Copyright
by
Ian Parker Roberts
2023

The Dissertation Committee for Ian Parker Roberts
certifies that this is the approved version of the following dissertation:

Full-Duplex Millimeter Wave Communication Systems: Theory and Practice

Committee:

Jeffrey G. Andrews, Co-Supervisor

Sriram Vishwanath, Co-Supervisor

Brian L. Evans

Hyeji Kim

Aditya Chopra

Full-Duplex Millimeter Wave Communication Systems: Theory and Practice

by

Ian Parker Roberts

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

The University of Texas at Austin

August 2023

Acknowledgments

My graduate studies have been immensely fulfilling and enjoyable. Five years ago, I would have never thought that things would have turned out the way they did. It is important to me that I take this opportunity to recognize the many people that have contributed to and supported my education and this journey. First, I would like to thank my advisors Jeffrey G. Andrews and Sriram Vishwanath. They continuously supported me and my graduate studies and provided me an abundance of freedom, allowing me to maintain independence and take ownership of my research. Jeff has been a consistent role model for me through his ability to deliver high-quality research and teaching while maintaining an unwavering positive and professional attitude. His optimistic yet pragmatic approach to wireless research is something I have appreciated and admired these past several years. Sriram has demonstrated an unyielding work ethic, extreme humility, and utmost kindness without exception throughout my graduate studies, and this has helped shape me both professionally and personally. I have learned a great deal from both of them and am truly fortunate to have been their student.

I would also like to thank my dissertation committee members: Brian L. Evans, Hyeji Kim, and Aditya Chopra. Each one of them has given me unique feedback, direction, and support in my research and career pursuits. I am incredibly fortunate to have collaborated with Aditya Chopra, Tom Novlan, and Milap Majmundar from AT&T Labs throughout my Ph.D.; it has been a true pleasure working with them on fulfilling research that I am extremely proud of. I would also like to thank Ahmed Alkhateeb at Arizona State University and Chan-Byoung Chae at Yonsei University for each hosting me as a visiting student during summer 2022; these experiences have broadened my research perspectives and portfolio and include some of my fondest memories of graduate school. I would also like to acknowledge Himal A. Suraweera at the University of Peradeniya for his consistent encouragement and support; he has played a central role in my service and involvement in the full-duplex research community at large.

I would especially like to thank those closest to me, in Austin and from afar, that have supported me during graduate school: Yesong Jeon, Tristan and CJ Amann, Eunsun Kim and her family, Rajesh Mishra, Agrim Bari, Mark Kamphoefner, Albert Reed, Logan Green, Amudheesan Nakkeeran, Kunal Sutaria, Ronald Palomares, and Gee Yong Suk. I have always considered myself extremely lucky to have so many people support and encourage me.

Being part of the Wireless Networking and Communications Group (WNCG) has been a highlight of my life; the students, faculty, and staff made it an incredibly welcoming, optimistic, and energetic atmosphere to do wireless research. I would like to thank many friends within WNCG not already mentioned: Ahmad AlAmmouri, Manan Gupta, Ethan Heng, Akash Doshi, Ezgi Tekgul, Nick Olson, Taekyun Lee, Juseong Park, Alperen Duru, Yunseong

Cho, Marius Arvinte, Kartik Patel, Nitin J. Myers, Yuyang Wang, Satyam Kumar, Sidharth Kumar, Rudrajit Das, Pooja Nuti, Hasan Beytur, Zach Clements, and Ali Lotfi Rezaabad. It would be difficult to overstate how the culture, precedent, and accomplishments of current and previous generations of WNCG-ers made it a truly special place to learn and work on wireless.

I would like to thank the professors at UT Austin who have provided me foundational knowledge and skills to conduct research, especially Jeff Andrews, Robert Heath, Sanjay Shakkottai, and Aryan Mokhtari. I would also like to recognize the WNCG staff and department staff: Karen Little, Jaymie Udan, Apipol Piman, Erin Salcher, Melanie Gulick, Barry Levitch, and Melody Singleton; they have truly made WNCG and Texas ECE exceptional.

I would very much like to thank the National Science Foundation (and the taxpayers!) for supporting my research through their Graduate Research Fellowship Program, which afforded me generous academic freedom to focus on problems I deemed important. Many thanks go to Mark Kamphoefner and Angel Syrett for helping me exhaustively refine my fellowship application.

I had no idea at the time that joining GenXComm for a summer internship before starting graduate school would have such lasting impacts; in fact, it is quite possible that I would not have done a Ph.D. if I had not had such an experience. I would like to thank the many employees at GenXComm that played unique roles introducing me to full-duplex, most especially: Sriram Vishwanath, Hardik Jain, Gautam Talreja, Anish Nair, Rajesh Mishra, Ronald Palomares, Paul Lim, Chechun Lin, Ousek Son, Sam Kokel, Farzad

Mokhtari-Koushyar, John Lane, and Marco Hernandez. I will always cherish the fun times we had working as a close-knit team, especially during the summer of 2018. I am extremely fortunate to have thoroughly enjoyed the topic of my Ph.D. research, and this is largely thanks to my time at GenXComm.

My undergraduate education from Missouri University of Science and Technology provided me a foundation on which I could pursue my career aspirations; for this I must thank Rosa Zheng and Reza Zoughi for exceptional fundamental courses in wireless. I would especially like to thank Rosa Zheng and her former Ph.D. student Mohammadhossein Behgam for fostering my interests in wireless through undergraduate research and for supporting my pursuit of a graduate degree. I would also like to thank those from Jackson R-2 School District for cultivating in me an enjoyment of science, math, and technology: Ms. Nancy, Kerri Middleton, Blake Miller, Tonya Skinner, and especially Lance McClard.

I would like to thank the friendly faces that organized and participated in Frisbee Fridays, along with our intramural teams Hammer Time and Yuka's Angels, for countless evenings of ultimate frisbee over the past few years—something I love even more than wireless! I will also always remember many fun afternoons playing cricket with far too many people to list here.

Finally, I would like to thank my family, especially my grandfather and my uncle, for instilling in me a foundation on which I have been able to pursue my career goals and much more.

Full-Duplex Millimeter Wave Communication Systems: Theory and Practice

Publication No. ________________

Ian Parker Roberts, Ph.D.
The University of Texas at Austin, 2023

Co-Supervisors: Jeffrey G. Andrews
Sriram Vishwanath

This dissertation is focused on studying, characterizing, and developing new solutions for full-duplex millimeter wave (mmWave) communication systems. Translating existing full-duplex solutions—which are largely for systems operating at lower carrier frequencies—to mmWave systems is not practically viable nor is it necessarily desirable. Instead, I propose several novel solutions for full-duplex mmWave systems, all of which harness beamforming with dense antenna arrays to cancel self-interference to levels near or even below the noise floor—unlocking the long-sought ability to simultaneously transmit and receive over the same frequency spectrum. Two trademarks of my proposed techniques are that they (i) account for a wide variety of key practical constraints and (ii) are informed by a campaign in which I collected and analyzed nearly 20 million measurements of self-interference using 28 GHz phased arrays. Numerical simulation and experimental evaluation are used to illustrate the effectiveness of

the proposed techniques, validating them as enablers of full-duplex mmWave systems. Based on the aforementioned measurements, this dissertation also extensively characterizes mmWave self-interference through statistical analysis and spatial modeling, both of which are the first of their kind. Together comprising this dissertation, the characterization of self-interference and the proposed solutions serve as strides toward transforming real-world mmWave communication systems by upgrading them with full-duplex capability.

Table of Contents

Acknowledgments **4**

Abstract **8**

List of Tables **15**

List of Figures **16**

List of Acronyms **26**

Chapter 1. Introduction **29**

1.1 Enhancements Introduced by Full-Duplex 31
1.2 Enabling Full-Duplex Millimeter Wave Systems 35
1.3 Summary of Contributions and Organization 37
1.4 Related Publications 39
1.5 Notation 42

Chapter 2. Background and Motivation **44**

2.1 Traditional Low-Frequency Full-Duplex Solutions 44
2.1.1 Analog Self-Interference Cancellation 45
2.1.2 Digital Self-Interference Cancellation 47
2.1.3 Circulators and Antenna Isolation 51
2.2 Millimeter Wave Communication Systems 53
2.2.1 Hybrid Beamforming and Analog Beamforming 54
2.2.2 Beam Alignment 55
2.3 Full-Duplex Solutions at Millimeter Wave 57
2.3.1 Extending Analog and Digital SIC 57
2.3.2 Tackling Self-Interference with Beamforming 60

Chapter 3. Hybrid Beamforming for Full-Duplex Millimeter Wave Communication Systems **63**
3.1 Introduction 63
3.1.1 Prior Work and Motivation 64
3.1.2 Contributions 67
3.2 System Model 68
3.3 Problem Formulation 73
3.4 Beam Candidate Acquisition 77
3.5 Hybrid Beamforming Design 83
3.5.1 Per-Antenna and Per-RF Chain Power Constraints . . . 84
3.5.2 Satisfying Our Constraints 86
3.5.3 Constraint Redundancy Conditions 92
3.5.4 Remainder of the Design 94
3.6 Numerical Results 98
3.7 Conclusion 108

Chapter 4. Analog Beamforming Codebooks for Full-Duplex Millimeter Wave Systems **110**
4.1 Introduction 111
4.2 System Model 116
4.3 Problem Formulation and Motivation 119
4.4 Beamforming Codebook Design 123
4.4.1 Quantifying Our Codebook Design Criteria 126
4.4.2 Assembling Our Codebook Design Problem 129
4.5 Solving for LoneSTAR Codebooks and Remarks 134
4.6 Simulation Setup and Performance Metrics 139
4.6.1 Baselines and Performance Metrics 141
4.6.2 Beam Patterns 142
4.6.3 Choosing Coverage Variances, σ^2_{tx} and σ^2_{rx} 143
4.7 Evaluating LoneSTAR Codebooks 146
4.7.1 For Various Link Qualities, $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$ 151
4.7.2 For Various Levels of Interference, $\left(\mathsf{INR}_{\mathrm{tx}}, \overline{\mathsf{INR}}_{\mathrm{rx}}\right)$ 151
4.7.3 Impact of Channel Estimation Error and Channel Structure 153
4.8 Conclusion and Future Directions 157

Chapter 5. Self-Interference Measurements at 28 GHz **159**
5.1 Introduction 159
5.1.1 Prior Work Characterizing mmWave Self-Interference . . 160
5.1.2 Contributions 163
5.2 Measurement Setup and Methodology 164
5.3 High-Level Summary and Spatial Insights 173
5.3.1 Is an Idealized Near-Field Channel Model Realistic? . . 175
5.3.2 Maximum, Median, and Minimum INR 177
5.3.3 Meeting INR Thresholds 180
5.3.4 INR for Particular Transmit-Receive Beam Pairs 182
5.4 Angular Spread of mmWave Self-Interference 185
5.4.1 INR Range Over Various Neighborhoods 189
5.4.2 Minimum INR Over Various Neighborhoods 191
5.4.3 Maximum INR Over Various Neighborhoods 198
5.5 Conclusion 203

Chapter 6. Spatial and Statistical Modeling of Millimeter Wave Self-Interference **209**
6.1 Introduction 209
6.1.1 Existing Models and Measurements 210
6.1.2 Contributions 212
6.1.3 Notation 214
6.2 Measurements of mmWave Self-Interference 214
6.3 Motivation and Roadmap 218
6.4 Self-Interference Over Small Spatial Neighborhoods 220
6.4.1 Confirming Log-Normal Nature Across All Measurements 222
6.4.2 What Causes This Log-Normal Nature? 224
6.5 Spatial and Statistical Modeling 226
6.5.1 Estimating Neighborhood Mean, μ_{ij} 227
6.5.2 Estimating Neighborhood Variance, σ_{ij}^2. 237
6.5.3 Summary 240
6.6 Generalization and Applications of Our Model 241

6.6.1 Fitting Our Model to Other Configurations 241
6.6.2 Applications of Our Model 245
6.7 Evaluating Our Statistical Model 246
6.7.1 Comparing Our Statistical Model to Measurements . . . 247
6.7.2 Using Our Model to Design and Evaluate Solutions . . . 251
6.8 Conclusion and Topics for Future Work 255

Chapter 7. Beam Selection for Full-Duplex Millimeter Wave Communication Systems 257
7.1 Introduction 257
7.2 System Model 261
7.3 Problem Formulation and Motivation 264
7.4 Initial Selection: Conventional Beam Alignment 267
7.5 Self-Interference Across Small Spatial Neighborhoods 270
7.6 STEER: Joint Transmit-Receive Beam Selection 276
7.6.1 Target Self-Interference Level, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ 277
7.6.2 Joint Transmit-Receive Beam Selection 277
7.7 Efficiently Implementing STEER in Real Systems 282
7.8 Evaluating STEER via Measurement and Simulation 287
7.8.1 Reducing Measurement Overhead via Algorithm 7.1 . . 289
7.8.2 Performance Metrics 290
7.8.3 Choosing the Self-Interference Target, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$. 292
7.8.4 Full-Duplex with STEER vs. Other Multiplexing Strategies 293
7.8.5 Impact of Neighborhood Size on STEER's Performance . 295
7.8.6 For Disparate Transmit and Receive Links 299
7.9 Conclusion 302

Chapter 8. Concluding Remarks 303
8.1 Summary of Contributions 303
8.2 Topics for Future Work 305
8.2.1 Extensions of STEER and LONESTAR 305
8.2.2 Shared-Antenna Full-Duplex mmWave Systems 306
8.2.3 Real-Time, Adaptive Full-Duplex Solutions 306

8.2.4 Network-Level Studies of Full-Duplex 307
8.2.5 Capitalizing on Advances in Machine Learning 307

Bibliography **308**

Vita **327**

List of Tables

5.1 Fitted $(\alpha_{\text{rng}}(\Delta\theta, \Delta\phi), \beta_{\text{rng}}(\Delta\theta, \Delta\phi))$. 204
5.2 Fitted $(\mu_{\text{min}}(\Delta\theta, \Delta\phi), \sigma^2_{\text{min}}(\Delta\theta, \Delta\phi))$. 205
5.3 Fitted $(\alpha_{\text{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\text{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}))$. 206
5.4 Fitted $(\mu_{\text{max}}(\Delta\theta, \Delta\phi), \sigma^2_{\text{max}}(\Delta\theta, \Delta\phi))$. 207
5.5 Fitted $(\alpha_{\text{max}}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\text{max}}(\Delta\theta, \Delta\phi, \mathsf{INR}))$. 208

6.1 Dominant cluster centers in our measurements ($N_{\text{clust}} = 4$). . . 231
6.2 Measurement system parameters and fitted model parameters. 243

7.1 Transmit and receive links during different full-duplexing modes. 263

List of Figures

1.1 FDD and TDD orthogonalize transmission and reception in the frequency and time domains, respectively. Full-duplex allows a device to transmit and receive on the same time-frequency resources. In practice, configuring a full-duplex solution incurs some overhead and FDD and TDD require a guard band/interval. 30

1.2 A full-duplex base station transmitting downlink and receiving uplink in-band with a full-duplex user. Self-interference manifests at both devices, requiring each to take measures to sufficiently mitigate such. 31

1.3 A full-duplex base station transmitting downlink to one half-duplex user while receiving uplink from another half-duplex user in-band. Self-interference manifests at the base station, whereas cross-link interference is inflicted onto the downlink user by the uplink user. 32

1.4 An access point equipped with full-duplex capability can sense the channel while transmitting, allowing it to reroute resources in the presence of interference that would otherwise cause collisions [1]. 33

1.5 (a) Without full-duplex capability, a sectorized base station must either transmit or receive across all sectors. (b) With inter-sector full-duplex, the base station is capable of transmitting or receiving with any of its sectors; self-interference is coupled across sectors. (c) With intra-sector full-duplex, the base station can transmit and receive within the same sector; self-interference is coupled within the same sector and potentially across sectors. 34

1.6 Traditionally, wireless networks like 4G cellular systems have made use of sub-6 GHz spectrum for reliable, widespread coverage. Modern 5G cellular systems have made use of higher, mmWave carrier frequencies to leverage wider bandwidths. Upper mid-band and sub-THz spectrum are being considered for future 6G cellular systems, largely for the same reason. 36

2.1 Self-interference manifests when a full-duplex transceiver attempts to simultaneously transmit and receive using the same frequency spectrum. Here, separate antennas are used for transmission and reception, with analog and digital SIC used to reconstruct and subsequently cancel self-interference incurred at the receiver. . . . 46

2.2 Self-interference combines with a desired receive signal at each receive antenna of a full-duplex transceiver. When self-interference is much stronger than the desired receive signal, it can saturate components that have a limited dynamic range, such LNAs and ADCs. This saturation should be prevented throughout the receive chain to ensure that digital SIC can be applied successfully and the desired signal can be received reliably. . . . 48

2.3 The transfer function of a practical LNA exhibiting gain compression. Amplifier gain is constant for low-power input signals, but beyond a certain input power, the amplifier is unable to maintain this gain. . . . 50

2.4 A circulator can be used to establish isolation between a transmitter and receiver sharing an antenna. Leakage through the circulator and reflections off the environment give rise to self-interference, however. . . . 51

2.5 A mmWave communication system where both the transmitter and receiver employ hybrid digital/analog beamforming. This staged beamforming architecture allows each to control dense antenna arrays with a reduced number of RF chains. . . . 54

2.6 A mmWave base station conducts beam alignment to close the link between it and the user it intends to serve. This is typically achieved by sweeping a set of beams (a codebook), measuring the received signal quality at the user for each beam, and then selecting the beam that maximized such. . . . 56

2.7 A full-duplex mmWave base station transmits downlink to a user with beamforming weights $\mathbf{f}$ while receiving from an uplink user with beamforming weights $\mathbf{w}$. In doing so, the transmit and receive beams couple over the self-interference MIMO channel $\mathbf{H}$ and the uplink user inflicts cross-link interference on the downlink user. . . . 59

3.1 A full-duplex mmWave device i transmitting to a device j as it receives from a device k in-band. . . . 69

3.2 Precoding-combining candidates exploit spatial diversity to deliver higher downlink and uplink quality while ensuring receive chain components are not saturated. . . . 80

3.3 Sum spectral efficiency as a function of SNR for various $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB, where $\kappa = 10$ dB, $b = 12$ bits, and $K_{ij} = K_{ki} = 3$. While not obvious from this figure, $\mathcal{R}_{ki}$ is fairly robust to self-interference with $b = 12$ bits. Stricter choices of $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB degrade $\mathcal{R}_{ij}$ and, thus, the sum $\mathcal{R}_{ij} + \mathcal{R}_{ki}$. 101

3.4 Spectral efficiency as a function of SNR for various K_{ij} and K_{ki}, where $\kappa = 10$ dB, $\eta_{\mathrm{LNA}} = 15$ dB, $\eta_{\mathrm{ADC}} = -5$ dB, and $b = 12$ bits. Increasing K_{ij} and K_{ki} can only improve $\mathcal{R}_{ij}$, whereas increasing K_{ki} may degrade $\mathcal{R}_{ki}$. The choice of K_{ij} and K_{ki} that maximizes the $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ varies with each realization. . . 102

3.5 Sum spectral efficiency as a function of η_{ADC} for various η_{LNA}, where $\kappa = 10$ dB, $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $b = 12$ bits. As η_{LNA} and η_{ADC} are relaxed, the system can achieve a greater sum spectral efficiency since transmit performance is less constrained by a limited receive dynamic range. . 104

3.6 Spectral efficiency as a function of η_{ADC} for various ADC resolutions, where $\kappa = 10$ dB, $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $\eta_{\mathrm{LNA}} = 20$ dB. Higher-resolution ADCs are more robust to self-interference whereas lower-resolution ADCs saturate if η_{ADC} is not properly chosen. The tradeoff associated with constraining transmit link performance and preventing ADC saturation is not so obvious in terms of sum spectral efficiency. 105

3.7 Sum spectral efficiency as a function of κ for various η_{LNA} and η_{ADC}, where $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $b = 12$ bits. 107

4.1 A full-duplex mmWave base station transmits to a downlink user while receiving from an uplink user in-band. In doing so, self-interference couples between its transmit beam $\mathbf{f}$ and receive beam $\mathbf{w}$ across the channel $\mathbf{H}$, and cross-link interference is inflicted on the downlink user by the uplink user. Transmit and receive beams are selected from codebooks $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ via beam alignment. In this work, we design these codebooks such that any selected $\mathbf{f} \in \bar{\mathcal{F}}$ and $\mathbf{w} \in \bar{\mathcal{W}}$ will couple low self-interference. 117

4.2 Following estimation of the self-interference channel $\mathbf{H}$, transmit and receive beam codebooks designed with LoneSTAR can be used for beam alignment and subsequently to serve any downlink and uplink users in a full-duplex fashion. 125

4.3 Most existing beamforming-based full-duplex mmWave solutions require estimation of the downlink and uplink channels and execution of their design for each user pair, leading to a total overhead that scales with the number of users served. 125

4.4 Using an estimate of the self-interference channel and knowledge of the estimation error variance, LoneSTAR produces transmit and receive beamforming codebooks that serve some coverage regions with quality based on tolerated coverage variances. . . 126

4.5 (a) The azimuth cut of a broadside beam from conventional CBF and Taylor codebooks. (b) The azimuth cuts of transmit beams produced by LoneSTAR that serve various azimuth directions at an elevation of 0°. 143

4.6 Normalized sum spectral efficiency γ_{sum} (a) as a function of the coverage variance $\sigma^2_{\mathrm{tx}} = \sigma^2_{\mathrm{rx}}$ for various $\overline{\mathsf{INR}}_{\mathrm{rx}}$; and (b) as a function of self-interference strength $\overline{\mathsf{INR}}_{\mathrm{rx}}$ for various $\sigma^2_{\mathrm{tx}} = \sigma^2_{\mathrm{rx}}$. In both, $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB, $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits, $\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB, and $\epsilon^2 = -\infty$ dB. Performance with the optimal $\sigma^2_{\mathrm{tx}} = \sigma^2_{\mathrm{rx}}$ is denoted with $\star$. 144

4.7 The CDFs of (a) self-interference $\mathsf{INR}_{\mathrm{rx}}$ where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB; and (b) uplink quality $\mathsf{SNR}_{\mathrm{rx}}$ where $\overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB. LoneSTAR ($b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits) can reduce self-interference between all beam pairs while also maintaining high uplink quality, netting it greater full-duplex performance than conventional codebooks (8 bits). 146

4.8 Normalized sum spectral efficiency γ_{sum} (a) as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$ where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB; (b) as a function of $\overline{\mathsf{INR}}_{\mathrm{rx}}$ where $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB. In both, $\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB, $\epsilon^2 = -\infty$ dB, and $\sigma^2_{\mathrm{tx}} = \sigma^2_{\mathrm{rx}}$ are tuned to maximize sum spectral efficiency. CBF and Taylor codebooks are each with $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits. 149

4.9 Normalized sum spectral efficiency γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ for the (a) Taylor codebook ($b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits) and (b) LoneSTAR codebooks, where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB, $\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB, $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits, and $\epsilon^2 = -\infty$ dB. LoneSTAR delivers higher γ_{sum} broadly across SNRs and demands lower SNRs to net $\gamma_{\mathrm{sum}} \geq 0.5$. 150

4.10 Normalized sum spectral efficiency γ_{sum} as a function of $\mathsf{INR}_{\mathrm{tx}}$ and $\overline{\mathsf{INR}}_{\mathrm{rx}}$ for the (a) Taylor codebook, where $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits; and (b) LoneSTAR codebooks, where $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits. In both, $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB and $\epsilon^2 = -\infty$ dB. LoneSTAR is more robust to self-interference and cross-link interference than conventional codebooks. 152

4.11 Normalized sum spectral efficiency as a function of (a) self-interference channel estimation error variance ϵ^2 and (b) mixing variance ζ^2. In both, $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} = 10$ dB, $\overline{\mathsf{INR}}_{\text{rx}} = 90$ dB, and $\mathsf{INR}_{\text{tx}} = -\infty$ dB. The true channel $\mathbf{H}$ is distributed the same in both, but in (b), we assume perfect channel knowledge, $\epsilon^2 = -\infty$ dB. 155

5.1 A full-duplex mmWave base station will incur different levels of self-interference depending on where its transmit beam and receive beam are steered. In this work, using the 28 GHz phased array platform in Figure 5.3, we swept the beams over a dense spatial profile, measuring self-interference for each transmit-receive beam pair. 165

5.2 A simplified block diagram of our measurement setup using phased arrays to inspect self-interference at 28 GHz. The degree of self-interference depends on the particular choice of transmit beam and receive beam. 165

5.3 The 28 GHz phased array measurement platform in an anechoic chamber; receive array on left and transmit array on right. Each array is a 16×16 planar array. Self-interference was measured at the receive array output by simultaneously transmitting and receiving over the same frequency spectrum. 166

5.4 The coordinate system observed by each phased array situated at its respective origin, facing toward the positive x axis. An azimuth-elevation of $(0°, 0°)$ corresponds to steering outward (broadside) along the x axis. An increase in azimuth is rightward, toward the y axis. An increase in elevation is upward, toward the z axis. 168

5.5 The idealized azimuth radiation pattern of our 16×16 transmit UPA steered broadside. The elevation pattern is identical. The EIRP is 60 dBm, and the 3 dB beamwidth is roughly 7°. The pattern of the receive UPA steered broadside is identical. . . . 169

5.6 The set of 2,541 transmit directions $\mathcal{A}_{\text{tx}}$ used during measurement, shown here as their respective projections onto the y-z plane, where looking into the page represents steering outward from the array (the positive x direction). The set of receive directions $\mathcal{A}_{\text{rx}}$ is identical, densely spanning $-60°$ to $60°$ in azimuth and $-10°$ to $10°$ in elevation, each in 1° steps. 171

5.7 A summary of the key power levels and power ratios associated with our measurement system. 172

5.8 The CDF of the nearly 6.5 million measured INR values and its fitted log-normal distribution. 174

5.9 (a) Measured INR as a function of azimuthal transmit and receive directions when $\phi_{\text{tx}} = \phi_{\text{rx}} = 0°$. (b) The simulated counterpart of (a) when using a spherical-wave near-field self-interference channel model [2], which clearly does not align well with what was measured. This motivates the need for a new, measurement-backed model of mmWave self-interference. . . . 176

5.10 For each transmit beam and receive beam, shown are the maximum, median, and minimum INR across all receive and transmit beams, respectively. 177

5.11 For each transmit beam and receive beam, shown is the fraction of receive beams and transmit beams, respectively, that meet various INR thresholds. 181

5.12 For each transmit beam and receive beam, shown is the measured INR across all receive and transmit beams, respectively. The red circle indicates the (a) receive direction and (b) transmit direction. 183

5.13 The transmit and receive steering neighborhoods $\mathcal{T}_{\text{tx}}^{(i)}$ and $\mathcal{T}_{\text{rx}}^{(j)}$, where the filled circles indicate the nominal steering directions $\left(\theta_{\text{tx}}^{(i)}, \phi_{\text{tx}}^{(i)}\right)$ and $\left(\theta_{\text{rx}}^{(j)}, \phi_{\text{rx}}^{(j)}\right)$ and the unfilled circles indicate directions comprising the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding each. In this case, the 25 directions in each yields $25^2 = 625$ transmit-receive steering combinations. 186

5.14 The number of transmit-receive beam pairs for a typical $(\Delta\theta, \Delta\phi)$-neighborhood under our $(1°, 1°)$ resolution. 188

5.15 (a) The CDF of $\mathsf{INR}_{ij}^{\text{rng}}$ across all nearly 6.5 million transmit-receive beam pairs for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. Fitted distributions for each are shown as dashed lines. (b) The fitted parameters $\alpha_{\text{rng}}(\Delta\theta, \Delta\phi)$ and $\beta_{\text{rng}}(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$. 190

5.16 For each transmit beam and receive beam, shown are the fraction of receive beams and transmit beams, respectively, that can offer an INR of 0 dB or less when allowed to deviate by various $(\Delta\theta, \Delta\phi)$. Low INR becomes reliably within arm's reach by increasing $(\Delta\theta, \Delta\phi)$. 192

5.17 (a) The CDF of $\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)$ for various $(\Delta\theta, \Delta\phi)$. For each, the fitted distribution is shown as a dashed line. (b) The fitted parameters $\mu_{\min}(\Delta\theta, \Delta\phi)$ and $\sigma_{\min}^2(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$. 193

5.18 The CDF of $\left\{\Delta\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\}$ for various INR, where $(\Delta\theta, \Delta\phi) = (2°, 2°)$. Their fitted counterparts are shown as dashed lines. 196

5.19 (a) The CDF of $\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)$ for various $(\Delta\theta, \Delta\phi)$. (b) The fitted parameters $\mu_{\max}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\max}(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$. 199

5.20 The CDF of $\left\{\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\}$ for various INR, where $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$. Their fitted counterparts are shown as dashed lines. Beam pairs with inherently low INR are typically at a greater risk of large increases in INR caused by small shifts in the transmit and receive beams. 201

6.1 A block diagram of our statistical model of self-interference. Based on system parameters, model parameters, and transmit and receive beams, our model produces realizations of self-interference that are statistically and spatially aligned with actual measurements. 219

6.2 (a) Three spatial neighborhoods whose measured self-interference follows a log-normal distribution, each with a unique fitted mean and variance. (b) The CDFs of the K-S statistic κ_{ij} across all nearly 6.5 million beam pairs (i, j) for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. The vast majority of measured beam pairs exhibit levels of self-interference that are log-normally distributed across the surrounding $(2^\circ, 2^\circ)$ spatial neighborhood. 223

6.3 (a) The measured azimuth cut of self-interference, where $\phi_{\mathrm{tx}} = \phi_{\mathrm{rx}} = 0^\circ$. (b) The simulated counterpart of (a) based on the commonly-used spherical-wave near-field channel model [2]. The stark difference between the two motivates the need for a new model of real-world self-interference, which we present herein. 228

6.4 An illustration of coupling clusters comprising the coarse geometric model of the self-interference channel $\bar{\mathbf{H}}$ between the transmit and receive arrays. From our measurements, we observed the presence of four dominant coupling clusters. This clustered coupling is a drastic deviation from the spherical-wave model [2] that has been widely used in the literature. 229

6.5 (a) The true fitted mean μ_{ij} for each of the nearly 6.5 million transmit-receive beam pairs (i, j). (b) The estimated counterpart of (a), where $\hat{\mu}_{ij}$ for each beam pair has been estimated using (6.24) with fitted ξ and $\bar{G}^2$. The strong structural similarity between the two confirms that $\bar{\mathbf{H}}$ captures large-scale spatial trends observed in our measurements. 234

6.6 (a) The CDF of the true fitted mean μ_{ij} and the estimated mean $\hat{\mu}_{ij}$ for all nearly 6.5 million beam pairs (i, j); in other words, the CDFs of Figure 6.5a and Figure 6.5b. (b) The CDF of the error between μ_{ij} and $\hat{\mu}_{ij}$ for all nearly 6.5 million beam pairs (i, j); in other words, the CDF of the difference between Figure 6.5a and Figure 6.5b. 235

6.7 (a) Pairs of true fitted $(\mu_{ij}, \sigma^2_{ij})$ and the linear estimator (dashed line) used to produce $\bar{\sigma}^2_{ij} = \mu_{ij} \cdot \alpha + \beta$. (b) The CDF of the linear estimator error $\sigma^2_{ij} - \bar{\sigma}^2_{ij}$, along with the corresponding fitted normal distribution $\mathcal{N}(0, \nu^2)$. 238

6.8 CDFs of the true fitted variance σ^2_{ij} and estimated variance $\hat{\sigma}^2_{ij}$. 239

6.9 (a) Measured self-interference $\mathcal{M}$ for $K = 40$ random transmit beams and $K = 40$ random receive beams. (b) A realized counterpart $\tilde{\mathcal{M}}$ of (a) using our statistical model. (c) A realized counterpart of (a) using the spherical-wave model [2]. The structural similarity between (a) and (b) illustrates that our model captures the dominant spatial components of self-interference in our measurements, whereas (c) does not. 248

6.10 The CDFs of the measured and realized self-interference in Figure 6.9a and Figure 6.9b, aligned closely in distribution. Also shown is the CDF of Figure 6.9c based on the spherical-wave model [2], which aligns neither structurally nor in distribution with measurements. 249

6.11 The CDFs of the K-S statistic $\kappa\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$ between measured and realized self-interference, where each CDF is taken across 5000 random $\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$ for each K. As K increases, the measured and realized self-interference more closely align in distribution. 250

6.12 Evaluating a beamforming-based full-duplex solution STEER [3] using actual measurements, using our proposed model, and using the spherical-wave model [2]. The (a) CDF of self-interference and (b) normalized sum spectral efficiency when using STEER versus conventional beam selection. Our proposed model more closely aligns with measurements and therefore serves as a more reliable evaluation tool. 253

7.1 A full-duplex IAB node receives backhaul from a fiber-connected donor node while simultaneously transmitting access to a UE, giving rise to self-interference at the IAB node and cross-link interference at the UE. We refer to this as the DL-DL operating mode. We also consider the UL-UL mode where the UE transmits uplink access and the IAB node transmits backhaul. 262

7.2 A block diagram summarizing beam selection via STEER, which jointly selects transmit and receive beams to reduce self-interference at a full-duplex mmWave transceiver. Conventional beam selection chooses transmit and receive beams independently, ignoring self-interference. 267

7.3 The spatial neighborhoods surrounding a given transmit direction and receive direction (shown as filled circles). The size of the neighborhoods is dictated by $(\Delta\theta, \Delta\phi)$ and their resolution by $(\delta\theta, \delta\phi)$. $(\Delta\vartheta, \Delta\varphi)$ will be relevant in Section 7.6. 271

7.4 (a) Illustrating the reduction in measurement overhead necessitated by STEER when using Algorithm 7.1. 65% of beam pairs only require 20% of measurements to locate a beam pair that meets a target of $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = 0$ dB. (b) The empirical probability of the resulting $(\Delta\vartheta, \Delta\varphi)$ after executing STEER on the collected INR measurements for $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB. 289

7.5 The fraction $\gamma_{\mathrm{sum}}^{\mathrm{ours}}$ of the sum capacity $C_{\mathrm{sum}}^{\mathrm{cb}}$ achieved by STEER as a function of the design parameter $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ for various $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. Heuristically, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB proves to be broadly optimal or near-optimal. 292

7.6 (a) The fraction of the sum capacity γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$ for various $\mathsf{INR}_{\mathrm{tx}}$. (b) The fraction of the sum capacity γ_{sum} as a function of $\mathsf{INR}_{\mathrm{tx}}$ for various $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. $\star$ markers indicate intersections with TDD-PC. 293

7.7 (a) The fraction of the sum capacity γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$ for various $(\Delta\theta, \Delta\phi)$, where cross-link interference $\mathsf{INR}_{\mathrm{tx}} = 0$ dB. (b) The unnormalized counterpart of (a). Enlarging $(\Delta\theta, \Delta\phi)$ offers noteworthy spectral efficiency gains, especially at high SNR. 296

7.8 (a) The CDF of $\mathsf{INR}_{\mathrm{rx}}$ for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. (b) The CDF of the gap between $\mathsf{SINR}_{\mathrm{rx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. STEER reliably reduces $\mathsf{INR}_{\mathrm{rx}}$, as evident in (a), while maintaining high beamforming gain, shifting $\mathsf{SINR}_{\mathrm{rx}}$ closer to $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ (its upper bound), as shown in (b). 297

7.9 The fraction of the sum capacity achieved by (a) conventional beam selection and (b) STEER as a function of $\left(\overline{\mathsf{SNR}}_{\mathrm{rx}}, \overline{\mathsf{SNR}}_{\mathrm{tx}}\right)$, where $\mathsf{INR}_{\mathrm{tx}} = 0$ dB. With STEER, greater sum spectral efficiency is achieved broadly across SNRs, and the region that nets $\gamma \leq 0.5$ (where half-duplexing is preferred) is greatly reduced with STEER. 300

7.10 For various cross-link interference levels, the $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$-region where $\gamma_{\mathrm{sum}} \leq 0.5$ (achieved by equal TDD). This shaded region shrinks dramatically with STEER, outside of which full-duplexing is justified. STEER generally demands lower SNRs to offer gains over half-duplex and can tolerate higher cross-link interference 301

List of Acronyms

5G fifth-generation

6G sixth-generation

ADC analog-to-digital converter

AGC automatic gain control

AoA angle of arrival

AoD angle of departure

CDF cumulative density function

DAC digital-to-analog converter

dB decibel

dBm decibel-milliwatt

DFT discrete Fourier transform

EIRP effective isotropic radiated power

FCC Federal Communications Commission

FDD frequency-division duplexing

i.i.d. independently and identically distributed

IAB integrated access and backhaul

INR interference-to-noise ratio

K-S Kolmogorov-Smirnov

LMMSE linear minimum mean square error

LNA low-noise amplifier

LOS line-of-sight

LSB least significant bit

MIMO multiple-input multiple-output

mmWave millimeter wave

MSE mean square error

PA power amplifier

RF radio frequency

RSRP reference signal received power

SDMA space-division multiple access

SDR software-defined radio

SIC self-interference cancellation

SINR signal-to-interference-plus-noise ratio

SNR signal-to-noise ratio

SVD singular value decomposition

TDD time-division duplexing

UE user equipment

UPA uniform planar array

VGA variable-gain amplifier

XDD cross-division duplexing

Chapter 1

Introduction

Until now, wireless communication systems have almost exclusively operated in a *half-duplex* fashion, where transmission and reception of radio waves are separated—or *orthogonalized*—in the frequency domain or time domain [4]. Put simply, signals transmitted and received by a traditional half-duplex wireless device exist in different frequency bands or at different times, referred to as frequency-division duplexing (FDD) and time-division duplexing (TDD), respectively [4, 5]. This half-duplex operation is necessitated by the manifestation of *self-interference* when a device attempts to receive signals while simultaneously transmitting in the same spectrum [6]. Self-interference stems from the fact that a signal transmitted by a device will inadvertently leak or couple into its own receiver when the transmitter and receiver are not orthogonalized in either time or frequency. In most cases, self-interference is many orders of magnitude stronger than the relatively weak signal a device is interested in receiving—the *desired receive signal*—which has presumably propagated tens or hundreds of meters [6]. This makes it virtually impossible to accurately recover the desired receive signal without taking measures to mitigate the effects of self-interference [6, 7]. By receiving in neighboring frequency spectrum or on a separate time slot as its transmissions, a half-duplex

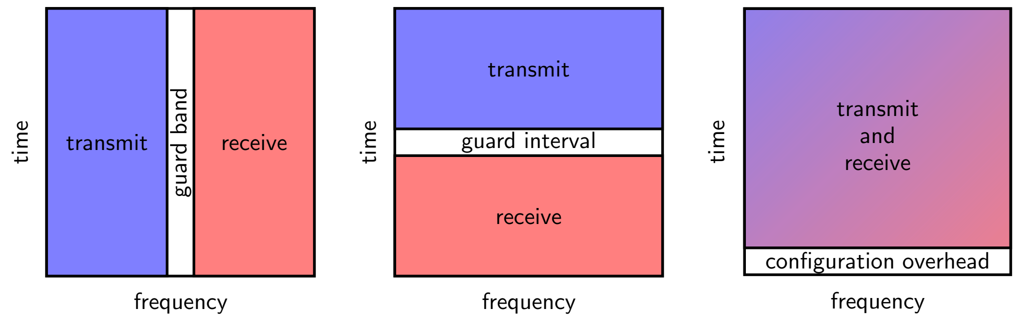

(a) FDD. (b) TDD. (c) Full-duplex.

Figure 1.1: FDD and TDD orthogonalize transmission and reception in the frequency and time domains, respectively. Full-duplex allows a device to transmit and receive on the same time-frequency resources. In practice, configuring a full-duplex solution incurs some overhead and FDD and TDD require a guard band/interval.

device avoids inflicting self-interference onto a desired receive signal, hence the long-standing use of FDD and TDD.

By their nature, FDD and TDD both consume radio resources by dedicating time-frequency resources to either transmission or reception, as illustrated in Figure 1.1. Of course, this would not be an issue if practical systems were not resource-constrained. In reality, all practical wireless communication systems operate on limited time-frequency resources. At the very least, most systems are confined to certain frequency spectrum by regulatory bodies, such as the Federal Communications Commission (FCC) in the United States. The consumption of radio resources by half-duplex operation has motivated researchers to explore upgrading devices with in-band *full-duplex*[1] ca-

[1]Throughout this dissertation, the term "full-duplex" is used to refer to in-band full-duplex operation, as opposed to out-of-band (or sub-band) full-duplex via FDD.

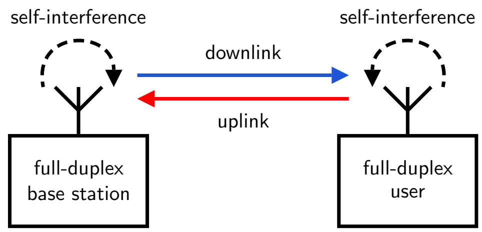

Figure 1.2: A full-duplex base station transmitting downlink and receiving uplink in-band with a full-duplex user. Self-interference manifests at both devices, requiring each to take measures to sufficiently mitigate such.

pability [6–18]. Starting in the late 2000s, researchers began heavily investigating and developing means to enable full-duplex operation. Since then, full-duplex has matured in the context of conventional, low-frequency communication systems, but it remains an active area of research in the context of higher-frequency communication systems [19–21], joint communication and sensing [22], spectrum sharing and cognitive radios [1, 23–25], and applied machine learning for wireless communications [26–29].

1.1 Enhancements Introduced by Full-Duplex

Full-duplex capability allows a transceiver to concurrently transmit and receive over the same frequency spectrum. In other words, transmission and reception can both make full use of the available frequency spectrum all the time, as illustrated in Figure 1.1c. As mentioned, when operating in this full-duplex fashion, self-interference is inflicted onto a desired receive signal. In pursuit of enabling full-duplex, researchers have therefore developed various ways to mitigate this self-interference using radio frequency (RF) components, analog

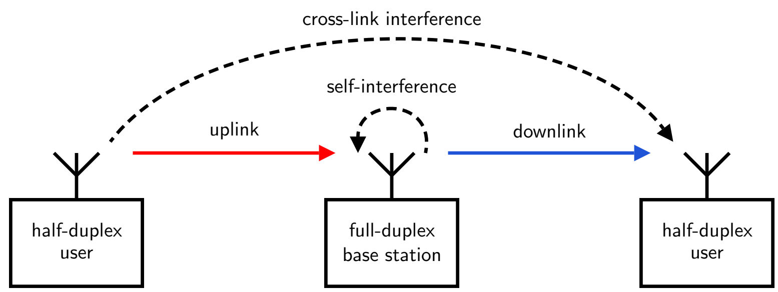

Figure 1.3: A full-duplex base station transmitting downlink to one half-duplex user while receiving uplink from another half-duplex user in-band. Self-interference manifests at the base station, whereas cross-link interference is inflicted onto the downlink user by the uplink user.

and digital filters, and other creative means [6, 7]. In general, most of these full-duplex solutions rely on the fact that a device can leverage knowledge of its own transmit signal to reconstruct and subsequently cancel self-interference. Assuming self-interference is sufficiently mitigated, a device could then reliably receive while transmitting in-band, enabling full-duplex operation and unlocking a number of enhancements. Methods to mitigate self-interference will be overviewed in more detail in Chapter 2.

As depicted in Figure 1.2, a full-duplex base station may transmit downlink and concurrently receive uplink with a neighboring full-duplex user. This makes better use of radio resources, and as intuition may suggest, leads to a potential doubling of spectral efficiency, as compared to half-duplex techniques. In other words, radio resources are being used twice as efficiently with full-duplex, since they are used for both transmission and reception, rather than divided as with FDD and TDD, as illustrated in Figure 1.1. Figure 1.3 depicts another full-duplex operating mode, where a full-duplex base station

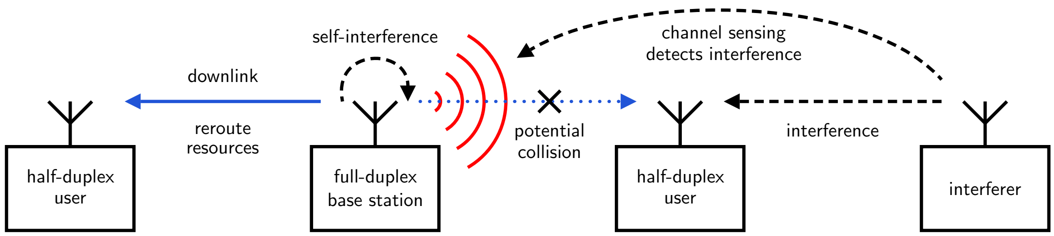

Figure 1.4: An access point equipped with full-duplex capability can sense the channel while transmitting, allowing it to reroute resources in the presence of interference that would otherwise cause collisions [1].

transmits downlink to a half-duplex user while receiving uplink from another half-duplex user. This mode of operation can also potentially double spectral efficiency, compared to traditional half-duplexing. It is important to note that, in this mode, cross-link interference is inflicted on the downlink user by the uplink user, the level of which depends on a number of factors, chiefly the users' locations and the environment.

Full-duplex capability can also enhance medium access and spectrum sharing. Half-duplex transceivers have been ubiquitous in wireless networks, and for good reason, communication standards and practices have been built on this half-duplex assumption. The ability to transmit and receive simultaneously and in-band is a transformative technology that can unlock new approaches to medium access [25, 30, 31] and spectrum sharing [1, 23, 24, 31] that are more efficient than those built on a half-duplex assumption. In turn, full-duplex can facilitate wireless networks that deliver higher throughput, inflict lower interference, and make better use of spectrum.

Upgrading transceivers with full-duplex capability can be felt at the

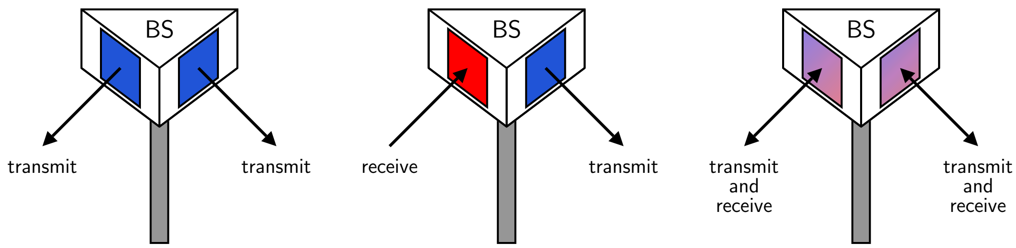

(a) Half-duplex. (b) Inter-sector full-duplex. (c) Intra-sector full-duplex.

Figure 1.5: (a) Without full-duplex capability, a sectorized base station must either transmit or receive across all sectors. (b) With inter-sector full-duplex, the base station is capable of transmitting or receiving with any of its sectors; self-interference is coupled across sectors. (c) With intra-sector full-duplex, the base station can transmit and receive within the same sector; self-interference is coupled within the same sector and potentially across sectors.

network level in a few ways. Even when only some devices are equipped with full-duplex—while the remainder are half-duplex-capable—a wireless network can enjoy improvements in throughput and latency [32, 33]. In fact, with full-duplex, network throughput can magnify beyond the doubling of spectral efficiency we are familiar with at the link level [32, 33]. This can be attributed to the fact that full-duplex can reduce multiplexing delays, reduce overhead associated with medium access control, and enable new ways to manage interference—all of which can improve network throughput.

As illustrated in Figure 1.5, upgrading a traditional *sectorized* base station with full-duplex can unlock scheduling constraints across its sectors. Today's sectorized base stations operate in a half-duplex fashion, either transmitting or receiving from all sectors simultaneously to avoid inflicting interference between sectors (i.e., self-interference). With *inter-sector* full-duplex capability, a base station could transmit or receive on any of its sectors at a given time, allowing it to more flexibly and thus more optimally schedule uplink and downlink across sectors. This is especially relevant to reduce the multiplexing delays associated with wireless backhauling and relaying. With the even more powerful *intra-sector* full-duplex, where each sector could transmit and receive simultaneously over the same spectrum, the more traditional gains of full-duplex would be enjoyed per sector. Beyond what was highlighted herein, there are applications of full-duplex technology in joint communication and sensing [22, 34], secure communications [35], military communications [36, 37], radar [22], and more [13, 14, 38].

1.2 Enabling Full-Duplex Millimeter Wave Systems

Wireless communication systems operating at millimeter wave (mmWave) carrier frequencies (roughly 30 GHz to 100 GHz) have begun playing a more prominent role in modern connectivity, such as in fifth-generation (5G) cellular systems and IEEE 802.11ay [39]. This shift to higher carrier frequencies—beyond traditional sub-6 GHz frequencies—is motivated by the opportunity to harness wide bandwidths that are available and meet the ever-growing demand for high-rate wireless access [40]. The key enabler to operating at mmWave frequencies is the use of dense antenna arrays—on the order of dozens or hundreds of antenna elements—to overcome high path loss and achieve sufficient link margin via highly directional *beamforming* [39, 40].

The focus of this dissertation is on upgrading mmWave communication

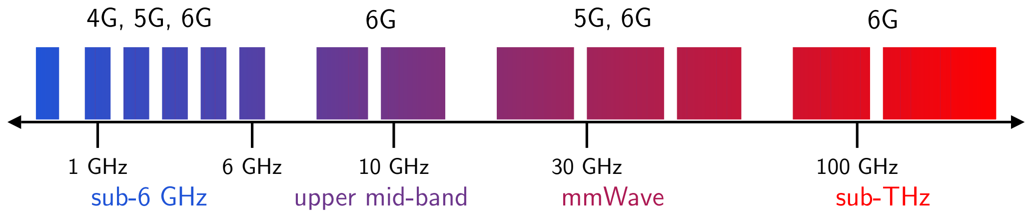

Figure 1.6: Traditionally, wireless networks like 4G cellular systems have made use of sub-6 GHz spectrum for reliable, widespread coverage. Modern 5G cellular systems have made use of higher, mmWave carrier frequencies to leverage wider bandwidths. Upper mid-band and sub-THz spectrum are being considered for future 6G cellular systems, largely for the same reason.

systems with full-duplex capability, typically in the context of next-generation cellular systems like 5G and future sixth-generation (6G) networks. As we will elaborate on in the next chapter, this dissertation aims to leverage beamforming with mmWave antenna arrays to mitigate self-interference *spatially* and enable full-duplex operation. This dissertation proposes solutions that accomplish this and also characterizes real-world mmWave self-interference itself through extensive measurement and modeling. With 6G cellular systems on the horizon, researchers are also actively considering the use of upper mid-band spectrum sandwiched between sub-6 GHz and mmWave, along with sub-terahertz (sub-THz) spectrum beyond mmWave, as illustrated in Figure 1.6. Although this dissertation focuses exclusively on mmWave systems, many of the proposed approaches and conclusions would likely extend to sub-THz and possibly upper mid-band communication systems.

There are a number of noteworthy challenges associated with creating beamforming solutions that mitigate self-interference and enable full-duplex mmWave systems. For instance, as discussed in the next chapter, beam alignment is a key component of current mmWave communication systems like 5G and IEEE 802.11ay and should therefore be accommodated by beamforming-based full-duplex solutions if possible. This is not a trivial task because it severely limits over-the-air channel knowledge. Also, transceiver-level factors—such as a limited receive dynamic range and digitally-controlled analog beamforming networks—can severely constrain the design of beamforming-based full-duplex solutions. This dissertation addresses and overcomes these key challenges, along with many others, in an effort to create full-duplex solutions that better translate to real-world mmWave communication systems.

1.3 Summary of Contributions and Organization

This dissertation aims to upgrade mmWave communication systems with full-duplex capability. In this pursuit, the work herein is particularly focused on addressing and incorporating noteworthy challenges and artifacts of real-world mmWave communication systems, such as those in 5G and future 6G cellular systems. Broadly speaking, the five principal contributions of dissertation span two high-level topics:

- (Contributions I, II, V) Creating full-duplex solutions that mitigate self-interference through beamforming while accounting for a variety of key practical considerations.
- (Contributions III, IV) Measurement and modeling of self-interference in real-world full-duplex mmWave communication systems.

Chapter 2 contains necessary and useful background on mmWave communication systems and on full-duplex and provides high-level motivation behind the scope of this dissertation. Chapters 3–7 contain the five principal contributions of this dissertation. Chapter 8 concludes this dissertation with summarizing remarks and directions for future work. The five principal contributions of this dissertation are summarized as follows.

I. (Chapter 3) A novel hybrid beamforming design for full-duplex mmWave systems that holistically accounts for a number of key practical considerations such as beam alignment, limited channel knowledge, a minimal number of RF chains, and a limited receive dynamic range.

II. (Chapter 4) The first beamforming codebook design specifically for full-duplex mmWave transceivers, called LoneSTAR. This is a novel enabler of full-duplex mmWave systems, allowing them to (i) close the link on the downlink and uplink via codebook-based beam alignment and (ii) simultaneously reduce self-interference through beamforming.

III. (Chapter 5) An extensive measurement campaign and statistical characterization of self-interference in a full-duplex mmWave system. Based on nearly 6.5 million measurements taken with 28 GHz phased arrays, this work sheds light on the degree of self-interference in real-world full-duplex mmWave systems, as well as its large-scale spatial characteristics and small-scale spatial variability.

IV. (Chapter 6) The first measurement-backed spatial and statistical model of self-interference in multi-panel, full-duplex mmWave transceivers. This model exhibits the large-scale and small-scale spatial characteristics of real-world self-interference observed in nearly 20 million collected measurements, allowing researchers to draw accurate realizations of self-interference and better evaluate full-duplex systems and solutions.

V. (Chapter 7) A measurement-driven beam selection methodology for full-duplex mmWave transceivers, called STEER, that delivers high beamforming gain while significantly reducing the self-interference coupled between transmit and receive beams. This novel solution allows a full-duplex mmWave transceiver to close the link with downlink and uplink users while simultaneously reducing self-interference—all without the need for over-the-air self-interference channel knowledge.

1.4 Related Publications

The content of this dissertation is based on several co-authored publications [3, 20, 38, 41–45], whose publishing agreements permit the inclusion of their content in this dissertation. The usage of each publication is detailed in the following and, for each, the contributions of each author are summarized.

- Chapter 1 and Chapter 2 include content from the following book chapter [38] (to be published) and magazine paper [20]. In [38], I. P. Roberts was the lead author, with H. A. Suraweera providing minor modifications and feedback on the outline and writing of the chapter. In [20], I. P. Roberts was the lead author, with J. G. Andrews, H. B. Jain, and S. Vishwanath providing feedback on the outline and writing of the article, along with minor modifications.

[38] I. P. Roberts and H. A. Suraweera, "Full-Duplex Transceivers for Next-Generation Wireless Communication Systems," in *Fundamentals of 6G Communications and Networking*, Y. Liu, J. Zhang, X. Lin, and J. Kim, Eds. Berlin: Springer Nature, 2023, (to be published).

[20] I. P. Roberts, J. G. Andrews, H. B. Jain, and S. Vishwanath, "Millimeter-Wave Full-Duplex Radios: New Challenges and Techniques," *IEEE Wireless Commun.*, vol. 28, no. 1, pp. 36–43, Feb. 2021.

- Chapter 3 includes content published in the following journal paper [41]. In [41], I. P. Roberts was the lead author, with J. G. Andrews and S. Vishwanath supervising the work and providing technical feedback.

[41] I. P. Roberts, J. G. Andrews, and S. Vishwanath, "Hybrid Beamforming for Millimeter Wave Full-Duplex under Limited Receive Dynamic Range," *IEEE Trans. Wireless Commun.*, vol. 20, no. 12, pp. 7758–7772, Dec. 2021.

- Chapter 4 includes content published in the following journal paper [42] and conference paper [43]. In [42, 43], I. P. Roberts was the lead author, with S. Vishwanath and J. G. Andrews supervising the work and providing technical feedback. H. B. Jain provided technical feedback in the development and motivation of [43].

[42] I. P. Roberts, S. Vishwanath, and J. G. Andrews, "LoneSTAR:

Analog Beamforming Codebooks for Full-Duplex Millimeter Wave Systems," *IEEE Trans. Wireless Commun.*, Jan. 2023, (early access).

[43] I. P. Roberts, H. B. Jain, S. Vishwanath, and J. G. Andrews, "Millimeter Wave Analog Beamforming Codebooks Robust to Self-Interference," *Proc. IEEE GLOBECOM*, Dec. 2021, pp. 1–6.

- Chapter 5 includes content published in the following journal paper [44]. In [44], I. P. Roberts processed and analyzed the measurements and prepared the paper; A. Chopra constructed the measurement platform and scripts; A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews offered feedback throughout the work and preparation of the paper.

[44] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "Beamformed Self-Interference Measurements at 28 GHz: Spatial Insights and Angular Spread," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9744–9760, Nov. 2022.

- Chapter 6 includes content in the following journal paper [45]. In [45], I. P. Roberts developed the proposed model and prepared the paper; A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews provided input on the motivation of the model, along with feedback on the paper.

[45] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "Spatial and Statistical Modeling of Multi-Panel Millimeter Wave Self-Interference," *IEEE J. Sel. Areas Commun.*, 2023, (to appear).

- Chapter 7 includes content published in the following journal paper [3]. In [3], I. P. Roberts developed the proposed beam selection methodology and was the lead author in preparing the paper; A. Chopra executed the measurements, which were then post-processed by I. P. Roberts to evaluate STEER; A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews provided feedback on the proposed methodology and paper.

  [3] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "STEER: Beam Selection for Full-Duplex Millimeter Wave Communication Systems," *IEEE Trans. Commun.*, vol. 70, no. 10, pp. 6902–6917, Oct. 2022.

1.5 Notation

The following notation and conventions are used throughout this dissertation. Standard roman fonts a, A are used to denote scalars. Bold lowercase $\mathbf{a}$ and bold uppercase $\mathbf{A}$ are used to denote column vectors and matrices, respectively. The element of a matrix $\mathbf{A}$ in the i-th row and j-th column is denoted $[\mathbf{A}]_{i,j}$; the entire i-th row or j-th column is denoted with $[\mathbf{A}]_{i,:}$ or $[\mathbf{A}]_{:,j}$, respectively. The vector of diagonal elements of a square matrix $\mathbf{A}$ is denoted $\mathrm{diag}(\mathbf{A})$. Sets are denoted with calligraphic fonts $\mathcal{A}$. The normal distribution having mean μ and variance σ^2 is denoted by $\mathcal{N}(\mu, \sigma^2)$. The complex normal distribution having mean μ and variance σ^2 is denoted by $\mathcal{N}_{\mathbb{C}}(\mu, \sigma^2)$. Directions in azimuth and elevation are often denoted θ and ϕ, respectively. Transmit and receive indexing are often denoted by i and j, respectively. A hat $\hat{a}$ is often used to denote the estimate of some ground truth a. All symbols have linear units by default; $[A]_{\mathrm{dB}} = 10 \cdot \log_{10}(A)$ is used to denote A in units of decibels (dB) and $[B]_{\mathrm{dBm}}$ to denote B in units of decibel-milliwatts (dBm).

Chapter 2

Background and Motivation

This chapter provides necessary background and motivation for the remainder of this dissertation. We begin by overviewing traditional full-duplex solutions that have proven to be effective for sub-6 GHz radios. Then, we provide brief, necessary background on mmWave communication systems. This chapter concludes by highlighting the challenges associated with translating low-frequency full-duplex solutions to higher-frequency, mmWave communication systems. We then call attention to the prospects of strategic beamforming to mitigate self-interference in full-duplex mmWave communication systems, motivating the principal contributions of this dissertation.

2.1 Traditional Low-Frequency Full-Duplex Solutions

Successfully equipping a device with full-duplex capability relies on mitigating—or *cancelling*—self-interference to levels that are sufficiently low [6, 7]. The amount of self-interference cancellation (SIC) needed depends on the particular application, but in most settings, full-duplex solutions aim to cancel self-interference to near or below the receiver noise floor. This ensures that the full-duplex resource gains are not eroded by the presence of high self-interference. In this section, we outline methods of SIC in multiple domains that have proven to be effective in low-frequency full-duplex radios.

Regardless of domain, the motivation behind SIC can largely be summarized as leveraging the fact that a transceiver has knowledge of its own transmitted signal and can therefore potentially reconstruct self-interference and subtract it from the received signal, leaving the desired portion virtually free of self-interference. In many cases, a staged approach to SIC is employed in low-frequency radios, where a portion of self-interference is cancelled using a tunable analog filter and a significant portion of the remaining *residual* self-interference is cancelled using digital filtering. This staged approach is depicted in Figure 2.1, where a full-duplex transceiver with separate antennas for transmission and reception employs an analog SIC filter between its antennas, along with a digital SIC filter.

