

Designing Sensing-Augmented Wireless Systems for the Internet of Things

by

Teng Wei

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

Doctor of Philosophy

(Electrical and Computer Engineering)

at the

UNIVERSITY OF WISCONSIN-MADISON

2018

Date of final oral examination: 07/23/2018

The dissertation is approved by the following members of the Final Oral Committee:

Xinyu Zhang, Associate Professor, Electrical and Computer Engineering
Parmesh Ramanathan, Professor, Electrical and Computer Engineering
Yu Hen Hu, Professor, Electrical and Computer Engineering
Kassem Fawaz, Assistant Professor, Electrical and Computer Engineering
Suman Banerjee, Professor, Computer Science

Designing Sensing-Augmented Wireless Systems for the Internet of Things

by Teng Wei

Submitted to the Department of Electrical and Computer Engineering
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in 2018

ABSTRACT

Wireless systems provide the most critical infrastructure for Internet-of-things (IoT) applications. Despite the fast-evolving technology, many new IoT applications, such as virtual reality and self-driving, are demanding for better connectivity and new passive sensing techniques to support the more intelligent functionalities. Although the forthcoming Millimeter-wave (mmWave) wireless systems can provide much higher throughput, its performance is sensitive to the mobility and hardly applicable to complicated application scenarios. In this thesis, I contend that many of such looming challenges can be resolved through an integrated design of sensing and communication capabilities. I propose a general design principle, called sensing-augmented wireless systems, which will enable new IoT sensing applications while resolving critical networking issues.

I will first introduce the design of two *sensing-oriented wireless systems* which target the sensing barrier. They achieve super-resolution passive sensing by re-purposing communication devices. The first system, Tagyro, is a 3D orientation tracking system for batteryless objects. It turns Radio-Frequency Identification (RFID) tags into passive sensors and transforms the runtime phase into real-time orientation estimation, enabling intelligent services for the smart home and supply chain. The second system, mTrack, is a fine-grained location tracking system for writing objects. It leverages the highly directional and steerable mmWave radios to pinpoint objects on any desktop

surfaces, realizing a virtual trackpad for touchless interaction.

I will further introduce two *sensing-assisted wireless systems*, targeting the networking demand. They exploit sensing information from external to overcome connectivity challenge in mmWave networks. Pia, a robust mmWave networking technique, can meet the needs of demanding Gbps applications such as wireless Virtual Reality (VR). The design harnesses the pose information on mobile devices to switch the link proactively and dynamically. Another system is, E-Mi, which can sense and reconstruct the dominant reflectors in an environment through 60 GHz radios' eyes, realizing efficacious Access Point (AP) deployment.

Thesis Supervisor: Xinyu Zhang

Title: Associate Professor of Electrical and Computer Engineering

To my parents

A C K N O W L E D G M E N T S

First and foremost, I would like to express my sincere gratitude to my advisor, Xinyu Zhang, for his consistent support and guidance for my years at UW-Madison. He set the best model example for me to follow in both academic research and personal life, let me understand how to work efficiently and effectively on critical problems, and light me through dilemmas and difficulties numerous times. He invested much time and efforts in every research projects I worked on, and guide me through every detail. I will never forget how many days and nights we spent together, brain-storming the hard and exciting directions, delving into the hard problems, and witnessing the breakthrough of outcomes. Without his keen suggestions and insights, I am not able to learn so quickly in my domain of knowledge and develop many technical and analytical skills in research. The skill set of critical thinking and capability of conducting meaningful research that he has taught will genuinely benefit my whole life.

I would also like to thank my thesis committee members Parmesh Ramanathan and Yu Hen Hu for their time and supports. Their early feedback has significantly helped to form and shape this thesis. Many technical designs in this dissertation have benefited from discussions with them.

I am sincerely grateful for Dinesh Bharadia's assistance. During the last year of my Ph.D. study, I worked as a visiting scholar at UCSD. Dinesh kindly provides me solid supports and guidance in both research and job career. I would also like to thank Vincent Leung for his generosity in sharing his hardware lab for my study. Getting access to the high-end testing equipment has significantly expanded my technical skills and accelerated my research progress, making many far-fetching ideas to become accessible.

I would also like to thank the Electrical and Computer Engineering staff, especially Lori Burrow. She helped to design the lab environment and organize deployment of hardware equipment. Her generous work has offered us a friendly and efficient working environment. I need to thank Michael Radloff and Hannah Roberg, who have been helping me to address academic confusions, from the first day at UW-Madison to the time I am submitting the thesis. I also own many thanks to Wyn Hughes at UCSD. She offered valuable help in ordering lab components and printing documents.

I would not have been able to explore many fascinating research ideas without the collaboration with my collaborators. During the first two years of my Ph.D. study, I had a fantastic time working with Sanjib Sur on multiple projects and participating in the Qualcomm competition. Thanks to our different backgrounds and skill sets, we together formed a strong team to target on challenging problems and move the progress forward quickly. I have learned a lot of good merits as well as technical knowledge from him. Without his harsh but helpful criticisms and suggestions, I would not have been able to publish many works at top conferences. It is my fortune to have such a great teammate.

I would also like to thank Shu Wang, who was a master student in my group and participated in my research project. He gave me the first and wonderful mentorship experience. Our collaboration on the vibrometry project has produced one of the top-rated papers I ever got. How much I wish to have more time to collaborate with him. Besides, working with Anfu Zhou on the 60 GHz system project was the most pleasant journey of my research. His high efficiency and hard working have accelerated the experiment evaluation like riding a rocket. I also have to mention Kun Qian, a visitor graduate in my lab. He has been the most clever student I ever met. It is my honor to have him participated in the same research project and hope we can work together again in the future.

The labmates at UW-Madison and UCSD together created a pleasant and comfortable working environment. Without their help and support, I won't be able to solve many bothering problems and to have so much fun. I would like to thank Jialiang Zhang, Chi Zhang, Chuhang Gao, Xiufeng Xie, Meng Tong, Shilin Zhu, Arunkumar Ravichan and Rana Hegazy. They have made my graduate school years exciting and enjoyable. I loved working with them and kept grateful for their great company and support.

Finally, yet most importantly, I would like to thank my family. My mom and dad have continuously provided faithful love, meticulous care, and unconditional support for all the years. My beloved girlfriend Yajin has always been inspiring and encouraging. I owe everything to their accompany. This thesis is dedicated to them.

This thesis produced from projects supported in part by the NSF under Grant CNS-1318292 CNS-1350039 CNS-1343363 CNS-1404613 CNS-1518728 CNS-1617321 and Google Faculty Research Award, HP Labs gift fund, Sony Research Award.

Previously Published Materials

Chapter 2 revises a previous publication [186]: Teng Wei and Xinyu Zhang, “*Gyro in the Air: Tracking 3D Orientation of Batteryless Internet-of-Things*,” In ACM International Conference on Mobile Computing and Networking (MobiCom), 2016.

Chapter 3 revises a previous publication [185]: Teng Wei and Xinyu Zhang, “*mTrack: High-Precision Passive Tracking Using Millimeter Wave Radios*,” In ACM International Conference on Mobile Computing and Networking (MobiCom), 2015.

Chapter 4 revises a previous publication [187]: Teng Wei and Xinyu Zhang, “*Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support*,” In ACM International Conference on Mobile Computing and Networking (MobiCom), 2017.

Chapter 5 revises a previous publication [188]: Teng Wei, Anfu Zhou, and Xinyu Zhang, “*Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors*,” In USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2017.

TABLE OF CONTENTS

Abstract	i
Dedication	iii
Acknowledgments	iv
Previously Published Materials	vii
List of Figures	xi
List of Tables	xiv
List of Algorithms	xv
List of Abbreviations	xvi
 Chapter	
1 Introduction	1
1.1 Ubiquitous Wireless Sensing for Batteryless IoT	4
1.2 Robust mmWave Networks for Connected IoT	7
1.3 Organization	9
 I Ubiquitous Wireless Sensing for Batteryless IoT	 10
2 Batteryless 3D Orientation Tracking using RFID Tags	11
2.1 Related Work	14
2.2 Preliminaries	17
2.2.1 Phase Model for an RFID System	17
2.2.2 Phase Difference of Arrival of a Tag Array	18
2.3 Tagyro Operation	20
2.4 Design	21
2.4.1 Tracking through Orientation Spectrum	21
2.4.2 Impact of Tag Radiation Pattern	23
2.4.3 Combating Tag Coupling Effect	28
2.4.4 Toward 3-DoF Orientation Tracking	33
2.4.5 Calibrating Frequency-Hopping Readers	38

2.4.6 Dealing with Asynchronous Phase Reading	39
2.5 Implementation	40
2.6 Evaluation	41
2.6.1 Micro Benchmarks	42
2.6.2 System Level Tests	46
2.6.3 Case Study	50
2.7 Discussion	52
3 Fine-grained Passive Motion Tracking Using mmWave Radios	53
3.1 Related Work	56
3.2 Understanding mmWave Passive Tracking	58
3.2.1 Reflection/Diffusion of Signals by Objects	59
3.2.2 Phase Variation Enables Fine-Grained Motion Tracking	60
3.2.3 Beam Steering Enables Localizing Absolute Position	62
3.3 An Overview of mTrack	63
3.4 Phase Tracking in mTrack	65
3.4.1 Basic Successive Tracking Algorithm	65
3.4.2 Tracking under Background Reflection	67
3.4.3 Phase Counting and Regeneration	72
3.5 Anchor Point Acquisition	76
3.6 Touch Event Detection	78
3.7 Implementation and Evaluation	80
3.7.1 Micro Benchmark	82
3.7.2 Performance on a Trackpad	88
3.7.3 Application of mTrack	91
3.8 Discussion	92
II Robust mmWave Networks for Connected IoT	93
4 Pose Information Assisted Robust 60 GHz Networks	94
4.1 Related Work	97
4.2 Background and Overview	100
4.2.1 Limited FoV of the 60 GHz Radio	100
4.2.2 Pia Operations	102
4.3 Pia Design	103
4.3.1 AP Selection	103
4.3.2 AP-Pose Sensing	108
4.4 Pose-Assisted Interference Management	111
4.4.1 Improving the Spatial Sharing Opportunity	112
4.4.2 Dealing With the Environmental Reflectors	118
4.4.3 Calibrating the Client's Pose Offset	122
4.5 Implementation	124
4.6 Evaluation	126
4.6.1 Micro Benchmarks	127

4.6.2 System Level Tests	133
4.7 Discussion	136
5 Environment Sensing for Optimizing mmWave Network Deployment	137
5.1 Introduction	137
5.2 Related Work	140
5.3 E-Mi: An Overview	142
5.4 Multipath Resolution Framework (MRF)	143
5.4.1 Estimate Path Angles Using Phased Arrays	144
5.4.2 Virtual Beamforming: Match Path Angles	146
5.4.3 Multi-Tone Ranging: Estimate Path Length	148
5.5 Dominant Reflector Reconstruction	150
5.5.1 Locating Reflecting Points in Environment	150
5.5.2 Reconstructing Dominant Reflector Layout and Reflectivity . . .	153
5.6 Parametric Ray-tracing: Predict Link Performance	156
5.7 Implementation	157
5.8 Experimental Validation	159
5.8.1 Effectiveness of Multipath Resolution	159
5.8.2 Effectiveness of Dominant Reflector Reconstruction	162
5.9 Case Study of E-Mi	167
5.9.1 Environment-Aware 60 GHz AP Deployment	167
5.9.2 Experimental Verification	169
5.10 Discussion	171
6 Conclusion	172
6.1 Lessons Learned	173
6.2 Future Work	175
Bibliography	179

LIST OF FIGURES

1.1	Thesis Contributions	3
2.1	Applications of Passive Orientation Sensing in IoT	12
2.2	Examples of Linear and Square Tag Arrays	12
2.3	Tagyro's Coordinate System	18
2.4	Conversion From PDoA to Orientation	19
2.5	Architecture of Tagyro	20
2.6	Example of Orientation Spectrum	22
2.7	Radiation Pattern of a RFID Tag	23
2.8	Phase and RSS Measurement of a Single RFID Tag	24
2.9	Setup of the Measurement Study	26
2.10	PDoA Measurement of a Tag Array	27
2.11	PDoA Simulation of a Tag Array	27
2.12	PDoA Deviation Over Tag Separation Distance	28
2.13	CDF of Residual PDoA Deviation	29
2.14	Illustration of Array Layout Sensing Algorithm	31
2.15	Setup of 3-DoF Orientation Sensing	34
2.16	Processing Blocks of 3-DoF Orientation Sensing	35
2.17	Phase Jumping Caused by Frequency Hopping	38
2.18	Phase Over Distance for a Fixed Frequency	38
2.19	Tagyro Implementation and Software Interface	40
2.20	Effective Layout vs Physical Layout	42
2.21	Orientation Spectrum for 4-Tag Array	42
2.22	Effectiveness of Array Layout Sensing	43
2.23	Accuracy of Selecting the Array in Valid Reading Zone	43
2.24	Effectiveness of Frequency Hopping Calibration	44
2.25	Residual Phase Error Caused by Asynchronous Phase Reading	44
2.26	Orientation Accuracy vs DoF	46
2.27	Orientation Estimation vs Ground-Truth	46
2.28	Orientation Error Over Tag Size	47
2.29	Orientation Error Under Human Activity	47
2.30	Orientation Error Over Distance	48
2.31	Orientation Error Under Blockage	48
2.32	Latency of Rotation Response	49
2.33	Tagyro Case Study Setup	49
2.34	Case Study of a 12-Side Dice	50

2.35	Case Study of Smart Home Applications	50
3.1	Setup of mTrack’s Feasibility Test	59
3.2	Feasibility Test of Target Detection	60
3.3	Feasibility Test of Target Tracking	61
3.4	Feasibility Test of Target Locating	62
3.5	mTrack’s System Setup	63
3.6	mTrack’s System Components	63
3.7	Model of the Multipath Reflection	67
3.8	Impact of Background Reflection on Phase Shift	68
3.9	Illustration of the Background Removal Algorithm	71
3.10	Illustration of the Phase Counting and Regeneration (PCR) Algorithm	74
3.11	Effectiveness of Spline Interpolation	77
3.12	Effectiveness of Anchor Point Acquisition (APA) Algorithm	77
3.13	Example Snapshot of Three Type Events	79
3.14	Hardware Architecture of mTrack	80
3.15	mTrack mmWave Testbed	81
3.16	Tracking Without Background Reflection	83
3.17	Tracking With Background Reflection	83
3.18	Tracking Error of Different Methods	83
3.19	Tracking Error Under Different Background Strength	84
3.20	Tracking Error Under Human Activity	84
3.21	Tracking Error Under Different Levels of Phase Noise	85
3.22	APA Performance With Spline Interpolation	86
3.23	Effectiveness of Background Subtraction	87
3.24	Locating Error Under Human Activity	87
3.25	Tracking Error With Opportunistic Calibration	88
3.26	Locating Error Under Different Target Materials	88
3.27	Tracking Error Under Different Target Materials	89
3.28	Map of APA Error	89
3.29	Map of Phase-Tracking Error	90
3.30	Touch Detection Accuracy for Different Users	90
3.31	Example Letter and Word by mTrack	91
3.32	Character and Word Recognition Accuracy	91
4.1	Multi-AP Network Architecture	95
4.2	Limited FoV of Commercial 60 GHz Radios	100
4.3	Performance Impact by the Limited FoV	102
4.4	Pia’s Modules and Operation Workflow	103
4.5	5-DoF Pose of the Phased Array	104
4.6	Link Status Prediction Based on the Poses	104
4.7	Link Status Prediction Error in a Sample Walking	106
4.8	AP Switching Based on the Link Prediction	107
4.9	Illustration of the AP-Pose Sensing Algorithm	110
4.10	Example of the Beam Strength Vector (BSV)	112

4.11	Samples of Measured Beam Patterns of the Phased Array	113
4.12	Example of the Beam Strength Map (BSM)	114
4.13	Compressive Angle Sensing and Reflector Estimation	119
4.14	Estimate the Initial Pose of the Client	123
4.15	Pia's Hardware Components and Configuration	124
4.16	System Initial Training Overhead	127
4.17	System Sensitivity to Client's Pose Error	128
4.18	Packet Latency During the AP Switching	129
4.19	Angle Sensing Error	130
4.20	Estimated Gain Error Over SNR of BSV	130
4.21	Pose Estimation Error Using Angle Sensing	131
4.22	Accuracy of Interference Prediction	131
4.23	SLR-Based Assignment vs SIR-Based Assignment	132
4.24	Link Availability and Hazard Times	133
4.25	Video Frame Latency	134
4.26	Throughput and Concurrency Opportunity	135
5.1	60 GHz vs 2.4 GHz Link Stability	138
5.2	Illustration of the MRF Procedure	143
5.3	Digital Array vs Analog Array	144
5.4	Locate the Rx and Reflecting Points Using a Single Tx	152
5.5	Reflecting Point and Virtual Source	154
5.6	E-Mi Implementation	158
5.7	Dominant Reflectors in an Office Environment	160
5.8	AoA/AoD Estimation Accuracy	160
5.9	Path Length Estimation Accuracy	161
5.10	Office Environment Reconstruction Accuracy	162
5.11	Corridor Environment Reconstruction Accuracy	164
5.12	AoA/AoD Prediction Accuracy	164
5.13	Path Length Prediction Accuracy	164
5.14	RSS Prediction Accuracy	165
5.15	Experiment in a Complicated Printing Room	165
5.16	Spatial Channel Profile	166
5.17	Reflector Location Error	166
5.18	Client Throughput CDF Under a 2-AP Architecture	170
5.19	Coverage and Throughput Enhancement	170

LIST OF TABLES

2.1	Query Frequency of RFID Tags	45
2.2	Efficiency of the Tracking Algorithm	45
3.1	Comparison of Different RF-Based Tracking Systems	58
3.2	Performance Profiles of mTrack's Phase Tracking Algorithms	75
3.3	Classification Accuracy of Touch Event	90
5.1	Estimated Reflection Loss on Different Reflectors	163

LIST OF ALGORITHMS

2.1	Sensing the Effective Layout of a Tag Array	32
3.1	Basic Phase Tracking Algorithm	66
3.2	Decision Tree for Touch Gesture Detection	80
4.1	AP-Pose Sensing	111
4.2	SLR-based AP and Beam Assignment	116
5.1	Virtual Beamforming	148

LIST OF ABBREVIATIONS

ALS	Array Layout Sensing
AoA	Angle of Arrival
AoD	Angle of Departure
AP	Access Point
APA	Anchor Point Acquisition
APS	AP-Pose Sensing
ASK	Amplitude-Shift Keying
BD	Background Dominating
BI	Beacon Interval
BSM	Beam Strength Map
BSV	Beam Strength Vector
CDF	Cumulative Distribution Function
CFO	Carrier Frequency Offset
CMDS	Classical Multi-Dimensional Scaling
COTS	Commercial Off-The-Shelf
CSI	Channel State Information
CWNA	Continuous White Noise Acceleration
DAC	Digital-to-Analog Converter
DDBR	Dual-Differential Background Removal
DoF	Degrees-of-Freedom
EPC	Electronic Product Code

FMCW Frequency-Modulation Continuous Wave

FoV Field of View

FPS Frames per Second

IoT Internet-of-things

ITM Intermediate

LOS Line-of-Sight

MCS Modulation and Coding Scheme

MIMO Multiple-input and Multiple-output

mmWave Millimeter-wave

MRF Multipath Resolution Framework

NLOS Non-Line-of-Sight

OFDM Orthogonal Frequency Division Multiplexing

PC Personal Computer

PCR Phase Counting and Regeneration

PDoA Phase Difference of Arrival

RF Radio Frequency

RFID Radio-Frequency Identification

RSS Received Signal Strength

RSSI Received Signal Strength Indicator

Rx Receiver

SIR Signal-to-Interference Ratio

SLAM Simultaneous Localization and Mapping

SLR Signal-to-Leakage Ratio

SNR Signal-to-Noise Ratio

SPS Samples per second

SSub Static Subtraction

TD Target Dominating

TDMA Time-Division Multiple Access

Tx Transmitter

UHF Ultra High Frequency

VCO Voltage-Controlled Oscillator

VR Virtual Reality

VS Virtual Source

WLAN Wireless Local Area Network

CHAPTER 1

Introduction

The rapid evolution of technology has proliferated numerous mobile devices, together with many exciting IoT services/applications. Today, we can already find many IoT products for the smart home in the market, *e.g.*, wireless door lock and voice assistant, enabling a safe and intelligent living environment. However, the trend of IoT will go beyond this, towards the smart city, smart industry, and connected vehicle. They together form a vast potential market that triggers tremendous revenue¹, which will continue to grow steadily in the next five years.

Among all the technologies, sensing and networking are two supporting techniques that are most demanded by the IoT, because many applications heavily rely on *information*, which is *exchanged and collected* from nearby objects, human, and environments, to make the computational decision. Traditionally, these functionalities are provided by versatile sensors (*e.g.*, light sensor, thermometer, microphone, *etc.*) and networking infrastructures (*e.g.*, Bluetooth, and WiFi). However, many emerging IoT applications, such as wireless VR and self-driving cars, are demanding supports beyond today sensors' capability and are quickly exhausting the capacity of current wireless infrastructures. We summarized the two demands as following:

- **Sensing Demand:** Due to the decent accuracy, using active sensors has become the predominant approach to collect sensing data in current IoT environments.

¹According to the Ericsson mobility report [23], the number of connected devices is estimated to reach \$20 billion in 2018, and this number will become 29 billion in 2022.

However, active sensors require battery power and need to be installed onto the target object, and thus become unsuitable for applications requiring long-term and automated monitoring. For example, users often find it intrusive and inconvenient to wear active sensors for activity sensing. Many IoT applications are still looking for new sensing capabilities to enable more intelligent functionalities in smart home and smart industry. Designing a batteryless and passive sensing scheme is critical for such IoT applications.

- **Networking Demand:** The next-generation wireless systems will face unprecedented challenges due to applications' surging demands for much higher throughput, lower latency, and better connectivity. mmWave networking has emerged as core technology to realize multi-Gbps wireless connectivity. However, due to intrinsic signal-strength attenuation, mmWave radios must use highly directional, electronically steerable phased-array antennas for communication purposes. But the high directionality makes the mmWave links highly unreliable — minor mobility, human blockage, and nearby reflectors can drastically affect the connectivity, making the link performance a sensitive function of the physical environment.

Past work in passive sensing has augmented cameras with advanced imaging processing as a primary mechanism. For example, Kinect [5], an input device for video game consoles, has demonstrated the capability of fine-grained human motion sensing. Nevertheless, the camera is sensitive to lighting condition, background variation, and occlusion. Therefore, they are unsuitable for many smart home applications that require working in the darkness or without light-of-sight. More importantly, massive use of cameras may cause a significant threat to privacy and safety². On the other hand, radio-based passive sensing has emerged as a promising alternative to overcome the above drawbacks in camera-based sensing. By applying sophisticated analytical algo-

²Researchers from the security firm Qualys demonstrated that Internet-based camera system could be hacked and cause privacy leakage [152].

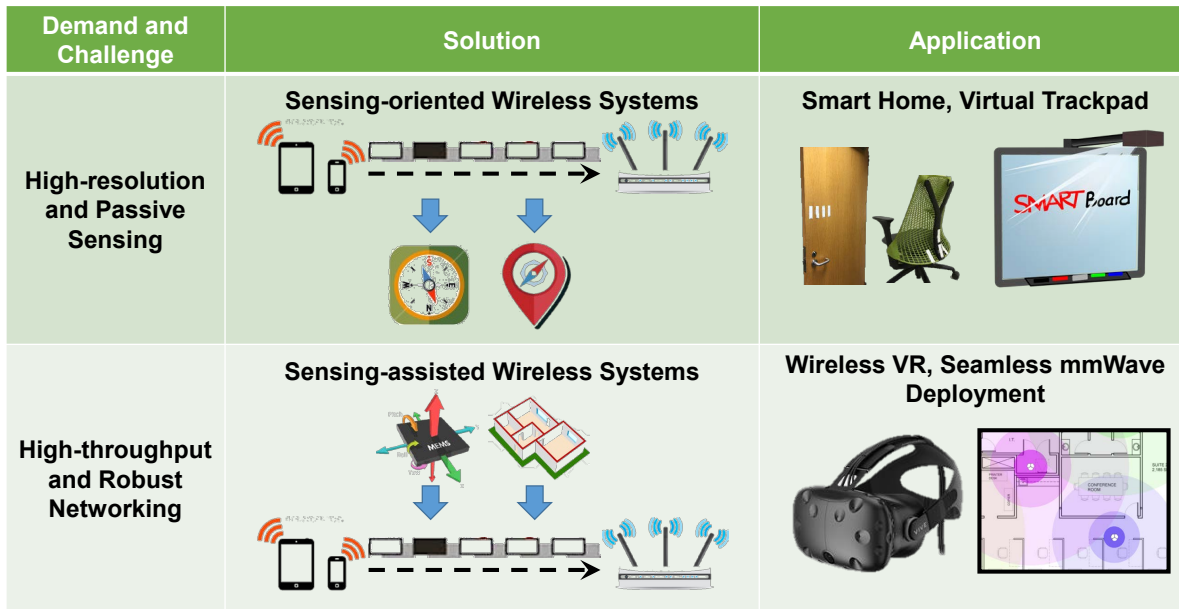


Figure 1.1: Thesis Contributions: This thesis introduces sensing-augmented wireless systems to meet the sensing and networking demands in IoT.

gorithms, we can extract critical sensing information from the wireless channel, which is disturbed by nearby objects. Existing works have demonstrated competitive capability in detecting and tracking a person. For example, radio tomography [184] detects human activity, WiFi imaging [80] senses object shape, and WiFi gesture control [111,130] recognizes different human body activities. However, unfortunately, today’s most wireless sensing methods are limited to tracking coarse motion or detecting a set of well-predefined patterns.

Past work in mmWave networking has designed algorithms and protocols to reduce the beam-searching overhead [26,160,210] and to improve the link reliability [25,166]. Nevertheless, these solutions mainly focused on the point-to-point link, and they are not scalable the dense mmWave network deployment. It is necessary to bring mmWave into large-scale mobile network settings, preparing this technology to meet the IoT traffic explosion.

To support the new demands from IoT, we propose the concept of *sensing-augmented wireless systems*. The key is to design computational models to process *wireless channel*

information and *sensor information* on wireless devices, with the end goal of enabling *new IoT sensing applications* and solve *critical network performance issues*. We no longer consider the radio as a standalone component, but instead, it will exploit and share intelligent sensing information with other system components. Depending on the flowing direction, we categorize it into two paradigms (Figure 1.1):

- **Sensing-Oriented Wireless Systems:** We turn the networking system into a “radar” that can provide information regarding nearby objects, human, and environments to the IoT applications.
- **Sensing-Assisted Wireless Systems:** The networking system takes sensing information from external sources, such as active sensors, to help optimize the network protocol and improve the networking performance.

This thesis presents multiple sensing-augmented systems for IoT that represent either new sensing schemes that enable radio-based super-resolution sensing of battery-less objects or new networking paradigms that demonstrate scalable mmWave network deployment under human mobility and blockage. The following sections briefly introduce these systems and present their core contributions.

1.1 Ubiquitous Wireless Sensing for Batteryless IoT

When the wireless signal propagates through space, it hits and bounces off nearby objects or human in the environment. As such, many emerging IoT networking paradigms can reveal sensitive information about the physical environment. The first part of this thesis enables the high-resolution passive sensing technique for today’s IoT applications. We achieve this goal by repurposing the functionalities of IoT wireless devices from legacy communications to ambient sensing. However, the signals are often covered by noise and mixed with many other disturbing components from unrelated tar-

gets. In response to these challenges, we apply a series of computational algorithms. We introduce the background signal cancellation technique for the mmWave radio and enable mm-level precision tracking of passive objects, e.g., pen or finger. The design can turn any desktop surface into a virtual trackpad. Besides location tracking, we further present a new system design that can sense the 3D orientation of batteryless objects in a smart-home environment. Our solutions have been implemented using Commercial Off-The-Shelf (COTS) communication radios, and we demonstrated they could achieve near-vision resolution.

This part makes the following key contributions:

- *Batteryless 3D Orientation Tracking using RFID Tags*: 3D orientation tracking is an essential ingredient for many IoT applications. For example, in a smart home environment, it can monitor the rotation and placement of daily objects. However, existing orientation tracking systems commonly require motion sensors like gyroscope and compass, which are only available on battery-powered devices. Tracking the orientation of passive batteryless objects is an equally important but more challenging problem.

To approach this objective, we designed an RFID-based passive orientation tracking system called Tagyro [186]. It attaches an array of RFID tags as orientation sensors on everyday objects. The tag-to-reader distance varies when the tag array rotates together with the object, which will alter the relative phase offset between tags. Based on this observation, we built a closed-form model to transform the runtime phase offsets into the orientation angle. However, the key challenge lies in the imperfect radiation pattern of practical RFID tag antennas, which substantially distort phase measurement. We addressed these challenges by designing a set of phase sampling and recovery algorithms, which together enable reliable orientation sensing with 3 degrees of freedom. We implemented a real-time version of Tagyro using the commercial RFID tags/antennas/reader (Impinj R420) and

demonstrated it could achieve high accuracy in orientation tracking. It is the first study of an RFID-based 3D orientation tracking system, which has a broad set of use cases, *e.g.*, smart home and mobile gaming.

- *Fine-grained Passive Motion Tracking Using mmWave Radios*: Radio-based passive sensing technique detects the motion of human body parts and objects through wireless signals, with the potential to spur a wide range of applications, *e.g.*, ubiquitous user-mobile interfaces. Prior work has employed various wireless signal features to sense objects under a set of predefined, coarse motion patterns. But an operational UI, like a trackpad, often needs to identify fine-grained, arbitrary motion.

To enable fine-grained passive motion tracking, we developed a practical system – mTrack [185]. It can track passive writing objects (*e.g.*, pen) at sub-centimeter precision. The key is to leverage the highly-directional 60 GHz millimeter-wave radio. Its miniaturized phased-array can form focused beam direction, creating stronger reflections from small objects. Moreover, the short wavelength of mmWave signals makes the phase feature highly sensitive to movement even at sub-centimeter scale. We designed the trajectory tracking algorithm using a signal-phase based model and incorporated novel mechanisms to suppress interference from background reflections to improve system reliability. We prototyped the mTrack system on a 60 GHz platform and demonstrated it could track a pen at sub-centimeter level accuracy, realizing the functionality of wireless transcription and virtual trackpad. Given the growing popularity of 60 GHz devices, this work potentially opens up many new pervasive sensing applications for everyday settings.

1.2 Robust mmWave Networks for Connected IoT

The second part of this thesis enables robust mmWave networking for throughput-demanding wireless IoT applications. Through experiment-based research exploration, we found these mmWave systems exhibit a significant dependence on the physical environment — their performance becomes a sensitive function of location, mobility, and interaction with human activities. The traditional approach in wireless protocols is to adapt link decision merely based on channel feedback. This approach works well for the point-to-point link setup but becomes inefficient under multi-AP deployment and user mobility. We believe if the radio can sense or recognize the physical space through multi-dimensional modalities of the sensing information, they will adapt the systems better to ever complicating application scenarios. However, developing smart wireless networks is challenging due to the intrinsic relationship between the target environment, abstract sensing data, and tremendous decision-space. We designed several sensing-augmented wireless systems that exploit computational techniques to improve the performance of emerging mmWave networks. Notably, the computational approach empowers radios to sense and reconstruct the radio environment, which in turn helps predict and optimize the network performance. The technique enables robust and highly efficient mmWave networks for new services, such as high-quality untethered VR.

This part makes the following key contributions:

- *Pose Information Assisted Robust 60 GHz Networks*: Although the beams emitted by mmWave radios are electronically steerable, their joint coverage, also known as Field of View (FoV), is still limited, making the links extremely sensitive to user location and even orientation change. Hence, achieving stable 60 GHz connectivity, even at room-level, has been a nontrivial task.

To overcome this challenge, we developed Pia [187], a robust 60 GHz mmWave system that can provide seamless coverage and mobility support at multi-Gbps

bitrate. It comprises multiple cooperating access points (APs), each covering a “picocell” region that can complement others’ blind spots. To facilitate the optimal AP selection and beam steering out of hundreds of directions, we harness the pose information from mobile client devices. It not only can make the network decision proactively (*i.e.*, ahead of time) to prevent link outage due to mobility, but also can manage cross-link interference, enabling high-efficient spatial reuse. To assist such decision making, we designed a set of AP-pose sensing and compressive angle estimation algorithms, which adopt a model-driven approach and fuse the pose information with link quality measurement. The entire system was implemented using commercial 60 GHz wireless adapters, and many-antenna phased array antenna (Qualcomm QCA9500). We demonstrated that the pose information, as the fundamental primitive, can improve the robustness of the mmWave network under mobility and capacity under multi-user spatial reuse. The design can be applied to meet the needs of demanding Gbps applications such as wireless VR, as well as to scale mmWave pico-cell/small-cell networks with dense AP deployment.

- *Environment Sensing for Optimizing mmWave Network Deployment:* Despite the potential of delivering multi-Gbps connectivity, 60 GHz millimeter-wave networks are vulnerable to propagation loss and obstacle blockage. To overcome such limitation, 60 GHz radios leverage reflections from object surfaces to steer around obstacles. Dependence on directivity and reflection, however, make the performance of 60 GHz networks a sensitive function of environment structure and reflectivity. Our measurement study manifested that a 60 GHz link can more efficiently detour blockage when placed near a concrete wall, which behaves like a mirror. Existing networking paradigms, which only focus on the cyberspace (network stack) rather than the physical environment, cannot handle such problems.

E-Mi [188] is a framework that harnesses 60 GHz radios’ sensing capabilities to

boost network performance. It uses a computational model to resolve the angular relation between 60 GHz transceivers, by simply extrapolating information out of the communication signals. From the geometrical model, we devised a reconstruction algorithm that can model the course outline of the environment through a sparse set of radio signal measurement. The environment model is then fed into a ray-tracer to predict the channel and network performance of arbitrarily located links. We prototyped the system using a customized 60 GHz software-defined radio testbed and demonstrated it could improve the network throughput and robustness by multiple folds. E-Mi represents a sensing-assisted paradigm that senses the environment from 60 GHz radios' eyes. It can not only facilitate mmWave networks but also augment a broad range of network planning and protocol reconfigurations.

1.3 Organization

The remainder of this thesis is organized as follows. In Chapter 2, we introduce Tagyro that turns batteryless RFID tags into passive sensors for orientation sensing of everyday objects. Next, Chapter 3 describes how mTrack localizes and tracks a passive object with sub-centimeter accuracy based on mmWave communication radios. Chapter 4 discusses how the pose information from portable devices can facilitate 60 GHz network operations, *e.g.*, AP switching and interference avoidance. Then, Chapter 5 presents a computational framework that optimizes the 60 GHz network deployment through environment sensing from mmWave radio's eyes. The final chapter concludes the thesis and presents our future work.

Part I

Ubiquitous Wireless Sensing for Batteryless IoT

CHAPTER 2

Batteryless 3D Orientation Tracking using RFID Tags

3D orientation is an essential input to many pervasive applications, such as motion tracking [177], mobile gaming, and activity recognition [31]. However, existing orientation tracking systems commonly require motion sensors like gyroscope and compass, which are only available on battery-powered devices. Tracking the orientation of passive batteryless objects is an equally important but more challenging problem.

A passive orientation sensing system can enable many new Internet-of-Things (IoT) applications (Figure 2.1). For example, in smart-home environment, it can monitor the rotation and placement of daily objects (*e.g.*, detecting door closure and monitoring activity of a rolling chair). In a supply chain, it can detect the tilt angle of orientation-sensitive cargo and warn the human workers when necessary. Orientation information of products on a conveyor belt can assist robotic arms to pick up an item using proper posture. Likewise, there exist numerous applications in mobile gaming (*e.g.*, guiding the positioning of LEGO pieces) and library management (*e.g.*, detecting upside down books). Existing passive orientation tracking solutions are predominantly based on computer vision [50, 77, 148, 153, 176]. Such systems often require extensive per-object per-angle training. They are sensitive to background variation, ambient light condition, and occlusion. Therefore, they are unsuitable for ad-hoc IoT applications.

In this chapter, we discuss *Tagyro*, which can track the 3D orientation of passive




Figure 2.1: Applications of passive orientation sensing in IoT.




Figure 2.2: Examples of linear and square tag arrays from the top view (showing cross sections of tags).

objects using RFID tags. The basic idea is to attach multiple tags to different spots on the object, which constitute a *tag array*. Some example array patterns are illustrated in Figure 2.2. An RFID reader can measure the *tag phase*, *i.e.*, the phase of signals returned from each tag. When the tag array rotates together with the object, the tag-to-reader distance among different tags varies, which in turn affects the relative phase offset between tags (referred to as Phase Difference of Arrival (PDoA)). Tagyro builds a closed-form model that transforms the measured PDoA into an *orientation spectrum*, which characterizes the likelihood of each orientation angle. It then derives the tag array's 2D orientation from the peak of orientation spectrum. By attaching an additional RFID antenna and tag array, Tagyro can aggregate the outputs from multiple 2D estimators and produce a 3D orientation estimation for the target object.

Realizing such an RFID-based 3D orientation tracking entails three key challenges that are unaddressed in previous RFID sensing systems [105, 151, 178, 181, 200]. *First*, to prevent phase ambiguity, adjacent tags in a tag array must be placed within close

proximity. However, our experiments reveal that the closely-placed tags tend to alter each other’s phases due to sophisticated electromagnetic interactions. Such *coupling effect* deviates the phase value from the traditional phase-distance model [200]. To address this issue, we propose an Array Layout Sensing (ALS) algorithm, which approximately models the coupling effect as if the tags repel each other in the physical space. ALS also enables Tagyro to automatically sense the tag array’s effective layout – a critical input for the orientation spectrum computation. Consequently, ALS also frees the user from knowing the array’s physical layout or placing the array with precise geometry.

Second, the radiation pattern of a practical RFID tag is far from the ideal isotropic model, *i.e.*, a point source that scatters signals uniformly in 3D space. Instead, it has blind directions with extremely weak emissions, rendering phase measurement infeasible. Tagyro addresses this challenge using a dual-array setup, where two arrays are arranged in orthogonal angles to compensate the blind directions. Tagyro can single out valid phase readings, and then unify the results to estimate the object’s orientation with 3 Degrees-of-Freedom (DoF).

Third, commercial FCC-compliant RFID readers run in the frequency-hopping mode and interrogate the tags sequentially, which cause phase discontinuity and significantly slow down the orientation sensing. We design a simple per-reader phase calibration mechanism, along with phase interpolation, which improves the phase-reading precision/granularity.

Tagyro, to our knowledge, represents the first wireless sensing system that can track the 3D orientation of passive objects. Its contributions can be summarized as follows.

(i) We propose an RFID tag-array setup for phase-based orientation tracking, and empirically study the impact of tags’ imperfect radiation patterns on phase reading. Our study reveals how the tag polarity, blind direction and coupling effect distort phase measurement, which has not been reported by previous RFID sensing/localization systems [96, 105, 116, 151, 178, 181, 200].

(ii) We propose an innovative algorithm to sense the effective layout of a tag array, which overcomes the tag coupling effect. We further propose mechanisms to combat the tags' non-isotropic radiation patterns, enabling orientation sensing in 3D space without blind direction. We have also developed phase calibration mechanisms and enabled fast orientation sensing on standard-compliant RFID devices.

(iii) We implement Tagyro on a COTS RFID system as a real-time 3D orientation sensing framework. Our experimental results show that Tagyro can track the 3-DoF orientation of an object at an average accuracy of 4° , with processing rate >37.7 samples/second and response latency <0.75 seconds. Tagyro achieves comparable accuracy and sampling rate as state-of-the-art gyroscope-based systems [212], but it needs no battery on the object, and evades the notorious drifting issue.

2.1 Related Work

Our work is most closely related with prior art in the following domains:

Sensor-equipped device. The most straightforward approach to sense orientation is to use active sensors, *e.g.*, magnetometer, gyroscope or accelerometer. Acube [212] combines these active sensors to estimate the orientation of a smartphone that is used for indoor localization and tracking. In industry setting [66], orientation sensors are deployed to guide robot arms to pick up objects. Active sensors are accurate but power hungry, and need frequent maintenance (*e.g.*, battery replacement or charging). RFID tags are batteryless and ideal for objects that cannot afford frequent maintenance. RFID-Die [45] and WISP [204] use customized RFID tags with embedded motion sensors that harvest energy from Radio Frequency (RF) signals. The power draw of sensors significantly reduces the tags' working range (down to 8 cm [45]) and reduces the duty cycle (*e.g.*, one sensing for every two-second charging [156]). In contrast, Tagyro uses COTS passive RFID tags, with a typical working range of around 3 m and sampling rate of

37.7 samples/second. Owing to low-cost of RFID tags (10 cents each [214]), Tagyro is readily usable in massive Internet-of-Things applications. In addition, mid-end smartphones with RFID readers are already available, which can bring Tagyro to home and in-situ usage. The Cricket Compass [129] tracks the orientation of a mobile device based on the distance difference of arrival between multiple pairs of ultrasonic transceivers. Although the high-level principle share similarities, Tagyro's RFID tag array tracking entails a unique set of challenges while bringing salient properties in terms of precision, cost, and energy efficiency.

Image-based orientation detection. Detecting the orientation of object is a general problem in computer vision [50,77,148,153,176]. Hinterstoisser *et al.* [77] proposed a template matching method to detect the tilt angle of objects in the image. Image-based approach is also used to detect human face, head and gaze orientation [48,89]. Saxena *et al.* [147,148] and Shimizu *et al.* [153] applied a pattern matching method to estimate the 3D orientation of objects in the image, which can guide robotic arms' grasping action. However, such image-based approaches only work in line-of-sight (LOS). They are very sensitive to lighting condition and fail even under partial occlusion by a paper box. On the other hand, Tagyro can accurately track object orientation even if the tags are occluded from the reader's LOS.

RFID tilt sensing. There has been preliminary work in RFID-based tilt sensing, which puts either the reader or tags on the object. Han *et al.* [74,154] used a site-survey approach to determine the position and orientation of a target (*e.g.* robot and chair) in a 2D space. However, they need to install the reader on the target, and require a dense deployment of RFID tags to cover the floor of interest. RF-Compass [178] uses a 2D-plane partitioning method to navigate a robot to gradually converge towards an object's direction, but cannot estimate the object's instantaneous 3D orientation. Krigslund *et al.* [72,95,96] estimate the 1D orientation of an RFID tag by tracking the RSS fluctuation based on the principle of antenna polarity. Due to polarity mismatch,

RSS decreases proportionally to the angular misalignment between the reader and the tag. However, the RSS is easily affected by distance and environment dynamics, and the polarity-based approach has 90° ambiguity in orientation estimation. Tagyro represents the first phase-based method to sense tag orientation without ambiguities, and in 3D space.

RF-based localization. RFID tag localization has been extensively studied recently [74, 76, 154, 179, 181, 200]. The state-of-the-art techniques, *e.g.*, Tagoram [200] and PinIt [179], can already locate the tag’s position at a cm-level or mm-level accuracy. On the other hand, estimating the orientation of RFID tag is a relatively new area that has not been well explored yet. The basic idea behind Tagyro’s tag array approach shares similar spirit with antenna-array based radio localization [197], which can track the relative angle between a multi-antenna receiver and an access point on the same 2D space. However, the unique properties of RFID, *e.g.*, tag coupling and non-isotropic radiation, bring unprecedented challenges to Tagyro.

One may think of applying them to estimate the object orientation: by attaching two tags to different parts of the object, its orientation can be derived from the estimated locations of tags. However, this approach has several limitations. *First*, existing schemes only address the 2D localization problem [74, 76, 154]. The 2D position (on a plane) limits the orientation estimation to only 1 DoF. *Second*, most state-of-the-art localization schemes have 10~20 cm errors [181]. They are unable to differentiate the position of two closely-placed tags, and thus cannot give precise orientation estimation. Although Tagoram [200] can estimate tag position at cm-level accuracy, it requires multiple readers that are precisely deployed at dedicated positions, which is not suitable for ad-hoc scenarios like smart home or mobile gaming. STPP [151] can estimate the relative position for closely-separated tags, but it can only discriminate the left/right/up/down ordering. Furthermore, the non-ideal radiation of RFID signal (Sec. 2.4.2) invalidates the basic assumption that the phase changes linearly over distance,

which requires us to rethink the effectiveness of existing phase-based RFID localization techniques [151, 200] when working in the 3D space.

RFID radiation pattern. Existing electromagnetic research has studied the imperfection of RFID tag radiation pattern. Zhang *et al.* [208] used two dipole antennas in orthogonal directions to achieve full RSS coverage over 3D space. Tagyro adopts a similar idea to address the blind direction issue, which deploys two tag arrays in perpendicular directions (Sec. 2.4.2). Yet, Tagyro addresses the more critical problem of polarity-induced phase-shifting, through adaptive array selection and polarity alignment. Bolotnyy *et al.* [42] studied the tag readability issue caused by coupling effect in a multi-tag RFID system. To our knowledge, Tagyro is the first to evaluate the impact of coupling effect on the phase of RFID tag array, and to counteract the effect by using the layout sensing technique.

2.2 Preliminaries

2.2.1 Phase Model for an RFID System

An RFID reader can transmit continuous-wave signals to interrogate a tag, and then receive backscattered signals from the tag which contains its unique identity information. The phase offset between the transmitted and received signals depends on the round-trip propagation distance d , as well as hardware-specific factors. Following a standard phase-distance model [12, 200], the phase value ϕ , also referred to as *tag phase*, equals:

$$\phi = \text{mod}\left(\frac{2\pi d}{\lambda} + \phi_{\text{Reader}} + \phi_{\text{Tag}}, 2\pi\right), \quad (2.1)$$

where λ denotes the signal wavelength, and ϕ_{Reader} and ϕ_{Tag} are the additional phase offsets induced by the reader's transmission/receiving circuits and tag's antenna response. Eq. (2.1) implies that a single tag is insufficient for 3D orientation sensing, because




Figure 2.3: Coordinate system refers to the reader antenna. Tag orientation is defined as the rotation offset $(\theta_x, \theta_y, \theta_z)$ along each axis.

simply changing the tag's orientation may not change the distance d and the phase ϕ .

2.2.2 Phase Difference of Arrival of a Tag Array

The basic idea of our approach is to form a tag array with multiple tags. Due to different distances to the reader, the tags cause different phase offsets, which we call the *Phase Difference of Arrival* (PDoA). Rotation of the tag array results in different PDoA values, from which we can estimate the orientation of the tag array, and hence the orientation of the object it is attached to.

