoofing in the background, enhancing mmSpoof 's robustness and flexibility.

■ **Asynchronous spoofing challenge:** A primary challenge for mmSpoof is to spoof distance and velocity without synchronizing with the victim's radar. The reflect array must

modulate the incident signal to spoof radar measurements appropriately. A naive solution involves applying a delay to the incident signal, increasing the time-of-flight detected by the victim's radar, thereby increasing perceived distance. However, reducing distance requires negative delays, which are non-causal. Our insight leverages the frequency-time relationship in radar chirp waveforms. By creating a frequency shift, which appears as a time shift, we can make the attacker's ghost object appear at arbitrary distances. A positive frequency shift makes the object appear closer, while a negative shift makes it appear farther.

■ **Decoupling distance and velocity challenge:** Another challenge is to spoof distance and velocity independently. Applying a frequency shift to the reflected signal affects both distance and velocity. To address this coupling, we fine-tune the frequency shift to control velocity spoofing without significantly altering distance spoofing. By exploiting the periodicity in velocity variations with frequency shift, we can vary the frequency shift with fine granularity to spoof velocity and with coarse granularity to spoof distance, enabling independent control.

■ **Realistic power variation challenge:** To spoof distances realistically, the attacker must control the received signal strength at the victim's radar, adhering to path loss regulations. Spoofing shorter distances requires higher signal strength than longer distances. We achieve this by controlling the gain of the reflected signal according to known path loss models at the operating frequency. Using a commercial phased array with a built-in gain control mechanism, we implement the required variable-gain spoofing.

■ **Parameter estimation challenge:** A final challenge is to estimate the frequency shift needed for flexible and robust spoofing. Important radar parameters like start frequency, chirp time, and chirp slope must be estimated in real-time, as security features may alter these parameters to thwart spoofing. We developed a digital signal processing mechanism to extract these parameters by analyzing the spectrogram of the received signal and detecting known chirp patterns in standard FMCW radars. Periodic estimation ensures accurate frequency shifts for effective spoofing.

■ **Demonstration and contributions:** We implement a proof-of-concept design of

mmSpoof's reflect array using commercial-of-the-shelf components such as mmWave phased arrays with inbuilt mixers, a clock generator, and software-defined radio (Figure 5.2). We performed various outdoor experiments with a commercial 24 GHz radar mounted on a vehicle. We show that mmSpoof can spoof the radar for up to 100 m, even under mobile scenarios. Experimental results show that mmSpoof indeed creates flexible distance and velocity spoofing under various types of spoofing scenarios. We also present several countermeasures against our spoofing. We summarize our contributions below.

- Design of reflective array-based attack that does not require synchronization with the victim's radar.
- A novel method to independently spoof distance and velocity at victim's radar.
- Procedure to estimate victim's radar parameter and spoof both distance and velocity in real-time, which makes the system robust to several counter-measures.
- A robust attack that can control received signal strength at the radar following the path loss regulations with the spoofed distance.
- Demonstrated our attack's feasibility in static and dynamic scenarios using COTS hardware and spoofing practicality beyond a 100 m range.

5.2 BeamArmor: Seamless Anti-Jamming in 5G Cellular Networks with MIMO Null-steering

Interference and Jamming pose a significant threat to both cellular networks and military communication devices. It involves intentionally or unintentionally disrupting wireless signals by transmitting powerful signals on the same frequency, rendering the original signals ineffective. In cellular networks, jamming can disrupt the normal operation of mobile phones and other wireless devices that rely on the network, which can lead to dropped calls and loss of data connectivity.

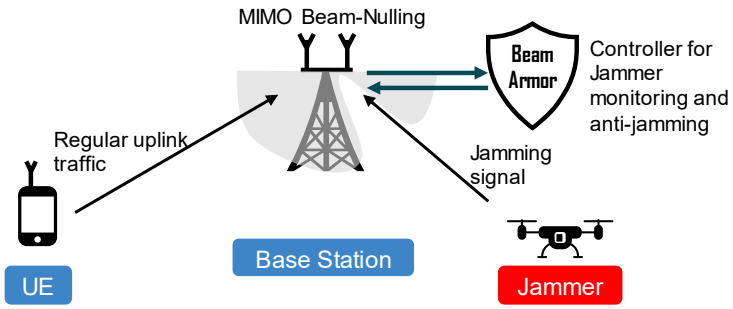

Figure 5.3. BeamArmor: a controller platform for jammer monitoring and MIMO null- steering anti-jamming solutions implemented with open-source srsRAN 5G cellular stack.

As per an estimate, the anti-jamming equipment market was more than 4 Billion in 2023 and is expected to increase by 10% in the next five years [19]. Among many anti-jamming techniques in literature, MIMO-based anti-jamming is attractive because it uses null-steering techniques to selectively suppress or cancel interfering signals while preserving desired signals [237, 68]. By adjusting beamforming weights adaptively, MIMO systems can nullify jamming signals dynamically, reducing their impact on signal quality and maintaining reliable communication.

The threat of jamming is relevant in the current and the next generations of cellular networks because of increased interference at the base station due to high densification for coverage [93, 119], simultaneous utilization of spectrum for cellular and non-terrestrial communication [200] and easy availability of cheap and configurable SDRs like HackRF [25] and bladeRF [20] for malicious jamming attacks. Given the dynamic nature of jamming signals, conventional one-time solutions, such as adjusting antenna orientation during deployment, prove insufficient. Consequently, real-time monitoring of jamming signals and anti-jamming becomes necessary. However, implementing real-time and seamless jammer monitoring and nulling techniques in operational cellular networks is challenging due to limited computation resources and strict deadlines for core PHY layer functionalities. Advanced anti-jamming techniques are computationally heavy and may break the tight time budget of RANs [68, 77]. To enable real-time monitoring, it is necessary to decouple core RAN functions from non-core functions like anti-jamming. The O-RAN initiative aims to disaggregate and virtualize RAN components,

allowing jammer monitoring and nulling to be implemented outside a RAN in a machine called RAN Intelligent Controller (RIC) [172].

We present BeamArmor [260], a practical real-time controller framework for monitoring and nulling jamming attacks in cellular networks. It is built on the open source RAN platform srsRAN as shown in Figure 5.3. We consider uplink (UL) traffic from srsUE to srsenb (base station) and a jammer that tries to jam the receiver at the base station. BeamArmor is a system for monitoring the presence of a jammer, estimating the jammer's channel parameters, and implementing null-beamforming towards the jammer in real time without compromising the communication link with the genuine user. We address the following challenges while building BeamArmor.

The key challenge for BeamArmor is to provide real-time monitoring and nulling in a tight latency budget while maintaining the core functionalities of the RAN. One effective approach is to deploy the BeamArmor on the RAN controller with near-real-time processing (10 ms to 1 sec) [172]. However, many RAN controllers such as Janus [77] are proprietary and not supported by srsRAN. Although FlexRIC [192] and EdgeRIC [126] are srsRAN compatible, they have limitations in supporting per-antenna channel exposure necessary for MIMO applications. To address this, we built a specialized RAN controller for MIMO applications. We create specific interfaces (called hooks) using the ZeroMQ [1] message passing between the RAN and the BeamArmor controller to frequently exchange bi-directional information relevant to hosted applications. The choice of the ZeroMQ publish/subscriber messaging pattern emerged as the optimal design decision, ensuring minimal impact on the timely execution of PHY layer processes.

The next challenge is to modify srsRAN and extract key information from the RAN that would be useful for MIMO applications Extracting too little information may reflect low accuracy, and extracting too much information may break the real-time nature. To address this issue, we identified specific locations in the srsRAN codebase that provide data relevant to MIMO applications, such as per-antenna IQ data. Subsequently, we implemented appropriate

downsampling techniques to reduce latency while maintaining high performance. This downsampled data is used at the controller to allow the applications to apply their MIMO-based controlling and monitoring. For the BeamArmor App's anti-jamming, the beam nulling has to be applied to the incoming UL traffic in real-time. Therefore, we provide additional hooks in the RAN stack to receive the beam-nulling control information from our controller.

Finally, we perform over-the-air experiments in a large conference room at 2.4 GHz ISM band with an end-end cellular stack running in real-time. Our setup consists of srsenb, srsue, and a jammer device implemented with USRP X300 SDRs. We characterize the impact of a wideband and narrowband jammer with varying power levels on the genuine communication link in terms of SINR (signal-to-interference and noise ratio), BLER (block error rate), and throughput metrics. We show that even with two antennas in srsRAN, effective beam nulling can mitigate the harmful effect of jamming with 10 dB of nulling, thus proving the practicality of our scheme. We believe our open-source controller has a broad impact in the research community to develop and implement more advanced MIMO beamforming techniques in real-time[1].

We summarize our contributions as follows:

- Building an open-source controller for MIMO applications such as jammer nulling.
- Demonstrating anti-jamming with full-stack srsRAN and over-the-air implementation in real environments.
- Evaluating narrowband and wideband jamming scenarios showcasing 10 dB of jammer nulling.
- Open sourcing artifacts for BeamArmor online.

5.2.1 MIMO null-steering model for jammer mitigation

This article discusses two jamming scenarios: narrowband jamming on a single frequency, typically the UL center frequency, and wideband jamming using random data as Orthogonal

[1]Artifact link https://wcsng.ucsd.edu/beamarmor.

Frequency Division Multiplex (OFDM) symbols. These jammers fall into the category of regular jammers. They do not follow any MAC layer protocol but instead inject the wireless channel with random signals transmitted on the carrier frequency of the UL traffic. The intention of the jamming device is to cause powerful interference, which acts as noise added to the legitimate signals and thus degrades the SINR (signal-to-interference and noise ratio) of the UL traffic, such that the BLER increases and the throughput of the cellular network drops significantly.

We consider two antenna base stations, as supported by srsRAN[2], and show that even two antennas are enough to create sufficiently wide nulls in the wave-front pattern. Increasing the number of antennas would create a larger number of nulls that could also be controlled to be of wider angle.

Our anti-jamming application controls the beamforming pattern in such a way that, ideally, zero signal power is received from the angle at which the jamming source is present to free the UL data from the jammer influence. To set the receiver beamforming pattern, we apply a weight α to the vector of received IQ samples y_1 and y_2 received by the first and second antenna of the radio front end, respectively (equation 5.1). We divide by the term $\sqrt{(1+\alpha^2)}$ to account for the signal power amplification resulting from applying α. Subsequently, y^* is passed to the PHY layer processing steps that follow the reception of raw IQ samples.

$$y^* = (y_1 + \alpha y_2)/\sqrt{(1+\alpha^2)} \tag{5.1}$$

where α resides in the time domain and is applied to all subcarrier frequencies equally, establishing the assumption of a frequency-flat channel. We compute the scalar factor α from the same objects y_1 and y_2, populated by IQ samples from the current or a prior Transmission Time Interval (TTI). α is computed from IQ samples of TTIs that do not have any users transmitting in UL. Additionally, the assumption is made that within those particular TTIs, all received signal power originates from the jammer source. Therefore, we compute α as a closed-form solution,

[2]The new version of srsRAN released in Sep 2023 supports 4 antennas

where α is the factor that results in the combined signal power to equal zero (equation 5.2).

$$|y_1 + \alpha y_2|^2 = 0 \tag{5.2}$$

This is a linear least square problem that can be solved in a computationally efficient way. Also, this method can be extended to a large N antenna array using a $N-1$-sized α vector using standard null-steering techniques [219]. Our method does not explicitly yield the angle at which the jammer is positioned to the base station. However, when applying the computed α to the UL IQ samples (equation 5.1), a null in the waveform pattern will be directed towards the jammer source, given our assumptions.

If the user position and jammer position are sufficiently wide apart in angle, which we assume to be reasonably likely, the user traffic is not affected by the nulling. Our null-steering method applies null-steering to the entire frequency band of the UL traffic by being a scalar factor in the time domain. While this is an imprecise method to mitigate wideband interference, it is nonetheless a computationally light implementation.

5.2.2 Practical, Real-time, and O-RAN-driven RAN Controller for BeamArmor

The BeamArmor applications are driven by data we collect from the cellular RAN through our custom RAN controller, which will be described in this subsection. Through collecting IQ sample streams from the active RX antennas, our platform enables not only null-steering for jammer mitigation but also other MIMO related applications. We built the RAN controller on top of the end-to-end software solution of a cellular stack, srsRAN, aiming at making our platform open-source. The controller uses one common interface realized through the ZeroMQ messaging library to exchange the IQ sample data in one direction and control information in the other. In our setup, the RAN controller is hosted on the same machine as the srsRAN base station application and a transport layer link between the two processes is created over a shared socket.