2.1.1 Analog Self-Interference Cancellation

As illustrated in Figure 2.1, analog SIC typically exists as a digitally-controlled analog filter placed between the transmitter and receiver of a full-duplex transceiver. Analog SIC filters come in many forms, existing as tunable RF, intermediate frequency, baseband, and even optical filters [7, 10, 46–48]. Analog SIC is often driven by tapping off a small portion of the upconverted and amplified RF transmit signal, which then undergoes filtering within the analog SIC filter before being injected at the receiver. Ideally, the injected signal is an inverted reconstruction of self-interference, which, when combined

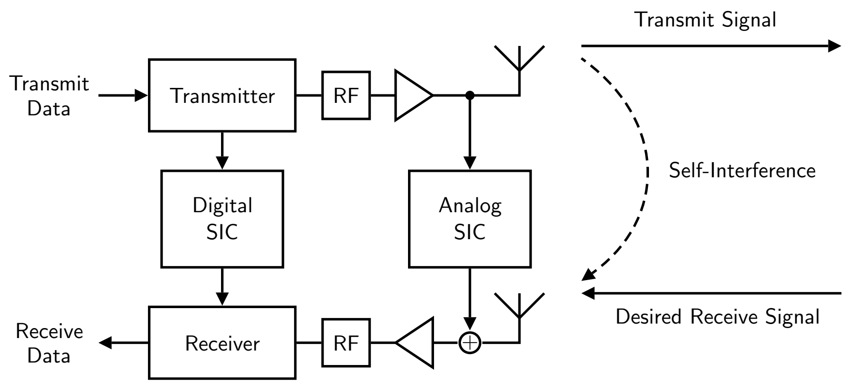

Figure 2.1: Self-interference manifests when a full-duplex transceiver attempts to simultaneously transmit and receive using the same frequency spectrum. Here, separate antennas are used for transmission and reception, with analog and digital SIC used to reconstruct and subsequently cancel self-interference incurred at the receiver.

with the received signal, destructively combines with self-interference. After this combining, there is practically some degree of residual self-interference due to imperfect reconstruction, which may stem from estimation errors, hardware limitations, and hardware imperfections [6, 7].

Tuning an analog SIC filter to effectively cancel self-interference largely consists of measuring self-interference and then configuring the filter to reconstruct its inverse. Analog SIC can be implemented as a time-domain filter or as a frequency-domain filter, meaning particular methods may vary but all tackle the same goal of reconstructing self-interference [7]. One method of time-domain analog SIC is to estimate the impulse response of the self-interference channel and then configure the analog filter to produce this (inverted) impulse response estimate, effectively equalizing self-interference. This estimation of the self-interference channel is typically executed by transmitting a pilot signal during a *quiet period*, when no desired receive signal is present. It is important to note that this consumes radio resources as configuration overhead, as illustrated in Figure 1.1c.

While it may seem fairly straightforward to implement analog SIC, it is practically quite challenging in most cases, especially outside of well-controlled lab settings. Most notably, there is small margin for error in SIC due to the often overwhelming strength of self-interference, reinforcing the need for extremely accurate, adaptable, and low-overhead self-interference reconstruction. Another challenging aspect is the need to miniaturize analog SIC filters into form-factors that integrate into devices such as cell phones, laptops, wireless access points, base stations, and the like [7]. This miniaturization is especially challenging in settings where the delay spread of self-interference has the potential to be high, since propagation delays may need to be physically realized within the analog SIC filter. Likewise, under rich frequency-selectivity, many taps may be needed to accurately reconstruct self-interference, complicating miniaturization.

2.1.2 Digital Self-Interference Cancellation

In addition to analog SIC, *digitally* cancelling self-interference is naturally an attractive option, thanks to incredible advances in digital signal processing capabilities in recent decades. The flexibility and sophistication of digital filtering can be leveraged to adaptively estimate and cancel self-

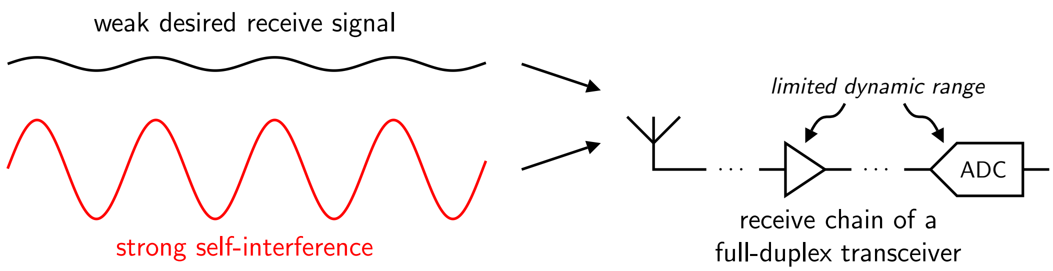

Figure 2.2: Self-interference combines with a desired receive signal at each receive antenna of a full-duplex transceiver. When self-interference is much stronger than the desired receive signal, it can saturate components that have a limited dynamic range, such LNAs and ADCs. This saturation should be prevented throughout the receive chain to ensure that digital SIC can be applied successfully and the desired signal can be received reliably.

interference in a similar vein as analog SIC, which has had impressive success [6, 7, 16–18, 49]. As depicted in Figure 2.1, digital SIC is executed after analog SIC and therefore aims to cancel residual self-interference that remains after prior efforts of cancellation. Naturally, one may ask whether digital SIC can cancel all self-interference, rendering analog SIC unnecessary. In general, this is not possible for a few reasons, stemming from hardware limitations and nonidealities, outlined as follows and illustrated in Figure 2.2 and Figure 2.3.

With reasonable resolution and appropriate gain control before analog-to-digital conversion, quantization noise is rarely a noteworthy issue in traditional half-duplex systems. In full-duplex systems, on the other hand, the strength of self-interference has the potential to *saturate* analog-to-digital converters (ADCs), even with ideal automatic gain control (AGC) and a reasonable number of bits [6, 17]. This is because the strength of quantization noise is dictated largely by that of self-interference since AGC acts on the combination of a desired signal plus self-interference at the ADC input. In such cases, only a portion of the ADC's total dynamic range is effectively used to quantize a desired signal. Consequently, even if self-interference could be completely reconstructed and cancelled digitally, its effects may remain in the form of increased quantization noise, which can severely and irreversibly degrade the quality of a desired receive signal. This reinforces the need to sufficiently cancel self-interference before it reaches the ADC input, which is often most reliably done via analog SIC in low-frequency radios.

Similarly, low-noise amplifiers (LNAs) in the receive chain of a full-duplex transceiver are subject to saturation due to self-interference [17]. Like ADCs, a practical LNA has a limited dynamic range due to *gain compression*, where, beyond a certain input power level, amplifier gain is no longer constant, as illustrated in Figure 2.3. Desired receive signals are typically well below this input power level, but self-interference can exceed this threshold since it can often be many orders of magnitude stronger than a desired receive signal. This saturation of LNAs can distort the received signal, making it difficult to apply digital SIC and reliably receive the desired receive signal. Like LNAs, other receive chain components may have a limited dynamic range that can be exceeded by self-interference, leading to undesirable saturation. Overall, this limited receive dynamic range is a fundamental motivator for analog SIC to ensure the strength of self-interference is reduced to maintain high linearity throughout the receive chain of a full-duplex transceiver.

Another challenge with relying solely on digital SIC is that practi-

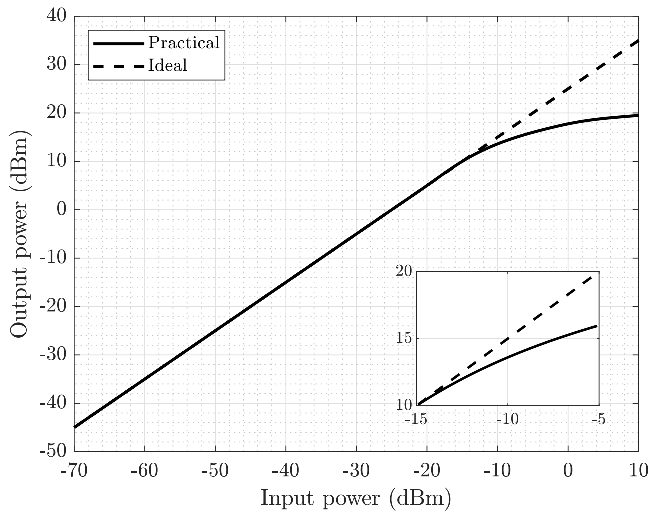

Figure 2.3: The transfer function of a practical LNA exhibiting gain compression. Amplifier gain is constant for low-power input signals, but beyond a certain input power, the amplifier is unable to maintain this gain.

cal transceivers introduce nonidealities in the transmit and receive chains, such as amplifier nonlinearities, I/Q imbalance, transmitter thermal noise, and phase noise. The digital domain does not have knowledge of these imperfections, since they manifest after digital-to-analog converters (DACs) and before ADCs [18]. When these nonidealities are not negligible, this requires digital SIC to accurately estimate and subsequently cancel them, which can be computationally complex. To reduce this burden, analog SIC can be well-positioned to cancel transmit-side impairments using the RF transmit signal as input to its cancellation filter, which inherently will include the nonidealities introduced by the transmit chain. Then, with successful cancellation, a large portion of self-interference will be mitigated before receiver-side impairments incur.

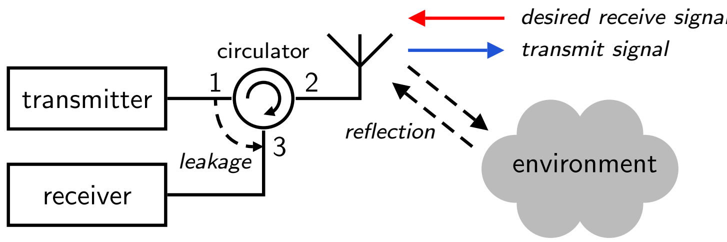

Figure 2.4: A circulator can be used to establish isolation between a transmitter and receiver sharing an antenna. Leakage through the circulator and reflections off the environment give rise to self-interference, however.

Nonetheless, classical digital signal processing approaches have been able to effectively estimate and account for transceiver impairments when reconstructing and subsequently cancelling self-interference [16]. This is done by modeling transceiver impairments with parameterized models, but the estimation of these parameters can be computationally expensive depending on transceiver nonlinearity and the SIC requirements. In addition to classical signal processing approaches, digital SIC solutions based on machine learning have been gaining traction and have shown impressive results [26–29, 50–52]. Machine learning has shown to offer comparable performance as classical digital SIC methods in capturing transceiver impairments when reconstructing self-interference but with reduced complexity, even in experimental demonstrations [26, 27, 50].

2.1.3 Circulators and Antenna Isolation

An RF component known as a *circulator* has been used in many monostatic radar and communication applications as a duplexer when a single antenna is used for simultaneous transmission and reception of RF signals [6,7]. In its simplest form, a circulator is a three-port, passive device where a signal entering a given port is "circulated" to the next port in the rotation. An example of this device being used with a single antenna shared by transmission and reception can be seen in Figure 2.4. Transmit signals enter port 1 of the circulator and exit at port 2, where they are radiated by the antenna. Signals received by the antenna enter port 2 and are circulated to port 3, where they exit the circulator and enter the receive chain. This establishes isolation between the transmitter and receiver of a full-duplex device using a shared antenna for transmission and reception.

One may reason that with an ideal circulator and with two radios operating in free space, full-duplex operation is trivial since perfect isolation is achieved between a radio's transmit signal and a desired receive signal. In reality, a circulator offers limited RF isolation between its ports, which introduces self-interference at the receiver. This is due to a number of factors, most notably the leakage between ports, reflections caused by imperfect matching at the antenna, and reflections off the environment. Circulators and other duplexers with small form-factors that offer high isolation for full-duplex are an active area of research with immense potential [53,54]. Nonetheless, analog and digital SIC can be used in conjunction with a circulator for single-antenna full-duplex transceivers. In such cases, SIC aims to cancel circulator leakage, as well as reflections off the environment and from the antenna.

In addition to circulators, increasing the inherent isolation between transmit and receive antennas through other means is an active area of research [7]. With separate transmit and receive antennas—as seen in Figure 2.1—there is the potential to position/orient the antennas to reduce their coupling, perhaps by exploiting polarization [9,55]. In addition, if space permits, placing absorbent material between separate transmit and receive antennas is also possible [7,14]. These passive methods of mitigating self-interference typically offer limited isolation and require additional measures of cancellation to enable full-duplex, especially when reflections off the environment are substantial.

2.2 Millimeter Wave Communication Systems

As mentioned in Chapter 1, modern wireless communication systems, such as 5G cellular, have turned to mmWave carrier frequencies (roughly 30 GHz to 100 GHz) to deliver high-rate connectivity by exploiting wide swaths of spectrum [40]. This shift from traditional sub-6 GHz carrier frequencies to higher mmWave frequencies involves noteworthy transceiver- and system-level challenges [39]. Most notably, mmWave signals suffer from more severe path loss than sub-6 GHz spectrum; recall, the Friis formula states that path loss increases quadratically with frequency. To overcome this path loss, mmWave devices employ antenna arrays with dozens or even hundreds of individual antenna elements. These dense antenna arrays can form highly directional beams that focus energy in particular directions, increasing received signal power to compensate for high path loss. These beams can electronically steered and shaped through *beamforming* by applying a complex weight to the

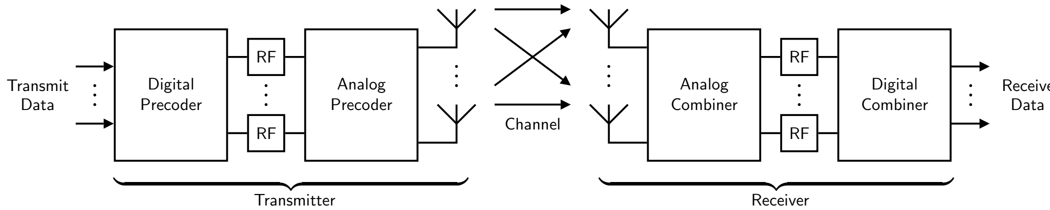

Figure 2.5: A mmWave communication system where both the transmitter and receiver employ hybrid digital/analog beamforming. This staged beamforming architecture allows each to control dense antenna arrays with a reduced number of RF chains.

signal transmitted/received at each antenna.

2.2.1 Hybrid Beamforming and Analog Beamforming

In traditional, lower-frequency multi-antenna communication systems, it has been common to use fully-digital beamforming. This requires a dedicated RF chain and DAC/ADC at *each antenna*. Scaling this to the dozens or hundreds of antennas at mmWave is typically not practically viable. Instead, a so-called *hybrid beamforming* architecture has often been adopted for mmWave transceivers, as illustrated in Figure 2.5 [39,56]. Through the combination of analog and digital beamforming, mmWave transcievers with hybrid beamforming can deliver high beamforming gain while also multiplexing more than one spatial stream—all with a reduced number of RF chains [56].

Digital beamforming takes place at baseband in software or digital logic, whereas analog beamforming is implemented at RF using a network of phase shifters, which are digitally controlled in practice. In addition to phase control, some analog beamforming networks also offer amplitude control through digitally-controlled attenuators or variable-gain amplifiers (VGAs) [57–59]. Given that phase and amplitude control are digitally controlled with finite resolution, there is a finite set of analog beamformers which can be realized physically, and the non-convexity of this discrete set has proven to complicate beamforming optimization in mmWave literature [39, 56]. In architectures without digital beamforming, such as when only one RF chain is present, *analog beamforming* can be used to deliver high gain but multiple spatial streams cannot be multiplexed.

2.2.2 Beam Alignment

Codebook-based analog beamforming is a critical component of today's mmWave communication systems [39, 60]. Rather than measure a high-dimensional multiple-input multiple-output (MIMO) channel and subsequently configure an analog beamforming network, modern mmWave systems instead rely on *beam alignment* (or *beam management*) procedures to identify promising beamforming directions, typically via exploration of a *codebook* of candidate beams (a set of beams), as illustrated in Figure 2.6 [39, 60, 61]. This offers a simple and robust way to configure an analog beamforming network without downlink/uplink MIMO channel knowledge *a priori*, which is not obtainable or is very resource-expensive to obtain in practice.

Let $\bar{\mathcal{F}}$ be a beamforming codebook of beams used by a mmWave base station to serve its users. Executing codebook-based beam alignment would aim to solve (or approximately solve) the following problem or one taking a

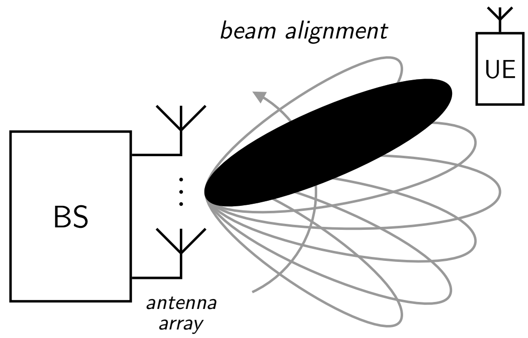

Figure 2.6: A mmWave base station conducts beam alignment to close the link between it and the user it intends to serve. This is typically achieved by sweeping a set of beams (a codebook), measuring the received signal quality at the user for each beam, and then selecting the beam that maximized such.

similar form.

$$\mathbf{f}^{\star} = \underset{\mathbf{f} \in \bar{\mathcal{F}}}{\operatorname{argmax}} \ \mathsf{SNR}(\mathbf{f}) \tag{2.1}$$

Practically, this problem is solved through a series of over-the-air downlink measurements, where the base station searches through its set of beams $\bar{\mathcal{F}}$ to find the beam $\mathbf{f}^{\star}$ which maximizes the received signal-to-noise ratio (SNR) at the user, for instance. As mentioned, today's real-world mmWave systems rely on beam alignment to close the link between mmWave devices. Therefore, as highlighted throughout this dissertation, it is imperative that beamforming-based full-duplex solutions accommodate beam alignment in order to more readily integrate into mmWave networks. This is a trademark of the solutions presented in this dissertation.

2.3 Full-Duplex Solutions at Millimeter Wave

Having outlined traditional full-duplex solutions for lower-frequency communication systems along with background on mmWave systems, we now overview potential solutions for full-duplex mmWave systems. We first consider the prospects of extending analog and digital SIC techniques to full-duplex mmWave transceivers. Then, we motivate the potential to mitigate self-interference purely through strategic beamforming at the transmitter and receiver of a full-duplex mmWave device.

2.3.1 Extending Analog and Digital SIC

When considering solutions for full-duplex mmWave transceivers, it is natural to first consider extending existing solutions, namely analog and digital SIC. To consider analog SIC as a potential full-duplex solution for mmWave systems, it is important to examine the placement of mmWave transceiver components, chiefly power amplifiers (PAs), LNAs, and ADCs. It is often the case that PAs and LNAs are per-antenna in mmWave transceivers, whereas ADCs are per-RF chain. For an analog SIC solution to capture nonlinear terms introduced by the transmit PAs—which have shown to be a bottleneck in sub-6 GHz full-duplex systems [17]—it will need to grow in one dimension with the number of transmit antennas. Likewise, to prevent LNA saturation, analog SIC will need to grow in its second dimension with the number of receive antennas. Its third dimension, the delay dimension (number of taps), will be based on the impulse response of the self-interference channel. Considering that mmWave systems will operate using dozens or hundreds of antennas over wide bandwidths, analog SIC solutions would presumably be relatively large in all three dimensions at mmWave and, as a result, likely prohibitive in size and complexity. The obstacles faced by extending analog SIC to mmWave systems motivate new approaches that address receiver-side saturation during full-duplex operation.

Digital SIC aims to mitigate residual self-interference—often both linear and significant nonlinear terms—after analog SIC. Fortunately, digital SIC remains a promising candidate for mmWave full-duplex since the number of RF chains remains low with hybrid beamforming, though wideband frequency-selectivity and exaggerated impairments in mmWave transceivers may drive up computational costs if not dealt with beforehand. A deeper investigation into digital SIC for mmWave full-duplex would likely yield many useful insights, particularly on transceiver nonlinearity and other impairments such as I/Q imbalance and phase noise, making it a good topic for future work. Ultimately, analog and digital SIC have proven to be effective in demonstrating full-duplex operation but their success has been limited in practice. This is in part due to the fact that analog SIC inherently necessitates hardware modifications and that the signal processing associated with digital SIC can be computationally complex. This motivates this dissertation's exploration of full-duplex solutions for mmWave transceivers that rely solely on strategic beamforming to mitigate self-interference spatially.

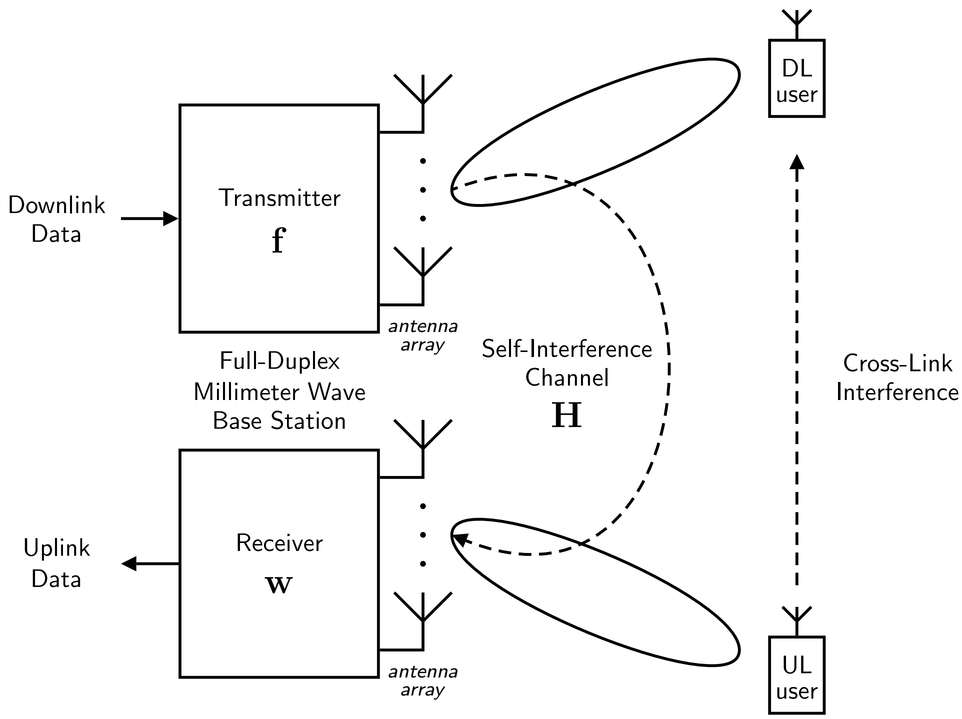

Figure 2.7: A full-duplex mmWave base station transmits downlink to a user with beamforming weights $\mathbf{f}$ while receiving from an uplink user with beamforming weights $\mathbf{w}$. In doing so, the transmit and receive beams couple over the self-interference MIMO channel $\mathbf{H}$ and the uplink user inflicts cross-link interference on the downlink user.

2.3.2 Tackling Self-Interference with Beamforming

As shown in Figure 2.7, a full-duplex mmWave base station equipped with separate, independently-controlled transmit and receive antenna arrays juggles a transmit beam $\mathbf{f}$ and a receive beam $\mathbf{w}$ to serve a downlink user and an uplink user simultaneously over the same frequency spectrum. In doing so, a self-interference MIMO channel $\mathbf{H}$ manifests between the transmit and receive arrays at the full-duplex base station. The degree of self-interference coupled depends on the coupling of the transmit beam $\mathbf{f}$ and the receive beam $\mathbf{w}$ with the self-interference channel $\mathbf{H}$. With dozens or even hundreds of antennas at the full-duplex base station, there is exciting potential to mitigate self-interference purely through strategic design of its transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$. This concept of spatial SIC through MIMO techniques was explored fairly extensively in lower-frequency radios, including in massive MIMO systems with many antennas [12, 62–67]. This dissertation focuses specifically on full-duplex mmWave systems, which are subject to unique transceiver- and network-level challenges and considerations, such as hybrid beamforming, beam alignment, and limited channel knowledge.

At mmWave, the prospect of such spatial SIC notably improves for a few reasons [19,20,68]. First of all, path loss and blockage increases at mmWave frequencies, compared to sub-6 GHz frequencies, which presumably weakens self-interference as it couples from a transmitter to receiver, both directly over-the-air and due to reflections off the environment; reflectivity increases at mmWave, however, which may lead to more self-interference from the environment. Second, with denser antenna arrays, the degrees-of-freedom available to cancel self-interference increases, compared to traditional sub-6 GHz MIMO systems. It is important to note that, as the number of antennas has increased immensely at mmWave, the number of data streams communicated has remained of similar order. Hybrid and purely-analog beamforming architectures, which are ubiquitous thus far in today's mmWave communication systems, fundamentally limit the number of data streams to the number of DACs and ADCs, which is comparable to that in conventional sub-6 GHz MIMO systems. With one or few data streams communicated and tens or hundreds of antennas, many degrees-of-freedom are potentially available for cancelling self-interference via strategic design of transmit and receive beamformers at a full-duplex mmWave transceiver.

Excitedly, if beamforming can cancel self-interference sufficiently, the need for analog and/or digital SIC vanishes, meaning there is the potential for full-duplex mmWave systems to operate without any additional hardware or complex signal processing. In other words, analog and digital SIC may not be needed for full-duplex mmWave systems with sufficient spatial SIC. Plenty of existing work has highlighted this by designing transmit and receive beamformers in such a way that self-interference is mitigated while maintaining downlink and uplink to users [19,20,41,43,63,67–78]. Interestingly, unlike traditional analog and digital SIC solutions, transmit beamforming introduces the unique opportunity to reduce the degree of self-interference that ever reaches the receive antennas. Like analog SIC, transmit and receive beamforming can be used to prevent self-interference from saturating receive chain components. We demonstrate this in Chapter 3 with a hybrid beamforming design that guarantees self-interference is below some power level at each LNA and each ADC at the receiver of the full-duplex mmWave device, ensuring they do not saturate. In addition to using beamforming alone to mitigate self-interference, researchers have also considered analog and/or digital SIC in conjunction, which relaxes the cancellation requirements of beamforming, allowing it to better serve uplink and downlink at the cost of added hardware or signal processing [55, 62, 78–83].

The remainder of this dissertation focuses on full-duplex mmWave systems that rely on beamforming to mitigate self-interference. The proposed solutions in Chapter 3, Chapter 4, and Chapter 7 incorporate a variety of key practical considerations and vary in their assumptions, leading to different solutions all with the same goal of enabling full-duplex mmWave systems. Chapter 5 and Chapter 6 extensively measure and model self-interference using 28 GHz phased arrays to better characterize the underlying coupling between the transmit and receive beams of a full-duplex mmWave system, with the hope that this leads to creating real-world beamforming-based solutions.

Chapter 3

Hybrid Beamforming for Full-Duplex Millimeter Wave Communication Systems

In the previous chapter, we introduced the potential for beamforming to cancel self-interference in full-duplex mmWave systems. To accomplish this, this chapter presents a hybrid beamforming design for full-duplex mmWave communication systems. Our design holistically accounts for the practical constraints of analog beamforming codebooks, a minimal number of RF chains, limited channel knowledge, beam alignment, and a limited receive dynamic range. To prevent self-interference from saturating receive components, such as LNAs and ADCs, our design framework limits the degree of self-interference on a per-antenna and per-RF chain basis.

3.1 Introduction

The ability for a transceiver to transmit and receive simultaneously in-band introduces exciting upgrades at the physical layer and in medium access when compared to existing half-duplex schemes such as TDD and FDD [84]. The gains supplied by full-duplex capability in mmWave systems are particularly attractive [19, 20], beyond the usual gains in spectral efficiency and latency. By full-duplexing access and backhaul, heterogeneous mmWave networks can be deployed with lower latency, higher spectral efficiency, and a reduced number of fiber drops. Key challenges in mmWave systems, such as beam alignment and beam tracking, could be transformed when devices can transmit and receive simultaneously, especially in highly dynamic applications like vehicular communication. The presence of communication, radar, and other incumbents in lightly regulated mmWave spectrum highlights the potential of novel strategies for medium access, in-band coexistence, and interference management via full-duplex.

3.1.1 Prior Work and Motivation

The majority of existing research on full-duplex has been in the context of lower carrier frequencies (e.g., sub-6 GHz). While many aspects of this existing work can be extended to mmWave, new approaches are necessary to enable full-duplex at mmWave [19, 20]. Dense antenna arrays and wide bandwidths are two key challenges for analog self-interference cancellation at mmWave in particular [20,39]. Furthermore, directly translating MIMO-based self-interference mitigation (e.g., [62, 67, 80]) to mmWave is complicated by hybrid digital/analog beamforming, propagation characteristics at mmWave, and system-level factors like beam alignment. While passive and polarization-based approaches have been proposed for mmWave [54,55,81], they are difficult to generalize to dense mmWave antenna arrays.

A number of recent bodies of work investigate methods where self-interference is mitigated by appropriately configuring the transmit and receive beamformers at a mmWave full-duplex device, sometimes termed *beamforming cancellation* [68–73, 77, 78, 85, 86]. Resembling MIMO-based approaches from sub-6 GHz full-duplex literature, existing beamforming cancellation designs suggest that mmWave full-duplex is theoretically possible without the hardware and computational costs associated with analog and digital self-interference cancellation. However, existing beamforming designs for mmWave full-duplex often fail to account for critical practical transceiver-level and system-level considerations. To start, real-world systems typically rely on codebook-based analog beamforming and beam alignment [39], meaning there is extremely limited freedom in the choice of analog beamformers. Designs such as those in [68–73] do not account for codebook-based analog beamforming and assume the ability to dynamically fine-tune each phase shifter in analog beamforming networks. Moreover, [68–73] assume infinite-precision phase shifters; in reality, phase shifters are almost certainly configured digitally, subjecting them to some degree of phase resolution. Almost all designs assume a lack of amplitude control in analog beamforming even though it is not uncommon to have both phase and amplitude control in practice. Those in [68–73, 85] do not account for beam alignment and assume full over-the-air channel knowledge, which is highly unlikely to have in practice.

Several designs [68–70, 73] involve analog-only beamforming, meaning they only support single-stream communication, which simplifies the design of beamforming-based self-interference mitigation. This is especially true in [68–

73] where the designs may be highly dependent on near-field self-interference channel conditions and are not shown to be robust against such. Some designs, such as those in [71, 78, 85, 86], take advantage of an increased number of RF chains that enable them to exploit the consequent dimensionality to mitigate self-interference in the digital domain. This is a strong assumption since the minimal number of RF chains necessary in hybrid beamforming is equal to the number of streams; increasing beyond this is undesirable in terms of financial cost and power consumption.

Finally, and perhaps most pertinent to this work, the majority of existing designs neglect the limited dynamic range of practical receivers [17, 87]. This is particularly important for full-duplex transceivers since self-interference—which is likely many orders of magnitude stronger than a desired receive signal—can overwhelm and saturate components of a receive chain if not sufficiently mitigated (hence, a limited dynamic range) [17, 66]. The work in [68–71, 73] accounts for ADC saturation by *completely* mitigating self-interference beforehand—which is not always possible—but do not account for other sources of saturation such as LNAs. In [78], the need to prevent ADC saturation is mentioned but is assumed to be satisfied without any mathematical basis. In [72, 85, 86], the need to prevent receiver-side saturation is ignored.

3.1.2 Contributions

In this work, we formulate mmWave MIMO expressions capturing practical receive dynamic range limitations *per-antenna* and *per-RF chain*. In particular, we motivate this work by the limited dynamic range of LNAs placed per-antenna and of ADCs placed per-RF chain, though it can be generalized to arbitrary dynamic range considerations. Using these formulations, we outline constraints to limit the self-interference power inflicted on each antenna and each RF chain at the receiver of the full-duplex device to prevent it from saturating. We outline conditions where meeting these per-antenna and per-RF chain self-interference power constraints are implicitly met, either by one another or other system factors.

Incorporating these constraints, we present a hybrid beamforming design that enables a mmWave transceiver to operate in a full-duplex fashion, serving two devices simultaneously in-band. Adhering to a multitude of practical considerations beyond a limited receive dynamic range, our design supports beam alignment schemes and codebook-based analog beamforming, rather than impractically assuming full knowledge of over-the-air channels and the ability to dynamically fine-tune phase shifters and attenuators. Furthermore, we limit the number of RF chains to the minimum necessary for multi-stream communication. To provide our design with freedom in the choice of its analog beamformers, we present a methodology for building sets of *candidate* analog beamformers based on measurements from codebook-based beam alignment. Our design is not self-interference channel model-dependent in that it does not exploit any particular structure.

Numerical results indicate scenarios where our design thrives, offering significant spectral efficiency gains over conventional half-duplex operation. These results also outline conditions under which per-antenna and per-RF chain self-interference power constraints restrict what is possible for mmWave full-duplex, providing useful insights to engineers on relationships between system parameters such as transmit power, RF isolation, ADC resolution, and the self-interference power reaching each antenna and each RF chain. Understanding the degree of self-interference mitigation required at specific points in the receiver can drive full-duplex system analyses, including those that may supplement beamforming-based approaches with analog and/or digital self-interference cancellation [78]. The goal of this work is not necessarily to produce a deployment-ready approach to mmWave full-duplex but rather to present a design methodology motivated by a select variety of practical considerations. We hope our approach and numerical results act as a step toward even more practically sound solutions and as a benchmark for systems constrained by a limited receive dynamic range.

3.2 System Model

This work considers the wireless system in Figure 3.1, where a mmWave transceiver i aims to transmit to a device j while receiving from a device k in the same band. Instead of turning to half-duplexing strategies like TDD or FDD to avoid self-interference, this work presents a design that enables

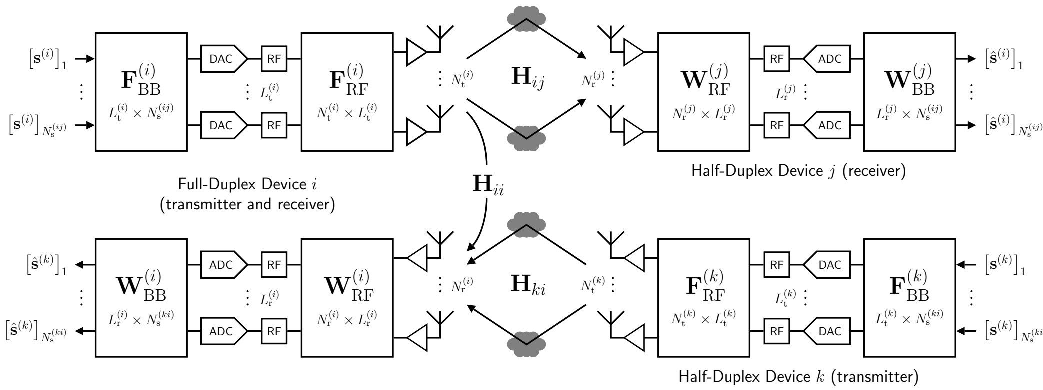

Figure 3.1: A full-duplex mmWave device i transmitting to a device j as it receives from a device k in-band.

in-band full-duplex operation by leveraging the spatial domain to mitigate self-interference. In this work, we consider the case where devices j and k are separate half-duplex devices, though many aspects of our contribution would extend naturally, or even simplify, when j and k are separate full-duplex devices or comprise a single full-duplex device.

Ubiquitous among practical mmWave transceivers is the use of hybrid digital/analog beamforming architectures where transmit precoding and receive combining are implemented by the combination of digital (baseband) and analog (RF) signal processing, as exhibited in Figure 3.1. We assume devices i, j, and k all employ hybrid beamforming in a fully-connected fashion where each antenna is connected to each RF chain via an analog beamforming network [39]. As illustrated in Figure 3.1, we assume that separate arrays are used at device i for transmission and reception and independent precoding and combining on the two is supported [20].

For devices $(m, n) \in \{(i, j), (k, i)\}$, we use the following notation. Let $N_{\mathrm{t}}^{(m)}$ and $N_{\mathrm{r}}^{(n)}$ be the number of transmit and receive antennas, respectively. Connecting the digital and analog stages, let $L_{\mathrm{t}}^{(m)}$ and $L_{\mathrm{r}}^{(n)}$ be the number of transmit and receive RF chains, respectively. Let $N_{\mathrm{s}}^{(mn)}$ be the number of symbol streams transmitted from device m to device n. Let $\mathbf{F}_{\mathrm{BB}}^{(m)} \in \mathbb{C}^{L_{\mathrm{t}}^{(m)} \times N_{\mathrm{s}}^{(mn)}}$ be the digital precoding matrix and $\mathbf{F}_{\mathrm{RF}}^{(m)} \in \mathbb{C}^{N_{\mathrm{t}}^{(m)} \times L_{\mathrm{t}}^{(m)}}$ be the analog precoding matrix, responsible for transmitting from m. Let $\mathbf{W}_{\mathrm{BB}}^{(n)} \in \mathbb{C}^{L_{\mathrm{r}}^{(n)} \times N_{\mathrm{s}}^{(mn)}}$ be the digital combining matrix and $\mathbf{W}_{\mathrm{RF}}^{(n)} \in \mathbb{C}^{N_{\mathrm{r}}^{(n)} \times L_{\mathrm{r}}^{(n)}}$ be the analog combining matrix, responsible for receiving at n.

For devices $(m, n) \in \{(i, j), (k, i)\}$, let $\mathbf{s}^{(m)} \in \mathbb{C}^{N_{\mathrm{s}}^{(mn)} \times 1}$ be the symbol vector transmitted by m intended for n, where the symbol covariance is

$$\mathbb{E}\left[\mathbf{s}^{(m)} \mathbf{s}^{(m)*}\right] = \frac{1}{N_{\mathrm{s}}^{(mn)}}\ \mathbf{I}. \tag{3.1}$$

We do not consider a specific signaling constellation, though we will evaluate our work assuming Gaussian signaling is employed. Let $\mathbf{n}^{(n)} \sim \mathcal{N}_{\mathbb{C}}(\mathbf{0}, \sigma_{\mathrm{n}}^2 \cdot \mathbf{I})$ be the $N_{\mathrm{r}}^{(n)} \times 1$ additive noise vector incurred at the receive array of n, where σ_{n}^2 represents a per-antenna noise power spectral density in watts/Hz and is assumed common across devices for simplicity. We denote the symbol period as T and the symbol bandwidth $B = T^{-1}$, which we assume to be constant across both links, being a full-duplex system. Let $\tilde{P}_{\mathrm{tx}}^{(m)}$ be the total transmit power of device m in watts and $P_{\mathrm{tx}}^{(m)}$ be the resulting transmit power in joules per symbol. We extend this convention, representing power quantities in watts using a tilde, $\tilde{P}$, and in joules per symbol without a tilde, P, which are linked via $P = \tilde{P} \cdot B^{-1} = \tilde{P} \cdot T$.

We impose the following digital precoding power constraint

$$\left\| \mathbf{F}_{\mathrm{BB}}^{(m)} \right\|_{\mathrm{F}}^{2} \leq 1 \tag{3.2}$$

and normalize the columns of $\mathbf{F}_{\mathrm{RF}}^{(m)}$ to have squared ℓ_2-norm $L_{\mathrm{t}}^{(m)}$ and of $\mathbf{W}_{\mathrm{RF}}^{(n)}$ to have squared ℓ_2-norm $N_{\mathrm{r}}^{(n)}$. Since this work is focused on receiver-side power levels, our analog combining normalization differs from the analog precoding one to ensure consistency with the physical combining taking place at receivers. As is common in practice, we will assume that columns of our analog precoders and combiners will come from analog beamforming codebooks that account for hardware constraints such as phase shifter resolution and amplitude control. That is, for $(m,n) \in \{(i,j),(k,i)\}$, we have

$$\left[\mathbf{F}_{\mathrm{RF}}^{(m)} \right]_{:,\ell} \in \mathcal{F}_{\mathrm{RF}}^{(m)}, \ \ell = 1, \ldots, L_{\mathrm{t}}^{(m)} \tag{3.3}$$

$$\left[\mathbf{W}_{\mathrm{RF}}^{(n)} \right]_{:,\ell} \in \mathcal{W}_{\mathrm{RF}}^{(n)}, \ \ell = 1, \ldots, L_{\mathrm{r}}^{(n)} \tag{3.4}$$

where $\mathcal{F}_{\mathrm{RF}}^{(m)}$ and $\mathcal{W}_{\mathrm{RF}}^{(n)}$ denote analog precoding and combining codebooks, respectively.

Now, let us consider $(m,n) \in \{(i,j),(k,i),(i,i)\}$. We assume that the large-scale power gain between devices m and n is given by G_{mn}^2. The $N_{\mathrm{r}}^{(n)} \times N_{\mathrm{t}}^{(m)}$ channel matrix between a transmitter m and receiver n is denoted $\mathbf{H}_{mn} \in \mathbb{C}^{N_{\mathrm{r}}^{(n)} \times N_{\mathrm{t}}^{(m)}}$. In this work, we consider the more straightforward case of frequency-flat MIMO channels, with frequency-selective ones necessary future work. Taking the perspective of our full-duplex device i, we term $\mathbf{H}_{ij}$ the *transmit* channel, $\mathbf{H}_{ki}$ the *receive* channel, and $\mathbf{H}_{ii}$ the *self-interference* channel. Note that we have not considered an inter-user interference channel between devices k and j since we assume that, with sufficient separation and/or user scheduling, the interference between the two to be negligible given the high path loss at mmWave and highly directional steering of energy that is typical.

We assume devices i and j as well as devices k and i are separated in a far-field fashion. As such, we model the transmit and receive channels as the composition of discrete rays with the Saleh-Valenzuela-based model [39]. Explicitly, channels $\mathbf{H}_{ij}$ and $\mathbf{H}_{ki}$ are of the form

$$\mathbf{H}_{mn} = \sqrt{\frac{1}{N_{\mathrm{rays}}^{(mn)}}} \sum_{u=1}^{N_{\mathrm{rays}}^{(mn)}} \beta_u \ \mathbf{a}_{\mathrm{rx}}^{(n)}(\phi_u) \ \mathbf{a}_{\mathrm{tx}}^{(m)}(\theta_u)^*. \tag{3.5}$$

where $(m, n) \in \{(i, j), (k, i)\}$. In each channel, $N_{\mathrm{rays}}^{(mn)}$ is a random variable dictating the number of rays in the channel. The complex gain of ray u is given as $\beta_u \sim \mathcal{N}_{\mathbb{C}}(0, 1)$. The u-th ray's angle of departure (AoD) and angle of arrival (AoA) are given as θ_u and ϕ_u, respectively. The transmit and receive array response vectors at these angles are given as $\mathbf{a}_{\mathrm{tx}}^{(m)}(\theta_u)$ and $\mathbf{a}_{\mathrm{rx}}^{(n)}(\phi_u)$, which have squared ℓ_2-norm $N_{\mathrm{t}}^{(m)}$ and $N_{\mathrm{r}}^{(n)}$, respectively. The coefficient in front of the summations handles a channel power normalization to ensure $\mathbb{E}\left[\|\mathbf{H}_{mn}\|_{\mathrm{F}}^2\right] = N_{\mathrm{t}}^{(m)} N_{\mathrm{r}}^{(n)}$.

A lack of measurements and characterization of self-interference at mmWave frequencies prevents us from confidently assuming a particular channel model for $\mathbf{H}_{ii}$ [20]. As such, our contribution herein does not rely on the self-interference channel's structure or properties. However, to evaluate our design, we employ a model that aims to capture the near-field nature of the transmit and receive arrays at i along with reflections that may stem from the environment [19, 72], which we explicitly state in Section 3.6. The large-scale power gain of the self-interference channel is represented by G_{ii}^2, which captures the RF isolation between the transmit and receive arrays at i. We define the SNR between two devices $(m, n) \in \{(i, j), (k, i)\}$ as

$$\mathsf{SNR}_{mn} \triangleq \frac{P_{\mathrm{tx}}^{(m)} \cdot G_{mn}^2}{\sigma_{\mathrm{n}}^2} = \frac{\tilde{P}_{\mathrm{tx}}^{(m)} \cdot G_{mn}^2}{\sigma_{\mathrm{n}}^2 \cdot B} \tag{3.6}$$

which captures the received power (without beamforming gains) versus the noise power.

3.3 Problem Formulation

This work is motivated by the fact that a receive chain of a full-duplex device—which practically has a limited dynamic range—is susceptible to saturation due to the overwhelming strength of self-interference [17, 87]. To highlight this, we consider two sources of limited receive dynamic range in this work: LNAs and ADCs. Like other amplifiers, LNAs begin saturating and introduce significant nonlinearities beyond a certain input power level, meaning only signals below some input power threshold see an approximately linear amplifier. The limited resolution of an ADC is a classical example of limited dynamic range; since the combination of a desired receive signal and self-interference enters the ADC, self-interference can drive up quantization noise and degrade the quality of the desired receive signal, even when digital cancellation is used [17]. If the LNAs, ADCs, and the rest of the receive chain had an infinite dynamic range and were fully linear, self-interference could be removed solely in the digital domain. This is not the case practically, meaning a portion of self-interference should be mitigated *via beamforming* before it can saturate the receiver and destroy any sense of linearity. If beamforming sufficiently mitigates self-interference to levels that prevent saturation, residual self-interference may then be cancelled digitally thanks to a well-preserved receive chain.

Consider Figure 3.1, where LNAs are placed per-antenna and ADCs are placed per-RF chain. To avoid saturating the LNAs at the full-duplex device, the self-interference power reaching each antenna must be mitigated to below some threshold. Similarly, to avoid saturating the ADCs, the self-interference reaching each RF chain must also be limited. This work investigates relying solely on beamforming to achieve this. Note that we are only concerned with preventing saturation at the receiver of the *full-duplex* device i—not that of the *half-duplex* device j—since the saturation we are considering stems from self-interference. There may exist other motivations (beyond LNAs and ADCs) for restricting the self-interference power at each antenna and at each RF chain. With this in mind, the design we present is not strictly for LNAs and ADCs but rather for meeting arbitrary per-antenna and per-RF chain self-interference power constraints; LNAs and ADCs are an important special case.

Using the previously defined system model, we can begin analyzing the signals that reach the LNAs and ADCs of our full-duplex device. The symbol vector at the input of the LNAs of i is

$$\mathbf{y}_{\mathrm{LNA}} = \mathbf{y}_{\mathrm{des,LNA}} + \mathbf{y}_{\mathrm{int,LNA}} + \mathbf{y}_{\mathrm{noise,LNA}} \in \mathbb{C}^{N_{\mathrm{r}}^{(i)} \times 1} \tag{3.7}$$

where the desired term is $\mathbf{y}_{\mathrm{des,LNA}} = \sqrt{P_{\mathrm{tx}}^{(k)}} G_{ki} \mathbf{H}_{ki} \mathbf{F}_{\mathrm{RF}}^{(k)} \mathbf{F}_{\mathrm{BB}}^{(k)} \mathbf{s}^{(k)}$, the self-interference term is $\mathbf{y}_{\mathrm{int,LNA}} = \sqrt{P_{\mathrm{tx}}^{(i)}} G_{ii} \mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{s}^{(i)}$, and the noise term is $\mathbf{y}_{\mathrm{noise,LNA}} = \mathbf{n}^{(i)}$. The signal at each antenna passes through its respective LNA, which we represent as $\mathcal{L}(\cdot)$ and model as $\mathcal{L}(x) = G_{\mathrm{rx}} \cdot x$ for $|x|^2 \leq P_{\mathrm{LNA}}^{\mathrm{max}}$, where the per-antenna symbol x undergoes a power gain of $G_{\mathrm{rx}}^2 > 0$ if its average power (over the symbol period) is below some threshold $P_{\mathrm{LNA}}^{\mathrm{max}}$. Otherwise, for $|x|^2 > P_{\mathrm{LNA}}^{\mathrm{max}}$, the LNA is saturated and x undergoes some nonlinear operation, often characterized by an established amplifier signal model (e.g., $P_{\mathrm{LNA}}^{\mathrm{max}}$ may be approximated by the P1dB point).

It is difficult to translate the effects of LNA saturation to the symbol level (the scope of this work) since it is inflicted instantaneously on time-domain signals. For this reason, we make no attempt to characterize LNA saturation with the understanding that linear LNA operation can be ensured by restricting the power of x up to $P_{\mathrm{LNA}}^{\mathrm{max}}$. We would like to point out that for a properly chosen $P_{\mathrm{LNA}}^{\mathrm{max}}$—which will likely include appropriate backoffs for the signal distribution and pulse shape—the *time-domain signal* will undergo linear amplification and, thus, so will the symbols. Since the gain of the LNA acts on signal-plus-noise, for conciseness, it can be abstracted out henceforth as $G_{\mathrm{rx}} = 1$ with appropriate scaling of the noise variance σ_{n}^2, which can simultaneously capture the noise figure of the LNA.

We overload the LNA transfer function $\mathcal{L}(\cdot)$ function to support vector input by the simple element-wise extension $[\mathcal{L}(\mathbf{x})]_\ell = \mathcal{L}(x_\ell)$, where $\mathbf{x} = [x_1, x_2, \ldots]^{\mathrm{T}}$. Following per-antenna LNAs, the signals are combined as

$$\mathbf{y}_{\mathrm{ADC}} = \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathcal{L}(\mathbf{y}_{\mathrm{LNA}}) \in \mathbb{C}^{L_{\mathrm{r}}^{(i)} \times 1} \tag{3.8}$$

where $\mathbf{y}_{\mathrm{ADC}}$ is the vector of per-RF chain symbols at the input of the ADCs. Under linear LNA operation, (3.8) can be written as

$$\begin{aligned} \mathbf{y}_{\mathrm{ADC}} &\overset{\mathrm{lin}}{=} \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{y}_{\mathrm{LNA}} && (3.9) \\ &= \mathbf{y}_{\mathrm{des,ADC}} + \mathbf{y}_{\mathrm{int,ADC}} + \mathbf{y}_{\mathrm{noise,ADC}} && (3.10) \end{aligned}$$

where $\mathbf{y}_{(\cdot),\mathrm{ADC}} = \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{y}_{(\cdot),\mathrm{LNA}}$. We use $\mathcal{Q}_b(\cdot)$ to represent a b-bit ADC, modeled as

$$\mathcal{Q}_b(x) = x + e_{\mathrm{quant}} \tag{3.11}$$

where x is the symbol reaching the ADC and e_{quant} is the error in perfectly digitizing x due to quantization noise, whose power under b-bit, uniform quantization can be approximated as [87]

$$\mathbb{E}\left[\left|e_{\mathrm{quant}}\right|^2\right] \approx \frac{|x|^2}{1.5 \cdot 2^{2b}}. \tag{3.12}$$

As the number of bits b increases, the magnitude of e_{quant} decreases. Similarly, as the magnitude of x increases, the quantization noise power increases. From this, one can see why mitigating self-interference before the ADC input is so important: increased self-interference can plague a desired receive signal with increased quantization noise. We overload $\mathcal{Q}_b(\cdot)$ to vectors as $[\mathcal{Q}_b(\mathbf{x})]_\ell = \mathcal{Q}_b(x_\ell)$, where $\mathbf{x} = [x_1, x_2, \ldots]^{\mathrm{T}}$. The symbol vector out of the ADCs is

$$\mathbf{y}_{\mathrm{dig}} = \mathcal{Q}_b(\mathbf{y}_{\mathrm{ADC}}) \in \mathbb{C}^{L_{\mathrm{r}}^{(i)} \times 1} \tag{3.13}$$

$$= \mathbf{y}_{\mathrm{des,ADC}} + \mathbf{y}_{\mathrm{int,ADC}} + \mathbf{y}_{\mathrm{noise,ADC}} + \mathbf{e}_{\mathrm{quant}}. \tag{3.14}$$

From this discussion, we can see that limiting the power into each of the LNAs and each of the ADCs is critical in preserving the linearity of the receive chain and reducing the effects of quantization. With these models and formulations in hand, we begin laying out our contribution, which aims to mitigate self-interference per-antenna and per-RF chain.

3.4 Beam Candidate Acquisition

Practical mmWave systems employ *beam alignment* schemes that aim to identify transmit-receive analog beamforming pairs that afford link margin sufficient for communication. Typically, these analog beamformers are chosen from a predetermined codebook of beams, which reduces complexity and offers robustness. Given that our full-duplex system entertains two independent links, beam alignment needs to be executed on the transmit link and on the receive link. In this section, we present a methodology for creating a set of beam *candidates* for transmission and reception on each link, offering our design in the next section greater freedom in its aim to reduce self-interference. This is motivated by the fact that some beam selections naturally afford more isolation at the full-duplex device than other pairs, meaning it may be preferable to use them for full-duplex, even if they are sub-optimal in a half-duplex sense.

We have not assumed knowledge of the over-the-air channels $\mathbf{H}_{ij}$ and $\mathbf{H}_{ki}$. Rather, we will be inspecting them using beam alignment, which we conduct in a half-duplex fashion on separate time-frequency resources on each link, meaning we avoid self-interference and any concern for saturation. For our design, we propose that candidate analog beamforming matrices be created as follows, which can be accommodated by existing beam alignment schemes. Let $\mathbf{F}_{\mathrm{tr}}^{(i)} \in \mathbb{C}^{N_{\mathrm{t}}^{(i)} \times M_{\mathrm{t}}^{(i)}}$ be a matrix whose $M_{\mathrm{t}}^{(i)}$ columns are *training* analog precoders used by i during candidate beam acquisition as it illuminates $\mathbf{H}_{ij}$. To observe these illuminations, let $\mathbf{W}_{\mathrm{tr}}^{(j)} \in \mathbb{C}^{N_{\mathrm{r}}^{(j)} \times M_{\mathrm{r}}^{(j)}}$ be a matrix whose $M_{\mathrm{r}}^{(j)}$ columns are *training* analog combiners used by j. We assume, for simplicity, that each of the $M_{\mathrm{t}}^{(i)}$ training precoders is observed by all $M_{\mathrm{r}}^{(j)}$ training combiners, though the ideas herein could be adapted when this is not the case. Analogously, let $\mathbf{F}_{\mathrm{tr}}^{(k)} \in \mathbb{C}^{N_{\mathrm{t}}^{(k)} \times M_{\mathrm{t}}^{(k)}}$ and $\mathbf{W}_{\mathrm{tr}}^{(i)} \in \mathbb{C}^{N_{\mathrm{r}}^{(i)} \times M_{\mathrm{r}}^{(i)}}$ be the training analog precoders and analog combiners used by k and i, respectively, used to inspect $\mathbf{H}_{ki}$. On both links, we assume the analog beamformers used during training come from their respective codebooks according to (3.3)–(3.4), though we do not suggest that it is necessary to sweep *all* beams (more beams will certainly help) nor do we specify particular beam codebooks. Instead, we present a methodology that can be tailored and applied to a variety of beam alignment schemes and codebooks. The collection of measurements for each link can be written in matrix form as

$$\mathbf{M}_{ij} = \sqrt{P_{\text{tx}}^{(i)}} G_{ij} \mathbf{W}_{\text{tr}}^{(j)*} \mathbf{H}_{ij} \mathbf{F}_{\text{tr}}^{(i)} \mathbf{I}_{M_{\text{t}}^{(i)}} \in \mathbb{C}^{M_{\text{r}}^{(j)} \times M_{\text{t}}^{(i)}} \tag{3.15}$$

$$\mathbf{M}_{ki} = \sqrt{P_{\text{tx}}^{(k)}} G_{ki} \mathbf{W}_{\text{tr}}^{(i)*} \mathbf{H}_{ki} \mathbf{F}_{\text{tr}}^{(k)} \mathbf{I}_{M_{\text{t}}^{(k)}} \in \mathbb{C}^{M_{\text{r}}^{(i)} \times M_{\text{t}}^{(k)}}. \tag{3.16}$$

Under channel reciprocity, $\mathbf{M}_{ij}$ may be measured in reverse from j to i to avoid feedback overhead. For consistency, we maintain notation as if measurements take place from i to j.

Given that the training precoders and combiners come from their respective codebooks, the strength of the measurements in $\mathbf{M}_{ij}$ and $\mathbf{M}_{ki}$ directly indicates which analog precoders (columns) and analog combiners (rows) are promising candidates on each link. Having multiple (say $L \geq 1$) RF chains, building a set of K candidates requires choosing K tuples of L strong beam pairs from the measurement matrix $\mathbf{M}$. In the particular case when $L = 1$ (i.e., analog-only beamforming), this process amounts to simply choosing the K strongest beam pairs from $\mathbf{M}$.

Let $\mathcal{T}_{ij}$ be a set of K_{ij} *candidate* analog precoding-combining pairs for communicating from i to j. Similarly, let $\mathcal{T}_{ki}$ be a set of K_{ki} *candidate* analog precoding-combining pairs for communicating from k to i.

$$\mathcal{T}_{ij} = \left\{ \left(\mathbf{F}_{\text{RF}}^{(i)}, \mathbf{W}_{\text{RF}}^{(j)} \right)_n : n = 1, \ldots, K_{ij} \right\} \tag{3.17}$$

$$\mathcal{T}_{ki} = \left\{ \left(\mathbf{F}_{\text{RF}}^{(k)}, \mathbf{W}_{\text{RF}}^{(i)} \right)_n : n = 1, \ldots, K_{ki} \right\} \tag{3.18}$$

Our goal is to form the candidate sets $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ with promising beamforming pairs for each link. We describe our method for doing so—which can be

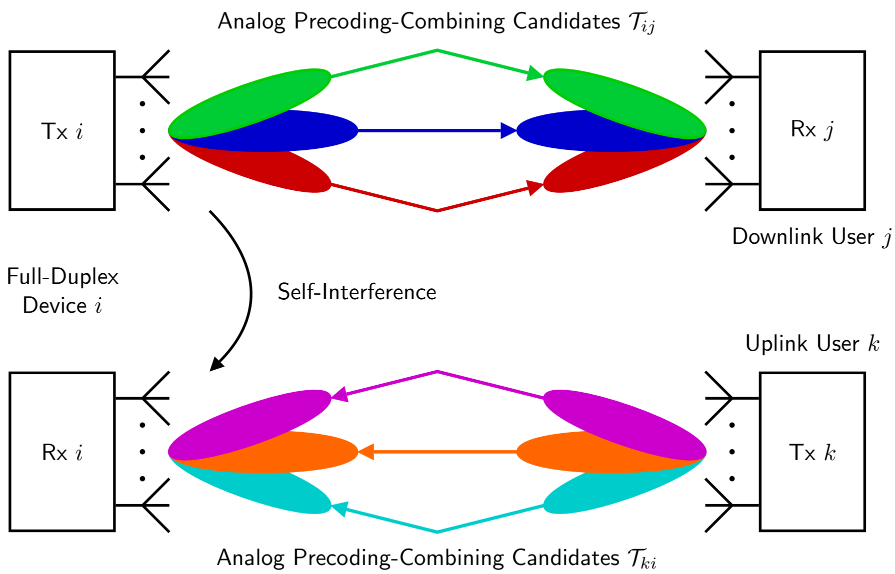

Figure 3.2: Precoding-combining candidates exploit spatial diversity to deliver higher downlink and uplink quality while ensuring receive chain components are not saturated.

replicated for $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ independently—by using generic notation (e.g., $\mathcal{T}$, K, $\mathbf{M}$) according to the summary shown in Algorithm 3.1.

We begin by finding the indices of the training precoders $\mathcal{M}_{\mathrm{tx}}$ and the training combiners $\mathcal{M}_{\mathrm{rx}}$ that revealed the top K strongest measurements in $\mathbf{M}$, sorted according to descending strength. We initialize $\mathcal{T}$ to an empty set. The first beam (column) of the n-th candidate is steered along the n-th strongest entry in $\mathbf{M}$. To ensure that we do not transmit or receive along directions more than once, we keep track of each beam's transmit and receive indices (i.e., t and r) in sets $\mathcal{J}_{\mathrm{tx}}$ and $\mathcal{J}_{\mathrm{rx}}$, respectively. Then, to choose the next column, we locate the strongest entry in $\mathbf{M}$ whose transmit or receive beam has not already been selected for this n-th candidate. When the columns of

Algorithm 3.1 Beam candidate acquisition algorithm.

Input: $\mathbf{M}$, $\mathbf{F}_{\mathrm{tr}}$, $\mathbf{W}_{\mathrm{tr}}$, L, K

  $\mathcal{T} = \{\varnothing\}$

  $\mathcal{M}_{\mathrm{rx}}, \mathcal{M}_{\mathrm{tx}} = \arg\mathrm{maxk}(\mathrm{abs}(\mathbf{M}), K)$

  for $n = 1 : K$ **do**

    $t = [\mathcal{M}_{\mathrm{tx}}]_n$

    $r = [\mathcal{M}_{\mathrm{rx}}]_n$

    $\mathbf{F}_{\mathrm{RF}} = [\mathbf{F}_{\mathrm{tr}}]_{:,t}$

    $\mathbf{W}_{\mathrm{RF}} = [\mathbf{W}_{\mathrm{tr}}]_{:,r}$

    $\mathcal{J}_{\mathrm{tx}} = \{t\}$

    $\mathcal{J}_{\mathrm{rx}} = \{r\}$

    for $\ell = 1 : L - 1$ **do**

      $r, t = \arg\max_{r,t} \left| [\mathbf{M}]_{r,t} \right|$ s.t. $t \notin \mathcal{J}_{\mathrm{tx}},\ r \notin \mathcal{J}_{\mathrm{rx}}$

      $\mathbf{F}_{\mathrm{RF}} = \begin{bmatrix} \mathbf{F}_{\mathrm{RF}} & [\mathbf{F}_{\mathrm{tr}}]_{:,t} \end{bmatrix}$

      $\mathbf{W}_{\mathrm{RF}} = \begin{bmatrix} \mathbf{W}_{\mathrm{RF}} & [\mathbf{W}_{\mathrm{tr}}]_{:,r} \end{bmatrix}$

      $\mathcal{J}_{\mathrm{tx}} = \mathcal{J}_{\mathrm{tx}} \cup t$

      $\mathcal{J}_{\mathrm{rx}} = \mathcal{J}_{\mathrm{rx}} \cup r$

    end for

    $\mathcal{T} = \mathcal{T} \cup (\mathbf{F}_{\mathrm{RF}}, \mathbf{W}_{\mathrm{RF}})$

  end for

Output: $\mathcal{T}$

$\mathbf{F}_{\mathrm{tr}}$ and $\mathbf{W}_{\mathrm{tr}}$ are linearly independent, this ensures that each analog precoding candidate and analog combining candidate are rank-L, which is necessary for multiplexing up to L streams. Note that, in this method, we have assumed that $L = L_{\mathrm{t}} = L_{\mathrm{r}} = N_{\mathrm{s}}$ for a given link, which is more practical than supplying devices with more RF chains than streams. This process is repeated until all L columns are populated, and then the candidate pair is appended to our candidate set $\mathcal{T}$. Once all K candidates have been generated, the set $\mathcal{T}$ is returned.