Consider Figure 2.3 where 4 tags form a square array. We first specify a 3D *coordinate system for the RFID reader*, in which the X-axis is parallel to the direction from the reader to the center of the tag array, and the X-Y plane is parallel to the ground. The *tag array's orientation* is defined as the rotation offset $(\theta_x, \theta_y, \theta_z)$ along each axis, relative to an initial state specified by the user. *This coordinate system allows us to track the tag array's orientation from the reader antenna's perspective.* Note that Tagyro is different from active motion sensors (gyroscope and magnetometer) [212], since the coordinate system will change *w.r.t.* to the position of the tag array. On the other hand, active motion sensors are location-independent because gravitational and magnetic fields on earth are used as the reference coordinate.

Figure 2.4 illustrates the rotation around the Z-axis by projecting 3D space on to the




Figure 2.4: Relation between the tags' PDoA and rotation angle. This example shows rotation along the Z-axis (by projecting to the X-Y plane).

X-Y plane. In this example, signals first reach tag 1 before tag i , and thus PDoA (tag i to tag 1) is positive. More generally, *the direction and amount of rotation determine the sign and value of PDoA, respectively*. Assuming signals come from far field, the PDoA $\Delta\phi_i$ for any tag i to the reference tag, say tag 1, can be derived from the tag separation distance projected to the reader's signal direction (*i.e.*, the red line in Figure 2.4):

$$\begin{aligned} \Delta\phi_i(\theta_Y, \theta_Z, d_i) &= \phi_i - \phi_1 \\ &= \text{mod}\left(\frac{4\pi d_i \sin(\theta_Y + \theta_Y^i) \sin(\theta_Z + \theta_Z^i)}{\lambda} + \Delta\phi_{\text{Tag}}^i, 2\pi\right), \end{aligned} \quad (2.2)$$

where d_i is the distance between the i^{th} and reference tag. θ_Y^i and θ_Z^i are initial angular offsets for tag i and reference tag in the initial state, which gauge the rotations along Y and Z axes respectively *w.r.t.* the directions of zero PDoA. They can be calculated from the array geometrical layout (to be addressed in Sec. 2.4.3). $\Delta\phi_{\text{Tag}}^i$ is the difference of antenna phase response (cf. Sec. 2.2.1) between tag i and the reference tag, which will be calibrated and removed by our array sensing module (Sec. 2.4.3.3). Note that PDoA cancels the unknown phase offset ϕ_{Reader} introduced by the reader's circuits.

In addition, note that the PDoA is irrelevant to the X-axis rotation in this example setup (Figure 2.3), since any rotation along X-axis does not change the distance offset between different tags and the reader. Therefore, *a single reader antenna can only track*

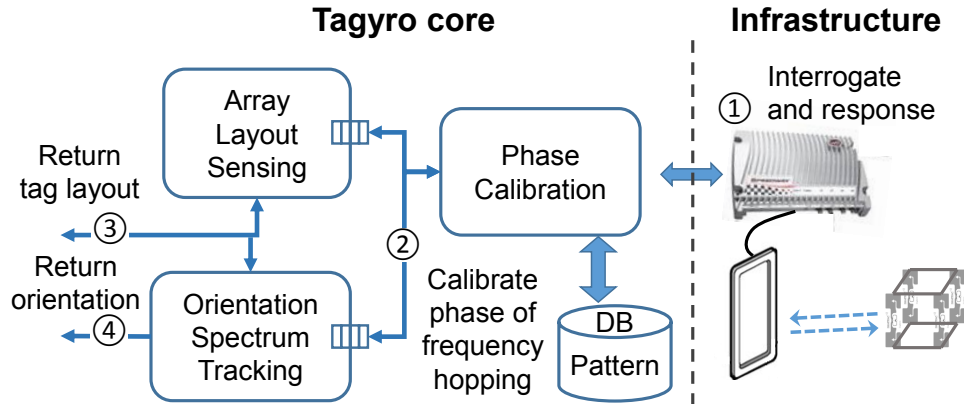


Figure 2.5: Tagyro's architecture and operation workflow.

the array rotation with 2 Degree-of-Freedoms (DoFs), i.e., along Y and Z axes. For 3-DoF tracking, Tagyro requires a reader with at least 2 antennas, which will be discussed in Sec. 2.4.4.

2.3 Tagyro Operation

Tagyro is an orientation sensing system that generalizes the aforementioned ideal PDoA tracking. To use Tagyro, a user needs to attach an array of RFID tags on the target object, following certain empirical guidelines and without knowing the precise geometrical relation between tags.

Figure 2.5 illustrates Tagyro's flow of operations, comprised of the following steps: (1) The reader interrogates the tag array following standard EPC protocol [140], which reads the tags sequentially and resolves collision via retransmission. Commercial readers (e.g., the Impinj R420 [21] that we use) can extract many pieces of information such as tag ID, phase and Received Signal Strength Indicator (RSSI), which are forwarded to a PC host running Tagyro. (2) Tagyro's *phase calibration* module pre-processes the phase information, and removes the phase discontinuity caused by the reader's channel hopping. (3) Tagyro's *Array Layout Sensing* (ALS) module estimates the tag array's effective layout, which differs from the geometrical layout due to mutual coupling between tags.

(4) Given the tag array layout, Tagyro measures the tag phase and runs an *orientation-spectrum tracking mechanism* to continuously track the tag array's orientation.

In what follows, we describe Tagyro's major design components and algorithms. We first describe how to compute orientation from the PDoA, assuming the tag array's layout is known (§ 2.4.1). Second, we explore the main challenges in PDoA and tag layout estimation (§ 2.4.2), particularly caused by the tags' imperfect radiation pattern. We introduce the ALS algorithm that estimates the tag array's effective layout, taking into account the coupling effect (§ 2.4.3). Third, we extend the single-array and single-antenna reader case to dual-array and dual-antenna reader, in order to realize 3-DoF orientation tracking (§ 2.4.4). We will also discuss phase calibration and run-time optimization mechanisms that improve the accuracy and speed of phase reading (§ 2.4.5 and § 2.4.6).

2.4 Design

2.4.1 Tracking through Orientation Spectrum

We first focus on the 2-DoF orientation, *i.e.*, sensing (θ_Y, θ_Z) , the array's angular rotation along Y and Z axis, relative to its initial state. Given the measured PDoA between each tag i and the reference tag, Tagyro estimates the tag array's orientation by computing an *orientation spectrum* – a 2D intensity graph in which each pixel indicates the likelihood of an orientation. We use the phaser-domain difference between the theoretical PDoA $\Delta\Phi_i$ (computed from Eq. (2.2)) and measured PDoA $\Delta\hat{\Phi}_i$ as pixel intensity, *i.e.*,

$$I(\theta_Y, \theta_Z) = \left| \sum_{i=1}^K e^{j(\Delta\Phi_i(\theta_Y, \theta_Z, d_i) - \Delta\hat{\Phi}_i)} \right| / K \quad (2.3)$$

where K is the tag number in the array. The phaser domain operation inherently handles the phase aliasing. When the theoretical PDoA matches the measured value, *e.g.*,




Figure 2.6: Example orientation spectrum from measurement.

$|\Delta\Phi_i - \Delta\hat{\Phi}_i| \approx 0$ for all i , the pixel will have the maximum intensity equal to 1. Otherwise the pixel intensity will be close to 0. The complexity of computing the spectrum is $O(N^2)$, where N is the number of values for each DoF rotation. In Tagyro, we discretize the rotation value θ_Y and θ_Z into 3° granularity, resulting in an orientation spectrum with 120×120 pixels.

The 2-DoF orientation is estimated by finding the highest peak in the orientation spectrum:

$$\{\theta_Y, \theta_Z\} = \arg \max_{\theta_Y \in [0, 2\pi), \theta_Z \in [0, 2\pi)} I(\theta_Y, \theta_Z). \quad (2.4)$$

Taking the 4-tag square array in Figure 2.3 again as an example, where neighboring tags are separated by the quarter wavelength of 8.2 cm, Figure 2.6 plots the orientation spectrum when the array rotates at $(\theta_Y = 0^\circ, \theta_Z = 80^\circ)$. The highest peak position occurs exactly at the same orientation. Besides, there are weaker grating lobes that are caused by the *spatial ambiguity* effect.

Spatial ambiguity is a well known effect when using multi-antenna to estimate a signal's angle of arrival – multiple angles may have similar likelihood of estimation. To minimize the spatial ambiguity, antennas need to be separated by less than half-wavelength ($\lambda/2$) away. For an RFID tag array, the equivalent separation should be $\lambda/4$ (i.e., 8.2 cm for 915 MHz UHF tags), because the signals traverse a round-trip




Figure 2.7: Measuring RSS and phase of a tag rotating along each axis. The 3D heat map shows radiation pattern of a typical tag.

propagation. Though grating lobes are generally weaker than the main lobe, they may occasionally obscure the detection of true peak because of residual phase measurement error and noise in RF channel. To ensure the robustness of Tagyro, instead of picking a single strongest peak, we search for the top three peaks in the orientation spectrum, and take the one that is closest to the previous estimation as final orientation result. This serves as a simple moving average outlier filter.

The above ideal orientation-spectrum tracking model requires three underlying conditions: (i) a tag backscatters wireless signals like an isotropic point source, so the signals cover 3D space without blind direction; (ii) the phase reading of a tag follows a predictable model over distance (Eq. (2.1)); (iii) tag array’s relative layout and the corresponding IDs are known. We next explore the practical challenges and solutions in Tagyro to meet these conditions.

2.4.2 Impact of Tag Radiation Pattern

We conduct empirical studies using a COTS RFID platform (Impinj R420 [21]) to better understand the radiation property of RFID tags. Although the first two of above basic assumptions (about RSS and phase) have been applied in previous works to localize tags in a 2D space [105, 151, 200], we are particularly interested in verifying whether they still hold under the tag array setup and the implications for 3D orientation tracking.




Figure 2.8: Measured phase (left) and RSS (right) of a single tag rotating along each of the three axes respectively.

2.4.2.1 Rotation of a Single Tag – Polarity and Blind Direction Issues

Observation 1 *The phase of a single tag is not only determined by the distance, but also affected by its relative orientation to the reader antenna.*

We first reinvestigate the standard phase-distance model (Eq. 2.1) under varying tag orientations. We place a slim tag 1.5 m in front of a flat-panel reader antenna (Figure 2.7), and rotate the tag center along each axis. Figure 2.8 plots the corresponding phase change. Ideally, the phase should remain constant, if it only depends on the tag-to-reader distance. Surprisingly, we find that the phase shifts linearly with the rotation along X-axis, and abrupt phase jumps occur with rotation along Y-axis. The phase only remains unchanged for the rotation along Z-axis.

This phenomenon is caused by the *polarity* of RFID antennas. Due to form-factor constraint, a tag’s antenna is commonly designed as a dipole, which is linear-polarized along the tag body (*i.e.*, Z-axis in Figure 2.7). On the other hand, to ensure they can read tags from a variety of angles, most RFID readers’ antennas are circular-polarized [100] – They comprise two perpendicular dipoles, fed with signals of 90° phase difference. When the tag rotates along the X-axis by 90° , the phase of signal received by the tag will change by $\frac{\pi}{2}$. *Since the signals traverse a round-trip propagation back to the reader, the corresponding phase shift measured by the reader doubles.* Similarly, the tag phase jump of 2π when rotating around Y-axis by 180° is also due to the abrupt change of polarity

direction. We repeat the experiment using other types of tags (Sec. 2.5) and observe a similar result.

Implication: Antenna polarity affects the tag phase, and may confuse the orientation spectrum computation which assumes phase only depends on the tag-to-reader distance. *To isolate the polarity effect, in Tagyro, we enforce a constraint on the tag placement, called polarity alignment – within a tag array, the tags’ bodies must be parallel to each other.* In this way, polarity introduces the same phase shift among all tags during rotation, and does not affect PDoA.

Observation 2 *An RFID tag has non-isotropic radiation: it mainly backscatters signals surrounding the tag body (XY-plane in Figure 2.7), and has blind direction along its radial direction (Z-axis).*

Under the same setup as above, Figure 2.8 plots the RSS with different tag rotations. We see that the tag reflects consistently strong RSS when rotating around X or Z axis. Yet, if it rotates around Y-axis, *the RSS drops significantly (by 20 dB) when the tag’s radial direction points to the reader.* The reader cannot reliably decode the tag’s signal at such blind direction, and hence the phase measurement becomes erroneous.

Implication: The blind direction of an RFID tag limits the range of orientation that Tagyro can track using a single tag array. To realize 3D orientation sensing, the blind direction of RFID tags must be compensated by other means.

2.4.2.2 Rotation of a Tag Array – Coupling Between Tags

Observation 3 *The phase of each tag in the tag array will be affected by neighboring tags through the coupling effect, which deviates actual PDoA from the standard phase-orientation model (Eq. 2.2).*

The coupling effect of RFID tag occurs in near-field where signals backscattered from a responding tag generate resonant voltage in the antennas of other tags [107]. As a result, besides the signal from responding tag, other nearby tags will also emit signals,




Figure 2.9: Measuring the PDoA of a two-tag array. One tag is in the rotation center and the other rotates along a circle $r=6$ cm.

which further alters the phase received by the reader. We now conduct two experiments to quantify such effect.

We now create a two-tag array following the polarity alignment requirement, and measure the PDoA between them (Figure 2.9). One tag is in the rotation center and the other rotates along a circle of radius $r = 6\text{cm}$. To evaluate the PDoA without coupling effect, we also collect the phase trace by rotating only one tag (removing the other), and then compute the PDoA at each of the sampling locations along the circle.

Figure 2.10 plots the PDoA for cases with (w) and without (w/o) coupling effect. We observe that when there is no coupling, the measured PDoA matches the standard model (Eq. (2.2)). However, when two tags rotate together, the *PDoA scope*, *i.e.*, difference between maximum and minimum values, has increased significantly. The result shows that PDoA value is scaled by the coupling effect between tags.

To better understand the root cause of PDoA deviation, we model the coupling effect of two tags as illustrated in Figure 2.9. We first consider the case without coupling, the carrier signal backscattered from Tag1 to the RFID reader is characterized by:

$$s_1 = A_1 \exp(j2\pi \frac{2d_1}{\lambda}) \quad (2.5)$$

where A_1 denotes the amplitude of received signal and $2d_1$ is the round-trip distance between tag and reader. Without loss of generality, we omit other constant phase offsets caused by tag/reader antennas. Eq. (2.5) is consistent with the standard phase-distance

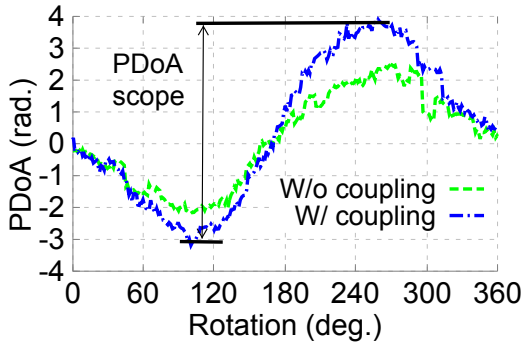


Figure 2.10: Measured PDoA value of a two-tag array.

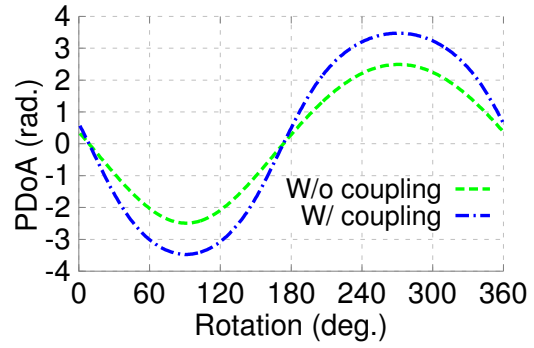


Figure 2.11: Simulated PDoA value of a two-tag array.

model (Eq. (2.1)) – phase is determined by the tag-to-reader distance.

When the coupling effect occurs, RFID reader essentially receives two copies of signals. One is the original signal coming from the responding tag, and the other copy is generated by the resonant voltage and coming from the other tag. Compared to the original one, the resonant copy experiences certain attenuation and phase shift, which are uniquely determined by the coupling coefficients between the two tags. We denote the resonant signal emitted from Tag2 as s_2^c , which is given by:

$$s_2^c = A_2^c \exp(j2\pi \frac{d_1 + d_2 + r}{\lambda} + \varphi^c) \quad (2.6)$$

where A_2^c and φ^c are the amplitude and phase offset of the resonant signal. d_2 and r are the Tag2-to-reader and Tag1-to-Tag2 distances respectively. In a similar way, when Tag2 is responding, we can model the signals from Tag1 and Tag2 as s_1^c and s_2 respectively.

When both copies of signals overlap and together reach the RFID reader, the measured phase value will deviate from the origin one, and thus the PDoA distorts. We plug in empirical parameters $A_1 = 1$, $A_2^c = 0.5$, $\varphi^c = 90^\circ$, and simulate the PDoA of a two-tag array based on the analytical model. Figure 2.11 shows that the simulation result is highly consistent with our experimental measurement.

We further evaluate the extent of coupling effect over tag separation distance. We vary the radius r from $3cm$ to $18cm$, at $3cm$ step. We use the PDoA deviation as metric,

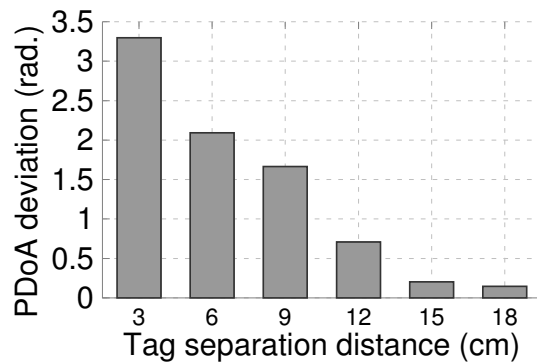


Figure 2.12: PDoA deviation over tag separation distance.

i.e., the difference of PDoA scope relative to the case without coupling. Figure 2.12 shows that *the closer two tags are placed to each other, the stronger the PDoA deviation will be*. Recall that computing orientation spectrum requires tags to be placed within quarter-wavelength distance (8.2cm) to minimize the spatial ambiguity (Sec. 2.4.1). Yet at this separation, coupling effect is strong and unavoidable.

Implication: The coupling effect, which deviates PDoA value from the phase-orientation model, will compromise the peak of orientation spectrum and consequently affect the performance of rotation tracking.

In summary, the above experiments reveal that the radiation pattern of RFID tags is far from the ideal model, and may compromise any RSS/phase-based sensing algorithms in 3D space. We now present the countermeasures in Tagyro.

2.4.3 Combating Tag Coupling Effect

2.4.3.1 Effective Distance Between Tags

Since tag coupling distorts the phase and PDoA, a natural countermeasure is to model the coupling effect and compensate the distortion accordingly. However, such a model must capture the mutual impedance between tags [81], which depends on elusive factors (*e.g.*, antenna shape and material) that are unlikely to be available to end-users.

To address this issue, we make an important observation from Figure 2.10: under

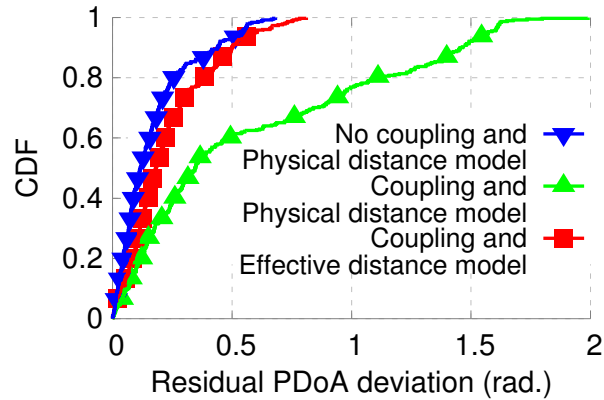


Figure 2.13: CDF of residual PDoA deviation. Physical distance and effective distance are 6 cm and 9 cm.

the coupling effect, the scope of PDoA is scaled, but the relation between rotation angle and PDoA maintains a similar trend as the no-coupling case. In other words, we can approximate the end-effect as virtually *rescaling* the tag separation to an appropriate value, which we call *effective distance*. Formally, the effective distance between the tag i and the reference tag is modeled as \hat{d}_i , which maximizes the alignment between measured PDoA and the theoretical value for any given orientation (θ_Y, θ_Z) :

$$\{\hat{d}_i, 1 \leq i \leq K\} = \arg \max_{d_i, 1 \leq i \leq K} \left| \sum_{i=1}^K e^{j(\Delta\Phi_i(\theta_Y, \theta_Z, d_i) - \Delta\hat{\Phi}_i)} \right|. \quad (2.7)$$

To verify this model, we conduct a benchmark experiment involving two tags with a physical distance of 6 cm. We measure the PDoA by rotating the array to three known orientations, and then find the effective distance through Eq. (2.7). Figure 2.13 shows the Cumulative Distribution Function (CDF) of PDoA deviation. We can see that, compared with the physical distance model that does not account for coupling, effective distance model reduces the maximum PDoA estimation error from 2 radians to 0.7 radians, and mean to as small as 0.1 radians. Overall, PDoA deviation is reduced to an extent comparable to the case without coupling effect (cf. Sec. 2.4.2.2). Note that, even without coupling, the PDoA model is imperfect and has minor residual errors due to the multipath and RF noise [105, 116, 181, 200]. They will be incorporated in our system

level experiments.

2.4.3.2 Sensing the Effective Layout of a Tag Array

However, in practice, it is difficult to compute the effective distance from Eq. (2.7) for more than two tags due to the huge search space. More importantly, Eq. (2.7) requires the knowledge of multiple rotation angles as input, which the user may not even be able to obtain. In Tagyro, we design an Array Layout Sensing (ALS) algorithm, which can automatically sense the effective distance between tags, and hence the layout of the entire tag array under coupling, which we refer to as *effective layout*.

First, to avoid requiring the knowledge of rotation angle, our key insight is that the scope of PDoA is determined by the effective distance of two coupled tags, and it is independent of the array's orientation. The PDoA value is bounded within $[-\frac{4\pi\hat{d}}{\lambda} + \Delta\phi_{\text{Tag}}, \frac{4\pi\hat{d}}{\lambda} + \Delta\phi_{\text{Tag}}]$ (cf. Eq. (2.2)). Thus, the effective distance can be determined by the PDoA scope as:

$$\hat{d} = \frac{\lambda PDoA_scope}{4 \cdot 2\pi}. \quad (2.8)$$

To obtain the *PDoA_scope*, Tagyro requires the user to rotate the tag array by more than one cycle (360°) roughly around each axis. Meanwhile, it measures the PDoA values and uses the difference between maximum and minimum PDoA as the *PDoA_scope*. This is a simple one-time calibration, needed only when the tag array is first formed.

Second, when measured phase value exceeds 2π , the RFID reader wraps it back to the $[0, 2\pi]$ range, which may cause ambiguity in estimating the max and min PDoA values. We remove the ambiguity through phase unwrapping [27]. Since the reader typically queries the tag at a much higher speed compared to the tag array rotation, the change between consecutive PDoA values is minor. Thus, any PDoA change greater than π or smaller than $-\pi$ indicates the occurrence of phase jumping, which can be removed by adding/subtracting π to the subsequent samples of PDoA values.

Third, computing the orientation spectrum requires the relative layout of tags in-

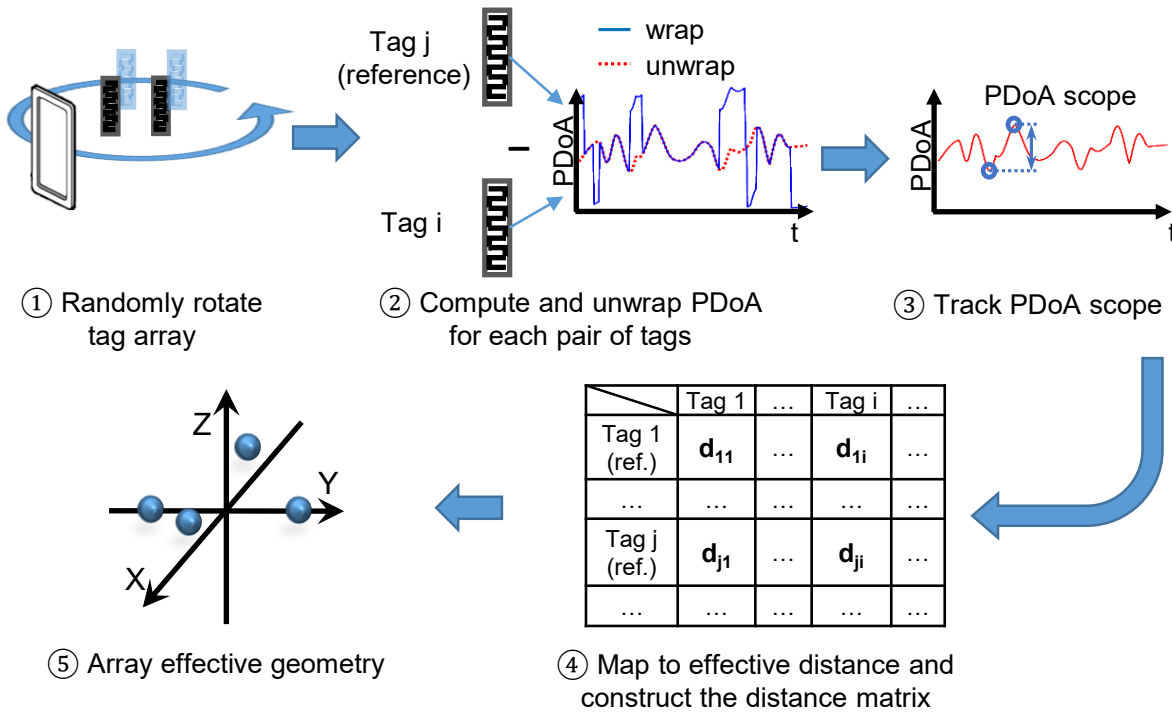


Figure 2.14: Array layout sensing algorithm to sense the effective geometry of a tag array.

stead of their separation. We approach this by mapping the tags' pairwise effective distance to the entire tag array's effective layout. This can be formulated as a Classical Multi-Dimensional Scaling (CMDS) problem [43], which takes a matrix that describes dissimilarity between points (e.g., effective distance in our case) as input, and outputs a coordinate matrix in the geometrical space that minimizes the residual distance error. An approximate algorithm [201] can solve this problem efficiently at complexity $O(K \log K)$ for K tags. We apply the CMDS function built-in MATLAB which by default projects the K points into a K -dimensional space. Since tags are located in a 3-D space, the coordinate values for the higher $K - 3$ dimensions are small and close to zero. So we can take outcomes in the first 3 dimensions as tag coordinates and ignore the others without much loss of accuracy. It is worth noting that using other implementations of CMDS algorithm will not affect the key design of Tagyro.

Figure 2.14 summarizes the operations of Tagyro's ALS algorithm: (i) When the

```

1: global  $p[1 \times K] \leftarrow -1$                                 ▷ Vector of historical phase
2: global  $a[K \times K] \leftarrow 0$                             ▷ Matrix of PDoA
3: global  $\mu[K \times K] \leftarrow -\text{inf}$                     ▷ Matrix of maximum PDoA
4: global  $\nu[K \times K] \leftarrow \text{inf}$                     ▷ Matrix of minimum PDoA
5: procedure SenseGeometry( $\sigma, \phi$ )
6:   if  $p[\sigma] == -1$  then                                ▷ Initialize new incoming tag
7:      $[p[\sigma], a[\sigma]] \leftarrow \text{initialize}(\phi)$ 
8:   end if
9:   for each tag i do                                       ▷ Reference tag
10:     $a[\sigma][i] = \phi - p[i]$                                ▷ Compute PDoA
11:     $a[i][\sigma] = p[i] - \phi$ 
12:    unwrap( $a$ )                                           ▷ Phase unwrapping
13:    if  $a[\sigma][i] > \mu[\sigma][i]$  then                   ▷ Update max PDoA
14:       $\mu[\sigma][i] = a[\sigma][i]; \nu[i][\sigma] = a[i][\sigma]$ 
15:    end if
16:    if  $a[\sigma][i] < \nu[\sigma][i]$  then                   ▷ Update min PDoA
17:       $\nu[\sigma][i] = a[\sigma][i]; \mu[i][\sigma] = a[i][\sigma]$ 
18:    end if
19:  end for
20:   $\hat{d} \leftarrow \frac{\lambda}{4} \frac{\mu - \nu}{2\pi}$                     ▷ Construct effective distance matrix
21:  return  $G \leftarrow \text{cmdscale}(\hat{d})$                     ▷ CMDS to effective geometry
22: end procedure

```

Algorithm 2.1: Sensing the Effective Layout of a Tag Array.

array randomly rotates, ALS computes the PDoA from a reference tag to every other tags; (ii) The unwrapping operation removes the aliasing phase; (iii) Then ALS determines the scope of PDoA by tracking the max and min values of phase, and maps it to effective distance; (iv) The above steps are repeated for another reference tag until all pairwise effective distances are obtained, which together form a matrix; (v) Finally, Tagyro computes the effective tag layout from the effective distance matrix using the CMDS algorithm.

Algorithm 2.1 presents our real-time implementation of the ALS scheme in Tagyro. It takes the tag ID σ and phase ϕ as input, and updates the instant layout G to the user as the reader continuously queries the tag array. The algorithm can be stopped by user when the sensed array layout stabilizes.

We emphasize that a tag array's effective layout may differ from the geometrical

layout. The difference reflects the overall impact of the coupling effect. Figure 2.14 illustrates the output from a microbenchmark experiment, where we run ALS to sense a 4-tag square array. All 4 tags are physically placed on the X-Y plane, but in the effective layout, 2 of them are repelled from the plane, which account for the PDoA deviation due to coupling.

It is worth noting that ALS is a one-time procedure as long as the geometry layout of the tag array remains unchanged. If the effective layout can be measured at factory calibration time, then the ALS procedure can be skipped. In case when the tag-attached object is not suitable for rotation, the user can hold the RFID reader and scan around the object to complete the ALS.

2.4.3.3 Removing Tag's Initial Phase Offset

The PDoA comprises $\Delta\phi_{Tag}$ – the difference of two tags' initial phase offset (Sec. 2.2.2). To find out $\Delta\phi_{Tag}$, recall that the PDoA of two tags is bounded by $[\nu, \mu]$, where $\nu = -\frac{4\pi\hat{d}}{\lambda} + \Delta\phi_{Tag}$ and $\mu = \frac{4\pi\hat{d}}{\lambda} + \Delta\phi_{Tag}$ are the min and max PDoA values respectively, which are estimated by the ALS algorithm. Therefore, we can compute $\Delta\phi_{Tag}$ as follows:

$$\Delta\phi_{Tag} = \frac{\mu + \nu}{2}. \quad (2.9)$$

This offset is a constant value for a pair of tags, which will be subtracted from the measured PDoA value (modeled by Eq. (2.2)) before computing the orientation spectrum.

2.4.4 Toward 3-DoF Orientation Tracking

The non-isotropic radiation of RFID tags limits the range of orientations that Tagyro can track using a single array. To realize 3-DoF orientation sensing, we propose a setup using dual-array and dual-antenna reader. As illustrated in Figure 2.15, our basic idea

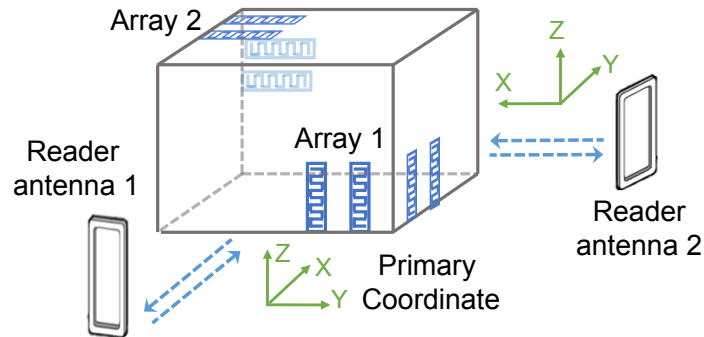


Figure 2.15: Example of dual-array, dual-antenna reader setup.

is to deploy another tag array along an orthogonal axis to the first one, so that it can compensate the blind direction, enabling full 2-DoF tracking. In addition, by deploying an additional reader antenna orthogonal to the first antenna, we can obtain total 4 DoFs (one being redundant). Compared to the 2-DoF case, 3-DoF orientation tracking requires an additional RFID antenna and tag array. Note that the dual-array can be attached not only to cuboid-shape objects, but also other arbitrary shapes (e.g., cylinder), as long as the two arrays can be arranged along perpendicular directions. We will present example use cases in Sec. 2.6.3.

Figure 2.16 illustrates Tagyro’s algorithmic modules under this new setup. During the initial setup, Tagyro needs to run the ALS algorithm for each tag array. At run time, it computes the orientation spectrum for each combination of tag array and reader-antenna, referred to as a *combo*. Meanwhile, it uses a *combo validator* to determine which combo evades the blind direction problem and provides a valid (low-noise) PDoA value. Each valid combo can track orientation along two axes. So, Tagyro runs an *orientation spectrum synthesis* scheme that unifies multiple valid combos’ sensing results to generate a 3-DoF orientation.

2.4.4.1 Combo Validator

Recall that, when a tag’s blind direction points near the reader antenna, the RSS of backscattered signals drops by more than 20 dB, and the phase values become erroneous

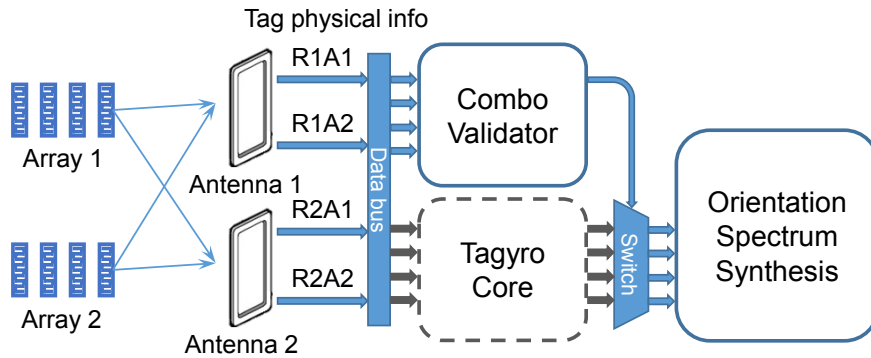


Figure 2.16: Processing blocks of the dual-array, dual-antenna reader setup for 3-DoF orientation sensing.

(Sec. 2.4.2.2). In addition, to satisfy the polarity alignment constraint (Sec. 2.4.2.1), all tags within an array have the same blind direction. So, as the arrays rotate, Tagyro's combo validator tries to exclude those combos in which the array's blind direction points to the reader antenna. RSS may seem a potential indicator – a combo that measures low RSS is likely to experience the blind direction problem. However, the absolute RSS depends on the tag-to-antenna distance, which is a random unknown factor. Therefore, it is impractical to use a fixed RSS threshold for combo validation.

Our solution builds on the observation that, since the two arrays are in orthogonal directions, for each antenna, *at most one of the two arrays is in the blind direction*. Let Ω_1 and Ω_2 denote the average RSS for tags in Array 1 and 2 respectively. By comparing their average RSS, Tagyro can determine which array is in the blind direction:

$$Blind_Direction_Array = \begin{cases} 1, & \Omega_2 - \Omega_1 > \delta \\ 2, & \Omega_1 - \Omega_2 > \delta \\ None, & \text{otherwise} \end{cases}$$

i.e., when the RSS of one array is lower than the other by δ , it is considered to be blinded. Note that the threshold δ only depends on the relative RSS between a valid and blinded array, and is independent of the tag-to-antenna distance. Based on the

empirical measurement in Figure 2.8, we set δ to a conservative value of 5 dB.

2.4.4.2 3D Orientation Spectrum Synthesis

In the dual-antenna reader setup, we take the first antenna's coordinate system as *primary*, and convert the estimation from the second antenna based on their geometrical relation. For example, in Figure 2.15, the orientation spectrum for the second antenna $I(\theta_Y, \theta_Z)$ in its own coordinate system becomes $I(\theta_X, \theta_Z)$ in the primary coordinate.

Tagyro merges the 2-DoF orientation spectrum (Eq. (2.3)) from different combos and produces a unified 3-DoF estimation. Let I_{RiAj} denote the 2-DoF spectrum of combo ij , comprised of Reader antenna i and Array j . Tagyro builds a new orientation spectrum in 3D space by summing the corresponding pixels across different combos. The synthesized spectrum for the setup in Figure 2.15 is:

$$I(\theta_X, \theta_Y, \theta_Z) = \sum_{j=1}^2 I_{R1Aj}(\theta_Y, \theta_Z) + \sum_{j=1}^2 I_{R2Aj}(\theta_X, \theta_Z).$$

Similar to the 2-DoF case (Sec. 2.4.1), Tagyro finds the 3-DoF rotation angles by searching the maximum value in the 3D orientation spectrum.

Here are a few aspects that worth further discussion:

Multipath Effects: Besides RFID tags, signals from the RFID reader may be reflected by other objects in the environment. Reflections from stationary objects do not disturb the RFID phase measurement [116], and hence will not affect Tagyro's performance. RFID reader estimates phase based on the modulated signals backscattered from the tag. Any signals reflected from static objects are unmodulated and are already removed by the reader's built-in high-pass filter.

Yet, multipath reflections from moving objects, e.g., human walking, may cause some adversary effects. In such cases, the reflected signals are dynamic, and thus may not be effectively removed by the high-pass filter. Fortunately, this adversary impact can

be effectively alleviated by increasing the tag number in the array (Figure 2.29).

Coupling Effect Caused by Surrounding Metallic Objects: The coupling effect occurs not only among RFID tags, but also between the tag and surrounding metallic materials in the environment, which we call *static* and *dynamic* coupling, respectively. The ALS algorithm is designed primarily to address the static coupling issue. The dynamic coupling cannot be effectively captured, because the metallic object does not rotate along with the tag array – when the tag array rotates, their coupling coefficients would change arbitrarily. Addressing the impact of dynamic coupling effect from surrounding metallic objects on the RFID phase measurement is still an open challenge. Fortunately, our observation (Figure 2.12) demonstrates that as long as the metallic object is not very close to the tag, e.g., $>$ the half wavelength ($\frac{\lambda}{2} \approx 16cm$), the coupling effect between them will almost vanish and will not affect the accuracy of phase measurement.

Size of Tag Array: Given the relatively large size of commodity Ultra High Frequency (UHF) RFID tags, it may not always be feasible to attach multiple tags to an object of small size. The size of tag array is primarily limited by the physical constraint, where the antenna size and separation have to match the wavelength of wireless frequency. As the RFID technology matures at higher frequency, e.g., 60 GHz, the size of tag array can be significantly reduced [127, 131], much like a finger nail. A smaller tag size will extend Tagyro towards more ubiquitous Internet-of-Things applications.

Tracking Orientation of Multiple Objects: When multiple objects exist in the same vicinity, Tagyro can track them independently given the tag IDs associated with each object. This is a huge advantage over vision-based approaches [50, 77, 148, 153, 176] which are easily confused by background and occlusion effects. Yet, the performance of Tagyro may be compromised due to the limited total query speed of the RFID reader, where the speed for each tag is reduced.

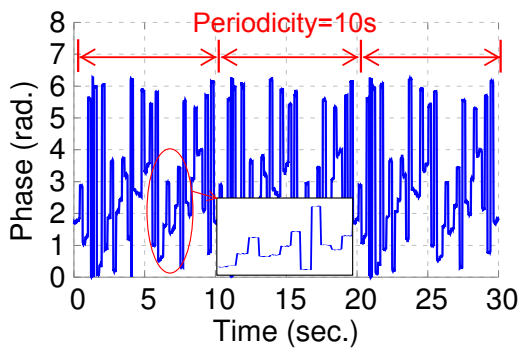


Figure 2.17: Phase jumping caused by frequency hopping.

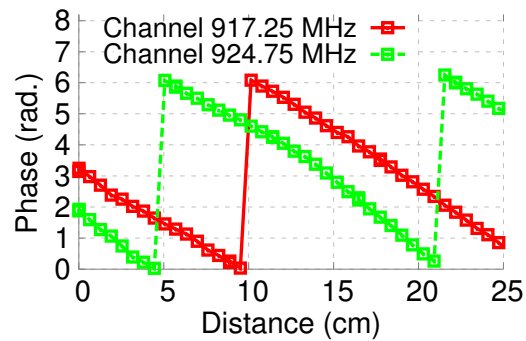


Figure 2.18: Phase over distance for a fixed frequency.

2.4.5 Calibrating Frequency-Hopping Readers

To limit co-channel interference, commercial UHF RFID readers must randomly hop to one of 50 center frequencies within the 902-928 MHz band every 200 ms, following FCC regulation [44]. Frequency hopping will cause phase discontinuity due to the phase difference of oscillator and non-uniform frequency response of the tags' antennas. To our knowledge, this issue has not been examined experimentally in prior phase-based RFID sensing systems. So we first conduct an empirical study using the Impinj R420 reader [21].

We first measure the phase of a *stationary* tag, which shows huge phase discontinuity over 30 seconds (Figure 2.17). When zooming into a short period, we can see that phase values are piecewise constant with each piece lasting around 200 ms, corresponding the sojourn time on each frequency. On the other hand, Figure 2.18 shows that if we select the phase values from a single frequency and filter out others, the phase-distance relation still follows the linear model. These experiments imply that *different frequencies induce different initial phase-offsets at the reader*. To overcome the issue, one may selectively read the phase from one frequency, but this will significantly reduce the phase reading rate, because reader needs 10 seconds to finish one round of frequency hopping.

In Tagyro, we instead calibrate the phase difference between frequencies, so that

the phase output looks like coming from a fixed frequency. The calibration is done by collecting an *initial phase measurement* that takes about 10s for an array before usage. Suppose during calibration, the tag is at distance d_0 to the reader, now the phase-distance relation in Eq. (2.1) also depends on frequency f_i , *i.e.*,

$$\phi(f_i, d_0) = \text{mod}\left(\frac{2\pi f_i d_0}{c} + \beta_i, 2\pi\right) \quad (2.10)$$

where c is light speed and β_i denotes the phase offset introduced at frequency f_i .

Given the initial phase measurement $\phi(f_i, d_0), \forall i \in [1, 50]$, we map the phase $\phi(f_i, d)$ for arbitrary distance d and frequency f_i to a common frequency f_r (default to 915.25 MHz). Substituting $\phi(f_i, d_0), \phi(f_r, d_0), \phi(f_i, d)$, and $\phi(f_r, d)$ into Eq. (2.10) and canceling the terms β_i and β_r , we have:

$$\phi(f_r, d) = \text{mod}\left([\phi(f_i, d) - \phi(f_i, d_0)] \frac{f_r}{f_i} + \phi(f_r, d_0), 2\pi\right). \quad (2.11)$$

Note that d and d_0 are not required to be known, because the phase value $\phi(f, d)$ is always read directly.

2.4.6 Dealing with Asynchronous Phase Reading

To compute the orientation of a tag array, the reader needs a synchronous phase snapshot of all tags. However, the Electronic Product Code (EPC) Gen2 RFID standard [140] adopts a Framed Slotted Aloha protocol which reads the tags asynchronously. To address this issue, Tagyro assumes the array's rotation speed remains similar over consecutive phase readings (typically a few tens of *ms*), and uses an interpolation method, which creates an instant phase snapshot for all tags at an interval of every T phase readings. In Tagyro, we set T to the number of tags in array. Let $\phi(t_i)$ denote the phase of a tag at time t_i . Tagyro computes the snapshot phase at time stamp t , given historical

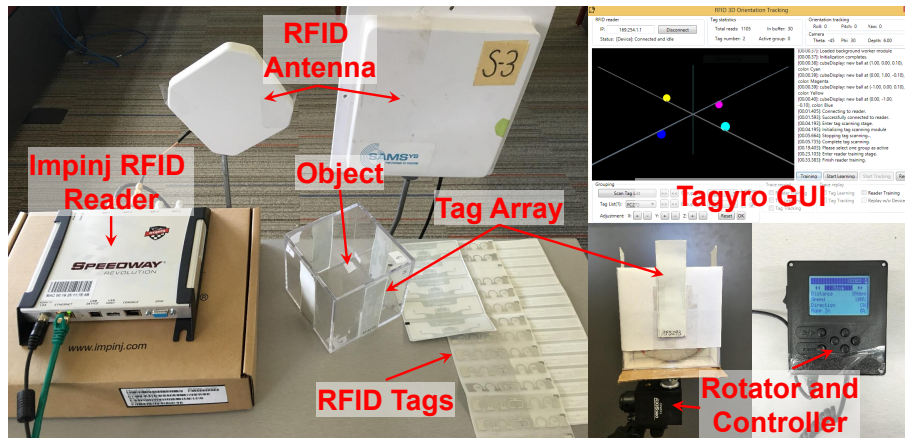


Figure 2.19: Tagyro implementation and software interface.

phase readings at time t_i and t_{i-1} :

$$\phi(t) = \phi(t_{i-1}) + (\phi(t_i) - \phi(t_{i-1})) \frac{t - t_{i-1}}{t_i - t_{i-1}}. \quad (2.12)$$

Note out that since the interpolated phase value will be eventually updated by new phase measurement, the error of the interpolation result will not accumulate or propagate over time.

2.5 Implementation

Hardware: We implement a prototype of Tagyro using the COTS RFID reader Impinj R420 [21], which is FCC-compliant and performs frequency hopping within $902.75 \sim 927.25$ MHz. Figure 2.19 illustrates our hardware components. Two circular polarized antennas are connected to the reader via its two RF ports. Their antenna gain and 3 dB-beamwidth are 5.5 dBi, 7 dBi and 100° and 80° [10, 19], respectively. We assemble the tag array using three different types of RFID tags: ALN-9740 [1], SMARTRAC DogBone [3], and SMARTRAC ShortDipole [6]. They are common label tag in flat, thin shape and commercially available [4]. Their maximum reading distance ranges from 15 ft to 20 ft. The reader interrogates RFID tags and sends query reports, containing information

of ID, RSS, phase, time stamp and channel, via Ethernet to a host PC that runs Tagyro.

Software: Tagyro’s software implementation contains three major modules: 3D GUI, RFID library and processing algorithms. The GUI module is developed in C#. It displays the tag array’s effective layout during initial setup, and instantaneous orientation during run-time tracking. The RFID library controls the reader using a Low-Level Reader Protocol (LLRP) ratified by EPCglobal. We implement Tagyro’s key processing algorithms (Sec. 2.4) in MATLAB, and then recompile them into standalone C libraries (DLL) using the Code Generation toolbox. The C libraries are imported into C# program and being invoked in real-time. To start, we first run the frequency calibration module for each array. Then we randomly rotate the tag array when Tagyro performs ALS to sense the tag array’s effective layout. Afterwards, Tagyro can begin its orientation tracking.