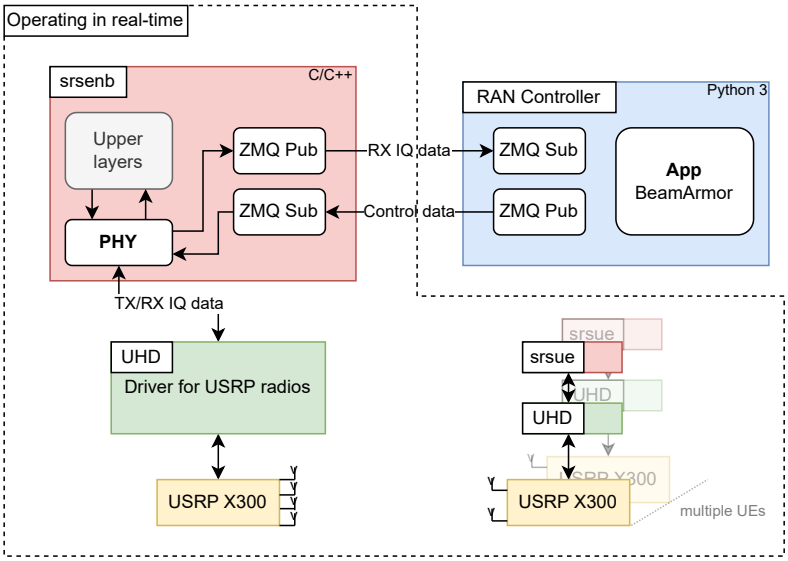

Figure 5.4. Core functional blocks of our RAN controller for BeamArmor in interaction with the srsenb application and the USRP radio peripheral.

To establish O-RAN's principle of desegregation of applications and RAN, the priority lies in providing a communication architecture between the cellular stack and the RAN controller. Figure 5.4 visualizes the operation of this communication with its relevant functional blocks. We chose to use the popular messaging library ZeroMQ [1] (ZMQ) to create two local communication links, although a link across separate machines in a network could be accommodated as well. To comply with the fine-grained clock of the physical (PHY) layer operations in the cellular stack, the outgoing IQ data and incoming control information must not interfere with the timely operation of the PHY layer. The ZMQ publish/subscriber messaging pattern emerged as the best design choice to meet this requirement.

Sufficiently accurate tracking of jamming activity required us to develop near-real-time control capability. In the process of creating the controller for BeamArmor, we have therefore refined the message data structures and their size sent via the ZMQ interface between the controller and RAN stack, so that the RAN can be probed for IQ data sufficiently frequently. In section 5.2.7, we evaluated various down-sampling rates for the IQ data at a RAN-probing periodicity of every 10 milliseconds.

5.2.3 BeamArmor Application for Jammer Monitoring and Nulling

The received IQ data handled by the RAN controller allows us to observe signal power across the UL spectrum and monitor SINR. During dedicated UL idle periods, i.e. users are not transmitting, peaks in received power and drops in UL SINR indicate UL interference that can be linked to jamming activity. Upon detecting jammer activity, we compute the beam-forming weight α from the received IQ data for null-steering. The estimated α is passed as control information to the srsenb using the ZMQ interface to apply beam-nulling with two antenna streams in real-time and successfully reduce the effects of jamming in the UL communication link between the base station and attached users.

In the process of this work, we evaluated applying null-steering to the PHY layer channels PUSCH and PUCCH separately. However, this approach revealed itself to be unreasonably tedious, as srsRAN does not support the PUSCH and PUCCH processing of multiple receiver channels in UL. The current version of srsRAN (Oct 2023) only supports DL 2 antenna MIMO with Alamounti codes, while UL, supports only SISO links. As a result, a separate PHY layer processing pipeline to process a second antenna port's input would be necessary to implement null-steering on the PHY channels.

Rather, we implemented a low-overhead solution in srsRAN to deploy null-steering: The unprocessed IQ data is weighed and added together in the time domain before it is passed to any further UL processing, like transformation to the time-frequency grid. Inside the PHY layer sub-process of srsRAN, the vector of RX IQ samples, denoted by *y1* and *y2* (antenna port 1 and 2 respectively), are processed as follows: Instead of passing the memory address of *y1* to the UL processing, as it is done by default by srsRAN, our modification writes $y^* = (y_1 + \alpha * y_2)/\sqrt{1+\alpha^2}$ to the same address, and all further processing is done to this new data y^*. This ensures the nulling of the jammer signal when both UE and jammer are present together and ensures unobstructed communication with the UE despite the jamming signal.

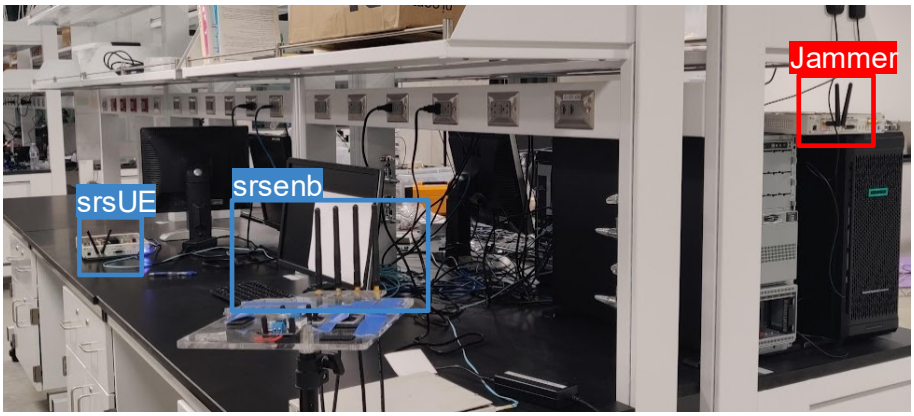

Figure 5.5. Experiment set up with the active elements (srsenb, srsUE, and jammer) in a lab environment. We tested with 3 different jammer locations.

5.2.4 Over-the-air Experiments

We build an over-the-air setup (Figure 5.5) to validate our jamming monitoring and nulling application. We use three SDR USRP X300 as srsenb, srsue, and jammer. The RAN controller runs on the same server as srsenb through a ZMQ-based transport socket.

5.2.5 Monitoring Jammer Interference with BeamArmor

We evaluate BeamArmor's jammer monitoring with varying jammer power as discussed in Section 5.2.3. Inside the BeamArmor App (see Figure 5.4), the received IQ data $y1$ and $y2$ can be processed to monitor the spectrum of the UL communication link in terms of received signal power. Table 5.1 shows the interfering effect of a 5 MHz wideband jammer by listing the SINR metric for various jammer power levels. Most notably in this table, the SINR suffers a reduction of 12 dB between a wideband jammer transmitting at 0 dBm and 30 dBm, respectively. The srsenb application reports UL SINR by default.

5.2.6 Mitigating Jammer with BeamArmor

First, we present BeamArmor's capability to cancel received signals from a certain direction, such as those from a jammer source. Figures 5.6(a) and 5.6(b) show how well BeamArmor can suppress signal power in the case of a single-tone source (single-tone jammer)

Table 5.1. UL SINR vs. increasing jammer TX power.

Jammer Power (dBm)	UL SINR (dB)
0	14.18
5	13.62
10	13.36
15	11.81
20	9.30
25	5.78
30	2.85

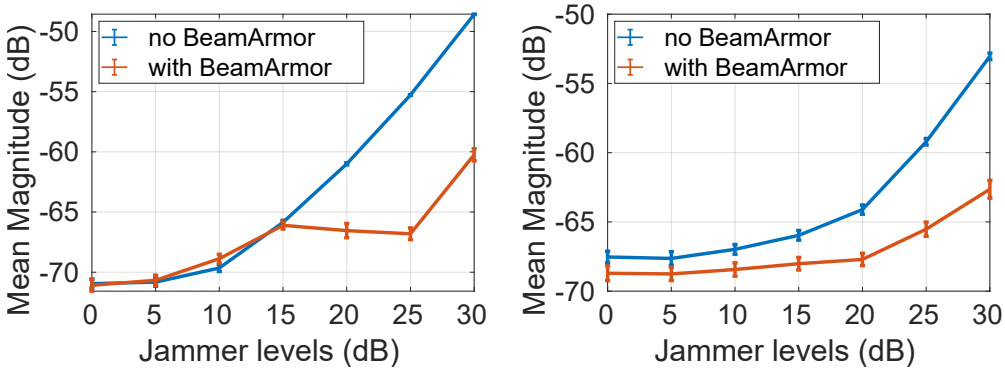

(a) Single-Tone jammer (b) 5 MHz Wideband jammer

Figure 5.6. BeamArmor's cancellation abilities: We present the mean received signal strength computed from received signal y in the form of $|y|^2$ as jammer power increases from 0 to 30 dBm. The red curve demonstrates how BeamArmor's null-steering decreases received jammer signal power.

and a 5 MHz random OFDM signal source (5 MHz wideband jammer). For a jammer source transmitting at 30 dBm gain, our anti-jamming solution achieves a 10 dB cancellation compared to the default base station operation with no active anti-jamming solution. The unexpected observation of baseline cancellation capability for the measurements belonging to a single-tone jammer between 0 dBm and 15 dBm transmit power could be caused by inaccurate null-steering factors α. As a reminder, our null-steering method assumes all signal power to be from the jammer source during the computation of α. In practice, at low jammer transmit powers, the received signal power at the base station will have relatively large parts of signal power originating from other sources than the jammer, resulting in a noisy value for α. Furthermore, the evaluation of cancellation capability would be more comprehensive if various jammer positions

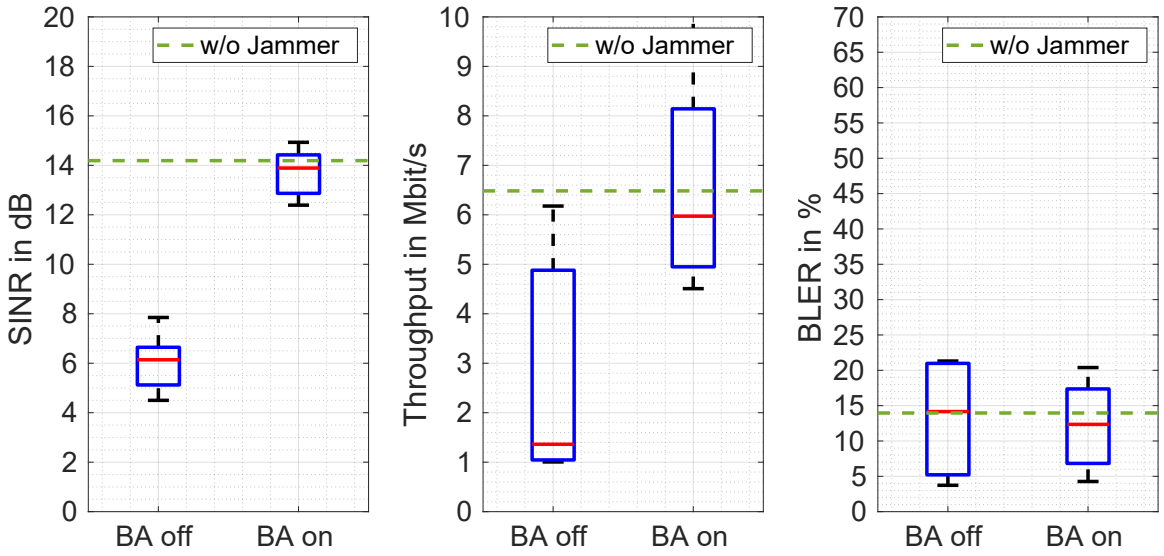

Figure 5.7. Metrics SINR, throughput, and BLER for **Single Tone jammer.** with and without BeamArmor's anti-jamming solution (BA on/off). The green dotted line shows the ground truth scenario without (w/o) any jammer present. The red line represents the median value of one boxplot.