To potentially reduce overhead, modifications to our proposed method can accommodate creative approaches to beam alignment, such as hierarchical beam search, compressed sensing, and partial codebook search. With slight modifications, our method could also accommodate cases where training measurements map to codebook candidates rather than measuring the candidates themselves. Additionally, we point out that L^2 beam pairs on a given link may be measured simultaneously with L RF chains. In general, to build promising candidates, it is critical that the number of training beams (M_{t} and M_{r}) be sufficiently large to locate at least N strong rays in each channel such that $K \leq \binom{N}{L}$.

Following the construction of $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ using this method, we now suggest that the following channel estimates be made, having not assumed knowledge of $\mathbf{H}_{ij}$ or $\mathbf{H}_{ki}$. To provide the design presented in the next section with channel information, we assume a set $\mathcal{H}_{ij}$ has been populated with estimates of the effective channel seen by each of the candidates in $\mathcal{T}_{ij}$ described as

$$\mathcal{H}_{ij} = \left\{ \mathbf{W}_{\mathrm{RF}}^{(j)*} \mathbf{H}_{ij} \mathbf{F}_{\mathrm{RF}}^{(i)} : \left(\mathbf{F}_{\mathrm{RF}}^{(i)}, \mathbf{W}_{\mathrm{RF}}^{(j)} \right) \in \mathcal{T}_{ij} \right\} \tag{3.19}$$

where $|\mathcal{H}_{ij}| = K_{ij}$. Note that each effective channel in $\mathcal{H}_{ij}$ is merely $L_{\mathrm{r}}^{(j)} \times L_{\mathrm{t}}^{(i)}$, a very small size relative to $\mathbf{H}_{ij}$, and is observed digitally. Also, under channel reciprocity, measurements in $\mathcal{H}_{ij}$ can be conducted from j to i to reduce overhead since only one of its entries needs to be fed back to j following the design, which takes place at i. By these accounts, it is our hope that the overhead associated with collecting the measurements in $\mathcal{H}_{ij}$ not be prohibitive.

Furthermore, we would like to point out that our design does not require executing this estimation on the link from k to i. This concludes beam candidate acquisition, having populated $\mathcal{T}_{ij}$, $\mathcal{T}_{ki}$, and $\mathcal{H}_{ij}$, which will enable the design presented in the next section.

3.5 Hybrid Beamforming Design

In this section, we present a hybrid beamforming design to enable mmWave full-duplex while accounting for per-antenna and per-RF chain power constraints at the receiver of a full-duplex device. The goal of our design is to achieve a high spectral efficiency on both links while ensuring that the power of self-interference reaching the full-duplex device's receiver is below some thresholds. To do so, our design leverages the analog beamforming candidate sets $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ found in the previous section to configure the analog beamformers at each device. The design that follows holds for general per-antenna and per-RF chain power constraints, even though we are considering LNA and ADC power constraints as particular motivators.

Our design supports spatial multiplexing multiple streams and importantly does not require more RF chains than necessary, where $L_{\mathrm{t}}^{(i)} = L_{\mathrm{r}}^{(j)} = N_{\mathrm{s}}^{(ij)}$ and $L_{\mathrm{t}}^{(k)} = L_{\mathrm{r}}^{(i)} = N_{\mathrm{s}}^{(ki)}$. Furthermore, we support the important case where $N_{\mathrm{s}}^{(ij)} = N_{\mathrm{s}}^{(ki)}$, implying $L_{\mathrm{t}}^{(i)} = L_{\mathrm{r}}^{(i)}$. This, along with the fact that we have not assumed anything about $\mathbf{H}_{ii}$, means our design's formulation does not rely on *completely* avoiding (i.e., zero-forcing) the over-the-air self-interference channel or even the *effective* self-interference channel, unlike several existing approaches. We assume we have perfect channel knowledge of $\mathbf{H}_{ii}$ but do not assume knowledge of $\mathbf{H}_{ij}$ or $\mathbf{H}_{ki}$. We motivate this assumption by presuming that the self-interference channel can be more reliably estimated given its strength and can be done so possibly through regular calibration. This work aims to provide an approximate "best-case" design under select practical considerations, chiefly a limited receive dynamic range, motivating us to ignore the impact of channel estimation error in this work; it would be valuable future work to investigate its impact. We assume large-scale quantities (e.g., transmit powers, large-scale channel gains, SNRs) are known.

3.5.1 Per-Antenna and Per-RF Chain Power Constraints

Let $\tilde{P}_{\mathrm{SI,LNA}}^{\mathrm{max}}$ and $\tilde{P}_{\mathrm{SI,ADC}}^{\mathrm{max}}$ be the maximum average self-interference power (in watts) over the symbol period allowed at each LNA and each ADC of the receiver of the full-duplex device i, respectively. Referring to the terms presented in our system model and problem formulation, let us form our LNA and ADC constraints using $\tilde{P}_{\mathrm{SI,LNA}}^{\mathrm{max}}$ and $\tilde{P}_{\mathrm{SI,ADC}}^{\mathrm{max}}$. A constraint bounding the symbol power (in watts) at each LNA can be written as

$$\left| \sqrt{\tilde{P}_{\mathrm{tx}}^{(i)}} G_{ii} [\mathbf{H}_{ii}]_{\ell,:} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{s}^{(i)} \right|^2 \leq \tilde{P}_{\mathrm{SI,LNA}}^{\mathrm{max}} \tag{3.20}$$

for all $\ell = 1, \ldots, N_{\mathrm{r}}^{(i)}$. Collecting these $N_{\mathrm{r}}^{(i)}$ constraints together, we can write

$$\mathrm{diag}\left(\tilde{P}_{\mathrm{tx}}^{(i)} G_{ii}^2 \mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{s}^{(i)} \mathbf{s}^{(i)*} \mathbf{F}_{\mathrm{BB}}^{(i)*} \mathbf{F}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ii}^{*} \right) \leq \tilde{P}_{\mathrm{SI,LNA}}^{\mathrm{max}} \cdot \mathbf{1}_{N_{\mathrm{r}}^{(i)}} \tag{3.21}$$

where $\mathbf{a} \leq \mathbf{b}$ denotes element-wise inequality. For a given channel realization, it is impractical to attempt to satisfy (3.21) on a per-symbol basis. Furthermore, it may be computationally expensive, severely sub-optimal, or potentially impossible to ensure (3.21) is met for all symbol vectors in a constellation. This motivates us to satisfy our constraint in expectation over $\mathbf{s}^{(i)}$ and, noting our defined symbol covariance (3.1), results in the constraint

$$\frac{\tilde{P}_{\mathrm{tx}}^{(i)} G_{ii}^2}{N_{\mathrm{s}}^{(ij)}} \cdot \mathrm{diag}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*}\mathbf{F}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}^{*}\right) \leq \tilde{P}_{\mathrm{SI,LNA}}^{\max} \cdot \mathbf{1}_{N_{\mathrm{r}}^{(i)}}. \tag{3.22}$$

In a similar fashion, by incorporating the analog combiner $\mathbf{W}_{\mathrm{RF}}^{(i)}$, we can express our per-RF chain received self-interference power constraint as

$$\frac{\tilde{P}_{\mathrm{tx}}^{(i)} G_{ii}^2}{N_{\mathrm{s}}^{(ij)}} \cdot \mathrm{diag}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*}\mathbf{F}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}^{*}\mathbf{W}_{\mathrm{RF}}^{(i)}\right) \leq \tilde{P}_{\mathrm{SI,ADC}}^{\max} \cdot \mathbf{1}_{L_{\mathrm{r}}^{(i)}} \tag{3.23}$$

which captures all $L_{\mathrm{r}}^{(i)}$ ADCs in a single expression.

To abstract out the impact $\tilde{P}_{\mathrm{tx}}^{(i)}$ and G_{ii}^2 have on meeting our per-antenna power constraint $\tilde{P}_{\mathrm{SI,LNA}}^{\max}$ and per-RF chain power constraint $\tilde{P}_{\mathrm{SI,ADC}}^{\max}$, we introduce the unitless variables

$$\eta_{\mathrm{LNA}} \triangleq \frac{\tilde{P}_{\mathrm{SI,LNA}}^{\max}}{\tilde{P}_{\mathrm{tx}}^{(i)} \cdot G_{ii}^2}, \quad \eta_{\mathrm{ADC}} \triangleq \frac{\tilde{P}_{\mathrm{SI,ADC}}^{\max}}{\tilde{P}_{\mathrm{tx}}^{(i)} \cdot G_{ii}^2} \tag{3.24}$$

which will provide more generalized analysis across combinations of $\tilde{P}_{\mathrm{SI,LNA}}^{\max}$, $\tilde{P}_{\mathrm{SI,ADC}}^{\max}$, $\tilde{P}_{\mathrm{tx}}^{(i)}$, and G_{ii}^2. Stricter constraints are when η_{LNA} and η_{ADC} are low (we permit little self-interference at the receiver), while relaxed constraints are when η_{LNA} and η_{ADC} are high (we permit high self-interference). Using (3.24), the constraints in (3.22) and (3.23) can be approximately expressed as

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{LNA}} \tag{3.25}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{ADC}} \tag{3.26}$$

respectively, where $\sigma_{\max}(\mathbf{A})$ denotes the maximum singular value of $\mathbf{A}$.

3.5.2 Satisfying Our Constraints

With our LNA and ADC constraints in hand, we turn our attention to producing a hybrid beamforming design that satisfies (3.25) and (3.26) while achieving an appreciable spectral efficiency on our two links. The mutual information (under Gaussian signaling), or spectral efficiency (in bps/Hz), of the link from i to j is referred to as $\mathcal{R}_{ij}$ and takes the familiar form [88].

$$\mathcal{R}_{ij} = \log_2 \left| \mathbf{I} + \frac{\mathsf{SNR}_{ij}}{N_{\mathrm{s}}^{(ij)}} \mathbf{W}_{\mathrm{BB}}^{(j)*} \mathbf{W}_{\mathrm{RF}}^{(j)*} \mathbf{H}_{ij} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)*} \mathbf{F}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ij}^{*} \mathbf{W}_{\mathrm{RF}}^{(j)} \mathbf{W}_{\mathrm{BB}}^{(j)} \left(\mathbf{Q}_{\mathrm{n}}^{(j)} \right)^{-1} \right|. \tag{3.27}$$

Treating the effects of self-interference as noise, the mutual information $\mathcal{R}_{ki}$ of the link from k to i is expressed as

$$\mathcal{R}_{ki} = \log_2 \left| \mathbf{I} + \frac{\mathsf{SNR}_{ki}}{N_{\mathrm{s}}^{(ki)}} \mathbf{W}_{\mathrm{BB}}^{(i)*} \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ki} \mathbf{F}_{\mathrm{RF}}^{(k)} \mathbf{F}_{\mathrm{BB}}^{(k)} \mathbf{F}_{\mathrm{BB}}^{(k)*} \mathbf{F}_{\mathrm{RF}}^{(k)*} \mathbf{H}_{ki}^{*} \mathbf{W}_{\mathrm{RF}}^{(i)} \mathbf{W}_{\mathrm{BB}}^{(i)} \left(\mathbf{Q}_{\mathrm{n}}^{(i)} + \mathbf{Q}_{\mathrm{int}}^{(i)} \right)^{-1} \right|. \tag{3.28}$$

We let $\mathbf{Q}_{\mathrm{n}}^{(n)}$ be the covariance of noise at the detector of $n \in \{i, j\}$ normalized to the noise power and let $\mathbf{Q}_{\mathrm{int}}^{(i)}$ be the covariance of self-interference at the detector of i normalized to the noise power. Since noise and the transmitted symbols from i to j are uncorrelated, the self-interference-plus-noise covariance can be written as $\mathbf{Q}_{\mathrm{n}}^{(i)} + \mathbf{Q}_{\mathrm{int}}^{(i)}$.

We design our full-duplex system with the objective to maximize the sum spectral efficiency $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ subject to our per-antenna and per-RF chain self-interference power constraints as well as the precoding power constraint in (3.2), described below in problem (3.29).

$$\max_{\substack{\left(\mathbf{F}_{\mathrm{RF}}^{(i)},\mathbf{W}_{\mathrm{RF}}^{(j)}\right)\in\mathcal{T}_{ij} \\ \left(\mathbf{F}_{\mathrm{RF}}^{(k)},\mathbf{W}_{\mathrm{RF}}^{(i)}\right)\in\mathcal{T}_{ki}}} \ \max_{\substack{\mathbf{F}_{\mathrm{BB}}^{(i)},\mathbf{W}_{\mathrm{BB}}^{(j)} \\ \mathbf{F}_{\mathrm{BB}}^{(k)},\mathbf{W}_{\mathrm{BB}}^{(i)}}} \mathcal{R}_{ij}+\mathcal{R}_{ki} \tag{3.29a}$$

$$\text{s.t.} \quad \left\|\mathbf{F}_{\mathrm{BB}}^{(i)}\right\|_{\mathrm{F}}^{2}\leq 1, \ \left\|\mathbf{F}_{\mathrm{BB}}^{(k)}\right\|_{\mathrm{F}}^{2}\leq 1 \tag{3.29b}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}}\cdot\sigma_{\max}^{2}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right)\leq\eta_{\mathrm{LNA}} \tag{3.29c}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}}\cdot\sigma_{\max}^{2}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right)\leq\eta_{\mathrm{ADC}} \tag{3.29d}$$

Note that we have restricted our choices for analog beamforming to the sets $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ supplied from beam candidate acquisition in Section 3.4. Solving problem (3.29) is difficult for a variety of reasons, chiefly the non-convexity arising from the interplay of the precoder at i in both transmit link performance and self-interference, meaning it impacts both $\mathcal{R}_{ij}$ and $\mathcal{R}_{ki}$. This motivates us to split our design into two stages. The first stage will be to configure a portion of our system subject to our constraints. Then, with the constraints met by the first stage, the second stage of our design will configure the remaining precoders and combiners.

The responsibility of satisfying the per-antenna self-interference power constraint lay solely in the transmitter of i as evidenced by (3.29c). In satisfying the per-RF chain constraint (3.29d), the analog combiner $\mathbf{W}_{\mathrm{RF}}^{(i)}$ can assist, though—like the analog precoder $\mathbf{F}_{\mathrm{RF}}^{(i)}$—it may be very limited in its ability to do so. For both constraints, it is expected that $\mathbf{F}_{\mathrm{BB}}^{(i)}$ will often carry most of the weight in meeting these two constraints, given its digital nature. Since the price of meeting the per-antenna constraint is paid solely by the transmitter in $\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}$ and since we expect the digital precoder $\mathbf{F}_{\mathrm{BB}}^{(i)}$ to incur even more sacrifice as it meets the per-RF chain constraint, we are motivated to prioritize the transmit link during the first stage of our design *while* meeting these constraints. Note that by meeting these constraints, we will be *inherently* improving receive link performance by reducing self-interference.

Suppose during beam candidate acquisition, we let $K_{ij} = K_{ki} = 1$ (we consider $K_{ij}, K_{ki} \geq 1$ shortly), leading to the analog beamformers $\mathbf{F}_{\mathrm{RF}}^{(i)}$, $\mathbf{W}_{\mathrm{RF}}^{(j)}$, $\mathbf{F}_{\mathrm{RF}}^{(k)}$, and $\mathbf{W}_{\mathrm{RF}}^{(i)}$ being fixed as the first and only candidates in $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$. In such a case, we formulate problem (3.31), where we aim to maximize the transmit link mutual information

$$\mathcal{I}_{ij} = \log_2 \left| \mathbf{I} + \frac{\mathsf{SNR}_{ij}}{N_{\mathrm{s}}^{(ij)}} \mathbf{W}_{\mathrm{RF}}^{(j)*} \mathbf{H}_{ij} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)*} \mathbf{F}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ij}^{*} \mathbf{W}_{\mathrm{RF}}^{(j)} \left(\mathbf{W}_{\mathrm{RF}}^{(j)*} \mathbf{W}_{\mathrm{RF}}^{(j)} \right)^{-1} \right| \tag{3.30}$$

subject to our per-antenna and per-RF chain constraints and the aforementioned precoding power constraint in (3.2).

$$\max_{\mathbf{F}_{\mathrm{BB}}^{(i)}} \ \mathcal{I}_{ij} \tag{3.31a}$$

$$\text{s.t.} \ \left\| \mathbf{F}_{\mathrm{BB}}^{(i)} \right\|_{\mathrm{F}}^{2} \leq 1 \tag{3.31b}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2} \left(\mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \right) \leq \eta_{\mathrm{LNA}} \tag{3.31c}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2} \left(\mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \right) \leq \eta_{\mathrm{ADC}} \tag{3.31d}$$

Solving problem (3.31) would ensure that transmission from i to j is prioritized while preventing receiver-side saturation. Problem (3.31) is technically not convex but can be easily recast as such.

Theorem 3.1. *Problem* (3.31) *can be recast as a convex problem.*

Proof. Noting that problem (3.31) is non-convex in $\mathbf{F}_{\mathrm{BB}}^{(i)}$ but is convex in the product $\mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*}$, we rewrite problem (3.31) using the positive semidefinite substitution $\mathbf{X}_{\mathrm{BB}}^{(i)} = \mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*} \succeq \mathbf{0}$ as

$$\max_{\mathbf{X}_{\mathrm{BB}}^{(i)} \succeq \mathbf{0}} \quad \mathcal{I}_{ij} \tag{3.32a}$$

$$\text{s.t.} \quad \left\|\mathbf{F}_{\mathrm{BB}}^{(i)}\right\|_{\mathrm{F}}^{2} \leq 1 \tag{3.32b}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{LNA}} \tag{3.32c}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{ADC}} \tag{3.32d}$$

$$\mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*} = \mathbf{X}_{\mathrm{BB}}^{(i)} \tag{3.32e}$$

With this simple restructuring, problem (3.32) is convex and can be solved efficiently using a convex solver (e.g., we used CVX [89]). Once solved, $\mathbf{X}_{\mathrm{BB}}^{(i)}$ can be factored to recover $\mathbf{F}_{\mathrm{BB}}^{(i)}$. □

Remark 3.1. The solution to problem (3.32) is unique but the factorization $\mathbf{X}_{\mathrm{BB}}^{(i)} = \mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*}$ is not. As such, our problem is a function of the digital precoder covariance and not of the digital precoder itself. In other words, the solution $\mathbf{X}_{\mathrm{BB}}^{(i)}$ to problem (3.32) can be factored arbitrarily, since $L_{\mathrm{t}}^{(i)} = N_{\mathrm{s}}^{(ij)}$, to retrieve a globally optimal $\mathbf{F}_{\mathrm{BB}}^{(i)}$.

Remark 3.2. Since $\mathbf{X}_{\mathrm{BB}}^{(i)} = \mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*} = \mathbf{0}$ is always a solution to problem (3.32), it is feasible. Intuitively, this can be attributed to the fact that it is always possible for the transmitter at i to shut off completely to ensure the LNAs and ADCs do not saturate.

Remark 3.3. At least one of the constraints (3.32b)–(3.32d) will be tight (have equality) under the optimal solution to problem (3.32). Intuitively, this can be attributed to the fact that additional power should be supplied to the digital precoder $\mathbf{F}_{\mathrm{BB}}^{(i)}$—which will increase the mutual information $\mathcal{I}_{ij}$—until it exceeds its power budget or self-interference power is too high at an LNA or ADC.

We now incorporate our sets of candidate beams $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ when $K_{ij}, K_{ki} \geq 1$, which offers the system more freedom in its design. By doing so, the diversity between candidate beams will generally allow our full-duplex transceiver i to better transmit to j while meeting the constraints. This is courtesy of the fact that some candidate beam pairs will naturally afford more isolation at the full-duplex device, reducing the costs paid by $\mathbf{F}_{\mathrm{BB}}^{(i)}$ as it attempts to reduce self-interference. This motivates us to wrap problem (3.31) with an outer maximization over $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$, resulting in the following optimization problem.

$$\max_{\substack{\left(\mathbf{F}_{\mathrm{RF}}^{(i)}, \mathbf{W}_{\mathrm{RF}}^{(j)}\right) \in \mathcal{T}_{ij} \\ \left(\mathbf{F}_{\mathrm{RF}}^{(k)}, \mathbf{W}_{\mathrm{RF}}^{(i)}\right) \in \mathcal{T}_{ki}}} \ \max_{\mathbf{F}_{\mathrm{BB}}^{(i)}} \mathcal{I}_{ij} \tag{3.33a}$$

$$\text{s.t.} \quad \left\| \mathbf{F}_{\mathrm{BB}}^{(i)} \right\|_{\mathrm{F}}^{2} \leq 1 \tag{3.33b}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2}\left(\mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{LNA}} \tag{3.33c}$$

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{ADC}} \tag{3.33d}$$

Having shown that the inner maximization can be solved via the convex reformulation in (3.32), we can solve problem (3.33) exhaustively over all possible candidate combinations in $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$. Solving (3.33) will prioritize the transmit link while meeting the self-interference power constraints and will leverage the candidates from $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ to do so. As a result, the receive link may incur costs if the candidates chosen from $\mathcal{T}_{ki}$ are not the optimal ones in a half-duplex sense. Choosing a smaller K_{ki} would constrain this receive link cost but would simultaneously increase transmit link cost. Increasing K_{ij} can only help transmit link performance but may incur additional complexity during beam candidate acquisition and/or solving problem (3.33).

Note that, as evidenced in (3.30), $\mathcal{I}_{ij}$ contains $\mathbf{H}_{ij}$, which we have not assumed explicit knowledge of. Instead, using $\mathcal{H}_{ij}$ from (3.19), we do have knowledge of the effective channel $\mathbf{W}_{\mathrm{RF}}^{(j)*}\mathbf{H}_{ij}\mathbf{F}_{\mathrm{RF}}^{(i)}$ for all $(\mathbf{F}_{\mathrm{RF}}^{(i)}, \mathbf{W}_{\mathrm{RF}}^{(i)}) \in \mathcal{T}_{ij}$, which can be used when solving problem (3.33). Furthermore, note that we do not require knowledge of $\mathbf{H}_{ki}$ or $\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ki}\mathbf{F}_{\mathrm{RF}}^{(k)}$ to solve (3.33). We do require knowledge of $\mathbf{H}_{ii}$ to construct constraints (3.33c) and (3.33d), which we have assumed knowledge of. Notice, however, that perhaps an estimate of $\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}$ can be used to compute the per-RF chain constraint (3.33d) since it may be estimated more reliably and frequently than that of $\mathbf{H}_{ii}$, given its relatively small size and fully-digital nature.

Solving problem (3.33) yields the design for five of the eight precoding and combining matrices: $\mathbf{W}_{\mathrm{BB}}^{(j)}$, $\mathbf{F}_{\mathrm{BB}}^{(k)}$, and $\mathbf{W}_{\mathrm{BB}}^{(i)}$ remain to be designed. Designing these will take place in the next stage of our design. Having prevented receiver-side components from saturating with appropriately chosen η_{LNA} and η_{ADC}, the receive chain at device i is approximately linear, allowing for a much more straightforward design of the receive link and preventing severe degradation of a desired receive signal.

3.5.3 Constraint Redundancy Conditions

We now pause from our design to more closely examine the interplay between our precoding power constraint, per-antenna self-interference power constraint, and per-RF chain self-interference power constraint. In doing so, we will see that under appropriate conditions the latter two constraints may be inherently met by the precoding power constraint and under other conditions, the per-RF chain self-interference power constraint may be inherently met by the per-antenna self-interference power constraint. Knowledge of these conditions can potentially accelerate solving problem (3.33) in a variety of ways by relieving us of one or more constraints. As such, the complexity associated with solving problem (3.33)—which depends heavily on one's choice of algorithm—is also dictated by the interplay between its three constraints. This interplay, as we will see, is characterized by a variety of factors including the realized self-interference channel $\mathbf{H}_{ii}$, η_{LNA}, η_{ADC}, and choice of analog beamformers/codebooks.

Theorem 3.2. *When condition* (3.34) *holds, the per-antenna constraint* (3.33c) *is implicitly met by the precoding power constraint* (3.33b).

$$\eta_{\mathrm{LNA}} \geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^{2}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) \tag{3.34}$$

Proof. Since the precoding power constraint (3.31b) is satisfied, we note that

$$\sigma_{\max}^2\left(\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \left\|\mathbf{F}_{\mathrm{BB}}^{(i)}\right\|_{\mathrm{F}}^2 \leq 1 \tag{3.35}$$

which allows us to see that

$$\begin{aligned}
\eta_{\mathrm{LNA}} &\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) && (3.36)\\
&\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) \cdot \sigma_{\max}^2\left(\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.37)\\
&\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.38)
\end{aligned}$$

where we have used $\sigma_{\max}^2(\mathbf{AB}) \leq \sigma_{\max}^2(\mathbf{A}) \cdot \sigma_{\max}^2(\mathbf{B})$. □

Theorem 3.3. *When condition* (3.39) *holds, the per-RF chain constraint* (3.33d) *is implicitly met by the precoding power constraint* (3.33b)*.*

$$\eta_{\mathrm{ADC}} \geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) \tag{3.39}$$

Proof. Using the fact from (3.35) and assuming (3.39) to be true, we can write

$$\begin{aligned}
\eta_{\mathrm{ADC}} &\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) && (3.40)\\
&\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\right) \cdot \sigma_{\max}^2\left(\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.41)\\
&\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma_{\max}^2\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.42)
\end{aligned}$$

□

Remark 3.4. When (3.34) and (3.39) hold, the (half-duplex) capacity-achieving strategy on the transmit link is the solution to problem (3.31) since the per-antenna and per-RF chain constraints vanish, leaving only the precoding power constraint.

Theorem 3.4. *When the per-antenna constraint* (3.33c) *is satisfied and condition* (3.43) *holds, the per-RF chain constraint* (3.33d) *is satisfied.*

$$\eta_{\mathrm{ADC}} \geq \eta_{\mathrm{LNA}} \cdot \sigma^2_{\mathrm{max}}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\right) \tag{3.43}$$

Proof. Starting with our assumption that the per-antenna constraint (3.33c) is satisfied, we have

$$\frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma^2_{\mathrm{max}}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) \leq \eta_{\mathrm{LNA}} \leq \frac{\eta_{\mathrm{ADC}}}{\sigma^2_{\mathrm{max}}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\right)} \tag{3.44}$$

which yields

$$\begin{aligned} \eta_{\mathrm{ADC}} &\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma^2_{\mathrm{max}}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\right) \cdot \sigma^2_{\mathrm{max}}\left(\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.45) \\ &\geq \frac{1}{N_{\mathrm{s}}^{(ij)}} \cdot \sigma^2_{\mathrm{max}}\left(\mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\right) && (3.46) \end{aligned}$$

which indicates directly that (3.33d) is satisfied. □

Corollary 3.1. *When* $\eta_{\mathrm{LNA}} = 0$ *and* $\eta_{\mathrm{ADC}} > 0$*, satisfying the per-antenna constraint* (3.33c) *also satisfies the per-RF chain constraint* (3.33d)*.*

Remark 3.5. A condition analogous to Theorem 3.4 stating that satisfying the per-antenna constraint (3.33c) based solely on meeting the per-RF chain constraint (3.33d) is not possible.

3.5.4 Remainder of the Design

We now complete our mmWave MIMO design. Solving (3.33) yields selections for $\mathbf{F}_{\mathrm{BB}}^{(i)}$, $\mathbf{F}_{\mathrm{RF}}^{(i)}$, $\mathbf{W}_{\mathrm{RF}}^{(j)}$, $\mathbf{F}_{\mathrm{RF}}^{(k)}$, and $\mathbf{W}_{\mathrm{RF}}^{(i)}$. Configuring $\mathbf{W}_{\mathrm{BB}}^{(j)}$, $\mathbf{F}_{\mathrm{BB}}^{(k)}$, and

$\mathbf{W}_{\text{BB}}^{(i)}$ remains, which we execute as follows. Having the rest of the transmit link configured, the optimal linear baseband combiner at j can be designed in a linear minimum mean square error (LMMSE) fashion as follows. Let $\tilde{\mathbf{H}}_{ij} \triangleq \mathbf{W}_{\text{RF}}^{(j)*}\mathbf{H}_{ij}\mathbf{F}_{\text{RF}}^{(i)}\mathbf{F}_{\text{BB}}^{(i)} \in \mathbb{C}^{L_{\text{r}}^{(j)} \times N_{\text{s}}^{(ij)}}$ be the effective transmit channel after solving (3.33). Note that this can computed from the product of $\mathbf{W}_{\text{RF}}^{(j)*}\mathbf{H}_{ij}\mathbf{F}_{\text{RF}}^{(i)}$, which is referenced from $\mathcal{H}_{ij}$, and $\mathbf{F}_{\text{BB}}^{(i)}$. Then, the LMMSE baseband combiner at j can be constructed as

$$\mathbf{W}_{\text{BB}}^{(j)} = \frac{1}{\sqrt{P_{\text{tx}}^{(i)} G_{ij}}} \left(\tilde{\mathbf{H}}_{ij}\tilde{\mathbf{H}}_{ij}^{*} + \frac{N_{\text{s}}^{(ij)}}{\mathsf{SNR}_{ij}} \mathbf{W}_{\text{RF}}^{(j)*}\mathbf{W}_{\text{RF}}^{(j)} \right)^{-1} \tilde{\mathbf{H}}_{ij}. \tag{3.47}$$

Since $\mathbf{F}_{\text{BB}}^{(i)}$ will not generally diagonalize the effective channel, an LMMSE combiner at j will aim to reduce inter-stream interference and reject noise. This concludes configuration of the transmit link, having set $\mathbf{F}_{\text{BB}}^{(i)}$, $\mathbf{F}_{\text{RF}}^{(i)}$, $\mathbf{W}_{\text{RF}}^{(j)}$, and $\mathbf{W}_{\text{BB}}^{(j)}$.

We now turn our attention to the receive link, where we need to configure $\mathbf{F}_{\text{BB}}^{(k)}$ and $\mathbf{W}_{\text{BB}}^{(i)}$. Now that we have made our selections of $\mathbf{F}_{\text{RF}}^{(k)}$ and $\mathbf{W}_{\text{RF}}^{(i)}$, we begin by estimating the relatively small channel $\tilde{\mathbf{H}}_{ki} \triangleq \mathbf{W}_{\text{RF}}^{(i)*}\mathbf{H}_{ki}\mathbf{F}_{\text{RF}}^{(k)} \in \mathbb{C}^{L_{\text{r}}^{(i)} \times L_{\text{t}}^{(k)}}$, which can be observed digitally and we assume is error-free. Taking the singular value decomposition (SVD) of this effective channel from k to i and accounting for noise coloring, we get $\mathbf{U}_{ki}\mathbf{\Sigma}_{ki}\mathbf{V}_{ki}^{*} = \text{SVD}\left(\left(\mathbf{W}_{\text{RF}}^{(i)*}\mathbf{W}_{\text{RF}}^{(i)} \right)^{-1/2} \tilde{\mathbf{H}}_{ki} \right)$. We then build the precoder that maximizes the mutual information offered to the RF chains of i as

$$\mathbf{F}_{\text{BB}}^{(k)} = \left[\mathbf{V}_{ki}\right]_{:,1:N_{\text{s}}^{(ki)}} \times \mathbf{P}^{(k)} \tag{3.48}$$

where $\mathbf{P}^{(k)}$ is a diagonal water-filling power allocation matrix [88].

Finally, we are left to configure $\mathbf{W}_{\mathrm{BB}}^{(i)}$. Before doing so, however, we notice that $\mathbf{W}_{\mathrm{BB}}^{(i)}$ and the received symbols it acts on exist in the digital domain, residing after the ADCs. Having knowledge of $\mathbf{H}_{ii}$, $\mathbf{s}^{(i)}$, and all other beamformers, we can synthesize the received self-interference and subtract it before applying our combiner $\mathbf{W}_{\mathrm{BB}}^{(i)}$. Recall from (3.14) that the symbol vector after the ADCs is $\mathbf{y}_{\mathrm{dig}} = \mathbf{y}_{\mathrm{des,ADC}} + \mathbf{y}_{\mathrm{int,ADC}} + \mathbf{y}_{\mathrm{noise,ADC}} + \mathbf{e}_{\mathrm{quant}}$. Since $\mathbf{y}_{\mathrm{dig}}$ is in the digital domain, we can compute $\mathbf{y}_{\mathrm{int,ADC}}$ as

$$\mathbf{y}_{\mathrm{int,ADC}} = \sqrt{P_{\mathrm{tx}}^{(i)}} G_{ii} \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ii} \mathbf{F}_{\mathrm{RF}}^{(i)} \mathbf{F}_{\mathrm{BB}}^{(i)} \mathbf{s}^{(i)} \tag{3.49}$$

and can subtract it from $\mathbf{y}_{\mathrm{dig}}$ before applying our combiner, which yields

$$\begin{aligned} \mathbf{y}^{(i)} &= \mathbf{y}_{\mathrm{dig}} - \mathbf{y}_{\mathrm{int,ADC}} && (3.50) \\ &= \mathbf{y}_{\mathrm{des,ADC}} + \mathbf{y}_{\mathrm{noise,ADC}} + \mathbf{e}_{\mathrm{quant}}. && (3.51) \end{aligned}$$

We can estimate the symbols intended for i from k by applying a combiner $\mathbf{W}_{\mathrm{BB}}^{(i)}$ to $\mathbf{y}^{(i)}$ as

$$\begin{aligned} \hat{\mathbf{s}}^{(k)} &= \mathbf{W}_{\mathrm{BB}}^{(i)*} \mathbf{y}^{(i)} && (3.52) \\ &= \mathbf{W}_{\mathrm{BB}}^{(i)*} (\mathbf{y}_{\mathrm{des,ADC}} + \mathbf{y}_{\mathrm{noise,ADC}} + \mathbf{e}_{\mathrm{quant}}) && (3.53) \end{aligned}$$

which will be corrupted by additive noise and the effects of quantization. At first glance, it appears that self-interference does not play a role our symbol estimate $\hat{\mathbf{s}}^{(k)}$. However, given that our ADCs have limited resolution, the power of the quantization error term $\mathbf{e}_{\mathrm{quant}}$ will increase with increased self-interference power. Recall that this is precisely the motivation for our per-RF chain self-interference power constraint.

As evidenced by (3.53), we can see that the linear combiner $\mathbf{W}_{\mathrm{BB}}^{(i)}$ acts on a desired signal plus two noise terms. Before proceeding, let us find the covariance of $\mathbf{e}_{\mathrm{quant}}$. Finding the covariance of the symbols reaching the ADCs, we get

$$\begin{aligned}\mathbb{E}[\mathbf{y}_{\mathrm{ADC}}\mathbf{y}_{\mathrm{ADC}}^{*}] = &\frac{P_{\mathrm{tx}}^{(k)} G_{ki}^{2}}{N_{\mathrm{s}}^{(ki)}} \cdot \mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ki}\mathbf{F}_{\mathrm{RF}}^{(k)}\mathbf{F}_{\mathrm{BB}}^{(k)}\mathbf{F}_{\mathrm{BB}}^{(k)*}\mathbf{F}_{\mathrm{RF}}^{(k)*}\mathbf{H}_{ki}^{*}\mathbf{W}_{\mathrm{RF}}^{(i)} \\ &+ \frac{P_{\mathrm{tx}}^{(i)} G_{ii}^{2}}{N_{\mathrm{s}}^{(ij)}} \cdot \mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}\mathbf{F}_{\mathrm{RF}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)}\mathbf{F}_{\mathrm{BB}}^{(i)*}\mathbf{F}_{\mathrm{RF}}^{(i)*}\mathbf{H}_{ii}^{*}\mathbf{W}_{\mathrm{RF}}^{(i)} \\ &+ \sigma_{\mathrm{n}}^{2} \cdot \mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{W}_{\mathrm{RF}}^{(i)}. \end{aligned} \tag{3.54}$$

Applying (3.12) per-ADC allows us to write the covariance matrix of $\mathbf{e}_{\mathrm{quant}}$ as

$$\mathbf{R}_{\mathrm{quant}} = \mathbb{E}\left[\mathbf{e}_{\mathrm{quant}}\mathbf{e}_{\mathrm{quant}}^{*}\right] \tag{3.55}$$

$$= \frac{1}{1.5 \cdot 2^{2b}} \cdot \mathbf{I} \odot \mathbb{E}[\mathbf{y}_{\mathrm{ADC}}\mathbf{y}_{\mathrm{ADC}}^{*}] \tag{3.56}$$

where $\odot$ denotes the Hadamard (element-wise) product. The linear design of $\mathbf{W}_{\mathrm{BB}}^{(i)}$ that minimizes the mean square error (MSE) of $\hat{\mathbf{s}}^{(k)}$ is

$$\mathbf{W}_{\mathrm{BB}}^{(i)} = \frac{1}{\sqrt{P_{\mathrm{tx}}^{(k)}} G_{ki}} \left(\tilde{\mathbf{H}}_{ki}\tilde{\mathbf{H}}_{ki}^{*} + \frac{N_{\mathrm{s}}^{(ki)}}{\mathsf{SNR}_{ki}} \mathbf{W}_{\mathrm{RF}}^{(i)*}\mathbf{W}_{\mathrm{RF}}^{(i)} + \frac{N_{\mathrm{s}}^{(ki)}}{P_{\mathrm{tx}}^{(k)} G_{ki}^{2}} \mathbf{R}_{\mathrm{quant}} \right)^{-1} \tilde{\mathbf{H}}_{ki}. \tag{3.57}$$

Now that our design is complete, we characterize the covariance of self-interference and of noise at the detector, which can be used to evaluate $\mathcal{R}_{ij}$

and $\mathcal{R}_{ki}$ in (3.27) and (3.28), respectively. Let the covariance of the received quantization noise due to self-interference at i, normalized to the noise power, be

$$\mathbf{Q}_{\mathrm{int}}^{(i)} = \sigma_{\mathrm{n}}^{-2} \cdot \mathbf{W}_{\mathrm{BB}}^{(i)*} \mathbf{R}_{\mathrm{quant}} \mathbf{W}_{\mathrm{BB}}^{(i)}. \tag{3.58}$$

Similarly, let the covariance of the received noise at $n \in \{i, j\}$, normalized to the noise power, be written as

$$\mathbf{Q}_{\mathrm{n}}^{(n)} = \mathbf{W}_{\mathrm{BB}}^{(n)*} \mathbf{W}_{\mathrm{RF}}^{(n)*} \mathbf{W}_{\mathrm{RF}}^{(n)} \mathbf{W}_{\mathrm{BB}}^{(n)}. \tag{3.59}$$

This concludes our design, which we evaluate in the following section.

3.6 Numerical Results

We have simulated our system in a Monte Carlo fashion with the following parameters. For simplicity, we use 32-element, half-wavelength uniform linear arrays with isotropic elements at all devices, where the horizontal transmit and receive array at i are separated vertically by 10 wavelengths. Each transmitter and receiver is equipped with 2 RF chains and accordingly multiplexes 2 streams. The transmit power at each device is 30 dBm, and the noise power is -85 dBm. To model the channel between the transmit and receive arrays of device i, we use the following summation [19, 72],

$$\mathbf{H}_{ii} = \sqrt{\frac{\kappa}{\kappa+1}} \mathbf{H}_{ii}^{\mathrm{NF}} + \sqrt{\frac{1}{\kappa+1}} \mathbf{H}_{ii}^{\mathrm{FF}} \tag{3.60}$$

where the Rician factor κ captures the amount of power in the near-field portion relative to the far-field portion. The near-field component is modeled using a spherical-wave model [2] as

$$\left[\mathbf{H}_{ii}^{\mathrm{NF}}\right]_{v,u} = \frac{\gamma}{r_{u,v}} \exp\left(-\mathrm{j}2\pi \frac{r_{u,v}}{\lambda}\right) \tag{3.61}$$

where $r_{u,v}$ is the distance between the u-th transmit antenna and the v-th receive antenna, λ is the carrier wavelength, and γ ensures that the channel is normalized such that $\mathbb{E}\left[\|\mathbf{H}_{ii}\|_{\mathrm{F}}^2\right] = N_{\mathrm{t}}^{(i)} N_{\mathrm{r}}^{(i)}$. Note that this near-field model is deterministic for a given relative array geometry at the full-duplex device i. The far-field component captures reflections from the environment and is modeled using (3.5) with $N_{\mathrm{rays}} \sim \mathrm{Unif}(1, 15)$. The transmit and receive channels are modeled with $N_{\mathrm{rays}} \sim \mathrm{Unif}(4, 15)$. For both channel models, each ray's AoD and AoA are drawn from $\mathrm{Unif}(-\pi/2, \pi/2)$. During beam candidate acquisition, we assume properly normalized discrete Fourier transform (DFT) codebooks are used. Furthermore, we assume the number of measurements taken during beam candidate acquisition is sufficiently large on each link such that $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$ are built using the strongest rays in their respective channels.

Let us define the transmit link capacity $\mathcal{C}_{ij}$ as the maximum spectral efficiency possible on the link from i to j when drawing the analog precoder $\mathbf{F}_{\mathrm{RF}}^{(i)}$ and analog combiner $\mathbf{W}_{\mathrm{RF}}^{(j)}$ from the beam candidate set $\mathcal{T}_{ij}$ as

$$\mathcal{C}_{ij} = \max_{\mathbf{F}_{\mathrm{BB}}^{(i)}, \left(\mathbf{F}_{\mathrm{RF}}^{(i)}, \mathbf{W}_{\mathrm{RF}}^{(j)}\right) \in \mathcal{T}_{ij}} \mathcal{I}_{ij} \quad \text{s.t.} \quad \left\|\mathbf{F}_{\mathrm{BB}}^{(i)}\right\|_{\mathrm{F}}^2 \leq 1 \tag{3.62}$$

which can be achieved using the well known method of water-filled eigenbeamforming. Let us define the receive link capacity $\mathcal{C}_{ki}$ as the maximum spectral efficiency possible on the link from k to i when drawing the analog precoder

$\mathbf{F}_{\mathrm{RF}}^{(k)}$ and the analog combiner $\mathbf{W}_{\mathrm{RF}}^{(i)}$ from the candidate set $\mathcal{T}_{ki}$ as

$$\mathcal{C}_{ki} = \max_{\mathbf{F}_{\mathrm{BB}}^{(k)}, \left(\mathbf{F}_{\mathrm{RF}}^{(k)}, \mathbf{W}_{\mathrm{RF}}^{(i)}\right) \in \mathcal{T}_{ki}} \mathcal{I}_{ki} \quad \text{s.t.} \quad \left\| \mathbf{F}_{\mathrm{BB}}^{(k)} \right\|_{\mathrm{F}}^{2} \leq 1 \tag{3.63}$$

where $\mathcal{I}_{ki}$ is

$$\mathcal{I}_{ki} = \log_2 \left| \mathbf{I} + \frac{\mathsf{SNR}_{ki}}{N_{\mathrm{s}}^{(ki)}} \mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{H}_{ki} \mathbf{F}_{\mathrm{RF}}^{(k)} \mathbf{F}_{\mathrm{BB}}^{(k)} \mathbf{F}_{\mathrm{BB}}^{(k)*} \mathbf{F}_{\mathrm{RF}}^{(k)*} \mathbf{H}_{ki}^{*} \mathbf{W}_{\mathrm{RF}}^{(i)} \left(\mathbf{W}_{\mathrm{RF}}^{(i)*} \mathbf{W}_{\mathrm{RF}}^{(i)} \right)^{-1} \right|. \tag{3.64}$$

While $\mathcal{C}_{ij}$ and $\mathcal{C}_{ki}$ are not the true channel capacities of $\mathbf{H}_{ij}$ and $\mathbf{H}_{ki}$, it is more meaningful when evaluating our results to use our codebook-based analog beamforming approach to accurately interpret the spectral efficiency gains (and costs) associated with our design versus a half-duplex system that is offered the same freedom in analog beamforming contained in $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$. Our design will, therefore, hope to achieve a sum spectral efficiency $\mathcal{R}_{ij}+\mathcal{R}_{ki} \geq \max\{\mathcal{C}_{ij}, \mathcal{C}_{ki}\}$ to justify operating in a full-duplex fashion rather than a half-duplex one.

As intuition suggests, **the stricter the LNA and ADC constraints, the greater the sacrifice made on the transmit link's spectral efficiency to meet these constraints**. Figure 3.3 confirms this, where the sum spectral efficiency $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ as a function of SNR is evaluated at various LNA and ADC power constraints. While not explicitly shown, the loss in sum spectral efficiency due to decreasing $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB is due to loss in $\mathcal{R}_{ij}$ in its attempt to prevent receiver-side saturation. Given that we are primarily preserving the receive link by limiting the self-interference power reaching it,

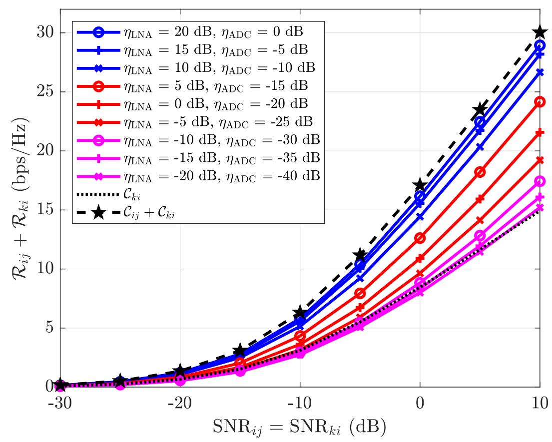

Figure 3.3: Sum spectral efficiency as a function of SNR for various $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB, where $\kappa = 10$ dB, $b = 12$ bits, and $K_{ij} = K_{ki} = 3$. While not obvious from this figure, $\mathcal{R}_{ki}$ is fairly robust to self-interference with $b = 12$ bits. Stricter choices of $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB degrade $\mathcal{R}_{ij}$ and, thus, the sum $\mathcal{R}_{ij} + \mathcal{R}_{ki}$.

the receive link sees little sacrifice as a function of $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB, having assumed the resolution of the ADC is $b = 12$ bits, which is fairly robust to these levels of $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB (more on this later). Therefore, the lower bound on $\mathcal{R}_{ij} + \mathcal{R}_{ki}$, as one would hope, is approximately the half-duplex receive capacity $\mathcal{C}_{ki}$. The small gap when $\mathcal{R}_{ij} + \mathcal{R}_{ki} < \mathcal{C}_{ki}$ at very low $\eta_{\mathrm{LNA}} = \eta_{\mathrm{ADC}} + 20$ dB in Figure 3.3 can be attributed to a small degree of quantization noise and, more significantly, the fact that the analog beamforming candidate chosen from $\mathcal{T}_{ki}$ is not always the $\mathcal{C}_{ki}$-achieving one, given that we have $K_{ki} = 3$. In other words, by design, the candidate from $\mathcal{T}_{ki}$ that maximizes transmit link performance subject to our constraints may not be the one that maximizes receive link performance.

Figure 3.4 exhibits the gains in spectral efficiency afforded by increas-

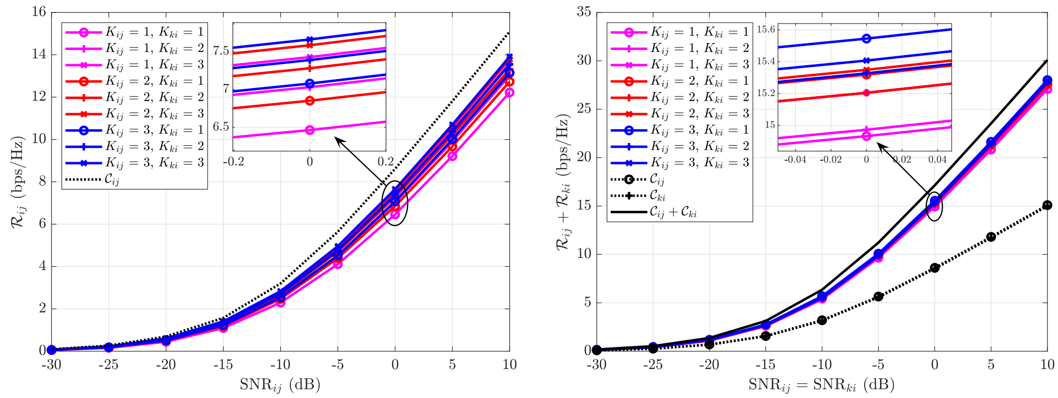

(a) Transmit link spectral efficiency.

(b) Sum spectral efficiency.

Figure 3.4: Spectral efficiency as a function of SNR for various K_{ij} and K_{ki}, where $\kappa = 10$ dB, $\eta_{\text{LNA}} = 15$ dB, $\eta_{\text{ADC}} = -5$ dB, and $b = 12$ bits. Increasing K_{ij} and K_{ki} can only improve $\mathcal{R}_{ij}$, whereas increasing K_{ki} may degrade $\mathcal{R}_{ki}$. The choice of K_{ij} and K_{ki} that maximizes the $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ varies with each realization.

ing the number of analog beamforming candidates to our design. **By increasing K_{ij} and K_{ki}, the system can improve its performance on the *transmit* link while meeting the per-antenna and per-RF chain constraints.** Referring to Figure 3.4a, when K_{ij} and K_{ki} increases from having only one candidate ($K_{ij} = K_{ki} = 1$) to having three candidates on each link ($K_{ij} = K_{ki} = 3$), we see a gain of approximately 1.25 bps/Hz in $\mathcal{R}_{ij}$ on average at $\mathsf{SNR}_{ij} = 0$ dB. This can be attributed to the fact that widening the search space will yield greater flexibility in meeting the constraints while maximizing performance on the transmit link. That is, rather than our optimization problem taking place over only $\mathbf{F}_{\text{BB}}^{(i)}$, it also takes place over the candidates in $\mathcal{T}_{ij}$ and $\mathcal{T}_{ki}$. Interestingly, we can see that, on average, supplying our design with increased K_{ki} has a relatively greater impact than K_{ij}; this can be seen by the fact that $K_{ij} = 1, K_{ki} = 3$ outperforms $K_{ij} = 3, K_{ki} = 1$ and even $K_{ij} = 3, K_{ki} = 2$ in terms of $\mathcal{R}_{ij} + \mathcal{R}_{ki}$. With increased K_{ki}, some sacrifices may be made on the receive link by choosing a candidate from $\mathcal{T}_{ki}$ that is not the $\mathcal{C}_{ki}$-achieving one. In general, K_{ij} and K_{ki} can be chosen to throttle performance between the transmit link and receive link. Choosing a small K_{ki}, for example, preserves the receive link but reduces the full-duplex device's flexibility in avoiding self-interference. For a fixed K_{ki}, choosing a large K_{ij} can generally only help the system by increasing $\mathcal{R}_{ij}$, though, increasing the number of candidates (K_{ij} or K_{ki}) adds to the overhead associated with our design. It is difficult to draw generalized conclusions on which beam pairs work well or even precisely why since this depends so heavily on the transmit and receive channels, the self-interference channel, the beam codebooks, η_{LNA}, η_{ADC}, and a number of other factors; such would be interesting future work.

In Figure 3.5, we evaluate the sum spectral efficiency $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ for various selections of η_{LNA} and η_{ADC}. To better illustrate this, we let $K_{ij} = K_{ki} = 1$, placing the sole responsibility of preventing saturation on $\mathbf{F}_{\mathrm{BB}}^{(i)}$ and making $\mathcal{R}_{ki}$ approximately constant (≈ 3.2 bps/Hz) across all η_{LNA}, η_{ADC} (since $b = 12$ bits). Intuitively, as η_{LNA} and η_{ADC} increase, optimizing transmission from i becomes more relaxed, allowing for higher $\mathcal{R}_{ij}$. For a given choice of η_{LNA}, we can see that at increasing η_{ADC} beyond some point has little to no effect on $\mathcal{R}_{ij}$. This can be attributed to the fact that **the LNA constraint begins to supersede the ADC constraint** at these points (recall Theorem 3.4). As η_{LNA} increases, the point at which η_{ADC} plays no role also increases.

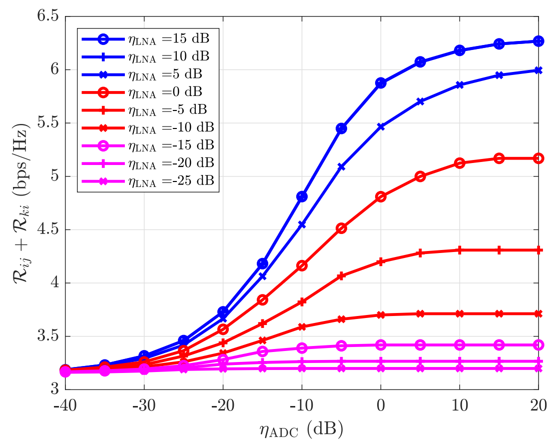

Figure 3.5: Sum spectral efficiency as a function of η_{ADC} for various η_{LNA}, where $\kappa = 10$ dB, $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $b = 12$ bits. As η_{LNA} and η_{ADC} are relaxed, the system can achieve a greater sum spectral efficiency since transmit performance is less constrained by a limited receive dynamic range.

Furthermore, beyond a certain η_{LNA} (e.g., $\eta_{\mathrm{LNA}} \geq 10$ dB), we can see that the LNA constraint becomes immaterial, suggesting that the precoding power constraint implicitly satisfies the LNA constraint (recall Theorem 3.2).

Now, we examine the importance of η_{ADC} for various ADC resolutions. A key motivator for this work is the fact that self-interference can increase quantization noise, degrading the effective SNR of a desired signal out of the ADC. ADCs having greater resolution have a higher dynamic range, allowing them to quantize signal-plus-interference-plus-noise without suffering from ADC saturation as severely as lower-resolution ADCs. This can be seen in Figure 3.6a, where we have evaluated $\mathcal{R}_{ki}$ as a function of η_{ADC} for various ADC resolutions, taking $\eta_{\mathrm{LNA}} = 20$ dB and $K_{ij} = K_{ki} = 1$ to reduce their impacts on interpreting these results. Higher-resolution ADCs are practically

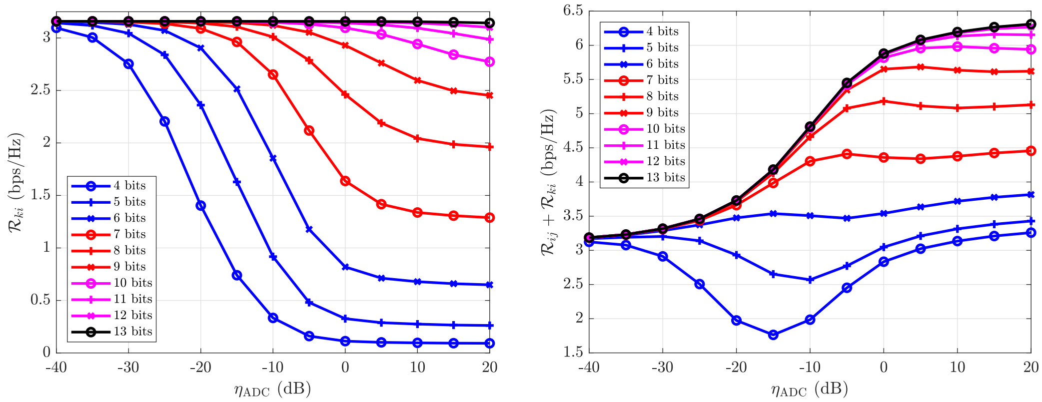

(a) Receive link spectral efficiency.

(b) Sum spectral efficiency.

Figure 3.6: Spectral efficiency as a function of η_{ADC} for various ADC resolutions, where $\kappa = 10$ dB, $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $\eta_{\mathrm{LNA}} = 20$ dB. Higher-resolution ADCs are more robust to self-interference whereas lower-resolution ADCs saturate if η_{ADC} is not properly chosen. The tradeoff associated with constraining transmit link performance and preventing ADC saturation is not so obvious in terms of sum spectral efficiency.

invariant across η_{ADC}, allowing them to achieve approximately the same $\mathcal{R}_{ki}$ regardless of the relative self-interference power at the ADCs. As the resolution decreases, we can see that quantization noise begins to take its toll on the spectral efficiency $\mathcal{R}_{ki}$, where it eventually plateaus beyond a certain η_{ADC} as the other constraints take effect. This highlights that an appropriate choice of η_{ADC} is intimately connected with the resolution of the ADCs and further justifies the motivation for this work: **under limited ADC resolution, limiting the self-interference power reaching the ADCs is critical**.

Interesting things happen in terms of the sum spectral efficiency $\mathcal{R}_{ij} + \mathcal{R}_{ki}$, with varying η_{ADC} and ADC resolutions, as depicted in Figure 3.6b. As discussed, high-resolution ADCs are relatively invariant to η_{ADC}, and therefore, changes in sum spectral efficiency can be attributed almost exclusively to changes in $\mathcal{R}_{ij}$ as a function of η_{ADC}. With low-resolution ADCs (e.g., $b = 4, 5$ bits), we see that $\mathcal{R}_{ij} + \mathcal{R}_{ki}$ initially decreases sharply as η_{ADC} increases. This is due to the falloff that we saw in Figure 3.6a. At very low (strict) η_{ADC}, $\mathcal{R}_{ij}$ is also very low, meaning the sharp falloff in $\mathcal{R}_{ki}$ drastically degrades the sum spectral efficiency. As η_{ADC} is increased (e.g., $\eta_{\text{ADC}} = -10$ dB), $\mathcal{R}_{ij}$ also increases while $\mathcal{R}_{ki}$ begins to plateau. Recall that $\mathcal{R}_{ij}$ is invariant to the ADC resolution at i. As η_{ADC} further increases (e.g., $\eta_{\text{ADC}} = 0$ dB), $\mathcal{R}_{ij}$ further increases and more rapidly so as its ADC saturation requirements become even more relaxed whereas $\mathcal{R}_{ki}$ further plateaus. Finally, as η_{ADC} increases further (e.g., $\eta_{\text{ADC}} \geq 10$ dB), $\mathcal{R}_{ki}$ remains plateaued and $\mathcal{R}_{ij}$ begins to also plateau as it sees less gain in $\mathcal{R}_{ij}$ as changes in η_{ADC} hold less meaning, given the presence of $\eta_{\text{LNA}} = 20$ dB and the precoding power constraint. For ADCs falling between very high resolutions and very low resolutions, the behavior can be explained in a similar fashion by this intertwining of $\mathcal{R}_{ij}$ and $\mathcal{R}_{ki}$, both of which begin to saturate beyond a certain η_{ADC}.

Figure 3.6b highlights an important fact: **an appropriate design is necessary to make full-duplex operation worthwhile over half-duplex**. This is evidenced by the fact that low-resolution ADCs (e.g., $b = 4, 5$ bits) demand such significant self-interference mitigation that the sacrifice made on the transmit link is not worthwhile. Furthermore, we can see that the $(\mathcal{R}_{ij} + \mathcal{R}_{ki})$-optimal degree of self-interference power permitted at the ADCs varies with resolution. At these optimal choices for η_{ADC}, introducing a modest

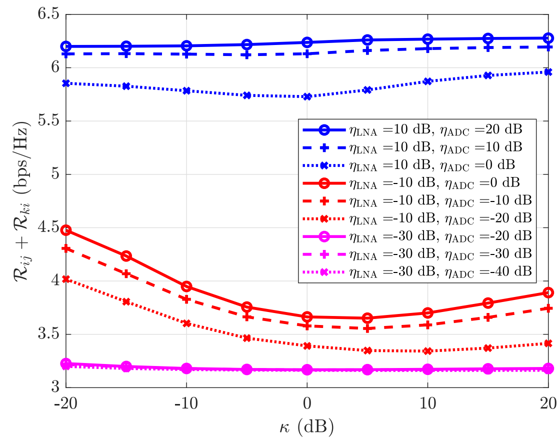

Figure 3.7: Sum spectral efficiency as a function of κ for various $\eta_{\rm LNA}$ and $\eta_{\rm ADC}$, where $K_{ij} = K_{ki} = 1$, $\mathsf{SNR}_{ij} = \mathsf{SNR}_{ki} = -10$ dB, and $b = 12$ bits.

amount of quantization noise to improve transmit link performance balances $\mathcal{R}_{ij}$ and $\mathcal{R}_{ki}$ such that their sum is maximized.