2.6 Evaluation

Methodology: We conduct experiments in an office environment to evaluate the effectiveness of Tagyro’s design components, accuracy of orientation sensing, as well as run-time latency. Figure 2.15 illustrates our experiments’ default setup, where 4 Dog-Bone tags are attached to a plastic cube forming a square tag array, separated roughly by quarter-wavelength (*i.e.*, 8.2 cm). Note that since our ALS algorithm can automatically sense the tag layout, it is not required to deploy the RFID tags with precise separation. Two reader antennas are placed 1 m away from the tag array at mutually perpendicular directions, both 70 cm above the ground. By default, the reader uses its maximum transmit power at 32 dBm.

Since we are unaware of any RFID-based system that tracks 3D orientation of an object, we evaluate accuracy of Tagyro against ground truth measurement. To obtain the ground-truth orientation, we mount the tag array (attached to an object) on a mechanical rotator [2] (Figure 2.19), which is controlled by the PC host to rotate at full

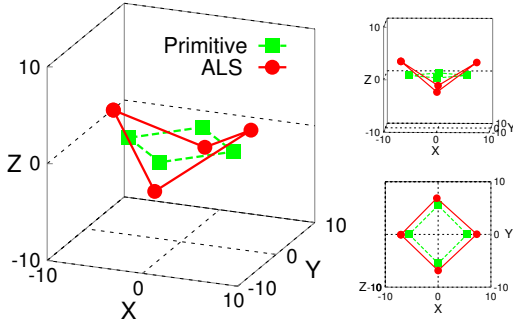


Figure 2.20: Effective and physical layout of the 4-tag array.

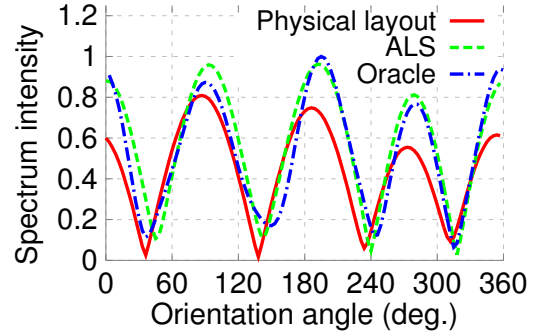


Figure 2.21: Normalized orientation spectrum for 4-tag array.

360-degree at maximum speed 1.71 rad/s with minimum granularity of 1° . We first rotate the tag array to a known orientation, and use it as the initial state. Then, the offset between the motor’s rotation angles and initial state is used as the ground-truth orientation. Since the motor can only rotate in one DoF, we manually change the pose of the tag array, so that different DoFs can be aligned to the motor’s rotation axis.

Metrics: To characterize Tagyro’s accuracy, we mainly focus on angular error, *i.e.*, the deviation of estimated orientation angle from the ground truth. Since Tagyro mainly relies on the PDoA information to pinpoint the tag orientation, in the micro-benchmark experiments we also use the phase error to quantify the effectiveness of our design.

2.6.1 Micro Benchmarks

Effectiveness of sensing effective layout. Recall a tag array’s effective layout sensed by ALS may differ from the physical geometry, due to tag coupling. Figure 2.20 illustrates the layout of a 4-tag square array for cases with and without the coupling effect. Note that in the effective layout, two tags deviate from the Z axis in order to virtually compensate the coupling effect. Yet, obtaining the ground truth of coupling effect requires the knowledge of tag’s physical parameters, *e.g.* coupling coefficient, which are not easily accessible in practice. On the other hand, the ALS accuracy will directly affect

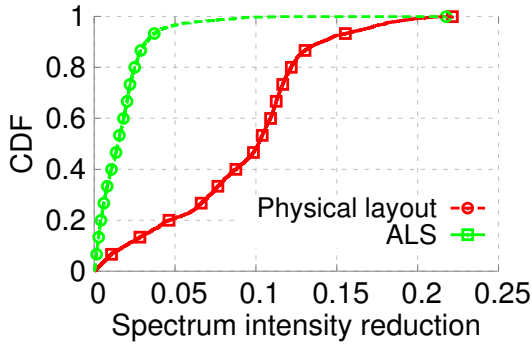


Figure 2.22: The CDF of spectrum intensity difference compared to the oracle.

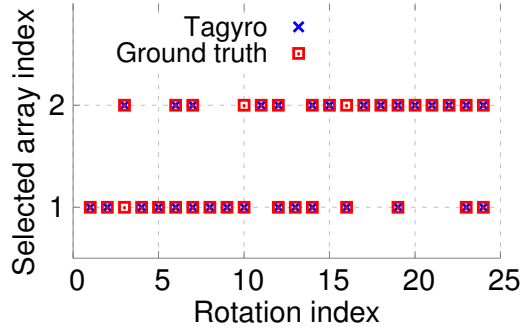


Figure 2.23: Accuracy of selecting the array in valid reading zone.

the ultimate performance of orientation tracking. Thus, we leave the evaluation of its effectiveness in the system level tests (Sec. 2.6.2).

To validate whether ALS can indeed counteract the coupling effect, we compute the orientation spectrum with effective layout. For contrast, we also simulate the oracle case using physical layout in a non-coupling scenario, which produces the ideal orientation spectrum. Generally, higher peak intensity means better spectrum (with higher contrast), because tracking will be more resilient to noise and spatial ambiguity. Figure 2.21 plots a snapshot spectrum in one dimension for θ_z by fixing θ_x and θ_y to zero. When the array is orientated at $(0, 0, 195^\circ)$, the orientation spectrum from ALS has similar peak intensity as oracle. Compared with the case when the physical layout is used directly, ALS significantly improves the contrast of spectrum. Figure 2.22 statistically plots the difference of highest intensity peak in spectrum relative to the oracle. The ALS algorithm reduces the intensity difference by $5.5\times$ compared with using physical layout.

Accuracy of selecting the valid combo. In 3D orientation tracking, Tagyro’s combo validator is critical in dealing with the blind direction issue. We verify this mechanism by rotating two arrays together in 3D space, and inspect if Tagyro can select the correct tag array. We determine the ground truth by checking if an array’s blind direction is within 40° pointing to the reader antenna.

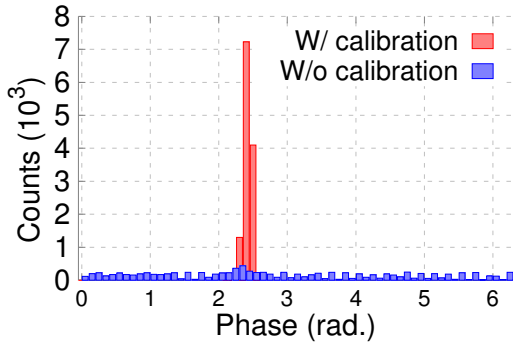


Figure 2.24: Histogram of phase distribution for frequency hopping calibration.

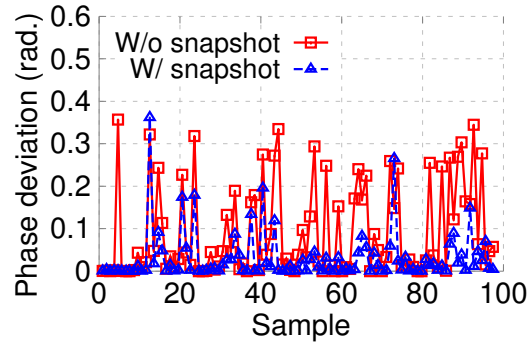


Figure 2.25: Residual phase error caused by asynchronous phase reading.

The experiment results in Figure 2.23 show that the array indices selected by Tagyro all match the ground truth. Since occasional blockage can also cause RSS offset between the two arrays to exceed the conservative threshold $\delta = 5dB$ (Sec. 2.4.4), some array indices in the ground truth may not be identified by Tagyro. However, this will not cause error because Tagyro can still use the other array to compute the orientation.

Channel hopping calibration. We next evaluate the impact of reader’s channel hopping on the phase value by measuring the phase of a *static* tag for 2 minutes. Figure 2.24 plots the histogram of phase values. Without calibration, phases are almost evenly distributed over all possible values with a *std.* of 1.78-radian. However, the *std.* is reduced by $46\times$, down to 0.039-radian after applying the calibration, which is very close to the noise floor 0.025-radian.

Effectiveness of the phase snapshot algorithm. Tags in an array are queried asynchronously by the reader, which makes the timestamps of phase reading misaligned. To evaluate this impact, we first rotate the tag array over 100 preset orientations and measure the phase values. Since the tag array is stationary during each measurement, the phase readings are consistent and used as ground-truth. Then we measure the phase while continuously rotating the array through these preset orientations. We observe that the phase error is substantially reduced after applying the phase snapshot algorithm (Figure 2.25). Mean error drops from 0.084 radian to 0.026 radian. We believe

Table 2.1: Query frequency of each tag in the array (SPS).

Tag #	4	6	8
Avg. spd.	43.1	40.0	37.7
10% spd.	32.1	26.8	25.1
90% spd.	50.0	47.9	46.6

Table 2.2: Efficiency of the tracking algorithm (FPS).

Tag #	4	6	8
1-DoF	4618	3653	3423
2-DoF	205	145	121
3-DoF	111	77	63

the reduction can be larger for a higher rotation speed. Note that the error can occasionally be magnified by the phase snapshot due to unexpected phase deviation, *e.g.* noise. But they will have marginal impact on orientation tracking since the occasionally magnified error does not accumulate or propagate over time.

Interrogating speed of tag array. For real-time orientation tracking, one key question is: how fast can the RFID reader query the tag array? The query speed determines how frequent Tagyro can update the orientation estimation. We evaluate the query speed by interrogating different sizes of tag arrays. The reader uses its default *AutoSet Dense* mode, which automatically adjusts its MAC-layer parameters to maximize the tag query speed. Table 2.1 summarizes the CDF of update frequency for a two-minute measurement when the array rotates at the maximum speed. Although the update frequency slightly decreases as we increase the array size, on average Tagyro can achieve 37.7 samples per second (SPS) even for the 8-tag array.

Efficiency of tracking algorithm. We further examine the real-time processing speed of our Tagyro implementation. We first collect a 2-minute trace data of the tag array’s phase and then replay the trace in our Tagyro program. In this way, we can evaluate the efficiency of our orientation tracking algorithm independently of the reader’s query speed. We run Tagyro on a i7-4770 PC utilizing one CPU core. Table 2.2 plots the frame rate (number of orientation sample outputs per second). Tagyro can maintain an average processing rate of 3500 FPS, 144.8 FPS and 76.7 FPS, for 1-DoF, 2-DoF and 3-DoF orientation tracking, respectively. The main bottleneck of Tagyro’s tracking efficiency lies in the RFID reader’s query speed (Table 2.1) which is much lower than




Figure 2.26: Orientation accuracy vs DoF. Error bar shows the 90th error.




Figure 2.27: Estimation and ground-truth for 1-DoF rotation.

the processing speed of Tagyro.

2.6.2 System Level Tests

We proceed to evaluate Tagyro’s overall performance in tracking 3D orientation under various practical settings.

Accuracy of 3D orientation tracking. We first evaluate the accuracy of orientation tracking under different DoFs: (i) *1-DoF*: tracking Z-axis rotation using a single tag array and single-antenna reader; (ii) *2-DoF*: tracking Y and Z axis rotations using dual-array and single-antenna reader; (iii) *3-DoF*: using the dual-array, dual-antenna reader setup as in Figure 2.15. Our experimental results (Figure 2.26) show that for both 1-DoF and 2-DoF, the average (90-percentile) tracking error is about 4° (8°). 2-DoF case has a similar accuracy with the 1-DoF since each of the two arrays independently tracks one DoF. Notably, the average orientation error of Z-axis is much smaller (2.2°) in 3-DoF case, because both of the reader’s antennas are outside the blind direction. Tagyro’s orientation spectrum synthesis can effectively unify the extra spatial diversity to produce a better estimate. Figure 2.27 plots 5 overlapped example tracking results for 360° rotation in 1-DoF. The figure depicts the orientation estimation of a tag array rotating at a constant speed and the ground truth value over the time, which shows that the estimated angle is highly consistent with the ground truth at every time stamp.

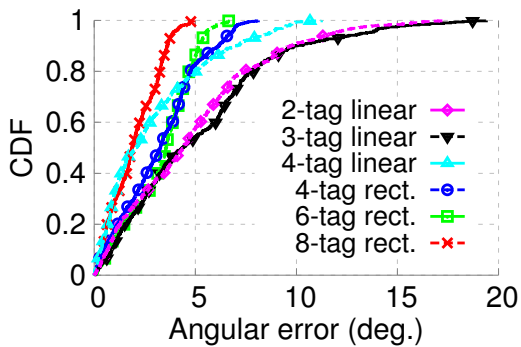


Figure 2.28: The CDF of orientation error vs. size of tag array.

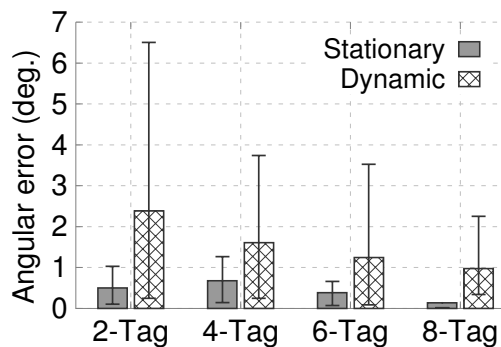


Figure 2.29: Orientation error under surrounding human activity. Error bar shows the 90th error.

Impact of array size. We next vary the number of tags in array from 2 to 8 and evaluate how it affects the accuracy of orientation tracking. Figure 2.2 shows the tag arrays' physical layouts. Figure 2.28 plots the CDF of angular error in 1-DoF tracking, which shows that *a larger tag array can effectively improve accuracy*. The average error is reduced from 4.80° for 2-tag array to 1.98° for 8-tag array. Moreover, the 90-percentile error is also reduced from 9.71° to 3.60° . The improvement gently diminishes as the size of tag array increases. In practice, an array of 4~6 tags may make a good balance between accuracy vs. form factor. Using high-frequency tags (*e.g.*, millimeter-wave) may significantly reduce the form factor due to shorter wave length and smaller antenna size. But this is beyond the scope of our work.

Impact of environmental dynamics. Like all other phase-based RFID sensing systems [105, 116, 181, 200], estimated phase value in Tagyro is subject to the channel noise and influenced by multipath environment. To evaluate their impacts, we measure the jitter of estimated orientation, while fixing the tags so that any jitter can be only caused by RF channel fluctuation. Figure 2.29 compares the angular jitter in a stationary environment with that where a human randomly walks around the reader and tags. The average jitter caused by the channel noise itself is small and typically less than 1° . On the other hand, the estimated orientation can deviate by 6.5° when there is significant human movement. This is mainly attributed to the intentionally created variation




Figure 2.30: Orientation error over tag-to-antenna distance.

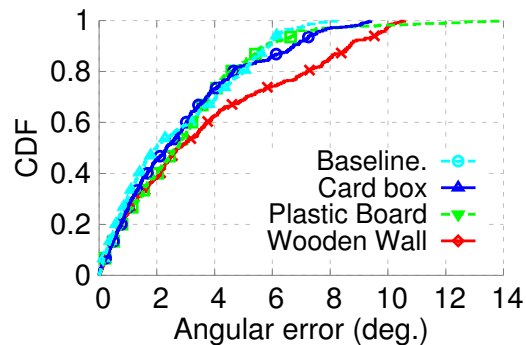


Figure 2.31: The CDF of orientation error under blockage.

of multipath reflections. Fortunately, a larger tag array can make Tagyro more resilient to environmental dynamics — with an 8-tag array, the angular jitter is reduced to about 1° .

Impact of tag-to-antenna distance. Typical RFID systems have a maximum reading range of 4~6 m [183]. Longer distance results in weaker signals, hence more noisy phase reading. To understand its impact on the orientation accuracy, we vary the tag-to-antenna distance from 1 m to 2.75 m. Figure 2.30 shows that the average angular error gently increases from 2.09° to 9.32° , with a maximum working distance about 3m in our setup. A closer examination reveals that the RSS drops significantly over distance. The lowest RSS across rotation angles decreases by more than 21 dBm, from -49 dBm at 1m to -70 dBm at 2 m. Since the sensitivity of the Impinj R420 reader is -82 dBm [21], its measured phase value starts to deviate from the phase-distance model (Sec. 2.2.1) when distance goes beyond 2 m. On the other hand, backscattered data bits can be decoded at a further distance because RFID adopts Amplitude-Shift Keying (ASK) modulation that is insensitive to the phase deviation. By deploying more RFID antennas, Tagyro can potentially cover a wider working area, which we will leave for our future work.

Under blockage. Tagyro’s usage scenarios may involve Non-Line-of-Sight (NLOS) environment, *e.g.*, in a supply chain, where the object (and tag array) may be placed




Figure 2.32: Latency of rotation response.



Figure 2.33: Eight tags are attached to the surface of a 12-side dice.

inside a card box (Figure 2.1). We mimic such NLOS cases by blocking the Line-of-Sight (LOS) between the tag array and reader using different materials: card box, plastic board and wooden wall, with thickness of 5 mm, 10 mm and 30 mm, respectively. Figure 2.31 shows that blockage has marginal impact on Tagyro's orientation tracking. The wooden wall slightly reduces the accuracy because its thickness attenuates the backscattered signal strength, similarly to increasing tag-to-reader distance. The experiment verifies that Tagyro can track orientation in NLOS scenarios where camera-based approaches [147] will fail.

Latency of rotation response. In the implementation, we buffer multiple latest phase values for algorithm processing and adopt a Kalman filter to reduce noise in orientation estimation. They will incur response latency in the tracking output. To evaluate the latency, we rotate the tag array by 90° and simultaneously measure the time it takes for the estimated orientation to be less than 5° close to the ground truth. Figure 2.32 plots the CDF of response latency from 50 measurements. The average values of response latency are 0.61s and 0.71s for 4-tag and 6-tag arrays respectively, which are sufficient for most application scenarios that do not have stringent timing constraint. Besides, the latency can be further reduced by using high-speed RFID readers.




Figure 2.34: Detecting which side of the dice faces up.

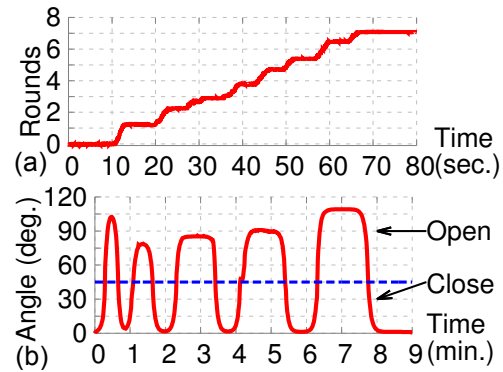


Figure 2.35: Detecting (a) usage of roll tissue and (b) open/close status of the door.

2.6.3 Case Study

In this section, we apply Tagyro in two case studies: a gaming scenario where Tagyro detects the orientation of a 12-side dice, and a smart home scenario where Tagyro monitors the rotation of the door and consumption of roll tissue.

Tracking a 12-side dice. To detect which side of a ball-like dice faces up, we need at least two DoFs – the horizontal rotation does not matter because it will not change the face-up number. Thus, we setup one reader in the ceiling, which projects signals toward the ground and attach two tag arrays to the dice in mutually perpendicular directions (Figure 2.33). We randomly toss the dice and record the estimated and ground-truth number that faces up. Figure 2.34 shows that Tagyro correctly identifies 33 out of 36 numbers, achieving 91.7% accuracy. The tracking accuracy of the 12-side dice is lower than the 3D case (Figure 2.26) due to two factors. *First*, the shape of dodecahedron does not allow us to attach tags completely following the constraint of polarity alignment (Section 4.2.1), which reduces the accuracy of combo validation (Section 4.4.1). *Second*, the two closely-placed tag arrays will create an inter-array coupling effect. Tagyro performs ALS independently for each array, which handles the intra-array coupling effect. Yet, it is unable to sense the inter-array coupling effect. On the other hand, Tagyro cannot perform a joint ALS for both arrays because their polarity

directions are orthogonal. One approach to mitigate the inter-array coupling effect is to separate the two arrays by a distance of greater than half-wavelength (Figure 2.12), such as attaching tags to different edges of the box (Figure 2.15).

Monitoring smart home. In the smart home scenario, Tagyro can monitor the status of daily objects, *e.g.*, whether the door is open or close, which direction a rolling chair (and a seated human) is facing, and how much roll tissue has been consumed. Many such problems can be addressed by sensing the rotation/orientation of the passive objects. Here we showcase the door and roll tissue scenarios. Since both of them have fixed rotation axes, (*i.e.*, rotation is restricted to 1-DoF), we attach a 4-tag array on each (Figure 2.1), and use a single-antenna reader as monitor. Figure 2.35(a) plots the accumulated rounds of rotations for a roll tissue, which unveils fine-grained information of the consumption over time. Figure 2.35(b) shows tracking result for a door, which measures not only its opening angle, but also the time duration when it stays open/close. These case studies also demonstrate Tagyro can be applied to a wide range of objects with different shapes.

Though these detection tasks may be also addressed by simple mechanical or IR sensor, Tagyro has several unique advantages over them. *(i)* Fine-grained rotation trajectory. Tagyro not only tells if the door opens or closes, but also the degree of opening, which is useful for applications, such as to determine if a robot or cargo can go through the door. *(ii)* Ad-hoc installation. Tagyro can instantly enable orientation tracking by attaching multiple tags to the target object, which does not involve heavy mechanical installation or altering the physical structure. *(iii)* Maintenance free. The batteryless RFID tags are continuously power by the RFID reader without the need for battery replacement.

2.7 Discussion

This chapter presents Tagyro, an RFID-based system to track the orientation of a batteryless object. By forming an array using multiple tags, Tagyro exploits their phase difference of arrival to find out the object orientation. Novel processing algorithms are proposed to address challenging issues such as coupling effect, which are raised in the tag-array RFID system. Tagyro is built upon COTS RFID reader and tags. We conduct system-level experiments and two case studies, which show that Tagyro can achieve high accuracy in orientation tracking. Tagyro represents the first study of an RFID-based 3D orientation tracking system, which can be applied to many practical scenarios including but not limit to smart home and mobile gaming.

CHAPTER 3

Fine-grained Passive Motion Tracking Using mmWave Radios

Radio-based passive object sensing is a rapidly developing technology that detects the motion of human body parts or associated objects through wireless signals. Compared with conventional vision-based approaches [174], e.g., LeapMotion [110] and Kinect [5], it is less intrusive and unaffected by ambient light conditions or the sunlight interference. Catalyzed by the proliferation of mobile devices, this technology holds potential to enable new ubiquitous user-mobile interfaces and spur a wide range of applications. The latest passive sensing technology can already accurately distinguish a prescribed set of body/limb movement [130] based on Doppler patterns or signal strength variation. Hand gesture sensing [111] is also achievable by training a pattern-matching model. To satisfy a broader range of applications, however, two design challenges remain open: fine resolution and unsupervised motion tracking.

We envision one such application scenario: tracking a writing object (e.g., a stylus pen or marker) wirelessly. Conceptually, empowered by multiple wireless devices, we can create an interactive trackpad on any conventional surface, e.g., whiteboard or desktop. This way, we can deploy large touch screens in a more flexible and economic way than traditional graphic tablets. This vision entails tracking a small writing object with at least sub-centimeter precision. Such precision has proven feasible in certain active RF sensing systems like Tagoram [200]. However, they are applicable only to

objects instrumented with an RFID tag or radio receiver.

This chapter introduces mTrack, which leverages 60 GHz millimeter wave (mm-Wave) radios to track the trajectory of a writing object at high precision. 60 GHz radios are standardized in IEEE 802.11ad [11], and anticipated to penetrate one third of wireless links by 2018 [53]. Commercial 802.11ad-capable smartphones have already been demonstrated [126]. Adopting mmWave instead of conventional microwave band (2.4 GHz or 5 GHz) brings multifold advantages. A shorter wavelength can create stronger reflections from small objects (*e.g.*, a pen), since wireless signals cannot easily bypass objects larger than wavelength. More importantly, mmWave devices like 802.11ad allow for miniaturized phase-array with dozens of antenna elements, which together create highly-directional “pencil-beams”. Such directional beams are electronically steerable, thereby creating a new dimension for object localization/tracking.

mTrack’s design harnesses the unique advantages of mmWave. It uses one Transmitter (Tx) to emit 60 GHz signals and illuminate a handheld object that roams on a trackpad area. Part of the signals are diffused by the object’s surface and captured by two directional receivers. mTrack leverages the Received Signal Strength (RSS) and phase (relative to Tx) to localize and track the object. Simply put, by measuring the arriving direction of reflected signals, each Receiver (Rx) can estimate the object’s relative angle, and pinpoint its initial *location* on the trackpad. We refer to this as *anchor point acquisition* (APA). When the object moves, the reflected signals’ path length varies, which alters the phase, providing salient hints for *tracking* the object with sub-wavelength resolution. Our feasibility study using a 60 GHz software-radio (Section 3.2) verifies these principles, and unveils the unique advantage of mmWave passive tracking over its microwave counterpart, in terms of potential precision, sensitivity to target size, *etc.* Yet in practical environment, these principles entail unique challenges, which we propose to solve in mTrack.

First, reflected signals from irrelevant background objects can severely distort the

target-reflected signal and affect phase-tracking accuracy. Prior passive tracking techniques use ultra-wideband [28], or measure and subtract static background reflection directly [62]. In contrast, mTrack’s mmWave uses single-carrier phase-tracking, and needs to handle background dynamics. To meet this challenge, we develop two algorithms: *dual-differential background removal*, and *phase counting and regeneration*, to recover the legitimate phase change induced by target movement only. These two techniques take unique advantage of short wavelength of mmWave signals, and stay at different vantage points when considering a tradeoff between tracking precision and resilience to the inherent phase noise in 60 GHz radios.

Second, although 60 GHz signals are commonly simplified as pseudo-optical [24], our measurement shows that they can be widely diffused after hitting a small object. Thus, one cannot migrate the specular reflection effect in laser/infrared tracking systems [8] for 60 GHz tracking. In mTrack, we observe that 60 GHz antenna response bears a roll-off response pattern, and reflected RSS is strongest when the object direction matches the peak response. mTrack’s APA mechanism thus leverages steerable 60 GHz Rx antenna to find the matching point, and identify the object’s relative angle. Since practical 60 GHz phase-array antennas can only switch between a discrete set of directions, mTrack reconstructs the ideal continuous scanning results from discrete sampling. Consequently, it can work even if the switching angle is much larger than beamwidth.

We have implemented mTrack on a custom-built 60 GHz software-radio platform. We evaluate mTrack’s APA and phase-tracking performance when user is navigating a pen over a $50\text{cm} \times 50\text{cm}$ virtual trackpad region. Experimental results show that mTrack can localize the pen’s angular position with error less than 1° , and track its motion trajectory with only 6 mm of median error and 8 mm of 90-percentile error. We also found that mTrack can be readily used to enable touch event detection (*i.e.*, pen clicking/leaving the trackpad), owing to the constrained beam pattern of mmWave

antennas. Using a simple feature-based detection algorithm, it achieves a detection accuracy of around 94%.

The main contributions of mTrack include the following:

(i) A feasibility study of fine-grained, sub-centimeter scale object localization/tracking using mmWave radios with steerable antennas, in contrast to microwave radios (Sec. 3.2).

(ii) A phase-based approach to track small objects like pens to high precision, thus realizing trackpad applications. The tracking scheme builds on two novel algorithms to counteract the impact of background interference (Sec. 3.4).

(iii) A localization mechanism that leverages 60 GHz antennas to pinpoint the object's initial location and complement phase-tracking via opportunistic calibration (Sec. 3.5).

(iv) Implementing mTrack on a reconfigurable 60 GHz radio testbed, and validating its localization/tracking performance in a practical wireless trackpad setup (Sec. 3.7).

3.1 Related Work

RF-based Active Tracking. Localization of radio-equipped objects has been explored extensively. RSS alone [177], or combined with phase [150], can serve as location signatures. But mapping channel information to location usually requires site survey and the accuracy is time and environment dependent [103]. In short-range, static environment, it is possible to locate a radio using a path-loss model (4cm error in an $1m^2$ area) [64]. However, for passive tracking, RSS-model becomes inaccurate due to multipath effects. Thus, mTrack only resorts to the variation of phase to realize passive tracking, which does not rely on any RSS-distance model.

Centimeter-scale RFID localization was achieved recently in RF-IDraw [181], which uses an interferometry technique to measure the relative phase between multiple RFID

readers. Tagoram [200] generates a phase hologram that statistically maps measured phase to a potential position, and computes moving trajectory through phase shifting. These tracking schemes require an RFID tag on the object, which is not readily available for writing objects in daily use. To enable high precision passive tracking, mTrack faces two new technical challenges unseen in [181,200]. First, reflected signal from irrelevant background objects can severely distort the target-reflected signal and affect phase-tracking accuracy (more details in Section 3.4.2.2). Second, current 60 GHz radio hardware still has non-negligible phase noise, which contaminates the phase shifting. To meet these challenges, mTrack incorporates two novel processing techniques that recover legitimate phase change (Section 3.4). In addition, mTrack leverages the unique feature of beam steerability in 60 GHz radios to realize precise localization.

RF-based Passive Tracking. Passive object tracking is reminiscent of the vast literature in radar systems [28, 113]. Conventional radar, however, mainly focused on tracking large moving objects using Doppler methods, and pushing the granularity *using wideband radios with GHz of sampling rate*. mTrack, in contrast, leverages single-carrier phased-based approach, taking advantage of electronically steerable 60 GHz antennas to track small objects with sub-centimeter precision, thus enabling near-field interactive applications like a trackpad. On the other hand, typical ranging radars require dedicated hardware, which is not readily available in 60 GHz communication system. Pulse radar [121] needs high-speed pulse generator while Frequency-Modulation Continuous Wave (FMCW) radar [28] requires a swiping carrier frequency controlled by Voltage-Controlled Oscillator (VCO). On the contrary, mTrack’s single-carrier design can be easily realized using 60 GHz communication hardware. Radar object tracking also faces interferences from background reflection. However, mTrack’s mmWave phase-tracking method encounters unique challenges that have not been addressed in existing radar tracking algorithms [28, 62, 198, 205] (more details in Sec. 3.4.2.2).

Radio tomography and imaging techniques [125, 184] deploy a mesh of sensors

Related Systems	Tagoram	RF-IDraw	Tomography	WiVi	WiTrack	mTrack
Work	[200]	[181]	[125, 198]	[29]	[28]	N/A
Object Type	Active	Active	Passive	Passive	Passive	Passive
Signal Feature	Phase	Phase	RSS	Phase	RSS	Phase/RSS
Methodology	Hologram	Interferometry	FP/AoA	ISAR	FMCW radar	PS/BS
Track Static Object	Yes	Yes	Yes	No	No	Yes
# of Tx&Rx	≥ 2	8	10 ~ 100	≥ 4	4	≥ 3
Tracking Range	1 ~ 10m	2 ~ 5m	3m ~ 20m	2 ~ 7m	3 ~ 9m	1m
Granularity	14mm	49mm	80cm ~ 3m	1m	30cm	8mm

Table 3.1: Comparison of different RF-based tracking systems. FP, PS and BS stand for fingerprint, phase shift and beam-steering. Data are from reported figures of cited works.

around sensing area and locate a person by identifying shadowed area that shows weak RSS or outstanding variation. mTrack is based on object reflection instead of blockage effect, and achieves a fine-grained tracking with much fewer sensing nodes. WiFi imaging (e.g., [80]) creates an image of sensing area through an antenna array, though with a low resolution. WiFi RSS and Doppler metric have also been leveraged in gesture recognition [111, 130], but the identification algorithm needs to be trained with known patterns. Besides tracking, millimeter wave has been used to infer object’s surface curvature and material [216], which can be applied to identify the target of interest.

To our knowledge, mTrack represents the first work that achieves sub-centimeter scale passive object tracking, taking advantage of the small wavelength and steerable antennas of mmWave radios. Table 3.1 compares mTrack with other recent RF-based active and passive tracking systems.

3.2 Understanding mmWave Passive Tracking

We first use our 60 GHz testbed to explore mmWave characteristics pertaining to high-precision passive object localization/tracking. Our experiments examine whether small writing objects can effectively reflect mmWave signals, and whether the reflected RSS and phase can serve as subtle location/motion hints. The experiments also reveal




Figure 3.1: Setup of feasibility test. (a) Reflection test (Sec. 3.2.1): Rx centrally rotates around the target. (b) Tracking test (Sec. 3.2.2): Rx points to a moving target. (c) Locating test (Sec. 3.2.3): Rx steers its beam toward various directions.

unique advantages of mmWave over 2.4 GHz microwave-band signals.

For a feasibility test, we create a simplified setup illustrated in Figure 3.1. Our testbed implementation is detailed in Section 3.7. By default, the receiver antenna is placed at view angle 0° co-located with the transmitter. Irrelevant reflections from background objects are reduced by placing RF absorbers [7] (with approximately 35 dB attenuation) near the boundary of testing region.

3.2.1 Reflection/Diffusion of Signals by Objects

To enable mmWave passive tracking, receiver must receive reflection signal from the target despite the view angle *w.r.t.* target. However, 60 GHz signals are often deemed to possess a pseudo-optical property. So, will they create mirror-like specular reflections when hitting glossy objects? We examine the reflective property by rotating the receiver's view angle around the location of the target – a stylus pen (0.4 cm radius) with a metal surface.




Figure 3.2: Small object (pen) causes diffusive 60 GHz signals that can be captured over different view angles.

Figure 3.2 plots the average RSS of 100 measured values at each view angle (*with* and *without* are abbreviated as *w* and *w/*). In the presence of the pen, the RSS remains consistently high from view angle 0° to 150° , indicating that *when reflecting 60 GHz signals, the target acts like a quasi-omni-directional antenna, rather than a mirror. Thus, even a highly directional 60 GHz receiver can capture the reflection from a wide range of view angles, as long as the target is illuminated by the transmitter.*

We also make two other observations. First, the transmitter’s beamwidth determines the illumination coverage. 90° to 180° beamwidth can ensure the target is illuminated as it moves across a wide region in front of the transmitter. Second, even without target, the RSS varies, partly due to residual reflections from background, and partly due to leakage signal from the transmitter especially at a wide view angle.

3.2.2 Phase Variation Enables Fine-Grained Motion Tracking

Moving distance. Under ideal propagation/reflection, the phase offset between transmitter and receiver should only depend on the signal path length (illumination plus reflection). Figure 3.3(a) shows the phase variation as the target moves at 24 mm/second away from the transmitter/receiver. Owing to extremely short wavelength, *i.e.*, around 5 mm, even a small change of total signal path length can vary the phase significantly. By unwrapping [27] the phase value, we observe an increasing trend, indicating increas-




Figure 3.3: Tracking under same 30° beamwidth (diameter: pen 1.2cm, bottle 5.8cm and can 10.3cm). (a) Small object (pen) can cause $[-\pi, \pi]$ phase variation in mmWave. (b) Microwave needs an object (e.g. can) of size $9\times$ larger.

ing path length, a consequence of the target’s actual movement pattern. In addition, we see that the unwrapped phase value increases by 230 radians, translating to a distance $\frac{230}{2\pi} \times 0.25\text{cm} \approx 9.15\text{cm}$, roughly matching the actual moving distance of 10 cm (Note that each phase cycle corresponds to both illuminating and reflecting paths). In effect, when zooming in one phase cycle, we can see that even sub-wavelength movement resolution is feasible.

These two observations hint that *phase change of reflected signals can indicate whether the target is moving towards/against the Tx/Rx, and the relative moving distance. Small wavelength of 60 GHz signals enables fine-grained distance resolution.* However, two challenges remain open: (i) A single receiver cannot resolve the target’s moving angle within a 2D plane. (ii) Phase has an inherent aliasing effect, and cannot indicate the absolute location of target. This can be clearly observed from Figure 3.3(a), where phase exhibits cyclic behavior, with many locations sharing the same phase value.

It is worthy to note that the phase value experiences a jump between 520ms and 540ms in Figure 3.3(a). This is caused by the signals from background reflection — the RF absorber in our feasibility test is not large enough to isolate all background objects. Our phase processing techniques (Section 3.4) will remove such abnormality and guarantee the correctness of accumulated phase.

Target size. Electromagnetic waves can easily bypass around obstacles whose size




Figure 3.4: Beam steering to locate target direction. Receiver gets strong RSS when steering toward the target.

is smaller than the wavelength, and will be blocked/reflected otherwise. Thus, we hypothesize that *passive tracking of small targets is a unique advantage of mmWave signals over microwave*. To verify the hypothesis, we place a pair of 2.4 GHz software radios following the same setup as above. Figure 3.3(b) plots the resulting phase variation as the target moves. For the stylus pen, the signal phase can hardly reach a full cycle of $-\pi$ to π , and shows no clear correlation with wavelength. This is because the majority of the microwave signals bypass around the pen, with a radius much smaller than the wavelength (12.5 cm). Hence, background reflections from afar tend to dominate and distort the phase. Only when the target (*e.g.*, a can) size is sufficiently large can the reflected signals dominate and manifest full cyclic behavior as in the mmWave case. This result confirms our hypothesis.

It is worth noting that mmWave signals may still suffer from distortion under strong background reflection. This is a unique challenge in comparison to tracking active objects (*e.g.*, those instrumented with RFID tags [181, 200]), and will be investigated in mTrack.

3.2.3 Beam Steering Enables Localizing Absolute Position

Despite the wide diffusion angle of the object (Figure 3.2), a 60 GHz receiver can have highly narrow beamwidth and will be able to capture the diffused signals only when it is




Figure 3.5: mTrack's system setup. Gray area denotes tracking region.

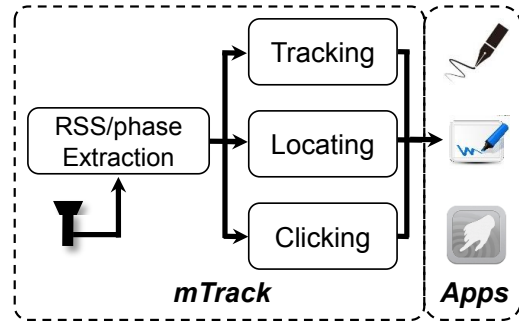


Figure 3.6: mTrack's system components and functionalities.

pointing to the target. Figure 3.4 verifies this property. We steer the receiver's pointing angle between -20° to 20° , at 1° granularity. The measured RSS peaks at around 0° when the receiver points towards the target. This indicates that *we can take advantage of 60 GHz beams' steerability to fix the target's absolute angle relative to the Tx/Rx.*

In practice, a 60 GHz antenna may not be steered continuously as in our experiment. In addition, the RSS may also suffer from background reflection effect, which causes multipath effect and RSS variation as shown in Figure 3.2. mTrack is designed to meet these challenges.

3.3 An Overview of mTrack

mTrack builds on the foregoing measurement observations to realize a high-precision mmWave trackpad. Figure 3.5 shows a typical mTrack setup. One quasi-omni Tx and two directional Rx antennas are placed to form a right angle, and track an object within a rectangular trackpad region ($0 < x < 2a$, $0 < y < 2b$). a and b are set to 50 cm by default.

Both the quasi-omni and directional beam patterns can be readily generated via 802.11ad devices' phased-array antennas [11, 180]. The quasi-omni transmitter beam in mTrack is used to illuminate the trackpad area. Typical quasi-omni beamwidth in 60

GHz ranges from 20° to 180° . The transmit signal is able to cover the tracking region, without the need to adjust the transmitter's beam direction for scanning target. On the other hand, the Rx adopts highly directional antennas, for two reasons. First, with a high antenna gain, it can substantially boost signal quality, enabling mTrack to capture the weak reflections from target. Second, a directional receiver can help isolate the leakage interference from the Tx, which may otherwise overwhelm the target-reflected signals. Using highly directional antennas reduces the beams' intersecting area. The receiver keeps track of the target by adaptively steering the beam direction. The system detects target's entrance and exit through the touch detection module (detailed in Section 3.6). It also knows when the target starts to move and the moving tendency by measuring the trend and amount of phase shifting (detailed in Section 3.4).

Figure 3.6 illustrates mTrack's design components and their interaction. The *phase tracking* module continuously tracks the target's location by measuring and manipulating phase incremental at each Rx. It resolves the grand challenge from background reflections and achieves *mm*-scale tracking accuracy. The *anchor point acquisition* (APA) module estimates the target's initial position through beam searching, leveraging our observation in Section 3.2.3. APA has lower location resolution than phase-tracking. Hence, it is only opportunistically called on to provide absolute location value for phase tracking and prevent its error accumulation. The *touch detection* module detects user lifting the writing object or reposition it on the trackpad. Finally, the tracking/locating/touching output can serve a variety of applications, including not only trackpad, but also handwriting transcription, gesture recording, *etc.*

3.4 Phase Tracking in mTrack

3.4.1 Basic Successive Tracking Algorithm

Translating phase change into path-length change. We first model the relation between target's moving distance and relative phase change between a pair of transmitter and receiver. The model formalizes our empirical observations in Section 3.2.2 and builds the foundation for phase-tracking.

Suppose the transmitter sends a single-tone signal of frequency f_c , represented as: $T_b(t) = Ae^{j(2\pi f_c t)}$, where A is the transmit signal magnitude. After target reflection and propagation, the receiver-captured signal becomes:

$$R_b(t) = \eta A e^{j[2\pi f_c(t-\tau)+\Phi_1]},$$

where η is the attenuation factor. Signal propagation time $\tau = \frac{D_t^T + D_t^R}{c}$, where c is the light velocity, and D_t^T and D_t^R are the distance from target to the transmitter and receiver, respectively. Φ_1 denotes the phase change caused by the reflection on the target's surface, which is a constant for the same material. Then the *phase shift*, defined as the received signal's phase change from time $t - 1$ to t , and labeled as $\Delta \arg(R_b)_{t-1}^t$, is proportional to the change of total path length due to movement:

$$\Delta \arg(R)_{t-1}^t = \frac{2\pi(D_t^T + D_t^R - D_{t-1}^T - D_{t-1}^R)}{\lambda_c}, \quad (3.1)$$

where λ_c is the wavelength corresponding to carrier frequency f_c .

From distance tracking to 2D tracking. Figure 3.5 illustrates mTrack's basic model setup. When deploying mTrack, the Tx and two Rx's are placed at $(2a, 2b)$, $(a, 2b)$ and $(2a, b)$, respectively. Both a and b are known.

Consider two consecutive signal samples $t - 1$ and t , while the target is moving.

```

1: Input: phase shift  $\Delta \arg(R1)_{t-1}^t$  and  $\Delta \arg(R2)_{t-1}^t$ ,
2: Output: new location  $o'(x_t, y_t)$ 
3: if isempty( $D_{t-1}^T, D_{t-1}^{R1}, D_{t-1}^{R2}$ )
4:    $o'(x_t, y_t) \leftarrow \text{APAQ}$  /*Locating module*/
5:    $\{D_{t-1}^T, D_{t-1}^{R1}, D_{t-1}^{R2}\} \leftarrow \{|\vec{o}' - \overrightarrow{o(Tx)}|, |\vec{o}' - \overrightarrow{o(Rx1)}|,$ 
6:      $|\vec{o}' - \overrightarrow{o(Rx2)}|\}$  /*Acquire initial path length*/
7: else /*Successive tracking*/
8:    $\{D_t^T, D_t^{R1}, D_t^{R2}\} \leftarrow \text{path}(D_{t-1}^T, D_{t-1}^{R1}, D_{t-1}^{R2}, \Delta \arg(R1)_{t-1}^t,$ 
9:      $\Delta \arg(R2)_{t-1}^t)$  /*Find new path lengths*/
10: /*Find new location and Update path lengths*/
11:  $o'(x_t, y_t) \leftarrow (\frac{D_t^{R1^2} - D_t^{T^2} + 3a^2}{2a}, \frac{D_t^{R2^2} - D_t^{T^2} + 3b^2}{2b})$ 
12:  $\{D_{t-1}^T, D_{t-1}^{R1}, D_{t-1}^{R2}\} \leftarrow \{D_t^T, D_t^{R1}, D_t^{R2}\}$ 
13: return  $o'(x_t, y_t)$ 

```

Algorithm 3.1: Basic Phase Tracking Algorithm.

Based on Eq. (3.1), we have:

$$(D_t^T + D_t^{R1}) = (D_{t-1}^T + D_{t-1}^{R1}) + \frac{\lambda_c \Delta \arg(R1)_{t-1}^t}{2\pi}$$

$$(D_t^T + D_t^{R2}) = (D_{t-1}^T + D_{t-1}^{R2}) + \frac{\lambda_c \Delta \arg(R2)_{t-1}^t}{2\pi}.$$

In addition with an inherent geometrical constraint of those signal path lengths (Figure 3.5), we can solve for three unknown path lengths D_t^T , D_t^{R1} and D_t^{R2} with three sets of equations, given that the corresponding path-lengths at $t - 1$ have already been obtained. This computation runs in a successive manner, with initial values D_0^T , D_0^{R1} and D_0^{R2} given by mTrack's anchor point acquisition (APA) module (to be discussed in Section 3.5). Finally, the target's absolute coordinate at time t is computed using the triangulation method, given the path length to the transmitter and both receivers' locations. Algorithm 3.1 summarizes the successive phase-tracking in mTrack.

This successive tracking is promising for real-time tracking because the solutions of path length and new location can be expressed in closed-form formulas and computed efficiently by single-step operations.

Notably, using highly directional receive antennas may result in a relatively small

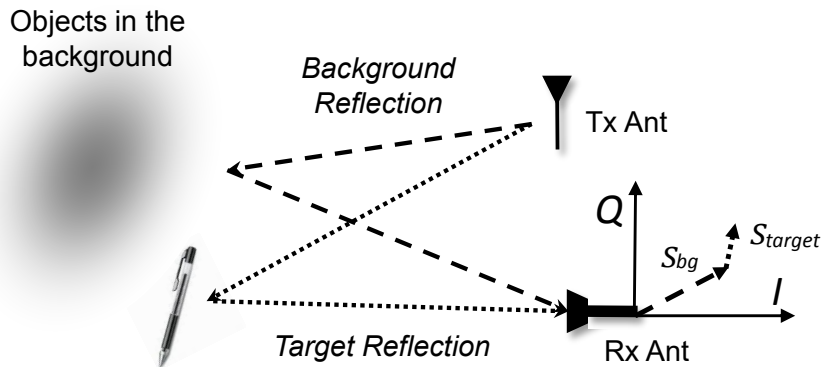


Figure 3.7: Model of the multipath reflection. We model the multipath reflections into two types: background reflection and target reflection. The sum of two reflection components is illustrated in phasor space.

detection region. For example, with two Rx antennas of beamwidth 4.5° and 0.5m away from the tracking region, the intersection area of two beams is only $5\text{cm} \times 5\text{cm}$. mTrack expands the detection region by adaptively steering the beam, whenever it gauges that the target may fall out of the Rx's beam coverage in the next 0.5 second, assuming its speed is consistent within this short period.