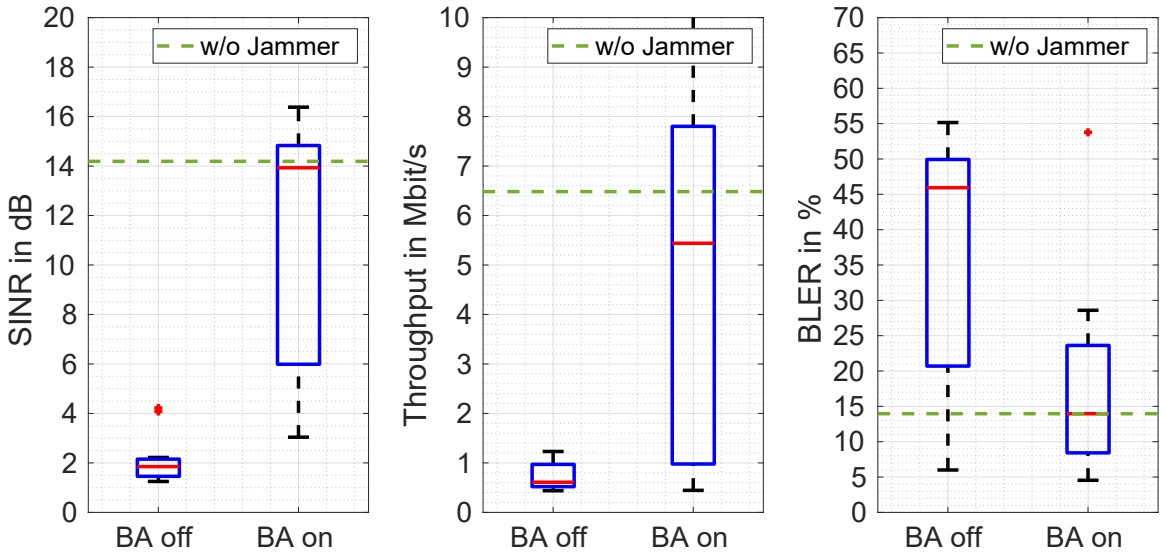

Figure 5.8. Metrics SINR, throughput, and BLER for **5 MHz Wideband jammer** with and without BeamArmor's anti-jamming solution (BA on/off). The green dotted line shows the ground truth scenario without (w/o) any jammer present. The red line represents the median value of one boxplot.

were measured, as opposed to one position, as done for Figure 5.6(a) and 5.6(b).

To evaluate the final performance of the RAN with BeamArmor's anti-jamming active, we conducted measurement runs as follows:

1. Turn srsenb and jammer on. Compute steering factor α.

2. Stop computing α and start applying α to UL processing at eNB side. This activates the anti-jamming.

3. Attach UE to eNB.

4. Start UL traffic from UE to eNB (10 Mbit/s UDP). Observe metrics.

In total, we evaluated our application from three different jammer positions, where each position's unique feature is the angle of the position to the position of the eNB, while the distance between the jammer and eNB remained at about 1 meter. As with our cancellation experiments, we evaluated BeamArmor's anti-jamming against a single-tone jammer and a wideband OFDM jammer.

The performance of our anti-jamming application was evaluated by comparing SINR, BLER, and throughput of the UL traffic received at the base station side. The communicated data is transmitted in blocks of fixed size, and BLER provides an indication of how reliable a communication link is. Figure 5.7 shows our results, where each boxplot bar contains 15 data points - the mean values of five individual measurement runs per each of the three jammer positions. As a benchmark, we added the mean of five measurements containing neither a jammer influence nor the application of BeamArmor's anti-jamming (dotted green line in Figure 5.7).

We conclude as the main takeaway: For all metrics and almost all combinations, our null-steering enabled anti-jamming solution provides on average (median) nearly benchmark performance if the jammer source is in close vicinity of the base station and therefore, the line-of-sight (LoS) path is very dominant. One exception is our measurements for BLER in the case of a single-tone jammer. Generally, the impact of a single-tone jammer is always weaker than a wideband jammer that interferes with several subcarrier frequencies. As mentioned in section 5.2.1, our null-steering method applies a scalar factor to all subcarriers. As a result, our method does not account for varying interference levels for jamming-affected subcarriers. This

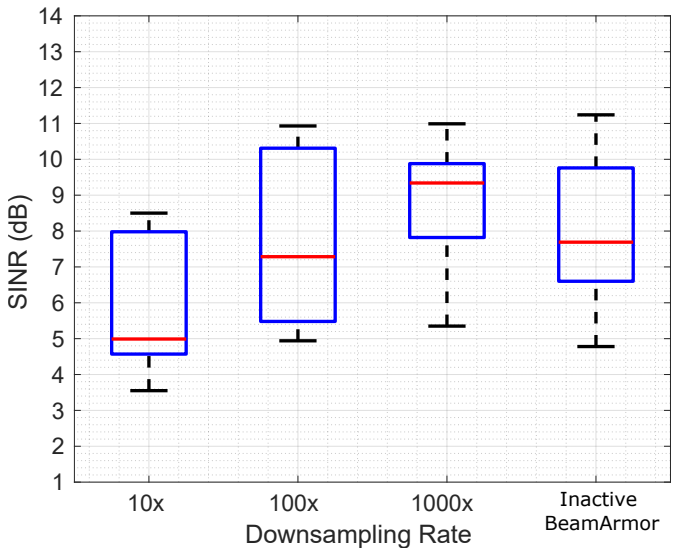

Figure 5.9. UL SINR during active BeamArmorcontroller communication at various down-sampling rates and at a fixed periodicity of every 10 TTI. UL SINR during an inactive BeamArmor controller serves as a benchmark performance.

circumstance can serve as an explanation for the large variance observed in our evaluation of BeamArmor in Figure 5.8.

5.2.7 Platform Evaluation

We decided that observing the UL traffic quality in the form of SINR (see Figure 5.9) would be sufficient to validate the functionality and performance of our RAN controller platform. In our measurements, we varied the rate at which we downsample y_1 and y_2 prior to sending, and fixed a send-periodicity of every 10 TTI. We compare the SINR metric with the case of "Inactive BeamArmor ", which is equivalent to infinite downsampling (no data is sent to the controller).

The result proves the functionality of BeamArmor's controller concept and does so without damage to UL traffic quality. Without downsampling, the RAN performance is affected, as evident from the low SINR of a median of 5 dB compared to the benchmark SINR of 7.5 dB. In our measurements of six runs per downsampling rate, 1000x downsampling performance is as good as the benchmark measurement. Thus, the RAN controller platform, deployed alongside an end-end RAN, is capable of providing PHY layer data in near real-time and enables MIMO

applications such as anti-jamming.

5.2.8 Related Work

The growing adoption of O-RAN standards has sparked a collective interest in RAN controller development across both industry and academia. This work advances the concept by focusing on the design of controllers and the development of applications tailored for MIMO applications.

■ **Controller platforms:** Janus [77] and FlexRIC [192] offer analytics and control capabilities for RANs in compliance with the O-RAN architecture. Janus [77] provides real-time control and inference over RAN, including support for streaming raw IQ samples to the controller. However, Janus [77] is built in an industry setting with a proprietary RAN stack, which is not compatible with open-source stacks such as srsRAN [201]. Nonetheless, our BeamArmor anti-jamming application can be demonstrated on Janus when this tool is available with open-source RANs such as srsRAN or Open-Air Interface [161] in the future.

■ **Testbeds for anti-jamming:** The study by [90] focuses on investigating the effects of jamming on the UL channel using a testbed powered by srsLTE and USRPs. However, it is important to note that the jammer is connected to the system using wired connections rather than a wireless link. Additionally, the study does not incorporate any anti-jamming mechanisms within the srsLTE framework. In [184], the researchers implemented jamming in the LTE UL using the Rohde & Schwarz radio communication tester CMW500 and [237] uses USRPs. They employed a signal generator and spectrum analyzer with a wired setup to investigate the power levels and EVM (Error Vector Magnitude) of the jammer, focusing on the PHY layer without considering the impact on the full cellular stack. Other testbeds are developed for WiFi [46, 91] or millimeter-wave devices [142]. In contrast, we study anti-jamming within a cellular full-stack srsRAN with wireless scenarios in real-time.

■ **Theoretical works on anti-jamming:** Several theoretical works on anti-jamming consider large antennas Massive MIMO systems [221, 36, 41, 214, 162, 234, 146, 32] For instance, the

research in [221] develops anti-jamming techniques for massive MIMO, utilizing random matrix theory. However, it should be noted that these advanced techniques are not required in small-scale MIMO systems that are commonly deployed in cellular stacks.

5.2.9 Limitations and Future Directions

BeamArmor demonstrates a real-time MIMO application for anti-jamming, offering ample opportunities for future expansion and development.

■ **Advanced jammers:** BeamArmor currently targets interference originating from a single LoS jammer due to the system utilizing only two antennas for UL reception. With the continuous evolution of open-source cellular software stacks, we envision an upgrade to BeamArmor with more than 2 antennas to handle multiple jammers or strong multi-paths NLoS scenarios. While we did not evaluate moving or duty-cycled jammer sources, BeamArmor could be operated such that the null-steering is continuously re-set towards a new jammer location. In any case, this is limited to the frequency at which the steering factor α can be computed, which is, at maximum, every TTI or every 1 ms.

■ **Frequency-selective Channel:** We assumed a frequency-flat channel resulting from a strong LoS component from a jammer in close proximity to the base station. To account for frequency-selective channel conditions, BeamArmor can be extended such that the null-steering factor α is computed for each subcarrier.

■ **Downlink Jammer:** We implemented BeamArmor for the UL traffic at the base station. Jammer interference on the DL at the user device would still harm or break a full duplex communication link. While some use cases of cellular communication might only depend on UL traffic, BeamArmor's null-steering technique could be implemented along the PHY layer processing of UEs (provided they are multi-antenna receivers) to provide jammer mitigation on the DL.

■ **O-RAN specifications:** While the controller implementation is motivated to comply with the O-RAN specifications, it remains an open piece of work to refine the utilized ZMQ interface to

match the O-RAN specifications exactly.

5.3 Conclusion and Outlook

With BeamArmor, we present a null-steering-enabled approach to mitigate the influence of jammer interference. The application, as well as the controller functioning as the host, are evaluated in over-the-air experiments for their viability. The possibility of utilizing the BeamArmor controller platform to demonstrate additional, innovative MIMO applications such as beamforming, multiplexing, antenna selection, user selection, power conservation, channel prediction, localization, and tracking based on recorded IQ samples poses a promising direction to extend this work.

We also briefly discussed mmSpoof which for the first time implements a practical and truly wireless attack on a mmWave FMCW radar. mmSpoof is designed to work without any prior knowledge about the victim radar and does not require any kind of synchronization. This makes mmSpoof robust to several of the current defense measures proposed by recent literature, such as frequency hopping. Further, with extensive experimentation, we show mmSpoof can independently spoof distance and velocity, creating any desired pattern for a controlled attack on victim radars. Finally, we discuss defense measures that would increase the safety against spoofing attacks like mmSpoof. We believe that the design principles described in mmSpoof highlight the fundamental challenges in securing any passive reflection based ranging and localization system.

5.4 Acknowledgements

BeamArmor research was supported in part by the NSF grant 2030245 and the U.S. Department of Defense N66001-23-F-0484 P00001. mmSpoof research was partially supported by NSF grants 2211805 and 2144914. mmSpoof is a joint work with my colleague Rohith Reddy Vennam as the lead author. We thank the anonymous reviewers and Shephard for

providing insightful feedback and the members of WCSNG lab, UC San Diego for their feedback throughout the process.

Chapter 5, in part, is a reprint of the material as it appears in F Zumegen, **IK Jain**, D Bharadia , “BeamArmor: Anti-Jamming in 5G Cellular Networks with MIMO Null-steering” accepted for publication in *ACM Hotmobile* 2024. The dissertation author was the primary investigator and author of this paper.

Chapter 6

Conclusion

6.1 Concluding Remarks

This thesis provides a vision for building the next generation of wireless networks towards 6G and beyond. Our founding fathers in wireless communication from Maxwell to Marconi used to work with circuits and devices. Then comes Shannon's theory and we simplified those complex circuits with simpler linear blocks and channel models and built the foundations of modern communication using OFDM and MIMO techniques. As we build the next-generation network, we need to understand the interplay between circuits and theory to push for even lower power, lower latency, higher reliability, and high scalability. Finally, when we deploy this network, we collect a lot of data. So we should take advantage of recent advances in data-driven techniques with advanced signal processing and machine learning to improve communication, sensing, and security for wireless networks. As a systems researcher, I have worked across all these domains connecting circuits, theory, and data. This approach has inspired me to develop new algorithms, new hardware designs, and new applications for the next generation of networks. Here I provide some remarks on our work in different areas of communication, sensing, and security and conclude this thesis.

- **mmReliable: Reliable and High-Throughput mmWave Links:** Millimeter-wave communication with high throughput and high reliability is poised to be a game-changer for V2X and VR applications. However, mmWave links are notorious for low reliability

since they suffer from frequent outages due to blockage and user mobility. We build mmReliable [111], a reliable mmWave system that implements multi-beamforming and user tracking to handle environmental vulnerabilities. It creates constructive multi-beam patterns and optimizes their angle, phase, and amplitude to maximize the signal strength at the receiver. Multi-beam links are reliable since they are resilient to occasional blockages of few constituent beams compared to a single-beam system. We implemented mmReliable on a 28 GHz testbed with 400 MHz bandwidth, and a 64-element phased array supporting 5G NR waveforms. Rigorous indoor and outdoor experiments demonstrated that mmReliable achieves close to 100% reliability and provides 2.3x improvement in the throughput-reliability product than traditional single-beam systems.