Finally, for the sake of completeness and to better understand the role κ (the self-interference channel's Rician factor) from (3.60) plays in our design, we have included Figure 3.7. We have evaluated the sum spectral efficiency for various pairs of $\eta_{\rm LNA}$ and $\eta_{\rm ADC}$ as a function of κ, fixing all other variables. Under relaxed conditions (i.e., high $\eta_{\rm LNA}$, $\eta_{\rm ADC}$), avoiding the self-interference channel becomes less of a priority, rendering κ less impactful. Under stringent conditions (i.e., low $\eta_{\rm LNA}$, $\eta_{\rm ADC}$), we can also see that κ does not play much of a role. This can be attributed to the fact that satisfying these strict LNA and ADC constraints is done so largely by power control of $\mathbf{F}_{\rm BB}^{(i)}$ rather than through steering strategies. This renders the underlying structure of $\mathbf{H}_{ii}$, and thus κ, less of a factor. In between, however, we see that modest choices of

η_{LNA} and η_{ADC} makes the role of κ more significant. When κ is very low, the self-interference channel is comprised primarily of far-field reflections. This leads to a self-interference channel that is spatially sparse, meaning avoiding $\mathbf{H}_{ii}$ becomes easier with highly directional DFT beams. This can be similarly stated that the inherent isolation between the rays of $\mathbf{H}_{ij}$ and $\mathbf{H}_{ii}$ and of $\mathbf{H}_{ki}$ and $\mathbf{H}_{ii}$ leads to more easily avoiding pushing self-interference onto the receiver of our full-duplex device. Similar behavior happens when κ is high, though the spatial sparsity of the self-interference channel stems from the near-field channel structure produced by our vertically-separated horizontal uniform linear arrays at the full-duplex device. Between, when κ approaches zero, the two spatially sparse channel components—the far-field portion and the near-field portion—mix relatively evenly, which leads to a self-interference channel that is less spatially sparse, making it more difficult to avoid pushing energy into with highly directional DFT beams. This leads to a relatively lower $\mathcal{R}_{ij}$ and, thus, lower sum spectral efficiency.

3.7 Conclusion

We have presented a hybrid beamforming design for mmWave full-duplex that holistically covers a number of practical considerations including codebook-based analog beamforming and beam alignment, a desirably low number of RF chains, and the need to prevent receiver-side saturation at the full-duplex device. Core to our design is its focus on limiting the self-interference power reaching each antenna and each RF chain to prevent saturating LNAs and ADCs. Our design utilizes sets of candidate analog beamformers to improve its flexibility in mitigating self-interference while maintaining service and to accommodate codebook-based analog beamforming. Numerical results highlight the costs and limitations associated with preventing receiver-side saturation, which can be used in system analyses when developing a mmWave full-duplex transceiver and determining what levels of self-interference mitigation should be aimed for. Potential future work includes prototyping mmWave full-duplex, characterization of mmWave self-interference, and integrating full-duplex into mmWave cellular standards.

Chapter 4

Analog Beamforming Codebooks for Full-Duplex Millimeter Wave Systems

In the previous chapter, we developed a hybrid digital/analog beamforming design to enable full-duplex mmWave communication systems. In this chapter, we introduce a design that instead relies solely on *analog* beamforming to enable full-duplex mmWave systems. This is largely motivated by the fact that analog-only beamforming is common in today's real-world mmWave systems. Rather than design particular transmit and receive beams to enable full-duplex, however, we instead design entire codebooks (sets) of beams that can be used for full-duplex operation; this allows the system to accommodate beam alignment and circumvents the need for downlink/uplink MIMO channel knowledge, both of which are practically desirable. It is important to note that the analog beamforming codebooks designed in this chapter could be directly used to enhance the analog beamforming stage of the hybrid beamforming design presented in the previous chapter.

4.1 Introduction

Codebook-based analog beamforming is a critical component of mmWave communication systems [39, 60]. Rather than estimate a high-dimensional MIMO channel and subsequently configure an analog beamforming network, modern mmWave systems instead rely on beam alignment procedures to identify promising beamforming directions, typically via exploration of a codebook of candidate beams [39, 60, 61]. This offers a simple and robust way to configure an analog beamforming network without downlink/uplink MIMO channel knowledge *a priori*, which is expensive to obtain in practice. In conventional half-duplex mmWave systems, it is relatively straightforward to design a beamforming codebook that delivers coverage over a desired region [61]. For instance, conjugate beamforming and DFT beamforming are simple ways to construct a set of beams that steer high gain in specific directions using digitally-controlled phase shifters [88]. Other codebook designs make use of amplitude control via digitally-controlled attenuators or VGAs to shape beams with suppressed side lobes in an effort to reduce interference [90–92].

The design of beamforming codebooks for in-band full-duplex mmWave communication systems, on the other hand, is not so straightforward. A full-duplex mmWave transceiver with separate transmit and receive arrays, for instance, will juggle a transmit beam and a receive beam at the same time over the same frequency spectrum, potentially serving two different half-duplex users. These transmit and receive beams couple over the MIMO channel that manifests between the arrays, plaguing a desired receive signal with self-interference. In a recent measurement campaign [44], we observed that the degree of this self-interference is typically prohibitively high for full-duplex operation in practical mmWave systems. This, along with the fact that modern mmWave systems rely on codebook-based beam alignment, motivates the need for beamforming codebooks designed specifically for full-duplex mmWave systems—as opposed to using conventional codebooks made for half-duplex systems [20]. In this work, we present LoneSTAR, the first framework for designing beamforming codebooks for full-duplex mmWave communication systems. LoneSTAR produces transmit and receive codebooks that deliver high beamforming gain and provide broad coverage to downlink and uplink users while coupling low self-interference between the transmit and receive beams of a full-duplex transceiver. By supporting beam alignment and reducing self-interference to levels near or below the noise floor regardless of which transmit and receive beams are selected from its codebooks, LoneSTAR can serve as a potential standalone solution for simultaneous transmission and reception, from which it derives its name.

Rather than using analog and digital self-interference cancellation to enable full-duplex as was commonly done in sub-6 GHz systems [6], recent work has proposed solely using beamforming to mitigate self-interference in mmWave communication systems [3, 20, 41, 68–77, 93]. Such methods largely aim to take advantage of the high-dimensional spatial domain afforded by dense mmWave antenna arrays to tackle self-interference via strategic design of transmit and receive beamformers. However, there are key practical hurdles neglected by existing designs that we address in this work holistically. Some designs [68–77] have neglected codebook-based analog beamforming, assuming the ability to fine-tune each phase shifter in real-time, and do not account for digitally-controlled phase shifters of practical analog beamforming networks. Most designs have assumed a lack of amplitude control, even though it is not uncommon to have both digitally-controlled phase shifters and attenuators or VGAs in practice [57–59]. In addition, existing solutions [68–73, 77] have assumed real-time knowledge of the downlink and uplink MIMO channels—along with that of the self-interference MIMO channel—to configure transmit and receive beams. As such, they neglect beam alignment and rely on frequent estimation of high-dimensional MIMO channels, both of which are practical shortcomings. Furthermore, the majority of existing designs are for hybrid digital/analog beamforming architectures and are not necessarily suitable for analog-only beamforming systems, which are common in practice for power and cost efficiency.

To our knowledge, the only existing beamforming-based full-duplex solution suitable for analog beamforming mmWave systems that supports codebook-based beam alignment is our recent work called STEER [3]. The work herein and that of STEER tackle a very similar problem but their approaches differ significantly due to differences in the assumption of self-interference channel knowledge. In STEER, we relied on measurements of candidate beam pairs to strategically select transmit and receive beams that couple low self-interference while maintaining high gain in desired directions. In this work, rather than use measurements of self-interference, we instead rely on (imperfect) estimation of the self-interference MIMO channel, which LoneSTAR uses to construct transmit and receive codebooks such that every transmit-receive beam pair couples low self-interference. Steer was validated using an actual 28 GHz phased array platform and proved to be a promising full-duplex solution, but it remains unclear if Steer's success will translate to other platforms. Presented herein, LoneSTAR is more flexible to the particular self-interference channel than Steer since it makes use of more degrees of freedom, meaning it has the potential to better generalize across mmWave platforms.

LoneSTAR: Beamforming codebooks for full-duplex mmWave systems. In this work, we present LoneSTAR, a novel approach to enable full-duplex mmWave systems through the use of analog beamforming codebooks strategically designed to reduce self-interference and simultaneously deliver high beamforming gain. We present an optimization framework for designing transmit and receive codebooks that minimize the self-interference coupled between all transmit-receive beam pairs while constraining the coverage they provide over defined service regions. A full-duplex mmWave transceiver can independently select any transmit beam and any receive beam from LoneSTAR codebooks and, with reasonable confidence, can expect low self-interference. LoneSTAR codebooks are designed at the full-duplex transceiver following estimation of the self-interference MIMO channel (which need not be perfect), do not require downlink/uplink MIMO channel knowledge, and do not demand any over-the-air feedback to/from users. After their construction, LoneSTAR codebooks can be used to conduct beam alignment and serve any downlink-uplink user pair in a full-duplex fashion thereafter with reduced self-interference. Importantly, LoneSTAR accounts for limited phase and amplitude control present in practical analog beamforming networks.

Validation of LoneSTAR as a full-duplex solution. We use extensive simulation to evaluate LoneSTAR, which highlights its potential to enable full-duplex operation in mmWave systems without any additional self-interference cancellation. We compare full-duplex performance with LoneSTAR codebooks against that with conventional codebooks to illustrate that LoneSTAR can supply an additional 10–50 dB of robustness to self-interference, depending on hardware constraints, without sacrificing high beamforming gain. This translates to impressive spectral efficiency gains that, with high-resolution phase shifters and attenuators, can approach the interference-free capacity of conventional codebooks. Simulation further illustrates LoneSTAR's robustness to self-interference channel estimation error and to cross-link interference between downlink and uplink users. Our results motivate a variety of future work including self-interference channel modeling and reliable self-interference channel estimation.

4.2 System Model

Illustrated in Figure 4.1, we consider a mmWave base station that employs separate, independently-controlled antenna arrays for transmission and reception that are used to serve a downlink user and uplink user in a full-duplex fashion (simultaneously and in-band). Note that these users may come from different distributions, depending on the array configuration (e.g., consider a sectorized base station). One may imagine full-duplex integrated access and backhaul (IAB), for instance, as an attractive application of this work in mmWave networks [32, 83, 94–96]. Let N_{t} and N_{r} be the number of transmit and receive antennas, respectively, at the arrays of the full-duplex base station. We denote as $\mathbf{a}_{\mathrm{tx}}(\theta) \in \mathbb{C}^{N_{\mathrm{t}} \times 1}$ and $\mathbf{a}_{\mathrm{rx}}(\theta) \in \mathbb{C}^{N_{\mathrm{r}} \times 1}$ its transmit and receive array response vectors, respectively, in the direction θ. These array response vectors for any θ can be computed based on the base station array geometries [88, 97], and we follow the convention $\|\mathbf{a}_{\mathrm{tx}}(\theta)\|_2^2 = N_{\mathrm{t}}$ and $\|\mathbf{a}_{\mathrm{rx}}(\theta)\|_2^2 = N_{\mathrm{r}}$. To simplify presentation, users are single-antenna devices, but this is not a necessary assumption.

We consider an analog beamforming system at the base station, though the codebook design herein could also be used by hybrid beamforming systems and even generalized to digital beamforming. Let $\mathbf{f} \in \mathbb{C}^{N_{\mathrm{t}} \times 1}$ be the beamforming vector used at the transmitter of the base station and $\mathbf{w} \in \mathbb{C}^{N_{\mathrm{r}} \times 1}$ be the beamforming vector used at its receiver. We consider analog beamforming networks equipped with phase shifters and attenuators which are digitally

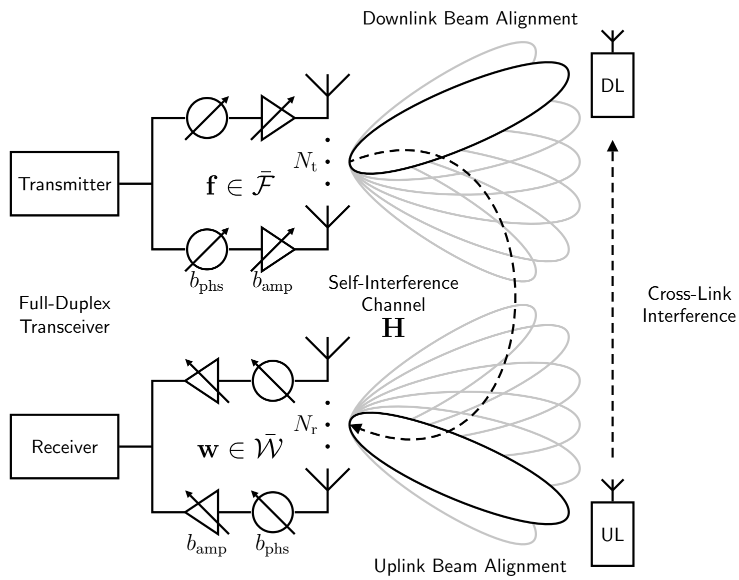

Figure 4.1: A full-duplex mmWave base station transmits to a downlink user while receiving from an uplink user in-band. In doing so, self-interference couples between its transmit beam $\mathbf{f}$ and receive beam $\mathbf{w}$ across the channel $\mathbf{H}$, and cross-link interference is inflicted on the downlink user by the uplink user. Transmit and receive beams are selected from codebooks $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ via beam alignment. In this work, we design these codebooks such that any selected $\mathbf{f} \in \bar{\mathcal{F}}$ and $\mathbf{w} \in \bar{\mathcal{W}}$ will couple low self-interference.

controlled with b_{phs} and b_{amp} bits, respectively.[1] Each beamforming weight must not exceed unit magnitude, understood by the fact that attenuators are used for amplitude control.

$$\left|[\mathbf{f}]_m\right| \leq 1 \ \forall \ m = 1, \ldots, N_{\mathrm{t}}, \qquad \left|[\mathbf{w}]_n\right| \leq 1 \ \forall \ n = 1, \ldots, N_{\mathrm{r}} \tag{4.1}$$

[1]Existing mmWave literature often considers networks with only phase shifters, but it is not uncommon for practical mmWave arrays to also be equipped with attenuators or VGAs for amplitude control [57–59]. Our formulation could also capture beamforming networks employing VGAs (rather than attenuators) by including their maximum gain settings in (4.1), for instance.

Directly from (4.1), we have $\|\mathbf{f}\|_2^2 \leq N_{\mathrm{t}}$ and $\|\mathbf{w}\|_2^2 \leq N_{\mathrm{r}}$. Capturing limited phase and amplitude control, let $\mathcal{F} \subset \mathbb{C}^{N_{\mathrm{t}} \times 1}$ and $\mathcal{W} \subset \mathbb{C}^{N_{\mathrm{r}} \times 1}$ be the sets of all possible physically realizable beamforming vectors $\mathbf{f}$ and $\mathbf{w}$, respectively; hence, it is practically required that $\mathbf{f} \in \mathcal{F}$ and $\mathbf{w} \in \mathcal{W}$. We assume the base station and uplink user transmit with powers $P_{\mathrm{tx}}^{\mathrm{BS}}$ and $P_{\mathrm{tx}}^{\mathrm{UE}}$, respectively. Additive noise powers per antenna at the base station and downlink user are respectively $P_{\mathrm{noise}}^{\mathrm{BS}}$ and $P_{\mathrm{noise}}^{\mathrm{UE}}$.

Taking the perspective of the full-duplex base station, let $\mathbf{h}_{\mathrm{tx}}^{*} \in \mathbb{C}^{1 \times N_{\mathrm{t}}}$ be the transmit channel vector from the base station to the downlink user. Let $\mathbf{h}_{\mathrm{rx}} \in \mathbb{C}^{N_{\mathrm{r}} \times 1}$ be the receive channel from the uplink user to the base station. We normalize each of these channels such that $\|\mathbf{h}_{\mathrm{tx}}\|_2^2 = N_{\mathrm{t}}$ and $\|\mathbf{h}_{\mathrm{rx}}\|_2^2 = N_{\mathrm{r}}$ and abstract out their large-scale gains as inverse path losses G_{tx}^2 and G_{rx}^2, respectively. Since transmission and reception happen simultaneously and in-band, a self-interference channel $\mathbf{H} \in \mathbb{C}^{N_{\mathrm{r}} \times N_{\mathrm{t}}}$ manifests between the transmit and receive arrays of the full-duplex base station. Currently, there is not a reliable, measurement-backed model for the self-interference channel $\mathbf{H}$ [20, 44]. As such, the design herein is not based on any assumption of $\mathbf{H}$, and we evaluate our design using multiple channel models in Section 4.7. To abstract out the large-scale gain of the self-interference channel from its spatial characteristics, we impose $\|\mathbf{H}\|_{\mathrm{F}}^2 = N_{\mathrm{t}} \cdot N_{\mathrm{r}}$ and let G^2 be its inverse path loss based on the inherent isolation between the arrays (practically, $G^2 \ll 0$ dB). With the downlink and uplink users operating simultaneously and in-band, a scalar cross-link interference channel h_{CL} arises between them. We normalize $|h_{\mathrm{CL}}|^2 = 1$ and let G_{CL}^2 be the inverse path loss of the cross-link interference channel, which practically depends on the users' locations and the environment.

4.3 Problem Formulation and Motivation

Using the established system model, we now further formulate and motivate this work. Again taking the perspective of the full-duplex base station, we refer to downlink as the *transmit link* and to uplink as the *receive link*. The signal-to-interference-plus-noise ratios (SINRs) of these two links can be expressed as

$$\mathsf{SINR}_{\mathrm{tx}} = \frac{\mathsf{SNR}_{\mathrm{tx}}}{1 + \mathsf{INR}_{\mathrm{tx}}}, \qquad \mathsf{SINR}_{\mathrm{rx}} = \frac{\mathsf{SNR}_{\mathrm{rx}}}{1 + \mathsf{INR}_{\mathrm{rx}}}, \tag{4.2}$$

each of which being dictated by its respective SNR and interference-to-noise ratio (INR). For some transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$, the SNRs of the transmit and receive links are respectively

$$\mathsf{SNR}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{BS}} \cdot G_{\mathrm{tx}}^2 \cdot |\mathbf{h}_{\mathrm{tx}}^* \mathbf{f}|^2}{N_{\mathrm{t}} \cdot P_{\mathrm{noise}}^{\mathrm{UE}}}, \qquad \mathsf{SNR}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{UE}} \cdot G_{\mathrm{rx}}^2 \cdot |\mathbf{w}^* \mathbf{h}_{\mathrm{rx}}|^2}{\|\mathbf{w}\|_2^2 \cdot P_{\mathrm{noise}}^{\mathrm{BS}}}, \tag{4.3}$$

where dividing by N_{t} in $\mathsf{SNR}_{\mathrm{tx}}$ handles power splitting within the transmit array and by $\|\mathbf{w}\|_2^2$ in $\mathsf{SNR}_{\mathrm{rx}}$ handles noise combining at the receive array. By virtue of the fact that $N^2 = \|\mathbf{h}\|_2^2 \cdot \|\mathbf{x}\|_2^2 = \max_{\mathbf{x}} \ |\mathbf{h}^* \mathbf{x}|^2$ s.t. $\|\mathbf{x}\|_2^2 \leq N, \|\mathbf{h}\|_2^2 = N$, the maximum SNRs are denoted with an overline as

$$\overline{\mathsf{SNR}}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{BS}} \cdot G_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}}{P_{\mathrm{noise}}^{\mathrm{UE}}} \geq \mathsf{SNR}_{\mathrm{tx}}, \tag{4.4}$$

$$\overline{\mathsf{SNR}}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{UE}} \cdot G_{\mathrm{rx}}^2 \cdot N_{\mathrm{r}}}{P_{\mathrm{noise}}^{\mathrm{BS}}} \geq \mathsf{SNR}_{\mathrm{rx}}. \tag{4.5}$$

In other words, these are achieved only when users are delivered maximum beamforming gain.

When juggling a transmit beam $\mathbf{f}$ and receive beam $\mathbf{w}$ in a full-duplex fashion, the base station couples self-interference, leading to a receive link INR of

$$\mathsf{INR}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{BS}} \cdot G^2 \cdot \left|\mathbf{w}^* \mathbf{H} \mathbf{f}\right|^2}{N_{\mathrm{t}} \cdot \|\mathbf{w}\|_2^2 \cdot P_{\mathrm{noise}}^{\mathrm{BS}}}, \tag{4.6}$$

where $\left|\mathbf{w}^* \mathbf{H} \mathbf{f}\right|^2$ captures the transmit and receive beam coupling over the self-interference channel. Strategically designing $\mathbf{f}$ and $\mathbf{w}$ according to $\mathbf{H}$ can therefore reduce self-interference. Based on our normalizations, the worst-case coupling between transmit and receive beams at the full-duplex base station is $N_{\mathrm{t}}^2 \cdot N_{\mathrm{r}}^2 = \max \left|\mathbf{w}^* \mathbf{H} \mathbf{f}\right|^2$, leading to a maximum possible $\mathsf{INR}_{\mathrm{rx}}$ of

$$\overline{\mathsf{INR}}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{BS}} \cdot G^2 \cdot N_{\mathrm{t}} \cdot N_{\mathrm{r}}}{P_{\mathrm{noise}}^{\mathrm{BS}}} \geq \mathsf{INR}_{\mathrm{rx}}, \tag{4.7}$$

which we use as a metric to capture the inherent strength of self-interference at the base station. Note that $\overline{\mathsf{INR}}_{\mathrm{rx}}$ depends solely on system parameters and artifacts at the base station. We denote as $\mathsf{INR}_{\mathrm{tx}}$ the INR incurred at the downlink user due to cross-link interference.

$$\mathsf{INR}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{UE}} \cdot G_{\mathrm{CL}}^2 \cdot \left|h_{\mathrm{CL}}\right|^2}{P_{\mathrm{noise}}^{\mathrm{UE}}} \tag{4.8}$$

Note that cross-link interference $\mathsf{INR}_{\mathrm{tx}}$ only depends on its inherent channel and path loss, not on $\mathbf{f}$ nor $\mathbf{w}$. In this work, we do not consider uplink power control as a means to reduce cross-link interference, though this could be useful in practice.

The achievable spectral efficiencies on each link, treating interference as noise, are

$$R_{\mathrm{tx}} = \log_2(1 + \mathsf{SINR}_{\mathrm{tx}}), \qquad R_{\mathrm{rx}} = \log_2(1 + \mathsf{SINR}_{\mathrm{rx}}), \tag{4.9}$$

whereas the capacities of the two links are $C_{\mathrm{tx}} = \log_2\left(1 + \overline{\mathsf{SNR}}_{\mathrm{tx}}\right)$ and $C_{\mathrm{rx}} = \log_2\left(1 + \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$. We assume no forms of active self-interference cancellation are employed at the base station. Instead, we will rely solely on beamforming to mitigate self-interference and enable full-duplex operation. The degree of self-interference coupled at the full-duplex base station in (4.6) depends on the transmit beam $\mathbf{f}$ and receive beam $\mathbf{w}$, each of which also dictates its respective link's SNR. As such, it is important that $\mathbf{f}$ and $\mathbf{w}$ couple low self-interference and deliver high $\mathsf{SNR}_{\mathrm{tx}}$ and $\mathsf{SNR}_{\mathrm{rx}}$ in an effort to achieve high sum spectral efficiency $R_{\mathrm{tx}} + R_{\mathrm{rx}}$.

In this work, we rely on estimation of the self-interference MIMO channel $\mathbf{H}$ at the full-duplex base station, but do not assume any knowledge of the downlink or uplink channels $\mathbf{h}_{\mathrm{tx}}$ and $\mathbf{h}_{\mathrm{rx}}$. We do not assume perfect estimation of $\mathbf{H}$ but rather model its imperfections as

$$\mathbf{H} = \bar{\mathbf{H}} + \mathbf{\Delta}, \qquad [\mathbf{\Delta}]_{i,j} \sim \mathcal{N}_{\mathbb{C}}\left(0, \epsilon^2\right) \ \forall \ i, j. \tag{4.10}$$

where $\bar{\mathbf{H}}$ is the estimate of $\mathbf{H}$ and $\mathbf{\Delta}$ captures independently and identically distributed (i.i.d.) Gaussian estimation error with mean 0 and variance ϵ^2. Self-interference channel estimation in full-duplex mmWave systems is an open research problem [20]. As such, the choice of a particular estimation error model is currently difficult to practically justify, but we remark that other error models (e.g., norm-based models) may be encompassed by our Gaussian model with arbitrarily high probability for appropriately chosen ϵ^2.

It is our hope that $\mathbf{H}$ can be estimated with high fidelity for three reasons. First, the self-interference channel is likely much stronger than traditional downlink/uplink MIMO channels, meaning it may be estimated with higher accuracy. Second, the estimation of $\mathbf{H}$ takes place across the arrays of the base station, meaning it does not consume resources for feedback and does not suffer from feedback quantization. Third, we expect that the self-interference channel will be fairly static on the timescale of data transmissions, especially the portion caused by direct coupling between the arrays[2], which will presumably dominate. Together, all of this suggests that the estimation of $\mathbf{H}$ may be more thorough and precise than that of traditional downlink/uplink MIMO channels. Nonetheless, our design in the next section incorporates channel estimation error. In our numerical evaluation in Sections 4.6–4.7, one self-interference channel model we consider is the spherical-wave model, which has been widely used in the related literature to capture near-field propagation between the transmit and receive arrays. Recent work [98] has presented a scheme for efficiently and accurately estimating such channels, especially considering the presumably high strength of self-interference $\overline{\mathsf{INR}}_{\mathrm{rx}}$. Self-interference channel estimation and its coherence time would be valuable

[2]In [44], we observed that 28 GHz self-interference in an anechoic chamber was nearly static on the order of minutes, at least.

directions for future work, especially in practical systems. The key practical advantages of LoneSTAR rely on the self-interference channel being sufficiently static over many downlink/uplink time slots, which we assume herein.

4.4 Beamforming Codebook Design

Practical mmWave systems—which have operated in a half-duplex fashion thus far—rely on beam alignment to identify beams that deliver high beamforming gain when serving a particular user. This is typically done by sweeping a set of candidate beams called a *codebook* and measuring the reference signal received power (RSRP) for each beam candidate in order to select which is used to close a link. This procedure allows a mmWave system to reliably deliver high beamforming gain without the need for downlink/uplink MIMO channel knowledge, which is expensive to obtain in practice.

As mentioned in the introduction, plenty of existing work has designed particular transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$ to use at the base station to reduce self-interference when full-duplexing downlink and uplink users, but most existing designs have a number of practical shortcomings. In this work, we instead design transmit and receive beam codebooks $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$, from which the base station can draw for $\mathbf{f}$ and $\mathbf{w}$ following conventional beam alignment. Note that $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ are not necessarily used directly for beam alignment but rather could be drawn from following whatever beam alignment procedure is performed. The goal of our design—which we call LoneSTAR—is to create codebooks that a full-duplex mmWave base station can use to deliver high beamforming gain to users while simultaneously mitigating self-interference. This will be achieved by strategically designing $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ such that their transmit and receive beams couple low self-interference, regardless of which $\mathbf{f} \in \bar{\mathcal{F}}$ and $\mathbf{w} \in \bar{\mathcal{W}}$ are selected following beam alignment. In contrast, codebooks created for conventional mmWave systems can deliver high beamforming gain but are not necessarily robust to self-interference if used in full-duplex systems [20]. In fact, this was confirmed in our recent measurement campaign [44] and in our recent work [3], both of which showed that conventional beams typically couple prohibitively high amounts of self-interference if used in a practical full-duplex mmWave system.

Our envisioned use of LoneSTAR is depicted in Figure 4.2, which illustrates that our codebook design need only be executed once per coherence time of $\mathbf{H}$, allowing the same full-duplex solution to be used over many downlink-uplink user pairs (i.e., data transmissions). This is a key practical advantage of LoneSTAR versus most existing mmWave full-duplex solutions, which we depict in Figure 4.3. Most often, existing designs tailor their solutions uniquely for each downlink-uplink user pair and warrant estimation of the downlink and uplink MIMO channels, which itself is impractical in today's systems. On the other hand, a single LoneSTAR design can serve a sequence of downlink-uplink user pairs in a full-duplex fashion without any additional overhead. Moreover, existing designs typically introduce over-the-air feedback to and from users, whereas LoneSTAR does not require this—its design taking place completely at the full-duplex transceiver. Naturally, the coherence

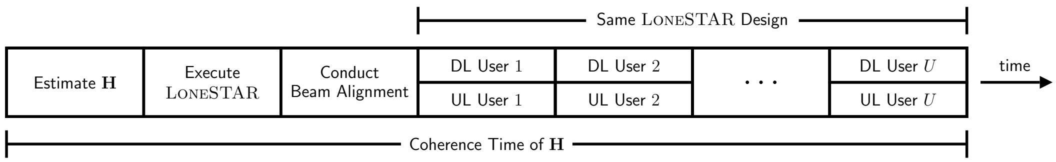

Figure 4.2: Following estimation of the self-interference channel $\mathbf{H}$, transmit and receive beam codebooks designed with LoneSTAR can be used for beam alignment and subsequently to serve any downlink and uplink users in a full-duplex fashion.

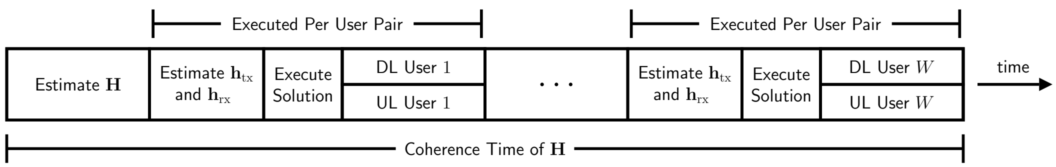

Figure 4.3: Most existing beamforming-based full-duplex mmWave solutions require estimation of the downlink and uplink channels and execution of their design for each user pair, leading to a total overhead that scales with the number of users served.

time of mmWave self-interference, which is not currently well understood, will be an important factor in justifying the use of LoneSTAR since it dictates how frequently $\mathbf{H}$ must be estimated and how often LoneSTAR needs to be re-executed. It should be noted, however, that existing designs also rely on estimation of the self-interference channel, meaning they too are subjected to its coherence time. If our design sufficiently mitigates self-interference between all beam pairs, a mmWave base station employing LoneSTAR codebooks can seamlessly operate in a full-duplex fashion without any additional self-interference cancellation—any transmit and receive beams selected via beam alignment will couple low self-interference.

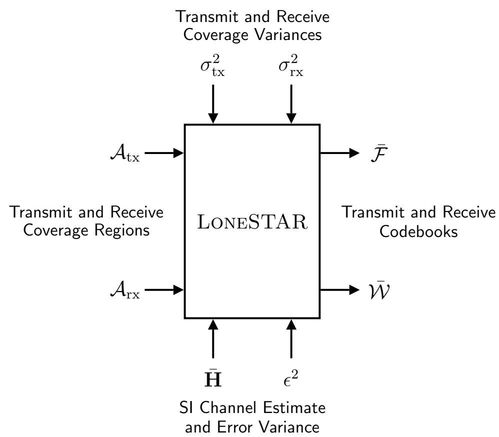

Figure 4.4: Using an estimate of the self-interference channel and knowledge of the estimation error variance, LONESTAR produces transmit and receive beamforming codebooks that serve some coverage regions with quality based on tolerated coverage variances.

4.4.1 Quantifying Our Codebook Design Criteria

In pursuit of transmit and receive beam codebooks that reliably offer high sum spectral efficiency $R_{\text{tx}} + R_{\text{rx}}$, we desire sets of transmit and receive beams that can deliver high SNR_{tx} and SNR_{rx} while coupling low self-interference INR_{rx}. As mentioned, cross-link interference INR_{tx} is fixed for a given downlink-uplink user pair since it does not depend on $\mathbf{f}$ nor $\mathbf{w}$. In the remainder of this section, we first quantify our codebook design criteria and then assemble a design problem that tackles both objectives jointly. As we proceed with outlining our design, please refer to Figure 4.4 for an illustration of the inputs and outputs of LONESTAR.

Definition 4.1: Transmit and receive coverage regions. Let $\mathcal{A}_{\text{tx}}$ be a set of M_{tx} directions—or a *transmit coverage region*—that our transmit codebook

$\bar{\mathcal{F}}$ aims to serve, and let $\mathcal{A}_{\mathrm{rx}}$ be a *receive coverage region* defined analogously.

$$\mathcal{A}_{\mathrm{tx}} = \left\{\theta_{\mathrm{tx}}^{(i)} : i = 1, \ldots, M_{\mathrm{tx}}\right\}, \qquad \mathcal{A}_{\mathrm{rx}} = \left\{\theta_{\mathrm{rx}}^{(j)} : j = 1, \ldots, M_{\mathrm{rx}}\right\} \tag{4.11}$$

Here, $\theta_{\mathrm{tx}}^{(i)}$ and $\theta_{\mathrm{rx}}^{(j)}$ each may capture azimuthal and elevational components of a given direction.

Definition 4.2: Transmit and receive beamforming codebooks. We index the M_{tx} beams in $\bar{\mathcal{F}}$ and the M_{rx} beams in $\bar{\mathcal{W}}$ according to $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$, respectively, as $\bar{\mathcal{F}} = \{\mathbf{f}_1, \ldots, \mathbf{f}_{M_{\mathrm{tx}}}\}$ and $\bar{\mathcal{W}} = \{\mathbf{w}_1, \ldots, \mathbf{w}_{M_{\mathrm{rx}}}\}$, where $\mathbf{f}_i \in \mathcal{F}$ is responsible for transmitting toward $\theta_{\mathrm{tx}}^{(i)}$ and $\mathbf{w}_j \in \mathcal{W}$ is responsible for receiving from $\theta_{\mathrm{rx}}^{(j)}$. Using this notation, we build the transmit codebook matrix $\mathbf{F}$ by stacking transmit beams $\mathbf{f}_i$ as follows and the receive codebook matrix $\mathbf{W}$ analogously.

$$\mathbf{F} = \begin{bmatrix} \mathbf{f}_1 & \mathbf{f}_2 & \cdots & \mathbf{f}_{M_{\mathrm{tx}}} \end{bmatrix} \in \mathbb{C}^{N_{\mathrm{t}} \times M_{\mathrm{tx}}}, \tag{4.12}$$

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}_1 & \mathbf{w}_2 & \cdots & \mathbf{w}_{M_{\mathrm{rx}}} \end{bmatrix} \in \mathbb{C}^{N_{\mathrm{r}} \times M_{\mathrm{rx}}} \tag{4.13}$$

Definition 4.3: Transmit and receive beamforming gains. Let $g_{\mathrm{tx}}^2\left(\theta_{\mathrm{tx}}^{(i)}\right)$ be the transmit beamforming gain afforded by $\mathbf{f}_i$ in the direction $\theta_{\mathrm{tx}}^{(i)}$ and $g_{\mathrm{rx}}^2\left(\theta_{\mathrm{rx}}^{(j)}\right)$ be the receive beamforming gain afforded by $\mathbf{w}_j$ in the direction $\theta_{\mathrm{rx}}^{(j)}$, expressed as

$$g_{\mathrm{tx}}^2\left(\theta_{\mathrm{tx}}^{(i)}\right) = \left|\mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}\right)^* \mathbf{f}_i\right|^2 \leq N_{\mathrm{t}}^2, \tag{4.14}$$

$$g_{\mathrm{rx}}^2\left(\theta_{\mathrm{rx}}^{(j)}\right) = \left|\mathbf{a}_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(j)}\right)^* \mathbf{w}_j\right|^2 \leq N_{\mathrm{r}}^2. \tag{4.15}$$

The transmit and receive beamforming gains are upper-bounded by N_{t}^2 and N_{r}^2, respectively, by virtue of normalizations presented in our system model.

Maximum beamforming gains can be achieved toward any direction θ with beams $\mathbf{f} = \mathbf{a}_{\text{tx}}(\theta)$ and $\mathbf{w} = \mathbf{a}_{\text{rx}}(\theta)$, respectively, often referred to as *conjugate beamforming* or *matched filter beamforming* [88].

Definition 4.4: Average self-interference coupled. The coupling between the i-th transmit beam $\mathbf{f}_i$ and the j-th receive beam $\mathbf{w}_j$ is captured by the product $\mathbf{w}_j^* \mathbf{H} \mathbf{f}_i$, which can be extended to encompass all transmit-receive beam pairs as $\mathbf{W}^* \mathbf{H} \mathbf{F} \in \mathbb{C}^{M_{\text{rx}} \times M_{\text{tx}}}$. LoneSTAR will aim to minimize the coupling across all beam pairs and does so by minimizing the average INR_{rx} coupled between beam pairs, which can be written as

$$\mathsf{INR}_{\text{rx}}^{\text{avg}} = \frac{1}{M_{\text{tx}} \cdot M_{\text{rx}}} \cdot \sum_{i=1}^{M_{\text{tx}}} \sum_{j=1}^{M_{\text{rx}}} \frac{P_{\text{tx}}^{\text{BS}} \cdot G^2 \cdot \left| \mathbf{w}_j^* \mathbf{H} \mathbf{f}_i \right|^2}{N_{\text{t}} \cdot \left\| \mathbf{w}_j \right\|_2^2 \cdot P_{\text{noise}}^{\text{BS}}} \tag{4.16}$$

$$\approx \frac{P_{\text{tx}}^{\text{BS}} \cdot G^2}{N_{\text{t}} \cdot N_{\text{r}} \cdot P_{\text{noise}}^{\text{BS}}} \cdot \frac{\left\| \mathbf{W}^* \mathbf{H} \mathbf{F} \right\|_{\text{F}}^2}{M_{\text{tx}} \cdot M_{\text{rx}}}, \tag{4.17}$$

where the approximation becomes equality when $\left\| \mathbf{w}_j \right\|_2^2 = N_{\text{r}}$ for all j. Minimizing $\mathsf{INR}_{\text{rx}}^{\text{avg}}$ can be accomplished (approximately) by minimizing $\left\| \mathbf{W}^* \mathbf{H} \mathbf{F} \right\|_{\text{F}}^2$ through the design of $\mathbf{F}$ and $\mathbf{W}$ for a given $\mathbf{H}$ since all other terms are constants. In LoneSTAR's attempt to minimize $\mathsf{INR}_{\text{rx}}^{\text{avg}}$, it must also ensure that the beams in $\mathbf{F}$ and $\mathbf{W}$ remain useful in serving users; of course, $\mathbf{F}$ and $\mathbf{W}$ minimize $\mathsf{INR}_{\text{rx}}^{\text{avg}}$ when driven to $\mathbf{0}$ if unconstrained. We now quantify a means to constrain LoneSTAR to ensure adequate service is maintained in its aim to minimize $\mathsf{INR}_{\text{rx}}^{\text{avg}}$.

Definition 4.5: Transmit and receive coverage variances. To offer flexibility in minimizing $\mathsf{INR}_{\text{rx}}^{\text{avg}}$, LoneSTAR tolerates some sacrifices in beamforming gain over the transmit and receive coverage regions. In other words, we allow for transmit and receive beamforming gains less than N_{t}^2 and N_{r}^2 over $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$ in pursuit of reduced self-interference. Let $\sigma_{\mathrm{tx}}^2 \geq 0$ be the maximum variance tolerated in delivering maximum transmit gain across $\mathcal{A}_{\mathrm{tx}}$, normalized to N_{t}^2. Defining $\sigma_{\mathrm{rx}}^2 \geq 0$ analogously, we express these tolerated coverage variances as

$$\frac{1}{M_{\mathrm{tx}}} \cdot \sum_{i=1}^{M_{\mathrm{tx}}} \frac{\left| N_{\mathrm{t}} - g_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}\right) \right|^2}{N_{\mathrm{t}}^2} \leq \sigma_{\mathrm{tx}}^2, \tag{4.18}$$

$$\frac{1}{M_{\mathrm{rx}}} \cdot \sum_{j=1}^{M_{\mathrm{rx}}} \frac{\left| N_{\mathrm{r}} - g_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(j)}\right) \right|^2}{N_{\mathrm{r}}^2} \leq \sigma_{\mathrm{rx}}^2. \tag{4.19}$$

Here, σ_{tx}^2 and σ_{rx}^2 are design parameters that can be tuned by system engineers; increasing them will offer LoneSTAR greater flexibility in reducing $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{avg}}$ but may degrade coverage since it permits more variability in beamforming gain over the service regions.

4.4.2 Assembling Our Codebook Design Problem

In the previous subsection, we outlined that LoneSTAR will attempt to minimize the average self-interference coupled by transmit and receive beams while constraining the variance of beamforming gain delivered to specified transmit and receive coverage regions. With expressions for our codebook design criteria in hand, we now turn our attention to assembling a formal design problem. Let $\mathbf{A}_{\mathrm{tx}} \in \mathbb{C}^{N_{\mathrm{t}} \times M_{\mathrm{tx}}}$ be the matrix of transmit array response vectors evaluated at the directions in $\mathcal{A}_{\mathrm{tx}}$, and let $\mathbf{A}_{\mathrm{rx}} \in \mathbb{C}^{N_{\mathrm{r}} \times M_{\mathrm{rx}}}$ be defined analogously.

$$\mathbf{A}_{\mathrm{tx}} = \left[\mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(1)}\right) \quad \mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(2)}\right) \quad \cdots \quad \mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(M_{\mathrm{tx}})}\right)\right] \tag{4.20}$$

$$\mathbf{A}_{\mathrm{rx}} = \left[\mathbf{a}_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(1)}\right) \quad \mathbf{a}_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(2)}\right) \quad \cdots \quad \mathbf{a}_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(M_{\mathrm{rx}})}\right)\right] \tag{4.21}$$

Recall, these can be computed directly based on the transmit and receive array geometries [88,97]. Using these, the expressions of (4.18) can be written equivalently (up to arbitrary phase shifts of $\mathbf{f}_i$ and $\mathbf{w}_j$) as follows, which we refer to as *coverage constraints*.

$$\left\| N_{\mathrm{t}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{tx}}^{*}\mathbf{F}) \right\|_2^2 \leq \sigma_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}^2 \cdot M_{\mathrm{tx}} \tag{4.22}$$

$$\left\| N_{\mathrm{r}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{rx}}^{*}\mathbf{W}) \right\|_2^2 \leq \sigma_{\mathrm{rx}}^2 \cdot N_{\mathrm{r}}^2 \cdot M_{\mathrm{rx}} \tag{4.23}$$

Note that the (i,i)-th entry of $\mathbf{A}_{\mathrm{tx}}^{*}\mathbf{F}$ has magnitude $g_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}\right)$ and the (j,j)-th entry of $\mathbf{A}_{\mathrm{rx}}^{*}\mathbf{W}$ has magnitude $g_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(j)}\right)$. By satisfying (4.22) and (4.23), engineers can ensure that LoneSTAR codebooks adequately serve their coverage regions for appropriately chosen σ_{tx}^2 and σ_{rx}^2.

With coverage constraints (4.22) and (4.23), we now formulate an analog beamforming codebook design in problem (4.24). The goal of LoneSTAR is to design transmit and receive codebooks that minimize the self-interference coupled between each transmit and receive beam pair while ensuring the codebooks remain useful in delivering service to users.

$$\min_{\mathbf{F},\mathbf{W}} \quad \|\mathbf{W}^*\mathbf{H}\mathbf{F}\|_{\mathrm{F}}^2 \tag{4.24a}$$

$$\text{s.t.} \quad \|N_{\mathrm{t}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{tx}}^*\mathbf{F})\|_2^2 \leq \sigma_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}^2 \cdot M_{\mathrm{tx}} \tag{4.24b}$$

$$\|N_{\mathrm{r}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{rx}}^*\mathbf{W})\|_2^2 \leq \sigma_{\mathrm{rx}}^2 \cdot N_{\mathrm{r}}^2 \cdot M_{\mathrm{rx}} \tag{4.24c}$$

$$[\mathbf{F}]_{:,i} \in \mathcal{F} \ \forall \ i = 1, \ldots, M_{\mathrm{tx}} \tag{4.24d}$$

$$[\mathbf{W}]_{:,j} \in \mathcal{W} \ \forall \ j = 1, \ldots, M_{\mathrm{rx}} \tag{4.24e}$$

By minimizing the Frobenius norm in the objective, our designed codebooks aim to minimize the average self-interference coupled across all transmit-receive beam pairs as evident in (4.17). While doing so, our codebooks must satisfy the coverage constraints (4.24b) and (4.24c) that were established in (4.22) and (4.23). Constraints (4.24d) and (4.24e) require our design to create codebooks whose beams are physically realizable with digitally-controlled phase shifters and attenuators.

Practical systems cannot solve problem (4.24) since they do not have knowledge of the true self-interference channel $\mathbf{H}$ but rather the estimate $\bar{\mathbf{H}}$. This motivates us to modify problem (4.24) to incorporate robustness to channel estimation error. It is common in robust communication system optimization to minimize the worst-case error inflicted by some uncertainty [99–102]. In our case, $\mathbf{\Delta}$ has i.i.d. Gaussian entries, meaning the worst-case error inflicted on the objective of problem (4.24) is unbounded. In light of this, we instead minimize the expectation of our objective over the distribution of $\mathbf{H}$. Forming problem (4.25), LoneSTAR aims to minimize the expected average self-interference coupled between beam pairs, given some estimate $\bar{\mathbf{H}}$.

$$\min_{\mathbf{F},\mathbf{W}} \ \mathop{\mathbb{E}}_{\mathbf{H}}\left[\|\mathbf{W}^*\mathbf{H}\mathbf{F}\|_{\mathrm{F}}^2 \mid \bar{\mathbf{H}}\right] \tag{4.25a}$$

$$\text{s.t.} \ \ \|N_{\mathrm{t}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{tx}}^*\mathbf{F})\|_2^2 \leq \sigma_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}^2 \cdot M_{\mathrm{tx}} \tag{4.25b}$$

$$\|N_{\mathrm{r}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{rx}}^*\mathbf{W})\|_2^2 \leq \sigma_{\mathrm{rx}}^2 \cdot N_{\mathrm{r}}^2 \cdot M_{\mathrm{rx}} \tag{4.25c}$$

$$[\mathbf{F}]_{:,i} \in \mathcal{F} \ \forall \ i = 1, \ldots, M_{\mathrm{tx}} \tag{4.25d}$$

$$[\mathbf{W}]_{:,j} \in \mathcal{W} \ \forall \ j = 1, \ldots, M_{\mathrm{rx}} \tag{4.25e}$$

To transform problem (4.25) into one more readily solved, we derive a closed-form expression of the objective (4.25a), courtesy of the fact that $\boldsymbol{\Delta}$ is Gaussian.

Theorem 4.1. *The objective of problem* (4.25) *is equivalent to*

$$\mathop{\mathbb{E}}_{\mathbf{H}}\left[\|\mathbf{W}^*\mathbf{H}\mathbf{F}\|_{\mathrm{F}}^2 \mid \bar{\mathbf{H}}\right] = \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \epsilon^2 \cdot \|\mathbf{F}\|_{\mathrm{F}}^2 \cdot \|\mathbf{W}\|_{\mathrm{F}}^2. \tag{4.26}$$

Proof. We begin by recalling (4.10) to write

$$\mathop{\mathbb{E}}_{\mathbf{H}}\left[\|\mathbf{W}^*\mathbf{H}\mathbf{F}\|_{\mathrm{F}}^2 \mid \bar{\mathbf{H}}\right] = \mathop{\mathbb{E}}_{\boldsymbol{\Delta}}\left[\left\|\mathbf{W}^*(\bar{\mathbf{H}}+\boldsymbol{\Delta})\mathbf{F}\right\|_{\mathrm{F}}^2\right] \tag{4.27}$$

$$\overset{(a)}{=} \mathop{\mathbb{E}}_{\boldsymbol{\Delta}}\left[\mathrm{tr}\big(\big(\mathbf{W}^*(\bar{\mathbf{H}}+\boldsymbol{\Delta})\mathbf{F}\big)\big(\mathbf{W}^*(\bar{\mathbf{H}}+\boldsymbol{\Delta})\mathbf{F}\big)^*\big)\right] \tag{4.28}$$

$$\overset{(b)}{=} \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \mathrm{tr}\left(\mathbf{F}^* \mathop{\mathbb{E}}_{\boldsymbol{\Delta}}[\boldsymbol{\Delta}^*\mathbf{W}\mathbf{W}^*\boldsymbol{\Delta}]\mathbf{F}\right) \tag{4.29}$$

$$\overset{(c)}{=} \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \mathrm{tr}\left(\mathbf{F}^* \mathop{\mathbb{E}}_{\boldsymbol{\Delta}}[\boldsymbol{\Delta}^*\mathbf{U}_{\mathbf{W}}\boldsymbol{\Lambda}_{\mathbf{W}}\mathbf{U}_{\mathbf{W}}^*\boldsymbol{\Delta}]\mathbf{F}\right) \tag{4.30}$$

$$\overset{(d)}{=} \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \mathrm{tr}\left(\mathbf{F}^* \mathop{\mathbb{E}}_{\tilde{\boldsymbol{\Delta}}}\left[\tilde{\boldsymbol{\Delta}}\boldsymbol{\Lambda}_{\mathbf{W}}\tilde{\boldsymbol{\Delta}}^*\right]\mathbf{F}\right) \tag{4.31}$$

$$\overset{(e)}{=} \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \mathrm{tr}(\mathbf{F}^*\mathbf{F}) \cdot \epsilon^2 \cdot \mathrm{tr}(\boldsymbol{\Lambda}_{\mathbf{W}}) \tag{4.32}$$

$$\overset{(f)}{=} \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \epsilon^2 \cdot \|\mathbf{F}\|_{\mathrm{F}}^2 \cdot \|\mathbf{W}\|_{\mathrm{F}}^2 \tag{4.33}$$

where (a) follows from $\|\mathbf{A}\|_{\mathrm{F}}^2 = \mathrm{tr}(\mathbf{A}\mathbf{A}^*)$; (b) is by the linearity of expectation and trace and the fact that $\mathbb{E}[\boldsymbol{\Delta}] = \mathbf{0}$; (c) is by taking the eigenvalue decomposition $\mathbf{W}\mathbf{W}^* = \mathbf{U}_{\mathbf{W}}\boldsymbol{\Lambda}_{\mathbf{W}}\mathbf{U}_{\mathbf{W}}^*$; (d) is by defining $\tilde{\boldsymbol{\Delta}} = \boldsymbol{\Delta}^*\mathbf{U}_{\mathbf{W}}$, which has the same distribution as $\boldsymbol{\Delta}$ by unitary invariance of i.i.d. zero-mean Gaussian matrices [88]; (e) is via the property [103]

$$\underset{\mathbf{X}}{\mathbb{E}}\left[\mathbf{X}\mathbf{A}\mathbf{X}^*\right] = \epsilon^2 \cdot \mathrm{tr}(\mathbf{A}^*) \cdot \mathbf{I}, \tag{4.34}$$

where $\mathbf{X}$ has entries i.i.d. Gaussian with mean 0 and variance ϵ^2; and (f) follows from the definition of Frobenius norm. $\square$

Using Theorem 4.1, problem (4.25) can be rewritten equivalently as follows.

$$\begin{aligned}
\min_{\mathbf{F},\mathbf{W}} \quad & \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \epsilon^2 \cdot \|\mathbf{F}\|_{\mathrm{F}}^2 \cdot \|\mathbf{W}\|_{\mathrm{F}}^2 && \text{(4.35a)}\\
\text{s.t.} \quad & \|N_{\mathrm{t}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{tx}}^*\mathbf{F})\|_2^2 \le \sigma_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}^2 \cdot M_{\mathrm{tx}} && \text{(4.35b)}\\
& \|N_{\mathrm{r}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{rx}}^*\mathbf{W})\|_2^2 \le \sigma_{\mathrm{rx}}^2 \cdot N_{\mathrm{r}}^2 \cdot M_{\mathrm{rx}} && \text{(4.35c)}\\
& [\mathbf{F}]_{:,i} \in \mathcal{F} \ \forall\ i = 1, \ldots, M_{\mathrm{tx}} && \text{(4.35d)}\\
& [\mathbf{W}]_{:,j} \in \mathcal{W} \ \forall\ j = 1, \ldots, M_{\mathrm{rx}} && \text{(4.35e)}
\end{aligned}$$

For some channel estimate $\bar{\mathbf{H}}$ and knowledge of the estimation error variance ϵ^2, engineers can aim to solve problem (4.35) for codebooks $\mathbf{F}$ and $\mathbf{W}$ that minimize the expected average self-interference coupled between transmit and receive beams while ensuring coverage is delivered with some quality parameterized by σ_{tx}^2 and σ_{rx}^2. This concludes the formulation of our codebook design LoneSTAR. In the next section, we present a means to solve problem (4.35), along with a summary of and commentary on our design.

4.5 Solving for LoneSTAR Codebooks and Remarks

In the previous section, we formulated problem (4.35), a codebook design problem that accounts for self-interference channel estimation error and digitally-controlled phase shifters and attenuators. In this section, we present an approach for (approximately) solving this design problem using a projected alternating minimization approach. We conclude this section with a summary of our design and remarks regarding design decisions and future research directions.

Problem (4.35) cannot be efficiently solved directly due to the non-convexity posed by digitally-controlled phase shifters and attenuators (i.e., $\mathcal{F}$ and $\mathcal{W}$ are non-convex sets). The joint optimization of $\mathbf{F}$ and $\mathbf{W}$ introduces further complexity, especially given their dimensionality. In order to approximately but efficiently solve problem (4.35), we take a projected alternating minimization approach, which is summarized in Algorithm 4.1. Let $\Pi_{\mathcal{F}}\left([\mathbf{F}]_{:,i}\right)$ be the projection of the i-th beam (column) of $\mathbf{F}$ onto the set of physically realizable beamforming vectors $\mathcal{F}$. We extend this projection onto each beam in $\mathbf{F}$ as $\Pi_{\mathcal{F}}(\mathbf{F})$ by applying it column-wise on $\mathbf{F}$; the projection of $\mathbf{W}$ onto $\mathcal{W}$ is defined analogously. With digitally-controlled phase shifters and attenuators, each element of $\mathbf{F}$ must be quantized in both phase and amplitude. Since the set of physically realizable beamforming weights is the same for each element in a beamforming vector, the projection of that vector reduces to projecting each of its elements. The (m, n)-th element following the projection of $\mathbf{F}$ onto

$\mathcal{F}$ can be expressed as

$$[\Pi_{\mathcal{F}}(\mathbf{F})]_{m,n} = A^\star \cdot \exp\left(\mathrm{j} \cdot \theta^\star\right), \tag{4.36}$$

where $A^\star$ and $\theta^\star$ are the quantized amplitude and phase of $[\mathbf{F}]_{m,n}$. These can be expressed as

$$A^\star = \underset{A_i \in \mathcal{A}}{\operatorname{argmin}} \ |A_i - A|, \tag{4.37}$$

$$\theta^\star = \underset{\theta_i \in \mathcal{P}}{\operatorname{argmin}} \ |\exp\left(\mathrm{j} \cdot \theta_i\right) - \exp\left(\mathrm{j} \cdot \theta\right)|, \tag{4.38}$$

where $A = \left|[\mathbf{F}]_{m,n}\right|$ and $\theta = \mathrm{angle}\left([\mathbf{F}]_{m,n}\right)$ are the magnitude and phase of $[\mathbf{F}]_{m,n}$. Here, $\mathcal{A}$ and $\mathcal{P}$ are the sets of physically realizable amplitudes (attenuator settings) and phases (phase shifter settings), respectively.

To approximately solve problem (4.35), we begin with $\mathbf{F} \leftarrow \Pi_{\mathcal{F}}(\mathbf{A}_{\mathrm{tx}})$ and $\mathbf{W} \leftarrow \Pi_{\mathcal{W}}(\mathbf{A}_{\mathrm{rx}})$, which initializes our beams as matched filters to achieve (near) maximum gains across $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$, respectively. Then, we fix the receive codebook $\mathbf{W}$ and solve problem (4.39) below.

$$\min_{\mathbf{F}} \ \left\|\mathbf{W}^*\bar{\mathbf{H}}\mathbf{F}\right\|_{\mathrm{F}}^2 + \epsilon^2 \cdot \|\mathbf{F}\|_{\mathrm{F}}^2 \cdot \|\mathbf{W}\|_{\mathrm{F}}^2 \tag{4.39a}$$

$$\text{s.t.} \ \ \|N_{\mathrm{t}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{tx}}^*\mathbf{F})\|_2^2 \leq \sigma_{\mathrm{tx}}^2 \cdot N_{\mathrm{t}}^2 \cdot M_{\mathrm{tx}} \tag{4.39b}$$

$$\left|[\mathbf{F}]_{i,j}\right| \leq 1, \ i = 1, \ldots, N_{\mathrm{t}}, \ j = 1, \ldots, M_{\mathrm{tx}} \tag{4.39c}$$

As a temporary convex relaxation, we have introduced constraint (4.39c) instead of requiring strict quantized phase and amplitude control. This enforces that each beamforming weight not exceed unit magnitude, understood by the fact that attenuators are used for amplitude control as outlined in (4.1). Problem (4.39) is convex and can be readily solved using convex optimization solvers (e.g., we used [89]). Dedicated algorithms more efficient than general solvers can likely be developed for LoneSTAR, but that is beyond the scope of this work, making it a good topic for future work. After solving problem (4.39), the solution $\mathbf{F}^{\star}$ is projected onto the set $\mathcal{F}$ as $\mathbf{F} \leftarrow \Pi_{\mathcal{F}}(\mathbf{F}^{\star})$ to ensure it is physically realizable. The projected $\mathbf{F}$ is then used when solving for the receive codebook in problem (4.40) below.

$$\min_{\mathbf{W}} \ \left\| \mathbf{W}^{*} \bar{\mathbf{H}} \mathbf{F} \right\|_{\mathrm{F}}^{2} + \epsilon^{2} \cdot \| \mathbf{F} \|_{\mathrm{F}}^{2} \cdot \| \mathbf{W} \|_{\mathrm{F}}^{2} \tag{4.40a}$$

$$\text{s.t.} \ \| N_{\mathrm{r}} \cdot \mathbf{1} - \mathrm{diag}(\mathbf{A}_{\mathrm{rx}}^{*} \mathbf{W}) \|_{2}^{2} \leq \sigma_{\mathrm{rx}}^{2} \cdot N_{\mathrm{r}}^{2} \cdot M_{\mathrm{rx}} \tag{4.40b}$$

$$\left| [\mathbf{W}]_{i,j} \right| \leq 1, \ i = 1, \ldots, N_{\mathrm{r}}, \ j = 1, \ldots, M_{\mathrm{rx}} \tag{4.40c}$$

Problem (4.40) has a form identical to problem (4.39) and is thus convex and can be solved using the same approach used for problem (4.39). Upon solving problem (4.40), we ensure the solution $\mathbf{W}^{\star}$ is physically realizable by projecting it onto $\mathcal{W}$ via $\mathbf{W} \leftarrow \Pi_{\mathcal{W}}(\mathbf{W}^{\star})$. This alternating minimization is only executed once since convergence is typically reached immediately due to limited phase and amplitude control; in other words, additional iterations simply project on to the same solutions as the first solution. This concludes our design of LoneSTAR, which has produced physically realizable transmit and receive codebook matrices $\mathbf{F}$ and $\mathbf{W}$. These can then be unpacked via Definition 4.2 to form the sets $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ that can be used for beam alignment and subsequent full-duplexing of downlink and uplink. In the next two sections, we evaluate LoneSTAR extensively through simulation.

Algorithm 4.1 Projected alternating minimization to approximately solve problem (4.35).

1. Initialize $\mathbf{F} \leftarrow \Pi_{\mathcal{F}}(\mathbf{A}_{\mathrm{tx}})$ and $\mathbf{W} \leftarrow \Pi_{\mathcal{W}}(\mathbf{A}_{\mathrm{rx}})$.
2. Solve problem (4.39) for the transmit codebook $\mathbf{F}^{\star}$.
3. Project $\mathbf{F}^{\star}$ to ensure it is physically realizable: $\mathbf{F} \leftarrow \Pi_{\mathcal{F}}(\mathbf{F}^{\star})$.
4. Solve problem (4.40) for the receive codebook $\mathbf{W}^{\star}$ using the updated $\mathbf{F}$.
5. Project $\mathbf{W}^{\star}$ to ensure it is physically realizable: $\mathbf{W} \leftarrow \Pi_{\mathcal{W}}(\mathbf{W}^{\star})$.

Remark 4.1: Design decisions. We present a summary of LoneSTAR in Algorithm 4.2, and provide a few remarks on our approach and communicate our motivations and design decisions. The objective of LoneSTAR is to create transmit and receive codebooks that minimize the average $\mathsf{INR}_{\mathrm{rx}}$ while still delivering high transmit and receive beamforming gain. It may seem preferred to maximize the average $\mathsf{SINR}_{\mathrm{rx}}$ or sum spectral efficiency, but this requires knowledge of the transmit and receive SNRs, which the system only has *after* conducting beam alignment, or real-time knowledge of $\mathbf{h}_{\mathrm{tx}}$ and $\mathbf{h}_{\mathrm{rx}}$, which is currently impractical. LoneSTAR only relies on knowledge of the self-interference MIMO channel and not that of the transmit and receive channels. By minimizing $\mathsf{INR}_{\mathrm{rx}}$ while constraining transmit and receive coverage quality (i.e., maintaining high $\mathsf{SNR}_{\mathrm{tx}}$ and $\mathsf{SNR}_{\mathrm{rx}}$), LoneSTAR can consequently achieve high $\mathsf{SINR}_{\mathrm{rx}}$ and $\mathsf{SINR}_{\mathrm{tx}}$. One may also suggest we instead minimize the codebooks' maximum $\mathsf{INR}_{\mathrm{rx}}$ (to minimize worst-case self-interference) rather than minimize the average $\mathsf{INR}_{\mathrm{rx}}$. We chose to minimize the average $\mathsf{INR}_{\mathrm{rx}}$ because this prevents LoneSTAR from being significantly hindered if there are select transmit-receive beam pairs that simply cannot offer low self-interference while delivering high beamforming gain. It is important to note that the Frobenius norm objective used by LoneSTAR will make it more incentivized to reduce high self-interference but does not exclusively prioritize worst-case beam pairs [99]. Individual beam pairs that do not yield sufficiently low $\mathsf{INR}_{\mathrm{rx}}$ when minimizing the average $\mathsf{INR}_{\mathrm{rx}}$ with LoneSTAR could potentially be avoided via scheduling or half-duplexed as a last resort.