3.4.2 Tracking under Background Reflection

3.4.2.1 Impact of background reflection on phase shift

Though directional antennas can alleviate multipath effect, antenna beamwidth still needs to be broad enough to cover a reasonably large detection region. Thus, background reflection is unavoidable. We model how background reflection affects mTrack's successive tracking and derive insights that help disentangle target reflection from background reflection. We use vector form on the I-Q plane to represent a signal sample (Figure 3.7). Since baseband signal is a single-carrier, background reflections can aggregate into a single signal sample \vec{S}_{bg} . Similarly, target-reflected signal sample is modeled as \vec{S}_{trg} . The demodulated signal at the receiver \vec{S}_{rec} at any time is the sum of two reflected

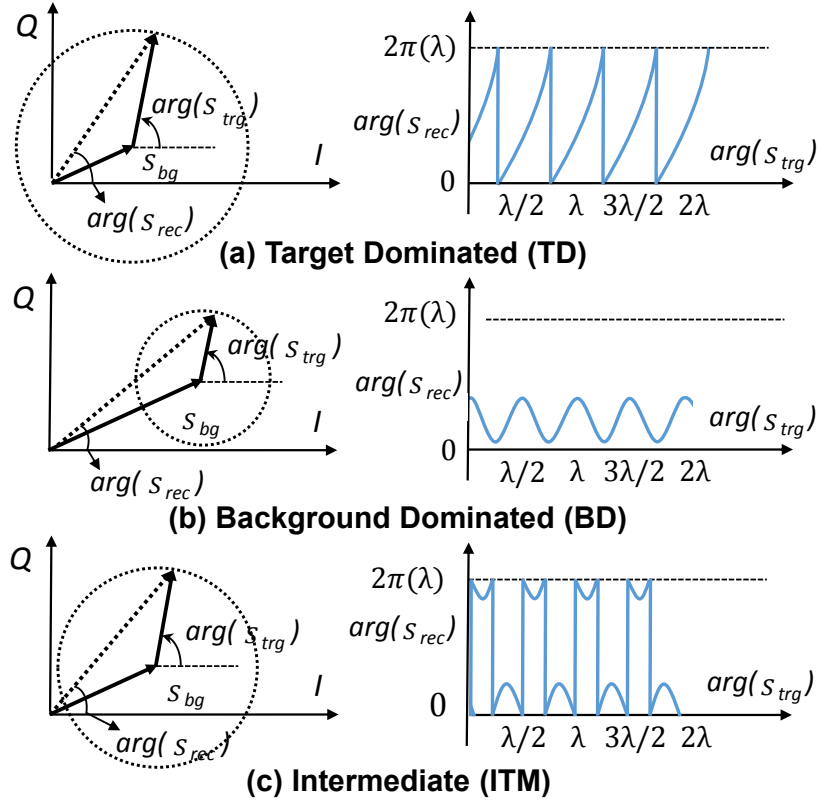


Figure 3.8: Impact of (static) background reflection on phase shift. (a) Target reflection dominates background. (b) Target reflection is much weaker than background. (c) Intermediate case.

components:

$$\vec{s}_{rec} = \vec{s}_{trg} + \vec{s}_{bg}. \quad (3.2)$$

Phase shift of \vec{s}_{trg} carries information about the moving distance of target, but the receiver can only measure the mixed signal \vec{s}_{rec} . So, can this mixed signal's phase still translate into target's moving distance? We answer the question by analyzing how the phase of \vec{s}_{rec} cycles as the target moves over half-wavelength (i.e., the phase of \vec{s}_{trg} cycles from 0 to 2π). The phase cycling pattern in turn depends on the magnitude of target reflection $|\vec{s}_{trg}|$, which we classify into three ranges, as shown in Figure 3.8.

(a) Target Dominating (TD): Suppose background reflection is stationary as the target moves across half-wavelength. Accordingly, as \vec{s}_{trg} rotates its phase through one

cycle, \vec{S}_{rec} crosses the I-plane once, resulting in phase shift of 2π . Thus, phase shift can still reveal the moving distance when target travels integer multiples of half-wavelength $\lambda_{0.5}$.

(b) Background Dominating (BD): As the phase of \vec{S}_{trg} goes through a 2π cycle, the measured phase of \vec{S}_{rec} does not cross the I-plane, and phase shift is smaller than 2π . The stronger the background reflection, the smaller it will be. Accordingly, the moving distance estimated from the phase shift of \vec{S}_{rec} is always shorter than actual value.

(c) Intermediate (ITM) case between above two: As shown in Figure 3.8(c), \vec{S}_{rec} crosses the I-plane twice, causing phase jump at each crossing point. Due to phase ambiguity of measured signal \vec{S}_{rec} , the measured phase shift also will not reveal the correct moving distance.

3.4.2.2 Dealing with Background Reflection

Why is it challenging? Multipath reflection from background is a unique challenge for mTrack which aims to track small passive objects at high precision. Unlike radar-like systems that track large targets (*e.g.*, human body [28, 198]), reflections from small objects tend to be weaker or comparable to background reflections. Due to the high sensitivity of mmWave, background reflection can disturb the signal phase even as the target moves over a short distance. Radar systems can use ultra-wideband signals to resolve the time-of-arrival of reflected signal paths, thus separating the foreground and background [113]. In contrast, mTrack relies on a narrow-band (single-tone) continuous wave that mixes background and target reflection.

Unlike active tracking systems [181, 200] where reflected paths originate coherently from the same radio transmitter, in passive tracking, only parts of the reflected signals contain desired information and need to be extrapolated. Background suppression in active tracking systems is much simpler, since the signal phase changes are caused by

baseband modulation, which operates at much higher rate (*e.g.*, multiple kHz in RFID [200]) than background variation. Hence, any background reflection can be considered as relatively static and filtered out by a DC filter [116]. This approach is not applicable to passive tracking systems where phase changes at slow rate due to target movement.

Dual-differential Background Removal. Since the phase shift directly from the received signal \vec{S}_{rec} is contaminated by background reflection, mTrack employs a *Dual-Differential Background Removal (DDBR)* algorithm to remove such impact, thus limiting the tracking error.

Underlying background removal are two assumptions: (i) Background reflection remains stationary across at least 3 consecutive signal samples. This holds in general because 60 GHz radios take RSS/phase samples (assuming one sample per packet) much more frequently than the change of background. (ii) Target reflection has the same decreasing or increasing trend of phase shift, and approximate RSS, across 3 consecutive samples. This holds again because of the high sampling rate relative to the writing object's slow motion.

Under these assumptions, the idea behind background cancellation is intuitive — we leverage the invariant behind signal differential:

$$\begin{aligned}\vec{S}_{rec}^t - \vec{S}_{rec}^{t-1} &= (\vec{S}_{trg}^t + \vec{S}_{bg}^t) - (\vec{S}_{trg}^{t-1} + \vec{S}_{bg}^{t-1}) \\ &\cong \vec{S}_{trg}^t - \vec{S}_{trg}^{t-1}.\end{aligned}\tag{3.3}$$

The average phase shift among three consecutive samples can be computed as:

$$\begin{aligned}&\frac{1}{2}[\Delta \arg(\vec{S}_{trg}^t)_{t-1} + \Delta \arg(\vec{S}_{trg}^t)_{t+1}] \\ &= \arg(\vec{S}_{trg}^{t+1} - \vec{S}_{trg}^t) - \arg(\vec{S}_{trg}^t - \vec{S}_{trg}^{t-1}).\end{aligned}\tag{3.4}$$

The proof for the above equation directly follows the geometrical relation shown




Figure 3.9: Illustration of signal vectors relation for differential background removal algorithm.

in Figure 3.9. For clarity, we introduce the following notations: $\theta_1 = \Delta \arg(\vec{S}_{trg})_{t-1}^t$, $\theta_2 = \Delta \arg(\vec{S}_{trg})_t^{t+1}$, and $\theta_3 = \arg(\vec{S}_{trg}^{t+1} - \vec{S}_{trg}^t) - \arg(\vec{S}_{trg}^t - \vec{S}_{trg}^{t-1})$. Since the target's reflection signal strength remains similar among three consecutive samples, the two triangles in solid line of Figure 3.9 can be considered as isosceles triangles. Following the geometrical relation, we have:

$$\frac{\pi - \theta_1}{2} + \frac{\pi - \theta_2}{2} + \theta_3 = \pi,$$

or equivalently $\theta_3 = \frac{\theta_1 + \theta_2}{2}$, which proves Equation (3.4). \square

From Eqs. (3.3) and (3.4), we get the phase shift with background reflection canceled by differentiating the phase of differential received signal:

$$\begin{aligned} \Delta \arg(\vec{S}_{trg})_{t-1}^{t+1} &= \Delta \arg(\vec{S}_{trg})_{t-1}^t + \Delta \arg(\vec{S}_{trg})_t^{t+1} \\ &\approx 2[\arg(\vec{S}_{rec}^{t+1} - \vec{S}_{rec}^t) - \arg(\vec{S}_{rec}^t - \vec{S}_{rec}^{t-1})]. \end{aligned} \quad (3.5)$$

The resulting $\Delta \arg(\vec{S}_{trg})_{t-1}^{t+1}$ is deemed as a single phase shift sample (with background canceled) to replace the one in the basic successive tracking Algorithm 3.1 to improve tracking accuracy.

Millimeter-wave radios typically have non-negligible phase noise which is inherent to its ultra-high frequency synthesizer. To prevent phase noise from contaminating the

phase shift measurement, mTrack adapts its phase sampling rate (by dropping intermediate samples), to ensure the phase shift between consecutive samples is not dominated by phase noise. We retain a phase sample only if its shift (relative to the previous sample) is above variance of the phase noise (available from hardware specifications).

3.4.3 Phase Counting and Regeneration

One limitation of the above DDBR algorithm is that it attempts to average out phase noise at the cost of lowering sampling rate. Consequently, between 3 consecutive sampling period, the background reflection may have already changed, thus violating DDBR's assumption.

To overcome this limitation, we design a complementary algorithm called *phase counting and regeneration* (PCR). PCR maintains a high sampling rate, but collects a sufficient number of samples (instead of 3) to recover the target-reflected signal from noise and background reflection. Under background reflection, although the absolute phase change of measured signal is no longer linear *w.r.t.* target's moving distance, its phase still exhibits a periodic pattern (Section 3.4.2.1). PCR harnesses this phenomenon to *identify/count the phase cycles* of measured signal, based on which it *regenerates* the phase change caused purely by target reflection. Then, it can simply run the basic phase-tracking Algorithm 3.1 to translate phase change into moving distance.

The three different patterns of background contamination (Section 3.4.2.1) impose significant challenges to this approach. First, the three cases may be mixed together, and the ITM case has severe phase aliasing effects. Second, we need to identify the periodicity based only on a few phase cycles to maintain the timeliness of tracking. Both challenges deter simple periodicity identification algorithms like Fast Fourier Transform. PCR addresses these challenges through three key steps (Figure 3.10 shows an example).

Reducing ITM to BD. We address the first issue by detecting the ITM case, and

convert it to the BD case which has no phase aliasing effect. For detection, PCR leverages the unique phase “jumps” in the ITM case caused by aliasing, *i.e.*, neighboring two jumps are in opposite directions (Figure 3.8). Formally, given a time series of N phase samples $x[1], x[2], \dots, x[N]$, it sequentially searches for a point m_i with phase jump:

$$m_i = \arg \min_{2 \leq i \leq N} \mathit{find}(|x[i] - x[i - 1]| > \pi). \quad (3.6)$$

It then checks if there exists a follow-on point that together constitutes the feature in ITM:

$$m_j = \arg \min_{m_i \leq j \leq N} \mathit{find}(|x[j] - x[j - 1]| > \pi) \quad (3.7)$$

$$\mathit{s.t.} \quad m_j - m_i < T_h \quad (3.8)$$

$$|x[m_i] - x[m_j]| > \pi, \quad (3.9)$$

where T_h is the largest periodicity that PCR can detect. Eq. (3.9) ensures the phase jumps at m_i and m_j are opposite.

Once an ITM case is detected, we can reduce it to the BD case by compensating the difference between discontinuous series:

$$x[i] = x[i] - (x[m_i] - x[m_i - 1]), \forall i \in [m_i, m_j]. \quad (3.10)$$

Periodicity Counting. To reliably detect the periodicity T of the time series of samples after alias removal, we adopt the phase coherence analysis in [102]. Let sequence $\mathbf{L}[i, d] = \{x[i + jd] \mid \forall 0 \leq j < \lfloor N/d \rfloor\}$, where $\{\}$ denotes concatenation of samples. Since $\mathbf{L}[i, d]$ are samples that are separated by d points in the given time series, by the periodicity definition, variance of $\mathbf{L}[i, d]$ will be minimized, *i.e.* 0, when $d = T$. Hence,




Figure 3.10: An example snapshot of using PCR to remove background reflection.

we find the periodicity T by iterating possible values in range $[T_l, T_h]$ for all sequence i :

$$T = \arg \min_t \text{var}(\{\overline{\mathbf{L}[i, t]} \mid \forall 1 \leq i \leq t\}), \forall t \in [T_l, T_h],$$

where $\overline{\mathbf{L}[i, t]}$ is $\mathbf{L}[i, t]$ with its mean value subtracted.

Regeneration. PCR will generate the counting sequence (red curve in Figure 3.10) based on the periodicity detected in the time series samples. The correct phase pattern is regenerated by creating 2π phase shift between any two neighboring counts. The trend of phase shift, either increasing or decreasing, is available from the basic successive tracking or DDBR algorithm.

We emphasize that *PCR* is uniquely designed for passive tracking using small-wavelength mmWave signals. It assumes the background reflection does not change as the target moves over $\lambda_{0.5}$, which may not hold when using microwave signals. In addition, its resolution is around $\lambda_{0.5}$, hence PCR's precision worsens to multiple centimeters when using microwave signals.

Table 3.2 summarizes and compares three processing algorithms. Since PCR algorithm works on time-series of samples spanning multiple phase cycles, it incurs proportionally higher complexity and tracking latency than Basic (successive tracking) and

Params	Basic	DDBR	PCR
Complexity	Low	Low	High
Latency	Small	Small	Large
BG Resilient	No	Yes	Yes
Error w/ BG	$\approx 30\text{mm}$	$\approx 5\text{mm}$	$\approx 8\text{mm}$
Noise Resilient	Yes	No	Yes
Error w/ Noise	$\approx 8\text{mm}$	$\approx 20\text{mm}$	$\approx 8\text{mm}$

Table 3.2: Performance profiles of mTrack’s phase tracking algorithms. Error statistics from Section 3.7.

DDBR, which work on 2 and 3 samples, respectively. However, it is more resilient to the phase noise, as will be verified in experiments.

Here are a few aspects that worth further discussion:

Object shape and size. mTrack is designed to track writing objects of small size, like pen or marker, which can be approximated as a point reflector, because their dimension (with diameter $\approx 1\text{cm}$) is much smaller than its distance to the antennas ($> 20\text{cm}$). Objects with deformable shape or large size, *e.g.* human hand and body, do not possess these properties, and hence cannot use mTrack’s model-driven tracking algorithm. Similarly, mTrack can not differentiate phase change due to large hand-shape deformation. Thus, users need to keep a relatively consistent hand shape when writing to avoid confusing the tracking algorithm. They also need to ensure a sizable part of the pen is exposed to the Tx/Rx to reflect the mmWave signals.

Phased-Array Antennas. Due to hardware limitation, we are only able to evaluate our concept with directional antennas. Beam pattern of a horn antenna differs slightly from that of an 802.11ad phase-array in two aspects: (i) The beamwidth of a phased-array is dependent on steering direction. For example, for a linear antenna array, beamwidth along 0° direction can be twice of that along 90° [180]. However, since beamforming codebook is predefined and known, we can still compensate the gain distortion of discrete sampling. The systematic solution to this problem will be left for our future work. (ii) Sidelobes of a phased array may be relatively large compared with that

of a directional horn antenna. Sidelobes may capture signals from undesired directions. However, since they usually have much smaller gain compared to the main lobe, as long as there is a RSS peak when pointing to the target, mTrack’s beam-steering based APA mechanism can still work.

Tracking multiple objects. mTrack can be potentially extended to track multiple targets simultaneously. When the objects are sufficiently separated, and not falling in the same beam direction, the receiver radios can steer the beam towards different objects in a time-interleaving manner and update their tracking trajectories sequentially. Further exploration of this idea is left for our future work.

3.5 Anchor Point Acquisition

APA complements phase tracking in two aspects: *(i)* estimating initial location, which serves as an anchor point to bootstrap the successive tracking; *(ii)* opportunistically calibrating the tracking results to prevent error accumulation.

Locating through discrete beam steering. Our feasibility study in Section 3.2 reveals a key challenge in locating the target: practical 60 GHz antennas can only be steered at discrete angles. We resolve this challenge by leveraging the roll-off pattern of the antenna gain. The gain of a 60 GHz directional Rx antenna (either horn antenna or phased-array [180]) decreases continuously near its beam edge. Accordingly, as a signal source moves from the beam center towards beam edge RSS falls off smoothly. The fall-off trend is preserved even if Rx steers itself in discrete steps.

Built on this observation, mTrack realizes discrete spatial sampling of the target’s direction, thereby relaxing the beam resolution requirement. In analog-to-digital signal conversion, a continuous signal can be preserved by interpolating it across discrete samples, using sinc as the basis function. Similarly for mTrack, the continuous spatial RSS variation can be reconstructed by interpolating the discrete samples, using antenna




Figure 3.11: Spline interpolation improves granularity of direction estimation from discrete steering.




Figure 3.12: APA improves locating accuracy that is disturbed by background reflection.

gain pattern (over angles) as basis function. Yet, since the exact gain pattern is unavailable, we choose to use *spline interpolation* instead. Spline interpolation performs piecewise polynomial approximation to the discrete samples with an objective that minimizes curve bending. It can thus recover the peak RSS position. To perfectly reconstruct continuous variation, angle of beam-steering for each step is required to be smaller than $2\times$ the beamwidth.

Figure 3.11 shows one example test result, with discrete RSS samples (blue dots) at 8° beam-steering steps, along with the interpolated curve. We can see that the interpolated maximum RSS position is very close to the ground truth.

Following mTrack’s setup (Figure 3.5), the target’s angle relative to each Rx antenna can be translated into its absolute location via simple geometrical manipulation. We omit the details for the sake of space.

Background RSS Subtraction. Similar to the successive phase tracking, the RSS values captured during APA may be contaminated by background reflection. However, RSS is much less sensitive than phase because the locating result will be affected *only if* the background change alters the peak in the RSS scanning. Thus, mTrack employs a simple iterative background RSS subtraction method in APA. Before use, each Rx collects the RSS of background reflection for each steering angle. When a new APA is invoked, the background RSS at each steering direction is subtracted from the measured

RSS. During the writing process, it knows the precise angular location of the target by combining phase tracking and initial APA. Hence, it periodically rescans the background RSS by steering off the target position.

During APA, the tracking module has to be temporarily blocked. However, since mmWave phase-array antenna can instantaneously switch its beam (at ns level latency as specified in 802.11ad [171]), the impact on tracking is negligible.

Figure 3.12 illustrates one test experiment, where a metal-surface cup is placed $25cm$ behind the detection region and a pen is used as tracking target. When there is only pen (blue curve) or only background (yellow curve) in the environment, mTrack can locate the direction of target accurately (in 40° and 30° respectively). When the tracking target is mixed with background object, the system is misled to take the later as locating result (in green curve) due to its stronger reflected signal strength. In the red curve, background subtraction cancels the background reflection from cup, and thus the target position can be accurately located.

Opportunistic Calibration. To curtail the error accumulation from successive phase tracking, APA is invoked opportunistically to recalibrate the absolute target position.

mTrack adopts a k/θ test to determine when to invoke a calibration. APA module periodically estimates anchor points. If past k anchor points all fall in a sector of angle θ originating from current position, the deviation of tracking trajectory can be considered as sufficiently large. In this case, the calibration will be applied by updating current position to the latest anchor point. In practice, mTrack empirically chooses $k = 3$ and $\theta = 120^\circ$.

3.6 Touch Event Detection

Besides tracking, mTrack detects touch-related gestures including: (i) “touch” – pen landing on writing region, (ii) “lift” – pen lifting above the region, and (iii) “click” –




Figure 3.13: Example snapshot of three type events. Phase shift variance and RSS are for event detection and classification.

essentially a sequential combination of “touch” and “lift” event. The touch gestures can serve as commands, *e.g.*, start/pause of tracking and segmentation of writing trajectory.

The key principle in mTrack’s gesture detection algorithm is to identify signal features that discriminate each gesture, and distinguish them from normal writing. Given the ultra-narrow beam, one may consider the touch gestures as “beam cutting” events, and use RSS variation patterns to detect them [111]. However, we find that RSS fluctuates significantly during writing, which incurs frequent false alarms. To address this problem, we resort to the *phase variance* as an additional feature along with RSS. Since phase shift reveals the target’s movement, its variance will characterize the abruptness of change in the pen’s moving speed, which is suitable for gesture detection.

Figure 3.13 illustrates a snapshot of signal features for the three types of gestures. A sliding window of duration $0.1s$ is used to smooth the phase variance. We can see that all gestures will create large phase variance, which is amenable for flagging a valid gesture event. To discriminate different types of gestures, we can resort to the RSS pattern, which tends to have higher value when the pen is in the tracking region. Based on above observations, mTrack identifies and classifies the gesture using a simple decision tree (Alg. 3.2).

In Algorithm 3.2, there are two parameters. s is the number of samples that are used to compute average RSS, which is empirically chosen as $0.1s$. H denotes the threshold

```

1: Input: phase variance  $\mathbf{P}[n]$  and RSS  $\mathbf{R}[n]$ 
2:  $\{V, m\} \leftarrow \max(\mathbf{P})$  /*max value and index*/
3: if  $V > H$ 
4:   if  $\mathbf{R}[m] > \text{mean}(\mathbf{R}[m : m + s])$  and
5:      $\mathbf{R}[m] > \text{mean}(\mathbf{R}[m - s : m])$  /*Check RSS*/
6:     return “Click” /*Click generates RSS peak*/
7:   else if  $\text{mean}(\mathbf{R}[m : m + s]) > \text{mean}(\mathbf{R}[m - s : m])$ 
8:     return “Touch” /*Touch increases RSS*/
9:   else
10:    return “Lift” /*Lift decreases RSS*/
11: return “NULL” /*Event not detected*/

```

Algorithm 3.2: Decision Tree for Touch Gesture Detection.

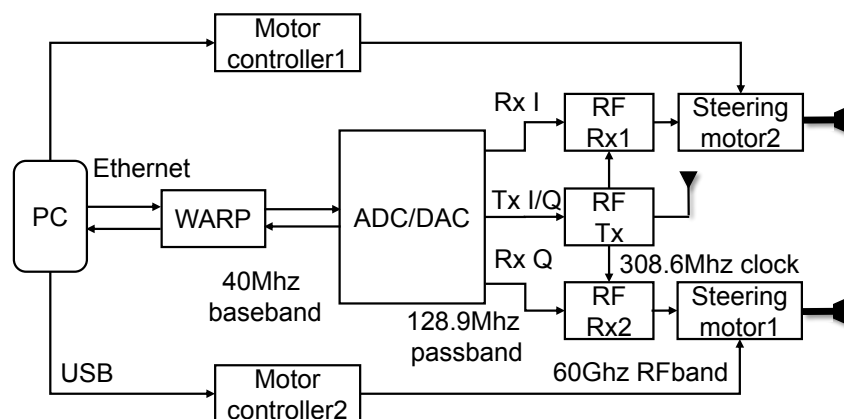


Figure 3.14: Hardware architecture of mTrack.

for a valid peak of phase variance. Since different users’ gesture styles (*e.g.*, touching speed) may differ, mTrack requests the user to provide a 10-sample training set before using the gesture detector. It then selects an H that minimizes the difference (computed using standardized Euclidean distance [195]) between the widths of the 10 peaks.

3.7 Implementation and Evaluation

We use a custom-built 60 GHz software-radio testbed (Figure 3.14) to evaluate the viability and performance of mTrack. The testbed uses WARP [138] as baseband processing unit, but extends its carrier frequency to 60 GHz. Baseband digital waveforms are gen-

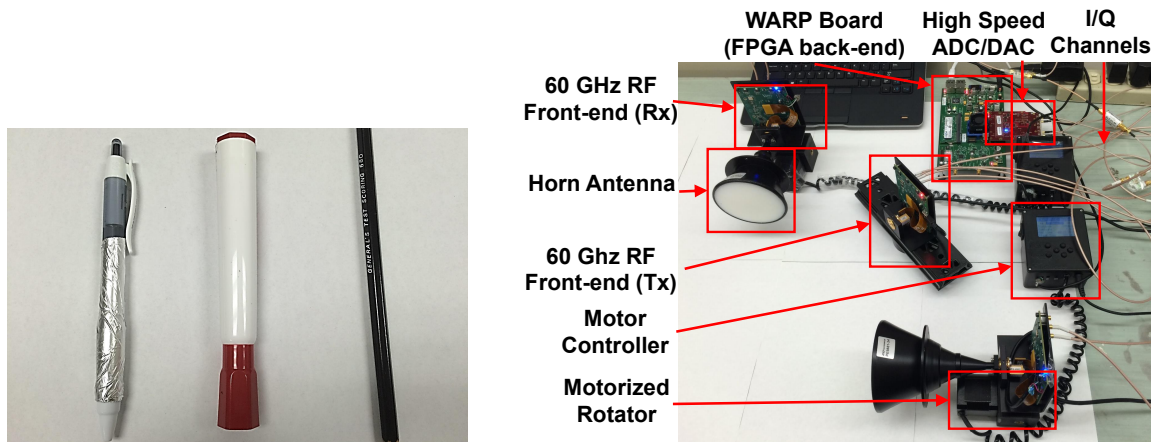


Figure 3.15: mmWave testbed. Testing metal-surfaced pen, plastic marker, wood pencil. Radios are placed 10cm beside 50cm \times 50cm writing region.

erated in a host PC, forwarded to WARP, and converted to analog using a high-speed Digital-to-Analog Converter (DAC). Resulting analog signals are carrier-modulated and transmitted by the Vubiq 60 GHz RF front-end [15]. Two receivers are synchronized by a 308.6 MHz external clock signal sourcing from transmitter oscillator. Received signals will be decoded following reverse path, and eventually reach the Personal Computer (PC) host which runs mTrack's algorithms.

Tx and Rx RF front-end each has a waveguide module as interface to 60 GHz antennas (Figure 3.15). Transmitter uses its original waveguide module PEM-001 as antenna, which is quasi-omni-directional (180° beamwidth). Tx sends 5 MHz single-carrier baseband signals modulated by 60 GHz carrier waves. Regarding receiver, we are unaware of any programmable 60 GHz phased-array antenna. Thus, in our prototype, signals are captured by a directional receiver with horn antenna PE9881-34 of 3.4° beamwidth. To emulate the effect of beam steering, the Rx antenna is mounted on a motion control system [2] that is connected to the PC host and allows programmable azimuth rotation at 1° granularity.

We also implement the RSS/phase extraction module on the 60 GHz testbed. It generates and sends single-tone frames (3000 samples each) through Tx, and estimates a pair of RSS and phase values from each received frame. To avoid phase ambiguity, the

change between any two phase values cannot exceed π . By default, the extraction module samples RSS/phase at 300 Hz. Thus, the highest tracking speed that our prototype can support is $2.5\text{mm} \times 300 \text{ Hz}/2 = 37.5\text{cm/s}$. The speed is sufficiently fast to capture normal hand writing. Note that 300 Hz phase sampling rate is far below our hardware limit. With 40 MHz bandwidth, the phase sampling rate can be up to $40 \text{ MHz}/3000 = 13 \text{ KHz}$, equivalently translated to a tracking speed of 26 m/s.

On top of the RSS/phase extraction, we implement mTrack’s major modules, *i.e.*, tracking, locating and touch detection. Core processing algorithms DDBR, PCR and APA are designed as middle-ware plugins that work transparently between RSS/phase extraction and these modules. Algorithms operate on samples that are loaded sequentially into a data buffer.

3.7.1 Micro Benchmark

3.7.1.1 Passive Tracking Performance

The experiments are conducted in an office environment with natural background: one Rx antenna faces a human and a metal cabinet 1.5 *m* away; the other one faces a drywall 2 *m* away. For a micro-benchmark test, a pen is first attached to the motion track and moves linearly under the control of the PC host. This setup allows us to obtain the pen’s location over time, and use it as ground truth to evaluate mTrack. Unless noted otherwise, we use the PCR algorithm in mTrack (Section 3.4.3) to combat background reflection.

Combating background reflection. We compare mTrack with the ground-truth (Oracle) and two alternative techniques: (i) *Basic*, which directly employs the phase-tracking algorithm in Section 3.4 without handling background reflection. (ii) *Static Subtraction (SSub)*, which surveys the background reflection offline when the target is out of tracking region. Then, it subtracts the measured background directly from




Figure 3.16: Tracking along an example linear trajectory(10 cm) without BG.




Figure 3.17: Tracking example with background reflection.




Figure 3.18: Error CDF of tracking methods under background reflection.

baseband signals, assuming the background is static and remains the same w/ or w/o target present. This is typically used in scenario where background signals have long coherent time, *e.g.*, radar-based object identification [29, 113].

Figure 3.16 and 3.17 depict example trajectory outputs of these schemes when the pen moves linearly across 10 cm w/ and w/o background reflection. Figure 3.18 plots the CDF of tracking errors across the trajectory and over 10 trials. When there is no background, all methods have little deviation. When background is present, the basic method exhibits large deviation (median error of 30 mm and 90-percentile 5 mm). Static subtraction leads to even larger error (median 40 mm and 90-percentile 70 mm), because mmWave’s phase is highly sensitive to environment. The background signals hardly remain static and subtracting an incorrect estimation incurs large penalty because of equivalently introducing an artificial background.

In contrast, mTrack’s DDBR and PCR algorithms both demonstrate *mm*-scale preci-




Figure 3.19: PCR performance under different background strength.




Figure 3.20: PCR performance under human as dynamic background.

sion, with median (90-percentile) error of 11 (18) *mm* and 6.5 (9) *mm*, respectively. PCR achieves higher accuracy than DDBR mainly because our 60 GHz software defined radio platform has relatively high phase noise (-80dBc at 5Mhz offset from carrier according to our measurement). It can lead to maximum deviation of up to 0.2 radian within one phase cycle. COTS 60 GHz radios have much lower phase noise (typically <110dBc at 5Mhz) [175]. Our trace-based simulation shows that DDBR can outperform PCR significantly under such noise level.

Different types of background. We now evaluate the performance of *Basic* and *mTrack* (using PCR) under 3 different background types: (i) Strong background: a metal cabinet is placed near the tracking region, facing one Rx antenna. (ii) Weak background: a 60 GHz RF absorber is placed at the end of tracking region to isolate the impact of cabinet on one Rx antenna, leaving background reflection from drywall to the other antenna. (iii) No background: RF absorbers are placed along both the x- and y-axis, protecting both Rx antennas from background reflection.

We also evaluate *mTrack* with human as dynamic background: (i) CS (close and stationary), MS (median and stationary), and FS (far and stationary): A human stands 1m, 2m and 3m facing one Rx antenna. Human breathing changes background reflection signals. (ii) WK1 and WK2: The human randomly walks in parallel or perpendicular at 2~3m to Rx antenna.

Figure 3.19 shows that *mTrack* maintains *mm*-scale precision across all background




Figure 3.21: Tracking trajectory and tracking error of the DDBR algorithm under different levels of phase noise.

types. Its PCR algorithm effectively eliminates the impact of background reflection, leaving median (90-percentile) residual error of only around 8 (12) *mm* compared with the perfect case (no background). In contrast, the basic phase tracking deviates by up to 30 (50) *mm*. Figure 3.20 shows that PCR can still effectively reduce tracking error under human walking as dynamic background, since human body will not strongly reflect mmWave signals.

Impact of phase noise on DDBR performance. We evaluate DDBR’s performance under different levels of phase noise through trace-based simulation. The phase shift of target reflection is generated by simulating the target moving along a triangle of size $5\text{cm} \times 5\text{cm}$ following the setup in Figure 3.5. Meanwhile, there is background reflection from a static reflector, whose reflectivity is $1/3$ of the target. The simulated phase shift is contaminated by the trace of phase noise collected from our platform.

Figure 3.21 shows the impact of phase noise. DDBR’s tracking error decreases by more than 80% when the phase noise is reduced from -80dBc to -86dBc , and after background removal, the trajectory is close to the oracle shape (triangle). Further reducing the phase noise (to $<-86\text{dBc}$) will not significantly improve tracking performance since phase noise is sufficiently small for DDBR to function correctly.




Figure 3.22: APA performance with spline interpolation.

3.7.1.2 Performance of APA

We proceed to verify mTrack’s APA algorithm with similar experiment setup as above. The target pen is randomly placed over 30 positions 40 to 50 *cm* away from the receivers. We evaluate mTrack’s error in locating the target’s angle relative to the transmitter/receiver (Section 3.5).

Benchmarking APA with discrete-step beam-steering. To emulate practical behavior of 60 GHz switched-beam antenna, we steer the antenna at discrete steps of 8° ($2\times$ the beamwidth). Figure 3.22 shows that the spline interpolation does play a crucial role in mitigating the effect of discrete switching between beam directions. Compared with the basic method that directly takes the RSS-maximizing direction (Section 3.5), it reduces the 90-percentile error from 2.5° to 1.5° , translating to positioning error of 10.5 *mm*.

Background subtraction for APA. In this experiment, in addition to the default background (cabinet and drywall), we intentionally place a metal cup near the end of writing region. Unlike large background like cabinet or drywall, small reflective object as such is more likely to create sharp peak along certain direction, thus disturbing APA.

Figure 3.23 shows the anchor positioning performance. When beam steering granularity is high (*e.g.*, $\leq 4^\circ$), mTrack has a small direction estimation error (around 1°) even without background subtraction. This is because the pen is closer to the antennas and generates stronger reflection than the background cup. Their reflection peaks can thus




Figure 3.23: Effectiveness of background subtraction in APA.




Figure 3.24: APA under dynamic background. Step of antenna steer 4° .

be reliably distinguished through fine-grained scanning. However, with coarser steering granularity, their peaks tend to merge, leading to an increasing error. Fortunately, background subtraction effectively mitigates the impact of background reflection, hence reducing the estimation error by 50%.

Figure 3.24 shows the APA performance under human movement as dynamic background. Human movement at 2m away from receiver does not affect positioning error, since reflecting RSS from human body is much weaker than pen. Estimation error without background subtraction increases to 2.8° when human stands close to receive antenna. However, background subtraction can still consistently reduce positioning error even under human movement.

3.7.1.3 Joint Performance of Tracking and APA.

Recall APA facilitates phase tracking through opportunistic calibration. In this experiment, we verify the effectiveness of this joint execution. The target pen moves along a *circular* trajectory of radius 7 cm. mTrack continuously runs phase tracking, and performs the k/θ -test (Section 3.5) every 2 seconds. It invokes APA calibration if the test dictates so. Figure 3.25 shows the tracking error at every 2-second check point. Without APA calibration, the phase tracking error steadily accumulates over time and reaches 46 cm when moving 150 cm continuously along the circle. In contrast, APA calibration caps the phase tracking error below 10 mm across 90% of the trajectory.

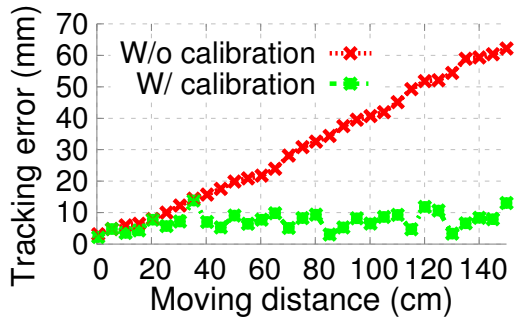


Figure 3.25: Error w/ and w/o opportunistic calibration.




Figure 3.26: Locating error under different target materials.

3.7.2 Performance on a Trackpad

We now evaluate mTrack’s performance in a real trackpad application. The experiments are conducted in an office environment with natural background (drywall, metal cabinet, a user, and occasional human walking by). A $50\text{cm} \times 50\text{cm}$ writing region is created on a wood table. To test the precision of APA in locating anchoring points, the user rests the pen tip on 40 random locations, ensuring the bottom part of the pen is exposed to the antennas. mTrack steers the antennas with granularity of 8° . To test phase tracking, the user draws 10 circles and 10 triangles (with 20 cm perimeter) following printed trajectories in the normal hand-writing speed. Since human-hand deformation will affect phase tracking, testers hold the middle portion of pen, while directional antennas point to the bottom portion. Due to lack of timing-synchronization between user writing trajectory and tracking estimation, we approximate the tracking error as the minimum projection distance from mTrack’s location estimation to the trajectory.

Types of writing objects. We evaluate APA for 3 writing objects of different reflectivity: metal-surfaced pen, plastic marker, and wood pencil (Figure 3.15(b)). Our benchmark measurement shows that, at 40 away from the transmit/receive antenna, the SNR of signals reflected by these objects are 12.3 dB, 10.1 dB and 4.7 dB, respectively. Figure 3.26 plots the APA error distribution, which shows 90-percentile error of 2 cm, 4 cm and 16 cm, respectively. Obviously, *object with strong reflectivity enables APA*




Figure 3.27: Tracking error under different target materials.




Figure 3.28: Map of APA Error

to easily combat noise, thus achieving higher precision. Note that the APA precision for pen is lower than the benchmark test in Section 3.7.1.2, mainly because the presence of user's hand creates more uncertainties.

Remarkably, mTrack's phase-tracking algorithm demonstrates the high accuracy in this trackpad application (Figure 3.27). The 90-percentile errors for pen, marker and pencil are 8 mm, 11 mm and 4.8 cm respectively.

Localization/tracking error across a large region. Distance between the target and the receiver determines the reflected signal strength and hence may affect mTrack's accuracy. The transmitter and receivers are placed at coordinates (100, 100) cm, (50, 100)cm and (100, 50)cm, respectively. To quantify such location-dependent error, we partition the writing area into 10cm × 10cm squares, and repeat the previous precision test on each square. Figure 3.28 and 3.29 plot the APA and phase-tracking error across all squares within a 90cm × 90cm region.

When the pen is close to both receivers (distance < 60cm), mTrack can achieve high accuracy with APA/tracking error of <1.5 cm and <8 mm, respectively. Accuracy starts degrading when the target moves over 70 cm away from receiver, and hence SNR drops. Nonetheless, the tracking error is still within 1.5 cm even when the target is 90 cm from the receiver. mmWave attenuates to almost noise floor at 100 cm owing to high pathloss of mmWave signals. We expect at least two ways of scaling the writing region: increasing the transmit power, and placing more receivers along the x- and y-axis. We




Figure 3.29: Map of phase-tracking error.




Figure 3.30: Touch detection accuracy for different users.

Event	Touch	Lift	Click	ND
Touch	94.0%	0	0	6.0%
Lift	0	93.5%	0	6.5%
Click	0	0	94.8%	5.2%

Table 3.3: Classification accuracy of touch event. ND means no-detection.

leave such exploration for future work.

Accuracy of touch event detection. To verify mTrack’s event detection algorithm, we recruit 7 users to perform the “touch”, “lift” and “click” gestures on the writing region using a pen. Each user first provides a 10-sample training set to find the best threshold H (Sec. 3.6), and then repeat each gesture for 50 times to test the detection algorithm. Table 3.3 shows the confusion matrix of gesture detection among all users, and Figure 3.30 plots the detection accuracy of each users over all events. The detection accuracy of all three types of events is around 94% and does not vary noticeably across users. Notably, the event misclassification rate is 0, indicating that the RSS based approach can reliably discriminate different gesture events owing to the highly directional beams. The miss detection of events is mainly due to the variation of users’ touching speed, which may not always be consistent with the training set.




Figure 3.31: Example letter and word by mTrack.




Figure 3.32: Character and word recognition accuracy.

3.7.3 Application of mTrack

We integrate mTrack with a word recognition application to showcase its potential in computer-human interaction. In specific, we record the writing trace from mTrack word by word, and then export the trace to control the mouse on a PC that runs the MyScript Stylus [18] word recognition software. mTrack’s event detector is used to segment the connections between characters in a word. Figure 3.31 shows two examples when we use mTrack to track the character/word trajectory, which exhibit high fidelity even if each character spans only 1 or 2 *cm*.

We further recruit 7 users to each write 100 random characters and 50 words using metal-surfaced pen on a 15 *cm* × 15*cm* writing region. The words are randomly picked from the standard MacKenzie set [108], which well represents the usage frequency of English words. Figure 3.32 plots the recognition accuracy across users. The accuracy for both character and word ranges from 81 to 89% across users. Users also perform the same test on a real trackpad, and the recognition accuracy is from 88 to 92%, which shows that mTrack’s performance is close to the trackpad. Considering the potential penetration of portable millimeter-wave radios [53, 126], we believe mTrack holds potential as a new form of “in-situ” transcription system.

3.8 Discussion

This chapter presents mTrack, a high precision passive object tracking system that uses 60 GHz signals as sensing medium. mTrack takes advantage of the short wavelength and steerable directional beams of 802.11ad-like 60 GHz radio devices, and uses subtle RSS/ phase variation to track passive writing objects on a trackpad-like area. It can track a pen at sub-centimeter level accuracy, which even outperforms existing systems that use radio-instrumented objects. Considering the growing popularity of 60 GHz devices, we believe mTrack can potentially open up a wide range of mobile sensing applications.

Part II

Robust mmWave Networks for Connected IoT

CHAPTER 4

Pose Information Assisted Robust 60 GHz Networks

The millimeter-wave (mmWave) wireless technology is emerging as a disruptive networking paradigm to provide multi-Gbps connectivity for demanding applications, such as wireless backhaul [58], cordless virtual reality (VR) [25], wireless fiber-to-home [54], mobile-to-screen video cast [22], *etc.* FCC's recent policy to release 14 GHz of unlicensed spectrum, along with standardization activities such as the IEEE 802.11ad [82], 802.15.3c [9] and 802.11ay [20], have spawned many consumer grade mobile devices on the 60 GHz mmWave band. For example, TPCAST [170], a 60 GHz adapter, can replace the cables between a VR headset and its PC host. Recent 60 GHz capable smartphones and laptops [144] can stream Gbps uncompressed videos to an external display.

Ideally, one would anticipate the 60 GHz technology as the next-frontier for mobile broadband, to replace the current WiFi. However, to date, the use cases of the 60 GHz wireless technology have mostly been focusing on point-to-point, stationary links. To enable seamless coverage and mobility support, 60 GHz networks need to overcome fundamental barriers that do not exist in the prior low-frequency counterparts. To compensate for the intrinsic attenuation losses due to short signal wavelengths, 60 GHz radios adopt phased-array antennas which comprise many planar patch elements to form directional beam patterns. Although these beams are electronically steerable, their joint coverage is still limited due to the inherent half-space coverage of patch antennas [134].

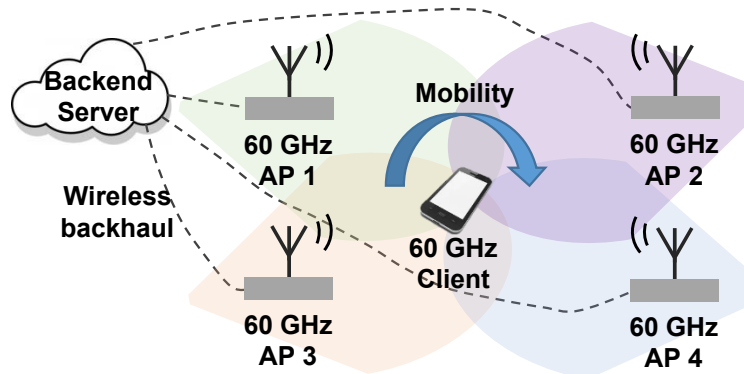


Figure 4.1: Multi-AP networks: APs are partially overlapped and connected to a backend server via wireless backhaul.

So the phased-array bears a limited field-of-view (FoV), much like a camera. And the signal strength largely depends on whether the receiver falls in the transmitter’s FoV. Consequently, achieving stable 60 GHz connectivity, even at room-level, becomes a non-trivial task.

This chapter introduces Pia, a robust 60 GHz network architecture that can provide room-scale coverage at multi-Gbps bit-rate. Pia is tailored for emerging applications such as wireless virtual-reality [25], augmented-reality [120], uncompressed miracast [22], all requiring Gbps connectivity between a mobile device and a backend server for computational offloading or graphical rendering. Our basic idea is to deploy multiple cooperating APs (*e.g.*, at the corners of the room), each covering a “picocell” region that can complement others’ blind spots. No matter how the client device moves and rotates, it is likely to fall in the FoV of at least one of the APs (Figure 4.1). The APs can be 802.11ad-compatible access points or dedicated relaying devices [25], which connect to a backend server through Ethernet cables or fixed-beam wireless backhaul [58].

The idea of multi-AP coordination has been investigated extensively in legacy WiFi or cellular networks [63, 115, 128, 159, 162], and the key problem is to determine which AP to connect to, based on signal strength, traffic load, *etc.* But the new characteristics of 60 GHz bring unique dimensions to the problem. Since each 60 GHz AP may possess hundreds of beam directions, a straightforward way of selecting the AP will entail trial-

and-error probing across all these directions for all APs. When multiple clients coexist, multiple iterations of probing are needed to negotiate the best spatial reuse, which compounds the overhead. Most critically, the probing needs to be done not only when a client changes its location, but also its orientation. Consequently, the probing overhead may overwhelm normal data transmission.

Pia overcomes this issue by leveraging the 5-DoF pose information (x, y, z position and polar/azimuth orientation) that is available on many mobile devices (*e.g.*, VR headsets [79] and smartphones/tablets with visual-inertial sensing capabilities [67, 101]). More specifically, Pia employs a model-driven approach to predict the best AP based on a client’s pose, and to assign the best beam pattern each AP should use to maximize the spatial reuse among multiple clients. To this end, we identify 4 unique design challenges.

First, our measurements indicate that the 60 GHz throughput often experiences a catastrophic drop as a client moves out of the AP’s FoV, so the conventional multi-AP protocols that react after throughput change can no longer sustain robust connectivity. We thus design a *pose-assisted link predictor* that allows the client to proactively switch to a new AP, before the current link quality drops to an intolerable level. This decision builds on a prediction of the client’s short-term pose change, paired with a simple model of the AP’s coverage which does not require probing all beam patterns.