- **mmFlexible: Flexible and Scalable Multi-user mmWave Networks:** Modern mmWave systems have limited scalability due to a lack of flexible frequency multiplexing. That is, all the frequency components in the signal are beamformed to one direction via pencil beams and cannot be streamed to other user directions. We presented a new flexible mmWave system called mmFlexible [109] that enabled flexible directional frequency multiplexing, where different frequency components of the mmWave signal can be beamformed in multiple arbitrary directions with the same pencil beam carrying the same power as the traditional signal beam. This system had two key contributions: (1) a novel mmWave front-end architecture called a delay-phased array that uses a variable delay and variable phase element to create the desired frequency-direction response. (2) a novel algorithm called FSDA (Frequency-space to delay-antenna) to estimate delay and phase values for the real-time operation of the delay-phased array. Through evaluations with mmWave channel traces, we showed that mmFlexible provides a 60-150% reduction in worst-case latency compared to baseline.

- **FlexLink: Decoupling Control and Data Beams in mmWave Networks:** Control signals are vital in mmWave to have accurate beamforming vectors and achieve high beamforming

gain. Data and control traffic require maintaining different directional beams, one for data and another for control. Operating two beams over time sequentially is not resource-efficient. Since mmWave and higher frequencies come with extremely high bandwidth, we could utilize orthogonal frequencies for control and data, but these frequency resources are not efficiently utilized with current antenna architectures. A phased array is used to create a directional beam that radiates all frequency subcarriers in the signal in one fixed direction. It means all the subcarriers for data and control are beamformed to one direction. Naturally, we could deploy our frequency-dependent multi-beamforming to create two beams, one for data frequencies and another for control frequencies. In FlexLink, we discussed the hardware details of the delay-phased array, a novel antenna architecture to enable frequency-dependent beamforming. We then create a fast O(1) closed-form formula for estimating delay and phase values. Finally, we achieve some important properties of the delay-phased array such as support for a wide fraction of bandwidth-parts. We built this antenna array and showed almost double spectral efficiency for the data signal when it is decoupled with the control signal.

- **CommRad: Collaborative Bi-Directional Sensing-Driven Millimeter-Wave Networks:** We addressed the challenge of maintaining directional mmWave links in dynamic mobile environments due to mobility-induced disruptions and blockage. we introduced mmReliable [108], a sensing-driven solution incorporating a radar sensor at the base station to track mobile users and maintain directional beams even under blockages. While radar provides high-resolution object tracking, it suffers from a fundamental challenge of lack of context, i.e., it cannot discern which objects in the environment represent active users, reflectors, or blockers. To obtain this contextual awareness, mmReliable unites wireless sensing capabilities of bi-static radio communication with the mono-static radar sensor, allowing radios to provide initial context to radar sensors. Subsequently, the radar aids in user tracking and sustains mobile links even in obstructed scenarios, resulting in robust

and high-throughput directional connections for all mobile users at all times. We evaluated this collaborative radar-radio framework using a 28 GHz mmWave testbed integrated with a radar sensor in various indoor and outdoor scenarios, demonstrating a 2.5x improvement in median throughput compared to a non-collaborative baseline.

- **Security Challenges in Communication and Sensing Systems:** Spoofing and Jamming attacks pose a serious threat to the normal functioning of communication and sensing networks. We briefly discussed mmSpoof [220], a demonstration of a spoofing attack on mmWave automotive radars using a reflect array. We then introduced BeamArmor [260], a practical real-time application to monitor and mitigate jamming attacks in cellular networks. Hosted on a custom RAN controller utilizing the open-source srsRAN stack, it mitigates jamming by deploying beam-nulling techniques with MIMO antennas. This ensures reduced interference impact on desired communication signals. The system ensures real-time operation of the RAN stack by efficiently managing computational resources, meeting the time requirements of the cellular network's physical layer functionalities, and seamlessly integrating with srsRAN. We evaluated BeamArmor App with over-the-air experiments that demonstrated over 10 dB of SNR improvement in real-time by suppressing the jamming signal through beam-nulling using only two antennas. This leads to a broader vision including expanding the BeamArmor platform to other real-time MIMO applications, such as advanced beamforming and localization techniques, and integrating it into O-RAN standards.

6.2 Future Directions

My research in this thesis has been focused on the advancement of mmWave technology, leading to a new era of communication and sensing. As we plan for the future, I envision a multi-modal approach, one that unites a multitude of sensors, from Radars and lidars to cameras, and a diverse array of protocols, including mmWave, WiFi, and UWB. Together, these elements will

collaboratively elevate both communication and sensing capabilities. In addition, our exploration extends beyond the boundaries of our terrestrial realm to embrace non-terrestrial mmWave networks, including low-earth orbit satellites, aimed at revolutionizing IoT, internet connectivity, and sensing on a larger scale. My vision also encompasses the creation and deployment of a network of reflective surfaces designed to enhance mmWave coverage sustainably. To this end, we can build an end-end FR2 cellular platform that connects antenna beam switching with higher layers of the cellular stacks, thus bringing opportunity in cross-layer optimization and application awareness in optimizing the cellular networks.

AI for Multi-modal Sensing and Communications: In this thesis, I have worked towards connecting two sensor modalities, the monostatic radar sensor, and bistatic radio sensor, into an integrated framework similar to CommRad [108]. This is just the beginning, and many research opportunities lie in sharing the same time-frequency-antenna resources optimally, designing joint waveforms for communication and sensing, and mitigating interference between radio receiver and sensing receiver for optimal performance. Additionally, adding more sensors such as cameras, Lidars, and IMU into a collaborative learning framework with artificial intelligence and diverse radio-frequency-based sensing protocols such as WiFi, UWB, and mmWave systems provides a unique opportunity to enhance sensing accuracy and resolution while making communication radio more efficient and low overhead.

Reliability in Non-Terrestrial mmWave Networks: My research on reliable and scalable communication has applications beyond the terrestrial realm to include non-terrestrial networks such as low-earth orbit satellites. The satellites in mmWave Ku/Ka bands (e.g., Starlink) suffer from similar problems of reliability with narrow 'pencil beams' and interference with terrestrial networks operating in the same bands [160]. These challenges are even magnificent in non-terrestrial links due to higher mobility, broader coverage, and a ground station that is as compact and efficient as today's smartphone. To this end, future research directions include building a testbed for the ground station using phased arrays and measurement campaigns to

understand reliability and interference problems and solve them through new architectures (e.g., true-time delay) and theory (e.g., multi-beamforming).

Low-power mmWave/THz Infrastructure: A key aspect of this thesis is to ensure sustainable and low-power mmWave and sub-THz infrastructure while delivering reliable and high-throughput performance. Due to high path loss and low coverage, mmWave infrastructure has to be deployed densely, which would increase the carbon footprint [93, HotCarbon 2022]. To improve range and reliability with low power, we can deploy repeaters and reflective intelligent surfaces (RIS) in the environment. Handling directional mmWave links with RIS, channel estimation, user mobility, and integration with our multi-beamforming architecture are challenging problems.

Machine Learning and Compute for NextG Cellular Networks: Cellular networks have become more accessible with the open-RAN initiative and the development of open-source software radio stacks such as srsRAN and Open-Air Interface enhanced by RAN intelligent controllers. Having built a BeamArmor controller [260, Hotmobile 2024] and mmWave testbed in this thesis, we can work towards building an end-end mmWave cellular stack with software radios aimed at optimizing the cellular network for various applications such as video streaming, AR/VR, and vehicular connectivity. We have already built an end-end VR streaming pipeline in [177, WPMC 2022]. Thanks to the decoupling of the RAN and the controller, computation-heavy and machine learning-driven inference can be adopted to optimize the cellular NextG network in real-time. Therefore, we can explore these methods to solve open problems in resource scheduling, channel prediction, RAN parameter tuning, and interference management, among others.

Bibliography

[1] ZeroMQ Documentation. http://zeromq.org. Accessed May 10, 2023.

[2] IEEE 802.11ad: Enhancements for Very High Throughput in the 60 GHz Band. *IEEE*, 2009.

[3] Verizon to test 5G at 28 GHz in Texas with Samsung. https://www.fiercewireless.com/tech/verizon-to-test-5g-at-28-ghz-texas-samsung, Feb 2016.

[4] 5G; Study on channel model for frequencies from 0.5 to 100 GHz. TR 138 901 V14.0.0, 3rd Generation Partnership Project (3GPP), 2017-05.

[5] Analog Devices, AD-FMCDAQ2-EBZ. https://wiki.analog.com/resources/eval/user-guides/ad-fmcdaq2-ebz, September 2019.

[6] Analog Devices, KCU 105. https://www.xilinx.com/products/boards-and-kits/kcu105.html, Sep 2019.

[7] Cinetics axis-360 motion control system. https://cinetics.com/axis360-pro/, Sep 2019.

[8] Qorvo, RFFC5071A: 85 - 4200 MHz Wideband Synthesizer / VCO with Integrated 6 GHz RF Mixer. https://www.qorvo.com/products/p/RFFC5071A, September 2019.

[9] Talon AD7200 Multi-Band Wi-Fi Router. https://www.tp-link.com/us/home-networking/wifi-router/ad7200/, Sep 2019.

[10] 5G NR (New Radio) Release 16. 3GPP. https://www.3gpp.org/release-16, Jul 2020.

[11] IBM device paves the way for 5G to reach full potential. https://www.ibm.com/blogs/research/2020/06/ibm-device-5g-full-potential/, Jun 2020.

[12] India's indigenous 5G testbed project. http://www.ee.iitm.ac.in/5g/, Jul 2020.

[13] Mmobile mobility dataset. https://wcsng.ucsd.edu/mmobile/, Jul 2020.

[14] Speech and multimedia Transmission Quality (STQ); QoS parameters and test scenarios for assessing network capabilities in 5G performance measurements. TR ETSI TR 103 702 V1.1.1 , 3rd Generation Partnership Project (3GPP), 2020-11.

[15] 12v battery. https://www.amazon.com/65Watts-Portable-Charger-Battery-Notebooks/dp/B07P87JTVT, Jul 2023.

[16] 5G NR Frequency Bands. https://www.5gradio.com/5g-technology/5g-nr-frequency-bands/, September 2023.

[17] Analog devices inc., adalm pluto sdr. https://www.analog.com/en/design-center/evaluation-hardware-and-software/evaluation-boards-kits/adalm-pluto.html, Jul 2023.

[18] Analog devices inc., adf5356 module. https://www.analog.com/en/products/adf5356.html, Jul 2023.

[19] Anti-jamming market size & share analysis - growth trends & forecasts (2023 - 2028). https://www.mordorintelligence.com/industry-reports/anti-jamming-equipment-market-growth, 2023.

[20] Bladerf. https://www.nuand.com/bladerf-1/, 2023.

[21] Digilent cmod artix a7-35t fpga. https://digilent.com/reference/programmable-logic/cmod-a7/start, Jul 2023.

[22] Ettus usrp b210. https://www.ettus.com/all-products/ub210-kit/, Jul 2023.

[23] Extreme waves 5g phased arrays. https://www.extreme-waves.com/5g, Jul 2023.

[24] Gnu radio. https://www.gnuradio.org/, Jul 2023.

[25] Hackrf one. https://greatscottgadgets.com/hackrf/one/, 2023.

[26] Inras radarbook at 24 ghz and 77 ghz. https://inras.at/en/radarbook2/, June 2024.

[27] 3GPP. 5G NR (New Radio) Release 16. 3GPP. https://www.3gpp.org/release-16, Oct 2019.

[28] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. Enabling high-quality untethered virtual reality. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 531–544, 2017.

[29] Omid Abari, Haitham Hassanieh, Michael Rodreguiz, and Dina Katabi. Poster: A millimeter wave software defined radio platform with phased arrays. In *Proceedings of*

the 22nd Annual International Conference on Mobile Computing and Networking, pages 419–420, 2016.

[30] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. Millimeter wave communications: From point-to-point links to agile network connections. In *HotNets*, pages 169–175, 2016.

[31] Iftikhar Ahmad, Ahsan Raza Khan, Rao Naveed Bin Rais, Ahmed Zoha, Muhammad Ali Imran, and Sajjad Hussain. Vision-assisted beam prediction for real world 6g drone communication. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–7. IEEE, 2023.

[32] Hossein Akhlaghpasand, Emil Björnson, and S Mohammad Razavizadeh. Jamming suppression in massive MIMO systems. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 67(1):182–186, 2019.

[33] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W Heath. Channel estimation and hybrid precoding for millimeter wave cellular systems. *IEEE Journal of Selected Topics in Signal Processing*, 8(5):831–846, 2014.