Remark 4.2: Complexity. It is difficult to comment on the complexity of our design for a few reasons. First and foremost, its complexity depends heavily on the solver used for solving problems (4.39) and (4.40). As mentioned, more efficient solving algorithms may be tailored for LoneSTAR, but that is out of the scope of this work. Additionally, total computational costs will depend on the initialization of $\mathbf{F}$ and $\mathbf{W}$. We have initialized them both as conjugate beamforming codebooks; in practice, it would likely be efficient to use previous solutions of $\mathbf{F}$ and $\mathbf{W}$ when rerunning LoneSTAR to account for drift in the self-interference channel. Note that our design takes place at the full-duplex base station, which presumably has sufficient compute resources to execute LoneSTAR. Practically, it may be preferred to conduct a thorough estimation of $\mathbf{H}$ and execution of LoneSTAR during initial setup of the base station and then perform updates to the estimate of $\mathbf{H}$ and possibly to our design over time; this would be interesting future work. Most existing beamforming-based full-duplex solutions need to be executed for each downlink-uplink user pair, whereas the same LoneSTAR design can be used across many user pairs, meaning it may offer significant computational savings in comparison.

Algorithm 4.2 A summary of our analog beamforming codebook design, LoneSTAR.

1. Obtain an estimate $\bar{\mathbf{H}}$ of the self-interference channel and the estimation error variance ϵ^2.
2. Define transmit and receive coverage regions $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$ and assemble $\mathbf{A}_{\mathrm{tx}}$ and $\mathbf{A}_{\mathrm{rx}}$.
3. Set transmit and receive coverage variances σ^2_{tx} and σ^2_{rx}.
4. Solve problem (4.35) for $\mathbf{F}$ and $\mathbf{W}$ using Algorithm 4.1.
5. Unpack $\mathbf{F}$ and $\mathbf{W}$ to form codebooks $\bar{\mathcal{F}}$ and $\bar{\mathcal{W}}$ using Definition 4.2.

4.6 Simulation Setup and Performance Metrics

We simulated the full-duplex system illustrated in Figure 4.1 at 30 GHz in a Monte Carlo fashion to extensively evaluate LoneSTAR against conventional codebooks [104]. The full-duplex base station is equipped with two 8×8 half-wavelength uniform planar arrays for transmission and reception; users have a single antenna for simplicity. Analog beamforming networks at the base station use log-stepped attenuators with 0.5 dB of attenuation per least significant bit (LSB) and phase shifters with uniform phase control over $[0, 2\pi)$. We consider downlink and uplink service regions spanning -60° to 60° in azimuth and -30° to 30° in elevation, each in 15° steps, resulting in codebooks of $M_{\mathrm{tx}} = M_{\mathrm{rx}} = 45$ beams. Users are distributed uniformly from -67.5° to 67.5° in azimuth and from -37.5° to 37.5° in elevation. The downlink and uplink channels are simulated as line-of-sight (LOS) channels for simplicity; similar conclusions would be expected when modeling them otherwise due to the directional nature of mmWave channels. To select which transmit and receive beams serve a given downlink and uplink user pair, we use exhaustive beam sweeping on each link independently and choose the beams that maximize $\mathsf{SNR}_{\mathrm{tx}}$ and $\mathsf{SNR}_{\mathrm{rx}}$. The choice of a particular self-interference channel model is a difficult one to justify practically, given there lacks a well-accepted, measurement-backed model. We have chosen to simulate the self-interference channel $\mathbf{H}$ using the spherical-wave channel model [2], which captures idealized near-field propagation between the transmit and receive arrays of the full-duplex device and has been used widely thus far in the related literature. This channel model is a function of the relative geometry between the transmit and receive arrays and is defined as

$$[\mathbf{H}]_{m,n} = \frac{\rho}{r_{n,m}} \exp\left(-\mathrm{j}2\pi \frac{r_{n,m}}{\lambda}\right), \tag{4.41}$$

where $r_{n,m}$ is the distance between the n-th transmit antenna and the m-th receive antenna, λ is the carrier wavelength, and ρ is a normalizing factor to satisfy $\|\mathbf{H}\|_{\mathrm{F}}^2 = N_{\mathrm{t}} \cdot N_{\mathrm{r}}$. To realize such a channel, we have vertically stacked our transmit and receive arrays with 10λ of separation. While this channel model may not hold perfectly in practice, it provides a sensible starting point to evaluate LoneSTAR, particularly when self-interference is dominated by idealized near-field propagation. Shortly, for more extensive evaluation, we consider a self-interference channel model that mixes this spherical-wave model with a Rayleigh faded channel. We do not model cross-link interference explicitly but instead evaluate performance over various $\mathsf{INR}_{\mathrm{tx}}$.

4.6.1 Baselines and Performance Metrics

Considering LoneSTAR is the first known codebook design specifically for mmWave full-duplex, we evaluate it against two conventional codebooks, both of which are used in practice by half-duplex systems to form sets of beams that provide broad coverage with high gain:

- A *CBF codebook* that uses conjugate beamforming [88], where

$$\mathbf{f}_i = \Pi_{\mathcal{F}}\left(\mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}\right)\right), \qquad \mathbf{w}_j = \Pi_{\mathcal{W}}\big(\mathbf{a}_{\mathrm{rx}}\big(\theta_{\mathrm{rx}}^{(j)}\big)\big). \tag{4.42}$$

  CBF codebooks can deliver (approximately) maximum gain to all directions in $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$. The projections here are simply to ensure the beams are physically realizable and do not have significant impacts on the beam shapes with modest phase shifter resolution b_{phs}.

- A *Taylor codebook* that uses conjugate beamforming with side lobe suppression as

$$\mathbf{f}_i = \Pi_{\mathcal{F}}\left(\mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}\right) \odot \mathbf{v}_{\mathrm{tx}}\right), \qquad \mathbf{w}_j = \Pi_{\mathcal{W}}\big(\mathbf{a}_{\mathrm{rx}}\big(\theta_{\mathrm{rx}}^{(j)}\big) \odot \mathbf{v}_{\mathrm{rx}}\big), \tag{4.43}$$

  where $\mathbf{v}_{\mathrm{tx}}$ and $\mathbf{v}_{\mathrm{rx}}$ are Taylor windows applied element-wise, providing 25 dB of side lobe suppression; we found this level of side lobe suppression competes best with LoneSTAR. Taylor windowing is an established means to reduce interference inflicted by side lobe levels at the cost of lessened beamforming gain and a wider main lobe [91, 92].

We assess LoneSTAR against our baseline codebooks in terms of sum spectral efficiency of the transmit and receive links. Specifically, we use the normalized sum spectral efficiency

$$\gamma_{\text{sum}} = \frac{R_{\text{tx}} + R_{\text{rx}}}{C_{\text{tx}}^{\text{cb}} + C_{\text{rx}}^{\text{cb}}}, \tag{4.44}$$

where $C_{\text{tx}}^{\text{cb}} = \log_2\left(1 + \mathsf{SNR}_{\text{tx}}^{\text{CBF}}\right)$ and $C_{\text{rx}}^{\text{cb}} = \log_2\left(1 + \mathsf{SNR}_{\text{rx}}^{\text{CBF}}\right)$ we term the *codebook capacities*; here, $\mathsf{SNR}_{\text{tx}}^{\text{CBF}}$ and $\mathsf{SNR}_{\text{rx}}^{\text{CBF}}$ are the maximum SNRs delivered to a given downlink-uplink user pair by the CBF codebook. In other words, $C_{\text{tx}}^{\text{cb}} + C_{\text{rx}}^{\text{cb}}$ is the maximum sum spectral efficiency when using the CBF codebook and interference is inherently not present (i.e., when $\overline{\mathsf{INR}}_{\text{rx}} = \mathsf{INR}_{\text{tx}} = -\infty$ dB). With this normalization, γ_{sum} will gauge how well a given codebook delivers transmission and reception in the face of self-interference and cross-link interference, relative to the interference-free capacity when beamforming with the CBF codebook. Note that $\gamma_{\text{sum}} = 0.5$ can be achieved with the CBF codebook by half-duplexing transmission and reception with equal TDD under an instantaneous power constraint. As such, LoneSTAR generally aims to exceed $\gamma_{\text{sum}} = 0.5$ and outperform that achieved by CBF and Taylor codebooks. Note that γ_{sum} is not truly upper-bounded by 1 since $C_{\text{tx}}^{\text{cb}} + C_{\text{rx}}^{\text{cb}}$ is not the true sum Shannon capacity; nonetheless, achieving $\gamma_{\text{sum}} \approx 1$ indicates near best-case full-duplex performance one can expect when using beamforming codebooks. Having simulated our system in a Monte Carlo fashion, in our results we will present γ_{sum} averaged over downlink-uplink user realizations.

4.6.2 Beam Patterns

In Figure 4.5, we compare conventional beams to those produced by LoneSTAR. Figure 4.5a depicts the shape of beams from our two baseline

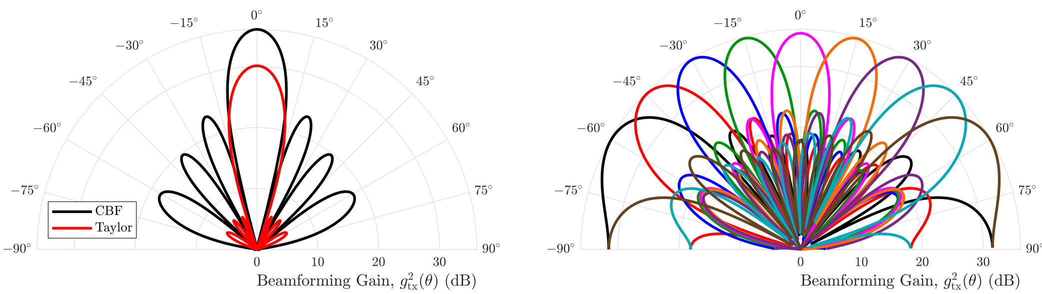

(a) Conventional broadside beams.

(b) LoneSTAR beams.

Figure 4.5: (a) The azimuth cut of a broadside beam from conventional CBF and Taylor codebooks. (b) The azimuth cuts of transmit beams produced by LoneSTAR that serve various azimuth directions at an elevation of 0°.

codebooks. CBF beams can each deliver maximum gain in their steered direction and exhibit a narrow main lobe and high side lobes. With the Taylor codebook, these side lobes are reduced to 25 dB below the main lobe, which itself sacrificed about 6 dB in maximum gain for this side lobe suppression. In Figure 4.5b, we plot transmit beams from a codebook produced by LoneSTAR; receive beams are similar. Clearly, unlike the Taylor codebook, LoneSTAR does not attempt to merely shrink side lobes to reduce self-interference. Instead, **LoneSTAR makes use of side lobes, shaping them to cancel self-interference spatially**, and while doing so, maintains near maximum gain across its beams.

4.6.3 Choosing Coverage Variances, σ^2_{tx} and σ^2_{rx}

In general, it is not possible to analytically state optimal choices for the LoneSTAR design parameters σ^2_{tx} and σ^2_{rx} which maximize normalized sum spectral efficiency γ_{sum}. This motivates us to examine heuristically how

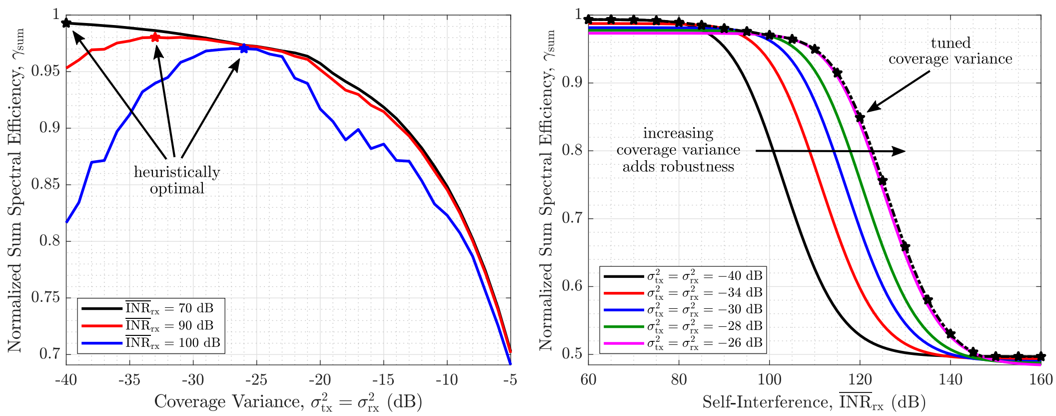

(a) Choosing a coverage variance, $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$. (b) As a function of self-interference, $\overline{\mathsf{INR}}_{\text{rx}}$.

Figure 4.6: Normalized sum spectral efficiency γ_{sum} (a) as a function of the coverage variance $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$ for various $\overline{\mathsf{INR}}_{\text{rx}}$; and (b) as a function of self-interference strength $\overline{\mathsf{INR}}_{\text{rx}}$ for various $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$. In both, $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} = 10$ dB, $b_{\text{phs}} = b_{\text{amp}} = 8$ bits, $\mathsf{INR}_{\text{tx}} = -\infty$ dB, and $\epsilon^2 = -\infty$ dB. Performance with the optimal $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$ is denoted with $\star$.

to choose σ^2_{tx} and σ^2_{rx}. In Figure 4.6a, we plot the normalized sum spectral efficiency γ_{sum} as a function of $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$ for various levels of self-interference $\overline{\mathsf{INR}}_{\text{rx}}$. Here, we let $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} = 10$ dB and have assumed no channel estimation error nor cross-link interference (i.e., $\epsilon^2 = -\infty$ dB, $\mathsf{INR}_{\text{tx}} = -\infty$ dB). At low self-interference $\overline{\mathsf{INR}}_{\text{rx}}$ (in black), there is inherently low coupling between the transmit and receive beams, meaning it is best to maintain high beamforming gain rather than reduce self-interference further. This makes a low coverage variance optimal on average, shown as the $\star$ at $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}} = -40$ dB. As $\overline{\mathsf{INR}}_{\text{rx}}$ increases, LoneSTAR benefits from having greater transmit and receive coverage variance since it affords more flexibility to reduce self-interference and increase the sum spectral efficiency.

In Figure 4.6b, we highlight this further by plotting the normalized sum spectral efficiency as a function of $\overline{\mathsf{INR}}_{\mathrm{rx}}$ for various $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$. Naturally, one may independently choose σ_{tx}^2 and σ_{rx}^2, rather than assume them equal as we have done here for simplicity; this can therefore be considered a conservative evaluation of LoneSTAR. The starred curve represents tuning $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ to maximize sum spectral efficiency at each $\overline{\mathsf{INR}}_{\mathrm{rx}}$. At very low $\overline{\mathsf{INR}}_{\mathrm{rx}}$, full-duplex operation "is free", as transmit and receive beams couple sufficiently low self-interference on their own. In this region, minimizing the coverage variance $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ is ideal since it preserves near-maximum $\mathsf{SNR}_{\mathrm{tx}}$ and $\mathsf{SNR}_{\mathrm{rx}}$. As $\overline{\mathsf{INR}}_{\mathrm{rx}}$ increases, **increasing $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ supplies the system with greater robustness to self-interference**, evidenced by the rightward shift of the curves. When $\overline{\mathsf{INR}}_{\mathrm{rx}}$ is overwhelmingly high, LoneSTAR cannot mitigate self-interference enough, succumbing to self-interference and offering at most around $\gamma_{\mathrm{sum}} = 0.5$ (essentially half-duplexing). Notice that in this region, it is optimal to return to using very low $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ in order to preserve half-duplex performance at worst by maximizing SNR. Clearly, choosing $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ is an important design decision and depends heavily on system factors (chiefly $\overline{\mathsf{INR}}_{\mathrm{rx}}$). In practice, we imagine a system relying on a lookup table to retrieve the optimal σ_{tx}^2 and σ_{rx}^2 based on its measured $\overline{\mathsf{INR}}_{\mathrm{rx}}$ (along with other system factors), which would likely depend on simulation and experimentation. Henceforth, we assume that the optimal $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ is used which maximizes γ_{sum}.

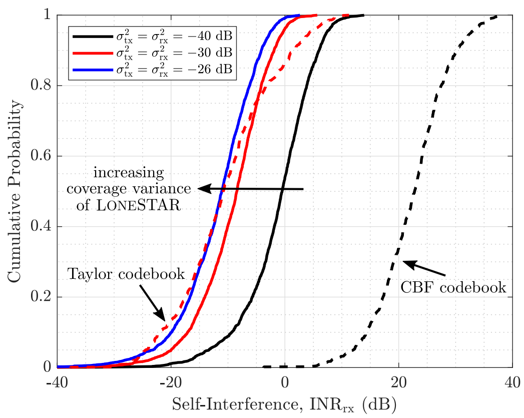

(a) CDF of self-interference, $\mathsf{INR}_{\mathrm{rx}}$.

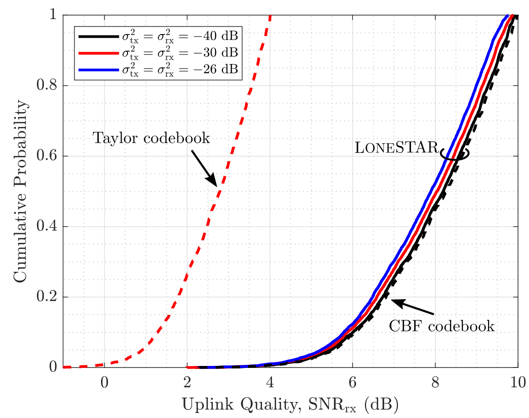

(b) CDF of uplink quality, $\mathsf{SNR}_{\mathrm{rx}}$.

Figure 4.7: The CDFs of (a) self-interference $\mathsf{INR}_{\mathrm{rx}}$ where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB; and (b) uplink quality $\mathsf{SNR}_{\mathrm{rx}}$ where $\overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB. LoneSTAR ($b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits) can reduce self-interference between all beam pairs while also maintaining high uplink quality, netting it greater full-duplex performance than conventional codebooks (8 bits).

4.7 Evaluating LoneSTAR Codebooks

Having presented the simulation setup and performance metrics in the previous section, we now compare performance with LoneSTAR versus that with conventional codebooks. Consider Figure 4.7a, which depicts the cumulative density function (CDF) of $\mathsf{INR}_{\mathrm{rx}}$ across all transmit-receive beam pairs within LoneSTAR codebooks (with $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits) versus conventional ones (with 8 bits), where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB[3]. For most beam pairs within the CBF codebook, self-interference is typically prohibitively high for full-duplex

[3]We commonly use $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB since this produces a median $\mathsf{INR}_{\mathrm{rx}}$ with the CBF codebook that aligns with our INR measurements using an actual 28 GHz phased array platform in [44].

operation, having $\mathsf{INR}_{\mathrm{rx}} \geq 10$ dB. Very few beam pairs yield $\mathsf{INR}_{\mathrm{rx}} \leq 0$ dB. Desirably, the Taylor codebook can reliably reduce $\mathsf{INR}_{\mathrm{rx}}$ to well below 0 dB for most beam pairs. Likewise, by increasing the coverage variance $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$, LoneSTAR can reduce self-interference between its beam pairs to levels comparable to the Taylor codebook. Note that the difference between LoneSTAR and the CBF codebook can be thought of as the amount of self-interference cancellation achieved by LoneSTAR.

Now, consider Figure 4.7b, which shows the CDF of the resulting uplink quality $\mathsf{SNR}_{\mathrm{rx}}$ across users, where $\overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB is its upper bound. The CBF codebook can deliver high gain broadly across users, with only a small fraction not within coverage. The Taylor codebook, on the other hand, delivers poor coverage since it sacrificed main lobe gain to shrink its side lobes. For the same three $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$ considered in Figure 4.7a, LoneSTAR can provide coverage comparable to the CBF codebook, falling short by only about 0.5 dB in distribution. Together, the two evaluations in Figure 4.7 highlight how **LoneSTAR is capable of offering increased SINR over conventional codebooks by making small sacrifices in coverage to significantly reduce self-interference.**

In Figure 4.8, we further compare LoneSTAR codebooks against our two baseline codebooks in terms of sum spectral efficiency. We assume no cross-link interference ($\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB) and no channel estimation error ($\epsilon^2 = -\infty$ dB); we will evaluate LoneSTAR against both of these shortly. In Figure 4.8a, we plot the normalized sum spectral efficiency γ_{sum} achieved by

LoneSTAR as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$ for various resolutions $b_{\mathrm{phs}} = b_{\mathrm{amp}}$, where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB. The CBF and Taylor codebooks were both configured with $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits. At low $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$, the CBF codebook can deliver $\gamma_{\mathrm{sum}} \approx 0.5$, since the receive link is overwhelmed by self-interference while the transmit link sees nearly full beamforming gain. The Taylor codebook sacrifices beamforming gain to reduce its side lobes and therefore achieves only $\gamma_{\mathrm{sum}} \approx 0.46$ at low $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. As $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$ increases, CBF and Taylor codebooks naturally both improve their performance as the impacts of self-interference diminish with increased SNR. Shown as solid lines, LoneSTAR improves significantly as $b_{\mathrm{phs}} = b_{\mathrm{amp}}$ increases from 4 bits to 8 bits. Even at low $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$, **LoneSTAR can deliver sum spectral efficiencies well above that of conventional codebooks by better rejecting self-interference while maintaining high beamforming gain.** This is especially true when $b_{\mathrm{phs}} = b_{\mathrm{amp}} \geq 6$ bits, which allows LoneSTAR to deliver upwards of 90% of the full-duplex codebook capacity. At high SNR, LoneSTAR naturally improves, saturating due to diminishing gains in $\log_2(1 + x)$. For $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} \geq 5$ dB, the Taylor codebook can significantly outperform the CBF codebook and can marginally outperform LoneSTAR having $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 4$ bits, especially at high SNR.

In Figure 4.8b, we now fix $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB and vary $\overline{\mathsf{INR}}_{\mathrm{rx}}$, keeping all other factors the same as in Figure 4.8a. CBF and LoneSTAR both can nearly reach the full-duplex codebook capacity at low $\overline{\mathsf{INR}}_{\mathrm{rx}}$, where self-interference is negligible and near maximal beamforming gain can be de-

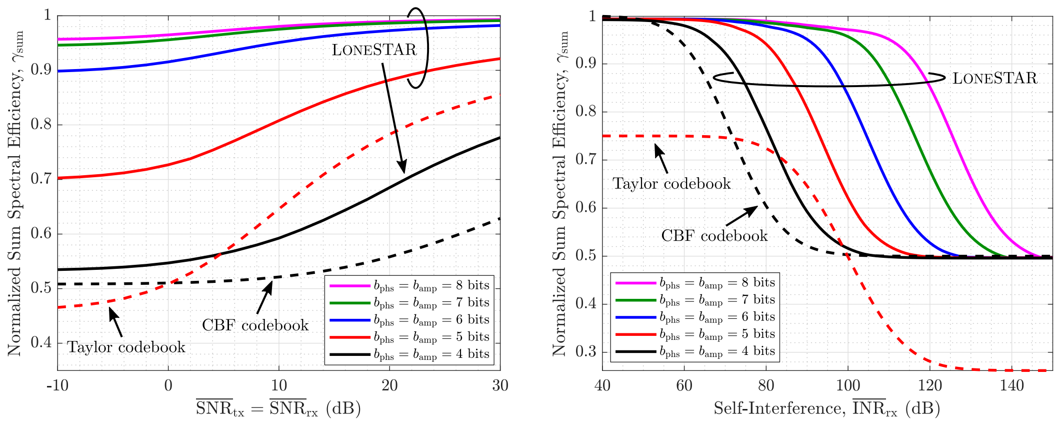

(a) As a function of link quality.

(b) As a function of self-interference, $\overline{\mathsf{INR}}_{\text{rx}}$.

Figure 4.8: Normalized sum spectral efficiency γ_{sum} (a) as a function of $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$ where $\overline{\mathsf{INR}}_{\text{rx}} = 90$ dB; (b) as a function of $\overline{\mathsf{INR}}_{\text{rx}}$ where $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} = 10$ dB. In both, $\mathsf{INR}_{\text{tx}} = -\infty$ dB, $\epsilon^2 = -\infty$ dB, and $\sigma^2_{\text{tx}} = \sigma^2_{\text{rx}}$ are tuned to maximize sum spectral efficiency. CBF and Taylor codebooks are each with $b_{\text{phs}} = b_{\text{amp}} = 8$ bits.

livered. The Taylor codebook sacrifices beamforming gain for reduced side lobes and therefore falls well short of $\gamma_{\text{sum}} = 1$, even in the presence of extremely weak self-interference. **As $\overline{\mathsf{INR}}_{\text{rx}}$ is increased, CBF quickly succumbs to self-interference, whereas LoneSTAR proves to be notably more robust to such.** In fact, with each unit increase in resolution $b_{\text{phs}} = b_{\text{amp}}$, LoneSTAR supplies around 10 dB of added robustness to self-interference. The Taylor codebook can outperform the CBF codebook but only for $\overline{\mathsf{INR}}_{\text{rx}} \in [72, 100]$ dB before also being overwhelmed by self-interference. At higher SNR, when $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} \geq 10$ dB, this $\overline{\mathsf{INR}}_{\text{rx}}$ range would widen, as evidenced by Figure 4.8a. As $\overline{\mathsf{INR}}_{\text{rx}}$ increases to extremely high levels, LoneSTAR converges to offer nearly the same sum spectral efficiency as CBF of

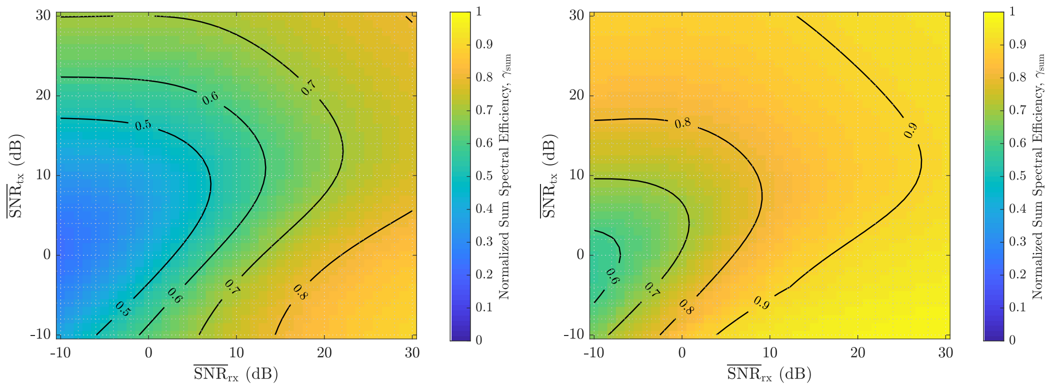

(a) Taylor codebook. (b) LoneSTAR codebooks.

Figure 4.9: Normalized sum spectral efficiency γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ for the (a) Taylor codebook ($b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits) and (b) LoneSTAR codebooks, where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB, $\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB, $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits, and $\epsilon^2 = -\infty$ dB. LoneSTAR delivers higher γ_{sum} broadly across SNRs and demands lower SNRs to net $\gamma_{\mathrm{sum}} \geq 0.5$.

$\gamma_{\mathrm{sum}} = 0.5$, essentially only serving the downlink user due to uplink being overwhelmed by self-interference. Each full-duplex mmWave base station will see a unique $\overline{\mathsf{INR}}_{\mathrm{rx}}$ depending on a variety of parameters including transmit power, noise power, and isolation between its transmit and receive arrays. This makes LoneSTAR an attractive codebook design since it is more robust than conventional codebooks over a wide range of $\overline{\mathsf{INR}}_{\mathrm{rx}}$ and provides greater flexibility when designing full-duplex mmWave platforms. With LoneSTAR, systems can operate at higher $\overline{\mathsf{INR}}_{\mathrm{rx}}$, meaning they can potentially increase transmit power without being severely penalized by increased self-interference, for instance.

4.7.1 For Various Link Qualities, $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$

Rather than assume symmetric downlink/uplink quality $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$, we now consider $\overline{\mathsf{SNR}}_{\mathrm{tx}} \neq \overline{\mathsf{SNR}}_{\mathrm{rx}}$. In Figure 4.9a, we show the normalized sum spectral efficiency achieved by the Taylor codebook as a function of $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$, where $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB and $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits. In Figure 4.9b, we show that of LoneSTAR with $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits. We chose to compare against the Taylor codebook rather than the CBF codebook because it better competes with LoneSTAR at $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB (see Figure 4.8a). Clearly, **LoneSTAR can deliver higher spectral efficiencies** than the Taylor codebook broadly over $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$. It is important to notice that the region where $\gamma_{\mathrm{sum}} \geq 0.5$—where full-duplexing is superior to half-duplexing with equal TDD—grows significantly with LoneSTAR versus the Taylor codebook. This means that full-duplex operation is justified over a much broader range of $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$ with LoneSTAR compared to with conventional codebooks. It is also important to notice that **LoneSTAR introduces a significant SNR gain versus conventional codebooks**, since it demands a much lower $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$ to achieve a desired sum spectral efficiency γ_{sum}. Note that we have tuned $\sigma_{\mathrm{tx}}^2 = \sigma_{\mathrm{rx}}^2$, meaning the performance of LoneSTAR may improve if σ_{tx}^2 and σ_{rx}^2 were tuned separately, especially with highly asymmetric SNRs.

4.7.2 For Various Levels of Interference, $\left(\mathsf{INR}_{\mathrm{tx}}, \overline{\mathsf{INR}}_{\mathrm{rx}}\right)$

Similar to Figure 4.9, we now compare the Taylor codebook against LoneSTAR for various $\left(\mathsf{INR}_{\mathrm{tx}}, \overline{\mathsf{INR}}_{\mathrm{rx}}\right)$ in Figure 4.10, where we have fixed

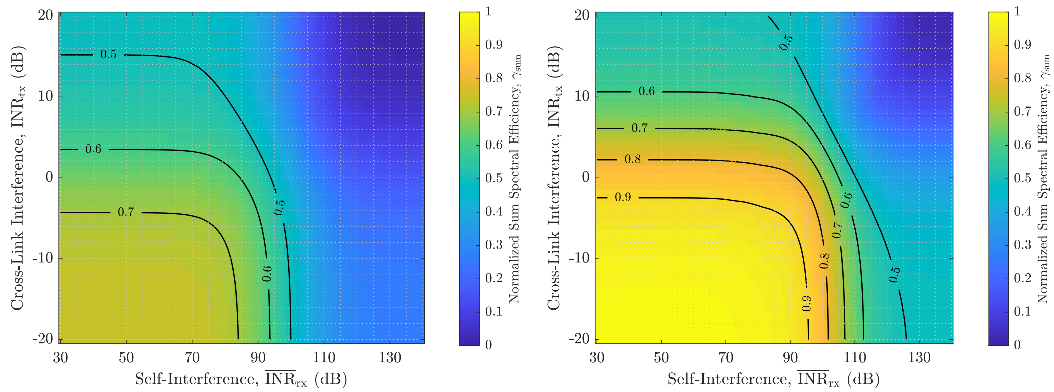

(a) Taylor codebook.

(b) LoneSTAR codebooks.

Figure 4.10: Normalized sum spectral efficiency γ_{sum} as a function of $\mathsf{INR}_{\mathrm{tx}}$ and $\overline{\mathsf{INR}}_{\mathrm{rx}}$ for the (a) Taylor codebook, where $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 8$ bits; and (b) LoneSTAR codebooks, where $b_{\mathrm{phs}} = b_{\mathrm{amp}} = 6$ bits. In both, $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB and $\epsilon^2 = -\infty$ dB. LoneSTAR is more robust to self-interference and cross-link interference than conventional codebooks.

$\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB. Again, LoneSTAR can clearly deliver much higher spectral efficiencies than the Taylor codebook, achieving nearly the full-duplex codebook capacity when cross-link interference and self-interference are both low. Even at very low INRs, the Taylor codebook cannot approach $\gamma_{\mathrm{sum}} = 1$ since it sacrificed beamforming gain for side lobe suppression. At high INRs, both LoneSTAR and the Taylor codebook struggle to deliver high γ_{sum}, largely due to the fact that neither can reduce cross-link interference. In between, at moderate INRs, LoneSTAR can sustain spectral efficiencies notably higher than the Taylor codebook. **LoneSTAR's added robustness to cross-link interference and self-interference is on display here**, as it can tolerate higher $\left(\mathsf{INR}_{\mathrm{tx}}, \overline{\mathsf{INR}}_{\mathrm{rx}}\right)$ than the Taylor codebook while delivering the same γ_{sum}. This is thanks to LoneSTAR's ability to net a higher receive link spectral efficiency, which can yield a higher sum spectral efficiency in the presence of high cross-link interference $\mathsf{INR}_{\mathrm{tx}}$. At an $\overline{\mathsf{INR}}_{\mathrm{rx}} = 50$ dB, for example, LoneSTAR provides around 10 dB of robustness to $\mathsf{INR}_{\mathrm{tx}}$ for $\gamma_{\mathrm{sum}} = 0.7$, compared to the Taylor codebook. At an $\mathsf{INR}_{\mathrm{tx}} = -10$ dB, LoneSTAR provides around 30 dB of robustness to self-interference $\overline{\mathsf{INR}}_{\mathrm{rx}}$ for $\gamma_{\mathrm{sum}} = 0.7$.

4.7.3 Impact of Channel Estimation Error and Channel Structure

Now, in Figure 4.11a, we examine the impact of imperfect self-interference channel knowledge, $\epsilon^2 > -\infty$ dB. To do so, we model $\bar{\mathbf{H}}$ as in (4.41), meaning the true channel $\mathbf{H}$ is Gaussian distributed about $\bar{\mathbf{H}}$ according to (4.10). We chose to model it in this fashion, rather than use $\mathbf{H}$ as the spherical-wave model in (4.41), for two reasons. First, it is more challenging for LoneSTAR to face a channel under this mixed model (as we will see shortly in Figure 4.11b). Second, it represents the potential case where the estimate $\bar{\mathbf{H}}$ is computed directly based on the relative geometry of the transmit and receive arrays using the spherical-wave model (which can be done *a priori*) and the estimation error captures deviations from this assumed spherical-wave behavior. In Figure 4.11a, we plot the normalized sum spectral efficiency as a function of self-interference channel estimation error variance ϵ^2. With accurate estimation of $\mathbf{H}$, LoneSTAR performs quite well as it has proven thus far. As the quality of estimation degrades, LoneSTAR has difficulty reliably avoiding self-interference. Its performance suffers but remains well above $\gamma_{\mathrm{sum}} = 0.5$ before converging to performance offered by the CBF codebook around $\epsilon^2 = -30$

dB. The CBF codebook sees mild degradation with increased estimation error. Clearly, the Taylor codebook, like LoneSTAR, suffers as ϵ^2 increases. It may seem strange that the CBF and Taylor codebooks both vary as a function of estimation error ϵ^2 since neither codebook depends on self-interference channel knowledge. This can be explained by the fact that the estimation error $\mathbf{\Delta}$ is distributed differently than the estimate $\bar{\mathbf{H}}$. In other words, the Taylor codebook is inherently more robust to the more structured spherical-wave self-interference channel than a Gaussian one. **Gaussian-distributed error proves to be especially difficult for LoneSTAR and Taylor to endure.** As such, it is evident that imperfect channel knowledge *and* the fact that estimation error is Gaussian both contribute to LoneSTAR's degradation as ϵ^2 increases.

We dissect the impacts of estimation error versus that of mixing the spherical-wave model with a Gaussian matrix in Figure 4.11b by considering the mixed self-interference channel model

$$\mathbf{H} = \frac{\mathbf{H}_{\mathrm{SW}} + \mathbf{H}_{\mathrm{Ray}}}{\left\| \mathbf{H}_{\mathrm{SW}} + \mathbf{H}_{\mathrm{Ray}} \right\|_{\mathrm{F}}} \cdot \sqrt{N_{\mathrm{t}} \cdot N_{\mathrm{r}}}, \qquad \left[\mathbf{H}_{\mathrm{Ray}} \right]_{m,n} \sim \mathcal{N}_{\mathbb{C}}\left(0, \zeta^2\right) \ \forall \ m, n, \quad (4.45)$$

where $\mathbf{H}_{\mathrm{SW}}$ is the spherical-wave model defined in (4.41) and $\mathbf{H}_{\mathrm{Ray}}$ is Rayleigh faded with variance ζ^2. The scaling in (4.45) simply ensures $\|\mathbf{H}\|_{\mathrm{F}}^2 = N_{\mathrm{t}} \cdot N_{\mathrm{r}}$. We evaluate LoneSTAR using this channel model for various ζ^2, which we term the *mixing variance*. In doing so, we assume perfect knowledge of the realized channel $\mathbf{H}$ (i.e., $\epsilon^2 = -\infty$ dB) to solely investigate the effects of this mixing. The spherical-wave channel $\mathbf{H}_{\mathrm{SW}}$ is a highly structured channel based

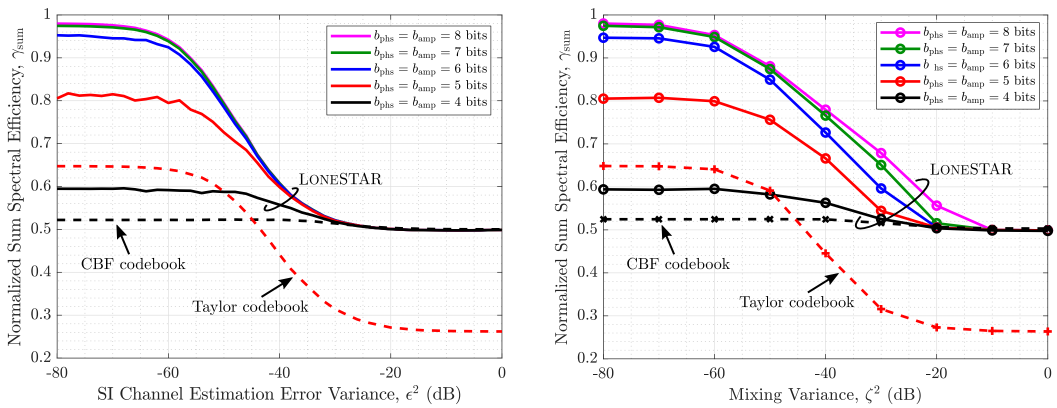

(a) Channel estimation error.

(b) Gaussian mixing.

Figure 4.11: Normalized sum spectral efficiency as a function of (a) self-interference channel estimation error variance ϵ^2 and (b) mixing variance ζ^2. In both, $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB, $\overline{\mathsf{INR}}_{\mathrm{rx}} = 90$ dB, and $\mathsf{INR}_{\mathrm{tx}} = -\infty$ dB. The true channel $\mathbf{H}$ is distributed the same in both, but in (b), we assume perfect channel knowledge, $\epsilon^2 = -\infty$ dB.

on the relative geometry between the transmit and receive arrays at the base station. The Rayleigh faded channel $\mathbf{H}_{\mathrm{Ray}}$ (a Gaussian matrix) lay at the other extreme, lacking any inherent structure. It is important to notice that the distribution of true channel $\mathbf{H}$ here is identical to that examined in Figure 4.11a. In this case, however, we consider $\epsilon^2 = -\infty$ dB to investigate the impacts of $\mathbf{H}$ becoming increasingly Gaussian, rather than those stemming from imperfect channel knowledge.

In Figure 4.11b, we plot the normalized sum spectral efficiency as a function of the mixing variance ζ^2. Having considered perfect channel knowledge $\epsilon^2 = -\infty$ dB here, Figure 4.11b is an upper bound on Figure 4.11a since they follow the same channel model. At very low ζ^2, $\mathbf{H} \approx \mathbf{H}_{\mathrm{SW}}$, allowing LoneSTAR to achieve appreciable spectral efficiencies above conventional codebooks, as we have seen thus far. As ζ^2 increases, the sum spectral efficiency offered by LoneSTAR tends to decrease, though so does the Taylor codebook and the CBF codebook slightly so. At high ζ^2, LoneSTAR suffers since it cannot find ways to pack the transmit and receive beams such that they do not couple across the self-interference channel. This is intimately related to the fact that Gaussian matrices (almost surely) tend to inflict self-interference across all dimensions of $\mathbf{H}$, making it a statistically more difficult channel for LoneSTAR to build codebooks that minimize transmit-receive beam coupling. **Thus, even with perfect channel estimation, LoneSTAR struggles with self-interference channels that are too heavily Gaussian.** However, with modest Gaussian mixing (e.g., $\zeta^2 = -40$ dB), LoneSTAR can deliver appreciable spectral efficiency with accurate channel estimation. From the results in Figure 4.11, we observe that there is the need for accurate and perhaps frequent self-interference channel estimation. Our hope is that this can be accomplished by taking advantage of the strength of the self-interference channel, along with a potentially known channel structure (e.g., [98]) and occasional—but thorough—measurements at the base station, as mentioned in Section 4.3. To overcome challenges associated with Gaussian channels, future work may extend LoneSTAR to the design of a single higher-dimensional codebook of transmit-receive beam pairs $(\mathbf{f}, \mathbf{w})$ that are jointly selected following beam alignment (e.g., similar to [3]).

4.8 Conclusion and Future Directions

Conventional beamforming codebooks are not suitable for full-duplex mmWave systems since they do not offer much robustness to self-interference. Meanwhile, existing beamforming-based solutions for full-duplex mmWave systems typically do not support codebook-based beam alignment, demand real-time knowledge of high-dimensional downlink/uplink MIMO channels, and do not scale well when serving many users over time—all of which are practical shortcomings. We have presented LoneSTAR, a novel approach to enable full-duplex mmWave communication systems through the design of analog beamforming codebooks that are robust to self-interference. Our design does not rely on downlink/uplink channel knowledge and therefore need only be executed once per coherence time of the self-interference channel, allowing the same LoneSTAR codebooks to be used for beam alignment and to serve many downlink-uplink user pairs thereafter. We illustrate that LoneSTAR can mitigate self-interference to levels sufficiently low for appreciable spectral efficiency gains compared to conventional codebooks by being more robust to self-interference and cross-link interference broadly across downlink and uplink SNRs.

This work has motivated a variety of future work for full-duplex mmWave systems. Extensive study and modeling of mmWave self-interference channels and their coherence time will be important to understanding the potentials of LoneSTAR—and full-duplex mmWave systems in general—along with developing efficient means for self-interference channel estimation. Also, methods to dynamically update LoneSTAR as the self-interference channel drifts would be valuable future work, along with implementations of LoneSTAR using real mmWave platforms. Naturally, there are other ways to improve upon this work, such as by better handling limited phase and amplitude control and by creating efficient, dedicated solvers for LoneSTAR. Beyond LoneSTAR, measurement-driven beamforming-based full-duplex solutions (potentially using machine learning) that do not require self-interference MIMO channel knowledge (e.g., Steer [3]) would be excellent contributions to the research community.

Chapter 5

Self-Interference Measurements at 28 GHz

In the previous chapters, we relied on proposed models of self-interference that have been widely used in the related literature to evaluate our designs. These highly-idealized models have yet to be validated with measurements, however, raising to question how applicable they are to real-world full-duplex mmWave systems. With this in mind, this chapter conducts the first extensive measurement campaign of mmWave self-interference and presents a thorough analysis of nearly 6.5 million collected measurements. The findings uncover large-scale and small-scale spatial characteristics of self-interference and motivate the need for new, measurement-backed models of mmWave self-interference.

5.1 Introduction

Research on full-duplex mmWave systems has been motivated by the potential for dense antenna arrays to strategically steer transmit and receive beams in a way that reduces self-interference [19, 20]. Proposed solutions for mmWave full-duplex (e.g., [19, 41, 68, 71–73, 77]) typically make use of hybrid digital/analog beamforming to shape transmit and receive beams that do not couple across the MIMO self-interference channel present between separate transmit and receive arrays of a full-duplex transceiver. With enough isolation between its transmit and receive beams, a full-duplex mmWave transceiver could simultaneously transmit and receive in-band, offering improvements at the physical layer, reduced latency, new approaches for medium access, and cost-effective solutions for network deployment [20]. Full-duplex IAB, for instance, could serve a downlink user while simultaneously receiving backhaul in-band, making better use of mmWave spectrum while reducing latency between the network core and its edge [32, 95, 105].

Given the highly directional nature of mmWave communication, understanding the spatial characteristics of self-interference is critical to evaluating proposed mmWave full-duplex solutions. This is especially true when considering analog-only beamforming systems where digital beamforming cannot be relied on to further mitigate self-interference. Currently, however, there is not a strong understanding of self-interference in mmWave systems. Research and development of full-duplex mmWave systems would benefit from measurement-backed insights on self-interference power levels and spatial characteristics, particularly when beamformed phased arrays are employed.

5.1.1 Prior Work Characterizing mmWave Self-Interference

One of the earliest known attempts to characterize mmWave self-interference was in [106], where a beam-sweeping approach was taken to measure the received self-interference power for a combination of transmit and receive beams.

A relatively low number of beams were swept using 28 GHz 8×8 uniform planar arrays (UPAs) in indoor and outdoor environments, which offered limited characterization of the spatial characteristics and the distribution of self-interference. Nonetheless, this work provided a valuable first look at the expected self-interference power levels seen by a multi-panel mmWave full-duplex system. In [107], using phased array transceivers, a proof-of-concept 60 GHz, short-range, full-duplex communication link was established. The authors observed performance improvements when varying the angular difference between the two transceivers, suggesting that there exist noteworthy spatial characteristics in 60 GHz self-interference. The work of [108] presents self-interference channel measurements at 28 GHz using a pair of directional horn antennas for transmission and reception as well as omnidirectional dipole antennas. Little was shown on the variability with changes in transmit and receive steering direction, though it was noted that some steering combinations offered starkly less self-interference than others.

In [109] and [110], the authors conducted measurements of the 60 GHz self-interference channel with a rotating channel sounder comprised of horn antennas used for transmission and reception, which were fixed to the sounder. Measurements in [109] and [110] showed that large variations in self-interference power can be seen as the sounder rotated in azimuth due to indoor features such as large furniture. Clear gains in isolation were had with cross polarization and the self-interference power delay profile saw variability across azimuth and elevation of the sounder. In [111], self-interference channel measurements at 70 GHz were conducted using lens antennas, which primarily provided insights on the degree of self-interference in various settings.

Measurements of mmWave self-interference in [106–111] are certainly useful but do not offer a means to evaluate proposed mmWave full-duplex solutions since they provide neither a MIMO self-interference channel model nor adequate beam-based measurements. To evaluate beamforming-based mmWave full-duplex solutions thus far, researchers have primarily used highly idealized channel models. For instance, the spherical-wave MIMO channel model [2] has been used most widely to capture coupling between the arrays of a full-duplex mmWave system. This extremely idealized geometric model is sensitive to small errors in the arrays' relative geometry due to the small wavelength at mmWave, and it does not capture significant artifacts of practical systems such as enclosures, mounting infrastructure, and non-isotropic antenna elements. In fact, we show herein that this model does not align with the measurements taken in this campaign. To account for environmental reflections—which were observed in [106]—it has been common for researchers to combine a ray-based model with the spherical-wave model in a Rician fashion [72, 112]. While this inches the model closer to a seemingly more practical one, it has yet to be verified with measurement. Finally, we remark that the severity of inaccurate self-interference channel models can lead to highly misleading results and conclusions for beamforming-based mmWave full-duplex solutions. This is because, in theory, only very few spatial degrees of freedom are needed to execute mmWave full-duplex and highly idealized channel models may readily offer these degrees of freedom whereas practical ones may not.

5.1.2 Contributions

Measuring and characterizing mmWave self-interference. Extending our work in [113], we present the first set of spatially dense measurements of self-interference at 28 GHz using finely steered phased arrays. Our nearly 6.5 million measurements shed light on the levels of self-interference that a practical mmWave full-duplex system could expect using conventional beam steering. This can drive the design requirements for full-duplex systems, including those relying solely on beam steering and those with additional self-interference cancellation measures. Importantly, we compare our measurements of self-interference against what one would expect based on the idealized spherical-wave model in [2] to show that such a channel does not align with reality—motivating the need for new, measurement-backed models. A spatial inspection of our measurements uncovers large-scale trends in self-interference based on general transmit and receive steering directions, as well as noteworthy small-scale variability when these steering directions undergo small shifts (on the order of one degree).

Quantifying and statistically modeling the angular spread of self-interference. To explore this small-scale variability further, we investigate the angular spread of self-interference. We examine how self-interference varies over small spatial neighborhoods to quantify the range in INR, minimum INR, and maximum INR when small shifts are made to the transmit and receive steering directions. We fit distributions to these quantities and tabulate the fitted parameters for various spatial neighborhood sizes to supply engineers with tools for conducting statistical analyses and evaluation of full-duplex mmWave systems. These findings on the angular spread of self-interference shed light on the efficacy of mmWave full-duplex and excitedly motivate a variety of future work including beam codebook design and beam selection.

This paper is organized as follows. In Section 5.2, we describe our 28 GHz self-interference measurement platform and methodology. In Section 5.3, we provide a summary of our measurements along with high-level spatial insights. In Section 5.4, we illustrate how small shifts in transmit and receive steering directions can lead to significant variations in self-interference. We conclude this paper in Section 5.5.

5.2 Measurement Setup and Methodology

In this section, we summarize the setup and methodology we used to collect measurements of the 28 GHz self-interference channel. This campaign sought to measure self-interference between the input of a transmitting phased array and the output of a colocated receiving phased array. Since the degree of self-interference coupled between the phased arrays depends on the steering directions of their beams, as illustrated in Figure 5.1, it was measured extensively across a variety of steering combinations.

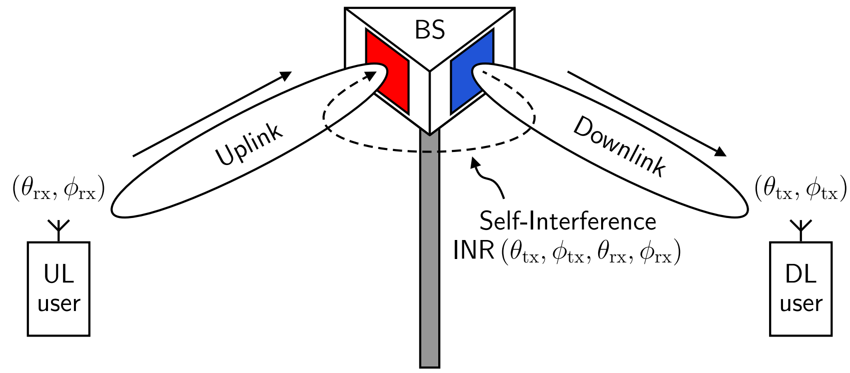

Figure 5.1: A full-duplex mmWave base station will incur different levels of self-interference depending on where its transmit beam and receive beam are steered. In this work, using the 28 GHz phased array platform in Figure 5.3, we swept the beams over a dense spatial profile, measuring self-interference for each transmit-receive beam pair.

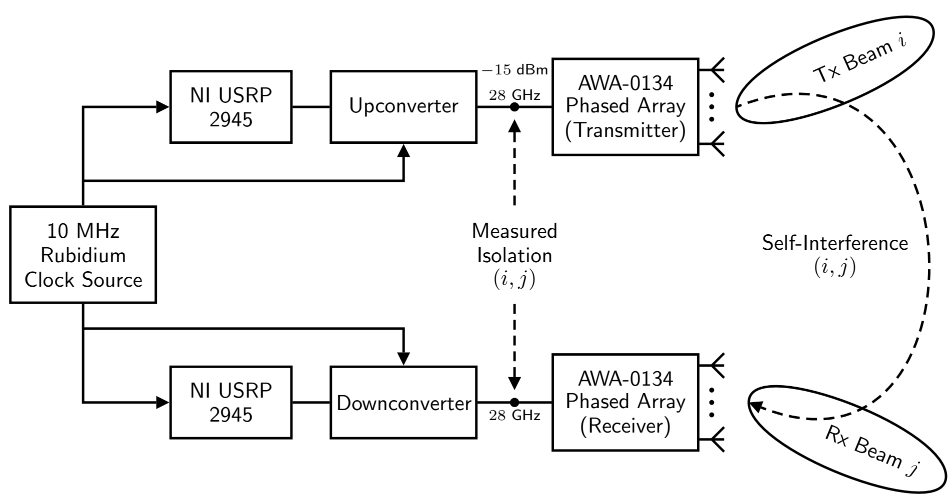

Figure 5.2: A simplified block diagram of our measurement setup using phased arrays to inspect self-interference at 28 GHz. The degree of self-interference depends on the particular choice of transmit beam and receive beam.

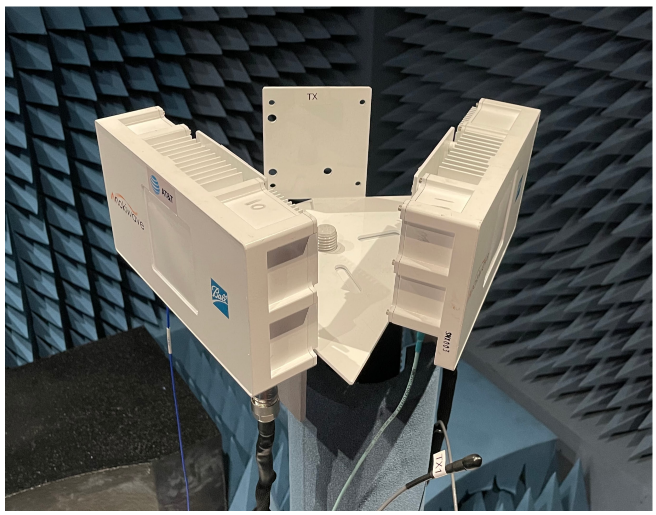

Figure 5.3: The 28 GHz phased array measurement platform in an anechoic chamber; receive array on left and transmit array on right. Each array is a 16×16 planar array. Self-interference was measured at the receive array output by simultaneously transmitting and receiving over the same frequency spectrum.

Our mmWave self-interference measurement system, illustrated as a block diagram in Figure 5.2, is comprised of two identical Anokiwave AWA-0134 28 GHz phased array modules [114]: one for transmission and one for reception. Each phased array module consists of a 16×16 half-wavelength UPA, offering high spatial resolution in azimuth and elevation. The transmit and receive arrays were mounted to separate sides of a 60° equilateral triangular platform, as shown in Figure 5.3, where the centers of the arrays were separated by 30 cm. This configuration aligns with practical, multi-panel (sectorized) full-duplex mmWave deployments, such as for full-duplex IAB as proposed in 3GPP [115]. The measurement platform was placed in an anechoic chamber free from significant reflectors; valuable future work would investigate the impact of reflections.

Each array can be electronically steered by a network of digitally controlled analog beamforming weights, allowing us to form narrow transmit and receive beams to inspect the directional characteristics of the direct coupling between the transmit and receive arrays. The transmit array is driven by an upconverted 28 GHz Zadoff-Chu sequence with 100 MHz of bandwidth and a power level of -15 dBm. The amplified and beamformed transmit signal is radiated by the transmit array at an effective isotropic radiated power (EIRP) of 60 dBm before coupling over the air with the receive array. The beamformed signal captured by the receive array is internally amplified and downconverted before being digitized. Separate software-defined radio platforms [116] were used to generate signals in the transmit chain and capture signals in the receive

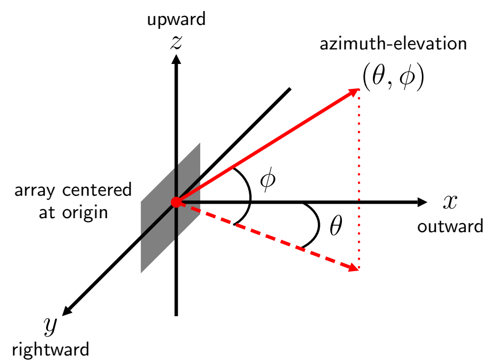

Figure 5.4: The coordinate system observed by each phased array situated at its respective origin, facing toward the positive x axis. An azimuth-elevation of $(0°, 0°)$ corresponds to steering outward (broadside) along the x axis. An increase in azimuth is rightward, toward the y axis. An increase in elevation is upward, toward the z axis.

chain. In baseband, the transmit and receive signals were processed to estimate the isolation from the input of the transmit array to the output of the receive array. In an effort to more reliably measure isolation, we employed a single Rubidium oscillator and a custom lossless dual upconversion/downconversion system built by Mi-Wave [117]. Correlation-based processing of the Zadoff-Chu signals was used to estimate isolation levels that extend well below the noise floor at the output of our receive array. We calibrated and verified our measurement capability using high fidelity test equipment [118] and stepped attenuators [119] to ensure our isolation measurements had low error (typically less than 1 dB) across a broad range of received power levels (roughly from -20 dBm to -110 dBm). Moreover, we confirmed the repeatability of our measurements in both the short-term (milliseconds) and long-term (minutes).

To describe the arrays' steering directions in 3-D, we use an azimuth-

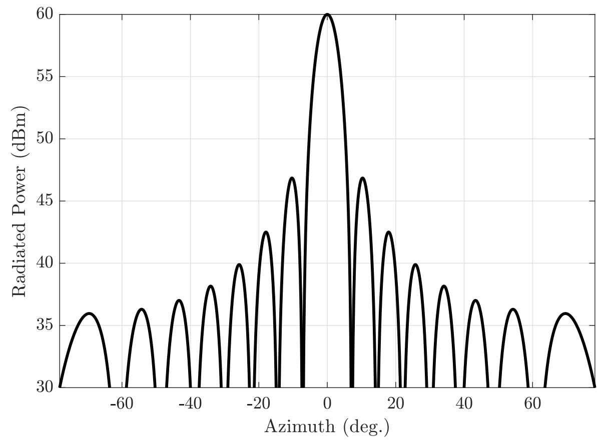

Figure 5.5: The idealized azimuth radiation pattern of our 16×16 transmit UPA steered broadside. The elevation pattern is identical. The EIRP is 60 dBm, and the 3 dB beamwidth is roughly 7°. The pattern of the receive UPA steered broadside is identical.

elevation convention as illustrated in Figure 5.4. We assume independent coordinate systems for each array, where each array is centered at the origin facing the positive x axis. From each array's perspective, broadside corresponds to 0° in azimuth and elevation (along the positive x axis), upward is an increase in elevation ϕ (positive z direction), and rightward is an increase in azimuth θ (positive y direction). Each array can be independently steered toward a relative azimuth-elevation (θ, ϕ) via beamforming weights. In this work, we employ conjugate beamforming (i.e., matched filter beamforming), where a beam is formed in a particular direction by setting beam weights equal to the conjugate of the array response in that direction. For context, the idealized pattern of a transmit beam steered broadside is shown in Figure 5.5, which has a 3 dB beamwidth of around 7° in both azimuth and elevation; the shape of a broadside receive beam is identical. Naturally, practical beam patterns will not be as well-defined nor exhibit the deep nulls as that shown in Figure 5.5.

The key power levels and power ratios associated with our measurement setup are summarized in Figure 5.7. When steering the transmit array toward some $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and receive array toward $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$, the power of self-interference coupled between the arrays is

$$P_{\mathrm{SI}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = P_{\mathrm{tx}} \cdot L(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})^{-1} \tag{5.1}$$

at the receive array output, where $P_{\mathrm{tx}} = -15$ dBm is the power into the transmit array and

$$L(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = \frac{1}{\left|\mathbf{w}(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})^{*}\mathbf{H}\mathbf{f}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})\right|^{2}} \tag{5.2}$$

is the effective isolation between the transmit array input and receive array output established by transmit beamforming weights $\mathbf{f}$ and receive weights $\mathbf{w}$; $\mathbf{H} \in \mathbb{C}^{256\times256}$ is the (unknown) over-the-air self-interference channel matrix between the arrays and has scaling that accounts for the inherent path loss along with transmit and receive gains. In other words, self-interference power P_{SI} includes the spatial coupling between the transmit and receive beams with the over-the-air channel along with large-scale power gains introduced by the transmit array (e.g., PAs) and receive array (e.g., LNAs).