Second, a vast literature in directional-antenna networking at low-frequency bands assumed cone-shaped beams [41]. In contrast, practical 60 GHz phased-arrays have imperfect directional beam patterns that often comprise multiple sidelobes, causing irregular interference patterns across spatial angles. Pia addresses this challenge using a *pose-assisted spatial sharing* mechanism, which jointly optimizes the AP selection and beam assignment, to maximize the concurrent transmission opportunity when multiple clients coexist.

Third, the above two mechanisms assume Pia knows the AP’s relative position and

orientation within the client’s coordinate. However, the actual APs are often placed in an ad-hoc manner. To meet this challenge, we design a statistical algorithm, called *AP-Pose Sensing (APS)*, that can estimate the AP’s pose based on a random set of link-quality measurement at the client side.

Finally, Pia’s interference management mainly models the line-of-sight (LOS) link quality. But the model can be occasionally disturbed by non-line-of-sight (NLOS) reflections from close-by objects (*e.g.*, concrete walls or metal cabinet). We design a novel *compressive angle estimation* method, that fuses the pose information with the link quality measurement, so as to discriminate the reflection paths and model their impacts separately.

The contributions of Pia can be summarized as follows.

(i) We propose to use pose information as a fundamental primitive, and identify the associated challenges/opportunities to facilitate robust connectivity within a 60 GHz multi-AP architecture.

(ii) We design pose-assisted link prediction, spatial sharing, and AP-pose sensing mechanisms, to enable efficient link/interference management for practical 60 GHz clients.

(iii) We implement Pia on a COTS testbed with 4 60 GHz radios each with a 4×8 phased array. Our experiments show that Pia’s pose-assisted AP switching effectively reduces the hazard times (catastrophic throughput drops) by $4.2 \sim 6.3 \times$, and the spatial sharing enhances the network capacity by 76% compared to the 802.11ad.

4.1 Related Work

Pia is most closely related with the following domains:

Robust 60 GHz networks. To make 60 GHz networks robust against human blockage and movement, prior research primarily focused on minimizing the beam search-

ing overhead, so that the 60 GHz links can efficiently recover from disruption. Existing systems leveraged the correlation between beams [166] or used out-of-band channel [25, 164] to estimate the best beam at low overhead. However, when blockage occurs, the reflection path may be either weak or non-existent [165]. To ensure better coverage, MoVR [25] adopted a customized 60 GHz relay to amplify and forward the AP's signals. To maximize the beam alignment, MoVR also leverages the pose to guide the AP/mirror's beam steering. Pia shares similar spirit—the multiple APs act as relays for the client. Yet, Pia focuses on the more general problems of AP selection/switching and spatial reuse, harnessing the pose information on client devices. Pia is compatible with 802.11ad or any future 60 GHz base stations, and does not need dedicated analog signal forwarders as in [25]. UbiG [145] discussed an alternative solution of 60 GHz AP switching based on the channel probing. The client can estimate its physical position in the network by analyzing multipath reflections from nearby APs. However, the intensive overhead of channel feedback and multipath computation make this approach not suitable for latency-sensitive applications, *e.g.*, wireless VR. In contrast, Pia can deliver the low-latency network and better capacity by leveraging the pose information and jointly minimizing the interference across multiple APs and clients.

The general idea of using smartphone sensors to improve network performance has been examined in [137], which proposed to synthesize various motion sensor hints to adapt protocol primitives such as bit-rate adaptation. Sani *et al.* [33] conducted an empirical study and showed that, by judiciously choosing among multiple directional antennas on a mobile device based on motion sensor information, the WiFi link SNR can be improved by 3 dB. In contrast, new challenges emerge in 60 GHz networks due to vulnerability to blockage and limited FoV. Meanwhile, the electronically steerable beams bring new opportunities to spatial reuse. Yang *et al.* [202] designed a sensor-assisted multi-level codebook for efficient beam searching under mobility. Pia, in contrast, harnesses the 5-DoF pose information to address the network-level problems arising from

multiple APs and clients.

AP handover and selection. Numerous solutions have been proposed to curtail the AP handover/reassociation overhead in WiFi networks. The most viable idea is to fork multiple virtual network interfaces [47, 90]. Although the switching mechanism itself is still applicable, the decision metrics for switching (*e.g.*, signal strength, distance [115, 162] and location [63, 128, 159]) in low-frequency networks are no longer applicable to the 60 GHz network due to orientation sensitivity. Alternative machine-learning based selection algorithms [55, 163, 196]) have been explored. However, such learning algorithms require dense training, and the accuracy is still much lower compared with Pia's geometry-model based algorithm (Sec. 4.6.1). Athanasiou *et al.* [37] proposed an asymptotically optimal algorithm to balance and ensure fair client association in 60 GHz networks. In contrast, Pia is the first work from an architecture/systems perspective that optimizes the 60 GHz link quality and reliability via AP/beam switching using the pose information.

Spatial reuse of 60 GHz networks. Directional antennas are known for achieving spatial reuse and alleviating interference. A vast literature [41] explored the MAC issues in directional networks. DIRC [106] further accounted for the impact of irregular interference with imperfect directionality. These systems focus on static radios, with a cone-shaped antenna model, which differs dramatically from the beam shape of practical 60 GHz phased arrays (Sec. 4.4.1). Park *et al.* [122] proposed a null-forming technique that can create specific null points to avoid interfering peer devices, but this requires arbitrary control of the antenna weights, whereas practical 60 GHz phased arrays can only use a set of fixed beamforming weights which form a codebook.

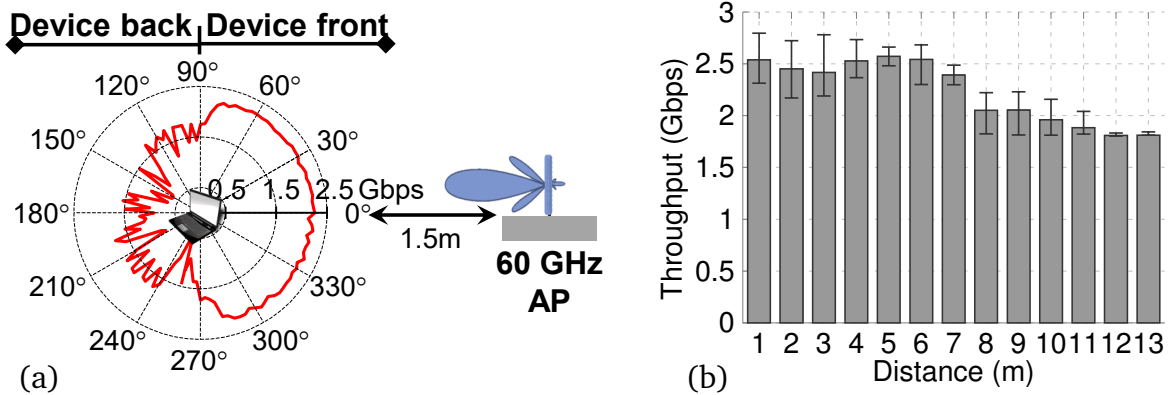


Figure 4.2: FoV and performance of the commercial 60 GHz laptop in our measurement. Varying (a) client device’s azimuth orientation, (b) AP-to-client distance (azimuth = 0°).

4.2 Background and Overview

4.2.1 Limited FoV of the 60 GHz Radio

Impact of FoV alignment on link quality. Early studies of 60 GHz phased-array [134] indicated that its coverage is often limited to less than half-space. To elucidate the problem, we set up two 802.11ad laptops, as AP and client, respectively; both equipped with the Qualcomm QCA6300 series chipset [133]. We measure the link throughput when rotating the client in front of the AP which is 1.5 m away and facing to the client directly (more details of the setup are in Sec. 4.5). From the results (Figure 4.2(a)), we can make two observations. *First*, the client’s high-throughput coverage area forms an angular sector (around 170°), which we refer to as its *FoV*. Once the AP falls outside the client’s FoV, the throughput drops sharply, from 2.2 Gbps to several hundred Mbps. Note that *the antenna’s FoV is different from its beam pattern*. Each beam pattern is much narrower than the FoV itself (more in Sec. 4.4.1). *Second*, there is no ‘notch’ inside the FoV—all the beam patterns together can fully cover the FoV, ensuring a consistently high throughput.

The limited FoV practically exists in all 60 GHz devices, because the phased array’s front-side comprises many planar antenna elements, each having a FoV of 80° to 180°

[134, 207]. The backside has very weak signal emissions, because it is grounded by a metal plane, and often faces inward the host device which causes strong attenuation. Following this common practice, 802.11ad-compatible laptops (e.g., Acer P446 and Dell Latitude E7240) typically install the phased array close to the outer surface of the lid. The TPCAST [170] 60 GHz adapter is mounted atop a VR headset, with the phased array pointing towards the ceiling. For smartphones, antennas are recommended to be placed on the top or bottom surface [142].

Note that, besides the host device itself, the human user can further block the phased array. In this work, we assume the user of the device always falls outside the phased array's FoV, *i.e.*, blocking the backside only. This is a valid assumption considering the antenna placement principle [168] (in the most unobstructed zone during regular use), and also the fact that the user moves synchronously with the device in practical applications, *e.g.*, mobile gaming and wireless VR. Blockage of other users can be alleviated by careful AP deployment/planning, which will be discussed in Sec. 4.7.

Figure 4.2 (b) also plots the measured throughput as link distance increases, with client and AP falling in each other's FoV. We find the link can sustain a high throughput over a reasonably long distance. Even at 13 m, the throughput remains at 1.8 Gbps, much higher than the out-of-FoV case. This range is consistent with previous measurements [117, 215]. It indicates that *the devices' orientation, which determines FoV alignment, plays a more critical role than link distance in maintaining robust room-level connectivity.*

Profiling the beam-steering under mobility. Can the beam steering on commercial 802.11ad radios be fast enough to adapt to the user movement? To answer this question, we walk inside the network area, holding the client device with natural pose variations, but keeping it within the FoV to an AP. Figure 4.3(a) further plots the throughput measurement over time. Although the throughput shows higher variations, the mobile in-FoV case can still maintain 50% at worst and 80% on average compared to the

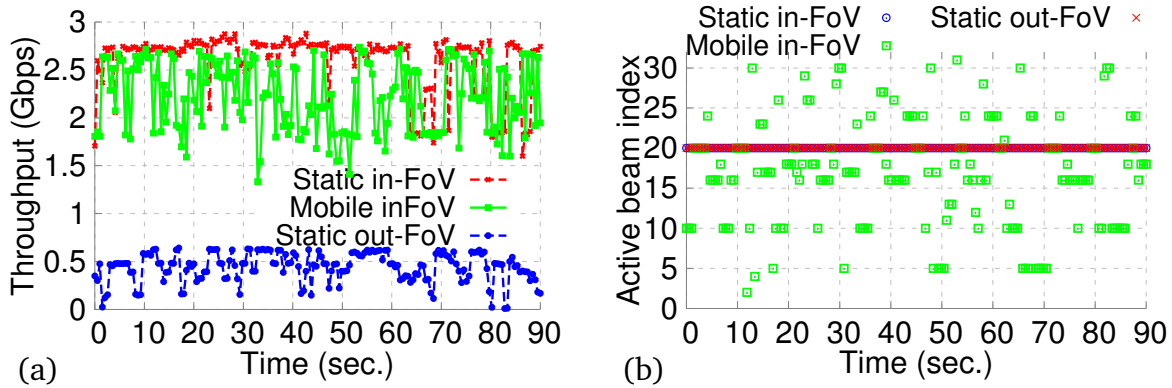


Figure 4.3: Performance impact by the limited FoV: (a) Throughput and (b) AP’s beam index under static/mobile in-FoV and static out-FoV.

stationary case, and $3\sim 5\times$ higher than the out-of-FoV case. The beam index changes rapidly in the mobile case (Figure 4.3 (b)), indicating the real-time beam adaption of 802.11ad can accommodate user mobility with reasonable efficiency, at least at walking speed. So the *in or out of FoV relation becomes the dominant factor that determines link throughput.*

4.2.2 Pia Operations

To facilitate seamless coverage and mobility support for 60 GHz networks, Pia adopts a two-stage workflow (Figure 4.4). The *sensing stage* takes the 60 GHz link status information, *i.e.*, Modulation and Coding Scheme (MCS) and beam index, along with the pose information from the mobile client as input, and runs the *AP-pose sensing (APS)* algorithm (Sec. 4.3.2) to estimate the APs’ pose within the client’s coordinate system. With the same input, Pia also runs a compressive angle estimation algorithm to estimate the location/orientation of major reflectors that may affect the network performance (Sec. 4.4.2). Both schemes are one-time initialization procedures that run before putting the network into use. They need to be repeated only when the APs are redeployed or the environment changes significantly.

During the *running stage*, each client periodically feeds back its pose information

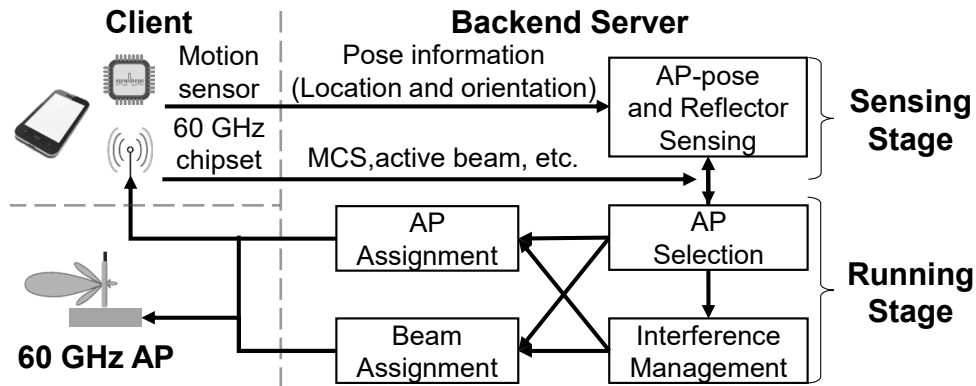


Figure 4.4: Pia’s modules and operation workflow.

to the *backend server*, using its current AP as a relay. The server runs the link prediction (Sec. 4.3.1) and interference management (Sec. 4.4) algorithms, and returns the decisions to the APs/clients. Pia’s decision-making mechanisms are proactive in nature, based on a prediction of the long-term (> 500 ms) pose (Sec. 4.3.1.2). The pose sampling period and feedback latency are relatively negligible compared with the look-ahead time.

4.3 Pia Design

4.3.1 AP Selection

We first consider a single mobile client and focus on how to dynamically select the AP to maintain high link throughput. This issue resembles the classical AP handoff problem in WiFi [63, 115, 128, 159, 162]. A common solution is to probe each AP and select the one with highest potential throughput. However, such approaches are reactive, and more critically, the many-beam phased-arrays in 60 GHz networks will cause huge probing overhead. For example, to probe all possible beams, an 802.11ad client needs to wait for the beam training beacon, which occurs only once per beacon period (typically 100 ms) [118]. By the time a proper AP is identified, the link may have already suffered from a long period of outage. Pia adopts a different design principle. It *predicts* the




Figure 4.5: 5-DoF pose of the phased array.




Figure 4.6: Link status prediction based on the poses.

best AP based on the client's pose, without explicit probing. It proactively decides on *which AP* the client should switch to, and *when*. The AP/beam selection algorithms will execute each time when a new pose data arrives. Pia will inform AP the new beam selection only when algorithm's output changes. We now describe the mechanisms in detail.

4.3.1.1 Pose-Assisted Link Predictor

In Pia, each pose sample, denoted as $\mathbf{P} = [x, y, z, \theta, \varphi]$, represents the 3D location, polar and azimuth angle of the norm of the phased array in a spherical coordinate system (Figure 4.5). We can ignore the rotation angle around the antenna's norm direction because such rotation does not change the FoV alignment relation between the client and AP.

Based on the empirical insights in Sec. 4.2.1, each client discriminates the APs using a binary metric 0/1, denoting whether an AP is estimated to be within or out of FoV. It should be noted that this binary metric is only applied for AP selection purposes. Finer-grained metrics will be used for multi-client interference management (Sec. 4.4.1). In addition, for picocell coverage within a room area, the power budget of commodity 60 GHz devices already provides consistently high throughput for the in-FoV case (Sec. 4.2.1). For larger scale networks such as outdoor small-cells, the AP-to-client distance may have a non-negligible impact, but can be easily modeled in Pia given the

location information.

The predictor takes the AP's and client's poses as input. It predicts the link status as 1 if they mutually fall within each other's FoV (Figure 4.6), and 0 otherwise. We define the FoV angle as from the antenna norm to the edge of its coverage (β_s and β_c in Figure 4.6). For example, the FoV of our device is 85° , measured in Figure 4.2. A device's FoV is generally available from its antenna specifications, or can be manually measured following Sec. 4.2.1.

Let $\mathbf{P}_c = [x_c, y_c, z_c, \theta_c, \varphi_c]$ and $\mathbf{P}_s = [x_s, y_s, z_s, \theta_s, \varphi_s]$ be poses of client and AP, respectively. Pia computes two vectors \mathbf{V}_s and \mathbf{V}_{sc} that point to the norm of AP's antenna and the direction from AP to client, following 3D geometry:

$$\mathbf{V}_s = \mathbf{R}_y(\theta_s)\mathbf{R}_z(\varphi_s)[0, 0, 1]^T \quad (4.1)$$

$$\mathbf{V}_{sc} = [x_c, y_c, z_c]^T - [x_s, y_s, z_s]^T \quad (4.2)$$

where $(\cdot)^T$ denotes transpose operator. $\mathbf{R}_y(\alpha)$ and $\mathbf{R}_z(\alpha)$ are the 3D rotation matrices [191] that rotate a point relative to the origin along Y-axis and Z-axis by an angle of α respectively. Eq. (4.1) converts the orientation of AP's antenna into a vector format. Then Pia can determine the angle ϕ_s between vectors \mathbf{V}_s and \mathbf{V}_{sc} by:

$$\phi_s = \text{atan2}(\|\mathbf{V}_s \times \mathbf{V}_{sc}\|_2, \mathbf{V}_s^T \mathbf{V}_{sc}) \quad (4.3)$$

where function $\text{atan2}(\cdot)$ calculates the four-quadrant inverse tangent [192]. Similarly, Pia computes the angle ϕ_c between the norm of client's antenna and $-\mathbf{V}_{sc}$, and checks if ϕ_s and ϕ_c are smaller than the AP's FoV β_s and client's FoV β_c :

$$\text{Link_State} = \begin{cases} 1, & \phi_s \leq \beta_s \quad \phi_c \leq \beta_c \\ 0, & \text{Otherwise} \end{cases} \quad (4.4)$$




Figure 4.7: Link status prediction error in a sample walking: false in-FoV prediction is highlighted by green triangle. The dense cluster implies multiple errors occur consecutively.

Figure 4.7 plots an example trace of the link predictor, as the client walks within a 2×3 m² region with natural orientation changes (More setup details are available in Sec. 4.6). To obtain the ground-truth link state, we measure the link throughput corresponding to each pose sample, and then convert it into 1/0, using a threshold of 1.6 Gbps that can reliably distinguish the in-FoV and out-FoV (Sec. 4.2.1). We highlight the false in-FoV predictions (*i.e.*, out-FoV predicted as in-FoV) which mislead the client to switch to a low-throughput AP. We find the prediction error is only 9.11% across all the 1324 pose samples. More importantly, the errors mostly concentrate on the boundary area *i.e.*, when transiting from in/out-FoV to out/in-FoV. We will show that such errors have marginal impacts because Pia reacts well before the transition occurs (Sec. 4.3.1.3).

4.3.1.2 Pose Predictor

Pia needs to estimate the link status for the near future based on a prediction of the client’s pose. For simplicity, it adopts a classical kinematic model, Continuous White Noise Acceleration (CWNA) [40], for pose prediction, although other kinematic models can be applied as well. CWNA assumes zero acceleration, *i.e.*, both the linear and angular velocities are stable within a very short duration. Thus, the pose \hat{P}_c at time

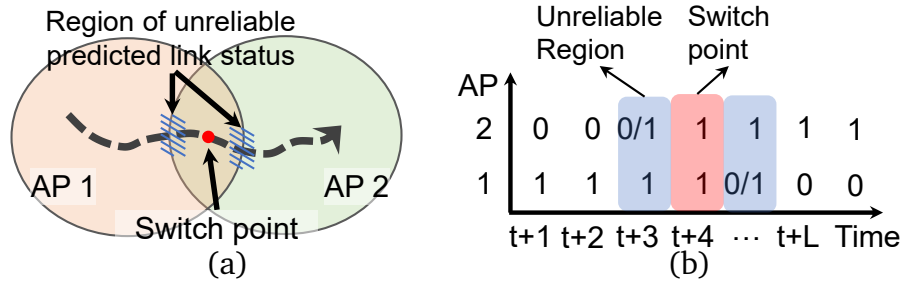


Figure 4.8: (a) AP switching during the transition area. (b) Table of predicted link status for two APs.

$t + 1$ can be predicted as:

$$\hat{\mathbf{P}}_c(t + 1) = \mathbf{P}_c(t) + \Delta\hat{\mathbf{P}}_c(t),$$

$$\Delta\hat{\mathbf{P}}_c(t + 1) = \Delta\mathbf{P}_c(t),$$

where $\Delta\mathbf{P}_c(t) = \mathbf{P}_c(t + 1) - \mathbf{P}_c(t)$ is the pose velocity.

A key question here is how far should Pia predict ahead of time. Let L denotes the look-ahead time. L should be long enough to ensure the AP switching can finish timely, yet not too long to make the prediction unreliable. In our Pia implementation (Sec. 4.5), the AP switching time is negligible (< 10 ms), so we empirically set L to 500 ms, roughly the scale where human movement velocity remains coherent [40].

4.3.1.3 AP Switcher

To avoid the adversarial impact of false in-FoV predictions (Sec. 4.3.1.1), Pia leverages the *transition area* where different APs' FoVs partially overlap (Figure 4.8 (a)). It makes AP switching decision ahead of time before the link status of current AP degrades, so as to avoid the boundary region vulnerable to false prediction. To realize seamless switching, the APs should be deployed with partially overlapping FoVs. However, we will empirically verify that, even when this requirement is not satisfied, Pia's AP switcher can still minimize the client's outage duration (Sec. 4.6.2).

To determine the desired AP and appropriate switching time, Pia's AP switcher takes

the predicted link status for all APs from time $t + 1$ to $t + L$ as input. Figure 4.8(b) shows an example time series of link status (0/1) prediction for two APs, computed based on the pose data of headset and AP using Eq. (4.4). Such information forms a table-like data structure. Pia then makes the switching decision following three basic rules: (i) Prepare a switching only if the current AP has at least one 0 value in the table. Otherwise, the AP will remain in the client's FoV in the look-ahead duration L and switching is unnecessary. (ii) Look for other APs that have overlapped 1s with current AP in the predicted time span, and select the AP with smallest AP-to-client distance as the AP to switch to. This essentially enables Pia to react within the transition area, and choose an alternative AP with potentially strongest RSS. (iii) Choose the center of overlapped 1s as the switching time, so as to avoid the boundary regions vulnerable to prediction errors (Figure 4.8).

There are two exceptional cases that Pia needs to handle in order to minimize the impact of potential link outage: (i) When no other APs have any 1s, Pia needs to stick to the current AP. (ii) When no overlapped 1s exist, Pia will switch to the AP that has the earliest 1 in the predicted time span.

4.3.2 AP-Pose Sensing

The above link predictor assumes the APs' poses are known. Ideally, a user can measure APs' poses during deployment. However, in practice, the phased array antenna is commonly sealed inside the device, and the exact direction where its FoV points to is invisible to the end user. Pia's APS algorithm overcomes this barrier by automatically estimating each AP's pose, which will account for the antenna array placement inside the device. Although certain localization schemes [84] may help obtain the AP's (x_s, y_s, z_s) coordinate, they require at least three reference devices. Pia's use cases typically do not satisfy this requirement because only one reference device, *i.e.*, the mobile client, exists with known position. Besides, these schemes need access to fine-grained PHY-layer in-

formation such as phase, and often require precise carrier clock calibration between the transmitter and receiver, which is not viable on commodity 60 GHz hardware. To enable 5-DoF AP-pose sensing for a wider range of 60 GHz devices, we devise a *statistical algorithm that only needs the link throughput measurement as input*.

Collecting the sensing data. APS is called during the sensing stage of Pia, when the user randomly walks within the network coverage. Meanwhile, the client device records its pose and link throughput with respect to a connected AP. Then, Pia converts the link throughput into binary link status in the same way as the experiments in Figure 4.7. Since a client device can estimate the link throughput for multiple APs simultaneously, a user only needs to collect the sensing data for all APs with a single walk. Besides, the user’s walking trace needs not cover every spot in the area.

Statistical estimation of the AP pose. Let S_s denote a set that contains all legitimate AP’s pose values, and S_t be the pose samples in the sensing trajectory. Then the optimal estimation of the AP pose should result in the minimum error in matching the measured in-FoV/out-FoV samples. We formulate this optimal estimation \mathbf{P}^* as one that minimizes the weighted false in-FoV errors (w_{FI}) and false out-FoV errors (w_{FO}) for a given set of samples:

$$\mathbf{P}^* = \arg \min_{\mathbf{P}_i \in S_s} w_{FI} \sum_{j=1}^{|S_t|} inFoV(\mathbf{P}_i, \mathbf{P}_j) * (1 - t_j) + w_{FO} \sum_{j=1}^{|S_t|} (1 - inFoV(\mathbf{P}_i, \mathbf{P}_j)) * t_j, \quad (4.5)$$

where $|S_t|$ represents the number of samples in S_t , and t_j is the ground-truth link status (0/1 value). Function $inFoV(\cdot)$ determines if the client and AP are within each other’s FoV (Eq. 4.4) and returns a binary (0/1) indicator. \mathbf{P}_i denotes the AP’s pose and is the only variable we need to solve. We choose to minimize the weighted error because the false in-FoV error is more harmful to the link stability than the false out-FoV error (Sec. 4.3.1). Thus, it deserves a larger weight in the optimization ($w_{FI}:w_{FO} = 4:1$ in Pia).




Figure 4.9: AP-pose sensing: Determine the orientation from a 3D ellipsoid that covers all client’s poses in the AP’s FoV.

However, Eq. (4.5) is non-convex, and solving it directly will be computationally expensive due to the large space of S_s . Suppose we partition locations in the unit of 10 cm and rotations in 3° . Then the number of candidate poses for a typical-size room ($120 m^3$) will be 2^{28} .

To prune the search space, we leverage a sweet spot where the axial direction of the AP’s FoV passes through. At this location spot, the in-FoV estimation should have a minimal error. To locate such a spot, consider an example trace of sensing data in Figure 4.9, where we mark all locations with “in-FoV” link status as red dots. Using the approach in [169], Pia finds a 3D ellipsoid of minimum volume covering all the red dots. The ellipsoid’s center is identified as the sweet spot because steering the AP’s FoV to it will maximize the number of red dots the AP covers. Afterward, Pia searches over the 3D space using Eq. (4.5) to determine the AP’s location. The search space becomes much smaller because it only involves 3-DoF. Since the sweet spot uniquely specifies the axial direction of the AP’s FoV, the AP’s orientation can be directly determined for a given location without searching. Algorithm 4.1 summarizes the above procedure of APS. Moreover, Pia can detect the AP position change after initial sensing, because an inaccurate AP pose will significantly reduce the AP selection accuracy (Sec. 4.6.1). When the AP selection reliability degrades below a certain threshold, Pia will notify the user to recalibrate the AP’s pose.

```

1: procedure APposeSensing( $\mathbf{S}_t, \mathbf{T}_t, w_{FI}, w_{FO}$ )
2:    $[x_d, y_d, z_d] = \text{ellipsoidCenter}(\mathbf{S}_t)$  ▷ Find the sweet spot
3:    $\mathbf{D} = \text{APLocList}(\mathbf{S}_t)$  ▷ Generate list of legitimate AP locations
4:    $vMin \leftarrow \text{inf}, \mathbf{P}^* \leftarrow \text{NULL}$  ▷ Initialize variables
5:   for  $[x_s, y_s, z_s] \in \mathbf{D}$  do
6:      $[\theta_s, \varphi_s] = \text{toAngle}([x_s, y_s, z_s], [x_d, y_d, z_d])$  ▷ To orientation
7:      $\mathbf{P}_i \leftarrow [x_s, y_s, z_s, \theta_s, \varphi_s], v \leftarrow 0$  ▷ Candidate AP pose
8:     for  $\mathbf{P}_j \in \mathbf{S}_t, t_j \in \mathbf{T}_t$  do
9:        $v = v + \text{infoV}(\mathbf{P}_i, \mathbf{P}_j) * (1 - t_j) * w_{FI}$ 
10:         $+ (1 - \text{infoV}(\mathbf{P}_i, \mathbf{P}_j)) * t_j * w_{FO}$  ▷ Eq. (4.5)
11:     end for
12:     if  $v < vMin$  then ▷ Minimize weighted error
13:        $vMin \leftarrow v, \mathbf{P}^* \leftarrow \mathbf{P}$ 
14:     end if
15:   end for
16:   return  $\mathbf{P}^*$  ▷ Return the estimated AP's pose
17: end procedure

```

Algorithm 4.1: AP-Pose Sensing.

4.4 Pose-Assisted Interference Management

Despite the 7 GHz of spectrum for 802.11ad, there only exist 4 orthogonal channels. Multiple clients sharing the same channel can become inevitable in dense networks, and when the AP-to-server wireless backhaul needs dedicated channels. The 802.11ad standard [82] specifies a spatial sharing mechanism, where directional links can sense their mutual interference. Concurrent transmissions are allowed only among interference-free directional links. However, to allow other links to sense its interference, each link needs to first transmit in an exclusive service period [82]. This repeats whenever its beam changes, which ultimately reduces spatial reuse, especially when there are multiple mobile links. We now describe how Pia overcomes the challenge.

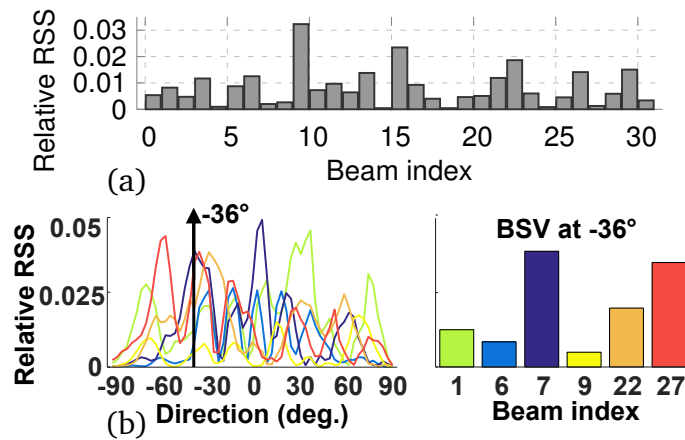


Figure 4.10: (a) Example measured BSV (RSS across all beam patterns for a given angle.) (b) Derive BSV from beam patterns at -36° .

4.4.1 Improving the Spatial Sharing Opportunity

4.4.1.1 Beam Strength Vector

Pia uses *beam strength vector* (BSV) as the core data structure to arbitrate the AP beam assignment for multi-client scenarios. BSV is defined as a vector characterizing the RSS values (*w.r.t.* the AP) across different transmit beam patterns. An example BSV is shown in Figure 4.10 (a), measured using our 60 GHz AP and a customized software radio (Sec. 4.5) separated by 2 m. To obtain the BSV in commodity 60 GHz devices, there are two approaches:

Measurement from radio. First, one can leverage 802.11ad’s built-in beam training protocol. At the beginning of each Beacon Interval (BI), the AP transmits multiple beacon frames, each using a different beam pattern. The client can record the RSS of each transmit beam pattern, which forms the BSV. This approach doesn’t rely on the pose information and can be applied to general scenarios.

Prediction from pose. However, an AP sends out beacon frames at a fixed time interval, e.g., 100 ms, which is too long for delay-sensitive VR applications. Instead of waiting until the beginning of a BI, Pia chooses to leverage the pose information to directly derive the BSV. It first calculates the polar/azimuth angle looking from the

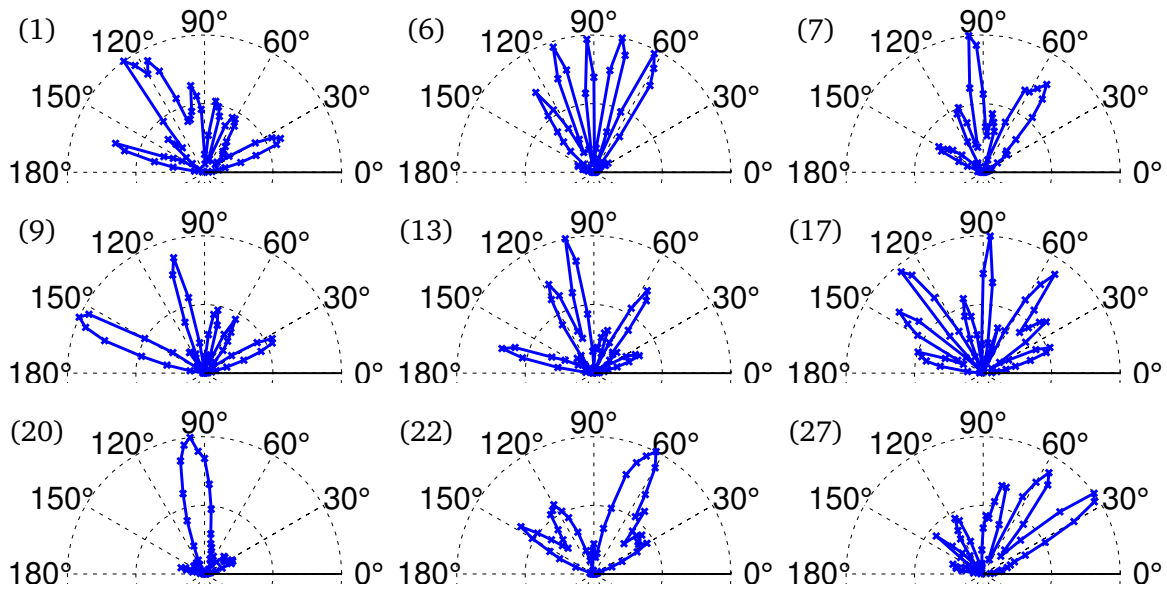


Figure 4.11: Samples of measured beam patterns of the phased array antenna. Left-top number is the beam index.

AP's pose to the client's position (similar to Eq. (4.2) and (4.3)). It then looks up the gains of all beam patterns in the view-angle direction. The angular beam pattern can be obtained from the phased-array datasheet or using a one-time measurement (Sec. 4.5). Figure 4.10 (b) illustrates a 2D example of this procedure, where the left plots the gains of different beam pattern indexes at different angles, and the right shows the corresponding BSV at angle -36° .

4.4.1.2 Pose-Assisted Spatial Sharing

To enhance the spatial reuse opportunity, we introduce a novel interference management framework, which leverages the pose information to simultaneously determine the *best AP and beam pattern* for each client. Our key insight is that practical phased-arrays are not perfectly directional. Due to the use of discretized beamforming weights (*i.e.*, codebook entries) [134], each beam pattern may bear a main lobe along with multiple weaker side lobes. Figure 4.11 illustrates example beam patterns of our AP device,




Figure 4.12: An example BSM for 3 APs and 3 clients.

measured using a software radio (Sec. 4.5)¹. The results are consistent with previous measurement study [117] using similar hardware. Therefore, instead of greedily choosing the beam with strongest main lobe for each link and then measure whether the links can coexist without interference as specified by 802.11ad [82], Pia can judiciously assign beams with weaker lobes, which may sacrifice the RSS of certain links but create more spatial reuse opportunities for them.

We formulate this basic idea as a joint AP selection and beam selection problem. For clarity, we will focus on the downlink and optimize the AP’s beam selection because the downlink traffic will be the dominance in the VR application. However, the extension to uplink and client-side optimization can follow the same principle since the client’s pose is available. Our objective is to find the optimal AP and beam assignment that maximizes the overall link quality using *beam strength map (BSM)*, a data structure that aggregates the BSVs between each pair of AP and client. Each row of BSM is a BSV for a client and an AP. Figure 4.12 plots an example measured using our 3-AP 3-client testbed (Sec. 4.5). To account for the signal attenuation over distance, Pia subtracts the

¹Practical phased-array beam patterns deviate from the horn shape because: (i) The OEM codebook does not account for the electromagnetic impact from peripheral electronic components on the phased array. (ii) To reduce the manufacturing cost, the 60 GHz phased-arrays only have limited phase-shift resolution (2-bit for each antenna element on the Qualcomm platform we use, *i.e.*, switching among 0° , 90° , 180° , and 270°). (iii) Our radio uses the same beam pattern for both beacon frames and data frames. The former is expected to have more diverse and wider beams [82].

predicted BSV by the pathloss of Frii's model [190] which takes the AP-to-client distance as input. Note that Pia does not need to predict client's absolute RSS because the spatial sharing mechanism only relies on the ratio between BSVs.

Suppose there are N_c clients and N_s APs. Let $\mathbf{A}(i)$ and $\mathbf{B}(i)$ be the AP and beam assignment for client i . Ideally, to maximize network capacity via spatial reuse, Pia should maximize the *Signal-to-Interference Ratio (SIR)* of all clients:

$$\max_{\mathbf{A}, \mathbf{B}} \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{BSM[\mathbf{A}(i), i, \mathbf{B}(i)]}{\sum_{j=1}^{N_s} INF_{max}(j, i)}, \quad (4.6)$$

where $BSM[\hat{a}, i, \hat{b}]$ denotes the signal strength from AP \hat{a} to client i using beam \hat{b} , and function $INF_{max}(\hat{a}, i)$ calculates the maximum interference that AP \hat{a} can cause to client i , which is given by:

$$INF_{max}(\hat{a}, i) = \max_{j, \mathbf{A}(j)=\hat{a}} INF(j, i),$$

$$INF(j, i) = \begin{cases} BSM[\mathbf{A}(j), i, \mathbf{B}(j)], & \mathbf{A}(j) \neq \mathbf{A}(i) \\ 0, & \mathbf{A}(j) = \mathbf{A}(i). \end{cases}$$

However, directly solving Eq. (4.6) entails a high runtime complexity — $\mathcal{O}((N_s N_b)^{N_c})$, where N_b is the number of beam patterns in the phased array. This above SIR-based solution is not scalable to the client number. Moreover, in our evaluation (Sec. 4.5), we found it cannot process the pose input in real time for even two clients. Instead, we design a lightweight algorithm that maximizes the *Signal-to-Leakage Ratio (SLR)*. The leakage is defined as the signal strength sent by an AP and received by an undesired client. SLR computes the ratio of desired client's RSS to leakage. Maximizing the SLR will also enhance the SIR because it minimizes the leakage that causes interference to others. SLR reduces the optimization complexity by decomposing the dependency between clients in AP assignment and beam assignment.

(i) *AP assignment.* To maximize the SLR, it is best to assign the client to an in-FoV

```

1: procedure assignSLR(BSM[], Nc)
2:   for i = 1 : Nc do
3:     A(i) ← indexOfMax(mean(BSM[:, i, :].'))           ▷ Assign AP
4:   end for
5:   for i = 1 : Nc do
6:     for j = 1 : Nc do
7:       SLR[j, :] ←  $\frac{BSM[A(i), i, :]}{BSM[A(i), j, :]}$            ▷ Leakage from client i to j
8:     end for
9:     B(i) ← indexOfMax(min(SLR[:, :]))           ▷ Assign beam
10:  end for
11:  return A and B           ▷ Return AP and beam assignment
12: end procedure

```

Algorithm 4.2: SLR-based AP and Beam Assignment.

AP at closest distance, because it maximizes the RSS while minimizing the leakage from other APs. Pia thus assigns the client to an AP with the maximum average BSV across all beam indexes (e.g., client 1 is assigned to AP 1 in Figure 4.12). When assigned to the same AP, clients can work under the 802.11ad Time-Division Multiple Access (TDMA) mode.

(ii) *Beam assignment.* Upon a new client assignment, the AP first iterates through each beam index, and calculates the RSS of current client and maximum leakage signal strength to other clients, which forms the SLR. The AP then chooses the beam assignment that has the maximum SLR value. Finally, it examines whether the resulting SIR for each client exceeds its packet decoding threshold; if not, it assigns a dedicated time slot for it. Algorithm 4.2 summarizes the SLR-based algorithm. Its runtime complexity can be straightforwardly derived as $\mathcal{O}(N_c N_b (N_s + N_c))$, and it is efficient enough to run in real time. Though the SLR-based algorithm is suboptimal, our evaluation shows its performance is close to the SIR-based algorithm.

Here are a few aspects that worth further discussion:

60 GHz AP deployment. Based on our experiments, we identified 3 general guidelines to maximize the coverage and effectiveness of Pia in a multi-AP network. (i) Deploying the APs higher than typical human heights, such as ceiling, that will maxi-

mize the AP's FoV to the client and minimize blockage of other passing-by people. (ii) However, placing the AP to the ceiling cannot prevent the link blockage. When a user walks away from the AP's beneath center, the client's FoV is still prone to moving out of the AP's FoV. To ensure multi-Gbps coverage everywhere, the APs should be densely deployed. (iii) The FoV of APs should be partially overlapped to minimize the number of blind spots. Besides, it is worth noting that Pia is not tied to any specific ways of AP deployment.

Pose information availability. Our evaluation results show Pia can tolerate 0.5 m location error and 0.35 radians orientation error. Current mobile devices, combining motion sensors, light sensor, camera, *etc.*, can already provide accurate pose tracking. A^3 [212] can estimate 3D orientation $<5^\circ$ using accelerometer and magnetometer. Many mobile localization systems [91, 92, 104, 155, 189] can track users at centimeter to decimeter accuracy. Commercial products such as Google Tango [67] can provide the 6-DoF tracking at cm-level accuracy, by using motion sensors and a depth-camera. VR headsets like HTC Vive [79] can tracking user motion at mm-level precision [59], which is the ideal candidate for our system. All of these systems suffice to support Pia's pose-assisted design principle.

Extension to general mobile 60 GHz networks. Here, we tailored the system design to match a typical wireless VR setup in a constrained room environment. However, the design principles can be extended to general millimeter-wave networks to support mobility and seamless coverage. Although the absolute location tracking technologies for mobile devices may not be sufficiently accurate and reliable, the relative motion tracking (*e.g.*, rotation change and moving offset) using built-in visual-inertial sensors can suffice for Pia. To alleviate the accumulated drift of relative tracking, Pia may periodically apply the pose calibration (Sec. 4.4.3). Moreover, to mitigate the initial training overhead, we could deploy APs to known locations and develop an online-sensing method to infer the phased array's orientation (inside the AP) and reflectors' poses dur-

ing usage. Though it may trigger many error predictions at the beginning, the accuracy could improve over more collected data, which we will leave for our future work.

4.4.2 Dealing With the Environmental Reflectors

The above scheme geometrically models the interference coming from transmitters within direct LOS. In practice, when the transmitters become close to strong ambient reflectors, their beams may be redirected, causing NLOS interference. To account for such effects, Pia incorporates a novel reflector sensing scheme which enhances the model in Sec. 4.4.1.2. It requires the user to conduct a one-time environment learning, by placing the 60 GHz client device near each strong reflector (typically concrete walls and large metal furniture [99]). Due to the well-known sparsity of 60 GHz channel [165], from 60 GHz radio's eyes, there are only a few dominant reflectors, which can account for more than 95% of the total energy that can be received by the radio in a typical indoor environment [188]. The selective scanning reduces the human labor and accelerates the sensing process by avoiding many unnecessary positions. However, we cannot let the user directly input the reflector coordinate since the reflectivity is unknown. Pia uses a novel signal angle sensing algorithm (Sec. 4.4.2.1), which fuses the pose sensor measurements with the BSV measurement, to estimate the relative location/orientation of reflectors. During the SLR estimation, Pia accounts for the NLOS interference from each reflector by modeling it as a virtual transmitter (Sec. 4.4.2.2). Moreover, strong reflectors unlikely move frequently due to the size and weight. We only need to rerun environment learning when their positions change.

4.4.2.1 Pose-Assisted Compressive Angle Estimation

The key idea in our NLOS interference model is to reverse engineer the reflectors' impact by tracing back NLOS signals' Angle of Departure (AoD) and Angle of Arrival (AoA). Such signal angles may be measured using a cone-shaped antenna with perfect direc-

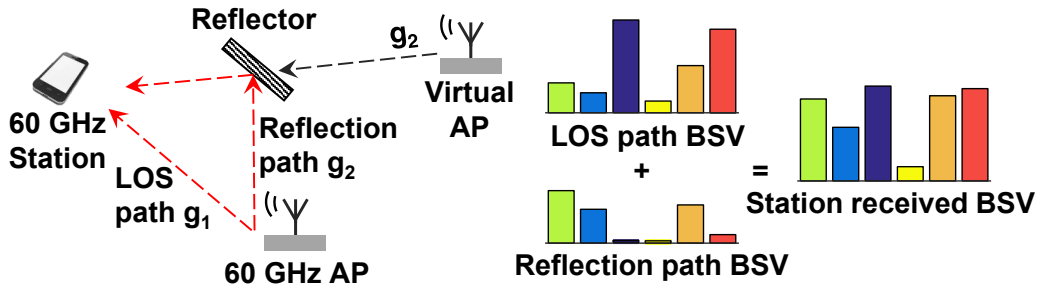


Figure 4.13: Compressive angle sensing and reflector estimation: measured BSV is the sum of BSVs along two paths.

tionality [217]. However, our goal is to enable angle sensing in practical 60 GHz devices which have imperfect beam patterns (Figure 4.11). Our solution originates from a key observation: Whenever the reflector’s impact becomes non-negligible, the measured BSV will deviate from the one predicted by using the LOS model (Sec. 4.4.1.2), and become the superposition of two BSVs along the LOS path and reflection path (Figure 4.13).

Following the same approach as Sec. 4.4.1.1, Pia first obtains the measured BSV, from which it then *estimates the number of paths and their angles*, based on a statistical optimization model. For simplicity of exposition, we focus on the AoD estimation, but the AoA applies in the same way. We first uniformly divide the AP-to-client view angles (θ, φ) into N_p pairs following the geodesic grid [193]. Let $\mathbf{V}(\theta_i, \varphi_i)$ denote the predicted BSV (Sec. 4.4.1.1) for the i^{th} pair of direction (θ_i, φ_i) , and g_i be the unknown channel gain along the BSV direction. The measured BSV \mathbf{V}_r can be expressed as the sum of BSVs from all directions multiplied by their associated gains: $\mathbf{V}_r = \sum_{i=1}^{N_p} \mathbf{V}(\theta_i, \varphi_i)g_i$. Thus, we may estimate the AoD by solving g_i in the system of equations:

$$\mathbf{V}_r - \sum_{i=1}^{N_p} \mathbf{V}(\theta_i, \varphi_i)g_i = 0 \quad (4.7)$$

where the direction (θ_i, φ_i) associated with a non-zero g_i will be the AoD of one path.