[34] Muhammad Alrabeiah, Andrew Hredzak, and Ahmed Alkhateeb. Millimeter wave base stations with cameras: Vision-aided beam and blockage prediction. In *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, pages 1–5. IEEE, 2020.

[35] alvarovalcarce. Nokia wireless suite - github reference. https://github.com/nokia/wireless-suite (2020/11/12).

[36] Youness Arjoune and Saleh Faruque. Smart jamming attacks in 5G new radio: A review. In *2020 10th annual computing and communication workshop and conference (CCWC)*, pages 1010–1015. IEEE, 2020.

[37] Canan Aydogdu, Fan Liu, Christos Masouros, Henk Wymeersch, and Mats Rydström. Distributed radar-aided vehicle-to-vehicle communication. In *2020 IEEE Radar Conference (RadarConf20)*, pages 1–6. IEEE, 2020.

[38] Irmak Aykin, Berk Akgun, Mingjie Feng, and Marwan Krunz. Mamba: A multi-armed bandit framework for beam tracking in millimeter-wave systems. In *2020 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2020.

[39] Irmak Aykin, Berk Akgun, and Marwan Krunz. Multi-beam transmissions for blockage resilience and reliability in millimeter-wave systems. *IEEE Journal on Selected Areas in Communications*, 37(12):2772–2785, 2019.

[40] Irmak Aykin, Berk Akgun, and Marwan Krunz. Smartlink: Exploiting channel clustering

effects for reliable millimeter wave communications. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 1117–1125. IEEE, 2019.

[41] Farhan M Aziz, Jeff S Shamma, and Gordon L Stüber. Jammer-type estimation in LTE with a smart jammer repeated game. *IEEE Transactions on Vehicular Technology*, 66(8):7422–7431, 2017.

[42] Kshitiz Bansal, Keshav Rungta, Siyuan Zhu, and Dinesh Bharadia. Pointillism: accurate 3d bounding box estimation with multi-radars. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 340–353, 2020.

[43] C Nicolas Barati, S Amir Hosseini, Marco Mezzavilla, Parisa Amiri-Eliasi, Sundeep Rangan, Thanasis Korakis, Shivendra S Panwar, and Michele Zorzi. Directional initial access for millimeter wave cellular systems. In *Signals, Systems and Computers, 2015 49th Asilomar Conference on*, pages 307–311. IEEE, 2015.

[44] C Nicolas Barati, S Amir Hosseini, Marco Mezzavilla, Thanasis Korakis, Shivendra S Panwar, Sundeep Rangan, and Michele Zorzi. Initial access in millimeter wave cellular systems. *IEEE Transactions on Wireless Communications*, 15(12):7926–7940, 2016.

[45] C Umit Bas, Rui Wang, Seun Sangodoyin, Thomas Choi, Sooyoung Hur, Kuyeon Whang, Jeongho Park, Charlie Jianzhong Zhang, and Andreas F Molisch. Outdoor to indoor propagation channel measurements at 28 ghz. *IEEE Transactions on Wireless Communications*, 18(3):1477–1489, 2019.

[46] Suzan Bayhan, Piotr Gawłowicz, Anatolij Zubow, and Adam Wolisz. Null-while-talk: Interference nulling for improved inter-technology coexistence in lte-u and wifi networks. *Pervasive and Mobile Computing*, 56:71–87, 2019.

[47] Jacob Benesty, Israel Cohen, and Jingdong Chen. *Array Processing*. Springer, 2019.

[48] Veljko Boljanovic and Danijela Cabric. Compressive estimation of wideband mmw channel using analog true-time-delay array. In *2021 IEEE Workshop on Signal Processing Systems (SiPS)*, pages 170–175. IEEE, 2021.

[49] Veljko Boljanovic, Han Yan, Chung-Ching Lin, Soumen Mohapatra, Deukhyoun Heo, Subhanshu Gupta, and Danijela Cabric. True-time-delay arrays for fast beam training in wideband millimeter-wave systems. *arXiv preprint arXiv:2007.08713*, 2020.

[50] Veljko Boljanovic, Han Yan, Chung-Ching Lin, Soumen Mohapatra, Deukhyoun Heo, Subhanshu Gupta, and Danijela Cabric. Fast beam training with true-time-delay arrays in wideband millimeter-wave systems. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 68(4):1727–1739, 2021.

[51] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[52] Derek Caudill, Jack Chuang, Sung Yun Jun, Camillo Gentile, and Nada Golmie. Real-time mmwave channel sounding through switched beamforming with 3-d dual-polarized phased-array antennas. *IEEE Transactions on Microwave Theory and Techniques*, 69(11):5021–5032, 2021.

[53] Gouranga Charan and Ahmed Alkhateeb. User identification: A key enabler for multi-user vision-aided communications. *IEEE Open Journal of the Communications Society*, 2023.

[54] Gouranga Charan, Muhammad Alrabeiah, Tawfik Osman, and Ahmed Alkhateeb. Camera based mmwave beam prediction: Towards multi-candidate real-world scenarios. *arXiv preprint arXiv:2308.06868*, 2023.

[55] Ruchir Chauhan. *A platform for false data injection in frequency modulated continuous wave radar*. Utah State University, 2014.

[56] Ruchir Chauhan, Ryan M Gerdes, and Kevin Heaslip. Demonstration of a false-data injection attack against an fmcw radar. *Embedded Security in Cars (ESCAR)*, 2014.

[57] Hung-Ruei Chen. Fmcw radar jamming techniques and analysis. Technical report, NAVAL POSTGRADUATE SCHOOL MONTEREY CA, 2013.

[58] Leyan Chen, Kai Liu, Zhibo Zhang, and Baoqi Li. Beam selection and power allocation: Using deep learning for sensing-assisted communication. *IEEE Wireless Communications Letters*, 2023.

[59] Tingjun Chen, Manav Kohli, Tianyi Dai, Angel Daniel Estigarribia, Dmitry Chizhik, Jinfeng Du, Rodolfo Feick, Reinaldo A Valenzuela, and Gil Zussman. 28 GHz channel measurements in the COSMOS testbed deployment area. In *Proceedings of the 3rd ACM Workshop on Millimeter-wave Networks and Sensing Systems*, pages 39–44, 2019.

[60] Zhanghua Chen and Jianxiao Xie. Radar aided active prediction based on deep learning. In *2023 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pages 1–5. IEEE, 2023.

[61] Moon-Kyu Cho, Ickhyun Song, and John D Cressler. A true time delay-based sige bi-directional t/r chipset for large-scale wideband timed array antennas. In *2018 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, pages 272–275. IEEE, 2018.

[62] Kaiyan Cui, Qiang Yang, Leming Shen, Yuanqing Zheng, and Jinsong Han. Integrated sensing and communication between daily devices and mmwave radars. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, pages 831–832,

2022.

[63] Nicholas C Currie and Charles E Brown. *Principles and applications of millimeter-wave radar*. Artech House Norwood, MA, 1987.

[64] Jian Dai, Xinhong Hao, Ping Li, Ze Li, and Xiaopeng Yan. Antijamming design and analysis of a novel pulse compression radar signal based on radar identity and chaotic encryption. *IEEE Access*, 8:5873–5884, 2020.

[65] Science Daily. New wireless system for greater 5G access. https://www.sciencedaily.com/releases/2023/05/230518120909.htm, May 2023.

[66] Umut Demirhan and Ahmed Alkhateeb. Radar aided 6g beam prediction: Deep learning algorithms and real-world demonstration. In *2022 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 2655–2660. IEEE, 2022.

[67] Umut Demirhan and Ahmed Alkhateeb. Radar aided proactive blockage prediction in real-world millimeter wave systems. In *ICC 2022-IEEE International Conference on Communications*, pages 4547–4552. IEEE, 2022.

[68] Tan Tai Do, Emil Björnson, Erik G Larsson, and S Mohammad Razavizadeh. Jamming-resistant receivers for the massive MIMO uplink. *IEEE Transactions on Information Forensics and Security*, 13(1), 2017.

[69] Manideep Dunna, Kshitiz Bansal, Sanjeev Anthia Ganesh, Eamon Patamasing, and Dinesh Bharadia. R-fiducial: Millimeter wave radar fiducials for sensing traffic infrastructure. In *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, pages 1–7. IEEE, 2023.

[70] Manideep Dunna, Chi Zhang, Daniel Sievenpiper, and Dinesh Bharadia. Scattermimo: Enabling virtual mimo with smart surfaces. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.

[71] Omar El Ayach, Robert W Heath, Shadi Abu-Surra, Sridhar Rajagopal, and Zhouyue Pi. Low complexity precoding for large millimeter wave MIMO systems. In *2012 IEEE international conference on communications (ICC)*, pages 3724–3729. IEEE, 2012.

[72] Omar El Ayach, Sridhar Rajagopal, Shadi Abu-Surra, Zhouyue Pi, and Robert W Heath. Spatially sparse precoding in millimeter wave MIMO systems. *IEEE transactions on wireless communications*, 13(3):1499–1513, 2014.

[73] Interesting Engineering. A New Laser-Like System Could Put 5G in Your Home. https://interestingengineering.com/innovation/a-new-laser-like-system-could-put-5g-in-your-home, Aug 2021.

[74] Tech Explore. New technology could bring the fastest version of 5G to your home and workplace. https://techxplore.com/news/2021-08-technology-fastest-version-5g-home.html, Aug 2021.

[75] FCC. Report and order and further notice of proposed rulemaking, 2016.

[76] Mingjie Feng, Shiwen Mao, and Tao Jiang. Dealing with link blockage in mmwave networks: D2d relaying or multi-beam reflection? In *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 1–5. IEEE, 2017.

[77] Xenofon Foukas, Bozidar Radunovic, Matthew Balkwill, and Zhihua Lai. Taking 5G RAN analytics and control to a new level. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–16, 2023.

[78] Santosh Ganji, Jaewon Kim, and PR Kumar. Terra: blockage resilience in outdoor mm-wave networks. In *Proceedings of the SIGCOMM'22 Poster and Demo Sessions*, pages 55–57. 2022.

[79] Venkata Siva Santosh Ganji, Tzu-Hsiang Lin, Francisco A Espinal, and PR Kumar. Unblock: Low complexity transient blockage recovery for mobile mm-wave devices. In *2021 International Conference on COMmunication Systems & NETworkS (COMSNETS)*, pages 501–508. IEEE, 2021.

[80] Feifei Gao, Liangyuan Xu, and Shaodan Ma. Integrated sensing and communications with joint beam-squint and beam-split for mmwave/thz massive mimo. *IEEE Transactions on Communications*, 71(5):2963–2976, 2023.

[81] Seyed Kasra Garakoui, Eric AM Klumperink, Bram Nauta, and Frank E van Vliet. Compact cascadable gm-c all-pass true time delay cell with reduced delay variation over frequency. *IEEE journal of solid-state circuits*, 50(3):693–703, 2015.

[82] Dolores Garcia, Jesus Omar Lacruz, Pablo Jiménez Mateo, and Joerg Widmer. Polar: Passive object localization with ieee 802.11 ad using phased antenna arrays. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 1838–1847. IEEE, 2020.

[83] Robin Garg, Gaurav Sharma, Ali Binaie, Sanket Jain, Sohail Ahasan, Armagan Dascurcu, Harish Krishnaswamy, and Arun S Natarajan. A 28-ghz beam-space mimo rx with spatial filtering and frequency-division multiplexing-based single-wire if interface. *IEEE Journal of Solid-State Circuits*, 2020.

[84] Zulkuf Genc, Umar H Rizvi, Ertan Onur, and Ignas Niemegeers. Robust 60 GHz indoor connectivity: Is it possible with reflections? In *2010 IEEE 71st vehicular technology*

conference, pages 1–5. IEEE, 2010.

[85] Erfan Ghaderi and Subhanshu Gupta. A four-element 500-mhz 40-mw 6-bit adc-enabled time-domain spatial signal processor. *IEEE Journal of Solid-State Circuits*, 56(6):1784–1794, 2020.

[86] Erfan Ghaderi, Ajith Sivadhasan Ramani, Arya A Rahimi, Deukhyoun Heo, Sudip Shekhar, and Subhanshu Gupta. An integrated discrete-time delay-compensating technique for large-array beamformers. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(9):3296–3306, 2019.

[87] Yasaman Ghasempour. Multi-Stream Beam-Training for mmWave MIMO Networks. In *Proceedings of the 24rd Annual International Conference on Mobile Computing and Networking*. ACM, 2018.

[88] Yasaman Ghasempour, Rabi Shrestha, Aaron Charous, Edward Knightly, and Daniel M Mittleman. Single-shot link discovery for terahertz wireless networks. *Nature communications*, 11(1):1–6, 2020.