Sets of N_{tx} transmit directions and N_{rx} receive directions

$$\mathcal{A}_{\mathrm{tx}} = \left\{\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right)\right\}_{i=1}^{N_{\mathrm{tx}}}, \qquad \mathcal{A}_{\mathrm{rx}} = \left\{\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)\right\}_{j=1}^{N_{\mathrm{rx}}} \tag{5.3}$$

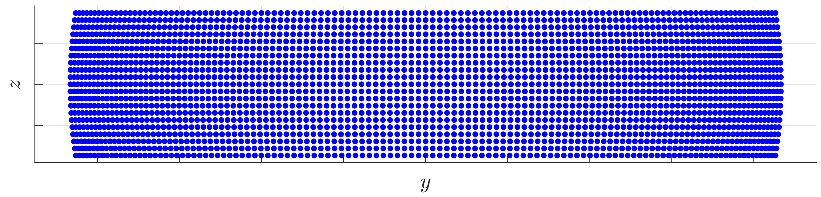

Figure 5.6: The set of 2,541 transmit directions $\mathcal{A}_{\text{tx}}$ used during measurement, shown here as their respective projections onto the y-z plane, where looking into the page represents steering outward from the array (the positive x direction). The set of receive directions $\mathcal{A}_{\text{rx}}$ is identical, densely spanning -60° to 60° in azimuth and -10° to 10° in elevation, each in 1° steps.

are specified prior to executing our measurement campaign. We measure the self-interference power between each transmit and receive steering combination for a total of $N_{\text{tx}} \times N_{\text{rx}}$ measurements. Depicted in Figure 5.6, the results in this work are based on measurements whose transmit and receive directions are each distributed uniformly in azimuth from -60° to 60° with 1° spacing and in elevation from -10° to 10° with 1° spacing. This amounts to $N_{\text{tx}} = N_{\text{rx}} = 121 \times 21 = 2541$ directions for transmission and for reception, totaling $2541 \times 2541 \approx 6.5$ million self-interference power measurements.

Referencing the measured self-interference power P_{SI} to the noise floor of the receive array, the INR for given transmit and receive directions can be written as

$$\mathsf{INR}(\theta_{\text{tx}}, \phi_{\text{tx}}, \theta_{\text{rx}}, \phi_{\text{rx}}) = \frac{P_{\text{SI}}(\theta_{\text{tx}}, \phi_{\text{tx}}, \theta_{\text{rx}}, \phi_{\text{rx}})}{P_{\text{noise}}} \tag{5.4}$$

where $P_{\text{noise}} = -68$ dBm is the noise power at the receive array output over 100 MHz. INR is an important quantity for describing if a full-duplex system is self-interference-limited ($\mathsf{INR} \gg 0$ dB), noise-limited ($\mathsf{INR} \ll 0$ dB), or

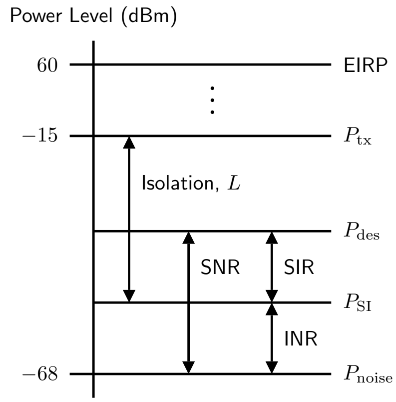

Figure 5.7: A summary of the key power levels and power ratios associated with our measurement system.

somewhere in between. For full-duplex, we desire low INR, roughly $\mathsf{INR} \leq 0$ dB in most cases, to ensure self-interference does not erode full-duplexing gain. The set of all nearly 6.5 million measured INR values we write as

$$\mathcal{I} = \left\{\mathsf{INR}(\theta_{\rm tx}, \phi_{\rm tx}, \theta_{\rm rx}, \phi_{\rm rx}) : (\theta_{\rm tx}, \phi_{\rm tx}) \in \mathcal{A}_{\rm tx}, (\theta_{\rm rx}, \phi_{\rm rx}) \in \mathcal{A}_{\rm rx}\right\} \tag{5.5}$$

and refer to the INR measured when transmitting with the i-th transmit beam and receiving with the j-th receive beam as

$$\mathsf{INR}_{ij} \triangleq \mathsf{INR}\left(\theta_{\rm tx}^{(i)}, \phi_{\rm tx}^{(i)}, \theta_{\rm rx}^{(j)}, \phi_{\rm rx}^{(j)}\right). \tag{5.6}$$

A *desired* signal having received power $P_{\rm des}$ (at the output of the receive array) would see a SINR of

$$\mathsf{SINR}(\theta_{\rm tx}, \phi_{\rm tx}, \theta_{\rm rx}, \phi_{\rm rx}) = \frac{P_{\rm des}}{P_{\rm noise} + P_{\rm SI}(\theta_{\rm tx}, \phi_{\rm tx}, \theta_{\rm rx}, \phi_{\rm rx})} \tag{5.7}$$

$$= \frac{\mathsf{SNR}}{1 + \mathsf{INR}(\theta_{\rm tx}, \phi_{\rm tx}, \theta_{\rm rx}, \phi_{\rm rx})} \tag{5.8}$$

where SNR is the SNR of the desired signal. Notice that SINR depends on the level of self-interference incurred when steering the transmitter toward $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and the receiver toward $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$; of course, SNR would practically also be a function of $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$. This work is solely concerned with measuring INR (self-interference), from which desired signals having some SNR can be evaluated in a full-duplex sense. We would like to point out that all measurements collected in this campaign are for a fixed setup as described; valuable future work would investigate the impact of system parameters such as beam shape (e.g., beamwidth and side lobe levels), array sizes, and panel geometries.

5.3 High-Level Summary and Spatial Insights

Perhaps the best summary of our measurements is the CDF of the nearly 6.5 million measured INR values in Figure 5.8. The maximum and minimum measured INR were nearly 46.99 dB and -44.57 dB, respectively. The measured INR typically falls between 0 dB and 40 dB, with median INR at 20.27 dB. Nearly 99% of beam pairs offer an INR greater than 0 dB, where self-interference power exceeds noise power. Around 90% of beam pairs yield an $\mathsf{INR} \geq 10$ dB, where self-interference is at least ten times as strong as noise. Just over 2% of beam pairs offer an $\mathsf{INR} \leq 3$ dB. Naturally, this CDF of INR would shift left/right if transmit power were to decrease/increase or noise power were to increase/decrease, for instance; those wishing to appropriately translate our measurements to systems with different power levels can refer to Figure 5.7.

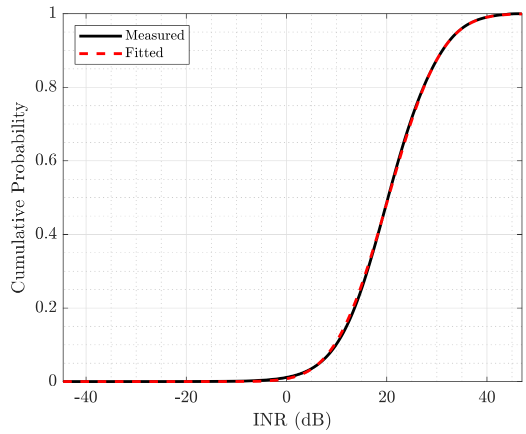

Figure 5.8: The CDF of the nearly 6.5 million measured INR values and its fitted log-normal distribution.

Remark 5.1. At first glance, the CDF of the measured INR values seems quite pessimistic from a full-duplex perspective, considering most beam pairs yield self-interference levels that are well above the noise floor. Multi-panel full-duplex mmWave systems similar to ours will typically be overwhelmed with self-interference when choosing a random transmit and receive beam, motivating the need for additional means to reduce self-interference.[1] Nonetheless, it is a welcome sight to observe that there exist select beam pairs that do in fact offer INR levels sufficiently low for full-duplex—we explore the directional nature of these beam pairs shortly in Subsection 5.3.2. Valuable future work would explore solutions to reduce self-interference and investigate how system design choices can potentially reduce INR levels while maintaining service to users.

[1]We explore one method for reducing self-interference via small shifts of the transmit and receive beams in Section 5.4.

We found that the CDF in Figure 5.8 can be well approximated by a log-normal distribution, shown as a dashed line in Figure 5.8. That is to say that the measured INR values in dB approximately follow a normal distribution as

$$[\mathcal{I}]_{\mathrm{dB}} \overset{\mathrm{fit}}{\sim} \mathcal{N}\left(\mu, \sigma^2\right) \tag{5.9}$$

where $\mu = 20.32$ and $\sigma^2 = 70.69$ are the fitted mean and variance of the normal distribution. Like the CDF in Figure 5.8, changes to large-scale parameters that impact INR—such as the transmit power or noise power—can be accounted for in μ to shift the fitted normal distribution left or right. Albeit limited, engineers can make first-order statistical approximations of self-interference via this log-normal distribution when drawing i.i.d. INR values as $[\mathsf{INR}]_{\mathrm{dB}} \sim \mathcal{N}(\mu, \sigma^2)$. For instance, the probability that the INR of a random transmit-receive beam pair will fall below γ (in linear units) can be well approximated as

$$\mathbb{P}[\mathsf{INR} \leq \gamma] = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{[\gamma]_{\mathrm{dB}} - \mu}{\sigma \cdot \sqrt{2}}\right)\right]. \tag{5.10}$$

5.3.1 Is an Idealized Near-Field Channel Model Realistic?

A natural question to ask before proceeding is if the aforementioned spherical-wave MIMO channel model [2]—an idealized near-field propagation model—aligns with our measurements. If it does, self-interference power values could be realized deterministically via the product of transmit and receive beamforming weights $\mathbf{f}$ and $\mathbf{w}$ with a MIMO channel matrix $\mathbf{H}$ based on the

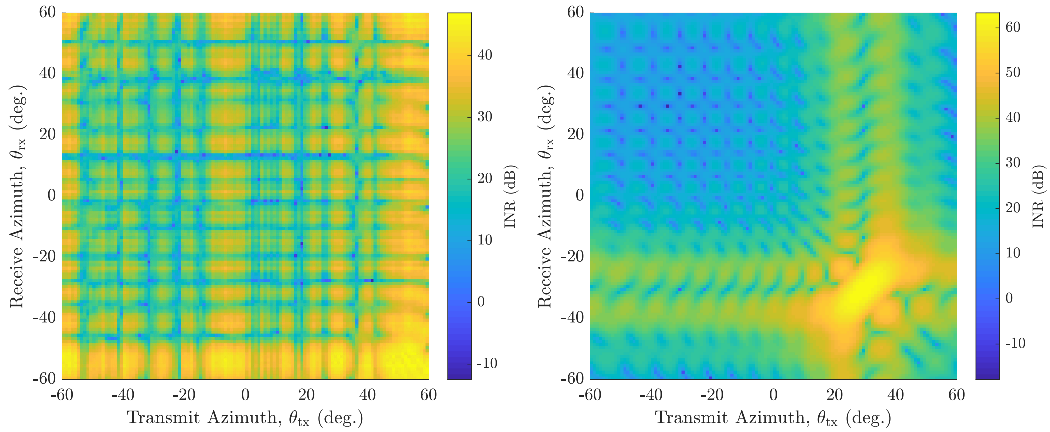

(a) Measured INR.

(b) Simulated INR based on [2].

Figure 5.9: (a) Measured INR as a function of azimuthal transmit and receive directions when $\phi_{\text{tx}} = \phi_{\text{rx}} = 0°$. (b) The simulated counterpart of (a) when using a spherical-wave near-field self-interference channel model [2], which clearly does not align well with what was measured. This motivates the need for a new, measurement-backed model of mmWave self-interference.

spherical-wave model. Unfortunately, however, we found that the spherical-wave MIMO channel model does not align with our measurements. Consider Figure 5.9, where we plot the measured INR values across the azimuth plane and the simulated counterpart using the spherical-wave MIMO channel model. Notice that the two are starkly different, indicating that this idealized near-field channel model—which has been used so frequently as a means to evaluate mmWave full-duplex—does not translate to practical systems, which pose a number of nonidealities stemming from array enclosures, mounting infrastructure, and non-isotropic antenna elements, for instance. This motivates the need for a practical, measurement-backed MIMO channel model for mmWave self-interference, which we plan to address in future work.

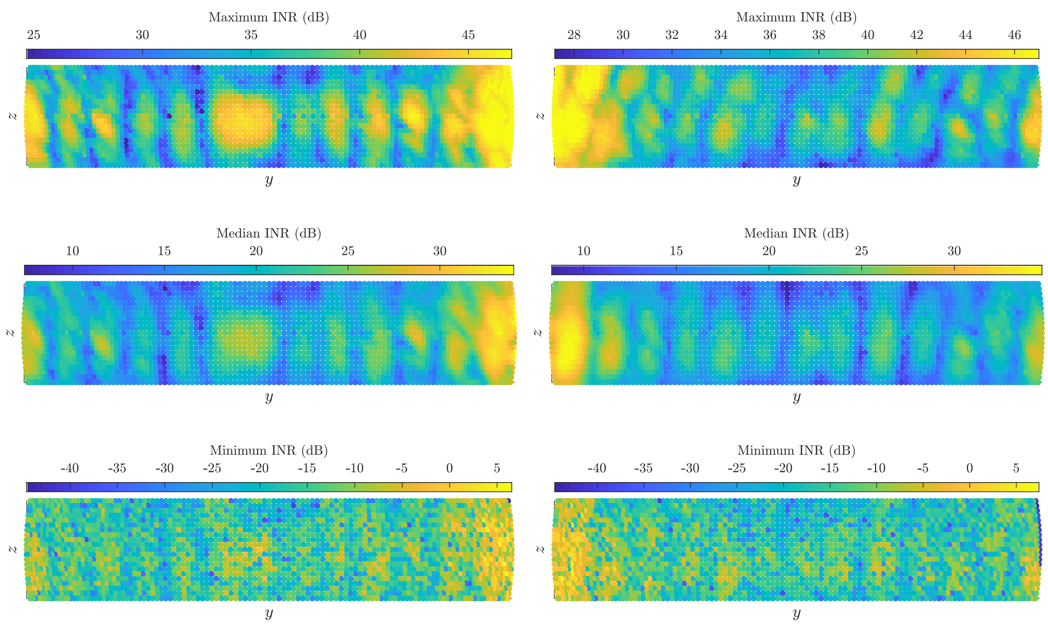

(a) Observed by each transmit beam. (b) Observed by each receive beam.

Figure 5.10: For each transmit beam and receive beam, shown are the maximum, median, and minimum INR across all receive and transmit beams, respectively.

5.3.2 Maximum, Median, and Minimum INR

The CDF in Figure 5.8 and its corresponding fitted distribution are certainly useful *statistically* but do not provide any *spatial* insight on self-interference. As such, we now hone in on narrower perspectives to better visualize and interpret our measurements spatially. First, let us begin by considering Figure 5.10a, which shows the maximum, median, and minimum INR observed by each transmit beam across all receive beams; each dot corresponds to a transmit beam's projection onto the y-z plane (i.e., its direction from the perspective of the transmit array). In other words, the maximum

INR observed by the i-th transmit beam (the i-th dot) is simply $\max_j \mathsf{INR}_{ij}$, for example, with median and minimum expressed analogously. Figure 5.10b similarly shows these statistics observed when receiving a particular direction. Referencing Figure 5.10b, we can see that the median INR per receive beam ranges from approximately 8 dB to 35 dB. The maximum INR observed at each receive beam is at least around 28 dB and at most over 46 dB, while the minimum INR is at least around -45 dB and at most around 7 dB. In a similar fashion, we examine these statistics for each transmit beam in Figure 5.10a, which tell a similar story both visually and numerically as the receive side.

Remark 5.2: Self-interference depends jointly on the transmit and receive beams. There are a few important things to take away from Figure 5.10a and Figure 5.10b. As intuition may suggest based on Figure 5.3, the results of Figure 5.10a and Figure 5.10b indicate that:

- transmitting to toward the receiver tends to couple more self-interference
- receiving toward the transmitter tends to couple more self-interference.

Considering minimum INR is at most around 7 dB in both, we see that (i) even when steering our transmitter toward the receiver, there exist some receive beam(s) that offer low INR (at most around 7 dB) and (ii) even when steering our receiver toward the transmitter, there exist some transmit beam(s) that offer low INR (at most around 7 dB). This suggests that—while transmitting toward the receiver and receiving toward the transmitter *generally* results in more self-interference—there exist receive beams and transmit beams that *can* offer low INR. In Section 5.4, we observe that low-INR beam pairs appear to be distributed throughout space. We have ongoing work that investigates if these low-INR beam pairs can in fact be used to serve users with high beamforming gain while simultaneously offering reduced self-interference, facilitating full-duplex operation. In a similar fashion, observing maximum INR illustrates that there also consistently exists transmit-receive combinations that can lead to high self-interference. From this, we can conclude that there are not transmit beams nor receive beams that *universally* offer low or high INR—though there exist those that tend to. Rather, the amount of self-interference coupled depends heavily on one's choice of transmit beam *and* receive beam.

Remark 5.3: Spatial insights. Additional takeaways include the fact that we observe strong similarities between the transmit and receive profiles, which validates some degree of channel symmetry. However, there do exist noteworthy differences, particularly the strong self-interference present when transmitting around broadside but not when receiving. The high self-interference coupled when transmitting around broadside is not necessarily expected nor easily explained; it can perhaps be attributed to coupling behind the arrays due to mounting hardware and array enclosures. Also, we see significantly more variation across y than z, suggesting that the azimuth of the steering direction plays a greater role than elevation, which one may expect since our transmitter and receiver are separated in azimuth but not in elevation. While it may seem obvious that transmitting toward the receiver and receiving toward the transmitter would couple the most self-interference, it was not clear that this would be the case since the transmit and receive arrays exist in the near-field of one another. The far-field distance of our arrays is approximately 2.4 meters based on the rule-of-thumb $2D^2/\lambda$ [97], while our arrays are separated by only 30 cm. The reactive/radiating near-field boundary, on the other hand, is around a mere 23 cm based on the rule-of-thumb $0.62\sqrt{D^3/\lambda}$ [97], suggesting that our arrays live just within the radiating near-field of one another. Operating in this near-field regime, the highly directional beams created by our UPAs are not necessarily "highly directional" from the perspective of one another [20].

5.3.3 Meeting INR Thresholds

Having looked at the maximum, median, and minimum INR for particular transmit beams and receive beams, we now examine the fraction of beams that offer certain levels of INR. In Figure 5.11a, for each transmit beam, we look at the fraction of receive beams that offer at most 10 dB, 20 dB, and 30 dB of INR. Similarly, in Figure 5.11b, for each receive beam, we look at the fraction of transmit beams that offer these same levels of INR. From the top plot of each, we see that a modest INR threshold of 10 dB (where self-interference is ten times stronger than noise) cannot be met very reliably by any transmit beam nor any receive beam. At best, select beams can only meet this target INR around 50% of the time, with the vast majority falling quite short of this. Naturally, as the INR threshold rises to 20 dB, the fraction of beams that can meet this threshold increases. Transmit and receive options emerge that offer an INR of at most 20 dB over 50% of the time, with some approaching 100%.

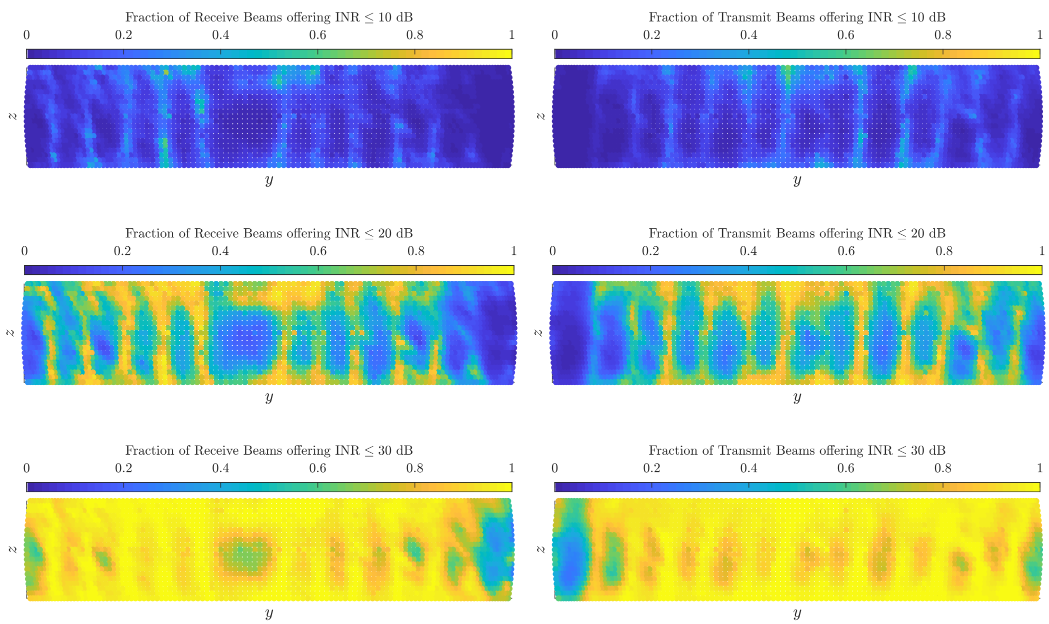

(a) Observed by each transmit beam.

(b) Observed by each receive beam.

Figure 5.11: For each transmit beam and receive beam, shown is the fraction of receive beams and transmit beams, respectively, that meet various INR thresholds.

Still, however, the transmit beams steering rightward toward the receiver and the receive beams steering leftward toward the transmitter struggle to offer an INR below 20 dB. Increasing the INR threshold further, nearly all transmit and receive beams can confidently offer an INR within 30 dB, though those least likely to do so are the rightward transmit beams, leftward receive beams, and broadside transmit beams.

Remark 5.4: Approximate transmit and receive nulls. The most promising transmit beams and receive beams in meeting an INR threshold of 20 dB, for example, can be seen as thin vertical strips of bright yellow. These vertical strips, which were also visible as low-INR beams in Figure 5.10, are likely attributed to nulls in the transmit and receive beam patterns. Notice, however, since the statistics of Figure 5.11 were taken over all transmit/receive beams, it shows that the transmit nulls are robust to some degree, somewhat reliably offering lower INR regardless of the receive beam being used (and vice versa). These transmit and receive beams offering lower INR across large fractions of receive beams and transmit beams, respectively, correspond to the approximate transmit and receive nulls at the channel input and output (i.e., approximate right and left null spaces of $\mathbf{H}$), respectively. Recall, from Figure 5.10a and Figure 5.10b, we did not see any transmit beams or receive beams that *universally* provided high isolation. If indeed these vertical strips are attributed to nulls in the beam patterns, this suggests that the self-interference channel between the transmit and receive arrays is quite directional, which somewhat further bucks the thought that near-field interaction dominates their coupling. Our future work will explore this to better understand the coupling nature of the arrays.

5.3.4 INR for Particular Transmit-Receive Beam Pairs

Honing in further, we now look at the isolation achieved at each transmit beam *for a particular receive beam* and at each receive beam *for a particular transmit beam*, as shown in Figure 5.12a and Figure 5.12b. Let us first consider the INR observed across receive beams for particular transmit beams; imagine fixing the transmit beam and sweeping the receive beam to measure

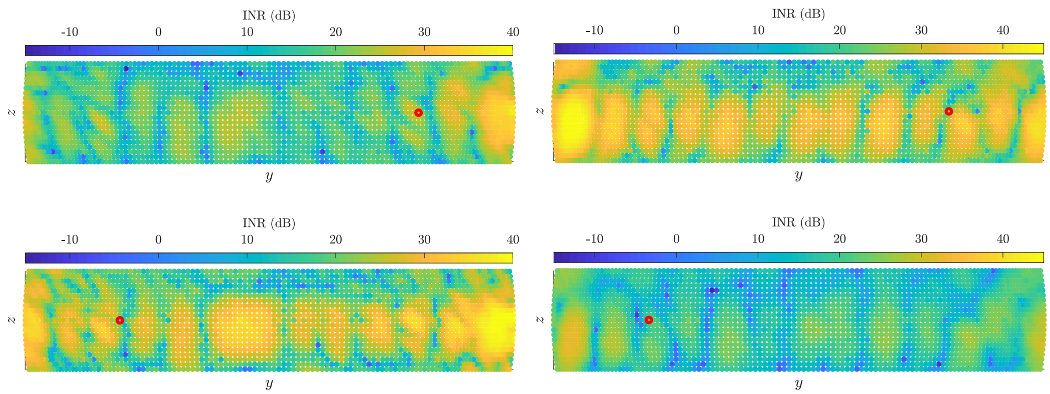

(a) Observed by each transmit beam.

(b) Observed by each receive beam.

Figure 5.12: For each transmit beam and receive beam, shown is the measured INR across all receive and transmit beams, respectively. The red circle indicates the (a) receive direction and (b) transmit direction.

INR at each. In Figure 5.12b, we have selected two transmit directions: toward the receive array (top plot) and away from the receive array (bottom plot). For each, have shown the INR measured between the transmit beam and each receive beam option. Shown in the top plot of Figure 5.12b, when transmitting rightward toward the receiver (whose direction is shown as a red circle), we see fairly high INR across the receive profile. Large orange/yellow spots make up most of the receive profile, highlighting just how difficult it may be to find a receive beam that offers low INR for this particular transmit beam. There exist some low-INR receive options narrowly in between large spots of orange or at high and low elevation.

Now, looking at the bottom plot of Figure 5.12b, when steering the transmitter leftward away from the receiver, the INR profile across receive beams expectedly changes. The INR profile sees a widespread decrease of about 10 dB or more and options for extremely low INR are more available. Still, receiving leftward toward the transmitter remains the least attractive option and reinforces that isolation may tend to be lower when transmitting rightward toward the receiver and when receiving leftward toward the transmitter—but this is not universally the case. Looking at both plots in Figure 5.12b, the receive beam that offers minimum INR varies with transmit beam, which further backs our claim that there are not receive beams that universally offer low INR. Moreover, the low-INR receive directions are typically quite narrow in the sense that small changes in receive direction can lead to significant changes in isolation. For instance, when transmitting rightward toward the receiver, the INR across receive beams varies by about 60 dB, and we see that shifting a receive beam by only 1° to 2° in azimuth and/or elevation can lead to changes of 20–30 dB or more in INR. Notice that this sensitivity to steering direction is much more apparent with low-INR beams than high-INR ones.

Similarly, in Figure 5.12a, we have selected two receive directions and, for each, have shown the INR measured between the receive beam and each transmit beam. Analogous conclusions can be drawn as with Figure 5.12b, though there are useful comments to make. Again, varying with each receive beam, there exists an INR-minimizing transmit beam. Notice that even when the receive beam is steered away from the transmit array (to the right; top plot), transmitting toward the receive array (to the right) still inflicts substantial self-interference. We can clearly see that simply steering the transmitter away from the receiver *or* steering the receiver away from the transmitter does not offer widespread low INR. Comparing Figure 5.12a and Figure 5.12b, we observe a certain degree of symmetry. Transmitting toward the receiver (top Figure 5.12b) is similar to receiving toward the transmitter (bottom Figure 5.12a). Transmitting away from the receiver (bottom Figure 5.12b) is similar to receiving away from the transmitter (top Figure 5.12a). This further verifies a sense of spatial symmetry of our self-interference channel $\mathbf{H}$.

Remark 5.5: Large-scale trends and small-scale variability. Figure 5.12a and Figure 5.12b highlight that there exist large-scale (global) trends in the amount of self-interference coupled between the transmit and receive arrays, since general steering direction of the transmitter and receiver can play a significant role in the INR profile. In addition, they also illustrate the local phenomena present in the INR profile: small shifts in steering direction can have drastic impacts on the degree of self-interference coupled. Figure 5.12a and Figure 5.12b showed this sensitivity of the transmit beam and receive beam separately—in the next section, we investigate this sensitivity when the transmit beam *and* receive beam both see small shifts in their steering direction.

5.4 Angular Spread of mmWave Self-Interference

In this section, we inspect how self-interference varies with small changes in transmit and receive directions. Let us begin by defining $\measuredangle(\theta_1, \theta_2)$ as the

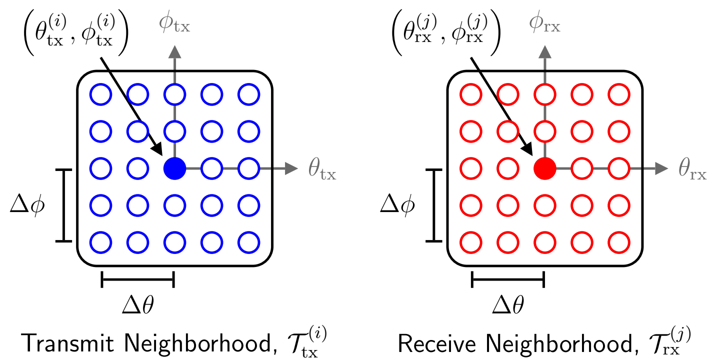

Figure 5.13: The transmit and receive steering neighborhoods $\mathcal{T}_{\mathrm{tx}}^{(i)}$ and $\mathcal{T}_{\mathrm{rx}}^{(j)}$, where the filled circles indicate the nominal steering directions $\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)$ and the unfilled circles indicate directions comprising the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding each. In this case, the 25 directions in each yields $25^2 = 625$ transmit-receive steering combinations.

absolute difference between two angles θ_1, θ_2 (in degrees), written as

$$\measuredangle(\theta_1, \theta_2) = \begin{cases} \zeta, & \zeta \leq 180^\circ \\ 360^\circ - \zeta, & \zeta > 180^\circ \end{cases} \tag{5.11}$$

where $\zeta = |\theta_1 - \theta_2| \bmod 360^\circ$ and mod is the modulo operator. Let $\mathcal{T}_{\mathrm{tx}}^{(i)}$ and $\mathcal{T}_{\mathrm{rx}}^{(j)}$ be the $(\Delta\theta, \Delta\phi)$-neighborhoods around the i-th transmit direction and j-th receive direction, respectively, defined as

$$\mathcal{T}_{\mathrm{tx}}^{(i)}(\Delta\theta, \Delta\phi) = \left\{(\theta, \phi) \in \mathcal{A}_{\mathrm{tx}} : \measuredangle\left(\theta, \theta_{\mathrm{tx}}^{(i)}\right) \leq \Delta\theta, \measuredangle\left(\phi, \phi_{\mathrm{tx}}^{(i)}\right) \leq \Delta\phi\right\} \tag{5.12}$$

$$\mathcal{T}_{\mathrm{rx}}^{(j)}(\Delta\theta, \Delta\phi) = \left\{(\theta, \phi) \in \mathcal{A}_{\mathrm{rx}} : \measuredangle\left(\theta, \theta_{\mathrm{rx}}^{(j)}\right) \leq \Delta\theta, \measuredangle\left(\phi, \phi_{\mathrm{rx}}^{(j)}\right) \leq \Delta\phi\right\} \tag{5.13}$$

and illustrated in Figure 5.13. For some $(\Delta\theta, \Delta\phi)$ in degrees, the cardinality of these sets is

$$\left|\mathcal{T}_{\mathrm{tx}}^{(i)}(\Delta\theta, \Delta\phi)\right|, \left|\mathcal{T}_{\mathrm{rx}}^{(j)}(\Delta\theta, \Delta\phi)\right| \leq (2 \cdot \Delta\theta + 1) \cdot (2 \cdot \Delta\phi + 1) \tag{5.14}$$

with equality when not at the edge of the measurement space (which is typical); the crude use of $\Delta\theta$ and $\Delta\phi$ here is thanks to our 1° spacing of $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$.

Then, let $\mathcal{I}_{ij}(\Delta\theta, \Delta\phi)$ be the set of measured INR values across the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding the (i, j)-th transmit-receive beam pair, expressed as

$$\mathcal{I}_{ij}(\Delta\theta, \Delta\phi) = \left\{ \mathsf{INR}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) : (\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}) \in \mathcal{T}_{\mathrm{tx}}^{(i)}(\Delta\theta, \Delta\phi), (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \in \mathcal{T}_{\mathrm{rx}}^{(j)}(\Delta\theta, \Delta\phi) \right\} \tag{5.15}$$

where $\mathcal{T}_{\mathrm{tx}}^{(i)}$ and $\mathcal{T}_{\mathrm{rx}}^{(j)}$ depends on the beam pair (i, j) and neighborhood size $(\Delta\theta, \Delta\phi)$. As the $(\Delta\theta, \Delta\phi)$-neighborhoods are widened, the cardinality of $\mathcal{I}_{ij}$ grows, which is simply the product of that of $\mathcal{T}_{\mathrm{tx}}^{(i)}$ and $\mathcal{T}_{\mathrm{rx}}^{(j)}$.

$$\left|\mathcal{I}_{ij}(\Delta\theta, \Delta\phi)\right| = \left|\mathcal{T}_{\mathrm{tx}}^{(i)}(\Delta\theta, \Delta\phi)\right| \cdot \left|\mathcal{T}_{\mathrm{rx}}^{(j)}(\Delta\theta, \Delta\phi)\right| \tag{5.16}$$

Based on (5.14), the upper bound of (5.16) is tabulated for various $(\Delta\theta, \Delta\phi)$ in Figure 5.14, which grows with order $\mathcal{O}(\Delta\theta^2 \cdot \Delta\phi^2)$.

The minimum INR and maximum INR offered by beam pairs across the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding the (i, j)-th beam pair can be expressed as simply

$$\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi) = \min\left\{\mathcal{I}_{ij}(\Delta\theta, \Delta\phi)\right\} \tag{5.17}$$

$$\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi) = \max\left\{\mathcal{I}_{ij}(\Delta\theta, \Delta\phi)\right\}. \tag{5.18}$$

Using these, the INR range we define as

$$\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi) = \frac{\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)}{\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)} \geq 1 \tag{5.19}$$

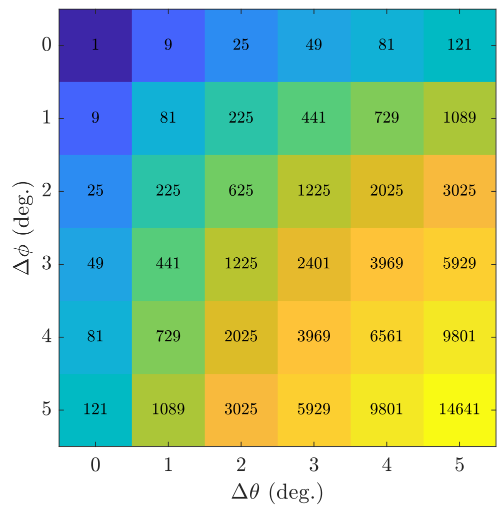

Figure 5.14: The number of transmit-receive beam pairs for a typical $(\Delta\theta, \Delta\phi)$-neighborhood under our $(1^\circ, 1^\circ)$ resolution.

which captures how much the INR can vary over the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding the (i, j)-th beam pair.

By examining $\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi)$, $\mathsf{INR}_{ij}^{\mathrm{max}}(\Delta\theta, \Delta\phi)$, and $\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ for each transmit-receive steering combination (i, j) and for variably sized neighborhoods, we can gain insight into the angular spread of self-interference. We point out that, since our measurements were taken with 1° resolution in azimuth and elevation, there exists the potential to see greater INR range, lower minimum INR, and/or higher maximum INR if sub-$(1^\circ, 1^\circ)$ resolutions were used; as such, the results herein can be considered a potentially conservative measure on these statistics over small neighborhoods.

5.4.1 INR Range Over Various Neighborhoods

In Figure 5.15a, we plot the CDF of INR range for variably sized $(\Delta\theta, \Delta\phi)$-neighborhoods across all measured direction pairs (i.e., each CDF contains nearly 6.5 million points). As shown in Figure 5.15a, moving a beam pair by only 1° in either azimuth *or* elevation can lead to notable changes in INR: around 25% of beam pairs observe over 10 dB of INR range in either case. As the neighborhood size increases, we naturally observe a wider range of INR. 50% of beam pairs see more than 17 dB of variability in INR across a $(1^\circ, 1^\circ)$-neighborhood. In other words, if we consider a beam pair at random and look around its $(1^\circ, 1^\circ)$-neighborhood, we would expect the INR to vary by 17 dB or more. Across a $(2^\circ, 2^\circ)$-neighborhood, 80% of beam pairs see around 25 dB or more of variability in INR. Notice that there exists slightly more variability in INR across azimuth than across elevation, evidenced by the $(1^\circ, 0^\circ)$- and $(0^\circ, 1^\circ)$-neighborhoods—perhaps due to the horizontal separation of our transmit and receive arrays.

To characterize $\mathsf{INR}_{ij}^{\mathrm{rng}}$ for variably sized $(\Delta\theta, \Delta\phi)$-neighborhoods, we have fit a distribution to $\left\{\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi)\right\}$. Specifically, we found that a Gamma distribution can be fitted to each of the CDFs in Figure 5.15a as follows

$$\left\{ \left[\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} \ \forall \ i,j \right\} \overset{\mathrm{fit}}{\sim} \mathrm{Gamma}(\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi), \beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi)) \quad (5.20)$$

where $\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi) > 0$ and $\beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi) > 0$ are the fitted shape and rate (inverse scale) of the Gamma distribution. The fitted Gamma distributions

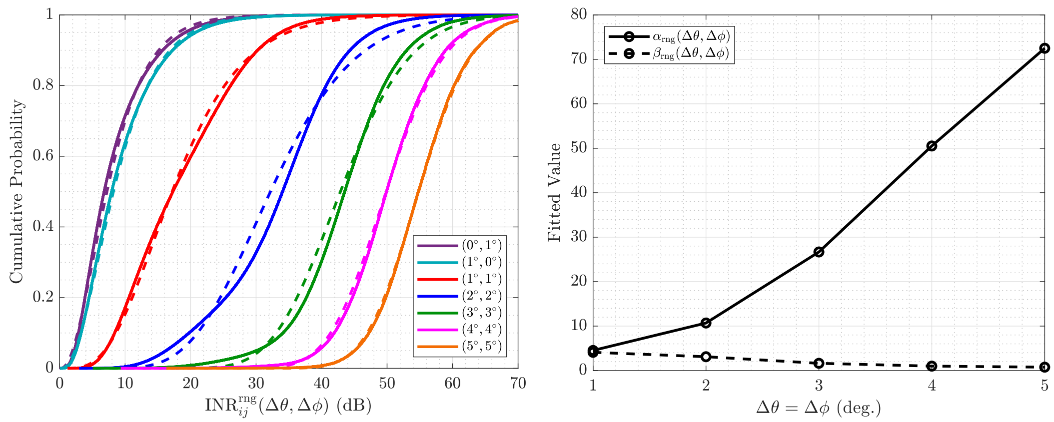

(a) $\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi)$.

(b) $\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ and $\beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$.

Figure 5.15: (a) The CDF of $\mathsf{INR}_{ij}^{\mathrm{rng}}$ across all nearly 6.5 million transmit-receive beam pairs for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. Fitted distributions for each are shown as dashed lines. (b) The fitted parameters $\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ and $\beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$.

for each neighborhood in Figure 5.15a are shown as dashed lines. In addition, we have plotted $\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ and $\beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi)$ as functions of $\Delta\theta = \Delta\phi$ in Figure 5.15b. As $\Delta\theta = \Delta\phi$ increases, the shape parameter α_{rng} drastically increases and the rate parameter β_{rng} decays toward zero, which is a reflection of the CDFs in Figure 5.15a shifting rightward.

In addition to those shown in Figure 5.15, we fitted unique Gamma distributions for $\Delta\theta, \Delta\phi \in \{0°, 1°, \ldots, 5°\}$ and tabulated the fitted parameters $(\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi), \beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi))$ for each in Table 5.1. Engineers wishing to realize the range in INR over a random $(\Delta\theta, \Delta\phi)$-neighborhood or conduct statistical analyses related to such can refer to Table 5.1 for adequate Gamma distribution parameters $(\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi), \beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi))$. Then, the expected range in INR in dB over some $(\Delta\theta, \Delta\phi)$-neighborhood, for instance, can be approximated as

$$\mathbb{E}\left[\left[\mathsf{INR}_{ij}^{\mathrm{rng}}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}}\right] = \frac{\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi)}{\beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi)} \tag{5.21}$$

based on the Gamma distribution, along with an assortment of other statistics readily computed.

Remark 5.6. Figure 5.15a highlights that the self-interference channel is not spatially smooth. Rather, small changes in steering direction can result in significant changes in the degree of self-interference coupled and, hence, significant changes in full-duplex performance. As such, mmWave full-duplex systems cannot expect to reliably avoid self-interference by broadly steering transmit and receive beams. Instead, transmit and receive beams will need to be carefully (and *jointly*) steered, as small errors in steering direction can lead to drastic changes in self-interference.

5.4.2 Minimum INR Over Various Neighborhoods

Now, we examine the minimum INR that can be reached by each beam pair if allowed to deviate around some spatial neighborhood. To illustrate this, we have included Figure 5.16. In Figure 5.16a, at each transmit beam, we show the fraction of receive beams that can offer a minimum INR of 0 dB or less. We do this for various neighborhood sizes $(\Delta\theta, \Delta\phi)$, where a $(0^\circ, 0^\circ)$-neighborhood is simply no deviation. With a $(0^\circ, 0^\circ)$-neighborhood, most transmit-receive beams cannot meet the INR threshold of 0 dB. As the neighborhood grows

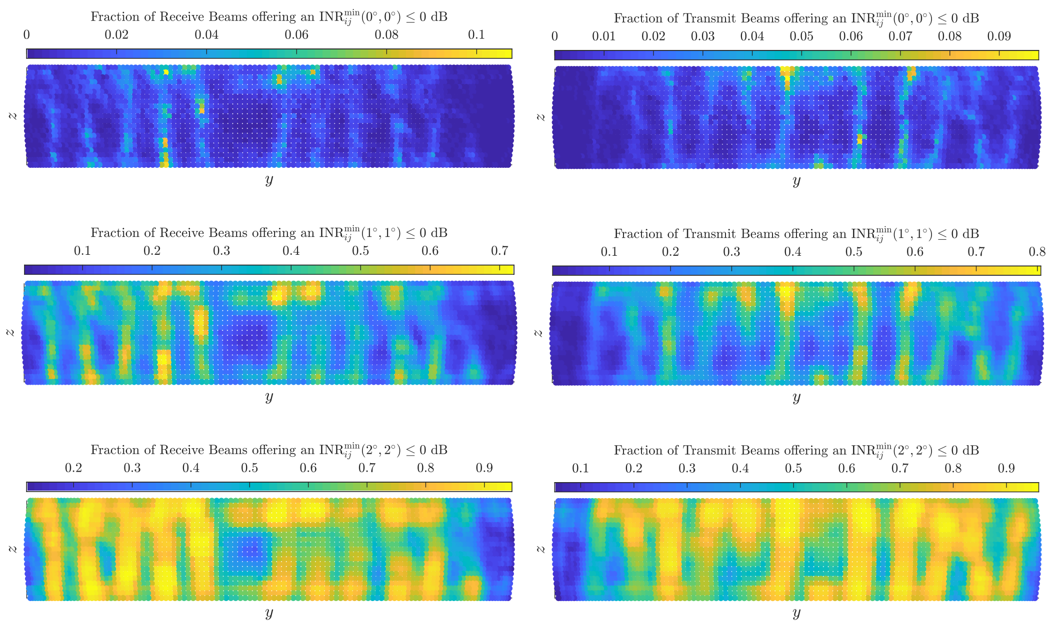

(a) Observed by each transmit beam. (b) Observed by each receive beam.

Figure 5.16: For each transmit beam and receive beam, shown are the fraction of receive beams and transmit beams, respectively, that can offer an INR of 0 dB or less when allowed to deviate by various $(\Delta\theta, \Delta\phi)$. Low INR becomes reliably within arm's reach by increasing $(\Delta\theta, \Delta\phi)$.

to $(1^\circ, 1^\circ)$, we see that, for some transmit beams, a large fraction of receive beams can reach an INR of 0 dB (and vice versa). With $(2^\circ, 2^\circ)$ of freedom, we see the clouds of yellow grow as more receive beams offer an INR of 0 dB for even more transmit beams. Figure 5.16 illustrates the significant changes observed in INR due to slight shifts of the transmit and receive beams and shows that INR levels suitable for full-duplex are in fact within arm's reach.

Consider Figure 5.17a, where we plot the CDF of $\left\{\mathsf{INR}_{ij}^{\min}\right\}$ for variably sized $(\Delta\theta, \Delta\phi)$-neighborhoods across all measured beam pairs. The dashed

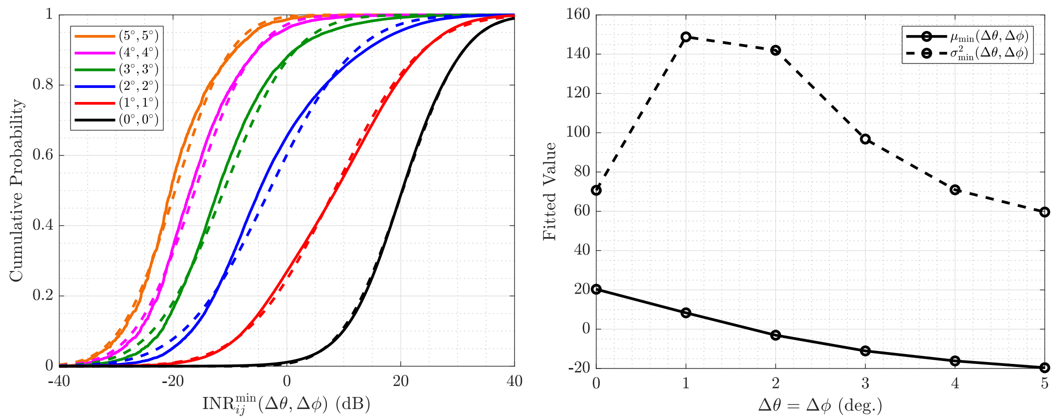

(a) $\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)$.

(b) $\mu_{\min}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\min}(\Delta\theta, \Delta\phi)$.

Figure 5.17: (a) The CDF of $\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)$ for various $(\Delta\theta, \Delta\phi)$. For each, the fitted distribution is shown as a dashed line. (b) The fitted parameters $\mu_{\min}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\min}(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$.

black line in Figure 5.17a is simply the CDF of $\{\mathsf{INR}_{ij}\}$ for each beam pair since $\Delta\theta = \Delta\phi = 0°$ (shown previously in Figure 5.8). When considering small neighborhoods around each beam pair, however, much more promising results are observed. When shifting beams by no more than 1° in azimuth and elevation, the probability of reaching $\mathsf{INR} \leq 0$ dB (where self-interference is no stronger than noise) grows to over 25% from around 1%. With 2°, it grows to over 65%; that is to say that over 65% of beam pairs are within $(2°, 2°)$ of a beam pair that offers $\mathsf{INR} \leq 0$ dB.

Remark 5.7: Slight shifts can significantly reduce self-interference. Clearly, INR can be greatly improved with slight shifts in the steering directions of the transmit and/or receive beams. From this, we draw an important conclusion: steering directions that do not *inherently* offer high isolation (i.e., low INR_{ij}) are likely spatially near ones that do. These are highly encouraging results for the potential of mmWave full-duplex since they suggest that self-interference can be greatly reduced while making very minor adjustments to the transmit and receive steering directions. This is reinforced further by the fact that our beams have a 3 dB beamwidth around 7°, meaning slight deviations will hopefully not sacrifice too much beamforming gain when making these adjustments. To reach these low-INR beam pairs, however, it may require searching over many beam pairs within a small spatial neighborhood, as highlighted by Figure 5.14. For instance, with our 1° resolution, a typical $(2^\circ, 2^\circ)$-neighborhood contains 625 transmit-receive beam pairs, highlighting that there may be practical hurdles in locating low-INR beam pairs within a given neighborhood. Exploring how side lobes, beamwidth, relative array geometry, and mounting infrastructure play a role in this small-scale variability would be valuable future work.

To conduct statistical analyses of $\mathsf{INR}_{ij}^{\min}$, we have fit a distribution to $\left\{\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)\right\}$ for variably sized $(\Delta\theta, \Delta\phi)$-neighborhoods. Specifically, we found that a normal distribution can be fitted to each of the CDFs in Figure 5.17a as follows

$$\left\{\left[\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} \ \forall \ i, j\right\} \overset{\mathrm{fit}}{\sim} \mathcal{N}\left(\mu_{\min}(\Delta\theta, \Delta\phi), \sigma_{\min}^2(\Delta\theta, \Delta\phi)\right) \qquad (5.22)$$

where $\mu_{\min}(\Delta\theta, \Delta\phi)$ and $\sigma_{\min}^2(\Delta\theta, \Delta\phi)$ are the fitted mean and variance of the normal distribution. The dashed lines in Figure 5.17a depict each neighborhood's fitted distribution, and Figure 5.17b show the fitted parameters for various $\Delta\theta = \Delta\phi$. Shown as the solid line in Figure 5.17b, the mean $\mathsf{INR}_{ij}^{\min}(\Delta\theta, \Delta\phi)$ steadily decreases by about 10 dB per unit increase in $\Delta\theta = \Delta\phi$ before beginning to saturate. The dashed line shows the variance of the fit, which initially increases and then decreases, which suggests that $(1^\circ, 1^\circ)$ and $(2^\circ, 2^\circ)$ of deviation can offer a reduction in INR but by highly variable amounts—evident also by their distributions in Figure 5.17a. The variance decreases as the majority of beam pairs can reach similar levels of INR with $(3^\circ, 3^\circ)$ or greater of deviation; the distributions in Figure 5.17a become more upright.

For neighborhood sizes where $\Delta\theta, \Delta\phi \in \{0^\circ, 1^\circ, \ldots, 5^\circ\}$, we tabulated the fitted parameters $(\mu_{\min}(\Delta\theta, \Delta\phi), \sigma^2_{\min}(\Delta\theta, \Delta\phi))$ in Table 5.2. To conduct statistical analyses related to minimum INR, engineers can refer to Table 5.2 for adequate distribution parameters $\mu_{\min}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\min}(\Delta\theta, \Delta\phi)$. For instance, to realize the minimum INR over a random neighborhood of size $(\Delta\theta, \Delta\phi)$, engineers can simply draw

$$\left[\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} \sim \mathcal{N}\big(\mu_{\min}(\Delta\theta, \Delta\phi), \sigma^2_{\min}(\Delta\theta, \Delta\phi)\big). \tag{5.23}$$

Plenty of statistics and statistical functions are readily available for the normal distribution which can be used to conduct system performance analyses. For instance, when $\left[\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} \sim \mathcal{N}(\mu_{\min}(\Delta\theta, \Delta\phi), \sigma^2_{\min}(\Delta\theta, \Delta\phi))$, the probability that the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding a random beam pair exhibits a minimum INR of at most ζ is

$$\mathbb{P}\left[\mathsf{INR}^{\min} \leq \zeta; \Delta\theta, \Delta\phi\right] = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{[\zeta]_{\mathrm{dB}} - \mu_{\min}(\Delta\theta, \Delta\phi)}{\sigma_{\min}(\Delta\theta, \Delta\phi) \cdot \sqrt{2}}\right)\right]. \tag{5.24}$$

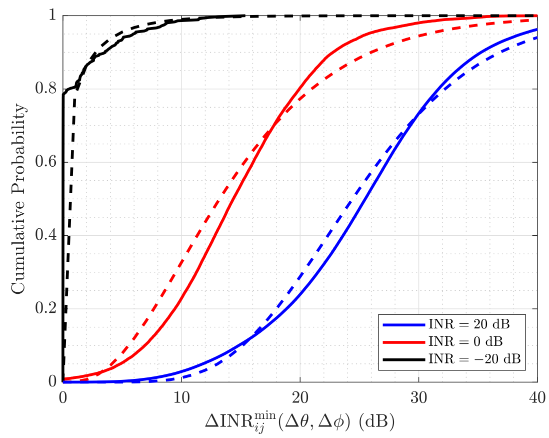

Figure 5.18: The CDF of $\left\{\Delta\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\}$ for various INR, where $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$. Their fitted counterparts are shown as dashed lines.

While fitting the CDFs in Figure 5.17a directly to normal distributions is useful, it does not capture how the distribution of $\mathsf{INR}_{ij}^{\mathrm{min}}$ varies as a function of INR_{ij}. In other words, it does not provide insight on if high-INR beam pairs are as likely to be within arm's reach of low INR as low-INR beam pairs are, for instance. To provide more detailed statistics on $\mathsf{INR}_{ij}^{\mathrm{min}}$ based on INR_{ij}, we begin by computing

$$\Delta\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) = \left[\mathsf{INR}_{ij}\right]_{\mathrm{dB}} - \left[\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} \geq 0 \tag{5.25}$$

which is the difference (in dB) between the inherent INR offered by the (i, j)-th beam pair and the minimum INR of its surrounding $(\Delta\theta, \Delta\phi)$-neighborhood. We subsequently form the set

$$\mathcal{D}_{\mathrm{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}) = \left\{\Delta\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\} \tag{5.26}$$

which is the set of all $\Delta\mathsf{INR}_{ij}^{\mathrm{min}}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij})$ values for beam pairs offering an

INR_{ij} of approximately INR. The approximation here is merely used to ensure $\mathcal{D}_{\min}$ has a sufficient number of points in it to successfully fit it, as will become clear. We found that the distribution of $\mathcal{D}_{\min}$ could be well approximated by a Gamma distribution as

$$\mathcal{D}_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}) \overset{\text{fit}}{\sim} \text{Gamma}(\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})) \quad (5.27)$$

where $\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ are the fitted shape and rate parameters, parameterized by INR in addition to $(\Delta\theta, \Delta\phi)$. Alongside the true CDFs of various $\mathcal{D}_{\min}$, we plotted their fitted distributions using dashed lines in Figure 5.18 for various INR and a $(2^\circ, 2^\circ)$-neighborhood. This illustrates that extremely low-INR beam pairs tend to see less reduction in INR over their neighborhoods compared to that of high-INR beam pairs. This is somewhat expected but also indicates that low-INR beam pairs are not congregated together but rather spread out throughout our transmit-receive space.

In Table 5.3, we tabulated $(\alpha_{\min}, \beta_{\min})$ for $\mathsf{INR} \in \{-20, -10, \dots, 40\}$ dB to provide engineers with $\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ for particular INR values. To approximate $\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ for any $\mathsf{INR} \in [-20, 40]$ dB, weighted interpolation can be used, for instance. It is our hope that this means to realize $\Delta\mathsf{INR}_{ij}^{\min}$ for particular INR_{ij} is useful for statistical analyses, simulation, and system evaluation. For instance, one may draw INR from the global distribution (i.e., the CDF in Figure 5.8) as $[\mathsf{INR}]_{\text{dB}} \sim \mathcal{N}(\mu, \sigma^2)$ and then use it when referencing Table 5.3 to fetch $\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ based on some neighborhood size

$(\Delta\theta, \Delta\phi)$. From there, a realization of $\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ can be drawn as

$$\left[\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})\right]_{\mathrm{dB}} \sim \\ [\mathsf{INR}]_{\mathrm{dB}} - \underbrace{\mathrm{Gamma}(\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}))}_{\Delta\mathsf{INR}^{\min}(\Delta\theta,\Delta\phi,\mathsf{INR})} \quad (5.28)$$

which is the minimum INR over the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding a beam pair offering a nominal INR of INR. Statistical analysis can be conducted using the Gamma distribution CDF, for instance, as follows. When $\Delta\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ follows a Gamma distribution with some parameters $\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$, the probability that the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding a random beam pair having an INR of INR exhibits a minimum INR of ζ or less is

$$\mathbb{P}\left[\mathsf{INR}^{\min} \leq \zeta; \Delta\theta, \Delta\phi, \mathsf{INR}\right] = \frac{\gamma(\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}) \cdot [\zeta]_{\mathrm{dB}})}{\Gamma(\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}))} \quad (5.29)$$

where $\Gamma(\cdot)$ is the Gamma function and $\gamma(\cdot, \cdot)$ is the lower incomplete Gamma function.

5.4.3 Maximum INR Over Various Neighborhoods

As was done for minimum INR over various neighborhoods, we have conducted an analysis and modeling for maximum INR over various neighborhoods. In Figure 5.19a, we plot the CDF of $\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)$ of all nearly 6.5 million beam pairs for variably sized neighborhoods. Shown in black is the $(0^\circ, 0^\circ)$-neighborhood, which is simply the CDF of INR_{ij}. Deviating by at

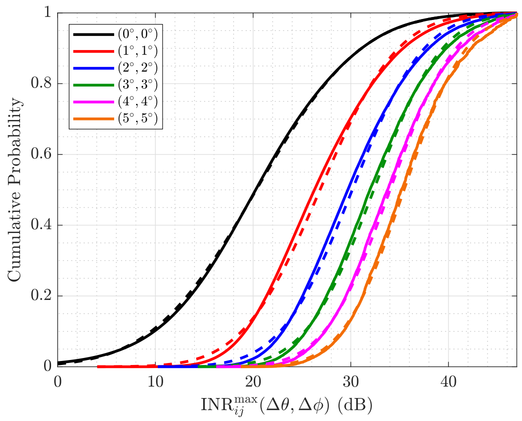

(a) $\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)$.

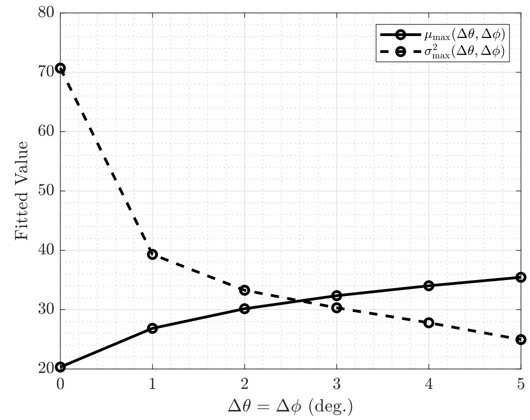

(b) $\mu_{\max}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\max}(\Delta\theta, \Delta\phi)$.

Figure 5.19: (a) The CDF of $\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)$ for various $(\Delta\theta, \Delta\phi)$. (b) The fitted parameters $\mu_{\max}(\Delta\theta, \Delta\phi)$ and $\sigma^2_{\max}(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$.

most $(1^\circ, 1^\circ)$, the median $\mathsf{INR}_{ij}^{\max}$ sees about a 6 dB increase from about 20 dB to 26 dB. Notice that the lower tail stops around 4 dB, meaning all nearly 6.5 million beam pairs are within $(1^\circ, 1^\circ)$ of a beam pair offering an INR of 4 dB or more. Deviating by at most $(2^\circ, 2^\circ)$, the median $\mathsf{INR}_{ij}^{\max}$ increases to 30 dB and the lower tail stops above 10 dB. This trend continues with diminishing gains as the neighborhood is widened.

Remark 5.8: Slight shifts can increase self-interference. Before, our analysis of $\mathsf{INR}_{ij}^{\min}$ highlighted that INR levels more attractive for full-duplex operation can be reached by small shifts in transmit and receive steering direction. Figure 5.19a highlights that small shifts in steering direction can likewise degrade (increase) INR. In fact, even beam pairs with very low INR inherently are highly sensitive, considering a $(1^\circ, 1^\circ)$ shift (at most) can lead to INR levels well above 0 dB, where full-duplex systems would be overwhelmingly self-interference-limited. As such, small shifts in the transmit and/or receive beams have the potential to drive self-interference to levels unfit for full-duplex. Therefore, if attempting to steer along high-isolation beam pairs to mitigate self-interference, there needs to be fairly high accuracy in doing so, potentially motivating high-resolution phase shifters, for instance.

Like before, we fit a distribution to the CDFs in Figure 5.19a to offer engineers a statistical tool for $\mathsf{INR}_{ij}^{\mathrm{max}}$. We fit a normal distribution as follows

$$\left\{ \left[\mathsf{INR}_{ij}^{\mathrm{max}}(\Delta\theta, \Delta\phi) \right]_{\mathrm{dB}} \ \forall \ i,j \right\} \overset{\mathrm{fit}}{\sim} \mathcal{N}\left(\mu_{\mathrm{max}}(\Delta\theta, \Delta\phi), \sigma_{\mathrm{max}}^2(\Delta\theta, \Delta\phi) \right) \tag{5.30}$$

where $\mu_{\mathrm{max}}(\Delta\theta, \Delta\phi)$ and $\sigma_{\mathrm{max}}^2(\Delta\theta, \Delta\phi)$ are the fitted mean and variance of the normal distribution. The dashed lines in Figure 5.19a depict each neighborhood's fitted distribution. Figure 5.19b shows the resulting fitted $\mu_{\mathrm{max}}(\Delta\theta, \Delta\phi)$ and $\sigma_{\mathrm{max}}^2(\Delta\theta, \Delta\phi)$ for various $\Delta\theta = \Delta\phi$. Naturally, the mean $\mu_{\mathrm{max}}(\Delta\theta, \Delta\phi)$ increases with neighborhood size but does so with diminishing gains. The variance $\sigma_{\mathrm{max}}^2(\Delta\theta, \Delta\phi)$ sees a sharp decrease with $\Delta\theta = \Delta\phi = 1^\circ$ from 0°. This highlights that so many of the nearly 6.5 million beam pairs are within a mere $(1^\circ, 1^\circ)$ of notably higher INR. The variance continues to trend down as the neighborhood widens since high-INR beam pairs can be more reliably reached. In addition to the select neighborhoods in Figure 5.19a, we have tabulated the fitted parameters $(\mu_{\mathrm{max}}(\Delta\theta, \Delta\phi), \sigma_{\mathrm{max}}^2(\Delta\theta, \Delta\phi))$ for a variety of $(\Delta\theta, \Delta\phi)$ in Table 5.4, as was done for $\mathsf{INR}_{ij}^{\mathrm{rng}}$ and $\mathsf{INR}_{ij}^{\mathrm{min}}$. Engineers can use these distributions to conduct a variety of statistical analyses related to $\mathsf{INR}_{ij}^{\mathrm{max}}$ (e.g., worst-case analyses analogous to (5.24)).