The problem (4.7) is under constrained because $N_p > N_b$. For example, to achieve

an AoD estimation granularity of 8° , we need to uniformly divide the 3D view angle into 337 directions according to the geodesic grid [193]. Yet, for a phased-array of 32 elements, the number of beam patterns is only 64 following the standard practice of codebook design [180]. To address this problem, we harness the channel sparsity of the 60 GHz channel [165], *i.e.*, most AoD directions have close to zero gain (*i.e.*, $g_i \approx 0$), and only those corresponding to the LOS and a few strong reflection paths are non-negligible. We thus stack the g_i into an $N_p \times 1$ column vector $\mathbf{G} = [g_1, \dots, g_{N_p}]^T$, which is now sparse. Thus, we can reformulate Eq. (4.7) into a compressive sensing problem that minimizes the l_1 norm of the sparse vector \mathbf{G} ,

$$\min \|\mathbf{G}\|_1 \quad \text{subject to} \quad \|\mathbf{V}_r - \hat{\mathbf{V}}\mathbf{G}\|_2 \leq \varepsilon \quad (4.8)$$

where $\hat{\mathbf{V}} = [\mathbf{V}(\theta_1, \varphi_1), \dots, \mathbf{V}(\theta_{N_p}, \varphi_{N_p})]$ is an $N_b \times N_p$ matrix, and ε is determined by the radio's noise power in BSV measurement. Eq. (4.8) can be solved efficiently in polynomial time using the ℓ_1 -MAGIC [46] or CVX [68] toolbox.

4.4.2.2 Modeling the Reflector Impact

Given the AoA/AoD, Pia can back-trace the departure and arrival paths, and locate the intersection points as the estimated reflector positions. Let $\mathbf{P}_c = [x_c, y_c, z_c, \theta_c, \varphi_c]$ and $\mathbf{P}_s = [x_s, y_s, z_s, \theta_s, \varphi_s]$ be the poses of client and AP, and (θ_a, φ_a) and (θ_d, φ_d) be the estimated AoA and AoD. The AP's pose can be estimated by the APS algorithm (Sec. 4.3.2). The AoA path travels along a series of points with coordinate:

$$\mathbf{K}_a(l_a) = \mathbf{R}_y(\theta_c)\mathbf{R}_z(\varphi_c) \begin{bmatrix} \cos(\varphi_a) \sin(\theta_a)l_a \\ \sin(\varphi_a) \sin(\theta_a)l_a \\ \cos(\theta_a)l_a \end{bmatrix} + \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix} \quad (4.9)$$

Eq. (4.9) first rotates a point at the AoA angle (θ_a, φ_a) *w.r.t.* the origin by an offset of

client's orientation (θ_c, φ_c) , and then translates the coordinate by an offset of client's location $[x_c, y_c, z_c]$. The only unknown variable in Eq. (4.9) is l_a that controls the distance between point $\mathbf{K}_a(l_a)$ and $[x_c, y_c, z_c]$. In a similar way, we back trace the AoD path and form a series of points with coordinate $\mathbf{K}_d(l_d)$ with l_d as the variable.

In practice, the AoA and AoD paths might not intersect in 3D space due to the residual error of angle sensing. Pia will select a point with closest distance to the two paths. Toward this goal, Pia determines the two points $\mathbf{K}_a(l_a^*)$ and $\mathbf{K}_d(l_d^*)$ with l_a^* and l_d^* that minimize the distance between them:

$$(l_a^*, l_d^*) = \arg \min_{(l_a, l_d)} \|\mathbf{K}_a(l_a) - \mathbf{K}_d(l_d)\|_2. \quad (4.10)$$

The closest intersection point, *i.e.*, the estimated reflector position, \mathbf{K}_r is their median center: $\mathbf{K}_r = (\mathbf{K}_a(l_a^*) + \mathbf{K}_d(l_d^*))/2$.

To capture the reflector's impact, Pia adopts the concept of virtual AP (vAP) – a virtual signal source that mirrors the position of real AP relative to the reflector (Figure 4.13). This idea is inspired by the 60 GHz ray-tracing method [56, 57, 93], which has proven to predict the 60 GHz signal propagation reasonably close to real measurement. A reflector will redirect the signals so that they look like emitting from the vAP. Given the client position $\mathbf{K}_c = [x_c, y_c, z_c]^T$ and reflector position \mathbf{K}_r , the vAP position \mathbf{K}_v can be derived from geometry (Figure 4.13), by moving from the reflector position along the AoA direction over distance l_a^* :

$$\mathbf{K}_v = \frac{l_d^*(\mathbf{K}_r - \mathbf{K}_c)}{l_a^*} + \mathbf{K}_r. \quad (4.11)$$

To determine the vAP's orientation (θ_v, φ_v) , Pia solves equation:

$$\mathbf{R}_y(\theta_v)\mathbf{R}_z(\varphi_v) \begin{bmatrix} \cos(-\varphi_d) \sin(\theta_d) \\ \sin(-\varphi_d) \sin(\theta_d) \\ \cos(\theta_d) \end{bmatrix} = \frac{\mathbf{K}_c - \mathbf{K}_r}{l_a^*}, \quad (4.12)$$

where the left side forms a unit vector from vAP's orientation (θ_v, φ_v) toward vAP's AoD direction $(\theta_d, -\varphi_d)$. The unit vector should be at the same direction as the one from reflector's position to client's position (*i.e.*, right side). The azimuth angle of vAP's AoD is negative because it is the mirror symmetry of the real AP.

Once the pose of vAP is determined, it will be used in Pia's spatial sharing module (Sec. 4.4.1.2) in the same way as a real AP, except that vAP should account for the reflection loss of reflector (estimated in the channel gain g_i). This procedure of vAP localization should be repeated for each reflector during the sensing stage.

4.4.3 Calibrating the Client's Pose Offset

Since the outputs of APS and reflector sensing will be used by other client devices, in order for Pia to operate correctly, each later joined client that does not run the sensing stage, needs a one-time calibration for its pose. The pose information on mobile devices may contain an unknown initial offset *w.r.t.* the AP's coordinate due to two reasons. (i) Not all devices are capable of reporting absolute location/orientation. Many accurate and mature indoor navigation technologies [67, 91, 92, 155, 189] can only measure relative location/orientation change based on visual-inertial sensors. (ii) The pose, especially orientation, of a phased array antenna may not be the same as the pose of its host device.

Therefore, Pia needs to estimate an initial pose for each newly joined client *w.r.t.* the AP's coordinate, so that it does not need to rerun the APS and reflector sensing. Our solution is designed upon the compressive angle estimation (Sec. 4.4.2.1), following




Figure 4.14: Estimate the initial pose of the client using compressive angle sensing.

two steps:

(i) *Determining the in-FoV APs.* When a client first enters the network area and moves around, it searches for two nearby APs within the FoV. There are multiple ways to determine if an AP is in the FoV. For example, the client can estimate the link throughput by sending data to the backend server and using the throughput thresholding heuristic in Sec. 4.2.1. Alternatively, the client can estimate the RSS from the beacon frames. Then, the client estimates the AoAs and AoDs *w.r.t.* the LOS APs (Figure 4.14). Note that the sensed angle is relative to the device's phased array antenna rather than the AP's coordinate. Then, the two APs share their pose information (estimated from APS) with the client. (ii) *Determining the initial pose.* The two LOS paths originating from two APs intersect at the client's location (Figure 4.14). Given the APs' locations and the AoD of the LOS paths, a simple geometrical calculation can pinpoint the client's location. Pia can determine the client's orientation in the AP's coordinate, in a similar way as finding the vAP's orientation in Sec. 4.4.2.2. The estimated location and orientation together form the initial pose of the client. It is worth noting that the initial pose estimation requires the AP fall in the client's FoV. Thus, it cannot be applied to assist the network operations, such as AP switching.




Figure 4.15: (Left) Pia’s hardware components, phased-array in the red circle. (Right) Multi-AP configuration.

4.5 Implementation

Multi-AP 60 GHz network testbed. We prototype Pia based on a multi-AP 60 GHz testbed, as shown in Figure 4.15. The testbed consists of three APs and one client, all equipped with the Qualcomm 60 GHz wireless network interface card (with QCA6310 32-element phased array/RF front-end and QCA6320 MAC/baseband). Each AP is a small PC which interfaces the card through an M.2-to-PCIe adapter. We mount the phased array antenna on a flat panel that can adjust its tilt angle. Each client is a laptop (Acer P446-M59BB) with the phased array built in the outer side of the laptop’s cover.

Although the server-to-AP backhaul links can use 60 GHz interfaces with fixed beams, our prototype implements them using 10 Gbps Ethernet with a NETGEAR XS708E switch. We configure the two-hop network into two subnets, and the APs simply act as relays between client and the backend server (Figure 4.15). To switch the AP, Pia modifies their gateway, and the data will be forwarded along the new AP. The switching latency is less than 5 ms (Sec. 4.6.2).

Extracting radio information from Linux driver. We use the latest wil6210 driver [98] for the 60 GHz wireless adapters. Since the firmware (wil6210.fw) and phased array codebook (wil6210.brd) are still unavailable in the Linux firmware library, we port them from a Windows driver (Atheros Sparrow 11ad) released by Acer. We modify the wil6210 driver and make it quickly (131 entries/second) export the low-layer informa-

tion to user space, including instant (Tx/Rx) throughput, signal quality, active (Tx/Rx) beam index, and (Tx/Rx) MCS. Such information is used as input to Pia’s modules.

Client pose. We attach a Google Tango tablet [67] to the laptop’s outer surface, which can track the laptop’s relative pose at cm-level and degree-level accuracy [52, 61] with less than 80 ms latency [16]. We will evaluate the impact of pose errors in Sec. 4.6. We develop an Android application that converts the quaternion rotation into pose (Sec. 4.3.1.1) and streams it to the laptop via USB at 297 entries/second.

Pia components. We implement the major design components (AP selection, APS, spatial sharing, signal angle sensing) in Matlab and C. The algorithms run on the back-end server (Figure 4.1) which is a commodity PC with i7-3770k CPU and 16 GB memory. To achieve millisecond AP switching latency, many legacy 802.11 drivers (*e.g.*, ath5k, ath9k, ath10k [14] and iwlwifi) support mac80211 VIF [17] which allows one client to maintain multiple virtual interfaces to different APs [47, 90], thus eliminating the re-authentication and re-association that account for the major overhead in AP switching. Since such support is still under development in wil6210, we achieve rapid AP switching through reverse-tethering, *i.e.*, configuring the AP into the managed mode and client into the AP mode, which allows the laptop to maintain connections to multiple senders.

Measuring beam patterns. Pia’s interference management requires the phased array’s beam patterns (Sec. 4.4). Since our 60 GHz radio’s hardware specifications are not public, we measure them using the same 60 GHz software radio as in [165, 188, 211]. We place a receiver radio (with a 20° horn antenna) 2 m away from an AP and point its antenna center to the AP’s phased array. We mount the AP on a 3D motion controller that can rotate its azimuth/elevation angle at a step of 3°, allowing us to measure the RSS from different directions (*i.e.*, angular beam pattern). Since the beam switching on the 60 GHz wireless card is controlled by closed-source firmware, we are currently unable to switch the phased array’s beam directly. In order to obtain the beam patterns for all beam indexes (31 in total), we measure the RSS values of beacon frames at the

beginning of each beacon interval, which are sent sequentially by the AP, each using a different beam pattern, following the 802.11ad standard [82, 118]. We predict the BSV using pose data from the measured beam patterns. Besides, we find that in our platform the beacon frames share the same codebook as the data frames. Even though they may be different in other radios, we can still follow the same principle and measure beam patterns of the data frame.

4.6 Evaluation

Methodology: We deploy the 60 GHz testbed in a 7×8 m² office environment, with 3 APs mounted on stands (2 m height) near the corner of the room, forming a triangle and facing 45° towards the ground. We choose the corner positions because of the physical constraints and convenience, just like the placement of typical home routers. Besides, it is also the best way to stress-test our system’s capability since the client needs to perform AP switching more frequently. Pia can easily scale to a larger space with more APs.

To bootstrap the system, we first generate a set of legitimate AP poses to feed APS (Sec. 4.3.2). Pia creates a bounding box using the rough room size input by the user, and then partitions the possible AP coordinate by a step of 10 cm. Further increasing the granularity will not improve the AP-pose sensing accuracy because it is already close to the system limit (Sec. 4.6.1). We then execute the AP switching in real time (Sec. 4.3.1). The current Qualcomm firmware does not support real-time beam control and RSS feedback for each beacon frame. To evaluate the interference management (Sec. 4.4.1), we collect BSV data over 100 random client poses for the commodity 60 GHz AP using our customized 60 GHz Rx (Sec. 4.5). We then use the BSV data as input to evaluate the performance of Pia’s spatial reuse, angle sensing, and reflector estimation, which best approximate the end-effect of an actual device.




Figure 4.16: System initial training overhead: (a) AP-pose sensing, (b) AP selection accuracy, over user’s training effort—walking distance.

Metrics: In the micro benchmarks, we focus on metrics directly related to the design modules’ performance, *e.g.*, AP selection accuracy, pose error, SIR, and angle sensing error. The system level tests will focus on network-level metrics: throughput, latency, and reliability. We measure the achievable throughput by iperf3 [83] and the network latency by sending back-to-back ping packets. By default, error bars in all results denote the 90-percentile error.

4.6.1 Micro Benchmarks

4.6.1.1 Effectiveness of APS and AP Selection

Impact of training. We first evaluate how the sensing stage affects APS (Sec. 4.3.2) by collecting training data while walking randomly over 15 meters, which only takes a user less than one minute to walk through. We partition the collected data into multiple sizes to evaluate the impact of data size on the sensing accuracy. Figure 4.16 (a) shows that both the location error and orientation error decrease dramatically over longer walking traces. When the walking distance is <5 meters, the estimation error tends to be large, because Pia doesn’t gather sufficient data to statistically filter out erroneous locations. With >11 m training, the AP’s pose error stabilizes at 0.29 m and 0.17 radian (*i.e.*, 9.7°), which indicates *the APS algorithm can quickly estimate AP’s pose with minimal training efforts.*

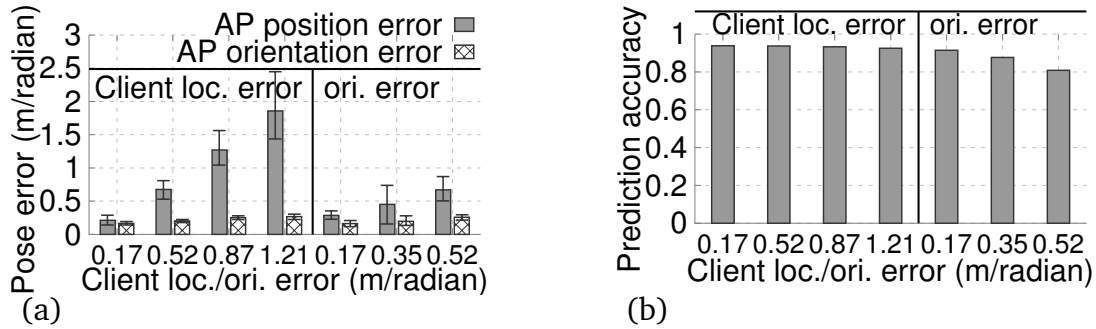


Figure 4.17: System sensitivity to client’s pose error: (a) AP-pose sensing, (b) AP selection accuracy, over client’s pose error.

We then use the collected data onwards for Pia’s link predictor. Figure 4.16 (b) shows that with distance >11 m, Pia’s AP selection reliability (i.e., choosing an in-FoV AP) can maintain at 94.4%. Pia can achieve a high prediction accuracy using relative small training efforts because it exploits the geometrical structure of the FoV.

Impact of client’s pose information error. We then inject Gaussian noise to the collected data, and vary the mean from mm to decimeter-level, consistent with the range of pose errors on mobile devices (Sec. 4.7). The resulting APS location estimation error (Figure 4.17 (a)) grows linearly with the client’s pose error. Fortunately, since many recent mobile location tracking schemes [173, 206, 213] can already provide at least 10~30 cm accuracy, the estimated AP pose is sufficient for reliable AP selection. Besides, since the APS is a one-time initialization, specialized tracking devices such as Tango can be used to ensure high precision. In addition, the APS orientation is not sensitive to client’s pose error because the sweet spot of ellipsoid center (Sec. 4.3.2) averages out the deviation.

Figure 4.17 (b) plots the impact of pose error on link prediction. The location error has a limited impact—even for 1.21 m client position error, Pia’s link prediction accuracy only decreases by 2.1%. Yet, 0.52 radians (30°) of orientation error reduces the prediction accuracy by more than 10%, because orientation error more easily deviates the radio’s FoV. Fortunately, recent 3D orientation tracking systems can maintain the error below $<5^\circ$ [212], as accurate as Tango. Therefore, Pia can run on mobile devices




Figure 4.18: Packet latency during the AP switching.

with more than 90% AP selection accuracy.

AP switching latency. To verify if the client can seamlessly switch across APs, we force it to switch between two APs per second. Figure 4.18 plots the round-trip latency between the client and backend server, from which we observe three patterns. The latency is less than 1 ms when measure occurs in a beacon interval (BI), and increases to 3 ms at the beginning of BI due to the beaconing overhead. The 3 ms latency follows the 100 ms periodicity that well matches with the AP's beacon interval. The AP switching causes additional 10~25 ms packet latency, because of the beam training overhead with the new AP. Yet, every time it only affects a single packet. Besides, the pose information can be exploited to reduce the beam adaptation overhead, which we leave for future work.

4.6.1.2 Effectiveness of Compressive Angle Estimation and Interference Prediction

Accuracy of signal angle sensing. The performance of Pia's compressive angle estimation (Sec. 4.4.2) is determined by the Signal-to-Noise Ratio (SNR) of BSV measurement. To evaluate the impact, we first measure BSV under higher SNR (>30 dB). Since BSVs are obtained by measuring the detailed waveform of beacon frames, we can add Gaussian noise to the original waveform and create BSVs for different SNRs. The CDF plot in Figure 4.19 (a) shows an average angle estimation error of less than 5° when SNR




Figure 4.19: Angle sensing error (a) over the SNR of BSV, (b) using multiple BSVs (SNR=6 dB).




Figure 4.20: Estimated gain error over SNR of BSV.

is high (e.g., >9 dB). Residual errors are caused by discretization of the view angles (Sec. 4.4.2). The angle error may drift a lot under low SNR, when the BSV is significantly distorted by noise. However, by simply repeating the BSV measurement by 5 times and averaging out, angle error can be reduced to less than 5° (Figure 4.19 (b)). Such repetitive measurements add little extra sensing time (only 5 BIs for 5 measurements).

Pia's angle sensing can also estimate each path's channel gain (Sec. 4.4.2.1). Figure 4.20 plots the CDF of estimation error, where we obtain the ground-truth by dividing the measured BSV by the predicted one (Sec. 4.4.1.1). A higher SNR leads to more accurate gain estimation (error <3.4 dB when SNR >9 dB). Notably, the estimated gain is generally smaller than the ground-truth because compressive sensing solver spreads the energy to other directions to minimize the residual error.

Accuracy of client-pose and reflector estimation. Since the client's pose estimation and reflector estimation rely on the compressive angle estimation, we run them




Figure 4.21: Pose estimation error using angle sensing.




Figure 4.22: Accuracy of interference prediction.

consecutively. We first repeat the former at 20 random spots and compare it with ground-truth that measures *w.r.t.* AP's coordinate using a laser range-finder. The result (Figure 4.21) shows a mean position offset of 0.5 m. We further place a strong reflector (metal sheet) and randomly vary its position around the Tx/Rx from 1 to 3 m. The reflector position estimation shows similar accuracy, albeit requiring 5 BSV measurements due to low SNR (reflected signal is 10~15 dB weaker than LOS). Notably, the positioning accuracy is comparable to state-of-the-art indoor localization systems [155, 206, 209, 213], and the orientation accuracy $< 5^\circ$ suffices for Pia to run reliably (Figure 4.17 (b)).

Interference prediction accuracy. Pia estimates the interference using predicted BSV (Sec. 4.4.1.2). To verify its accuracy, we calculate the signal strength difference between the collected BSV data and the predicted one subtracted by the pathloss of Frii's model (Sec. 4.4.1.2). We found Pia can reliably predict the strength of LOS paths based on their poses and active beam indexes regardless of the client-to-AP distance, with a small average error of 0.74 dB (Figure 4.22). The prediction error over reflected paths is larger (< 4 dB) due to extra errors from reflector estimation. Yet, Pia can tolerate this residual prediction error by using a more conservative decodable threshold (Sec. 4.4.1.2).

Performance of SLR-based assignment. The SLR-based algorithm evades exponential searching space for optimal AP and beam assignment (Sec. 4.4.1.2). Figure 4.23 (a)




Figure 4.23: (a) SIR difference, (b) Fairness difference, between SLR-based and SIR-based assignment.

and Figure 4.23 (b) plot the SIR and Jain’s fairness [85] difference between the SLR algorithm and the optimum (*i.e.*, maximizing average SIR). We found at least 50% of the SLR outputs are the optimal assignment (*i.e.*, zero SIR difference), and 90% have <3 dB difference, implying the SLR-based algorithm performs close to the optimum. Furthermore, for those suboptimal assignments, the SLR-based algorithm more likely gives better fairness (*i.e.*, negative fairness difference), because the optimal solution may sacrifice fairness to maximize the SIR.

Since we do not have real-time control over the radio’s transmit beam index, we cannot quantify the impact of client’s moving speed. The beam selection (Sec. 4.4.1.2) may not catch up, for instance, if the large pose estimation delay (*i.e.*, time between when pose changes and when estimation changes) is large compared to moving speed. In this case, the radio may have to fall back to the 802.11ad standard. However, for the VR system, since pose information feedback is also used for graphics rendering, we believe the pose estimation delay is sufficiently small for Pia to operate correctly under the VR-motion speed. We will have more exploration in the future work once the hardware grants us a better control.




Figure 4.24: (a) Link availability, (b) Hazard times, of five AP switching schemes.

4.6.2 System Level Tests

4.6.2.1 Network Robustness by AP Switching

We conduct a system-level test of the AP switching by emulating the wireless VR scenario, where the client follows a random walk and pan/tilt change. The backend server streams a real-time uncompressed video (1280×720 at 75 FPS, approximately 1.58 Gbps bit-rate) via the APs to clients. Owing to the channel fading, throughput measurement will show small fluctuations even when the link is stationary. The small throughput variations do not affect the video transmission reliability, but will falsely trigger many bad link indications if we threshold the link state by the video bitrate. Thus, we choose the threshold as 95% of the video bitrate and consider video frames cannot be delivered reliably when the network throughput drops below $T = 1.50$ Gbps. We run Pia against 4 schemes:

(i) *noSwitch*: Client connects to a single fixed AP. (ii) *hardProbe*: Client sequentially probes each AP (once per second) by temporarily switching to it and sending a few packets to have the MCS stabilized. If the probed AP's link quality is better, the client will stay with it. (iii) *softProbe*: Client measures network throughput every 0.5s, and probes for the best one only if the throughput drops below T . (iv) *oracle*: Switch AP based on the known throughput in the trace data without counting any overhead.

Link availability. Figure 4.24 (a) plots the *link availability*, i.e., percentage of time that throughput exceeds T . Pia's availability is 97.3%, close to the *oracle* (99.6%) whose




Figure 4.25: Video frame latency of five AP switching schemes.

availability is below 100% due to certain blind spots not covered by any AP. This can be addressed by proper deployment or adding more APs. The availabilities of other schemes are much lower ($< 84.9\%$). Even *hardProbe* and *softProbe* show low availability due to the AP probing overhead.

Hazard times. Figure 4.24 (b) plots hazard times, *i.e.*, the number of occurrences that link throughput drops below T in a 5-minute test. Compared with *noSwitch*, *softProbe*, and *hardProbe*, *Pia* reduces the hazard times by $6.3\times$, $4.5\times$ and $4.2\times$. The extra hazard times over the *oracle* is caused by occasional wrong AP selection. The unavailability, *i.e.*, 1-availability, divided by hazard times gives the average duration of each drop. *Pia can not only quickly recover from wrong predictions, but also prevent most throughput drops by proactively switching the AP.*

Video frame latency. Figure 4.25 plots the measured video frame latency. The mean and 90-percentile latency of *Pia* (11.0 ms and 17.8 ms) are close to the *oracle* (10.6 ms and 16.4 ms), and much lower other schemes (90-percentile about 40 ms). Note that the latency can be reduced substantially under higher PHY bit-rates (802.11ad supports up to 6.7 Gbps, whereas the Qualcomm radio can only reach 2.5 Gbps).

4.6.2.2 Multi-client Spatial Sharing

We conduct trace-based emulation to evaluate *Pia*'s spatial sharing mechanism. We reuse the BSV trace collected in Sec. 4.6.1, calculate the SIRs of clients in concurrent




Figure 4.26: (a) Throughput, and (b) concurrency opportunity, by four spatial sharing schemes.

transmission, and map them to the achievable bitrate following a standard 802.11ad rate table [165]. The impact of reflectors is modeled through the vAPs (Sec. 4.4.2). We determine the AP and beam selection for each client using Pia and 3 baselines: (i) *noShare*: The transmission opportunity is randomly assigned to one client at a time. Each client finds the best AP and beam based on the 802.11ad AP discovery and beam training. (ii) *802.11ad*: Clients run interference sensing, each in a dedicated BI, to determine the concurrent transmission feasibility (Sec. 4.4.1). Afterward, they transmit concurrently in next 5 BIs. (iii) *oracle*: Maximizing the clients' SIRs (Eq. (4.6)).

Throughput gain. Figure 4.26 (a) plots the achievable throughput over the number of clients. All clients achieve similar throughput in *noShare*, since it roughly splits the time equally. *802.11ad* improves throughput by 17.3% compared to *noShare*. Yet, clients suffer from unfairness—client1's throughput is 75.7% higher than others, because client1 is at a position causing asymmetric interference that starves others, a well-known problem in directional networks [106]. Pia performs closely to the oracle, with $1.47\times$ and $2\times$ gain over *noShare* for 2 and 3 client cases, respectively, and offers higher fairness. Note that the throughput stops increasing as client number reaches 4 because spatial reuse saturates.

Concurrent transmission opportunity. Figure 4.26 (b) plots percentage of successful concurrent transmissions for client1 and client2. Although the percentage gradually decreases as more clients join, Pia improves the concurrent transmission opportunity

by 28% and 36.7% for 2-client and 4-client respectively compared to *802.11ad*. The stacked bars represent a breakdown of the source of gain, from selecting beams with weaker signal leakage (bottom) and eliminating interference sensing overhead (top), which together verify the effectiveness of Pia's pose-assisted spatial reuse mechanism.

4.7 Discussion

Whereas the “laser-like” beam of millimeter-wave radios has triggered many doubts about their coverage, we have demonstrated the high potential of Pia to address the concern. Pia employs the 5-DoF pose information on mobile devices, together with a multi-AP network architecture, to achieve seamless coverage under mobility and user's orientation change. Pia can be applied to meet the needs of demanding Gbps applications such as wireless VR. Its design principle of pose-assisted mobility 60 GHz networking can also be extended to general mmWave pico-cell/small-cell networks with dense AP deployment.

CHAPTER 5

Environment Sensing for Optimizing mmWave Network Deployment

5.1 Introduction

The unlicensed millimeter wave (mmWave) band, centered at 60 GHz and spanning 14 GHz spectrum [194], represents the most promising venue to meet the massive surge in mobile data. Recently proposed mmWave network standards, like 802.11ad [82], provision multi-Gbps connectivity for a new wave of applications such as cordless computing and wireless fiber-to-home.

Despite the huge potential, 60 GHz networks face a number of challenges unseen in conventional low-frequency networks: due to ultra-high carrier frequency, the 60 GHz radios are extremely vulnerable to propagation loss and obstacle blockage. To overcome such limitation, 60 GHz radios commonly adopt many-element *phased-array antennas* to form highly directional, steerable beams, which leverage reflections to steer around obstacles. Dependence on directivity and reflection, however, makes the network performance a sensitive function of node placement and environmental characteristics (e.g., geometrical layout and reflectivity of ambient surfaces).

To elucidate the challenge, we set up two laptops with Qualcomm tri-band QCA6500 chip (2.4/5/60GHz) [133], and randomly place them over 16 different locations. For each link, we measured the normalized throughput, *i.e.*, the none-line-of-sight (NLOS)




Figure 5.1: Normalized throughput of same-distance 60 GHz and 2.4 GHz links over different node placements.

throughput when a human obstacle stands in between, divided by the LOS throughput. Figure 5.1 plots the CDF across links. We observe up to $3\times$ throughput gap when the *same-distance* 60 GHz link is placed at different locations, versus $1.4\times$ for the 2.4 GHz link, implying that 60 GHz’s NLOS performance is much more sensitive to environment. We found that the 60 GHz transmitter can more effectively detour blockage, if it is placed near a concrete wall that acts like a mirror. Obviously, reflectors in the environment have a crucial impact on 60 GHz performance.

Of course, one may not always be able to alter the environment. However, we argue that, by judiciously placing 60 GHz access points (APs) within a given environment, we can substantially improve network coverage and robustness to blockage. To this end, one may conduct a blanket site-survey and search for the capacity-maximizing AP location, but the search space becomes formidable because of the numerous beam directions and human blockage patterns. In this chapter, we discuss *E-Mi*, a system that can automatically “sense” (model) the major reflectors in the environment from a 60 GHz radio’s eyes, and predict the performance of arbitrarily located links. The prediction can in turn help optimize AP placement to maximize network capacity and robustness.

The core challenge in E-Mi is: how to sense the environment using mmWave radios which can only measure the received signal strength (RSS) and phase between each other? Conventional environment mapping approaches (*e.g.*, stereo camera and laser

radar [35, 38, 65, 67, 217]) need dedicated hardware and do not capture environment properties specific to mmWave. In contrast, E-Mi leverages the well known sparsity of mmWave channels [157, 165, 199]: from the 60 GHz radios' eyes, there are usually only a few *dominating reflectors* in practical environment. E-Mi samples the environment by fixing the Tx radio, and moving the Rx to a few locations. At each location, the radio channel comprises one LOS path, and many NLOS ones. By measuring the RSS/phase, E-Mi traces back all NLOS propagation paths, uses a geometrical model to locate where the paths hit reflectors, and eventually reverse-engineers the location and reflectivity of dominating reflectors. Such environment information is then fed into a ray tracing engine, which can predict the wireless channel quality of arbitrarily located Tx/Rx.

E-Mi's reflector learning is predicated on the accurate tracing of propagation paths, which itself is an open challenge. Specifically, E-Mi needs to disentangle all the NLOS paths for each Rx location, and estimate each path's angle of arrival (AoA), angle of departure (AoD) and length. This differs from the vast literature of phased-array localization algorithms that only exploit the LOS path [88, 197]. E-Mi solves the problem using a *Multipath Resolution Framework (MRF)*, which resolves different paths' angles/lengths by creating "virtual beams" by post-processing the measured RSS/phase.

We have implemented E-Mi on a 60 GHz testbed. Our experiments demonstrate that E-Mi can accurately resolve NLOS propagation paths, with an average error of 3.5° , 3.5° , and 0.4 m , for AoA, AoD and path length, respectively. By simply sampling 15 receiver locations in an office environment, E-Mi can effectively predict the link quality of other unobserved locations, with median RSS error of 2.8 dB and AoA(AoD) error $4.5^\circ(5.7^\circ)$.

E-Mi can be a convenient toolset to predict site-specific RSS distributions and assist 60 GHz network deployment and configuration. In this section, we apply E-Mi to one case study to answer the following question: How to deploy the 60 GHz APs to maximize the average network capacity and improve resilience to blockage? Our experiments show that an E-Mi-augmented deployment obtains $2.2\times$ to $4.5\times$ median throughput

gain over empirical approaches. E-Mi also makes the 60 GHz network more robust, reducing median throughput loss from around 700 Mbps to 20 Mbps under random human blockage.

To summarize, the main contributions of E-Mi include:

(i) A multipath resolution framework that allows a pair of 60 GHz Tx and Rx to trace back the $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$ of all NLOS paths, simply via RSS/phase measurement.

(ii) An reflector localization scheme that can locate where the reflectors “bend” propagation paths, and then recover the layout/reflectivity of dominant reflectors.

(iv) Applying the sensing information to predicting the channel quality of arbitrarily located Tx and Rx, which in turn helps optimize the AP deployment, achieving multi-fold capacity gain and robustness under human blockage.

5.2 Related Work

Wireless network planning/profiling. Wireless network planning is a classical problem that has been relying on empirical solutions for decades. RF site survey, despite its tedious war-driving procedure, is still widely adopted by enterprise Wireless Local Area Network (WLAN) and cellular network planning tools [49, 132]. Recent work used roaming robots [94] or sparse sampling [141] to access the RSS distribution under a given AP/basestation deployment. But these approaches hardly shed lights on how to plan a new/better deployment.

Active planning can overcome the limitation by using ray-tracing. Earlier study of 60 GHz channel statistical characterization [158] unveiled that 60 GHz signals have predictable spatial structure in an environment. But they require precise mapping of dominant reflectors. Such a map is not always available and is sensitive to environmental change (*e.g.*, placing a new cabinet). E-Mi essentially circumvents this hurdle by allowing mmWave radios to directly construct the environment map.

Radio-based environment sensing. The Simultaneous Localization and Mapping (SLAM) problem has been extensively studied in robotics [60, 139, 182]. Typical SLAM systems need to roam a robot, and map the environment based on dead-reckoning and visual images. Such systems are predicated on two factors: (i) precisely controlled robotic movement and blanket coverage, to generate an extensive point-cloud representation of the survey area. (ii) environment sensors, such as sonar, stereo camera and LIDAR [70], to explicitly locate landmarks or obstacles. The elusive nature of wireless signals prohibits us from meeting the same requirement. It involves nontrivial human efforts to label the reflectivity of each reflector. Besides, the reflectivity may be hardly available if an object contains compound materials. State-of-the-art radio-based SLAM [65, 203] can only achieve localization accuracy of around 5 meters, far from enough to predict the spatial performance of a wireless network.

Recent work [69, 87, 217] adopted mmWave radars to explicitly *scan* the environment. By continuously moving the radar in front of the obstacle's body, they can identify the shape/reflectivity. In contrast, E-Mi leverages the sparsity of 60 GHz signal structure, so as to locate all dominant reflectors with only a few sampling locations.

Localization using antenna arrays. Antenna array has demonstrated tremendous potential in localization, especially because it can identify AoA using signal processing algorithms like MUSIC [149] and ESPRIT [143]. Recent systems [78, 88, 197] renovated such algorithms to localize a client via multi-AP triangulation. In contrast, E-Mi uses mmWave phased-array to handle the more challenging problem of recovering NLOS propagation paths and locating reflectors. A side benefit of E-Mi is that it can locate a node using a single AP (Sec. 5.5) and therefore build a spatial distribution map of possible client locations, which can in turn help optimize the AP deployment (Sec. 5.9).

Sensor-assisted protocol adaptation. E-Mi is partly inspired by the principle of sensor-assisted communications. Nanthapuri *et al.* [146] proposed to discriminate various networking context (*e.g.*, mobile vs. static) using external sensors, and adapt the

protocols accordingly. Ravindranath *et al.* [137] applied a similar principle to assisting link-level rate adaptation, *etc.*. BBS [119] leveraged a WiFi antenna array to estimate the signal's AoA and facilitate the 60 GHz radio beam steering. BeamSpy [166] detects human blockage and adapts its beam to a new direction without beam scanning. In contrast to this line of research, E-Mi uses 60 GHz radios themselves as sensors to reconstruct the reflectors and predict the site-specific RSS distribution to guide AP deployment.

5.3 E-Mi: An Overview

E-Mi samples the RSS/phase between a pair of 60 GHz AP and client (also denoted as Tx and Rx), and uses the samples as input to two major modules: (i) *Multipath resolution framework* (Sec. 5.4), which estimates the geometry, *i.e.*, $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$, of each propagation path and also discriminates their RSS/phase. (ii) *Dominant reflector reconstruction* (Sec. 5.5): which locates the reflecting points (*i.e.*, spots where the paths hit the reflector), and reconstructs the layout/reflectivity of dominating reflectors, forming a coarse environment map. A network planner can feed E-Mi's reconstruction result to a 60 GHz ray-tracing engine, and identify the AP locations that lead to higher capacity/robustness (Sec. 5.9). This essentially supersedes the laborious war-driving in traditional wireless site survey [132].

E-Mi requires the Tx and Rx to be equipped with phased-arrays of practical size (default to 16-element, as in typical 802.11ad radios [13]). It does not need a customized PHY layer—It only requires the channel state, which is a portable function on many commodity WiFi devices [73] and expected to be available in the 802.11ad products. Although E-Mi works in a constrained environment that can be illuminated by Tx's signals, the Tx can be moved to different positions to extend its coverage.

When scanning the wireless channel, E-Mi places the Tx and Rx well above the

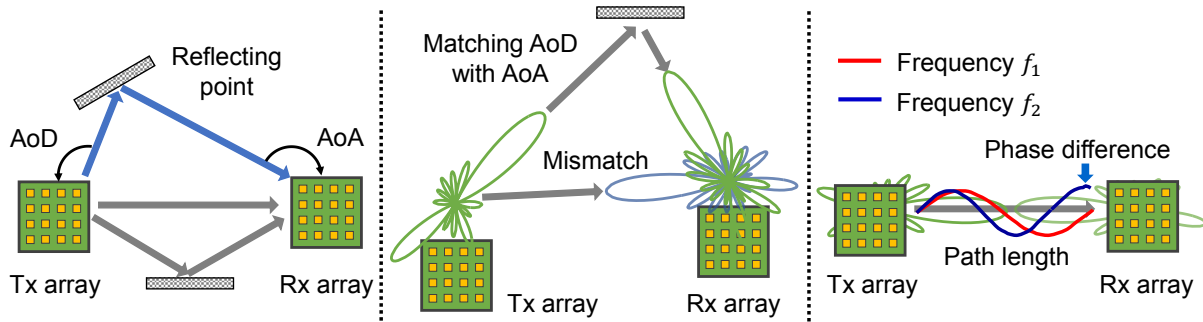


Figure 5.2: The MRF identifies the $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$ in three steps: (1) Identify the dominant AoA and AoD; (2) Pair AoA and AoD directions that belong to the same path; (3) Estimate the length of each path.

ground, so that they only “see” dominant reflectors like walls and furnitures. They affect the average-case network performance which are of utmost interest for network planners. In case such reflectors change their locations, we can accommodate the changes by rerunning E-Mi.

5.4 Multipath Resolution Framework (MRF)

The MRF estimates the $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$ of dominating paths between the Tx and Rx. As illustrated in the Figure 5.2, the AoA and AoD are determined by the relative positions of dominant reflector, Tx and Rx, and independent of the beam pattern of phased arrays. To estimate these intrinsic parameters, a naive solution is to use *beam scanning*: the Tx/Rx may steer over all possible combinations of beam directions, and find the ones with high RSS. However, a 60 GHz phased-array can only steer to a set of discrete directions (*e.g.*, a 16-element one can only steer between beams with 22.5° separation [180]). The discrete beam scanning prevents us from measuring the signal angle precisely. Moreover, unlike horn antennas, phased-arrays have imperfect directionality — besides the main beam, their antenna pattern bears many sidelobes which interfere the AoA/AoD estimation.

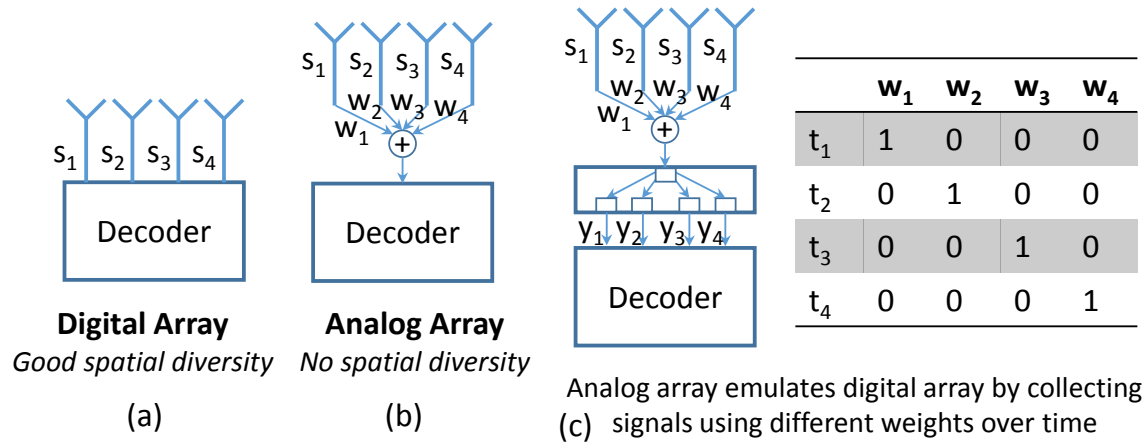


Figure 5.3: Isolating individual antenna's signals on an analog phased array, by switching across different weight vectors.

E-Mi's MRF introduces three mechanisms (Figure 5.2) to meet the challenges. (i) We first estimate the dominating AoD and AoA directions, originating from Tx and ending at Rx, respectively. We adapt a classical signal angle estimation algorithm to 60 GHz phased arrays, which enables super-resolution (*i.e.*, finer resolution than discrete steering by generating a continuous angular spectrum, and unaffected by the imperfect beam shape of phased arrays). (ii) We design a *virtual beamforming* (vBeam) scheme that pairs the AoD and AoA directions belonging to the same NLOS path. (iii) We employ a *multi-tone ranging* scheme to estimate the total length of each path.

5.4.1 Estimate Path Angles Using Phased Arrays

Conventional multi-antenna receivers can estimate signal AoA using *angular spectrum analysis*, which singles out the arrival angles with strong signal strength [88, 197]. However, such analysis needs to *isolate* the signals on each antenna element using digital phased arrays (Figure 5.3(a)). Practical 60 GHz radios use *analog phased arrays* (Figure 5.3(b)), which have a single input/output, comprising a weighted sum of individual antenna's signals that obfuscate each antenna's signals.

To overcome this limitation, a natural way is to vary the weights and obtain a sys-

tem of equations to solve for individual signals. Suppose there are N elements on the receiving phased array, and $\mathbf{S} = [S_1, S_2, \dots, S_N]^T$ denotes the signals on the N individual antenna elements. When the phased array imposes row vector of weight \mathbf{w}_1 on its antenna elements, the received signal becomes $y_1 = \mathbf{w}_1 \mathbf{S}$. Suppose the array switches across N different weight vectors to receive the same signals by N times. The weights constitute a 2-D matrix $\mathbf{W}_R = [\mathbf{w}_1; \mathbf{w}_2; \dots; \mathbf{w}_N]$ with each row being a weight vector. Then, the N output signals form a vector: $\mathbf{Y} = \mathbf{W}_R \mathbf{S}$. So, one can simply use $\mathbf{S} = \mathbf{W}_R^{-1} \mathbf{Y}$ to recover \mathbf{S} , and hence isolate the signals on each antenna element. Figure 5.3(c) illustrates an example where $N = 4$.

In practice, the weight vectors in a 60 GHz phased array are built into hardware and can only be selected from a predefined group, called *codebook*. The key question is: can we find a set of weight vectors to form a matrix \mathbf{W}_R that is invertible? The answer is positive: we can find the weight matrix from 60 GHz codebook that is orthogonal (*i.e.*, $\mathbf{W}_R \mathbf{W}_R^H = \mathbf{I}$, where $(\cdot)^H$ denotes the conjugate transpose), and hence invertible. The beamforming codebook ensures orthogonality between weight vectors because it will maximize the isolation across different beam patterns [136, 180].

To estimate the AoD, a symmetrical operation is needed at the Tx. Suppose the Tx phased array has M antenna elements, then it uses M different sets of weights to transmit the signals by M times, which similarly constitute a transmit matrix \mathbf{W}_T . We populate \mathbf{S} into an N -by- M matrix. Each element (i, j) of \mathbf{S} represents the signals on i -th Rx antenna, when the j -th Tx antenna element is triggered. Then, the received signals of analog-array becomes: $\mathbf{Y} = \mathbf{W}_R \mathbf{S} \mathbf{W}_T^H$. Each column/row in matrix \mathbf{Y} contains received signals measured using a specific transmitting/receiving weight vector. By way of a similar orthogonality argument as the Rx, we can recover \mathbf{S} as follows: $\mathbf{S} = \mathbf{W}_R^H \mathbf{Y} \mathbf{W}_T$.

Isolation of individual antenna's signals allows us to apply MUSIC [75], an eigen angle analysis algorithm to jointly estimate the AoA/AoD, in the same way as in digi-

tal phased arrays [197]. MUSIC can achieve a scalable resolution with more antenna elements and extricate the discrete beam shape of phased arrays. Specifically, we measure the preamble signals sent/received by standard 60 GHz radios [82], which are sent repeatedly across packets, and across different Tx/Rx beam patterns. We isolate the preamble signals sent/received by different Tx/Rx antenna elements. Then, we run MUSIC to compute the angular spectrum, essentially the likelihood of signals coming from different angles. Finally, we find the peaks in the angular spectrum that are larger than the noise floor and take the corresponding directions as AoAs/AoDs of dominating paths.

A few additional operations are worth noting: (i) Since each AoD is pairwise to an AoA *w.r.t.* the same dominant reflector, we remove the excessive AoA/AoD estimations of smaller eigenvalue, and make sure the number of AoA and AoD values are equal. (ii) To ensure the consistency of reference direction, *i.e.*, 0 degree, in the measurement, the antenna's orientation can be simply kept at a fixed direction, or be compensated by motion sensors in the mobile device. (iii) MUSIC is adopted only for AoA and AoD estimation. The RSS estimation and AoA/AoD pairing of each path is accomplished by the virtual beamforming, which will be detailed next.

5.4.2 Virtual Beamforming: Match Path Angles

E-Mi's virtual beamforming (vBeam) algorithm serves two purposes: *First*, the AoAs and AoDs identified above do not have a pairwise mapping. The vBeam can pair up the AoA and AoD values that belong to the same path. *Second*, the received signals \mathbf{S} are a mix from all propagation paths. To estimate the length of each path (Sec. 5.4.3), their signals have to be separated from each other.

The basic idea is to process the received signal matrix \mathbf{S} offline and emulate Tx/Rx beamforming towards specific directions. This allows us to generate arbitrary beam patterns, bypassing the codebook constraint of phased arrays. Then, vBeam uses a

beam matching metric to single out each pair of AoD and AoA directions that belong to the same path.