[89] Marco Giordani, Michele Polese, Arnab Roy, Douglas Castor, and Michele Zorzi. A tutorial on beam management for 3GPP NR at mmWave frequencies. *IEEE Communications Surveys & Tutorials*, 21(1):173–196, 2018.

[90] Felix Girke, Fabian Kurtz, Nils Dorsch, and Christian Wietfeld. Towards resilient 5G: Lessons learned from experimental evaluations of LTE uplink jamming. In *2019 IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1–6. IEEE, 2019.

[91] Shyamnath Gollakota, Samuel David Perli, and Dina Katabi. Interference alignment and cancellation. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 159–170, 2009.

[92] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, March 2014.

[93] Agrim Gupta, Ish Jain, and Dinesh Bharadia. Multiple smaller base stations are greener than a single powerful one: Densification of wireless cellular networks. In *1st Workshop on Sustainable Computer Systems Design and Implementation (HotCarbon)*, 2022.

[94] Ismail Guvenc, Farshad Koohifar, Simran Singh, Mihail L Sichitiu, and David Matolak. Detection, tracking, and interdiction for amateur drones. *IEEE Communications Magazine*, 56(4):75–81, 2018.

[95] Muhammad Kumail Haider, Yasaman Ghasempour, Dimitrios Koutsonikolas, and Ed-

ward W Knightly. Listeer: Mmwave beam acquisition and steering by tracking indicator leds on wireless aps. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 273–288. ACM, 2018.

[96] Muhammad Kumail Haider and Edward W Knightly. Mobility resilience and overhead constrained adaptation in directional 60 ghz wlans: protocol design and system implementation. In *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 61–70, 2016.

[97] Jürgen Hasch, Eray Topak, Raik Schnabel, Thomas Zwick, Robert Weigel, and Christian Waldschmidt. Millimeter-wave technology for automotive radar sensors in the 77 ghz frequency band. *IEEE transactions on microwave theory and techniques*, 60(3):845–860, 2012.

[98] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. Fast millimeter wave beam alignment. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 432–445. ACM, 2018.

[99] Robert W Heath, Nuria Gonzalez-Prelcic, Sundeep Rangan, Wonil Roh, and Akbar M Sayeed. An overview of signal processing techniques for millimeter wave MIMO systems. *IEEE journal of selected topics in signal processing*, 10(3):436–453, 2016.

[100] Yuqiang Heng and Jeffrey G Andrews. Grid-free mimo beam alignment through site-specific deep learning. *IEEE Transactions on Wireless Communications*, 23(2):908–921, 2023.

[101] Yuqiang Heng, Yu Zhang, Ahmed Alkhateeb, and Jeffrey G Andrews. Site-specific beam alignment in 6g via deep learning. *arXiv preprint arXiv:2403.16186*, 2024.

[102] Feng Hu and Koen Mouthaan. A 1–20 ghz 400 ps true-time delay with small delay error in 0.13 μm cmos for broadband phased array antennas. In *2015 IEEE MTT-S International Microwave Symposium*, pages 1–3. IEEE, 2015.

[103] Jingqi Huang. *Demystifying Millimeter-Wave V2X: Testbed Measurement and Large-Scale Simulation*. PhD thesis, UC San Diego, 2020.

[104] Shoaib Imran, Gouranga Charan, and Ahmed Alkhateeb. Environment semantic aided communication: A real world demonstration for beam prediction. *arXiv preprint arXiv:2302.06736*, 2023.

[105] Shantanu Ingle and Madhuri Phute. Tesla autopilot: semi autonomous driving, an uptick for future autonomy. *International Research Journal of Engineering and Technology*, 3(9):369–372, 2016.

[106] Rafail Ismayilov, Megumi Kaneko, Takefumi Hiraguri, and Kentaro Nishimori. Adaptive beam-frequency allocation algorithm with position uncertainty for millimeter-wave mimo systems. In *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, pages 1–5. IEEE, 2018.

[107] Ish Kumar Jain, Rajeev Kumar, and Shivendra S Panwar. The impact of mobile blockers on millimeter wave cellular systems. *IEEE Journal on Selected Areas in Communications*, 37(4):854–868, 2019.

[108] Ish Kumar Jain, Suriyaa MM, and Dinesh Bharadia. Commrad: Context-aware sensing-driven millimeter-wave networks. *arXiv preprint arXiv:2407.08817*, 2024.

[109] Ish Kumar Jain, Rohith Reddy Vennam, Raghav Subbaraman, and Dinesh Bharadia. mmflexible: Flexible directional frequency multiplexing for multi-user mmwave networks, IEEE INFOCOM 2023. https://ieeexplore.ieee.org/document/10229065, 2023.

[110] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Demo and dataset for mmwave multi-beam tracking using mmobile 28 ghz testbed. In *Proceedings of the 12th ACM Wireless of the Students, by the Students, and for the Students (S3) Workshop*, pages 8–8, 2021.

[111] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Two beams are better than one: Towards reliable and high throughput mmwave links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM '21, page 488–502, New York, NY, USA, 2021. Association for Computing Machinery.

[112] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Two beams are better than one: towards reliable and high throughput mmwave links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 488–502, 2021.

[113] Ish Kumar Jain, Raghav Subbaraman, Tejas Harekrishna Sadarahalli, Xiangwei Shao, Hou-Wei Lin, and Dinesh Bharadia. mmobile: Building a mmwave testbed to evaluate and address mobility effects. In *Proceedings of the 4th ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, pages 1–6, 2020.

[114] Cheol Jeong, Jeongho Park, and Hyunkyu Yu. Random access in millimeter-wave beamforming cellular networks: issues and approaches. *IEEE Communications Magazine*, 53(1):180–185, 2015.

[115] Hyoungju Ji, Sunho Park, Jeongho Yeo, Younsun Kim, Juho Lee, and Byonghyo Shim. Ultra-reliable and low-latency communications in 5G downlink: Physical layer aspects. *IEEE Wireless Communications*, 25(3):124–130, 2018.

[116] Shuaifeng Jiang, Gouranga Charan, and Ahmed Alkhateeb. Lidar aided future beam

prediction in real-world millimeter wave v2i communications. *IEEE Wireless Communications Letters*, 12(2):212–216, 2022.

[117] Feng Jianli. Wideband fm anti-jamming method of fmcw radar. In *2021 6th International Conference on Intelligent Computing and Signal Processing (ICSP)*, pages 777–780. IEEE, 2021.

[118] Suraj Jog, Jiaming Wang, Junfeng Guan, Thomas Moon, Haitham Hassanieh, and Romit Roy Choudhury. Many-to-many beam alignment in millimeter wave networks. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 783–800, 2019.

[119] Mahmoud Kamel, Walaa Hamouda, and Amr Youssef. Ultra-dense networks: A survey. *IEEE Communications surveys & tutorials*, 18(4):2522–2545, 2016.

[120] Shinpei Kato, Eijiro Takeuchi, Yoshio Ishiguro, Yoshiki Ninomiya, Kazuya Takeda, and Tsuyoshi Hamada. An open approach to autonomous vehicles. *IEEE Micro*, 35(6):60–68, 2015.

[121] Fotios Katsilieris, Paolo Braca, and Stefano Coraluppi. Detection of malicious ais position spoofing by exploiting radar information. In *proceedings of the 16th international conference on information fusion*, pages 1196–1203. IEEE, 2013.

[122] Muhammad Asif Khan, Hamid Menouar, Aisha Eldeeb, Adnan Abu-Dayya, and Flora D Salim. On the detection of unauthorized drones-techniques and future perspectives: A review. *IEEE Sensors Journal*, 2022.

[123] Wahab Khawaja, Ozgur Ozdemir, Yavuz Yapici, Ismail Guvenc, and Yuichi Kakishima. Coverage enhancement for mm wave communications using passive reflectors. In *2018 11th Global Symposium on Millimeter Waves (GSMM)*, pages 1–6. IEEE, 2018.

[124] Kerim Kibaroglu, Mustafa Sayginer, Thomas Phelps, and Gabriel M Rebeiz. A 64-element 28-GHz phased-array transceiver with 52-dBm EIRP and 8–12-Gb/s 5G link at 300 meters without any calibration. *IEEE Transactions on Microwave Theory and Techniques*, 66(12):5796–5811, 2018.

[125] Aldebaro Klautau, Nuria González-Prelcic, and Robert W Heath. Lidar data for deep learning-based mmwave beam-selection. *IEEE Wireless Communications Letters*, 8(3):909–912, 2019.

[126] Woo-Hyun Ko, Ujwal Dinesha, Ushasi Ghosh, Srinivas Shakkottai, Dinesh Bharadia, and Raini Wu. Edgeric: Empowering realtime intelligent optimization and control in nextg networks. *arXiv preprint arXiv:2304.11199*, 2023.

[127] Rony Komissarov and Avishai Wool. Spoofing attacks against vehicular fmcw radar. In *Proceedings of the 5th Workshop on Attacks and Solutions in Hardware Security*, pages 91–97, 2021.

[128] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 269–282, 2015.

[129] Hansol Ku, Jinyue Song, Ding Zhang, Prasant Mohapatra, and Parth Pathak. Characterizing real-time radar-assisted beamforming in mmwave v2v links. In *2023 20th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 168–176. IEEE, 2023.

[130] Jesus O Lacruz, Rafael Ruiz Ortiz, and Joerg Widmer. A real-time experimentation platform for sub-6 ghz and millimeter-wave mimo systems. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 427–439, 2021.

[131] Jesus Omar Lacruz et al. mm-FLEX: an open platform for millimeter-wave mobile full-bandwidth experimentation. In *MobiSys*, pages 1–13, 2020.

[132] Min Li, Nayu Li, Huiyan Gao, Zijiang Zhang, Shaogang Wang, Yen-Cheng Kuan, Chunyi Song, Xiaopeng Yu, Qun Jane Gu, and Zhiwei Xu. An 800-ps origami true-time-delay-based cmos receiver front end for 6.5–9-ghz phased arrays. *IEEE Solid-State Circuits Letters*, 3:382–385, 2020.

[133] Ruifu Li, Han Yan, and Danijela Cabric. Rainbow-link: Beam-alignment-free and grant-free mmw multiple access using true-time-delay array. *IEEE Journal on Selected Areas in Communications*, 2022.

[134] Tianxiang Li, Haofan Lu, Reza Rezvani, Ali Abedi, and Omid Abari. Bringing wifi localization to any wifi devices. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks*, pages 46–52, 2022.

[135] Tianxiang Li, Mohammad Hossein Mazaheri, Kalaivani Kamalakannan, Haofan Lu, and Omid Abari. Can iot devices be powered up by future indoor wireless networks? In *Proceedings of the 25th International Workshop on Mobile Computing Systems and Applications*, pages 73–78, 2024.

[136] Chung-Ching Lin, Qiuyan Xu, Huan Hu, and Subhanshu Gupta. Design considerations of time-interleaved discrete-time beamformers toward wideband communications. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 2023.

[137] Enrique M. Lizarraga, Gabriel N. Maggio, and Alexis A. Dowhuszko. Hybrid beamform-

ing algorithm using reinforcement learning for millimeter wave wireless systems. In *2019 XVIII Workshop on Information Processing and Control (RPIC)*, pages 253–258, 2019.

[138] Adrian Loch, Hany Assasa, Joan Palacios, Joerg Widmer, Hans Suys, and Björn Debaillie. Zero overhead device tracking in 60 GHz wireless networks using multi-lobe beam patterns. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, pages 224–237, 2017.

[139] Hao Luo, Umut Demirhan, and Ahmed Alkhateeb. Millimeter wave v2v beam tracking using radar: Algorithms and real-world demonstration. *arXiv preprint arXiv:2308.01558*, 2023.

[140] Qian Ma, Hyunchul Chung, Yusheng Yin, Tumay Kanar, Samet Zihir, and Gabriel M Rebeiz. A 5G 24-30 GHz 2x32 Element Dual-Polarized Dual-Beam Phased Array Base-Station for 2x2 MIMO System. In *2019 IEEE Global Communications Conference (GLOBECOM)*, pages 1–5. IEEE, 2019.

[141] George R MacCartney, Theodore S Rappaport, and Sundeep Rangan. Rapid fading due to human blockage in pedestrian crowds at 5G millimeter-wave frequencies. In *GLOBECOM 2017-2017 IEEE Global Communications Conference*, pages 1–7. IEEE, 2017.

[142] Sohrab Madani, Suraj Jog, Jesus O. Lacruz, Joerg Widmer, and Haitham Hassanieh. Practical null steering in millimeter wave networks. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 903–921. USENIX Association, April 2021.

[143] Sohrab Madani, Suraj Jog, Jesús Omar Lacruz, Joerg Widmer, and Haitham Hassanieh. Practical null steering in millimeter wave networks. In *NSDI*, pages 903–921, 2021.