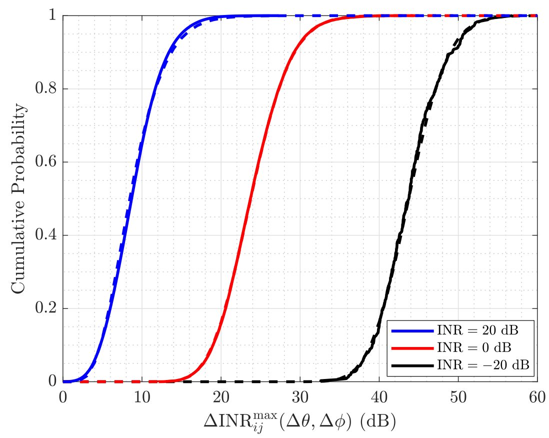

Figure 5.20: The CDF of $\left\{\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\}$ for various INR, where $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$. Their fitted counterparts are shown as dashed lines. Beam pairs with inherently low INR are typically at a greater risk of large increases in INR caused by small shifts in the transmit and receive beams.

As was done with minimum INR, we conduct a statistical fit of $\mathsf{INR}_{ij}^{\max}$ that depends on INR_{ij}. We define $\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)$ as

$$\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) = \left[\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi)\right]_{\mathrm{dB}} - \left[\mathsf{INR}_{ij}\right]_{\mathrm{dB}} \geq 0 \tag{5.31}$$

which is the difference (in dB) between the maximum INR within the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding the (i, j)-th transmit-receive beam pair and the INR offered by that beam pair. We form the set

$$\mathcal{D}_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}) = \left\{\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij}) : \mathsf{INR}_{ij} \approx \mathsf{INR}\right\} \tag{5.32}$$

by collecting all $\Delta\mathsf{INR}_{ij}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}_{ij})$ for beam pairs offering an INR_{ij} of approximately INR. We again use a Gamma distribution to approximate the distribution of $\mathcal{D}_{\max}$ as

$$\mathcal{D}_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}) \overset{\text{fit}}{\sim} \mathrm{Gamma}(\alpha_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})) \tag{5.33}$$

where $\alpha_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})$ are the fitted shape and rate parameters, parameterized by INR in addition to $(\Delta\theta, \Delta\phi)$. We plotted fitted distributions using dashed lines in Figure 5.20 for various INR and a $(2^\circ, 2^\circ)$-neighborhood.

In addition, we tabulated $(\alpha_{\max}, \beta_{\max})$ for $\mathsf{INR} \in \{-20, -10, \ldots, 40\}$ dB to provide engineers with statistical tools. Again, weighted averaging may be used to interpolate between INR values listed in Table 5.5. Like for minimum INR, $\Delta\mathsf{INR}_{ij}^{\max}$ for particular INR_{ij} is can be realized using these fitted distributions as follows. One may draw INR from the global distribution (i.e., the CDF in Figure 5.8) as $[\mathsf{INR}]_{\mathrm{dB}} \sim \mathcal{N}(\mu, \sigma^2)$ and then use it when referencing Table 5.5 to fetch $\alpha_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\beta_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})$ based on some neighborhood size $(\Delta\theta, \Delta\phi)$. A realization of the maximum INR over the $(\Delta\theta, \Delta\phi)$-neighborhood surrounding a beam pair offering a nominal INR of INR can be drawn as

$$[\mathsf{INR}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})]_{\mathrm{dB}} \sim [\mathsf{INR}]_{\mathrm{dB}} + \underbrace{\mathrm{Gamma}(\alpha_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}))}_{\Delta\mathsf{INR}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})} \quad (5.34)$$

which can facilitate statistical analyses; (5.29) can be straightforwardly translated from $\mathsf{INR}^{\min}$ to $\mathsf{INR}^{\max}$, for instance. Note that INR range as a function of INR can be realized using $\mathsf{INR}^{\min}(\Delta\theta, \Delta\phi, \mathsf{INR})$ and $\mathsf{INR}^{\max}(\Delta\theta, \Delta\phi, \mathsf{INR})$.

5.5 Conclusion

We have collected nearly 6.5 million measurements of multi-panel self-interference at 28 GHz to better understand its spatial and statistical characteristics, providing the most comprehensive examination of such to date. Our measurements illustrate that the degree of self-interference coupled between colocated transmitting and receiving phased arrays tends to be higher when the transmit and receive beams are steered toward one another but small shifts in steering direction (on the order of one degree) can lead to significant changes in such. We have analyzed and statistically modeled this sensitivity, providing engineers with useful insights and statistical tools that can drive system design and evaluation, including those that may use analog and/or digital self-interference cancellation. This measurement campaign sheds light on the efficacy of multi-panel mmWave full-duplex systems, such full-duplex IAB proposed in 3GPP, and motivates strategic beam steering as a potential route to mitigate self-interference without prohibitively compromising beamforming gain. Valuable future work would investigate the impacts of beam shape, array size, environmental reflections, and relative array geometry on self-interference. Future directions capitalizing on this campaign include beam selection for mmWave full-duplex, proposing a practically sound MIMO channel model for mmWave self-interference, and prototyping full-duplex mmWave systems.

Table 5.1: Fitted $(\alpha_{\mathrm{rng}}(\Delta\theta, \Delta\phi), \beta_{\mathrm{rng}}(\Delta\theta, \Delta\phi))$.

$\Delta\phi$ \ $\Delta\theta$	0°	1°	2°	3°	4°	5°
0°	—	(2.74, 3.40)	(4.42, 3.64)	(6.73, 3.15)	(9.50, 2.62)	(12.50, 2.22)
1°	(2.59, 3.19)	(4.52, 4.10)	(6.90, 3.82)	(10.90, 2.94)	(16.04, 2.25)	(21.69, 1.80)
2°	(4.04, 3.48)	(6.57, 3.85)	(10.69, 3.11)	(17.57, 2.21)	(25.96, 1.64)	(34.67, 1.30)
3°	(5.80, 3.22)	(9.63, 3.16)	(16.31, 2.35)	(26.67, 1.63)	(38.06, 1.23)	(48.77, 1.01)
4°	(7.98, 2.81)	(13.91, 2.49)	(23.80, 1.77)	(37.32, 1.26)	(50.52, 0.99)	(61.80, 0.85)
5°	(10.39, 2.44)	(18.85, 2.00)	(31.67, 1.42)	(47.24, 1.05)	(61.17, 0.86)	(72.50, 0.76)

Table 5.2: Fitted $(\mu_{\min}(\Delta\theta, \Delta\phi), \sigma^2_{\min}(\Delta\theta, \Delta\phi))$.

$\Delta\phi$ \ $\Delta\theta$	$0°$	$1°$	$2°$	$3°$	$4°$	$5°$
$0°$	$(20.32, 70.69)$	$(15.04, 102.75)$	$(10.34, 114.62)$	$(6.47, 112.94)$	$(3.53, 107.37)$	$(1.31, 101.93)$
$1°$	$(15.58, 98.86)$	$(8.32, 148.79)$	$(2.30, 152.39)$	$(-2.34, 135.13)$	$(-5.62, 118.00)$	$(-7.98, 105.59)$
$2°$	$(11.58, 109.04)$	$(3.15, 153.04)$	$(-3.07, 141.96)$	$(-7.56, 117.66)$	$(-10.61, 99.81)$	$(-12.80, 88.34)$
$3°$	$(8.23, 106.83)$	$(-0.88, 137.50)$	$(-6.93, 119.30)$	$(-11.07, 96.78)$	$(-13.87, 82.49)$	$(-15.89, 73.80)$
$4°$	$(5.50, 99.76)$	$(-3.98, 118.15)$	$(-9.74, 99.32)$	$(-13.57, 81.48)$	$(-16.16, 70.95)$	$(-18.05, 64.68)$
$5°$	$(3.39, 93.25)$	$(-6.27, 103.55)$	$(-11.76, 86.70)$	$(-15.37, 72.63)$	$(-17.82, 64.54)$	$(-19.63, 59.62)$

Table 5.3: Fitted $(\alpha_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\min}(\Delta\theta, \Delta\phi, \mathsf{INR}))$.

	INR (dB)						
$(\Delta\theta, \Delta\phi)$	-20	-10	0	10	20	30	40
$(1^\circ, 1^\circ)$	$(0.19, 1.34)$	$(0.21, 11.37)$	$(0.89, 10.27)$	$(3.66, 4.05)$	$(3.26, 3.91)$	$(3.22, 2.80)$	$(4.09, 1.47)$
$(2^\circ, 2^\circ)$	$(0.18, 4.91)$	$(0.41, 13.93)$	$(3.06, 4.83)$	$(10.28, 2.21)$	$(8.67, 2.93)$	$(5.26, 4.18)$	$(4.15, 4.03)$
$(3^\circ, 3^\circ)$	$(0.21, 9.31)$	$(0.87, 10.20)$	$(6.49, 2.86)$	$(17.38, 1.58)$	$(19.94, 1.66)$	$(12.18, 2.74)$	$(7.96, 3.79)$
$(4^\circ, 4^\circ)$	$(0.26, 12.40)$	$(1.86, 6.27)$	$(10.40, 2.07)$	$(24.37, 1.26)$	$(31.12, 1.21)$	$(22.78, 1.77)$	$(20.63, 1.94)$
$(5^\circ, 5^\circ)$	$(0.39, 11.99)$	$(2.98, 4.66)$	$(15.20, 1.57)$	$(31.79, 1.05)$	$(40.47, 1.01)$	$(30.88, 1.45)$	$(36.15, 1.27)$

Table 5.4: Fitted $(\mu_{\max}(\Delta\theta, \Delta\phi), \sigma^2_{\max}(\Delta\theta, \Delta\phi))$.

$\Delta\phi$ \ $\Delta\theta$	$0°$	$1°$	$2°$	$3°$	$4°$	$5°$
$0°$	$(20.32, 70.69)$	$(24.36, 44.95)$	$(26.42, 39.02)$	$(27.63, 36.81)$	$(28.41, 35.64)$	$(29.06, 34.66)$
$1°$	$(23.84, 49.63)$	$(26.85, 39.31)$	$(28.64, 35.37)$	$(29.71, 33.95)$	$(30.41, 33.26)$	$(31.00, 32.60)$
$2°$	$(25.63, 45.34)$	$(28.42, 37.00)$	$(30.15, 33.27)$	$(31.18, 31.99)$	$(31.85, 31.51)$	$(32.41, 31.00)$
$3°$	$(26.89, 43.38)$	$(29.61, 35.48)$	$(31.33, 31.57)$	$(32.35, 30.31)$	$(33.01, 29.93)$	$(33.56, 29.52)$
$4°$	$(27.89, 41.95)$	$(30.61, 33.81)$	$(32.35, 29.51)$	$(33.38, 28.13)$	$(34.02, 27.79)$	$(34.57, 27.40)$
$5°$	$(28.72, 40.61)$	$(31.46, 32.02)$	$(33.23, 27.27)$	$(34.26, 25.73)$	$(34.90, 25.36)$	$(35.45, 24.95)$

Table 5.5: Fitted $(\alpha_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}), \beta_{\max}(\Delta\theta, \Delta\phi, \mathsf{INR}))$.

	INR (dB)						
$(\Delta\theta, \Delta\phi)$	-20	-10	0	10	20	30	40
$(1°, 1°)$	$(97.36, 0.39)$	$(53.69, 0.53)$	$(22.22, 0.83)$	$(8.27, 1.23)$	$(3.94, 1.44)$	$(3.65, 1.09)$	$(3.73, 0.72)$
$(2°, 2°)$	$(125.30, 0.35)$	$(74.38, 0.46)$	$(35.07, 0.68)$	$(14.26, 1.04)$	$(6.18, 1.44)$	$(4.83, 1.28)$	$(4.76, 0.84)$
$(3°, 3°)$	$(132.69, 0.36)$	$(83.02, 0.45)$	$(42.63, 0.64)$	$(18.35, 0.97)$	$(7.76, 1.44)$	$(5.40, 1.41)$	$(5.37, 0.88)$
$(4°, 4°)$	$(143.23, 0.35)$	$(92.33, 0.43)$	$(49.46, 0.60)$	$(22.61, 0.89)$	$(9.53, 1.36)$	$(5.96, 1.46)$	$(6.10, 0.84)$
$(5°, 5°)$	$(160.40, 0.32)$	$(107.86, 0.39)$	$(58.95, 0.53)$	$(28.38, 0.78)$	$(12.12, 1.19)$	$(6.53, 1.47)$	$(6.94, 0.79)$

Chapter 6

Spatial and Statistical Modeling of Millimeter Wave Self-Interference

In Chapter 5, we conducted an extensive measurement campaign of mmWave self-interference and analyzed nearly 6.5 million collected measurements. We statistically characterized self-interference in a variety of ways based on our measurements, but we made no attempt to *spatially* model self-interference. In this chapter, we spatially and statistically model self-interference, shedding light on the underlying spatial composition of the self-interference channel.

6.1 Introduction

Full-duplex mmWave communication systems have drawn increased attention recently due to their potential enhancements and applications in next-generation wireless networks [19, 20, 32, 68, 95]. The prospects of full-duplex in mmWave systems are particularly exciting thanks to ongoing efforts demonstrating the cancellation of self-interference purely through beamforming [41, 42, 68, 72, 74–77]. By strategically steering or shaping its transmit and receive beams, it has been shown that a full-duplex mmWave system can reduce self-interference to levels near or below the noise floor, facilitating simultaneous transmission and reception across the same frequency spectrum—without requiring analog or digital SIC. The success of such beamforming-based approaches, however, has almost exclusively been validated through simulation using highly-idealized self-interference channel models that have not been verified through measurement. In fact, until now, very little has been understood about the spatial composition of mmWave self-interference in real-world systems, raising questions regarding the efficacy of these proposed solutions and of full-duplex mmWave systems altogether. A measurement-backed model of mmWave self-interference would facilitate accurate design and evaluation of practical full-duplex mmWave systems.

6.1.1 Existing Models and Measurements

The most popular way to model self-interference in mmWave systems thus far has been based on near-field propagation via the spherical-wave MIMO channel model [2]. Albeit sensible, this channel model is extremely idealized and ignores a multitude of real-world factors including non-isotropic antenna array elements, array enclosures, mounting infrastructure, and a variety of other nonidealities. Furthermore, this channel model is deterministic for a given relative transmit and receive array geometry, which prevents engineers from capturing artifacts unique to individual full-duplex mmWave systems. This spherical-wave model has often been used in conjunction with a ray-based model [72, 112] to capture reflections off the environment in addition to the direct coupling between the transmit and receive arrays. Through simulation, these models of self-interference have been used widely to develop and validate beamforming-based solutions for full-duplex mmWave systems. However, these approaches to model self-interference have not been verified through measurement, meaning it is unclear if proposed solutions evaluated under such models would actually see success in practice.

Nevertheless, there have been efforts to measure and characterize self-interference at mmWave frequencies. In [106], for instance, measurements of 28 GHz self-interference power were reported for 8×8 UPAs in indoor and outdoor environments. These measurements are certainly valuable but offer limited insights on the spatial and statistical characteristics of self-interference. The works of [108–111] also collected measurements of mmWave self-interference using horn and lens antennas, which shed light on environmental factors and the power delay profile of self-interference. However, the measurements in [106, 108–111] did not provide extensive characterization or modeling of self-interference (not a small task), limiting their usefulness in the development and evaluation of full-duplex mmWave systems—especially those employing phased arrays and relying on beamforming to cancel self-interference.

Thus far, the most extensive measurement campaign and characterization of mmWave self-interference is our recent work [44]. Therein, we collected nearly 6.5 million measurements of 28 GHz self-interference using 16×16 UPAs, whose analysis revealed high-level spatial trends and noteworthy small-scale variability. We provided a thorough statistical characterization of measured self-interference to facilitate the design and evaluation of full-duplex mmWave systems. In [44], however, we made no attempt to model the *spatial* composition of self-interference but did show that the spherical-wave channel model [2] does not align with our measurements—motivating the need for a new model fit for real-world full-duplex mmWave systems.

6.1.2 Contributions

In this paper, we introduce the first measurement-backed model of self-interference in multi-panel[1] full-duplex mmWave systems. Using measurements of 28 GHz self-interference collected in our campaign [44], we construct a novel stochastic model that aligns both *statistically* and *spatially* with real-world self-interference. Our model is relatively simple, has theoretical foundations, and is built on system and model parameters with physical meaning. This gives our model the potential to generalize to systems beyond our own 28 GHz phased array platform via unique parameterization. We demonstrate this by fitting our model to nearly 13 million additional measurements collected from two other system configurations. To construct our model, we uncovered a coarse geometric approximation of the self-interference channel from within our measurements, which suggests that the dominant coupling between the transmit and receive arrays manifests as clusters of rays in a far-field manner, rather than in a near-field, spherical-wave fashion. This is a novel finding that can steer future work aiming to model self-interference MIMO channels

[1]We use the term "multi-panel" to refer to separate transmit and receive arrays.

in full-duplex mmWave systems.

By comparing statistical realizations of self-interference from our model to actual measurements, we show that our model exhibits the large-scale and small-scale spatial characteristics of self-interference observed in real full-duplex mmWave systems. In fact, to further confirm this, we evaluate an existing full-duplex solution [3] using our model and repeat this evaluation using actual measurements. The coinciding results of these two evaluations indicates that our model can indeed serve as a useful tool in the design and validation of full-duplex mmWave systems by giving engineers the ability to draw accurate realizations of self-interference and conduct statistical analyses. Simultaneously, we also illustrate that the commonly-used spherical-wave model [2] neither spatially nor statistically aligns with our measurements of self-interference.

The present work and our prior work [44] can both be used to draw statistical realizations of mmWave self-interference, but the two differ notably. With [44], statistical realizations must simply be drawn from the global distribution of our measurements and therefore lack any spatial consistency. With the model presented herein, realizations can be drawn that are both statistically and spatially consistent with our measurements, courtesy of the spatial model of self-interference that we have uncovered. This spatial modeling sheds light on the underlying nature of self-interference, with the consequent spatial consistency facilitating more accurate design and evaluation of full-duplex mmWave systems.

6.1.3 Notation

We use $\mathcal{N}(\mu, \sigma^2)$ to denote the normal distribution having mean μ and variance σ^2. We use θ and ϕ as directions in azimuth and elevation, respectively. We use i and j for transmit and receive indexing, respectively. We use hat $\hat{a}$ to denote the estimate of some ground truth a. We use bar $\bar{A}$ to denote a model parameter having connections to some theoretical system parameter A. All symbols have linear units by default; we use $[A]_{\mathrm{dB}} = 10 \cdot \log_{10}(A)$ to denote A in units of decibels (dB) and $[B]_{\mathrm{dBm}}$ to denote B in units of decibel-milliwatts (dBm).

6.2 Measurements of mmWave Self-Interference

The work herein is based on data collected in our recent measurement campaign of 28 GHz self-interference [44], which was highlighted in Chapter 5. In this section, we will summarize our measurement platform and methodology, highlighting key and necessary information. We collected measurements of self-interference incurred by a multi-panel full-duplex mmWave system employing two phased arrays mounted on separate sides of a sectorized triangular platform, as shown previously in Figure 5.3 in Chapter 5. One phased array acted as a transmitter while the other served as a receiver, with their centers separated by 30 cm. The transmitter and receiver are identical Anokiwave AWA-0134 28 GHz phased arrays [114], each of which contains a 16×16 half-wavelength UPA. This sectorized configuration with industry-grade phased arrays aligns with realistic multi-panel full-duplex mmWave deployments, such as for IAB as proposed in 3GPP [95]. Measurements were conducted in an anechoic chamber, free from significant reflectors, to inspect the direct coupling between the arrays; investigating self-interference in real-world environments is necessary future work.

Each phased array can be electronically steered toward some azimuth and elevation via digitally-controlled analog beamforming weights. With 256 antenna elements, our phased arrays can produce highly directional beams, enabling us to inspect self-interference with fine spatial granularity. To steer our transmit beam or receive beam toward some azimuth-elevation (θ, ϕ), we use conjugate beamforming [88], where the transmit and receive beamforming vectors are

$$\mathbf{f}(\theta, \phi) = \mathbf{a}_{\mathrm{tx}}(\theta, \phi), \qquad \mathbf{w}(\theta, \phi) = \mathbf{a}_{\mathrm{rx}}(\theta, \phi). \tag{6.1}$$

Here, $\mathbf{a}_{\mathrm{tx}}(\theta, \phi)$ and $\mathbf{a}_{\mathrm{rx}}(\theta, \phi)$ are the array response vectors of our transmit array and receive array in some direction (θ, ϕ). These array response vectors are solely a characteristic of the array geometries and carrier frequency and can be computed in a closed-form [88, 97]. We employ the same 3-D geometry shown in Figure 5.4 and used throughout Chapter 5. Our array response vectors and beamforming vectors are normalized as

$$\|\mathbf{f}(\theta, \phi)\|_2^2 = \|\mathbf{a}_{\mathrm{tx}}(\theta, \phi)\|_2^2 = N_{\mathrm{t}}, \qquad \|\mathbf{w}(\theta, \phi)\|_2^2 = \|\mathbf{a}_{\mathrm{rx}}(\theta, \phi)\|_2^2 = N_{\mathrm{r}}, \tag{6.2}$$

where $N_{\mathrm{t}} = N_{\mathrm{r}} = 256$ is the number of transmit and receive elements in our particular arrays.

Assuming a linear setting, when transmitting toward some direction $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and receiving toward $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$, the coupled self-interference power at the receiver output can be expressed as

$$P_{\mathrm{SI}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = \mathsf{EIRP} \cdot G^2 \cdot \underbrace{\left|\mathbf{w}(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})^* \mathbf{H} \mathbf{f}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})\right|^2}_{\text{coupling factor}}, \tag{6.3}$$

where $\mathbf{H} \in \mathbb{C}^{N_{\mathrm{r}} \times N_{\mathrm{t}}}$ is the self-interference MIMO channel manifesting between the transmit and receive arrays and $\mathsf{EIRP} = 60$ dBm is the EIRP of our transmitting phased array. Here, the scalar G^2 captures the inherent path loss between the arrays, along with amplification and losses internal to the transmit and receive array modules. Rather than raw power, it is perhaps more useful to use an INR to characterize self-interference. As such, we normalize received self-interference power to the noise floor as

$$\mathsf{INR}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = \frac{P_{\mathrm{SI}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})}{P_{\mathrm{noise}}}, \tag{6.4}$$

where $P_{\mathrm{noise}} = -68$ dBm is the integrated noise power at our receive array output over our 100 MHz measurement bandwidth. We used Zadoff-Chu sequences and correlation-based processing to reliably measure received self-interference power well below the noise floor. Using high-fidelity test equipment and stepped attenuators, we calibrated our measurement capability and verified that it had low error (typically less than 1 dB) broadly across received power levels.

With measures of INR, full-duplex system performance can be quantified by computing the received SINR as

$$\mathsf{SINR}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = \frac{\mathsf{SNR}}{1 + \mathsf{INR}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})}, \tag{6.5}$$

where SNR is the SNR of a desired receive signal at the output of the receive array, which would practically depend on $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$. To solely understand the degree of self-interference at the receive array output, this work is concerned with measuring and modeling INR when no further cancellation is employed (i.e., no analog or digital SIC). INR alone is useful in indicating if a full-duplex system is noise-limited ($\mathsf{INR} \ll 0$ dB) or self-interference-limited ($\mathsf{INR} \gg 0$ dB). Such systems desire a low INR, say $\mathsf{INR} < 0$ dB in most cases, to ensure self-interference does not erode full-duplexing gains.

To collect INR measurements over a broad and dense spatial profile, we swept the transmit beam and receive beam across a number of directions in azimuth and elevation. Prior to measurement, we specified sets of N_{tx} transmit directions and N_{rx} receive directions

$$\mathcal{A}_{\mathrm{tx}} = \left\{ \left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)} \right) \right\}_{i=1}^{N_{\mathrm{tx}}}, \qquad \mathcal{A}_{\mathrm{rx}} = \left\{ \left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)} \right) \right\}_{j=1}^{N_{\mathrm{rx}}} \tag{6.6}$$

over which the phased arrays were electronically swept, with self-interference measured for each transmit-receive combination. In this work, we densely swept the transmit and receive beams from -60° to 60° in azimuth and from -10° to 10° in elevation, both in 1° steps.

$$\mathcal{A}_{\mathrm{tx}} = \mathcal{A}_{\mathrm{rx}} = \underbrace{\{-60^\circ, -59^\circ, \ldots, 60^\circ\}}_{\mathsf{azimuth}} \times \underbrace{\{-10^\circ, -9^\circ, \ldots, 10^\circ\}}_{\mathsf{elevation}} \tag{6.7}$$

$$= \{(-60^\circ, -10^\circ), \ldots, (-60^\circ, 10^\circ), \ldots, (60^\circ, 10^\circ)\} \tag{6.8}$$

This amounts to $N_{\text{tx}} = N_{\text{rx}} = 121 \times 21 = 2541$ steering directions for a total of $N_{\text{tx}} \times N_{\text{rx}} \approx 6.5$ million INR measurements. The set of nearly 6.5 million INR measurements we denote as

$$\mathcal{I} = \{\mathsf{INR}(\theta_{\text{tx}}, \phi_{\text{tx}}, \theta_{\text{rx}}, \phi_{\text{rx}}) : (\theta_{\text{tx}}, \phi_{\text{tx}}) \in \mathcal{A}_{\text{tx}}, (\theta_{\text{rx}}, \phi_{\text{rx}}) \in \mathcal{A}_{\text{rx}}\}. \tag{6.9}$$

We spatially and statistically model self-interference from this set of measurements in the sections that follow. For a summary of these collected measurements, please see [44] or refer to Chapter 5.

6.3 Motivation and Roadmap

The goal of this work is to produce a stochastic model of mmWave self-interference that is statistically and spatially well-aligned with our collected measurements. Such a model would allow engineers to draw countless realizations of self-interference and facilitate accurate simulation, development, and evaluation of full-duplex mmWave systems. Even if one were to assume the same setup as ours, it is unrealistic to assume our measurements could be directly extended to other systems due to beam-steering and positioning differences, among numerous other practical artifacts. As such, being able to draw *realizations* of self-interference—as opposed to estimating deterministic values—will be more generalizable and useful to the research community.

As we will elaborate on shortly, our statistical model of self-interference is based on two characteristics observed in our measurements:

1. On a large scale (at a high level), there is a connection between the

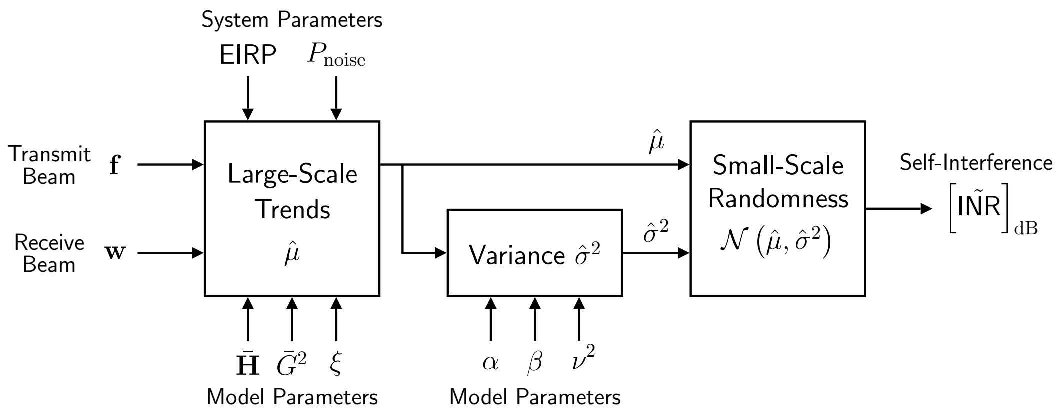

Figure 6.1: A block diagram of our statistical model of self-interference. Based on system parameters, model parameters, and transmit and receive beams, our model produces realizations of self-interference that are statistically and spatially aligned with actual measurements.

steering directions of the transmit and receive beams and the degree of self-interference incurred. Broadly speaking, some transmit and receive directions tend to incur high self-interference while others tend to incur low self-interference.

2. On a small scale (within small spatial neighborhoods), the system incurs seemingly random amounts of self-interference. Slightly shifting the transmit and receive steering directions can dramatically alter the degree of self-interference coupled.

We leverage these large-scale and small-scale characteristics to construct a stochastic model of self-interference that both statistically and spatially aligns with our measurements.

A block diagram summarizing our model is shown in Figure 6.1, which will become clear as we proceed in its presentation. For particular transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$, a mean parameter $\hat{\mu}$ is estimated, which dictates the location of the distribution from which self-interference is drawn. The variance of this distribution $\hat{\sigma}^2$ is dictated by $\hat{\mu}$ and other model parameters. This approach allows our model to capture the large-scale spatial trends in self-interference along with the small-scale variability observed over small spatial neighborhoods. With appropriate parameterization, our model has the potential to be extended to other systems and environments beyond our own, which we demonstrate in Section 6.6 by fitting our model to measurements collected with two other system configurations.

6.4 Self-Interference Over Small Spatial Neighborhoods

In pursuit of a model of self-interference that is stochastic, we make use of a small-scale phenomenon present in our measurements: seemingly random amounts of self-interference are observed over small spatial neighborhoods. To quantify this small-scale randomness, we introduce the concept of a *spatial neighborhood*, illustrated previously in Figure 5.13. Recall that in (5.11) we defined $\measuredangle(\theta_1, \theta_2)$ as the absolute difference between two angles θ_1, θ_2. Let $\mathcal{T}_{\mathrm{tx}}^{(i)}(\Delta\theta, \Delta\phi)$ be the spatial neighborhood of size $(\Delta\theta, \Delta\phi)$ in azimuth-elevation surrounding the i-th measured transmit direction $\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right) \in \mathcal{A}_{\mathrm{tx}}$, defined as follows, and let $\mathcal{T}_{\mathrm{rx}}^{(j)}(\Delta\theta, \Delta\phi)$ analogously be the spatial neighborhood surrounding the j-th measured receive direction $\left(\theta_{\text{rx}}^{(j)}, \phi_{\text{rx}}^{(j)}\right) \in \mathcal{A}_{\text{rx}}$.

$$\mathcal{T}_{\text{tx}}^{(i)}(\Delta\theta, \Delta\phi) = \left\{(\theta, \phi) \in \mathcal{A}_{\text{tx}} : \measuredangle\left(\theta, \theta_{\text{tx}}^{(i)}\right) \leq \Delta\theta, \measuredangle\left(\phi, \phi_{\text{tx}}^{(i)}\right) \leq \Delta\phi\right\} \tag{6.10}$$

$$\mathcal{T}_{\text{rx}}^{(j)}(\Delta\theta, \Delta\phi) = \left\{(\theta, \phi) \in \mathcal{A}_{\text{rx}} : \measuredangle\left(\theta, \theta_{\text{rx}}^{(j)}\right) \leq \Delta\theta, \measuredangle\left(\phi, \phi_{\text{rx}}^{(j)}\right) \leq \Delta\phi\right\} \tag{6.11}$$

In other words, $\mathcal{T}_{\text{tx}}^{(i)}(\Delta\theta, \Delta\phi)$ is the set of measured transmit directions within $\Delta\theta$ in azimuth and within $\Delta\phi$ in elevation of the i-th transmit direction $\left(\theta_{\text{tx}}^{(i)}, \phi_{\text{tx}}^{(i)}\right)$. Using these, let $\mathcal{I}_{ij}(\Delta\theta, \Delta\phi)$ be the set of measured INR values over transmit and receive spatial neighborhoods surrounding the (i, j)-th beam pair.

$$\begin{aligned}&\mathcal{I}_{ij}(\Delta\theta, \Delta\phi) = \\ &\left\{\mathsf{INR}(\theta_{\text{tx}}, \phi_{\text{tx}}, \theta_{\text{rx}}, \phi_{\text{rx}}) : (\theta_{\text{tx}}, \phi_{\text{tx}}) \in \mathcal{T}_{\text{tx}}^{(i)}(\Delta\theta, \Delta\phi), (\theta_{\text{rx}}, \phi_{\text{rx}}) \in \mathcal{T}_{\text{rx}}^{(j)}(\Delta\theta, \Delta\phi)\right\}\end{aligned} \tag{6.12}$$

We observed that INR values over neighborhoods of size $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ were approximately log-normally distributed throughout our measurements. Put simply, we found that

$$\left[\mathsf{INR} \in \mathcal{I}_{ij}(2^\circ, 2^\circ)\right]_{\text{dB}} \overset{\text{approx.}}{\sim} \mathcal{N}\left(\mu_{ij}, \sigma_{ij}^2\right) \ \forall \ i, j, \tag{6.13}$$

where μ_{ij} and σ_{ij}^2 are the mean and variance of the normal distribution fitted to $\left[\mathcal{I}_{ij}(2^\circ, 2^\circ)\right]_{\text{dB}}$. Note that μ_{ij} and σ_{ij}^2 are specific to the neighborhood surrounding a particular transmit-receive beam pair (i, j), meaning there are nearly 6.5 million fitted $\left(\mu_{ij}, \sigma_{ij}^2\right)$. Also, note that μ_{ij} implicitly has units of decibels. Three instances of $\mathcal{I}_{ij}(2^\circ, 2^\circ)$ depicting this log-normal nature can be observed in Figure 6.2a, each of which has a unique fitted mean μ_{ij} and variance σ_{ij}^2.

Each of these three CDFs in Figure 6.2a illustrates that slightly shifting the transmit and receive beams leads to wide variability in self-interference.

6.4.1 Confirming Log-Normal Nature Across All Measurements

To quantitatively evaluate this log-normal nature over our entire set of measurements, we computed the Kolmogorov-Smirnov (K-S) statistic of the fit for each (i, j) as follows. Henceforth, under the assumption of a $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ spatial neighborhood, we employ the shorthand

$$\mathcal{I}_{ij} \triangleq \left[\mathcal{I}_{ij}(2^\circ, 2^\circ)\right]_{\mathrm{dB}}. \tag{6.14}$$

Let $\mathbb{F}_{\mathcal{I}_{ij}}(X)$ be the empirical CDF of $\mathcal{I}_{ij}$—that is, the fraction of elements in $\mathcal{I}_{ij}$ less than or equal to X, defined as

$$\mathbb{F}_{\mathcal{I}_{ij}}(X) = \frac{1}{|\mathcal{I}_{ij}|} \cdot \sum_{\mathsf{INR} \in \mathcal{I}_{ij}} 1\{\mathsf{INR} \leq X\}, \tag{6.15}$$

where $1\{\cdot\}$ is the indicator function. Similarly, let $\mathbb{F}_{\mathcal{N}_{ij}}(X)$ be the CDF of the normal distribution $\mathcal{N}_{ij} \triangleq \mathcal{N}\left(\mu_{ij}, \sigma_{ij}^2\right)$ evaluated at X. For each (i, j), the K-S statistic $\kappa_{ij} \in [0, 1]$ can be written as

$$\kappa_{ij} = \max_{X} \ \left|\mathbb{F}_{\mathcal{I}_{ij}}(X) - \mathbb{F}_{\mathcal{N}_{ij}}(X)\right|, \tag{6.16}$$

which describes the maximum absolute difference in cumulative density between two distributions and therefore serves as a measure of similarity between the distributions. A $\kappa_{ij} \leq 0.1$ indicates that self-interference over the (i, j)-th neighborhood closely follows a log-normal distribution, whereas higher κ_{ij} indicates poorer fitting.

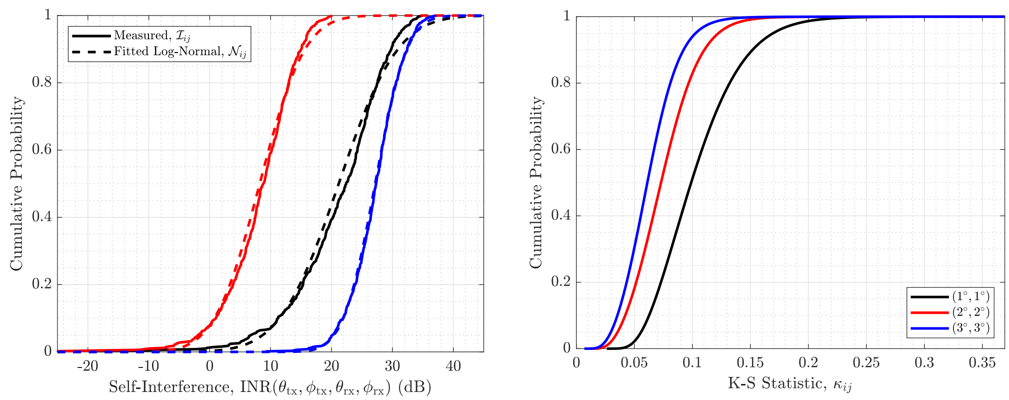

(a) Examplar log-normal neighborhoods. (b) K-S statistic κ_{ij} for various $(\Delta\theta, \Delta\phi)$.

Figure 6.2: (a) Three spatial neighborhoods whose measured self-interference follows a log-normal distribution, each with a unique fitted mean and variance. (b) The CDFs of the K-S statistic κ_{ij} across all nearly 6.5 million beam pairs (i, j) for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. The vast majority of measured beam pairs exhibit levels of self-interference that are log-normally distributed across the surrounding $(2°, 2°)$ spatial neighborhood.

In Figure 6.2b, we plot the CDF of κ_{ij} over the nearly 6.5 million neighborhoods (i.e., for each (i, j)). We have plotted this for a neighborhood size of $(2°, 2°)$, as well as $(1°, 1°)$ and $(3°, 3°)$; for now, consider $(2°, 2°)$. We see that nearly 50% are within $\kappa_{ij} \leq 0.07$, over 80% are within $\kappa_{ij} \leq 0.1$, and nearly 95% are within $\kappa_{ij} \leq 0.12$. This suggests that the vast large majority of neighborhoods are well approximated by their fitted log-normal distributions. There exists a noteworthy upper tail with less than 1% of neighborhoods yielding $\kappa_{ij} \geq 0.15$, which indicates that select neighborhoods do not follow a log-normal distribution as closely as the overwhelming majority. Nonetheless, the log-normal nature of local neighborhoods is the foundation for the statistical model of self-interference we present henceforth, and we acknowledge that it may limit how well our model can capture the outliers of our measurements.

Compared to $(1°, 1°)$, a neighborhood size of $(2°, 2°)$ more often closely follows a log-normal distribution, a notable improvement of about 0.03 in κ_{ij} in distribution. A neighborhood size of $(3°, 3°)$ indeed offers improvements over $(2°, 2°)$ in terms of log-normal fit, but these are marginal, making it more desirable to use a smaller neighborhood to allow our model to capture finer large-scale spatial variability. This motivates our use of a neighborhood size of $(2°, 2°)$ henceforth with the understanding that others may be suitable as well. In addition, distributions other than log-normal could also used to capture this small-scale variability; we chose a log-normal because it fit well broadly across our measurements and because its statistical parameters can be captured succinctly with our proposed model, as we will see.

6.4.2 What Causes This Log-Normal Nature?

It is difficult to precisely describe the roots of this log-normal nature, but we can make educated speculations based on an understanding of our measurement system and wireless fundamentals. First of all, self-interference is presumably coupled behind the arrays by their side lobes and back lobes, and when slightly shifting the steering directions of the main lobes, these side and back lobes fluctuate and the location of nulls move—this can lead to destructive combining when nulls manifest in directions of strong interference. Moreover, in theory, arrays exhibit idealized beam patterns with well-defined nulls, which would make self-interference more straightforwardly understood.

In practice, however, beam shapes can significantly deviate from their familiar idealized patterns, exhibiting irregular lobes and shallower nulls. This irregularity is even more apparent behind the arrays, in part due to phase shifts and attenuation introduced by their enclosures, adding further complexity.

Other factors likely at play are near-field effects between the transmitter and receiver. The far-field distance from either of our arrays is $2D^2/\lambda \approx 2.5$ meters, and the reactive/radiating near-field boundary is $0.62\sqrt{D^3/\lambda} \approx 23$ cm, where each array's largest dimension is $D \approx 11$ cm [97]. The separation of our arrays is 30 cm, meaning they reside just within the radiating near-field of one another based on these rules-of-thumb. Electromagnetic propagation in the near-field is less predictable than that in the far-field and bucks typical assumptions in array theory, most notably the plane wave assumption; recall, the beam patterns we are familiar with are typically from a far-field perspective. In near-field settings, side lobes can grow and main lobes can widen, for instance [20,97], further complicating the interaction between transmit and receive beams. In addition, due to these near-field effects, there is the potential that self-interference is in fact not linearly related to the transmit and receive beams as expressed in (6.3), but this requires further investigation to flesh out. As we will see, the dominant source of self-interference actually resembles far-field interaction, raising questions about the near-field nature of the underlying self-interference channel. To summarize, we conjecture that the combination of irregular beam patterns and near-field effects atop far-field interaction leads to seemingly random fluctuations in self-interference with slight shifts of the transmit and receive beams. Detailed electromagnetic simulation would likely yield more concrete explanations for this log-normal behavior but would demand a significant investment of time and resources, making it a good topic for dedicated future work.

6.5 Spatial and Statistical Modeling

In this section, we present our model of self-interference for multi-panel full-duplex mmWave systems. Our method is relatively simple, has theoretical foundations, and is built on system and model parameters having real-world meaning. Naturally, we cannot guarantee that our model—nor the measurements it is based on—will perfectly translate to other systems. We suspect that it can indeed be generalized to systems beyond our own, however, through appropriate parameterization, which we show in Section 6.6. At a high level, our model of mmWave self-interference accomplishes two goals:

1. It is stochastic, allowing researchers to draw countless realizations of self-interference, rather than a single deterministic quantity.
2. Its realizations align spatially and in distribution with real-world self-interference.

As mentioned in the previous section, the foundation of our model is the seemingly random nature of self-interference across small spatial neighborhoods. Having fit a log-normal distribution to the neighborhood surrounding each beam pair (i, j), we have nearly 6.5 million fitted mean μ_{ij} and variance σ_{ij}^2. In this section, we reverse-engineer a methodology to estimate these fitted parameters $\left(\mu_{ij}, \sigma_{ij}^2\right)$ based on transmit beam $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right)$ and receive beam $\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)$. Then, using this estimation methodology, a realization of self-interference for general transmit beam $\mathbf{f}$ and receive beam $\mathbf{w}$ can be drawn as

$$\left[\widetilde{\mathsf{INR}}\right]_{\mathrm{dB}} \sim \mathcal{N}\left(\hat{\mu}, \hat{\sigma}^2\right) \tag{6.17}$$

where $\hat{\mu}$ and $\hat{\sigma}^2$ are an estimated mean and variance based on $\mathbf{f}$ and $\mathbf{w}$.

6.5.1 Estimating Neighborhood Mean, μ_{ij}

With nearly 6.5 million true fitted μ_{ij}, we seek a method to reliably estimate each based on the known transmit and receive beams used during measurement. If successful, the same approach can then be used to estimate the mean for general transmit and receive beams. To estimate the mean μ_{ij} for a given transmit-receive pair (i, j), we found that a coarse approximation of the over-the-air self-interference channel between the transmit and receive arrays can be used. As mentioned in the introduction, the spherical-wave MIMO channel model [2]—an idealized near-field model—has been widely used in research on full-duplex mmWave systems. However, as pointed out in our prior work [44] and illustrate in Figure 6.3, this spherical-wave channel model does not align with our measurements of self-interference. In light of this, we now propose a different channel model to coarsely describe the spatial structure of self-interference.

Coarse geometric modeling of the self-interference MIMO channel. It is practically difficult to precisely measure the over-the-air self-interference

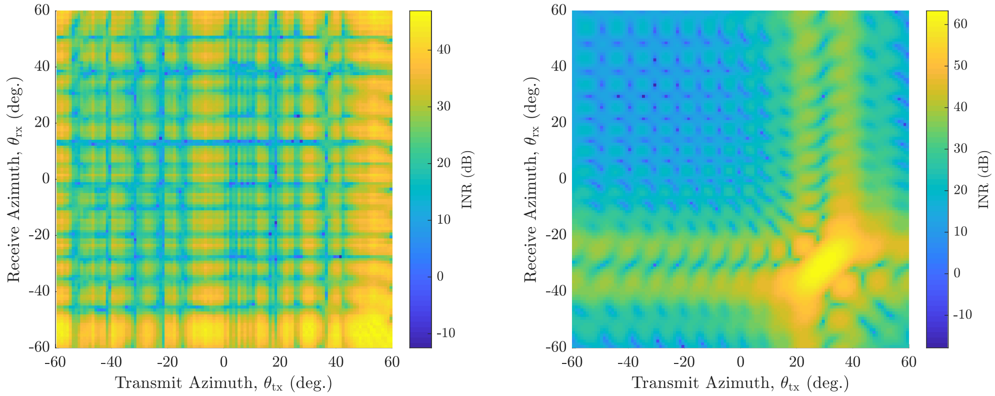

(a) Actual measurements.

(b) Spherical-wave model [2].

Figure 6.3: (a) The measured azimuth cut of self-interference, where $\phi_{\rm tx} = \phi_{\rm rx} = 0°$. (b) The simulated counterpart of (a) based on the commonly-used spherical-wave near-field channel model [2]. The stark difference between the two motivates the need for a new model of real-world self-interference, which we present herein.

MIMO channel (i.e., the channel matrix $\mathbf{H}$ shown in (6.3)) due to its high dimensionality and the fact that it is observed through the lens of the analog beamforming networks in the transmit and receive phased arrays. Fortunately, we have uncovered from our measurements a coarse geometric model of the self-interference channel. This model does not capture the significant variability observed locally but rather the large-scale spatial components of the self-interference channel. We now outline this proposed geometric model, which we will use to estimate μ_{ij}.

Let us begin by inspecting the azimuth cut of our INR measurements (i.e., $\phi_{\rm tx} = \phi_{\rm rx} = 0°$) plotted in Figure 6.3a. From the bright clouds in Figure 6.3a, we can readily see that there is high self-interference when $(\theta_{\rm tx}, \theta_{\rm rx})$ is approximately $(0°, -60°)$, $(60°, -60°)$, $(60°, 60°)$, and $(-60°, -60°)$. This hints

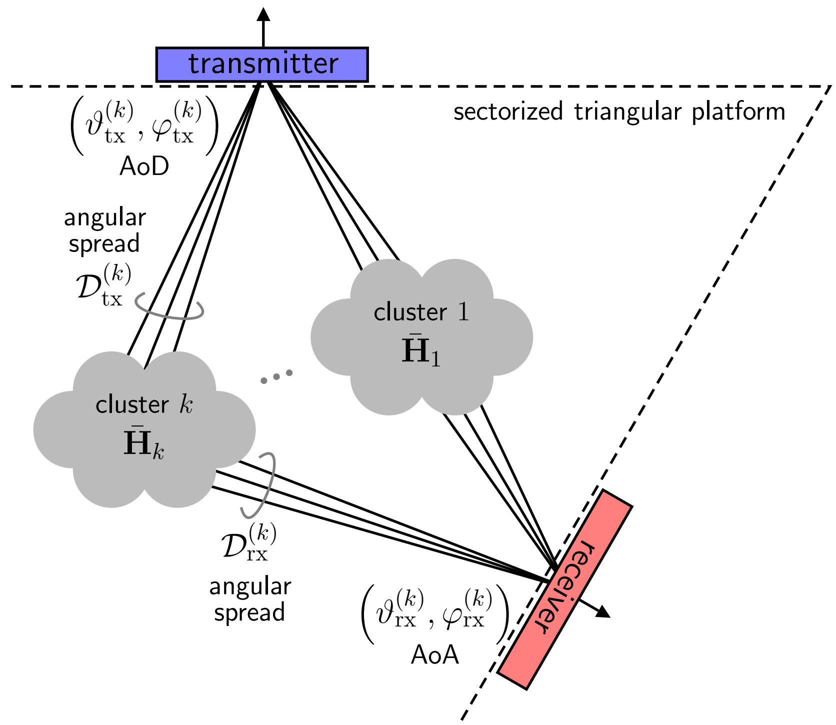

Figure 6.4: An illustration of coupling clusters comprising the coarse geometric model of the self-interference channel $\bar{\mathbf{H}}$ between the transmit and receive arrays. From our measurements, we observed the presence of four dominant coupling clusters. This clustered coupling is a drastic deviation from the spherical-wave model [2] that has been widely used in the literature.

at which angular components couple high self-interference and paves the way for our proposed geometric model. Rather than model this coupling as single rays extending from the transmit array to the receive array, we instead model them as clouds of rays—or *coupling clusters*—with some non-zero angular spread, as illustrated in Figure 6.4. Suppose N_{clust} coupling clusters exist between the transmit and receive arrays, where each cluster contains the same number of rays for simplicity. The center of the k-th cluster is defined by some AoD from the transmit array $\left(\vartheta_{\mathrm{tx}}^{(k)}, \varphi_{\mathrm{tx}}^{(k)}\right)$ and some AoA at the receive array $\left(\vartheta_{\mathrm{rx}}^{(k)}, \varphi_{\mathrm{rx}}^{(k)}\right)$. These cluster centers correspond to the bright clouds identified by visual inspection of Figure 6.3a.

We assume the angular spread of each cluster is quantified in azimuth and elevation by some $\Delta\vartheta$ and $\Delta\varphi$ and is common across clusters for simplicity. We use $\mathcal{D}_{\vartheta}$ and $\mathcal{D}_{\varphi}$ defined as

$$\mathcal{D}_{\vartheta} = \{-\Delta\vartheta, \ldots, 0, \ldots, +\Delta\vartheta\}, \qquad \mathcal{D}_{\varphi} = \{-\Delta\varphi, \ldots, 0, \ldots, +\Delta\varphi\} \tag{6.18}$$

to discretize this angular spread into individual rays with some resolution in azimuth and elevation, respectively. For instance, discretizing both $\mathcal{D}_{\vartheta}$ and $\mathcal{D}_{\varphi}$ with 1° resolution has been sufficient for us. We model a coupling cluster as the collection of all AoD-AoA combinations based on the angular spread $(\Delta\vartheta, \Delta\varphi)$. In this fashion, the AoDs and AoAs of the rays comprising the k-th

Table 6.1: Dominant cluster centers in our measurements ($N_{\text{clust}} = 4$).

Cluster Index k	**AoD** $\left(\vartheta_{\text{tx}}^{(k)}, \varphi_{\text{tx}}^{(k)}\right)$	**AoA** $\left(\vartheta_{\text{rx}}^{(k)}, \varphi_{\text{rx}}^{(k)}\right)$
1	$(-174^\circ, 0^\circ)$	$(-122^\circ, 0^\circ)$
2	$(126^\circ, 0^\circ)$	$(-122^\circ, 0^\circ)$
3	$(-118^\circ, 0^\circ)$	$(-122^\circ, 0^\circ)$
4	$(126^\circ, 0^\circ)$	$(118^\circ, 0^\circ)$

coupling cluster can be expressed as

$$\mathcal{D}_{\text{tx}}^{(k)} = \left(\vartheta_{\text{tx}}^{(k)}, \varphi_{\text{tx}}^{(k)}\right) + \mathcal{D}_{\vartheta} \times \mathcal{D}_{\varphi}, \tag{6.19}$$

$$\mathcal{D}_{\text{rx}}^{(k)} = \underbrace{\left(\vartheta_{\text{rx}}^{(k)}, \varphi_{\text{rx}}^{(k)}\right)}_{\text{cluster center}} + \underbrace{\mathcal{D}_{\vartheta} \times \mathcal{D}_{\varphi}}_{\text{angular spread}}. \tag{6.20}$$

The channel matrix produced by the k-th cluster is then written as

$$\bar{\mathbf{H}}_k = \sum_{(\vartheta_{\text{tx}}, \varphi_{\text{tx}}) \in \mathcal{D}_{\text{tx}}^{(k)}} \ \sum_{(\vartheta_{\text{rx}}, \varphi_{\text{rx}}) \in \mathcal{D}_{\text{rx}}^{(k)}} \mathbf{a}_{\text{rx}}(\vartheta_{\text{rx}}, \varphi_{\text{rx}}) \ \mathbf{a}_{\text{tx}}(\vartheta_{\text{tx}}, \varphi_{\text{tx}})^{*}, \tag{6.21}$$

which is simply the sum of rank-1 matrices produced by each of the rays in the cluster; recall, $\mathbf{a}_{\text{tx}}(\vartheta_{\text{tx}}, \varphi_{\text{tx}})$ and $\mathbf{a}_{\text{rx}}(\vartheta_{\text{rx}}, \varphi_{\text{rx}})$ are the array responses of the transmit array and receive array for some AoD and AoA, respectively. As implicitly done here, assuming all rays to have unit gain has proven sufficient for us but this could be generalized straightforwardly. The resulting channel matrix comprised of all N_{clust} coupling clusters is then

$$\bar{\mathbf{H}} = \sum_{k=1}^{N_{\text{clust}}} \bar{\mathbf{H}}_k \cdot \frac{\sqrt{N_{\text{t}} \cdot N_{\text{r}}}}{\left\| \sum_{k=1}^{N_{\text{clust}}} \bar{\mathbf{H}}_k \right\|_{\text{F}}}, \tag{6.22}$$

where the scaling is to ensure the channel matrix has fixed energy $\left\|\bar{\mathbf{H}}\right\|_{\text{F}}^2 = N_{\text{t}} \cdot N_{\text{r}}$.

Inspection of our measurements revealed the presence of $N_{\text{clust}} = 4$ dominant clusters whose centers are listed in Table 6.1 and were mentioned before based on Figure 6.3a. Notice that the azimuthal components of the AoDs and AoAs are beyond $\pm 90^\circ$ rather than within the $[-60^\circ, 60^\circ]$ that was measured, as was illustrated in Figure 6.4. This decision is an arbitrary one since the array response of a UPA is symmetric about $\pm 90^\circ$; in other words, an azimuth of 60° induces the same array response as an azimuth of 120°. We made this choice to follow the intuition that self-interference is presumably coupled *behind* the arrays since measurements took place in an anechoic chamber free of any significant reflectors in front of the arrays. We have observed that the dominant coupling exists in and around the azimuth plane, meaning $\varphi_{\text{tx}}^{(k)} = \varphi_{\text{rx}}^{(k)} = 0^\circ$ for all $k = 1, \ldots, N_{\text{clust}}$. This can likely be attributed to the fact that our transmit and receive arrays are aligned in elevation, though in other configurations, components beyond an elevation of 0° may play a more significant role. From our measurements, we empirically found a good angular spread to be $\Delta\vartheta = 4^\circ$ and $\Delta\varphi = 3^\circ$. Note that other cluster parameters could be used for systems that see other self-interference profiles, potentially also to account for reflections off the environment.

Using our geometric channel to estimate μ_{ij}. With our geometric channel model $\bar{\mathbf{H}}$ in hand, we now aim to estimate μ_{ij} for particular transmit and receive beams. To do so, we developed the estimator

$$\hat{\mu}_{ij} = \left[\frac{\mathsf{EIRP} \cdot \bar{G}^2 \cdot \left| \mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)^* \bar{\mathbf{H}} \mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right) \right|^{2\cdot\xi}}{P_{\mathrm{noise}}} \right]_{\mathrm{dB}}, \tag{6.23}$$

which is rooted in using $\bar{\mathbf{H}}$ and the beamforming weights $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right)$ and $\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)$ to capture coupling trends as a function of transmit and receive beamforming. The expression of (6.23) is quite similar to and is inspired by that of (6.3)–(6.4). The scalars ξ and $\bar{G}^2$ serve as model parameters that we will fit based on our measurements using the true $\{\mu_{ij}\}$. Along with normalizing the energy of the channel matrix as $\left\|\bar{\mathbf{H}}\right\|_{\mathrm{F}}^2 = N_{\mathrm{t}} \cdot N_{\mathrm{r}}$, we also remind the reader that we normalize our beams as $\|\mathbf{f}\|_2^2 = N_{\mathrm{t}}$ and $\|\mathbf{w}\|_2^2 = N_{\mathrm{r}}$. These normalizations are crucial for proper scaling of $\hat{\mu}_{ij}$. Recall that (6.23) is used to estimate the mean μ_{ij} of the (i, j)-th neighborhood and not the actual INR of the (i, j)-th beam pair. For this reason, $\bar{G}^2$ and $\bar{\mathbf{H}}$ should not be thought of as equal to G^2 and $\mathbf{H}$ in (6.3), though they may closely align.

Fitting model parameters ξ and $\bar{G}^2$ to measurements. The model parameters ξ and $\bar{G}^2$ introduced in (6.23) allow us to account for distributional discrepancies in $\{\mu_{ij}\}$ and the INR one would theoretically expect based on our coarsely-approximated channel $\bar{\mathbf{H}}$. The role of these model parameters can be better observed by expressing (6.23) in log-form as

$$\hat{\mu}_{ij} = \underbrace{\xi}_{\text{scale}} \cdot \underbrace{[\Gamma_{ij}]_{\mathrm{dB}}}_{\text{shape}} + \underbrace{\left[\bar{G}^2\right]_{\mathrm{dB}}}_{\text{location}} + \underbrace{[\mathsf{EIRP}]_{\mathrm{dBm}} - [P_{\mathrm{noise}}]_{\mathrm{dBm}}}_{\text{system parameters}}, \tag{6.24}$$

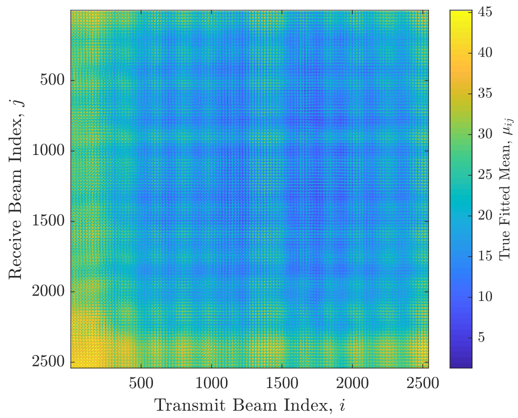

(a) True fitted mean μ_{ij}.

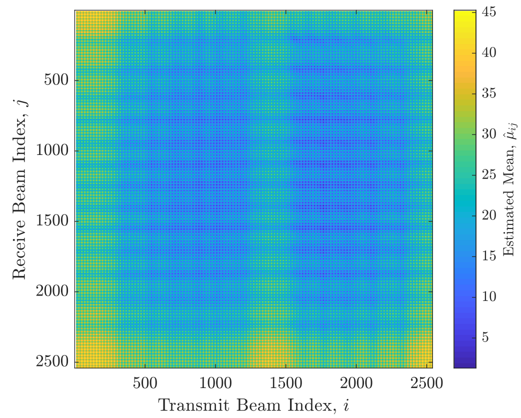

(b) Estimated mean $\hat{\mu}_{ij}$.

Figure 6.5: (a) The true fitted mean μ_{ij} for each of the nearly 6.5 million transmit-receive beam pairs (i, j). (b) The estimated counterpart of (a), where $\hat{\mu}_{ij}$ for each beam pair has been estimated using (6.24) with fitted ξ and $\bar{G}^2$. The strong structural similarity between the two confirms that $\bar{\mathbf{H}}$ captures large-scale spatial trends observed in our measurements.

where

$$\Gamma_{ij} = \left| \mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)^* \bar{\mathbf{H}} \mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right) \right|^2. \tag{6.25}$$

From (6.24), we see that the shape of the distribution of $\{\hat{\mu}_{ij}\}$ is dictated by the coupling factors $\{\Gamma_{ij}\}$, while its scale and location are controlled by ξ and $\bar{G}^2$, respectively.

In order to ensure the distribution of our estimates $\{\hat{\mu}_{ij}\}$ matches that of the true $\{\mu_{ij}\}$ in terms of scale and location, we fit ξ and $\bar{G}^2$ as follows. We first fit ξ as

$$\xi = \sqrt{\frac{\mathrm{var}(\{\mu_{ij}\})}{\mathrm{var}\left(\left\{[\Gamma_{ij}]_{\mathrm{dB}}\right\}\right)}} \tag{6.26}$$

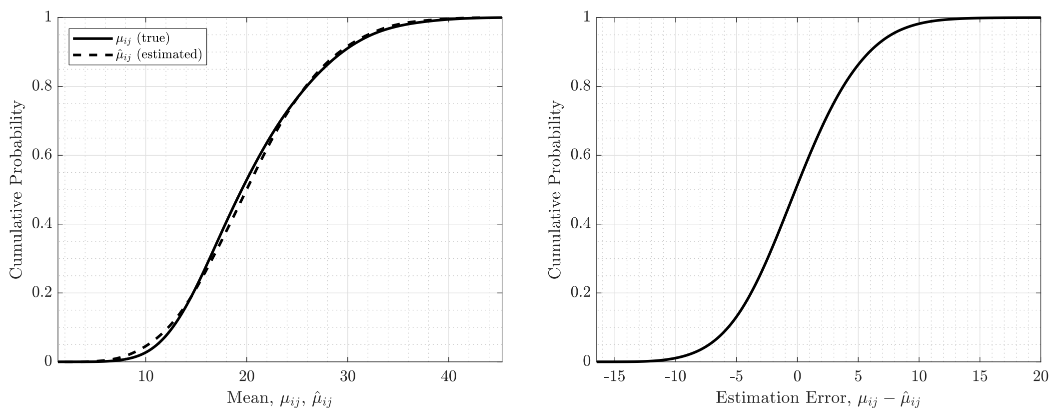

(a) True fitted mean μ_{ij} and estimated $\hat{\mu}_{ij}$.