Beam generation: vBeam generates weight vectors of specific beam patterns and applies them to signals from different antenna elements. Whereas the weight vectors can be computed using conventional delay-sum beamformer [172], vBeam applies a beam-nulling technique instead, which beamforms to the desired AoA/AoD directions while nulling signals from other AoA/AoD directions. *This effectively steers the phased-array's sidelobes toward directions from which there is no signal coming, and thus helps suppress irrelevant signals.*

Suppose the AoA and AoD identified above are denoted by vectors $\Theta = [\theta_1, \theta_2, \dots, \theta_K]$ and $\Phi = [\phi_1, \phi_2, \dots, \phi_K]$, with K being the number of dominant paths. Denote $\mathbf{a}_r(\theta_i)$ and $\mathbf{a}_t(\phi_i)$ as column weight vectors that beamform toward AoA/AoD angle θ_i and ϕ_i . Take the Rx-side as an example, the nulling beam vector $\mathbf{a}_r^{null}(\theta_i)$, which beamforms to θ_j for $j = i$ and nulls other θ_j for $j \neq i$, can be directly derived from $\mathbf{a}_r(\theta_i)$ by [161].

Beam matching: Suppose vBeam beamforms towards AoA angle θ_i and AoD angle ϕ_j using the foregoing approach. In order to determine whether θ_i and ϕ_j belong to the same propagation path, we design a *beam matching metric* \mathbf{F} , which manifests a high value if and only if θ_i and ϕ_j match to the same path. \mathbf{F} is computed by:

$$\mathbf{F}[i, j] = \mathbb{E}[|\mathbf{a}_r^{null}(\theta_i) \mathbf{S} \mathbf{a}_t^{null}(\phi_j)^H|^2], \forall 1 \leq i, j \leq K,$$

where the inner part of above equation applies the virtual beams to signal matrix \mathbf{S} , and outer expectation computes the corresponding RSS. Since it is difficult to find an absolute gauge threshold, E-Mi adopts an iterative algorithm (Algorithm 5.1). It starts with the largest metric and takes the corresponding AoA/AoD as a pair. Then it removes values of pairwise AoA/AoD from the row and column of the matching matrix \mathbf{F} and

```

1: procedure  $vBeam(\Theta, \Phi, \mathbf{S})$ 
2:   for  $i = 1:K, j = 1:K$  do ▷ Loop for Rx and Tx arrays
3:      $F[i, j] = \mathbb{E}[|\mathbf{a}_r^{null}(\hat{\theta}_i)\mathbf{S}\mathbf{a}_t^{null}(\hat{\phi}_j)^H|^2]$  ▷ Beamform RSS
4:   end for
5:   for  $i = 1:K$  do ▷ Beam matching
6:      $[I\_row, I\_col] \leftarrow \max\_subscript(\mathbf{F})$  ▷ Subscript of maximum
7:      $\mathbf{F}[I\_row, :] \leftarrow 0; \mathbf{F}[:, I\_col] \leftarrow 0;$ 
8:      $\hat{\Theta}[i] \leftarrow \Theta[I\_row]; \hat{\Phi}[i] \leftarrow \Phi[I\_col]$ 
9:   end for
10:  return  $\hat{\Theta}$  and  $\hat{\Phi}$  ▷ pairwise AoA and AoD
11: end procedure

```

Algorithm 5.1: Virtual Beamforming.

repeats above procedure to find the next largest matching metric. This approach works well for paths of different signal strengths.

Once vBeam identifies all the pairwise AoA/AoD, it can *isolate path i 's signal* \mathbf{S}_i^{path} by projecting the entire signal matrix \mathbf{S} towards path i 's AoA and AoD:

$$\mathbf{S}_i^{path} = \mathbf{a}_r^{null}(\hat{\Theta}[i])\mathbf{S}\mathbf{a}_t^{null}(\hat{\Phi}[i])^H, \quad (5.1)$$

where $\hat{\Theta}$ and $\hat{\Phi}$ are matrices of the pairwise AoA and AoD. E-Mi then further estimates the RSS of signal isolated from each path.

5.4.3 Multi-Tone Ranging: Estimate Path Length

E-Mi estimates each path i 's length by processing its signals \mathbf{S}_i^{path} , using a *multi-tone ranging* mechanism. Modern communication systems such as 60 GHz 802.11ad commonly adopt Orthogonal Frequency Division Multiplexing (OFDM), which modulates signals across different frequency tones (called *subcarriers*). The phase of each subcarrier can be measured using built-in channel estimators. Suppose a subcarrier has frequency f_1 , then its phase increases linearly with propagation path length d , following $2\pi f_1 d/c$. Our multi-tone ranging leverages the *phase divergence* among OFDM subcarri-

ers, caused by their frequency difference. Given two subcarriers with frequency f_1 and f_2 , their phase divergence at distance d equals $\Delta\varphi = 2\pi(f_2 - f_1)d/c$, where c is light speed. f_1, f_2 are known and $\Delta\varphi$ can be measured. So we can easily map $\Delta\varphi$ back to d .

To improve resilience to channel noise, E-Mi harnesses diversity from many subcarriers in 802.11ad-like communication systems. Suppose we have isolated the preamble signal S_i^{path} along path i (Sec. 5.4.2). Suppose L subcarriers are located at frequencies f_1, f_2, \dots, f_L , and the Rx-measured phase values are $\varphi_1, \varphi_2, \dots, \varphi_L$. Then we estimate path i 's length via the following optimization framework:

$$d^* = \arg \max_d \left| \sum_{i=1}^L e^{j(\varphi_i - \frac{2\pi f_i d}{c})} \right|. \quad (5.2)$$

The RHS of Eq. (5.2) computes the difference between measured phase and theoretical phase over distance in the phaser domain. Essentially, the optimal distance estimation d^* leads to the closest match between these two sets of phase values.

In practice, due to the Carrier Frequency Offset (CFO) between Tx and Rx, the subcarriers will bear unknown phase shifts, which contaminate the phase measurement. We cannot apply the standard CFO compensation technique [167] in this case because it will simultaneously nullify the phase divergence. We address this problem using a *reference calibration* scheme. Specifically, we first separate the Tx and Rx by a known distance d_0 and measure the phase value $\varphi_i(d_0)$ of each subcarrier i . When the Tx/Rx moves to a new (unknown) distance d , the CFO can be canceled by computing their phase difference: $\varphi_i(d) - \varphi_i(d_0) = \frac{2\pi f_i(d-d_0)}{c}$. Substituting φ_i by $\varphi_i - \varphi_i(d_0)$, Eq. (5.2) can be reformulated as:

$$d^* = \arg \max_d \left| \sum_{i=1}^L e^{j(\varphi_i - \varphi_i(d_0) - \frac{2\pi f_i(d-d_0)}{c})} \right|. \quad (5.3)$$

We note that the phase divergence has an aliasing effect: if the phase difference between two subcarriers exceeds 2π , it will wrap and cause ambiguity. To maximize the

unambiguous ranging distance, we should maximize the cycle length of phase divergence, or equivalently minimize the frequency separation between subcarriers (denoted as f_{min}). The unambiguous range is thus determined by $\frac{c}{f_{min}}$. For 802.11ad, f_{min} equals the separation between adjacent subcarriers, *i.e.*, 5.156 MHz [82], or equivalent to up to 58.18 *m* unambiguous ranging distance, which is sufficient for indoor scenarios.

5.5 Dominant Reflector Reconstruction

The $\langle AoA, AoD, length \rangle$ of all propagation paths form a set of spatial constraints, allowing E-Mi to locate the *reflecting points*, *i.e.*, points where dominant reflectors “bend” the propagation paths. Consequently, E-Mi can geometrically reconstruct reflectors’ position, orientation, and reflectivity.

5.5.1 Locating Reflecting Points in Environment

To locate the reflecting points, E-Mi first *pinpoints the Rx* relative to the Tx, based on the $\langle AoA, AoD, length \rangle$ constraints. Figure 5.4 shows an example. Suppose we obtained the $\langle AoA, AoD, length \rangle$ of a single path. Then any point along *line segment AB* satisfies the same $\langle AoA, AoD, length \rangle$ constraint, and is likely to be the Rx position. Therefore, a single path cannot pinpoint the Rx. But we can resolve the ambiguity by adding another path: the intersection between line segment *AB* of one path and segment *A'B'* of another path pinpoints the Rx location.

Practical environment may encounter more than two paths. Denote the path *i*’s $\langle AoA, AoD, length \rangle$ as $\hat{\theta}_i$, $\hat{\phi}_i$ and \hat{d}_i . Following Figure 5.4, we can use a simple geometry to represent the intersection formed by all line segments:

$$\mu_i x + \nu_i y = \gamma_i, \quad (5.4)$$

$$\begin{aligned}
\text{with } \mu_i &= \sin(\hat{\theta}_i) + \sin(\hat{\phi}_i), \quad \nu_i = -(\cos(\hat{\theta}_i) + \cos(\hat{\phi}_i)) \\
\gamma_i &= x_{tx}(\sin(\hat{\theta}_i) + \sin(\hat{\phi}_i)) - y_{tx}(\cos(\hat{\theta}_i) + \cos(\hat{\phi}_i)) \\
&\quad + \hat{d}_i(\sin(\hat{\phi}_i)\cos(\hat{\theta}_i) - \cos(\hat{\phi}_i)\sin(\hat{\theta}_i))
\end{aligned}$$

where (x_{tx}, y_{tx}) and (x_{rx}, y_{rx}) are the Tx and Rx position, respectively. Assuming the Rx position is intersected by N line segments, Eq. (5.4) can be rewritten in a matrix format:

$$\Gamma X = P, \quad (5.5)$$

where $X = [x_{rx}, y_{rx}]^T$, $\Gamma = [\mu_1, \dots, \mu_K; \nu_1, \dots, \nu_K]^T$, and $P = [\gamma_1, \dots, \gamma_K]^T$. In practice, due to residual error of measurement, line segments may not intersect on a single point. We thus reformulate Eq. (5.5) as a least-square optimization problem:

$$X^* = \arg \min_X \|P - \Gamma X\|^2, \quad (5.6)$$

where X^* estimates the Rx position with minimum error.

Unfortunately, our initial experimental tests found the optimization alone works poorly due two practical factors: (i) The $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$ estimation (Sec. 5.4) contains residual errors, especially for long-range and weak-RSS paths. Such errors may cause intersecting line segments to be close to parallel, which significantly amplifies the Rx location error. (ii) The MRF (Sec. 5.4) may capture high-order reflections that do not follow the model in Figure 5.4. Such mismatch may deviate the estimation arbitrarily away from the real position.

E-Mi introduces two mechanisms to overcome above challenges.

(1) Weighting the residual error. We first reformulate the optimization problem in Eq. (5.6) to account for MRF's residual errors. The inner term of Eq. (5.6) calculates the difference between matrices ΓX and P . Since different paths' geometries have different residual errors, we weight the paths according to the *confidence level* in MRF's




Figure 5.4: Using a single Tx to locate the Rx and reflecting points. Solid green line represents one propagation path. Dotted green lines are alternative paths with the same $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$ constraint.

estimation, based on the following observation: those paths of shorter length (thus higher RSS) tend to have less error. Thus, we can use the inverse of path length as the weight. Besides, we need to minimize the sum of distance from Rx position to all line segments, which requires normalization by a coefficient $\sqrt{\mu_i^2 + \nu_i^2}$ [39]. The final weight value for path i is:

$$w_i = \frac{1}{d_i \sqrt{\mu_i^2 + \nu_i^2}},$$

which consists of the inverse of path length and the normalization factor. Stacking the weights into a vector $W = \text{diag}(w_1, \dots, w_K)$, the optimization problem Eq. (5.6) can be rewritten as $X^* = \arg \min_X \|W(P - \Gamma X)\|^2$, which can be solved by standard least-square algorithms.

(2) Filtering higher-order reflection. To constrain the problem within the geometrical model of first-order reflection (Figure 5.4), we should exclude any higher order reflection paths in the optimization. Our key observation is that line segment intersections of higher-order reflections tend to randomly distribute and exhibit a larger deviation since they do not fit into the geometry model for the first-order reflection. Therefore, we apply a K-means clustering algorithm to filter out the $p\%$ most signif-

icantly deviated line segments that most likely belong to the higher-order reflection. The choice of p value depends on the amount of higher-order reflections. We prefer a larger p value in a highly reflective environment, and otherwise a smaller p . Yet, we find E-Mi is not sensitive to it because most indoor environments have comparable number of dominant reflectors, and an empirical value (e.g., $p = 20$) would suffice.

After determining the client position, the *locations of reflecting points* (x_{ref}^i, y_{ref}^i) can be estimated by:

$$\frac{x_{ref}^i - x_r}{y_{ref}^i - y_r} = \tan(\hat{\theta}_i), \quad \frac{x_{ref}^i - x_t}{y_{ref}^i - y_t} = \tan(\hat{\phi}_i), \quad (5.7)$$

which solves a set of equations following simple geometry in Figure 5.4.

5.5.2 Reconstructing Dominant Reflector Layout and Reflectivity

E-Mi reconstructs the dominant reflector geometry (orientation/location/length) and reflectivity, by sampling the wireless channel across a *sparse* set of Rx locations, and locating the corresponding reflecting points following the above steps. It creates a 2-D cross section of the environment corresponding to the horizontal plane of the Tx/Rx. Extension to the 3-D case will be discussed in Sec. 5.10.

Reconstructing dominant reflectors' geometry. Ideally, we can move the Rx to many positions, each helping to locate multiple reflecting points. A sufficient number of reflecting points can form a *pixel cloud* that outlines the reflector geometry. However, due to sparsity of the propagation paths [157, 165, 199], collecting a dense pixel cloud requires hundreds of Rx positions even for a small office.

To avoid such war-driving, we design a sparse reconstruction method which only samples at a few positions. We abstract the reflectors into two categories. *Specular reflectors*: a large continuous surface, e.g., wall and cabinet, upon which mmWave signals experience specular reflection [217]; *Diffusive reflectors*: small-size objects, e.g., pillar




Figure 5.5: Reflecting point (VP) and corresponding virtual source (VS), used to recover the geometry and location of specular/diffusive reflectors.

and computer screen, which scatter signals towards a wide range of angles [185].

(i) *Locating specular reflectors.* We model a specular reflector as a continuum of planar segments, each with different reflectivity or orientation. According to the law of reflection, signals that are reflected by a specular reflector look like emitting from a virtual source (Figure 5.5). Given the $\langle AoA, AoD, length \rangle$ and reflecting point location of one propagation path, we can use simple geometry to pinpoint the Virtual Source (VS) position relative to the real TX position, following the law of reflection.

Ideally, each specular segment should have a single VS even as the Rx moves and creates different reflection paths. Yet due to the MRF's residual errors, the VSs estimated at different Rx positions do not exactly overlap, but fortunately they tend to form a cluster (Figure 5.5). We apply the clustering algorithm [36] to isolate the clusters and use each cluster center to represent one VS. Once a reflector segment's VS is identified, its position and orientation is readily determined via a mirror-partitioning between the real source and VS. Further, we extend the reflector segment, and take the size of geometrical shape from pixel cloud as constraint to determine the length of each reflector segment. An example experiment will be provided in Sec. 5.8.2 to elucidate the procedure.

(ii) *Locating diffusive reflectors.* Observing that a diffusive reflector corresponds to densely concentrated reflecting points (Figure 5.5), we can also apply the clustering

algorithm on the reflecting points and use the centers of resulting clusters to represent the diffusive reflectors. Yet this will be interfered by dispersive reflecting points created by specular reflectors. Fortunately, based on the previous step, we can identify and exclude such interfering points. Specifically, we identify them based on their Euclidean distance to the specular reflectors. Threshold is set to $3\times$ the variance of reflector position error, which can isolate a majority of specular reflecting points.

Estimating reflection loss. We now describe how E-Mi models reflection loss, the major distorting factor when signals hit the reflector. Other factors such as diffraction may also vary the signal strength but the effect is minor [109]. E-Mi separately models the reflection loss of each path it has identified. Three factors contribute to the propagation loss: free-space pathloss, oxygen absorption O_l and reflection loss R_l , *i.e.*,

$$RSS = P_t + G_t + G_r - (20 \log_{10}(d) + O_l + R_l),$$

where d is the path length. P_t, G_t, G_r represent the Tx power and Tx/Rx antenna gain. The O_l almost remains a constant for distance of tens of meters [136]. For each path, d and RSS are known from MRF (Sec. 5.4). Thus, to obtain R_l , we need to obtain the constant parameters P_t, G_t, G_r , which may not be available in practice. In addition, the constant value O_l is unknown either.

We address this issue by using the LOS path as reference calibration to cancel those unknown factors. First, we can isolate the LOS path's signals from NLOS paths' signals by metrics such as shortest path length and strongest RSS. We then estimate the reflection loss of each NLOS reflecting path via a simple subtraction:

$$R_l = RSS_{LOS} - RSS_{ref} - 20 \log_{10}(d_{ref}/d_{LOS}), \quad (5.8)$$

where RSS_{LOS} , RSS_{ref} and d_{LOS} , d_{ref} are RSS and path length for LOS and reflected path. Since each reflecting segment/point may have multiple estimations corresponding to

multiple reflection paths, we take the average as its final reflection loss.

5.6 Parametric Ray-tracing: Predict Link Performance

Ray-tracing is a fine-grained way to model wireless signal propagation in both indoor and outdoor environment [57,86]. It tracks the details of how *each signal path* is attenuated over distance and reshaped by reflectors. Measurement studies demonstrated that, given a precise physical description of reflectors, the signal pattern predicted by ray-tracing is reasonably close to real measurement in both LOS and NLOS scenarios [114].

E-Mi employs a parametric ray-tracing engine, whose input is the aforementioned layout/reflectivity for dominant reflectors constructed directly from the 60 GHz radios' eyes. We develop E-Mi's ray-tracer following the classical approaches in [56,93], which models the signal propagation in a 2-D domain using a geometrical/optical tracing. The ray-tracer captures the attenuation and reflection effects along all paths, and recursively traces a path until it attenuates by more than 30 dB. In addition, the ray-tracer accounts for the angle-dependent antenna gain patterns from phased-arrays. The gain patterns can be obtained from either phased-array simulators or hardware specification. After synthesizing signals from all paths, the ray-tracer outputs the final RSS and converts it to bit-rate following a standard 802.11ad rate table [165].

Here are a few aspects that worth further discussion:

Using E-Mi in commodity phased-arrays: Our evaluation used a virtual array of 16 omni-directional antenna elements to synthesize a phased-array. Commodity phased-arrays may have a limited field-of-view angle, and their beams reside within half-space (180°) [13]. However, as long as the codebook and gain pattern are available (usually specified by device manufacturers), E-Mi's multipath resolution framework is applicable. In addition, we can flip the phased-arrays' orientation to ensure full-space coverage.

From 2-D to 3-D sensing: Our E-Mi design places the Tx/Rx on the same height

and senses a 2-D cross-section. Extending E-Mi to the 3-D case involves some new challenges, *e.g.*, resolving AoA/AoDs along both azimuth and elevation dimensions. However, the main design principles of E-Mi can still apply. Notably, the geometry of dominant reflectors along the vertical dimension (mostly floors and ceilings) is much simpler and easier to estimate. A user can even directly provide the height information to assist E-Mi in estimating the dominant reflectors in 3-D. We leave such exploration to future work.

Sensing complicate-structured environment: E-Mi abstracts the 2-D environment as a composition of lines (for specular reflectors) and spots (for diffusive reflectors). The abstraction is accurate if the radio environment is sparse, *i.e.*, dominated by large reflectors like walls and furnitures (*e.g.*, cabinet, bookshelf, and refrigerator). Environmental sparsity in turn causes channel sparsity, which has been observed by many 60 GHz measurement studies [34, 135, 158, 199]. Nonetheless, E-Mi cannot capture complicated structures. These structure may violate the channel sparsity assumption and exacerbate higher order reflections. E-Mi does not attempt to capture mobile structures either, as clarified in Sec. 5.3.

5.7 Implementation

We prototype E-Mi on a custom-built 60 GHz software-radio testbed, which uses WARP [138] to generate and process baseband waveforms under the control of a host PC. The digital signals are converted to/from analog through a high-speed DAC/ADC, and carrier-modulated by the Pasternack PEM-003 60 GHz RF front-end [123].

The physical-layer metrics (*e.g.*, channel state information) in commercial 60 GHz radios, though internally available to the radio vendors, are still not opened to the public yet. The recently developed phased-arrays in 60 GHz software radios [207] can only support a small number of elements. Therefore, we reproduce the effects of a 60




Figure 5.6: Left: Custom-built 60 GHz software-radio platform. Top-right: Antenna placement of a 16-element phased-array. Bottom-right: Example beam pattern.

GHz circular phased array using the time-lapse synthesis method, which follows the way that a phased array modulates, weights and transmits wireless signals. We mount a 60 GHz omni-directional antenna MI-WAVE 267V (with 360° azimuth and 40° elevation beamwidth) [112] onto Axis360 [2], a programmable motion control system (Figure 5.6). The Axis360 rotates the antenna to discrete positions along a circle, each position corresponding to one antenna element on a real circular array. The array dimension follows empirical recommendations in antenna design [161]. The radius of a 16-element circular array is 6.4 mm , with 2.5 mm (half-wavelength) between adjacent elements. This time-lapse synthesis approach has been adopted and verified by previous works [30, 97].

To synthesize a pair of Tx/Rx phased-arrays, we apply beamforming weights (based on a standard 802.11ad codebook [180]) and then combine the measurement from all elements within the Rx circular array. Since E-Mi runs in static environment, this time-lapse approach can realistically reproduce a real phased-array where all elements are excited concurrently. Besides, each element of a phased array antenna is expected to have a close to omnidirectional coverage in horizon plane [71]. Figure 5.6 depicts

an example Tx phased-array gain pattern generated by this time-lapse approach, and measured using a 3.4° horn receiver. Despite the measurement speed of our platform is slow currently owing to the mechanical antenna rotation, a full-fledged 60 GHz device, that has the electronic phased array antenna, can transmit a wireless packet at tens of microseconds. The overall sensing time of each location will be at millisecond-level.

We implement E-Mi’s major modules (Sec. 5.3) within the software-radio’s host PC. Due to limited bandwidth, our platform cannot send 802.11ad-compatible preambles for channel estimation. Instead, the Tx sends five orthogonal tones from 3 MHz to 15 MHz as baseband signals, modulated by 60.48 GHz carrier frequency. This does not obstruct our validation because the narrow band implementation can be considered as only utilizing a few subcarriers in the 2 GHz wide band. When an 802.11ad-compatible device is available, E-Mi can be easily extended to conduct MRF and dominant reflector reconstruction across orthogonal subcarriers over a wideband.

For evaluation purpose, we also use the beam-scanning method to acquire the ground-truth $\langle AoA, AoD, length \rangle$, similar to Rappaport *et al.* [199]. We use a pair of Tx/Rx radios equipped with directional horn antennas of 3.4° beamwidth [124]. With Axis360, the Tx antenna sweeps the horizontal plane at a step of 3° . Meanwhile, the Rx measures the wireless channel and steers to next step after the Tx completes a full scanning.

5.8 Experimental Validation

5.8.1 Effectiveness of Multipath Resolution

To verify the MRF (Sec. 5.4), we set up a pair of Tx and Rx, each synthesizing a 16-element phased-array. We conduct experiments in a $90 m^2$ office environment, which represents a typical indoor environment. The dominating reflectors involve 2 drywalls, 2 concrete walls and 1 pillar (Figure 5.7). We fix the Tx and randomly move Rx over multiple locations. The result is compared against the ground-truth AoA, AoD (mea-




Figure 5.7: Dominant reflectors in an office environment.




Figure 5.8: Left: Estimated AoA/AoD from MRF v.s. ground truth; Right: CDF of AoA/AoD error from MRF.

sured using the oracle beam scanning, Sec. 5.7) and path length (measured using a laser ranger). The ground-truth measurement reveals each link has 3 to 5 dominating propagation paths.

Success rate of AoA/AoD detection: Recall that MRF needs to *detect* and then *pair* each AoA/AoD that belongs to the same path (Sec. 5.4). Our measurement shows that MRF correctly detects 89% and 82% of AoAs and AoDs, and almost 100% of the correctly detected AoAs/AoDs are correctly paired. MRF fails to detect AoA/AoD of some paths primarily because their reflected signal strength is too weak – We find the RSS of unidentified paths is typically 16 dB lower than the LOS path. In other words, the detection failure is not critical since they will have a limited impact on link performance.

Accuracy in resolving AoA/AoD: Figure 5.8 depicts how accurately E-Mi can resolve the AoA/AoD of all the multipaths that it successfully detects. We observe that *the estimated direction closely matches the true direction*. On average, the AoA/AoD error




Figure 5.9: Left: Estimated path length v.s. ground truth; Right: CDF of path length error from MRF.

is only 3.5° and the 90-th error is 8° . Considering that the beam-switching granularity of a 16-element phased-array is 22.5° [180], the *accuracy of AoA/AoD estimation indeed bypasses the granularity constraint of codebook-based beamforming and achieves super-angular resolution*. We expect the accuracy can be improved by larger phased-arrays, due to more entries in the received signal matrix (Sec. 5.4).

For comparison, we also run the codebook-based beam-scanning method (Sec. 5.4). We found its success rates in detecting AoA, AoD and pairing the AoA/AoD is only 64.4%, 66.7% and 53.3% respectively. And the average estimation error of AoA and AoD are 21.0° and 22.1° , respectively. The fundamental reason lies in the aforementioned sidelobe problem (Sec. 5.4). *This experiment further verifies the necessity and effectiveness of the virtual beamforming method in MRF, which pairs up AoA with AoD while nullifying sidelobes.*

Accuracy of path length estimation: We run E-Mi’s multi-tone ranging mechanism over all detected paths. The scatter plot in Figure 5.9 shows the estimated length v.s. true length. E-Mi achieves an average error of only 0.4 m and 90-th error of 1 m. The LOS paths (typically < 4 m) tend to have smaller estimation error (0.23 m on average) than NLOS reflection paths due to better RSS. E-Mi can achieve this ranging accuracy using even a relative small bandwidth because the vBeam algorithm (Sec. 5.4.2) can isolate the signals from different paths. This accuracy is sufficient for most 60 GHz applications since the prediction metrics (AoA/AoD and signal strength) are not very

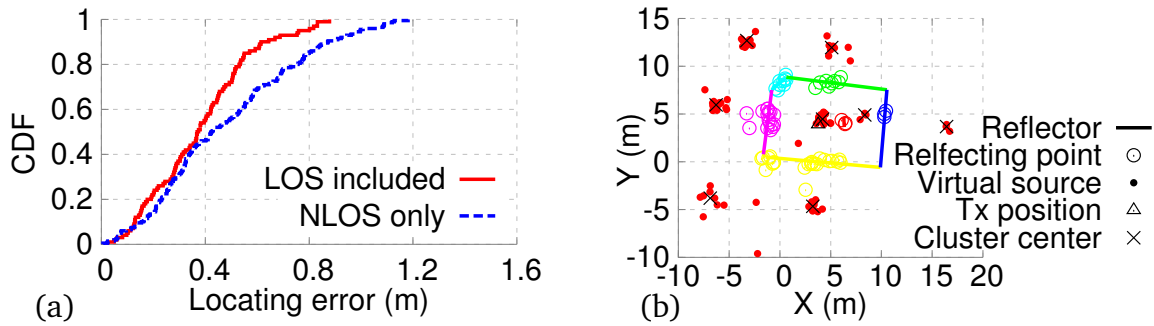


Figure 5.10: Environment reconstruction accuracy: (a) Error CDF in the reflecting point localization algorithm. (b) Breakdown of reconstruction output for an office environment.

sensitive to the range measurement error (by the Friis law, 1 m error only causes less than 2 dB path loss deviation). It is expected that the ranging error will further reduce (Sec. 5.4.3) in practical 802.11ad radios with 1.7 GHz bandwidth. To summarize, this microbenchmark verifies that *the multi-tone ranging precision in E-Mi is sufficient to support dominant reflector reconstruction, even when using our low-end communication hardware.*

5.8.2 Effectiveness of Dominant Reflector Reconstruction

Following the setup in Sec. 5.8.1, a static Tx executes the algorithm in Sec. 5.5.1 to locate the reflecting points for each given Rx position. We move the Rx to 15 uniform positions to reconstruct the dominating reflectors.

Accuracy in localizing reflecting points. Since the reflecting point location has a linear, deterministic relation with the Rx position (Sec. 5.5.1), we mainly focus on evaluating the latter, whose ground truth is obtained via a laser range finder BOSCH DLE40.

Figure 5.10 (a) plots the CDF of localization error. E-Mi can locate the Rx position with mean/90-th error of 0.38 m and 0.6 m, which is even smaller than the path length estimation, because we apply the minimum least square method that leverages the redundant information of multiple reflected paths to reduce the estimation error. Indeed,

Reflector color	Reflection loss	Reflector color	Reflection loss
— drywall	15.87 dB	— concrete wall	8.60 dB
— drywall	13.83 dB	— drywall	20.56 dB
— concrete wall	5.50 dB	— pillar/monitor	1.24 dB

Table 5.1: Estimated reflection loss on different reflectors.

when we intentionally eliminate the LOS paths, the performance (“NLOS only”) drops due to lower path diversity. The results verify that E-Mi’s *reflecting point localization algorithm indeed achieves high precision based on the MRF. More paths provide more diversity and hence higher accuracy.*

Performance of dominant reflector reconstruction. Recall that, given an estimation of the Rx’s and reflecting points’ positions, the dominant reflector reconstruction locates the virtual source, specular reflector, and diffusive reflector, respectively. Figure 5.10 (b) puts together the output from each step, and shows the final reconstruction output based on 15 Rx sampling positions. We observe that the output closely matches the ground truth: 4 walls (specular reflectors) and one pillar/monitor (diffusive reflector). In effect, even the geometrical size of the reflectors matches the ground truth well with less than 0.3 *m* error. Table 5.1 lists the estimated reflection loss. The two concrete walls have 10+ dB lower loss than the two drywalls. And the metal pillar/monitor shows even lower loss. Intuitively, this matches with the reflectivity of materials at 60 GHz [99]. We have also evaluated E-Mi in the cross of a corridor (Figure 5.11) and observed a similar level of accuracy in positioning the dominant reflectors.

In summary, *E-Mi can effectively identify the geometry, locations and distinguish the reflectivity of major reflectors from the 60 GHz radio’s eyes.*

Accuracy of link performance prediction. We now feed the reconstructed reflector geometry/reflectivity into the parametric ray-tracing engine, and predict the spatial channel for another set of 15 randomly located Tx/Rx pairs. Following the ground-truth measurement (Sec. 5.7), we found these links have 66 paths in total. Figure 5.12 showcases example results from two randomly selected links. We observe that the predicted

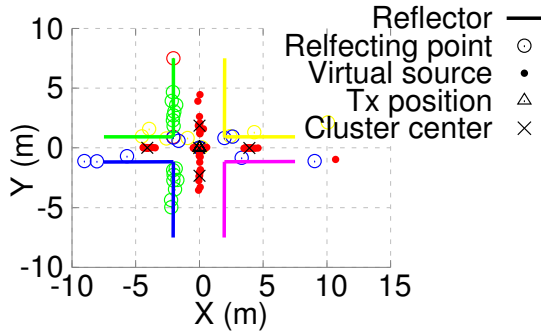


Figure 5.11: Breakdown of reconstruction output for the corridor environment.




Figure 5.12: Predicted v.s. measured AoA/AoD patterns for two locations.





Figure 5.13: Left: Error in predicting the AoA/AoD of unobserved locations. Right: Error in predicting the path length of unobserved locations.

AoA/AoD patterns are highly consistent with the ground truth. The missing spot (e.g., AoA 30° /AoD 40°) is caused by reflection that is not captured by the Tx/Rx during MRF. We found such spots correspond to signal paths with negligible RSS and hence little impact on network performance. Also, adding more Rx position samples can incrementally reduce the probability of prediction loss.

Figure 5.13 plots the channel prediction error over all propagation paths among all links. The average and 90-th errors of path length are 0.64 m and 1.41 m . The average AoA and AoD prediction errors are 4.5° and 5.7° , and 90-th errors are 12.2° and 10.0° , respectively. These results verify that *E-Mi* can accurately predict the AoA/AoD and path length of unobserved locations, based on a number of sparse samples.

Figure 5.14 further shows the predicted v.s. measured RSS among all paths and the corresponding CDF. We observe that among all paths and links, *E-Mi*'s median RSS




Figure 5.14: Left: Predicted RSS v.s. measured RSS; Right: CDF of RSS difference between prediction and measurement.



Figure 5.15: Experiment in a complicated printing room.

prediction error is only 2.8 dB. The scatter plot in Figure 5.14 further shows that locations with higher RSS benefit from higher prediction accuracy, since it mainly involves LOS paths that follow the Friis pathloss model more closely. Even though E-Mi’s prediction is imperfect, we show that it can already become a salient tool for network deployment and protocol optimization (Sec. 5.9).

Scalability in complicated environment. We explore the scalability and generality of E-Mi in two steps. *First*, we collect ground-truth channel profile in a printing room – a more complicated environment (Figure 5.15) than the office. Figure 5.16 (a) and (b) compare the channel profiles (high RSS corresponding to the AoAs/AoDs created by dominant reflectors). Although the printing room hosts a much larger number of objects, the number of dominant reflectors remains similar (~ 6). Our close examination reveals that the dominant directions mainly come from the metal shelf, glass window




Figure 5.16: Spatial channel profile of (a) a printing room with many reflectors and (b) an office environment.




Figure 5.17: Reflector location error vs. number of Rx sampling positions.

and concrete wall, *etc.*, *i.e.*, large objects with strong reflectivity. On the other hand, wooden shelf, desk, and small structures on the wall, barely contribute to the RSS. In both environments, the 6 most dominant reflections account for $> 95\%$ of the total RSS. This experiment demonstrates that *even in a complicated environment, the 60 GHz channel remains sparse, i.e., only a few dominant reflectors determine the channel quality.* Thus, to predict the channel, E-Mi only needs to capture and model a few dominant reflectors.

Second, we investigate how many sampling locations are needed to reconstruct the dominant reflectors. We simulate an $8 \times 10 \text{ m}^2$ room environment with four concrete walls. The Tx is placed at the center and Rx randomly moves over $5 \sim 30$ locations in each test. Across tests, different number of planar reflectors are placed randomly around the Tx/Rx locations. We use the ray tracing to obtain the ground-truth $\langle \text{AoA}, \text{AoD}, \text{length} \rangle$. Figure 5.17 depicts the average error between E-Mi's reconstructed reflector

positions and ground truth. Generally, *more dominant reflectors requires more sampling positions*. Yet, even for an 8-reflector environment, E-Mi only needs 20 sampling locations to ensure an accuracy of around 0.2 m. Since the number of dominating reflectors tends to be sparse, the amount of on-site sampling needed in E-Mi is still substantially lower than that of a site-survey. Besides, the reflector accuracy depends on the radio's sensing position, and it tends to be more accurate if the radio can see stronger reflected paths from a reflector. Fortunately, this location sensitivity will be averaged out when the number of sensing positions is sufficiently large.

5.9 Case Study of E-Mi

In this section, we present an example application which uses E-Mi to predict 60 GHz network performance and optimize the AP placement. E-Mi can also be applied to other scenarios, *e.g.*, device localization and environment mapping. Yet, exploring E-Mi in a broader range of areas is beyond the scope of this work.

5.9.1 Environment-Aware 60 GHz AP Deployment

Measurements (Sec. 5.1 and [165]) showed that *the performance of 60 GHz networks is a sensitive function of location and reflector position, specifically w.r.t.*

(i) *Coverage*: The spatial RSS distribution of a 60 GHz AP tends to be unevenly distributed, even among same-distance links, due to two unique factors: (a) *High directionality*: The AP's phased-array antenna can only generate a discrete set of directional beam patterns, which typically point to unevenly distributed spatial angles [136, 165]. (b) *Ambient reflections*: Different reflectors cause RSS to distribute unevenly, *e.g.*, receivers close to strong reflectors tend to benefit from high RSS [165].

(ii) *Robustness*: *i.e.*, resilience of the network under blockage. 60 GHz links tend to be frequently disrupted due to inability to diffract/penetrate human body. Beam-

steering alleviates the problem, but whether the resulting reflection path can detour blockage highly depends on the geometry/reflectivity of environment [51, 165].

To address the environment sensitivity, we propose to use a multi-AP architecture to cover a constrained environment. Through a central controller, the APs can tightly cooperate with each other. When experiencing poor throughput or blockage, a client can immediately switch to an alternative AP. The architecture itself is not new, but in 60 GHz environment, it needs to meet one key challenge: *For a set of APs under a given environment, how to deploy them optimally so as to maximize the coverage and robustness to blockage?*

We employ E-Mi to answer this question. The basic idea is to predict the AP locations that provide best coverage and robustness to a typical set of client spots. The *client spots* are locations at which clients tend to concentrate. The client spots can be manually specified by users within E-Mi's environment map (e.g., Figure 5.10 (b)). Alternatively, the E-Mi AP can divide the environment map into grids, implicitly sense clients' locations over time (Sec. 5.5.2), and incrementally build clients' spatial distribution. Client spots can be defined as grids where the dwelling probability exceeds a threshold (e.g., 0.05).

To maximize coverage, we define the performance metric \mathcal{D} as the *mean bitrate* among W given client spots. Suppose each client is associated to the AP with highest bitrate. Then, for each combination of AP locations,

$$\mathcal{D} = \frac{1}{W} \sum_{i=1}^W \arg \max_{\{j=1, \dots, A\}} T(RSS)_{ij}, \quad (5.9)$$

where RSS_{ij} represents the RSS from AP j to client i . A is the number of APs, and $T(\cdot)$ maps RSS to the achievable bitrate using the 802.11ad rate table [166]. The candidate AP position can either fall in the grids or be manually specified by users within E-Mi's environment map. To find the best multi-AP deployment, we simply use E-Mi to predict the performance of all candidate position combinations and select the one that

maximizes \mathcal{D} .

For a given multi-AP deployment, we further define the *robustness metric* as $\mathcal{E} = \mathcal{D}'/\mathcal{D}$, where \mathcal{D}' is the mean bitrate under blockage. Using E-Mi, we can “rehearse” the impact of human body blockage without field war-driving. Since human body is aquaphobic [32, 136, 165], the blockage may annihilate one or more paths. So we can use ray-tracing to derive \mathcal{D}' by averaging the bitrate resulting from blockage of random movement. We repeat the procedure over candidate AP position combinations and single out the best. The optimized AP location essentially maximizes the number of paths to each client. The more reachable paths will make the network more resilient to blockage because radios can immediately reestablish the link upon blocking via another path.

5.9.2 Experimental Verification

We evaluate the E-Mi-based AP deployment with dominant reflectors in Figure 5.10 (b). We set 24 random client spots and divide the environment into 20 equal sized grids. The center of each grid is considered as a candidate AP position. We compare E-Mi with *random* deployments and an empirical approach that puts APs in four *corners* to maximize coverage. Since our experiments have already validated the accuracy of the ray tracing method, we reuse ray tracing to evaluate the RSS from an AP to client. The RSS is then mapped to bit-rate and link throughput following [165]. Consistent with Sec. 5.8.1, each AP/client has a 16-element phased array, with 32 codebook entries and beamwidth of 22.5° .

Coverage: Figure 5.18 plots the CDF of clients’ throughput under 2-AP deployment. For *Random*, the throughput is averaged across all AP locations. We observe that E-Mi gains substantial advantage from its environment-awareness, with median throughput improvement of $2.24\times$ over *Random*, and $4.54\times$ over *Corner*. Moreover, E-Mi consistently delivers higher capacity for all clients, *i.e.*, it does not sacrifice fairness. The




Figure 5.18: Client throughput CDF under a 2-AP architecture.

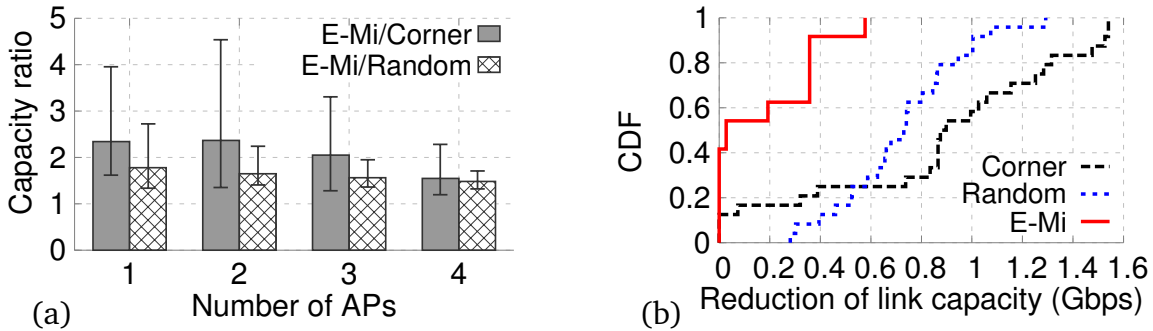


Figure 5.19: (a) Coverage improvement over alternative strategies. (b) Throughput loss under random blockage.

results manifest the ineffectiveness of empirical approaches, which are unaware of the impact of dominant reflectors on 60 GHz network performance. Even with 2 APs, E-Mi can optimize AP placement and boost network capacity to $1.4\times$ compared with a single-AP deployment.

We further evaluate E-Mi in 30 different environment topologies, created by intentionally adding reflectors (up to 10), with random orientations, inside the environment. Figure 5.19 (a) shows that, given a single AP, E-Mi has $2.2\times$ average capacity gain over the empirical deployment, and up to $4\times$ gain in certain environment that is observed to feature heterogeneous reflector placement. More APs can offset the environment heterogeneous, and hence degrade E-Mi's gain slightly.

Robustness: Under the same topology as Figure 5.18, Figure 5.19 (b) plots the CDF of throughput loss across all clients under random human blockages (created by randomly moving at different locations inside the environment). Owing to its awareness

of reflectors, E-Mi's median throughput loss is only around 20 Mbps, in comparison to 700 Mbps to 830 Mbps in the empirical approaches. The normalized throughput gain of our optimization is consistent with the measurement (Figure 5.1) using commercial 60 GHz devices.

We also found that the coverage-maximizing deployment differs from the robustness-maximizing one. In practice, one may use a weighted balance between the metrics, based on how likely the blockage is to happen.

5.10 Discussion

We present E-Mi as a sensing-assisted paradigm to facilitate 60 GHz networks, whose performance is highly sensitive to reflectors. E-Mi senses the environment from 60 GHz radios' eyes. It "reverse-engineers" the geometry/reflectivity of dominant reflectors, by tracing back the LOS/NLOS paths between a pair of 60 GHz nodes. Through case studies and testbed experiments, we have demonstrated how such environment information can be harnessed to predict 60 GHz network performance, which can in turn augment a broad range of network planning and protocol reconfigurations.

CHAPTER 6

Conclusion

In conclusion, this dissertation proposes to augment the wireless systems with sensing information, with the end goal of enabling new IoT applications and solving critical networking performance issues. Conventional wireless systems are primarily standalone, executing prescribed network protocols in the cyberspace — without exploiting or sharing any intelligent sensing information with other system parts. Though such design philosophy works well in the past, it becomes challenging to meet the new IoT demands and fails to scale with the complexity of mmWave networks, *e.g.*, 60 GHz radios suffering from significant performance loss under user mobility and human blockage.

This dissertation is inspired by such an observation — wireless channels exhibit a strong dependence on the physical environment. As such, many emerging IoT networking paradigms can reveal sensitive information about the physical environment. Thus, we take an unconventional approach and show that many open-challenging problems in wireless systems can be readily solved by exploring sensing capabilities from wireless channels and signals. Following this theme, we have proposed two design paradigms — sensing-oriented intelligent systems and sensing-assisted networking protocols. The former enables passive sensing of the physical space through a new modality, which is complementary to active sensors and will trigger many exciting IoT applications, *e.g.*, batteryless orientation and location tracking. The latter paradigm, reversing the above principle, demonstrates that external sensors can significantly improve the networking

reliability for applications such as wireless VR. Besides, we also show that wireless system can benefit from passive sensing information, which allows them to see the nearby environment from the radios' eyes and helps optimize the network deployment. The environment awareness offers the systems superior performance in comparison with conventional wireless protocols which reside only in the cyber space.

6.1 Lessons Learned

This thesis elaborates on contributions that span both sensing and networking areas. Traditional research in radio sensing often focuses on designing new hardware that realizes the radar functionality, which however limits the potential application scenarios. On the other hand, many wireless protocols currently running treat the physical environment as a black box, which tends to be unreliable when the channel becomes highly dynamic. In contrast, we bridge the gap by jointly considering the sensing and networking in the wireless system design. We demonstrate that general communication devices can be re-purposed and achieve the radar-like sensing capability. Moreover, it becomes possible to crack many open challenges in the wireless networking area by exploiting external sensing information for assistance. Below are some of the lessons we have learned.

- **Re-purposed functionalities of Communication Hardware:** The first challenge in designing sensing-augmented wireless systems stems from the non-ideal hardware of communication radios. Traditional radars allow us to process the received signal in the analog domain. However, wireless networking systems tend to digitalize the data symbols, since they are built only for communication purpose. As a result, they often expose limited physical-layer information, such as RSS and phase, for public access. Critical information regarding the physical environment will get obfuscated. Therefore, novel processing algorithms are needed

to reconstruct the meaningful environment information. Besides, compared with dedicated radar hardware, wireless communication circuits are far from ideal, with limitations such as narrow bandwidth and uncalibrated antenna response, which make the signal processing difficult. However, we can compensate for such hardware difference by designing additional training procedure to calibrate the systems.

Performing the preliminary experimental study is a helpful way to identify such problems in an early stage. For instance, before Tagyro, RFID localization research adopts an ideal phase-distance model, which assumes the phase changes linearly over the tag-to-reader distance. Despite the simplicity, my experimental study shows that the model hardly characterizes the phase change in actual RFID systems, since the phase is also a sensitive function of the tag's orientation and the electric-magnetic coupling of antennas. Thus, the ideal model only holds under stringent setups. In contrast, Tagyro adopts a layout-calibration procedure to address the phase distortion caused by the coupling effects and can estimate the rotation of RFID tags. In this thesis, we demonstrated how experimental study could uncover the possible imperfection of communication hardware and bring it into more practical applications through innovative designs.

- **Computational Modeling of Physical Environment:** Another challenge in designing sensing-augmented wireless systems comes from the intrinsic relationship between the environment (target), abstract sensing data, and tremendous decision-space. The channel data often contains incomplete information about the target environment and thus does not directly lead to the decision space. Moreover, the signals are covered by noise and mixed with many other disturbing components from unrelated targets. Making sense of the data requires a deep understanding of the interaction between the sensing signals and the physical environment, distilling internal structural patterns, and designing analytical algo-

rithms.