[144] Simone Mangiante, Guenter Klas, Amit Navon, Zhuang GuanHua, Ju Ran, and Marco Dias Silva. Vr is on the edge: How to deliver 360 videos in mobile networks. In *Proceedings of the Workshop on Virtual Reality and Augmented Reality Network*, pages 30–35, 2017.

[145] Shah Marjan, Lin Bai, and Chao Han. Split antenna array in millimeter wave for secure vehicular communication. In *MATEC Web of Conferences*, volume 173, page 02024. EDP Sciences, 2018.

[146] Gian Marti, Torben Kölle, and Christoph Studer. Mitigating Smart Jammers in Multi-User MIMO. *IEEE Transactions on Signal Processing*, 71:756–771, 2023.

[147] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. Compressive channel estimation and tracking for large arrays in mm-wave picocells. *IEEE Journal of Selected Topics in Signal Processing*, 10(3):514–527, 2016.

[148] Inceptive Mind. Researchers introduce new wireless system for greater 5G access. https://www.inceptivemind.com/researchers-introduce-new-wireless-system-greater-5g-access/30839/, May 2023.

[149] Noriyuki Miura, Tatsuya Machida, Kohei Matsuda, Makoto Nagata, Shoei Nashimoto, and Daisuke Suzuki. A low-cost replica-based distance-spoofing attack on mmwave fmcw radar. In *Proceedings of the 3rd ACM Workshop on Attacks and Solutions in Hardware Security Workshop*, pages 95–100, 2019.

[150] Imon Mondal and Nagendra Krishnapura. A 2-ghz bandwidth, 0.25–1.7 ns true-time-delay element using a variable-order all-pass filter architecture in 0.13μm cmos. *IEEE Journal of Solid-State Circuits*, 52(8):2180–2193, 2017.

[151] Thomas Moon, Jounsup Park, and Seungmo Kim. Bluefmcw: Random frequency hopping radar for mitigation of interference and spoofing. *EURASIP Journal on Advances in Signal Processing*, 2022(1):1–17, 2022.

[152] Yasser Nabil, Hesham ElSawy, Suhail Al-Dharrab, Hassan Mostafa, and Hussein Attia. Beamwidth design tradeoffs in radar-aided millimeter-wave cellular networks: A stochastic geometry approach. *IEEE Access*, 2024.

[153] Aravind Nagulu, Aditya Gaonkar, Sohail Ahasan, Sasank Garikapati, Tingjun Chen, Gil Zussman, and Harish Krishnaswamy. A full-duplex receiver with true-time-delay cancelers based on switched-capacitor-networks operating beyond the delay–bandwidth limit. *IEEE Journal of Solid-State Circuits*, 56(5):1398–1411, 2021.

[154] Arvind Narayanan, Jason Carpenter, Eman Ramadan, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. A first measurement study of commercial mmwave 5G performance on smartphones. *arXiv preprint arXiv:1909.07532*, 2019.

[155] Shoei Nashimoto, Daisuke Suzuki, Noriyuki Miura, Tatsuya Machida, Kohei Matsuda, and Makoto Nagata. Low-cost distance-spoofing attack on fmcw radar and its feasibility study on countermeasure. *Journal of Cryptographic Engineering*, 11(3):289–298, 2021.

[156] Dang Qua Nguyen and Taejoon Kim. Joint delay-phase precoding under true-time delay constraints in wideband sub-thz hybrid massive mimo systems. *IEEE Transactions on Communications*, 2024.

[157] Jiali Nie, Quan Zhou, Junsheng Mu, and Xiaojun Jing. Vision and radar multimodal aided beam prediction: Facilitating metaverse development. In *Proceedings of the 2nd Workshop on Integrated Sensing and Communications for Metaverse*, pages 13–18, 2023.

[158] Thomas Nitsche, Carlos Cordeiro, Adriana B Flores, Edward W Knightly, Eldad Perahia, and Joerg C Widmer. Ieee 802.11 ad: directional 60 ghz communication for multi-gigabit-

per-second wi-fi. *IEEE Communications Magazine*, 52(12):132–141, 2014.

[159] Thomas Nitsche, Adriana B Flores, Edward W Knightly, and Joerg Widmer. Steering with eyes closed: mm-wave beam steering without in-band measurement. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 2416–2424. IEEE, 2015.

[160] NSF. Spectrum and Wireless Innovation enabled by Future Technologies - Satellite-Terrestrial Coexistence (SWIFT-SAT). https://www.nsf.gov/pubs/2023/nsf23567/nsf23567.htm, June 2023.

[161] OAI. Open Air Interface 5G Radio Access Network Project. https://openairinterface.org/oai-5g-ran-project/, 2023.

[162] Garret Okamoto. Jammer nulling via low complexity blind beamforming algorithm. In *2007 IEEE Antennas and Propagation Society International Symposium*, pages 25–28. IEEE, 2007.

[163] Mihai Ordean and Flavio D Garcia. Millimeter-wave automotive radar spoofing. *arXiv preprint arXiv:2205.06567*, 2022.

[164] Sophocles J Orfanidis. Electromagnetic waves and antennas. 2002.

[165] Joan Palacios. Adaptive Codebook Optimization for Beam Training on Off-the-Shelf IEEE 802.11ad Devices. In *Proceedings of the 24rd Annual International Conference on Mobile Computing and Networking*. ACM, 2018.

[166] Joan Palacios, Paolo Casari, Hany Assasa, and Joerg Widmer. Leap: Location estimation and predictive handover with consumer-grade mmwave devices. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 2377–2385. IEEE, 2019.

[167] Junhyeong Park, Seungwoon Park, Do-Hoon Kim, and Seong-Ook Park. Leakage mitigation in heterodyne fmcw radar for small drone detection with stationary point concentration technique. *IEEE Transactions on Microwave Theory and Techniques*, 67(3):1221–1232, 2019.

[168] Andre Pearce, J Andrew Zhang, Richard Xu, and Kai Wu. Multi-object tracking with mmwave radar: A review. *Electronics*, 12(2):308, 2023.

[169] Ioannis Pefkianakis and Kyu-Han Kim. Accurate 3d localization for 60 GHz networks. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, pages 120–131, 2018.

[170] Guillermo Pocovi, Hamidreza Shariatmadari, Gilberto Berardinelli, Klaus Pedersen, Jens Steiner, and Zexian Li. Achieving ultra-reliable low-latency communications: Challenges

and envisioned system enhancements. *IEEE Network*, 32(2):8–15, 2018.

[171] Richard Poisel. *Modern communications jamming principles and techniques*. Artech house, 2011.

[172] Michele Polese, Leonardo Bonati, Salvatore D'Oro, Stefano Basagni, and Tommaso Melodia. Understanding o-ran: Architecture, interfaces, algorithms, security, and research challenges. *IEEE Communications Surveys & Tutorials*, 2023.

[173] Michele Polese et al. Millimetera: Toward a large-scale open-source mmwave and terahertz experimental testbed. In *ACM Workshop on Millimeter-wave Networks and Sensing Systems*, pages 27–32, 2019.

[174] pSemi. pSemi 28 GHz phased arrays. https://www.psemi.com/.

[175] Feng Qian, Bo Han, Qingyang Xiao, and Vijay Gopalakrishnan. Flare: Practical viewport-adaptive 360-degree video streaming for mobile devices. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 99–114, 2018.

[176] Jian Qiao, Xuemin Shen, Jon W Mark, and Yejun He. Mac-layer concurrent beamforming protocol for indoor millimeter-wave networks. *IEEE Transactions on Vehicular Technology*, 64(1):327–338, 2014.

[177] Tian Qiu, Ish Kumar Jain, Raini Wu, Dinesh Bharadia, and Pamela Cosman. Delivering 360-degree video with viewport-adaptive truncation. In *2022 25th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, pages 7–11. IEEE, 2022.

[178] Qualcomm. Linux wil6210 driver. https://github.com/torvalds/linux/tree/master/drivers/net/wireless/ath/wil6210, August 2017.

[179] ABI Rsearch Qualcomm. Augmented and virtual reality: the first wave of 5g killer apps. https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/cr-qcom-140.pdf, 2017.

[180] Theodore S Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N Wong, Jocelyn K Schulz, Mathew Samimi, and Felix Gutierrez. Millimeter wave mobile communications for 5g cellular: It will work! *IEEE access*, 1:335–349, 2013.

[181] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. Noncoherent mmwave path tracking. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications*, pages 13–18, 2017.

[182] Vishnu V Ratnam, Jianhua Mo, Ahmad Alammouri, Boon Loong Ng, Jianzhong Zhang, and Andreas F Molisch. Joint phase-time arrays: A paradigm for frequency-dependent analog beamforming in 6g. *IEEE Access*, 10:73364–73377, 2022.

[183] Dipankar Raychaudhuri, Ivan Seskar, Gil Zussman, Thanasis Korakis, Dan Kilper, Tingjun Chen, Jakub Kolodziejski, Michael Sherman, Zoran Kostic, Xiaoxiong Gu, et al. Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–13, 2020.

[184] Grecia Romero, Virginie Deniau, and Olivier Stienne. LTE Physical layer vulnerability test to different types of jamming signals. In *2019 International Symposium on Electromagnetic Compatibility-EMC EUROPE*, pages 1138–1143. IEEE, 2019.

[185] Ruth Rotman, Moshe Tur, and Lior Yaron. True time delay in phased arrays. *Proceedings of the IEEE*, 104(3):504–518, 2016.

[186] Bodhisatwa Sadhu, Arun Paidimarri, Mark Ferriss, Mark Yeck, Xiaoxiong Gu, and Alberto Valdes-Garcia. A 128-element dual-polarized software-defined phased array radio for mm-wave 5G experimentation. In *Proceedings of the 2nd ACM Workshop on Millimeter Wave Networks and Sensing Systems*, pages 21–25, 2018.

[187] Hooman Saeidi, Suresh Venkatesh, Xuyang Lu, and Kaushik Sengupta. Thz prism: One-shot simultaneous localization of multiple wireless nodes with leaky-wave thz antennas and transceivers in cmos. *IEEE Journal of Solid-State Circuits*, 56(12):3840–3854, 2021.

[188] Swetank Kumar Saha, Yasaman Ghasempour, et al. X60: A programmable testbed for wideband 60 GHz WLANS with phased arrays. *Computer Communications*, 133:77–88, 2019.

[189] Alphan Şahin, Mihail L Sichitiu, and İsmail Guvenç. A millimeter-wave software-defined radio for wireless experimentation. *arXiv preprint arXiv:2302.08444*, 2023.

[190] Samsung. Samsung 5G phased array. http://bit.ly/2oNSbnx, 2020.

[191] Akbar M Sayeed and Vasanthan Raghavan. Maximizing MIMO capacity in sparse multipath with reconfigurable antenna arrays. *IEEE Journal of Selected Topics in Signal Processing*, 1(1):156–166, 2007.

[192] Robert Schmidt, Mikel Irazabal, and Navid Nikaein. Flexric: An sdk for next-generation sd-rans. In *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '21, page 411–425, New York, NY, USA, 2021. ACM.

[193] S Sekretarov, D Vavriv, V Vinogradov, A Kravtsov, Y Bulakh, and V Zolotarev. Frequency scanning x-band antenna for 3d radar systems. In *2023 53rd European Microwave Conference (EuMC)*, pages 959–962. IEEE, 2023.

[194] Sina Shaham, Ming Ding, Matthew Kokshoorn, Zihuai Lin, Shuping Dang, and Rana Abbas. Fast channel estimation and beam tracking for millimeter wave vehicular communications. *IEEE Access*, 7:141104–141118, 2019.

[195] Sihua Shao, Hanbin Zhang, Dimitrios Koutsonikolas, and Abdallah Khreishah. Two-dimensional reduction of beam training overhead in crowded 802.11 ad based networks. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 680–685. IEEE, 2018.

[196] Joseph A Shaw. Radiometry and the friis transmission equation. *American journal of physics*, 81(1):33–37, 2013.

[197] Li-Hsiang Shen and Kai-Ten Feng. Mobility-aware subband and beam resource allocation schemes for millimeter wave wireless networks. *IEEE Transactions on Vehicular Technology*, 69(10):11893–11908, 2020.

[198] Min Soo Sim, Yeon-Geun Lim, Sang Hyun Park, Linglong Dai, and Chan-Byoung Chae. Deep learning-based mmwave beam selection for 5g nr/6g with sub-6 ghz channel information: Algorithms and prototype validation. *IEEE Access*, 8:51634–51646, 2020.

[199] Dmitrii Solomitckii, Margarita Gapeyenko, Vasilii Semkin, Sergey Andreev, and Yevgeni Koucheryavy. Technologies for efficient amateur drone detection in 5g millimeter-wave cellular infrastructure. *IEEE Communications Magazine*, 56(1):43–50, 2018.

[200] IEEE Spectrum. Starlink and other LEO constellations face a new set of security risks. https://spectrum.ieee.org/satellite-jamming, May 2023.

[201] srsRAN. srsran: Open source 4g/5g software radio access network. https://github.com/srsran/srsRAN, 2023.

[202] Matthias Steinhauer, Hans-Oliver Ruoß, Hans Irion, and Wolfgang Menzel. Millimeter-wave-radar sensor based on a transceiver array for automotive applications. *IEEE transactions on microwave theory and techniques*, 56(2):261–269, 2008.

[203] Shu Sun and Theodore S Rappaport. Multi-beam antenna combining for 28 GHz cellular link improvement in urban environments. In *2013 IEEE Global Communications Conference (GLOBECOM)*, pages 3754–3759. IEEE, 2013.

[204] Shunqiao Sun, Athina P Petropulu, and H Vincent Poor. Mimo radar for advanced driver-assistance systems and autonomous driving: Advantages and challenges. *IEEE Signal*

Processing Magazine, 37(4):98–117, 2020.

[205] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. WiFi-assisted 60 GHz wireless networks. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 28–41. ACM, 2017.

[206] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. Towards scalable and ubiquitous millimeter-wave wireless networks. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 257–271. ACM, 2018.

[207] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 60 GHz indoor networking through flexible beams: A link-level profiling. In *ACM SIGMETRICS Performance Evaluation Review*, volume 43, pages 71–84. ACM, 2015.

[208] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. Beamspy: enabling robust 60 GHz links under blockage. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 193–206, 2016.

[209] Zenon Szczepaniak and Waldemar Susek. Microwave delay lines for analogue signal correlation. In *Radioelectronic Systems Conference 2019*, volume 11442, pages 505–516. SPIE, 2020.

[210] Jingbo Tan and Linglong Dai. Delay-phase precoding for thz massive mimo with beam split. In *2019 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6. IEEE, 2019.

[211] Jingbo Tan and Linglong Dai. Wideband beam tracking in thz massive mimo systems. *IEEE Journal on Selected Areas in Communications*, 39(6):1693–1710, 2021.

[212] Xin Tan, Zhi Sun, Dimitrios Koutsonikolas, and Josep M Jornet. Enabling indoor mobile millimeter-wave networks based on smart reflect-arrays. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 270–278. IEEE, 2018.

[213] Inc Telecom Infra Project. Analysis of 28GHz and 60GHz channel measurements in an indoor environment. https://telecominfraproject.com/, August 2019.

[214] Jose A Torres, Richard M Davis, J David R Kramer, and Ronald L Fante. Efficient wideband jammer nulling when using stretch processing. *IEEE Transactions on Aerospace and Electronic Systems*, 36(4):1167–1178, 2000.

[215] Y Ming Tsang, Ada SY Poon, and Sateesh Addepalli. Coding the beams: Improving beamforming training in mmwave communication system. In *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, pages 1–6. IEEE, 2011.

[216] David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.

[217] Michael Ulrich, Sascha Braun, Daniel Köhler, Daniel Niederlöhner, Florian Faion, Claudius Gläser, and Holger Blume. Improved orientation estimation and detection with hybrid object detection networks for automotive radar. *arXiv preprint arXiv:2205.02111*, 2022.

[218] Vutha Va, Xinchen Zhang, and Robert W Heath. Beam switching for millimeter wave communication to support high speed trains. In *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, pages 1–5. IEEE, 2015.

[219] Harry L Van Trees. *Optimum array processing: Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 2002.

[220] Rohith Reddy Vennam, Ish Kumar Jain, Kshitiz Bansal, Joshua Orozco, Puja Shukla, Aanjhan Ranganathan, and Dinesh Bharadia. mmspoof: Resilient spoofing of automotive millimeter-wave radars using reflect array. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 1807–1821. IEEE, 2023.

[221] Julia Vinogradova, Emil Björnson, and Erik G Larsson. Detection and mitigation of jamming attacks in massive MIMO systems using random matrix theory. In *2016 IEEE 17th International Workshop on SPAWC*, pages 1–5. IEEE, 2016.

[222] Aditya Wadaskar, Veljko Boljanovic, Han Yan, and Danijela Cabric. 3D Rainbow Beam Design for Fast Beam Training with True-Time-Delay Arrays in Wideband Millimeter-Wave Systems. In *2021 55th Asilomar Conference on Signals, Systems, and Computers*, pages 85–92. IEEE, 2021.

[223] Aditya Wadaskar, Ding Zhao, Ibrahim Pehlivan, and Danijela Cabric. Structured two-stage true-time-delay array code book design for multi-user data communication. In *GLOBECOM 2023-2023 IEEE Global Communications Conference*, pages 3378–3384. IEEE, 2023.

[224] Guolong Wang, Wei Hong, Hui Zhang, and Lixia Cao. Investigations on anti-jamming method for 77ghz automotive millimeter-wave radar. In *2016 IEEE International Conference on Ubiquitous Wireless Broadband (ICUWB)*, pages 1–3. IEEE, 2016.

[225] Peilan Wang, Jun Fang, Xiaojun Yuan, Zhi Chen, and Hongbin Li. Intelligent reflecting surface-assisted millimeter wave communications: Joint active and passive precoding design. *IEEE Transactions on Vehicular Technology*, 69(12):14960–14973, 2020.

[226] Song Wang, Jingqi Huang, and Xinyu Zhang. Demystifying millimeter-wave V2X: Towards robust and efficient directional connectivity under high mobility. In *Proceedings*

of the 26th Annual International Conference on Mobile Computing and Networking, pages 1–14, 2020.

[227] Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. *X-Array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays*. Association for Computing Machinery, New York, NY, USA, 2020.

[228] Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. X-array: approximating omnidirectional millimeter-wave coverage using an array of phased arrays. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.

[229] Teng Wei and Xinyu Zhang. Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, pages 42–55. ACM, 2017.

[230] Teng Wei, Anfu Zhou, and Xinyu Zhang. Facilitating robust 60 GHz network deployment by sensing ambient reflectors. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 213–226, 2017.

[231] Timothy Woodford, Xinyu Zhang, Eugene Chai, Karthikeyan Sundaresan, and Amir Khojastepour. Spacebeam: Lidar-driven one-shot mmwave beam management. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 389–401, 2021.

[232] Xilinx. Cmod A7-35T. https://www.xilinx.com/products/boards-and-kits/1-f3zdsm.html, Jul 2021.

[233] Hao Xu, Vikas Kukshya, and Theodore S Rappaport. Spatial and temporal characteristics of 60-ghz indoor channels. *IEEE Journal on selected areas in communications*, 20(3):620–630, 2002.

[234] Shengbo Xu, Weiyang Xu, Cunhua Pan, and Maged Elkashlan. Detection of jamming attack in non-coherent massive simo systems. *IEEE Transactions on Information Forensics and Security*, 14(9):2387–2399, 2019.

[235] Chen Yan, Wenyuan Xu, and Jianhao Liu. Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle. *Def Con*, 24(8):109, 2016.

[236] Han Yan, Veljko Boljanovic, and Danijela Cabric. Wideband millimeter-wave beam training with true-time-delay array architecture. In *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pages 1447–1452. IEEE, 2019.

[237] Qiben Yan, Huacheng Zeng, Tingting Jiang, Ming Li, Wenjing Lou, and Y Thomas

Hou. MIMO-based jamming resilient communication in wireless networks. In *IEEE INFOCOM*, pages 2697–2706. IEEE, 2014.

[238] Ibrar Yaqoob, Latif U Khan, SM Ahsan Kazmi, Muhammad Imran, Nadra Guizani, and Choong Seon Hong. Autonomous driving cars in smart cities: Recent advances, requirements, and challenges. *IEEE Network*, 34(1):174–181, 2019.

[239] E Yeh, Junil Choi, N Prelcic, C Bhat, and Robert W Heath Jr. Security in automotive radar and vehicular networks. *submitted to Microwave Journal*, 2016.

[240] Ozlem Yildiz, Ahmad AlAmmouri, Jianhua Mo, Younghan Nam, Elza Erkip, et al. 3d beamforming through joint phase-time arrays. *arXiv preprint arXiv:2401.00819*, 2024.

[241] Ding Zhang, Mihir Garude, and Parth H Pathak. mmchoir: Exploiting joint transmissions for reliable 60GHz mmwave wlans. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 251–260. ACM, 2018.

[242] J Andrew Zhang, Xiaojing Huang, Y Jay Guo, Jinhong Yuan, and Robert W Heath. Multibeam for joint communication and radar sensing using steerable analog antenna arrays. *IEEE Transactions on Vehicular Technology*, 68(1):671–685, 2018.

[243] Jialiang Zhang, Xinyu Zhang, Pushkar Kulkarni, and Parameswaran Ramanathan. Openmili: a 60 GHz software radio platform with a reconfigurable phased-array antenna. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 162–175, 2016.

[244] Jian A Zhang, Xiaojing Huang, Val Dyadyuk, and Y Jay Guo. Massive hybrid antenna array for millimeter-wave cellular communications. *IEEE Wireless Communications*, 22(1):79–87, 2015.

[245] Wenbin Zhang, Yuming Wei, Shaochuan Wu, Weixiao Meng, and Wei Xiang. Joint beam and resource allocation in 5g mmwave small cell systems. *IEEE Transactions on Vehicular Technology*, 68(10):10272–10277, 2019.

[246] Ding Zhao, Ibrahim Pehlivan, Aditya Wadaskar, and Danijela Cabric. Fast frequency-direction mapping design for data communication with true-time-delay array architecture. In *2024 International Conference on Computing, Networking and Communications (ICNC)*, pages 1071–1076. IEEE, 2024.

[247] Man Zhao, Danpu Liu, and Feng Yu. Improving the robustness of 60 GHz indoor connectivity by deployment of mirrors. In *2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 188–193. IEEE, 2018.

[248] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. M-cube: A millimeter-wave massive mimo software radio. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.

[249] Kai Zheng, Kun Qian, Timothy Woodford, and Xinyu Zhang. Neuroradar: A neuromorphic radar sensor for low-power iot systems. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*, pages 223–236, 2023.

[250] Kai Zheng, Wuqiong Zhao, Timothy Woodford, Renjie Zhao, Xinyu Zhang, and Yingbo Hua. Enhancing mmwave radar sensing using a phased-mimo architecture. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*, pages 56–69, 2024.

[251] Anfu Zhou, Leilei Wu, Shaoqing Xu, Huadong Ma, Teng Wei, and Xinyu Zhang. Following the shadow: Agile 3-D beam-steering for 60 GHz wireless networks. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 2375–2383. IEEE, 2018.

[252] Anfu Zhou, Shaoqing Xu, Song Wang, Jingqi Huang, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. Robot navigation in radio beam space: Leveraging robotic intelligence for seamless mmwave network coverage. In *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 161–170. ACM, 2019.

[253] Anfu Zhou, Shaoqing Xu, Song Wang, Jingqi Huang, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. Robotic millimeter-wave wireless networks. *IEEE/ACM Transactions on Networking*, 28(4):1534–1549, 2020.

[254] Anfu Zhou, Shaoyuan Yang, Yi Yang, Yuhang Fan, and Huadong Ma. Autonomous environment mapping using commodity millimeter-wave network device. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 1126–1134. IEEE, 2019.

[255] Anfu Zhou, Xinyu Zhang, and Huadong Ma. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.

[256] Liang Zhou and Yoji Ohashi. Efficient codebook-based MIMO beamforming for millimeter-wave WLANs. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, pages 1885–1889. IEEE, 2012.

[257] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y Zhao, and Haitao Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. *ACM SIGCOMM Computer Communication Review*, 42(4):443–454, 2012.

[258] Dalin Zhu, Junil Choi, Qian Cheng, Weimin Xiao, and Robert W Heath. High-resolution angle tracking for mobile wideband millimeter-wave systems with antenna array calibration. *IEEE Transactions on Wireless Communications*, 17(11):7173–7189, 2018.

[259] Yibo Zhu, Zengbin Zhang, Zhinus Marzi, Chris Nelson, Upamanyu Madhow, Ben Y Zhao, and Haitao Zheng. Demystifying 60GHz outdoor picocells. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 5–16. ACM, 2014.

[260] Frederik Jonathan Zumegen, Ish Kumar Jain, and Dinesh Bharadia. Beamarmor: Seamless anti-jamming in 5g cellular networks with mimo null-steering. In *Proceedings of the 25th International Workshop on Mobile Computing Systems and Applications*, pages 121–126, 2024.