Building signal-level computational simulation model is a quick way to explore the underlying interaction between the sensing information and physical environment. As demonstrated in this thesis, the computational approach provides us the capability to synthesize wireless signal distortions created by physical environments, human activity, and daily objects, allowing us to verify the algorithm design even before implement the whole systems. For example, prior to our mTrack design, passive mmWave localization mainly leverages the narrow beam and reflection signal strength to determine the object's direction. Nevertheless, this approach can only achieve low precision since the RSS is not a reliable function of the distance. In contrast, mTrack proposed to leverage the phase information for fine-grained passive tracking, but it suffers from the background reflection which distorts the signals. To overcome this problem, we leveraged the computational approach to simulate how the signal deviates under background reflection, which enabled us to identify the inherent property and reconstruct the original pattern without distortion. E-Mi represents another similar example, where we leverage a 60 GHz ray-tracing engine to model how the dominant reflectors in a given environment reshape the signal at the receiver side. It allows us to design the algorithm with better accuracy than previous approaches.

6.2 Future Work

This dissertation has demonstrated that the concept of sensing-augmented wireless systems can help create the more reliable and capable cyber-system as well as new sensing applications/services, laying down the foundation for the advent of IoT. However, we believe the design principle can have broader implications and will have long-lasting impacts. An important future direction that builds upon this thesis is to continue fo-

cusing on developing new computational algorithms and networking architecture for smart wireless-connected devices, to meet the needs of demanding applications such as wireless VR and autonomous driving. This design philosophy opens up many new research opportunities:

Multi-band Content Delivery Framework: Over the past decades, the wireless technology has been mainly focusing on delivering high bit-rate to meet the tremendous growth of mobile devices and video streaming services. The latest mmWave and massive Multiple-input and Multiple-output (MIMO) technologies can already provide multi-Gbps to a single user. However, besides bit-rate, many emerging applications have strict requirements in low-latency and high-reliability. Although mmWave links can meet the throughput requirements, they are vulnerable to user mobility and obstacle blockage. Solving these new challenges requires us to not only redesign new networking protocols/algorithms but also reconsider existing network architecture. In this thesis, the Pia design marks one first attempt in improving the 60 GHz link robustness for wireless VR applications. However, it only focuses on a small mmWave network. An immediate next step is to extend the design to enterprise scale with dense node deployment. The goal is to use the WiFi interface as the sensing-assistive channel and enable robust mmWave spatial sharing under user mobility. For example, one potential direction to design a multipath content delivery framework that can unify multiple radios that are currently available in mobile devices, such as WiFi and WiGig, to create a single low-latency, mobility-robust and energy-efficient virtualized link for demanding applications.

Scalable mmWave Sensing Systems: Existing sub-6 GHz wireless sensing applications are commonly limited to a set of well-predefined gestures or activity patterns. However, mmWave radios can change this situation and enable powerful wireless sensing tools due to the directional beam, ultra-short wavelength, and tremendous bandwidth. It can provide super-fine sensing resolution in both spatial and temporal do-

mains. The mTrack system design demonstrated that mmWave signals could be easily exploited to track human finger movement at sub-centimeter granularity. As future work, this can be extended to a larger scale and realize an accurate indoor positioning system using ubiquitous mmWave radios. Moreover, in the 5G, the millimeter-wave cellular base-stations can be turned into a fine-grained location tracking system to coordinate the scheduling of self-driving cars, *e.g.*, determine who should leave first at a four-way stop intersection. It will be especially useful in an urban area because the traffic is dense and the GPS system deteriorates.

Sensing-assisted Codebook Generation: In E-Mi, we demonstrate that smart wireless sensing can benefit mmWave networks performance by “seeing” the surrounding environment. However, it doesn’t account for the reflection and coupling effects from nearby objects, which will also distort mmWave beam patterns, making the factory-default codebook less effective. To overcome this problem, we need to design an automated codebook generation mechanism that can learn the environmental and electrical impacts, and automatically generate/optimize beamforming codebooks that will provide better signal quality/coverage. Enabling such intelligent wireless sensing systems requires synergistic solutions between signal processing and machine learning.

Towards Seamlessly-connected mmWave Networks: Besides high throughput, wireless network connectivity is a focusing area of research in both academia and industry, which aims to address the coverage and availability problems. It is an important research direction because, in the US, 1.4 million people living in rural areas still have no LTE access. Moreover, existing wireless network infrastructures are fragile in disasters, *e.g.*, it took more than two weeks for primary cellular service to recover after Hurricane Irma. However, solving this problem is fundamentally challenging due to the centralized legacy cellular-network architecture. We believe the mmWave networks can change this situation owing to two unique advantages: directional transmission and vast spectrum. As a result, it can enable efficient peer-to-peer communication via spa-

tial reuse. Besides, through directional wireless backhaul and relays, mmWave radios can instantly connect millions of devices throughout the space yet still maintaining high throughput. To expand the network coverage, one solely deploys a node in the desired area and doesn't need to worry about wired backhaul deployment or complex network configurations. The node would automatically discover nearby peers and together form a high-speed network. Enabling such smart and self-dominated mmWave nodes not only requires exchanging external sensing information, such as geo-location and mobility, but will also necessitate new designs of hardware, protocols, and algorithms.

Protecting User Privacy Under Ubiquitous Sensing: Even though the sensing-oriented wireless systems can enable many new IoT applications, it will expose significant threats to privacy and security. For example, mobile devices, actively transmitting radio signals, *e.g.*, from WiFi or Bluetooth, can be exploited to recover sensitive information such as movement trajectory and sound. Thus, it is crucial to detect and prevent the malicious users, and only permit authorized ones to access meaningful sensing information. To realize these objectives, a secured wireless physical-layer that can "encrypt" the Channel State Information (CSI) will become necessary. A feasible solution, for instance, is to let the transmitter create artificial distortion patterns in the amplitude, phase, and Doppler shift in each packet. Such distortion will not affect the receiver's packet decoding. However, it will prevent the malicious attacker from extracting any useful sensing information from the CSI. An authorized user can compensate the distortion patterns and obtain valid sensing results by using encryption information. Moreover, this solution can also be extended to protect against passive wireless sensing, in which the target object (*e.g.*, human body) does not transmit signals actively. Nearby mobile devices, *e.g.*, smartphone, can assist to create such distortion patterns as jamming signals to confuse the attacker.

BIBLIOGRAPHY

- [1] ALN-9740 Squiggle Inlay. <http://www.alientechnology.com/wp-content/uploads/Alien-Technology-Higgs-4-ALN-9740-Squiggle.pdf>.
- [2] Axis360 Motion Control System. <http://cinetics.com/two-axis360/>.
- [3] Dogbone - Smartrac. https://www.smartrac-group.com/files/content/Products_Services/PDF/DogBone_M4.pdf.
- [4] RFID Tag Type Options. http://www.idtechex.com/research/articles/rfid_tag_type_options_00000041.asp.
- [5] RoboRealm. Microsoft Kinect, 2013. [http://www.roborealm.com/help/Microsoft Kinect.php](http://www.roborealm.com/help/Microsoft%20Kinect.php).
- [6] ShortDipole - Smartrac. https://www.smartrac-group.com/files/content/Products_Services/PDF/ShortDipole_M5.pdf.
- [7] TDK IS-005A RF Absorber. <http://www.tdk.com/>.
- [8] VICON. <http://www.vicon.com/>.
- [9] IEEE Standards 802.15.3c, 2009.
- [10] Mini RFID Panel Antenna S9025P. <http://www.es.co.th/Schematic/PDF/S9025PR-S9025PL.PDF>, 2009.
- [11] IEEE 802.11ad PHY Specifications: Enhancements for Very High Throughput in the 60 GHz Band.
- [12] Low Level User Data Support. *Speedway Revolution Reader Application Note* (2013).
- [13] Wilocity 802.11ad Multi-Gigabit Wireless Chipset. <http://wilocity.com>, 2013.
- [14] Virtual STA and AP Interfaces with ath5k, ath9k and ath10k. <http://www.candelatech.com/vsta.php>, 2014.
- [15] Vubiq 60GHz System. <http://vubiq.com/v60wgd03/>, 2014.
- [16] Google Wants to Use Tango Tech for VR, But Admits Current Dev Kits Aren't Optimized. <http://www.roadtovr.com/google-wants-to-use-tango-tech-for-vr-but-admits-current-dev-kits-arent-optimized/>, 2015.
- [17] mac80211 Multiple Virtual Interface (vif) Support. <https://wireless.wiki.kernel.org/en/users/documentation/iw/vif>, 2015.

- [18] MyScript technology. <http://myscript.com/technology/>, 2015.
- [19] Patch Antenna (Circular) A0003-02. <http://www.atlasrfidstore.com/invengo-xc-1003-uhf-rfid-mobile-phone>, 2015.
- [20] IEEE 802.11ay Study Group. http://www.ieee802.org/11/Reports/tgay_update.htm, 2016.
- [21] Impinj RFID Reader. <http://www.impinj.com/products/readers/>, 2016.
- [22] Miracast. <http://www.wi-fi.org/wi-fi-certified-miracast/>, 2016.
- [23] Internet of Things forecast – Ericsson. <https://www.ericsson.com/en/mobility-report/internet-of-things-forecast>, 2018.
- [24] A. MALTSEV AND R. MASLENNIKOV AND A. SEVASTYANOV AND A. LOMAYEV AND A. KHORYAEV. Statistical Channel Model for 60 GHz WLAN Systems in Conference Room Environment. In *Proc. of European Conference on Antennas and Propagation* (2010).
- [25] ABARI, O., BHARADIA, D., DUFFIELD, A., AND KATABI, D. Cutting the Cord in Virtual Reality. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)* (2016).
- [26] ABARI, O., HASSANIEH, H., RODRIGUEZ, M., AND KATABI, D. Millimeter wave communications: From point-to-point links to agile network connections. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks* (2016).
- [27] ABBAS, K. A New Recurrent Approach for Phase Unwrapping. *International Journal of Applied Science and Engineering* (2005).
- [28] ADIB, F., KABELAC, Z., KATABI, D., AND MILLER, R. C. 3D Tracking via Body Radio Reflections. In *Proc. of USENIX NSDI* (2014).
- [29] ADIB, F., AND KATABI, D. See Through Walls with WiFi! In *Proc. of ACM SIGCOMM* (2013).
- [30] ADIB, F., KUMAR, S., ARYAN, O., GOLLAKOTA, S., AND KATABI, D. Interference Alignment by Motion. In *ACM MobiCom* (2013).
- [31] AGRAWAL, S., CONSTANDACHE, I., GAONKAR, S., ROY CHOUDHURY, R., CAVES, K., AND DERUYTER, F. Using Mobile Phones to Write in Air. In *Proc. of ACM MobiSys* (2011).
- [32] ALEKSEEV, S., RADZIEVSKY, A., LOGANI, M., AND ZISKIN, M. Millimeter Wave Dosimetry of Human Skin. In *Bioelectromagnetics* (2008).
- [33] AMIRI SANI, A., ZHONG, L., AND SABHARWAL, A. Directional Antenna Diversity for Mobile Devices: Characterizations and Solutions. In *Proceedings of the International Conference on Mobile Computing and Networking (MobiCom)* (2010).
- [34] ANDERSON, C., AND RAPPAPORT, T. In-Building Wideband Partition Loss Measurements at 2.5 and 60 GHz. *IEEE Transactions on Wireless Communications* (2004).
- [35] APPLEBY, R., AND ANDERTON, R. Millimeter-Wave and Submillimeter-Wave Imaging for Security and Surveillance. *Proceedings of the IEEE* (2007).

- [36] ARTHUR, D., AND VASSILVITSKII, S. K-means++: The Advantages of Careful Seeding. In *ACM-SIAM Symposium on Discrete Algorithms* (2007).
- [37] ATHANASIOU, G., WEERADDANA, P. C., FISCHIONE, C., AND TASSIULAS, L. Optimizing Client Association for Load Balancing and Fairness in Millimeter-Wave Wireless Networks. *IEEE/ACM Transactions on Networking* (2015).
- [38] AULINAS, J., PETILLOT, Y., SALVI, J., AND LLADÓ, X. The SLAM Problem: A Survey. In *Proc. of the International Conference on Artificial Intelligence Research and Development* (2008).
- [39] BALLANTINE, J. P., AND JERBERT, A. R. Distance from a Line, or Plane, to a Point. *Mathematical Association of America* (1952).
- [40] BAR-SHALOM, Y., LI, X. R., AND KIRUBARAJAN, T. *Estimation With Applications to Tracking and Navigation: Theory Algorithms and Software*. John Wiley & Sons, 2004.
- [41] BAZAN, O., AND JASEEMUDDIN, M. A Survey On MAC Protocols for Wireless Adhoc Networks with Beamforming Antennas. *IEEE Communications Surveys Tutorials* 14, 2 (2012).
- [42] BOLOTNYI, L., KRIZE, S., AND ROBINS, G. The Practicality of Multi-Tag RFID Systems. In *IWRT* (2007).
- [43] BORG, I., AND GROENEN, P. J. *Modern Multidimensional Scaling: Theory and Applications*. Springer Science & Business Media, 2005.
- [44] BUETTNER, M., AND WETHERALL, D. A Gen 2 RFID Monitor Based on the USRP. *ACM SIGCOMM Computer Communication Review* (2010).
- [45] BÜTHE, L., HARDEGGER, M., BRÜLISAUER, P., AND TRÖSTER, G. RFID-Die: Battery-Free Orientation Sensing Using an Array of Passive Tilt Switches. In *Proceedings of ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (2014).
- [46] CANDÈS, E., AND ROMBERG, J. L1-Magic: Recovery of Sparse Signals via Convex Programming. URL: www.acm.caltech.edu/l1magic/downloads/l1magic.pdf (2005).
- [47] CHANDRA, R., AND BAHL, P. MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card. In *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies* (2004), IEEE.
- [48] CHANG, C.-C., AND AGHAJAN, H. Collaborative Face Orientation Detection in Wireless Image Sensor Networks. In *Proceedings of ACM SenSys Workshop on Distributed Smart Cameras* (2006).
- [49] CISCO INC. Site Survey Guidelines for WLAN Deployment, 2013.
- [50] COLLET, A., BERENSON, D., SRINIVASA, S. S., AND FERGUSON, D. Object Recognition and Full Pose Registration from a Single Image for Robotic Manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)* (2009).

- [51] COLLONGE, S., ZAHARIA, G., AND ZEIN, G. Influence of the Human Activity on Wide-Band Characteristics of the 60 GHz Indoor Radio Channel. *IEEE Trans. on Wireless Comm.* (2004).
- [52] COSTI, M., AND DABROWSKI, R. Mobile motion capture.
- [53] DAILYWIRELESS. 60 GHz Backhaul for Small Cells, 2013.
- [54] DANG, B. L., LARRODE, M. G., PRASAD, R. V., NIEMEGERES, I., AND KOONEN, A. Radio-Over-Fiber Based Architecture for Seamless Wireless Indoor Communication in the 60GHz Band. *Computer communications* (2007).
- [55] DANG, M. S., PRAKASH, A., ANVEKAR, D. K., KAPOOR, D., AND SHOREY, R. Fuzzy Logic Based Handoff in Wireless Networks. In *Vehicular Technology Conference Proceedings, 2000. VTC 2000-Spring Tokyo. 2000 IEEE 51st* (2000), vol. 3, IEEE.
- [56] DE GROOT, E., BOSE, T., COOPER, C., AND KRUSE, M. Remote Transmitter Tracking with Raytraced Fingerprint Database. In *IEEE MILCOM* (2014).
- [57] DEGLI-ESPOSTI, AND ETC. Ray-Tracing-Based mm-Wave Beamforming Assessment. *IEEE Access* (2014).
- [58] DEHOS, C., GONZÁLEZ, J. L., DE DOMENICO, A., KTÉNAS, D., AND DUSSOPT, L. Millimeter-Wave Access and Backhauling: The Solution to the Exponential Data Traffic Increase in 5G Mobile Communications Systems? *IEEE Communications Magazine* (2014).
- [59] DOC-OK.ORG. Lighthouse Tracking Examined. <http://doc-ok.org/?p=1478>, 2016.
- [60] DURRANT-WHYTE, H., AND BAILEY, T. Simultaneous Localization and Mapping: Part I. *IEEE Robotics Automation Magazine* (2006).
- [61] EBERHARD GÜLCH. Investigations on Google Tango Development Kit for Personal Indoor Mapping. https://agile-online.org/conference_paper/cds/agile_2016/posters/102_Paper_in_PDF.pdf, 2016.
- [62] EDELSTEIN, A., AND RABBAT, M. Background Subtraction for Online Calibration of Baseline RSS in RF Sensing Networks. *CoRR abs/1207.1137* (2012).
- [63] EVENSEN, K., PETLUND, A., RIISER, H., VIGMOSTAD, P., KASPAR, D., GRIWODZ, C., AND HALVORSEN, P. Mobile Video Streaming Using Location-based Network Prediction and Transparent Handover. In *International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)* (2011).
- [64] FANG, H. *60 GHz RSS Localization with Omni-directional and Horn Antennas*. PhD thesis, 2010.
- [65] FERRIS, B., FOX, D., AND LAWRENCE, N. WiFi-SLAM Using Gaussian Process Latent Variable Models. In *Proc. of IJCAI* (2007).
- [66] FORTINO, G., AND TRUNFIO, P. Internet of Things Based on Smart Objects. *Fortino & P. Trunfio, eds., Cham: Springer International Publishing* (2014).

- [67] GOOGLE. Tango. <https://get.google.com/tango/>, 2014.
- [68] GRANT, M., AND BOYD, S. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [69] GUIDI, F., GUERRA, A., AND DARDARI, D. Millimeter-wave Massive Arrays for Indoor SLAM. In *IEEE International Conference on Communications Workshops (ICC)* (2014).
- [70] GUIZZO, E. How Google's Self-Driving Car Works. *IEEE Spectrum*, 2011.
- [71] GULBRANDSEN, F. Design and Analysis of an X-band Phased Array Patch Antenna. Master's thesis, Norwegian University of Science and Technology, 2013.
- [72] GUPTA, G., SINGH, B. P., BAL, A., KEDIA, D., AND HARISH, A. Orientation detection using passive uhf rfid technology [education column]. *IEEE Antennas and Propagation Magazine* (2014).
- [73] HALPERIN, D., HU, W., SHETH, A., AND WETHERALL, D. Tool Release: Gathering 802.11n Traces with Channel State Information. *ACM SIGCOMM CCR* (2011).
- [74] HAN, S., LIM, H., AND LEE, J. An Efficient Localization Scheme for a Differential-Driving Mobile Robot Based on RFID System. *Industrial Electronics, IEEE Transactions on* (2007).
- [75] HAYES, M. H. *Statistical Digital Signal Processing and Modeling*. John Wiley & Sons, 2009.
- [76] HINSKE, S. Determining the Position and Orientation of Multi-Tagged Objects Using RFID Technology. In *IEEE Pervasive Computing and Communications Workshops* (2007).
- [77] HINTERSTOISSER, S., LEPETIT, V., ILIC, S., FUA, P., AND NAVAB, N. Dominant Orientation Templates for Real-Time Detection of Texture-less Objects.
- [78] HONG, S. S., AND KATTI, S. R. DOF: A Local Wireless Information Plane. In *Proc. of ACM SIGCOMM* (2011).
- [79] HTC. Vive. <https://www.vive.com/us/>, 2016.
- [80] HUANG, DONNY AND NANDAKUMAR, RAJALAKSHMI AND GOLLAKOTA, SHYAMNATH. Feasibility and Limits of Wi-Fi Imaging. In *ACM SenSys* (2014).
- [81] HUI, H., AND LU, S. Receiving Mutual Impedance between Two Parallel Dipole Antennas. In *Proceedings of the IEEE TENCON* (2006).
- [82] IEEE STANDARDS ASSOCIATION. IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band, 2012.
- [83] IPERF. The Ultimate Speed Test Tool for TCP, UDP and SCTP. <https://iperf.fr/>, 2015.
- [84] JAFARI, A., SARRAZIN, J., LAUTRU, D., BENLARBI-DELAÏ, A., PETRILLO, L., AND DE DONCKER, P. NLOS Influence on 60 GHz Indoor Localization Based on a New TDOA Extraction Approach. In *Microwave Conference (EuMC), 2013 European* (2013), IEEE.
- [85] JAIN, R., CHIU, D.-M., AND HAWES, W. R. *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System*. Eastern Research Laboratory, Digital Equipment Corporation Hudson, MA, 1984.

- [86] JI, Y., BIAZ, S., PANDEY, S., AND AGRAWAL, P. ARIADNE: A Dynamic Indoor Signal Map Construction and Localization System. In *Proc. of ACM MobiSys* (2006).
- [87] JOSE, E., AND ADAMS, M. An Augmented State SLAM Formulation for Multiple Line-of-Sight Features with Millimetre Wave RADAR. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2005).
- [88] JOSHI, K., HONG, S., AND KATTI, S. PinPoint: Localizing Interfering Radios. In *Proc. of USENIX NSDI* (2013).
- [89] KAMINSKI, J. Y., TEICHER, J., AND SHAVIT, A. Head Orientation and Gaze Detection from a Single Image. In *International conference of computer vision theory and applications* (2006).
- [90] KANDULA, S., LIN, K. C.-J., BADIRKHANLI, T., AND KATABI, D. FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In *USENIX NSDI* (2008).
- [91] KANG, W., AND HAN, Y. SmartPDR: Smartphone-Based Pedestrian Dead Reckoning for Indoor Localization. *IEEE Sensors journal* (2015).
- [92] KANG, W., NAM, S., HAN, Y., AND LEE, S. Improved Heading Estimation for Smartphone-Based Indoor Positioning Systems. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on* (2012), IEEE.
- [93] KHAFAJI, A., SAADANE, R., EL ABBADI, J., AND BELKASMI, M. Ray Tracing Technique based 60 GHz Band Propagation Modeling and Influence of People Shadowing. *IJECSE* (2008).
- [94] KIM, K.-H., MIN, A. W., AND SHIN, K. G. Sybot: An Adaptive and Mobile Spectrum Survey System for WiFi Networks. In *Proc. of ACM MobiCom* (2010).
- [95] KRIGSLUND, R., DOSEN, S., POPOVSKI, P., DIDERIKSEN, J. L., PEDERSEN, G. F., AND FARINA, D. A Novel Technology for Motion Capture Using Passive UHF RFID Tags. *IEEE Transactions on Biomedical Engineering* (2013).
- [96] KRIGSLUND, R., POPOVSKI, P., PEDERSEN, G. F., AND BANK, K. Potential of RFID Systems to Detect Object Orientation. In *IEEE International Conference on Communications (ICC)* (2011).
- [97] KUMAR, S., GIL, S., KATABI, D., AND RUS, D. Accurate Indoor Localization with Zero Start-up Cost. In *ACM MobiCom* (2014).
- [98] KVALO. Linux wil6210 Driver. <http://git.kernel.org/pub/scm/linux/kernel/git/kvalo/ath.git/>, 2016.
- [99] LANGEN, B., LOBER, G., AND HERZIG, W. Reflection and Transmission Behaviour of Building Materials at 60 GHz. In *Proc. of IEEE PIMRC* (1994).
- [100] LEHPAMER, H. *RFID Design Principles*. Artech House, 2012.
- [101] LENOVO. Phab 2 Pro With Tango. <http://shop.lenovo.com/us/en/tango/>, 2016.
- [102] LINDSTRÖM, J., KOKKO, H., AND RANTA, E. Detecting Periodicity in Short and Noisy Time Series Data. *Oikos* 78, 2 (1997).

- [103] LIU, H., GAN, Y., YANG, J., SIDHOM, S., WANG, Y., CHEN, Y., AND YE, F. Push the Limit of WiFi Based Localization for Smartphones. In *ACM MobiCom* (2012).
- [104] LIU, K., LIU, X., AND LI, X. Guoguo: Enabling Fine-Grained Indoor Localization via Smartphone. In *MobiSys* (2013), ACM.
- [105] LIU, T., YANG, L., LIN, Q., GUO, Y., AND LIU, Y. Anchor-Free Backscatter Positioning for RFID Tags With High Accuracy. In *Proc. of IEEE INFOCOM* (2014).
- [106] LIU, X., SHETH, A., KAMINSKY, M., PAPAGIANNAKI, K., SESHAN, S., AND STEENKISTE, P. DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas. *ACM SIGCOMM Computer Communication Review* (2009).
- [107] LUI, H. S. Mutual Coupling in Antenna Arrays.
- [108] MACKENZIE, I. S., AND SOUKOREFF, R. W. Phrase Sets for Evaluating Text Entry Techniques. In *ACM CHI Extended Abstracts* (2003).
- [109] MALTSEV, A., MASLENNIKOV, R., SEVASTYANOV, A., LOMAYEV, A., AND KHORYAEV, A. Statistical Channel Model for 60 GHz WLAN Systems in Conference Room Environment. In *Antennas and Propagation (EuCAP), 2010 Proceedings of the Fourth European Conference on* (2010), IEEE.
- [110] MCELHEARN, K. Leap motion controller fails in normal conditions. <http://www.mcelhearn.com/not-a-review-leap-motioncontroller-fails-in-normal-conditions/>, 2014.
- [111] MELGAREJO, P., ZHANG, X., RAMANATHAN, P., AND CHU, D. Leveraging Directional Antenna Capabilities for Fine-grained Gesture Recognition. In *ACM UbiComp* (2014).
- [112] MI-WAVE INC. Omni-Directional Antenna Series 267. <http://www.miww.com/millimeter-wave-products/antenna-products/>, 2014.
- [113] NANZER, J. *Microwave and Millimeter-Wave Remote Sensing for Security Applications*. Artech House, 2012.
- [114] NEEKZAD, B., SAYRAFIAN-POUR, K., PEREZ, J., AND BARAS, J. S. Comparison of Ray Tracing Simulations and Millimeter Wave Channel Sounding Measurements. In *Proc. of IEEE PIMRC* (2007).
- [115] NICHOLSON, A. J., CHAWATHE, Y., CHEN, M. Y., NOBLE, B. D., AND WETHERALL, D. Improved Access Point Selection. In *Proceedings of the 4th international conference on Mobile systems, applications and services* (2006), ACM.
- [116] NIKITIN, P. V., MARTINEZ, R., RAMAMURTHY, S., LELAND, H., SPIESS, G., AND RAO, K. Phase Based Spatial Identification of UHF RFID Tags. In *RFID, 2010 IEEE International Conference on* (2010).
- [117] NITSCHKE, T., BIELSA, G., TEJADO, I., LOCH, A., AND WIDMER, J. Boon and Bane of 60 GHz Networks: Practical Insights Into Beamforming, Interference, and Frame Level Operation. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies* (2015), ACM.

- [118] NITSCHKE, T., CORDEIRO, C., FLORES, A., KNIGHTLY, E. W., PERAHIA, E., AND WIDMER, J. C. Ieee 802.11 ad: Directional 60 ghz communication for multi-gbps wi-fi, 2014.
- [119] NITSCHKE, T., FLORES, A. B., KNIGHTLY, E. W., AND WIDMER, J. Steering With Eyes Closed: Mm-Wave Beam Steering Without In-Band Measurement. In *Proc. of IEEE INFOCOM* (2015).
- [120] ORTS-ESCOLANO, S., RHEMANN, C., FANELLO, S., CHANG, W., KOWDLE, A., DEGTYAREV, Y., KIM, D., DAVIDSON, P. L., KHAMIS, S., DOU, M., TANKOVICH, V., LOOP, C., CAI, Q., CHOU, P. A., MENNICKEN, S., VALENTIN, J., PRADEEP, V., WANG, S., KANG, S. B., KOHLI, P., LUTCHYN, Y., KESKIN, C., AND IZADI, S. Holoportation: Virtual 3D Teleportation in Real-time. In *Proceedings of ACM Symposium on User Interface Software and Technology (UIST)* (2016).
- [121] OSSBERGER, G., BUCHEGGER, T., SCHIMBACK, E., STELZER, A., AND WEIGEL, R. Non-invasive respiratory movement detection and monitoring of hidden humans using ultra wideband pulse radar. In *Conference on Ultrawideband Systems and Technologies* (2004).
- [122] PARK, M., GOPALAKRISHNAN, P., AND ROBERTS, R. Interference Mitigation Techniques in 60 GHz Wireless Networks. *IEEE Communications Magazine* (2009).
- [123] PASTERNAK INC. 60 GHz Transmit/Receive (Tx/Rx) Development System. <http://www.pasternack.com/60-ghz-test-development-system-pem003-kit-p.aspx>, 2015.
- [124] PASTERNAK INC. WR-15 Waveguide Horn Antenna Operating From 56 GHz to 66 GHz. <http://www.pasternack.com/horn-antenna-50-75-ghz-nominal-34-dbi-gain-wr-15-pe9881-34-p.aspx>, 2015.
- [125] PATWARI, N., BREWER, L., TATE, Q., KALTIOKALLIO, O., AND BOCCA, M. Breathfinding: A Wireless Network That Monitors and Locates Breathing in a Home. *JSAC* 8, 1 (2014).
- [126] PC MAGAZINE. Wilocity Unveils Blazing Fast 802.11ad Smartphone Wi-Fi Chip, 2014.
- [127] PELLERANO, S., ALVARADO, J., AND PALASKAS, Y. A mm-wave power-harvesting rfid tag in 90 nm cmos. *IEEE Journal of Solid-State Circuits* (2010).
- [128] PICHAPATI, V., KOWSHIK, H., SUBRAMANIAN, A. P., KOKKU, R., AND CHETLUR, M. Location Assisted Handoffs in Dense Cellular Networks. In *IEEE International Conference on Sensing, Communication, and Networking (SECON)* (2014).
- [129] PRIYANTHA, N. B., MIU, A. K., BALAKRISHNAN, H., AND TELLER, S. The Cricket Compass for Context-Aware Mobile Applications. In *Proc. of ACM MobiCom* (2001).
- [130] PU, Q., GUPTA, S., GOLLAKOTA, S., AND PATEL, S. Whole-home Gesture Recognition using Wireless Signals. In *ACM MobiCom* (2013).
- [131] PURSULA, P., KARTTAAVI, T., KANTANEN, M., LAMMINEN, A., HOLMBERG, J., LAHDES, M., MARTTILA, I., LAHTI, M., LUUKANEN, A., AND VAHA-HEIKKILA, T. 60-GHz Millimeter-wave Identification Reader on 90-nm CMOS and LTCC. *IEEE Transactions on Microwave Theory and Techniques* (2011).
- [132] PUTHENPURA, S. Understanding the Science Behind Small Cell Deployment, 2013.

- [133] QUALCOMM. Qualcomm 802.11ad Products to Lead the Way for Multi-band Wi-Fi Ecosystem. <https://www.qualcomm.com/news/releases/2016/01/05/qualcomm-80211ad-products-lead-way-multi-band-wi-fi-ecosystem>, 2016.
- [134] RAMACHANDRAN, K., PRASAD, N., HOSOYA, K., MARUHASHI, K., AND RANGARAJAN, S. Adaptive beamforming for 60 ghz radios: Challenges and preliminary solutions. In *Proceedings of the ACM International Workshop on mmWave Communications: From Circuits to Networks (mmCom)* (2010).
- [135] RAPPAPORT, T., GUTIERREZ, F., BEN-DOR, E., MURDOCK, J., QIAO, Y., AND TAMIR, J. Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications. *IEEE Transactions on Antennas and Propagation* (2013).
- [136] RAPPAPORT, T. S., JR., R. W. H., DANIELS, R. C., AND MURDOCK, J. N. *Millimeter Wave Wireless Communications*. Prentice Hall, 2014.
- [137] RAVINDRANATH, L., NEWPORT, C., BALAKRISHNAN, H., AND MADDEN, S. Improving Wireless Network Performance Using Sensor Hints. In *Proc. of USENIX NSDI* (2011).
- [138] RICE UNIVERSITY. Wireless Open-Access Research Platform. <http://warp.rice.edu/trac/wiki>, 2013.
- [139] RIISGAARD, S., AND BLAS, M. R. SLAM for Dummies: A Tutorial Approach to Simultaneous Localization and Mapping, 2003. MIT 16-412j Notes.
- [140] ROBERTI, M. EPCglobal Ratifies Gen 2 Standard. *RFID Journal* (2004).
- [141] ROBINSON, J., SWAMINATHAN, R., AND KNIGHTLY, E. W. Assessment of Urban-scale Wireless Networks with a Small Number of Measurements. In *Proc. of ACM MobiCom* (2008).
- [142] ROWELL, C., AND LAM, E. Y. Mobile-phone antenna design. *IEEE Antennas and Propagation Magazine* (2012).
- [143] ROY, R., AND KAILATH, T. ESPRIT-Estimation of Signal Parameters Via Rotational Invariance Techniques. *IEEE Transactions on Acoustics, Speech and Signal Processing* (1989).
- [144] SAHA, S. K., VIRA, V. V., GARG, A., AND KOUTSONIKOLAS, D. 60 GHz Multi-Gigabit Indoor WLANs: Dream or Reality? *arXiv preprint arXiv:1509.04274* (2015).
- [145] SANJIB SUR, IOANNIS PEFKIANAKIS, X. Z. K.-H. K. Ubiquitous Gbps Wireless Access Network. In *Proceedings of ACM MobiCom* (2018).
- [146] SANTHAPURI, N. K., MANWEILER, J., SEN, S., BAO, X., CHOUDHURY, R. R., AND NELAKUDITI, S. Sensor Assisted Wireless Communication. In *Proc. of IEEE Workshop on Local and Metropolitan Area Networks (LANMAN)* (2010).
- [147] SAXENA, A., DRIEMEYER, J., AND NG, A. Y. Robotic Grasping of Novel Objects Using Vision. *The International Journal of Robotics Research* (2008).
- [148] SAXENA, A., DRIEMEYER, J., AND NG, A. Y. Learning 3-D Object Orientation From Images. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on* (2009).

- [149] SCHMIDT, R. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Transactions on Antennas and Propagation* (1986).
- [150] SEN, S., RADUNOVIC, B., AND ET AL. Spot Localization using PHY Layer Information. In *ACM MobiSys* (2012).
- [151] SHANGGUAN, L., YANG, Z., LIU, A. X., ZHOU, Z., AND LIU, Y. Relative Localization of RFID Tags Using Spatial-Temporal Phase Profiling. In *Proc. of USENIX NSDI* (2015).
- [152] SHEKYAN, S., AND HARUTYUNYAN, A. To watch or be watched: Turning your surveillance camera against you.“. *Hack In The Box. Lecture conducted in Amsterdam* (2013).
- [153] SHIMIZU, S., NAGAHASHI, T., AND FUJIYOSHI, H. Robust and Accurate Detection of Object Orientation and ID Without Color Segmentation. In *RoboCup 2005: Robot Soccer World Cup IX*. 2005.
- [154] SHIREHJINI, A. A. N., YASSINE, A., AND SHIRMOHAMMADI, S. An RFID-Based Position and Orientation Measurement System for Mobile Objects in Intelligent Environments. *Instrumentation and Measurement, IEEE Transactions on* (2012).
- [155] SHU, Y., SHIN, K. G., HE, T., AND CHEN, J. Last-Mile Navigation Using Smartphones. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), ACM.
- [156] SMITH, J. R., SAMPLE, A. P., POWLEDGE, P. S., ROY, S., AND MAMISHEV, A. A Wirelessly-Powered Platform for Sensing and Computation. In *ACM UbiComp*. 2006.
- [157] SMULDERS, P. Exploiting the 60 GHz Band for Local Wireless Multimedia Access: Prospects and Future Directions. *IEEE Communications Magazine* (2002).
- [158] SMULDERS, P. F. M. Statistical Characterization of 60-GHz Indoor Radio Channels. *IEEE Transactions on Antennas and Propagation* (2009).
- [159] SOH, W.-S., AND KIM, H. S. QoS Provisioning in Cellular Networks Based on Mobility Prediction Techniques. *IEEE Communications Magazine* 41, 1 (2003).
- [160] STEINMETZER, D., WEGEMER, D., SCHULZ, M., WIDMER, J., AND HOLLICK, M. Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11ad devices. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies* (2017).
- [161] STEVANOVIC, I., SKRIVERVIK, A., AND MOSIG, J. R. Smart Antenna Systems for Mobile Communications. Tech. rep., 2003.
- [162] STOJMENOVIC, I. *Handbook of Wireless Networks and Mobile Computing*. John Wiley & Sons, 2003.
- [163] SULEESATHIRA, R., AND KUNARAK, S. Neural Network Handoff in Shadow-Rayleigh Fading. In *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on* (2005), IEEE.
- [164] SUR, S., PEFKIANAKIS, I., ZHANG, X., AND KIM, K.-H. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom* (2017).

- [165] SUR, S., VENKATESWARAN, V., ZHANG, X., AND RAMANATHAN, P. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proceedings of ACM SIGMETRICS* (2015).
- [166] SUR, S., ZHANG, X., RAMANATHAN, P., AND CHANDRA, R. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)* (2016).
- [167] TAN, K., FANG, J., ZHANG, Y., CHEN, S., SHI, L., ZHANG, J., AND ZHANG, Y. Fine-grained Channel Access in Wireless LAN. In *SIGCOMM* (2011).
- [168] TAPAN PATTNAYAK, G. T. Antenna Design and RF Layout Guidelines. Cypress Semiconductor, 2016.
- [169] TODD, M. J., AND YILDIRIM, E. A. On Khachiyan’s Algorithm for the Computation of Minimum-Volume Enclosing Ellipsoids. *Discrete Applied Mathematics* (2007).
- [170] TPCAST. HTC VIVE Wireless Upgrade Kit. http://www.tpcast.cn/h_en/index.html, 2016.
- [171] VALDES-GARCIA, A., REYNOLDS, S., NATARAJAN, A., KAM, D., LIU, D., LAI, J.-W., HUANG, Y.-L., CHEN, P.-Y., TSAI, M.-D., ZHAN, J.-H., ET AL. Single-element and phased-array transceiver chipsets for 60-GHz Gb/s communications. *Communications Magazine, IEEE* (2011).
- [172] VAN VEEN, B. D., AND BUCKLEY, K. M. Beamforming: A Versatile Approach to Spatial Filtering. *IEEE ASSP Magazine* (1988).
- [173] VASISHT, D., KUMAR, S., AND KATABI, D. Decimeter-level Localization With a Single WiFi Access Point. In *Proceedings of Usenix Conference on Networked Systems Design and Implementation (NSDI)* (2016).
- [174] WACHS, J. P., KÖLSCH, M., STERN, H., AND EDAN, Y. Vision-based Hand-gesture Applications. *Communications of the ACM* 54, 2 (2011).
- [175] WAHEED, A., AND ASHIK, A. Mitigation of Phase Noise at Millimeter-Wave Frequencies for Wireless Personal Area Network Applications. *Masters Theses* (2008).
- [176] WAN, W., LU, F., WU, Z., AND HARADA, K. Teaching Robots to Do Object Assembly Using Multi-Modal 3D Vision. *arXiv preprint arXiv:1601.06473* (2016).
- [177] WANG, H., FARID, M., SEN, S., YOUSSEF, M., ELGOHARY, A., AND CHOUDHURY, R. R. No Need to War-drive: Unsupervised Indoor Localization. In *Proceedings of ACM MobiSys* (2012).
- [178] WANG, J., ADIB, F., KNEPPER, R., KATABI, D., AND RUS, D. RF-Compass: Robot Object Manipulation Using RFIDs. In *Proc. of ACM MobiCom* (2013).
- [179] WANG, J., AND KATABI, D. Dude, where’s my card?: RFID Positioning That Works With Multipath and Non-Line of Sight. *ACM SIGCOMM Computer Communication Review* (2013).

- [180] WANG, J., LAN, Z., SUM, C.-S., PYO, C.-W., GAO, J., BAYKAS, T., RAHMAN, A., FUNADA, R., KOJIMA, F., LAKKIS, I., ET AL. Beamforming Codebook Design and Performance Evaluation for 60GHz Wideband WPANs. In *IEEE Trans. on Vehicular Technology* (2009).
- [181] WANG, J., VASISHT, D., AND KATABI, D. RF-IDraw: Virtual Touch Screen in the Air using RF Signals. In *ACM SIGCOMM* (2014).
- [182] WANG, J. J., HU, G., HUANG, S., AND DISSANAYAKE, G. 3D Landmarks Extraction from a Range Imager Data for SLAM. In *Proc. of Australasian Conference on Robotics and Automation* (2009).
- [183] WANT, R. An Introduction to RFID Technology. *Pervasive Computing, IEEE* (2006).
- [184] WEI, B., VARSHNEY, A., HU, W., PATWARI, N., AND CHOU, C. T. dRTI: Directional RadioTomographic Imaging. *CoRR abs/1402.2744* (2014).
- [185] WEI, T., AND ZHANG, X. mTrack: High-Precision Passive Tracking Using Millimeter Wave Radios. In *Proc. of ACM MobiCom* (2015).
- [186] WEI, T., AND ZHANG, X. Gyro in the Air: Tracking 3D Orientation of Batteryless Internet-of-Things. In *Proc. of ACM MobiCom* (2016).
- [187] WEI, T., AND ZHANG, X. Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support. In *Proc. of ACM MobiCom* (2017).
- [188] WEI, T., ZHOU, A., AND ZHANG, X. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *NSDI* (2017).
- [189] WERNER, M., KESSEL, M., AND MAROUANE, C. Indoor Positioning Using Smartphone Camera. In *Indoor Positioning and Indoor Navigation (IPIN), 2011 International Conference on* (2011), IEEE.
- [190] WIKIPEDIA. Friis Transmission Equation. https://en.wikipedia.org/wiki/Friis_transmission_equation, 2016.
- [191] WIKIPEDIA. Rotation Matrix in 3D. https://en.wikipedia.org/wiki/Rotation_matrix, 2016.
- [192] WIKIPEDIA. Four-Quadrant Inverse Tangent. <https://en.wikipedia.org/wiki/Atan2>, 2017.
- [193] WIKIPEDIA. Geodesic Grid. https://en.wikipedia.org/wiki/Geodesic_grid, 2017.
- [194] WIRE, B. FCC Adopts New Rules for Unlicensed V-Band, Extending the Coverage to 57-71 GHz, 2016.
- [195] WITTEN, I. H., AND FRANK, E. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann Publishers, 2005.
- [196] WONG, K. D., AND COX, D. C. A Pattern Recognition System for Handoff Algorithms. *IEEE Journal on selected areas in communications* (2000).
- [197] XIONG, J., AND JAMIESON, K. ArrayTrack: A Fine-Grained Indoor Location System. In *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)* (2013).

- [198] XU, C., FIRNER, B., MOORE, R. S., ZHANG, Y., TRAPPE, W., HOWARD, R., ZHANG, F., AND AN, N. SCPL: Indoor Device-free Multi-subject Counting and Localization Using Radio Signal Strength. In *Proc. of ACM/IEEE IPSN* (2013).
- [199] XU, H., KUKSHYA, V., AND RAPPAPORT, T. Spatial and Temporal Characteristics of 60-GHz Indoor Channels. *IEEE Journal on Selected Areas in Communications* (2002).
- [200] YANG, L., CHEN, Y., LI, X.-Y., XIAO, C., LI, M., AND LIU, Y. Tagoram: Real-Time Tracking of Mobile RFID Tags to High Precision Using COTS Devices. In *ACM MobiCom* (2014).
- [201] YANG, T., LIU, J., MCMILLAN, L., AND WANG, W. A Fast Approximation to Multidimensional Scaling. In *IEEE workshop on Computation Intensive Methods for Computer Vision* (2006).
- [202] YANG, Z., PATHAK, P. H., ZENG, Y., AND MOHAPATRA, P. Sensor-Assisted Codebook-Based Beamforming for Mobility Management in 60 Ghz WLANs. In *IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS)* (2015).
- [203] YANG, Z., WU, C., AND LIU, Y. Locating in Fingerprint Space: Wireless Indoor Localization with Little Human Intervention. In *Proc. of ACM MobiCom* (2012).
- [204] YEAGER, D. J., SAMPLE, A. P., SMITH, J. R., AND SMITH, J. R. Wisp: A Passively Powered UHF RFID Tag With Sensing and Computation. *RFID Handbook: Applications, Technology, Security, and Privacy* (2008).
- [205] YOUSSEF, M., MAH, M., AND AGRAWALA, A. Challenges: Device-free Passive Localization for Wireless Environments. In *Proc. of ACM MobiCom* (2007).
- [206] ZHANG, C., AND ZHANG, X. LiTell: Robust Indoor Localization Using Unmodified Light Fixtures. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016).
- [207] ZHANG, J., ZHANG, X., KULKARNI, P., AND RAMANATHAN, P. OpenMili: A 60 GHz Software Radio Platform with a Reconfigurable Phased-Array Antenna. In *Proceedings of ACM MobiCom* (2016).
- [208] ZHANG, Y., WANG, B., ZHANG, H., AND KONG, W. True 3D Antenna for UHF RFID Application. In *6th Asia Symposium on Quality Electronic Design (ASQED)* (2015).
- [209] ZHENG, Y., SHEN, G., LI, L., ZHAO, C., LI, M., AND ZHAO, F. Travi-Navi: Self-Deployable Indoor Navigation System. In *Proceedings of the 20th annual international conference on Mobile computing and networking* (2014), ACM.
- [210] ZHOU, A., WU, L., XU, S., MA, H., WEI, T., AND ZHANG, X. Following the shadow: Agile 3-d beam-steering for 60 ghz wireless networks. In *IEEE Conference on Computer Communications* (2018).
- [211] ZHOU, A., ZHANG, X., AND MA, H. Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering. In *Proc. of IEEE INFOCOM* (2016).
- [212] ZHOU, P., LI, M., AND SHEN, G. Use It Free: Instantly Knowing Your Phone Attitude. In *Proc. of ACM MobiCom* (2014).

- [213] ZHU, S., AND ZHANG, X. Enabling High-Precision Visible Light Localization in Today's Buildings. In *Proc. of ACM MobiSys* (2017).
- [214] ZHU, X., MUKHOPADHYAY, S. K., AND KURATA, H. A Review of RFID Technology and Its Managerial Applications in Different Industries. *Journal of Engineering and Technology Management* (2012).
- [215] ZHU, Y., ZHANG, Z., MARZI, Z., NELSON, C., MADHOW, U., ZHAO, B. Y., AND ZHENG, H. Demystifying 60GHz Outdoor Picocells. In *Proceedings of ACM MobiCom* (2014).
- [216] ZHU, Y., ZHU, Y., ZHANG, Z., ZHAO, B. Y., AND ZHENG, H. 60 GHz Mobile Imaging Radar. HotMobile '15.
- [217] ZHU, Y., ZHU, Y., ZHAO, B. Y., AND ZHENG, H. Reusing 60GHz Radios for Mobile Radar Imaging. In *Proc. of ACM MobiCom* (2015).