(b) Estimation error $\mu_{ij} - \hat{\mu}_{ij}$.

Figure 6.6: (a) The CDF of the true fitted mean μ_{ij} and the estimated mean $\hat{\mu}_{ij}$ for all nearly 6.5 million beam pairs (i, j); in other words, the CDFs of Figure 6.5a and Figure 6.5b. (b) The CDF of the error between μ_{ij} and $\hat{\mu}_{ij}$ for all nearly 6.5 million beam pairs (i, j); in other words, the CDF of the difference between Figure 6.5a and Figure 6.5b.

to align the variance of our estimates $\{\hat{\mu}_{ij}\}$ with the variance of the true $\{\mu_{ij}\}$. Then, $\bar{G}^2$ is fitted to align their locations as

$$\left[\bar{G}^2\right]_{\mathrm{dB}} = \mathrm{mean}(\{\mu_{ij}\}) - \mathrm{mean}\big(\{\xi \cdot [\Gamma_{ij}]_{\mathrm{dB}}\}\big) - [\mathsf{EIRP}]_{\mathrm{dBm}} + [P_{\mathrm{noise}}]_{\mathrm{dBm}}, \tag{6.27}$$

which ensures that the mean of our estimates $\{\hat{\mu}_{ij}\}$ matches that of the true $\{\mu_{ij}\}$. Following this fitting from our measurements and $\bar{\mathbf{H}}$, we have $\xi = 0.502$ and $\bar{G}^2 = -129.00$ dB.

We evaluate our estimation of μ_{ij} in Figure 6.5 and Figure 6.6. First, consider Figure 6.5a and Figure 6.5b, which depict all nearly 6.5 million μ_{ij} and the corresponding estimates $\{\hat{\mu}_{ij}\}$, respectively. Clearly, our estimates $\{\hat{\mu}_{ij}\}$ align visually with the true $\{\mu_{ij}\}$. Our ability to closely reproduce the structure of $\{\mu_{ij}\}$ via $\bar{\mathbf{H}}$ suggests that the dominant source of self-interference stems from coupling clusters (implicitly a far-field model) as opposed to the aforementioned highly-idealized near-field model [2]. The magnitude of these estimates $\{\hat{\mu}_{ij}\}$ follows that of the true $\{\mu_{ij}\}$, courtesy of appropriately fitted ξ and $\bar{G}^2$. This can be further observed in Figure 6.6a, which compares the CDFs of $\{\mu_{ij}\}$ and $\{\hat{\mu}_{ij}\}$ and illustrates that our estimates align closely in distribution with the ground truth. Here, it can be thought that $\bar{\mathbf{H}}$ is responsible for aligning the *shape* of the two distributions, while $\bar{G}^2$ and ξ align the *location* and *scale*. Finally, Figure 6.6b shows the empirical CDF of the error associated with our estimation of μ_{ij}. As desired, this error is centered about 0, and the vast majority of its density lay within $|\mu_{ij} - \hat{\mu}_{ij}| \leq 10$. Approximately 75% of our estimates are within 5 dB and 90% are within 7.5 dB. Altogether, this evaluation confirms that our approach accurately captures the spatial and distributional characteristics of μ_{ij}.

It is natural to ask whether ξ and $\bar{G}^2$ hold any physical meaning. It is difficult to state precisely, but we can interpret ξ and $\bar{G}^2$ as accomplishing a few things. First, most naturally, $\bar{G}^2$ can largely be interpreted as capturing the inverse path loss and isolation between the transmit and receive arrays (e.g., due to enclosures, free space path loss), along with amplification therein. The scale parameter ξ plays an important role in correcting for deviations from the theoretical coupling between transmit and receive beams. It can account for small errors in our coarse geometric self-interference channel model $\bar{\mathbf{H}}$, along with irregular beam patterns produced by practical arrays. Practical beam patterns, for instance, typically do not exhibit the clean shapes and deep nulls we are accustomed to from array theory. ξ also may account for the fact that self-interference is potentially not linearly related to $\mathbf{f}$ and $\mathbf{w}$. In addition, ξ also importantly corrects for the fact that we are using Γ_{ij} to estimate the *mean* of the (i, j)-th neighborhood, rather than directly the INR of the (i, j)-th beam pair. The parameters $\bar{G}^2$ and ξ—along with $\bar{\mathbf{H}}$—allow our statistical model to be extended to other systems by tailoring them appropriately and presumably contain additional underlying physical meaning beyond what has already been mentioned, even if it is not yet fully fleshed out. Note that the estimator in (6.23) and (6.24) is presented for measured beam pairs but can be naturally extended to more general $\mathbf{f}$ and $\mathbf{w}$ by dropping (i, j) indexing, using fitted ξ and $\bar{G}^2$.

6.5.2 Estimating Neighborhood Variance, σ_{ij}^2.

Having estimated the mean μ_{ij}, we now focus on estimating the neighborhood variance σ_{ij}^2. Estimating μ_{ij} involved model and system parameters and fairly intuitive methods with sound connections to the physical world. The estimation of the variance σ_{ij}^2, on the other hand, proved to be much more difficult, with less pronounced dependence on the transmit and receive beams. Consider Figure 6.7a, which shows scattered pairs of the true fitted $\left(\mu_{ij}, \sigma_{ij}^2\right)$. We can observe that the variance σ_{ij}^2 varies significantly with the mean but that it tends to be smaller with higher μ_{ij}. Conceptually, this suggests that beam pairs seeing high self-interference tend to observe less variability across

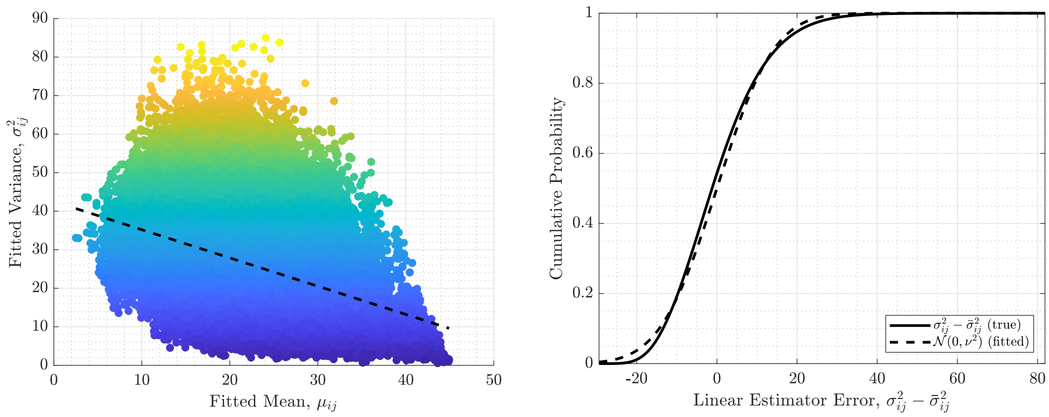

(a) Linear estimation of σ^2_{ij} based on μ_{ij}. (b) Linear estimator error $\sigma^2_{ij} - \bar{\sigma}^2_{ij}$.

Figure 6.7: (a) Pairs of true fitted $(\mu_{ij}, \sigma^2_{ij})$ and the linear estimator (dashed line) used to produce $\bar{\sigma}^2_{ij} = \mu_{ij} \cdot \alpha + \beta$. (b) The CDF of the linear estimator error $\sigma^2_{ij} - \bar{\sigma}^2_{ij}$, along with the corresponding fitted normal distribution $\mathcal{N}(0, \nu^2)$.

their neighborhood—in line with conclusions drawn in [44]. Fitting a simple linear estimator to the true $\left(\mu_{ij}, \sigma^2_{ij}\right)$, shown as the dashed line, allows us to make the linear estimate $\bar{\sigma}^2_{ij}$ of the variance σ^2_{ij} using the mean μ_{ij} via

$$\bar{\sigma}^2_{ij} = \mu_{ij} \cdot \alpha + \beta \tag{6.28}$$

where $\alpha = -0.733$ and $\beta = 42.53$ are the slope and bias of the linear estimator fitted to our measurements. While this simple estimator captures the general trend in variance as a function of the mean, it does not capture its broad and seemingly random variability observed in Figure 6.7a. To estimate $\hat{\sigma}^2_{ij}$, we therefore include randomness from a normal distribution in conjunction with the linear estimator to increase its variability. Following the linear estimator, $\hat{\sigma}^2_{ij}$ is drawn via

$$\hat{\sigma}^2_{ij} \sim \left[\bar{\sigma}^2_{ij} + \mathcal{N}\left(0, \nu^2\right)\right]^+ \tag{6.29}$$

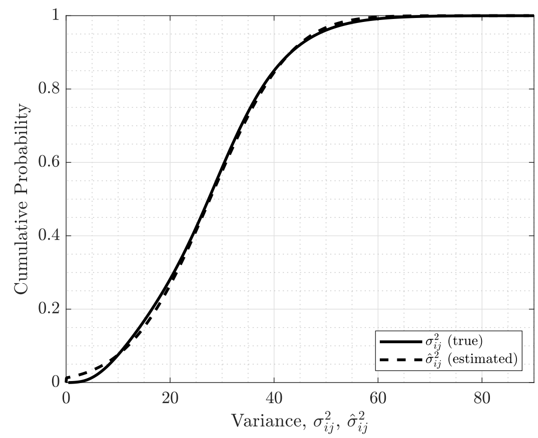

Figure 6.8: CDFs of the true fitted variance σ_{ij}^2 and estimated variance $\hat{\sigma}_{ij}^2$.

where the variance of the normal distribution is fitted as $\nu^2 = 126.091$ based on the distribution of the error $\left\{\sigma_{ij}^2 - \bar{\sigma}_{ij}^2\right\}$ illustrated in Figure 6.7b; here, $[x]^+ = \max(x, 0)$ simply ensures $\hat{\sigma}_{ij}^2 \geq 0$.

By supplementing the linear estimate $\bar{\sigma}_{ij}^2$ with normally-distributed noise, the distribution of $\left\{\hat{\sigma}_{ij}^2\right\}$ closely aligns with that of the true fitted $\left\{\sigma_{ij}^2\right\}$, as illustrated in Figure 6.8. Note that this approach aligns their *distributions* rather than aiming to perfectly estimate $\hat{\sigma}_{ij}^2 = \sigma_{ij}^2$. This was done deliberately due to the simple fact that there was no reliable way to true estimate σ_{ij}^2; rather, it appeared to be extremely random. However, by aligning the distributions of $\left\{\hat{\sigma}_{ij}^2\right\}$ and $\left\{\sigma_{ij}^2\right\}$, our model is able to produce realizations of self-interference with spatial richness comparable to the true measurements. This importantly allows our model to capture the variability in self-interference over small spatial neighborhoods that was observed in our measurements.

Algorithm 6.1 A statistical realization of self-interference.

Input: $\mathbf{f}$, $\mathbf{w}$; $\bar{\mathbf{H}}$, $\bar{G}^2$, ξ, α, β, ν^2; EIRP, P_{noise}

Estimate mean using (6.23)**:**

$\Gamma = \left| \mathbf{w}^* \bar{\mathbf{H}} \mathbf{f} \right|^2$

$\hat{\mu} = \xi \cdot [\Gamma]_{\text{dB}} + \left[\bar{G}^2\right]_{\text{dB}} + [\mathsf{EIRP}]_{\text{dBm}} - [P_{\text{noise}}]_{\text{dBm}}$

Draw variance using (6.28) **and** (6.29)**:**

$\bar{\sigma}^2 = \hat{\mu} \cdot \alpha + \beta$

$\hat{\sigma}^2 \sim \left[\bar{\sigma}^2 + \mathcal{N}(0, \nu^2)\right]^+$

Draw self-interference:

$[\tilde{\mathsf{INR}}]_{\text{dB}} \sim \mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$

Bound INR to prevent extreme outliers (optional):

$\tilde{\mathsf{INR}} \leftarrow \min\left(\max\left(\tilde{\mathsf{INR}}, \mathsf{INR}^{\min}\right), \mathsf{INR}^{\max}\right)$

Calculate self-interference power (if needed):

$[P_{\text{SI}}]_{\text{dBm}} = [P_{\text{noise}}]_{\text{dBm}} + [\tilde{\mathsf{INR}}]_{\text{dB}}$

6.5.3 Summary

A summary of our statistical model of self-interference is shown in Algorithm 6.1. For particular transmit and receive beamforming weights $\mathbf{f}$ and $\mathbf{w}$, the mean $\hat{\mu}$ is computed using model parameters $\bar{\mathbf{H}}$, $\bar{G}^2$, and ξ and system parameters EIRP and P_{noise}. Then, the variance $\hat{\sigma}^2$ can be realized using $\hat{\mu}$ and model parameters α and β, along with normally-distributed noise having variance ν^2. Recall, these system and model parameters can be unique to a particular full-duplex mmWave platform; they may be fitted directly from measurements or referenced from future wireless standards or the literature. A realization of self-interference can be drawn from the normal distribution $\mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, which can be bounded if desired to prevent extreme outliers.

6.6 Generalization and Applications of Our Model

In the previous section, we detailed the construction of our model based on a set of nearly 6.5 million measurements. In this section, we highlight how this model can generalize to other systems by collecting nearly 13 million additional measurements from two different system configurations. Then, we highlight applications of our proposed model to overview how it can be used in full-duplex research and development.

6.6.1 Fitting Our Model to Other Configurations

Naturally, we cannot guarantee how well our statistical model will translate to other systems, carrier frequencies, and environments, but it does have the potential to via custom system and model parameters. Our model readily accommodates general system EIRP and noise power. Model parameters can be uniquely tailored to individual platforms by collecting measurements and subsequently fitting them using the methods described in Section 6.5 or by using parameters published in future literature or wireless standards. Other systems will presumably see unique spatial profiles of self-interference, which can be accounted via $\bar{\mathbf{H}}$ through a different number of clusters, different cluster AoDs/AoAs, and a different angular spread. Note that our model can directly accommodate arbitrary array geometries and sizes by constructing $\bar{\mathbf{H}}$ with the appropriate array responses. However, it is not clear how different array geometries may affect our model and its fitted parameters. For instance, perhaps the small-scale variability observed over small spatial neighborhoods decreases with fewer array elements. Investigating this could make for interesting future work.

To more concretely confirm that our model can indeed generalize beyond the single set of measurements it was constructed on, we repeated our nearly 6.5 million measurements of self-interference in two additional system configurations. We fit our model of self-interference to each set of measurements and tabulated their fitted parameters in Table 6.2:

- The *default* column corresponds to the aforementioned setup and nearly 6.5 million measurements used extensively herein and presented in [44].
- The *vertical* column corresponds to a vertically-polarized version of *default*, where we repeated the measurements with both arrays physically rotated 90°.
- The *tapered* column corresponds to repeating the measurements of *default*, except with transmit and receive beams having tapered side lobes.

When fitting our model to each of these, we found that the spatial profile of self-interference was common across all three configurations. In other words, the same coarse model of the self-interference channel $\bar{\mathbf{H}}$ was observed in all three cases. This exciting finding reinforces that the coupling between the arrays occurs in the form of clusters of rays and that this coupling is a function of the relative array geometry and of the mounting infrastructure.

We can observe that ξ is approximately 0.5 across all three configurations, suggesting it too is tied to the physical setup and/or the arrays them-

Table 6.2: Measurement system parameters and fitted model parameters.

Parameter(s)	**Symbol(s)**	**Default**	**Vertical**	**Tapered**
EIRP	EIRP	60 dBm	60 dBm	54 dBm
Noise Power	P_{noise}	-68 dBm	-68 dBm	-68 dBm
Cluster Centers of $\bar{\mathbf{H}}$	$\left(\vartheta_{\text{tx}}^{(k)}, \varphi_{\text{tx}}^{(k)}\right), \left(\vartheta_{\text{rx}}^{(k)}, \varphi_{\text{rx}}^{(k)}\right)$	See Table 6.1	See Table 6.1	See Table 6.1
Angular Spread of $\bar{\mathbf{H}}$	$(\Delta\vartheta, \Delta\varphi)$	$(4, 3)$	$(4, 3)$	$(4, 3)$
Location Parameter of $\hat{\mu}$	$\bar{G}^2$	-129.00 dB	-141.58 dB	-144.58 dB
Scale Parameter of $\hat{\mu}$	ξ	0.502	0.527	0.498
Estimator Parameters of $\hat{\sigma}^2$	(α, β)	$(-0.733, 42.53)$	$(-0.588, 29.71)$	$(-0.822, 25.42)$
Variance Parameter of $\hat{\sigma}^2$	ν^2	126.091	75.794	110.391

selves. Relative to the default configuration, $\bar{G}^2$ decreases in the vertical and tapered configurations. In the vertical configuration, this is perhaps courtesy of extra isolation offered by the array enclosures and elements when rotated sideways. In the tapered configuration, this decrease in $\bar{G}^2$ is perhaps due to the extra isolation offered by reduced side lobe and back lobe levels. Connections between the other model parameters across the three measurements is less obvious, and fleshing out how system factors impact all model parameters is an attractive topic for future work.

While not explicitly shown due to space constraints, the additional two measurement sets yielded fitted models that offered similar results as those presented for the default setup in Section 6.5. This confirms that our model can indeed be generalized to other systems through unique parameterization. Furthermore, the fact that the fitted model parameters of all three measurement sets are of the same order—and even closely align in some cases—strengthens the belief that our model and its parameters have strong underpinnings to the real world, even if they are not yet fully understood. In general, however, the viability of and methodology for extending our model to systems beyond our own requires contributions and measurements from the research community at large. Our hope is that the model presented herein is a first step toward even more rigorous modeling of self-interference in full-duplex mmWave systems, with the end goal being a practical model of the self-interference channel matrix $\mathbf{H}$, if possible.

6.6.2 Applications of Our Model

Our model allows engineers to draw statistical realizations of self-interference that a practical full-duplex mmWave system would incur when using particular transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$. A full-duplex mmWave system employing codebook-based beamforming, for example, will see unique levels of self-interference between each transmit-receive beam pair—our model can be used to realize these levels of self-interference. In such a case, transmit-receive beam pairs offering low self-interference will be more desirable in a full-duplex sense. This simple fact raises interesting questions about the potential for user selection to facilitate full-duplex operation, which can be further investigated using our statistical model. Also, note that beams $\mathbf{f}$ and $\mathbf{w}$ do not necessarily have to lay on our grid of measured beams. They can potentially be within our 1° resolution or beyond our measured spatial profiles defined in (6.7), but we cannot be sure of the reliability of such with verifying it through measurement. Our model also readily accommodates the case where transmit and receive beams steer in the same direction—i.e., $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}) = (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$—such as when communicating with another full-duplex device.

Another important application of our model is its role in creating and evaluating beamforming-based solutions for full-duplex mmWave systems. While we have evaluated our model using conjugate beams and beams with tapered side lobes, we cannot confidently expect our model to produce accurate realizations of self-interference with *arbitrary* transmit and receive beams $\mathbf{f}$ and $\mathbf{w}$; this would require extremely precise modeling of the self-interference channel $\mathbf{H}$. As a result, we can only be confident in our model when considering highly directional beams, especially the conjugate beams we have used in (6.1). Therefore, our model should not be used directly to evaluate beamforming-based full-duplex solutions that tailor transmit and receive beams based on knowledge of the true self-interference channel $\mathbf{H}$.

Nonetheless, our model is suitable for evaluating some beamforming-based solutions, especially those that rely purely on beam steering to reduce self-interference. STEER [3], for instance, is a beam refinement methodology that leverages the variability of self-interference over small spatial neighborhoods to significantly reduce self-interference without compromising downlink and uplink beamforming gain. Our statistical model can be used to evaluate STEER, as we will show in the next section, and can facilitate the creation and evaluation of other beamforming-based solutions for full-duplex beyond STEER. Since our model is a stochastic one, it can be drawn from indefinitely to conduct research over a broad variety of realizations of self-interference—perhaps network-wide—which can shed light on potential edge-cases and lower-tail performance. Finally, our model's realizations of self-interference can also be used to drive requirements on the degree of additional cancellation required by full-duplex mmWave systems employing analog or digital SIC.

6.7 Evaluating Our Statistical Model

In this section, we evaluate our model of self-interference by comparing realizations drawn from our model against those observed in measurements.

We aim to show that our model is capable of capturing the large-scale spatial trends and the small-scale spatial variability seen in our measurements. As part of this, we use our model to evaluate an existing full-duplex solution and compare the results to an evaluation using actual measurements.

6.7.1 Comparing Our Statistical Model to Measurements

Ideally, values of self-interference drawn from our model should align spatially and in distribution with what was measured. To evaluate such quantitatively, consider the following. Let $\mathcal{K}_{\mathrm{tx}}$ be a random subset of K elements from $\{1, \ldots, N_{\mathrm{tx}}\}$ without replacement. In other words, $\mathcal{K}_{\mathrm{tx}}$ contains K random transmit indices. Let $\mathcal{K}_{\mathrm{rx}}$ be defined analogously as a set of K random receive indices.

$$\mathcal{K}_{\mathrm{tx}} = \mathrm{randk}(\{1, \ldots, N_{\mathrm{tx}}\}, K), \qquad \mathcal{K}_{\mathrm{rx}} = \mathrm{randk}(\{1, \ldots, N_{\mathrm{rx}}\}, K) \tag{6.30}$$

Then, let $\mathcal{M}$ be the set of measured INR values for each transmit-receive beam pair whose indices are elements of $\mathcal{K}_{\mathrm{tx}}$ and $\mathcal{K}_{\mathrm{rx}}$, defined as

$$\mathcal{M} = \{[\mathsf{INR}_{mn}]_{\mathrm{dB}} : m \in \mathcal{K}_{\mathrm{tx}},\ n \in \mathcal{K}_{\mathrm{rx}}\}, \tag{6.31}$$

where we employ the shorthand $\mathsf{INR}_{mn} \triangleq \mathsf{INR}\left(\theta_{\mathrm{tx}}^{(m)}, \phi_{\mathrm{tx}}^{(m)}, \theta_{\mathrm{rx}}^{(n)}, \phi_{\mathrm{rx}}^{(n)}\right)$. We can draw a statistical realization of $\mathcal{M}$ by forming $\tilde{\mathcal{M}}$ as

$$\tilde{\mathcal{M}} = \left\{ \left[\tilde{\mathsf{INR}}_{mn}\right]_{\mathrm{dB}} \sim \mathcal{N}\left(\hat{\mu}_{mn}, \hat{\sigma}^2_{mn}\right) : m \in \mathcal{K}_{\mathrm{tx}},\ n \in \mathcal{K}_{\mathrm{rx}} \right\}, \tag{6.32}$$

where $\tilde{\mathsf{INR}}_{mn}$ is drawn from our statistical model. By comparing $\mathcal{M}$ and $\tilde{\mathcal{M}}$, we can gain insight into how well our model can realize self-interference that is statistically and spatially similar to our measurements.

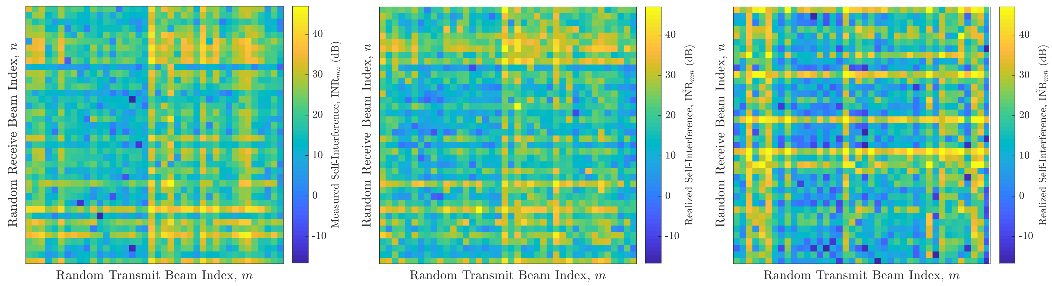

(a) Actual measurements. (b) Proposed model. (c) Spherical-wave model [2].

Figure 6.9: (a) Measured self-interference $\mathcal{M}$ for $K = 40$ random transmit beams and $K = 40$ random receive beams. (b) A realized counterpart $\tilde{\mathcal{M}}$ of (a) using our statistical model. (c) A realized counterpart of (a) using the spherical-wave model [2]. The structural similarity between (a) and (b) illustrates that our model captures the dominant spatial components of self-interference in our measurements, whereas (c) does not.

Consider Figure 6.9, which compares measurements $\mathcal{M}$ and realizations $\tilde{\mathcal{M}}$ of self-interference for $K = 40$ random transmit beams and receive beams; also shown is self-interference one would expect based on the spherical-wave model [2] that has been widely used in the literature.[2] Notice the color scale is common across all three. Visually, one can observe the strong structural similarity between Figure 6.9a and Figure 6.9b. The transmit beams (columns) that tend to inflict higher self-interference in our measurements also tend to when using our model to realize self-interference. As observed in our measurements, these transmit beams do not yield high self-interference universally. Rather, there are select receive beams that couple low self-interference, even

[2]To model self-interference under the spherical-wave channel [2], we constructed $\mathbf{H}$ based on the geometry of our phased array platform and used (6.3), where G^2 was fitted to align in median with our measurements of self-interference.

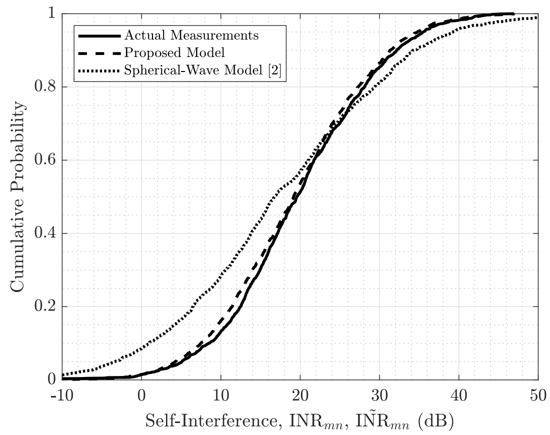

Figure 6.10: The CDFs of the measured and realized self-interference in Figure 6.9a and Figure 6.9b, aligned closely in distribution. Also shown is the CDF of Figure 6.9c based on the spherical-wave model [2], which aligns neither structurally nor in distribution with measurements.

if the majority couple high self-interference. Likewise can be said about the receive beams (rows) that tend to couple higher degrees of self-interference. Similar conclusions are drawn when considering beams that tend to yield low self-interference, though these are less pronounced in general. Comparing Figure 6.9a and Figure 6.9b importantly demonstrates that our model is capable of realizing self-interference that coincides spatially with our measurements. Figure 6.9c, on the other hand, illustrates that the spherical-wave model [2] does not at all align with our measurements of real-world self-interference—observed also in Figure 6.3.

The CDFs of the three figures in Figure 6.9 are depicted in Figure 6.10, which shows that our model produces self-interference that is statistically similar to that of measurements. While our model may not *exactly* replicate

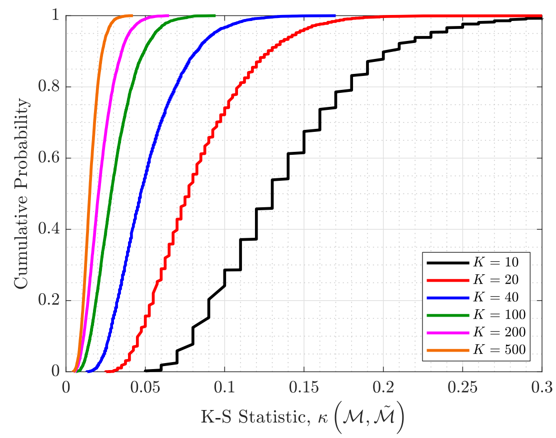

Figure 6.11: The CDFs of the K-S statistic $\kappa\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$ between measured and realized self-interference, where each CDF is taken across 5000 random $\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$ for each K. As K increases, the measured and realized self-interference more closely align in distribution.

what has been measured (by design), it importantly does yield *distributions* of self-interference in line with measurements while also capturing dominant spatial characteristics. The spherical-wave model [2], however, does not align closely in distribution with actual measurements, exhibiting heavier upper and lower tails.

Now, consider Figure 6.11, where we extend this comparison of measurements and our proposed model over many realizations and for various K. For a given K, we draw 5000 random measurements $\mathcal{M}$ and realizations $\tilde{\mathcal{M}}$ with our model, and for each, we compute the K-S statistic $\kappa\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$. The K-S statistic here is the difference in density between the CDFs of $\mathcal{M}$ and $\tilde{\mathcal{M}}$ and desirably is less than 0.1. Figure 6.11 depicts the resulting CDF of the K-S statistic over all 5000 random $\left(\mathcal{M}, \tilde{\mathcal{M}}\right)$. With $K = 10$, the K-S statistic lay below 0.2 for 90% of the random realizations. As K increases, our model produces realizations of self-interference $\tilde{\mathcal{M}}$ that more closely align in distribution to actual measurements $\mathcal{M}$. With $K = 40$, for instance, over 96% of the time, our model produces realizations of $\tilde{\mathcal{M}}$ whose distribution differs from measurements $\mathcal{M}$ by at most 0.1 in density. This suggests that our model can be reliably used to realize self-interference that aligns statistically with our measurements, even with a fairly modest number of samples.

The comparison of our statistical model to measurements in Figure 6.9 and Figure 6.10 is particularly applicable in the context of codebook-based full-duplex mmWave systems, as mentioned in Subsection 6.6.2. A system employing transmit and receive codebooks, each with K beams for instance, will see a unique degree of self-interference between each of the K^2 transmit-receive beam pairs. With our model, engineers can draw realizations of self-interference between each beam pair, producing a matrix of self-interference that aligns spatially and statistically with our measurements, especially if the number of beams is at least around $K = 20$.

6.7.2 Using Our Model to Design and Evaluate Solutions

Recent work [3] proposes a beamforming-based solution for full-duplex mmWave systems called Steer. Steer is a measurement-driven beam refinement methodology for full-duplex mmWave systems that can be applied atop conventional codebook-based beam selection. Suppose initial transmit and receive beam selections (i.e., steering directions) have been made at a full-duplex mmWave transceiver to serve some downlink user and uplink user. Surrounding these initial transmit and receive steering directions, STEER constructs spatial neighborhoods based on some resolution and size. Each transmit-receive steering combination within the spatial neighborhood is measured until a beam pair offering sufficiently low self-interference is found, minimizing the distance the selected beam pair is from the initial beam selection. As such, the success of STEER as a full-duplex solution relies heavily on the variability of self-interference over small spatial neighborhoods. Therefore, to confirm our model can be used to accurately evaluate STEER, it must produce realizations of self-interference that offer variability over small spatial neighborhoods that is of a similar degree as that seen in reality.

To investigate if our model can be used to evaluate full-duplex solutions, we replicate the evaluation process of STEER described in [3], with minor changes stated as follows; we refer the reader to [3] for details not explicitly stated herein. We first use measurements of self-interference to evaluate STEER and then use realizations of self-interference from our statistical model. If the two evaluations align, this will provide further confidence in the relevance and practicality of our model. When running STEER, we use a spatial resolution of $(\delta\theta, \delta\phi) = (1^\circ, 1^\circ)$ and a neighborhood size of $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$, along with a self-interference target of $\mathsf{INR} \leq 0$ dB. On the downlink and uplink, we perform initial beam selection using a codebook of beams distributed in azimuth from -56° to 56° with 8° spacing and in elevation from -8° to 8° with 8° spacing, for a total of 45 beams. We use exhaustive beam search

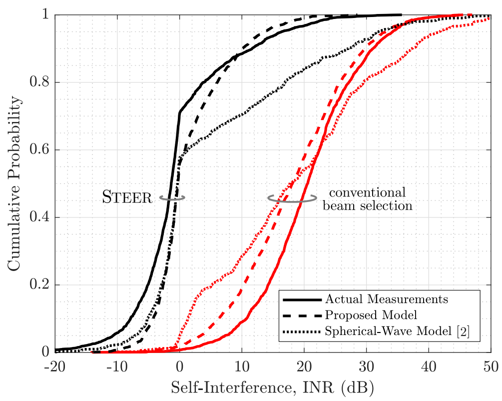

(a) CDF of self-interference.

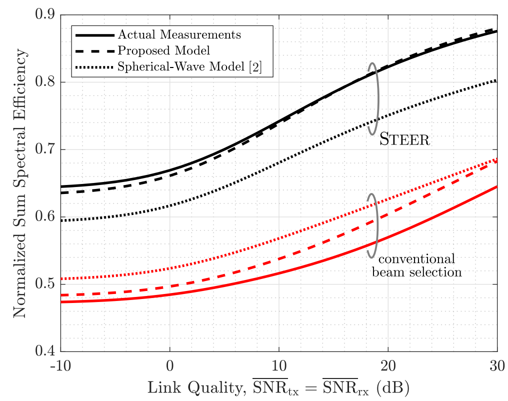

(b) Sum spectral efficiency.

Figure 6.12: Evaluating a beamforming-based full-duplex solution STEER [3] using actual measurements, using our proposed model, and using the spherical-wave model [2]. The (a) CDF of self-interference and (b) normalized sum spectral efficiency when using STEER versus conventional beam selection. Our proposed model more closely aligns with measurements and therefore serves as a more reliable evaluation tool.

to maximize downlink and uplink SNR when conducting conventional beam selection.

The results of our evaluations of STEER are shown in Figure 6.12. First, consider Figure 6.12a, which depicts the distribution of self-interference with and without STEER, when using actual measurements for evaluation (solid lines), when using our statistical model (dashed lines), and when using the spherical-wave model [2] (dotted lines). Without STEER (red lines), the distribution of self-interference across transmit-receive beam pairs following conventional beam selection lay largely above 0 dB for the most part, with most beam pairs yielding self-interference levels prohibitively high for full-duplex. The similarity between the solid and dashed red lines confirms that our model can realize levels of self-interference comparable to reality (i.e., akin to Figure 6.9 and Figure 6.10). With Steer (black lines), lower levels of self-interference can be reached much more reliably: a significant fraction of the time, self-interference is near or below the noise floor. Albeit not perfect, the solid and dashed black lines closely coincide, which importantly shows that our statistical model produces small-scale variability comparable to that seen in real systems. The spherical-wave model [2] falls short in both cases, yielding distributions that do not as closely follow the actual measurements.

We extend this evaluation a step further by inspecting the resultant sum spectral efficiency in Figure 6.12b as a function of downlink and uplink quality. Again, evaluation with our statistical model closely aligns with that using actual measurements, differing in sum spectral efficiency only marginally, while the spherical-wave model [2] differs more substantially. From Figure 6.12a and Figure 6.12b, we do note, however, that our model exhibits slightly more optimistic levels of self-interference with conventional beam selection and slightly more pessimistic levels with Steer, compared to measurements. Nonetheless, it is clear that our proposed model captures the large-scale trends in self-interference and the small-scale variability observed in practice, confirming that it can indeed be used to develop and evaluate solutions for full-duplex mmWave systems—more closely aligning with actual measurements than the spherical-wave model [2].

6.8 Conclusion and Topics for Future Work

In this paper, we presented the first measurement-backed model of self-interference in multi-panel full-duplex mmWave systems. The proposed model can produce realizations of self-interference that are statistically and spatially aligned with our collected measurements. We have effectively reduced the presumably quite complex self-interference channel to a model parameterized by a few values, offering convenience and intuition on the underlying nature of the channel. With appropriate parameterization, our model can translate to systems beyond our own, and we encourage others to fit, evaluate, and publish model parameters to the research community. Importantly, we uncovered a coarse self-interference channel model that suggests that energy propagates from the transmit array to the receive array along clusters of rays in a far-field fashion—a drastic deviation from the spherical-wave model commonly used in the literature thus far. We showed that our model can produce realizations of self-interference that exhibit the large-scale and small-scale trends observed in real systems, making it a useful tool in the design and evaluation of full-duplex solutions and systems—especially for those without access to phased array platforms.

This work motivates and enables a wide variety of future research. First, independent work confirming that our model is applicable to a variety of full-duplex systems at various carrier frequencies and in a multitude of environments would be an essential next step, along with reporting fitted model parameters for such. From this, it would be useful to establish connections between system parameters and model parameters (e.g., how array size impacts α, β, and ν^2). Naturally, we hope our model can serve as useful tool in the advancement of full-duplex mmWave systems by allowing researchers to conduct statistical performance analyses and design full-duplex solutions. Training machine learning models to estimate or produce realizations of self-interference would also be interesting future work. A lofty and important goal for the research community is to develop a model of the self-interference channel (i.e., $\mathbf{H}$) that is backed by measurements; it is our hope that the findings in this paper—especially the coarse channel $\bar{\mathbf{H}}$—can be useful leads in this pursuit.

Chapter 7

Beam Selection for Full-Duplex Millimeter Wave Communication Systems

In Chapter 5, we conducted an extensive measurement campaign of mmWave self-interference and analyzed the nearly 6.5 million collected measurements. One exciting finding from the analysis of these measurements was that slight shifts of the transmit and receive beams can greatly reduce self-interference. In this chapter, we leverage this observed phenomenon to construct a full-duplex solution—called STEER—that relies on strategic beam selection to enable full-duplex operation in mmWave systems. We implement STEER on an actual phased array platform and validate its effectiveness by combining actual measurements of self-interference with extensive simulation.

7.1 Introduction

To equip mmWave transceivers with full-duplex capability, recent work has proposed leveraging dense antenna arrays to mitigate self-interference spatially via strategic beamforming [19, 20, 41, 43, 63, 67–78]. For instance, in [41, 71–77], designs are presented that tailor hybrid beamformers at a mmWave transceiver to mitigate self-interference while maintaining transmission and reception. Proposed solutions [55, 78–83] have also suggested using analog and digital self-interference cancellation to enable mmWave full-duplex; solutions like these require additional hardware, demand complex digital signal processing, and/or do not scale well to systems with many antennas. For these reasons, the scope of the present paper focuses on using beamforming alone to mitigate self-interference spatially, which does not require dedicated hardware nor complex signal processing. If successfully equipped with full-duplex capability, mmWave communication systems could see impressive throughput and latency enhancements, which magnify at the network level and facilitate IAB network deployments [32].

Most proposed beamforming-based solutions to mitigate self-interference are not well-suited for practical systems for a few reasons. First, many practical mmWave systems are equipped with analog beamforming but lack digital beamforming, rendering proposed hybrid beamforming designs (e.g., [41, 71–77]) unfit for such systems. This is especially problematic since digital beamforming mitigates the large majority of self-interference in most of these designs. Proposed designs [43, 68–70] that rely solely on analog beamforming (rather than hybrid beamforming) are also impractical for a few reasons. These rely on instantaneous knowledge of the self-interference channel—a high-dimensional MIMO channel—whose real-time estimation is currently impractical due to complications posed by analog beamforming and its sheer size. Furthermore, [68–77] require instantaneous knowledge of the downlink and uplink MIMO channels between a full-duplex transceiver and the devices it serves; even this is currently impractical in mmWave networks, which circumvent MIMO channel estimation via beam alignment. Some designs [68–73, 77] do not account for phase shifters and/or attenuators with limited resolution that practical analog beamforming networks are subjected to.

Many proposed beamforming designs for full-duplex mmWave systems do not accommodate beam alignment and subsequent analog beam selection [68–76]. Beam alignment is a critical component of practical mmWave systems to provide sufficient link margin to sustain communication without the need for uplink/downlink MIMO channel knowledge [39, 60]. Using measurements from beam alignment, a mmWave system can configure its analog beamformers through *beam selection*. Conventional half-duplex systems typically aim to overcome severe path loss by selecting beams that maximize received SNR. Like half-duplex mmWave systems, a full-duplex one will presumably execute beam selection and will do so on its transmit link and receive link. Naïvely applying conventional beam selection on the two links independently, however, does not account for self-interference that couples between the transmit and receive beams when operating in a full-duplex fashion. This has motivated us to create the first beam selection methodology for full-duplex systems. The two principal contributions of this paper are summarized as follows.

A beam selection methodology for full-duplex systems. We present Steer, the first beam selection methodology for jointly selecting the transmit and receive beams of a full-duplex mmWave transceiver. To do so, we leverage our observations from a recent measurement campaign of mmWave self-interference [44, 113], which showed that small shifts in the steering directions of the transmit and receive beams (on the order of one degree) can lead to noteworthy reductions in self-interference. Steer makes use of self-interference measurements across small spatial neighborhoods to jointly select transmit and receive beams at the full-duplex device that offer reduced self-interference while delivering high beamforming gain on the uplink and downlink. Following its formulation, we present an algorithm for executing Steer with a minimal number of self-interference measurements. The execution of Steer takes place only at the full-duplex device, introducing no changes to the devices being served nor additional over-the-air feedback, making it compatible with existing beam alignment schemes.

Validation of Steer through measurement and simulation. We validate Steer by combining simulation with self-interference measurements from an industry-grade 28 GHz phased array platform. These experimental results illustrate that full-duplex operation with Steer can offer a sum spectral efficiency notably higher than both half-duplex and full-duplex operation with beams from conventional beam alignment. In fact, in most cases, Steer can reduce self-interference to levels such that no additional cancellation is warranted (i.e., near or below the noise floor). With Steer, beamforming alone can deliver self-interference mitigation sufficient for full-duplex while importantly accommodating beam alignment, even in the presence of cross-link interference arising when serving devices simultaneously and in-band.

7.2 System Model

As a relevant application of this work, we consider the mmWave communication system shown in Figure 7.1, where a sectorized, multi-panel IAB node maintains backhaul and serves access in a full-duplex fashion (i.e., on the same time-frequency resources) [120]. This is particularly attractive application of full-duplex in mmWave cellular systems [94, 115], but the contributions of this work are not limited to IAB. Illustrated in Figure 7.1 is the downlink-downlink (DL-DL) operating mode, where an IAB node receives backhaul from a fiber-connected donor and transmits access to a user equipment (UE). In this work, we present a formulation and design that generalizes to both the DL-DL operating mode and the analogous uplink-uplink (UL-UL) mode where the IAB node receives access and transmits backhaul. We are particularly interested in these full-duplex operating modes since they unlock scheduling opportunities that can reduce latency in IAB networks while also increasing throughput [32]. It is important to consider and evaluate both full-duplexing modes since there may exist significant disparities between the donor and UE—most notably transmit power, noise power, and number of antennas.

Separate UPAs are present at the IAB node, each of which can either transmit or receive and can be independently configured via a network of analog beamforming weights. In this multi-panel full-duplex setting, one array will transmit while the other receives. To simplify notation between the DL-DL and UL-UL modes, we assume each array at the IAB node is equipped with N_{a} antennas. We denote the vector of transmit beamforming weights

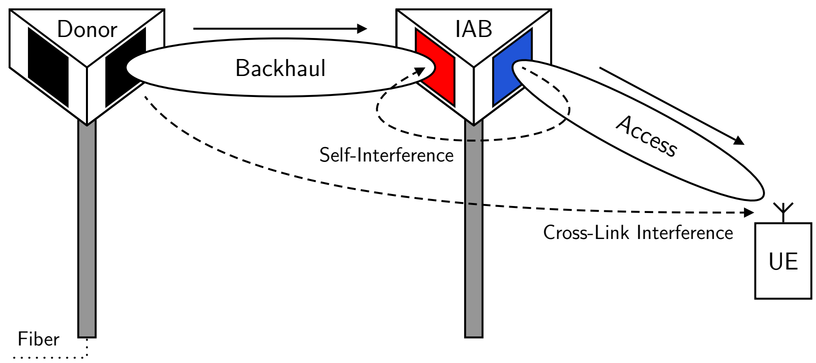

Figure 7.1: A full-duplex IAB node receives backhaul from a fiber-connected donor node while simultaneously transmitting access to a UE, giving rise to self-interference at the IAB node and cross-link interference at the UE. We refer to this as the DL-DL operating mode. We also consider the UL-UL mode where the UE transmits uplink access and the IAB node transmits backhaul.

at the IAB node as $\mathbf{f} \in \mathbb{C}^{N_\mathrm{a} \times 1}$. Likewise, the receive beamforming vector at the IAB node is denoted $\mathbf{w} \in \mathbb{C}^{N_\mathrm{a} \times 1}$. For transmit power and noise power normalizations, we assume that the beamforming weights have unit power as $\|\mathbf{f}\|_2^2 = \|\mathbf{w}\|_2^2 = 1$. Extending this work to systems with multiple beams at the transmitter and receiver would be great future work.

For simplicity, we assume the UE is a single-antenna device, though the work herein could extend naturally to those with multiple antennas. Let the row vector $\mathbf{h}_{\mathrm{AC}}^* \in \mathbb{C}^{1 \times N_\mathrm{a}}$ be the access channel between the transmit array of the IAB node and the UE. Practically, the donor will have an antenna array through which it serves backhaul. With an array at the donor and an array at the IAB node, we use $\mathbf{H}_{\mathrm{BH}}$ to denote the MIMO backhaul channel matrix between the donor and the receive panel of the IAB node. We assume the donor transmits with some beamforming weights $\mathbf{v}$ and instead consider

Table 7.1: Transmit and receive links during different full-duplexing modes.

Mode	**Transmit Link**	**Receive Link**
Downlink-Downlink	Access	Backhaul
Uplink-Uplink	Backhaul	Access

henceforth the column vector

$$\mathbf{h}_{\mathrm{BH}} \triangleq \mathbf{H}_{\mathrm{BH}}\mathbf{v} \in \mathbb{C}^{N_{\mathrm{a}} \times 1} \tag{7.1}$$

which is the effective backhaul channel from the beamformed donor to the IAB node—abstracting out beamforming at the donor. We normalize the access and backhaul channel vectors as $\|\mathbf{h}_{\mathrm{AC}}\|_2^2 = \|\mathbf{h}_{\mathrm{BH}}\|_2^2 = N_{\mathrm{a}}$ and abstract out their large-scale path gains (inverse path loss) as G_{AC}^2 and G_{BH}^2, respectively. We invite readers to assume access and backhaul channels that are LOS for simplicity, but this work does not depend on such.

In the DL-DL operating mode—when transmitting access and receiving backhaul from the IAB node in a full-duplex fashion—a MIMO self-interference channel $\mathbf{H} \in \mathbb{C}^{N_{\mathrm{a}} \times N_{\mathrm{a}}}$ manifests between the transmit and receive arrays of the IAB node. We similarly abstract out its large-scale path gain as G^2 and enforce $\|\mathbf{H}\|_{\mathrm{F}}^2 = N_{\mathrm{a}}^2$. In addition, during DL-DL, since the donor transmits backhaul while the UE receives access, a cross-link interference channel $\mathbf{h}_{\mathrm{CL}}^*$ (a row vector) manifests between the donor's antenna array and the UE. Having conditioned on some beamforming weights $\mathbf{v}$ at the donor, the effective cross-link interference channel is the scalar

$$h_{\mathrm{CL}} \triangleq \mathbf{h}_{\mathrm{CL}}^* \mathbf{v} \in \mathbb{C}^{1 \times 1}. \tag{7.2}$$

We similarly abstract out its large-scale gain as G_{CL}^2 by letting $|h_{\mathrm{CL}}|^2 = 1$. Symbols are transmitted by the donor, IAB node, and UE with powers (in watts) $P_{\mathrm{tx}}^{\mathrm{Don}}$, $P_{\mathrm{tx}}^{\mathrm{IAB}}$, and $P_{\mathrm{tx}}^{\mathrm{UE}}$, respectively. Additive noise incurred at the donor, IAB node, and UE have respective powers (in watts) $P_{\mathrm{noise}}^{\mathrm{Don}}$, $P_{\mathrm{noise}}^{\mathrm{IAB}}$, and $P_{\mathrm{noise}}^{\mathrm{UE}}$. With these definitions in hand, we can formulate the quality of each link in the DL-DL and UL-UL modes.

7.3 Problem Formulation and Motivation

We leverage our formulations of the DL-DL and UL-UL modes to present a general formulation of the system. Taking the perspective of the full-duplex IAB node, we introduce the terms *transmit link* and *receive link*, which correspond to either access or backhaul depending on if the system operates in DL-DL or UL-UL mode, as summarized in Table 7.1. In the DL-DL mode, for instance, the SNR of the transmit link and receive link can be expressed respectively as

$$\mathsf{SNR}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{IAB}} \cdot G_{\mathrm{AC}}^2 \cdot |\mathbf{h}_{\mathrm{AC}}^* \mathbf{f}|^2}{P_{\mathrm{noise}}^{\mathrm{UE}}}, \tag{7.3}$$

$$\mathsf{SNR}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{Don}} \cdot G_{\mathrm{BH}}^2 \cdot |\mathbf{w}^* \mathbf{h}_{\mathrm{BH}}|^2}{P_{\mathrm{noise}}^{\mathrm{IAB}}}. \tag{7.4}$$

By virtue of the fact that $N_{\mathrm{a}} = \max_{\mathbf{x}} \ |\mathbf{h}^* \mathbf{x}|^2$ s.t. $\|\mathbf{x}\|_2^2 = 1, \|\mathbf{h}\|_2^2 = N_{\mathrm{a}}$, the maximum SNRs during the DL-DL mode, denoted with an overline, are

$$\overline{\mathsf{SNR}}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{IAB}} \cdot G_{\mathrm{AC}}^2 \cdot N_{\mathrm{a}}}{P_{\mathrm{noise}}^{\mathrm{UE}}} \geq \mathsf{SNR}_{\mathrm{tx}}, \tag{7.5}$$

$$\overline{\mathsf{SNR}}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{Don}} \cdot G_{\mathrm{BH}}^2 \cdot N_{\mathrm{a}}}{P_{\mathrm{noise}}^{\mathrm{IAB}}} \geq \mathsf{SNR}_{\mathrm{rx}}. \tag{7.6}$$

These are achieved by beamforming directly toward the UE and donor to deliver maximum gain. In the DL-DL mode, access is corrupted by cross-link interference and backhaul is corrupted by self-interference, leading to the INR of the transmit and receive links being respectively

$$\mathsf{INR}_{\mathrm{tx}} = \frac{P_{\mathrm{tx}}^{\mathrm{Don}} \cdot G_{\mathrm{CL}}^2 \cdot \left|h_{\mathrm{CL}}\right|^2}{P_{\mathrm{noise}}^{\mathrm{UE}}}, \tag{7.7}$$

$$\mathsf{INR}_{\mathrm{rx}} = \frac{P_{\mathrm{tx}}^{\mathrm{IAB}} \cdot G^2 \cdot \left|\mathbf{w}^*\mathbf{H}\mathbf{f}\right|^2}{P_{\mathrm{noise}}^{\mathrm{IAB}}}. \tag{7.8}$$

Notice that the degree of self-interference depends on its channel $\mathbf{H}$ and the beamformers $\mathbf{f}$ and $\mathbf{w}$ at the IAB node. Cross-link interference, on the other hand, does not depend on $\mathbf{f}$ nor $\mathbf{w}$ and is fixed for a given setting, having conditioned on the donor's beamformer $\mathbf{v}$. All terms presented here for the DL-DL mode can be defined analogously for the UL-UL mode, with backhaul as the transmit link and access as the receive link. Note that, regardless of operating mode, the transmit link is plagued by cross-link interference and the receive link is plagued by self-interference.

With these formulations of the transmit and receive links, their SINRs can be written as

$$\mathsf{SINR}_{\mathrm{tx}} = \frac{\mathsf{SNR}_{\mathrm{tx}}}{1 + \mathsf{INR}_{\mathrm{tx}}}, \qquad \mathsf{SINR}_{\mathrm{rx}} = \frac{\mathsf{SNR}_{\mathrm{rx}}}{1 + \mathsf{INR}_{\mathrm{rx}}}. \tag{7.9}$$

Treating interference as noise, the maximum achievable spectral efficiencies on each link are

$$R_{\mathrm{tx}} = \log_2(1 + \mathsf{SINR}_{\mathrm{tx}}), \qquad R_{\mathrm{rx}} = \log_2(1 + \mathsf{SINR}_{\mathrm{rx}}) \tag{7.10}$$

whereas the individual Shannon capacities of each of these links are

$$C_{\mathrm{tx}} = \log_2\left(1 + \overline{\mathsf{SNR}}_{\mathrm{tx}}\right), \qquad C_{\mathrm{rx}} = \log_2\left(1 + \overline{\mathsf{SNR}}_{\mathrm{rx}}\right). \tag{7.11}$$

In this paper, we seek a means to strategically select beams $\mathbf{f}$ and $\mathbf{w}$ such that self-interference can be significantly reduced and spectral efficiency can be improved over conventional/naïve approaches. Taking a full-duplex perspective, we desire a sum spectral efficiency $R_{\mathrm{tx}} + R_{\mathrm{rx}}$ that approaches the full-duplex capacity $C_{\mathrm{tx}} + C_{\mathrm{rx}}$. In this pursuit, we account for a number of practical considerations in this work, most notably limited channel knowledge and practical codebook-based beam alignment.

In this work, we aim to leverage small-scale phenomena observed in our recent measurement campaign [44,113] by slightly shifting transmit and receive beams so that self-interference (i.e., $\mathsf{INR}_{\mathrm{rx}}$) can be reduced while still delivering high beamforming gain (i.e., high $\mathsf{SNR}_{\mathrm{tx}}$ and $\mathsf{SNR}_{\mathrm{rx}}$) in the desired directions. To do this in a systematic manner, we introduce the first beam selection methodology specifically for full-duplex mmWave systems, called STEER, that makes use of self-interference measurements to choose beams that reduce self-interference and facilitate full-duplex operation. In fact, we show that STEER can reduce self-interference to levels sufficiently low for full-duplex operation without the need for any additional analog nor digital cancellation. This is particularly desirable because it removes the need for additional hardware and signal processing that conventional self-interference cancellation strategies demand. In the sections that follow, we present the three components of STEER:

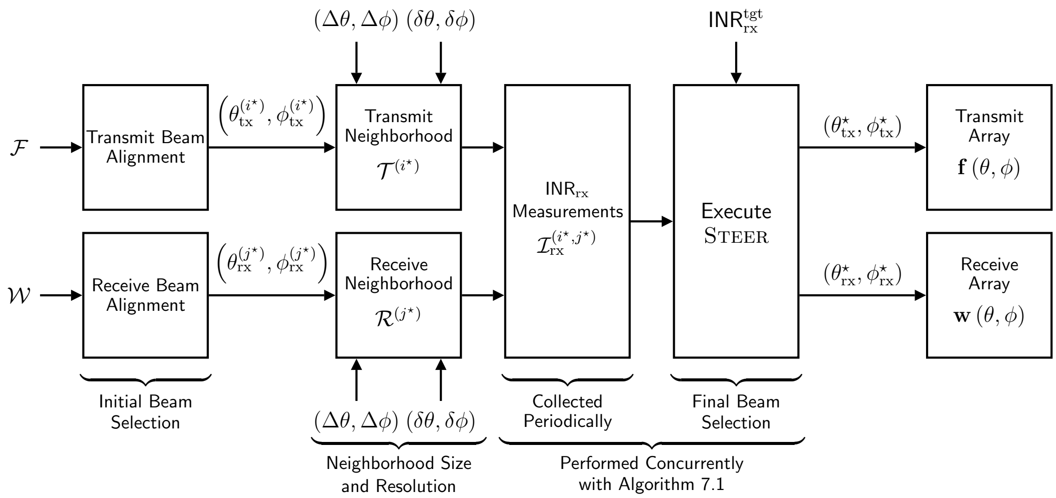

Figure 7.2: A block diagram summarizing beam selection via STEER, which jointly selects transmit and receive beams to reduce self-interference at a full-duplex mmWave transceiver. Conventional beam selection chooses transmit and receive beams independently, ignoring self-interference.

1. Initial beam selection via conventional beam alignment (Section 7.4).
2. Measurements of self-interference across small spatial neighborhoods (Section 7.5); these are not necessarily taken in real-time but rather periodically as needed.
3. Final beam selection to minimize self-interference (Section 7.6).

A block diagram summarizing STEER is shown in Figure 7.2, whose details will become clear as we present our design in the next three sections.

7.4 Initial Selection: Conventional Beam Alignment

Practical mmWave communication systems rely on beam alignment schemes to deliver high beamforming gain. These schemes typically involve sweeping candidate beams, measuring the RSRP for each candidate, and delivering feedback before determining the beam(s) for the mmWave link [39,60]. Candidate beams often come from a codebook, which is constructed by first defining a service region (some portion of space based on an assumed user distribution) and then discretizing it based on the desired number of beams in the codebook or their beamwidth.

In a traditional half-duplex fashion, we suppose the IAB node conducts beam alignment on its transmit link with N_{tx} beams and on its receive link with N_{rx} beams. The N_{tx} transmit beams and N_{rx} receive beams are spatially distributed over their desired coverage regions, where each beam is responsible for serving some portion of its respective region. Describing the steering direction of each beam in an azimuth-elevation fashion, the collection of transmit directions $\mathcal{A}_{\mathrm{tx}}$ and receive directions $\mathcal{A}_{\mathrm{rx}}$ we write as

$$\mathcal{A}_{\mathrm{tx}} = \left\{ \left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)} \right) : i = 1, \ldots, N_{\mathrm{tx}} \right\}, \tag{7.12}$$

$$\mathcal{A}_{\mathrm{rx}} = \left\{ \left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)} \right) : j = 1, \ldots, N_{\mathrm{rx}} \right\}. \tag{7.13}$$

For instance, suppose $\mathcal{A}_{\mathrm{tx}}$ and $\mathcal{A}_{\mathrm{rx}}$ are each comprised of $N_{\mathrm{tx}} = N_{\mathrm{rx}}$ directions distributed uniformly from -60° to 60° in azimuth and from -30° to 30° in elevation.

Practical phased array systems are often equipped with a mapping from desired steering direction to beamforming weights based on some beam design methodology.[1] As such, let $\mathcal{F} = \{\mathbf{f}(\theta, \phi) : (\theta, \phi) \in \mathcal{A}_{\mathrm{tx}}\}$ and $\mathcal{W} = \{\mathbf{w}(\theta, \phi) : (\theta, \phi) \in \mathcal{A}_{\mathrm{rx}}\}$ be the transmit and receive codebooks used for conventional beam alignment, where $\mathbf{f}(\theta, \phi)$ and $\mathbf{w}(\theta, \phi)$ are transmit and receive weights designed to steer toward some (θ, ϕ). Let $\mathbf{h}_{\mathrm{tx}}^{*}$ and $\mathbf{h}_{\mathrm{rx}}$ be the transmit and receive channels corresponding to the particular operating mode. Conventional beam alignment on the transmit link aims to solve (or approximately solve)

$$i^{\star} = \underset{i \in \{1,\ldots,N_{\mathrm{tx}}\}}{\mathrm{argmax}} \ \left| \mathbf{h}_{\mathrm{tx}}^{*} \mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right) \right|^{2} \tag{7.14}$$

to identify the transmit beam $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right) \in \mathcal{F}$ that maximizes beamforming gain delivered on the transmit link. In other words, the transmit link user—either the donor or the UE depending on the mode—is approximately located in the direction $\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right)$ from the transmit panel of the IAB node. Likewise, beam alignment on the receive link aims to solve

$$j^{\star} = \underset{j \in \{1,\ldots,N_{\mathrm{rx}}\}}{\mathrm{argmax}} \ \left| \mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)^{*} \mathbf{h}_{\mathrm{rx}} \right|^{2} \tag{7.15}$$

which identifies the approximate direction of the receive link user. In practice, solving these optimization problems is typically done through a series of of RSRP measurements and feedback between the IAB node and the user it aims to serve; recall, we do not have knowledge of $\mathbf{h}_{\mathrm{tx}}$ nor $\mathbf{h}_{\mathrm{rx}}$. As the first stage of our design, we propose that beam alignment be executed in a half-duplex fashion to yield some $\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^{\star})}, \phi_{\mathrm{rx}}^{(j^{\star})}\right)$, though we do not suggest any

[1] It is not uncommon for this mapping to be proprietary and to account for nonidealities in the array pattern.

particular scheme for doing so. As such, our design can accommodate existing beam alignment schemes without changes (including hierarchical schemes) and does not introduce any additional over-the-air feedback. If using the beams from conventional beam alignment, the nominal SNRs of the transmit and receive links are

$$\mathsf{SNR}_{\mathrm{tx}}^{\mathrm{nom}} \triangleq \overline{\mathsf{SNR}}_{\mathrm{tx}} \cdot \frac{\left|\mathbf{h}_{\mathrm{tx}}^{*}\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right)\right|^{2}}{N_{\mathrm{a}}}, \tag{7.16}$$

$$\mathsf{SNR}_{\mathrm{rx}}^{\mathrm{nom}} \triangleq \overline{\mathsf{SNR}}_{\mathrm{rx}} \cdot \frac{\left|\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j^{\star})}, \phi_{\mathrm{rx}}^{(j^{\star})}\right)^{*}\mathbf{h}_{\mathrm{rx}}\right|^{2}}{N_{\mathrm{a}}}. \tag{7.17}$$

These SNRs are some fraction of the maximum link SNRs based on how effectively the selected beams from conventional beam alignment steer toward the transmit and receive users, which naturally depends on their locations, the environment, and the beam codebooks. The beams output by conventional beam selection will initialize STEER's pursuit to find beams that offer high SNR and reduced self-interference. In doing so, STEER relies on measurements outlined in the following section.

7.5 Self-Interference Across Small Spatial Neighborhoods

Our recent measurement campaign [44, 113] illustrated that slightly shifting the steering directions of the transmit and receive beams (on the order of one degree) can greatly reduce self-interference. To find promising transmit and receive beams (ones that offer reduced self-interference) that steer in approximately the same directions as those identified by conventional

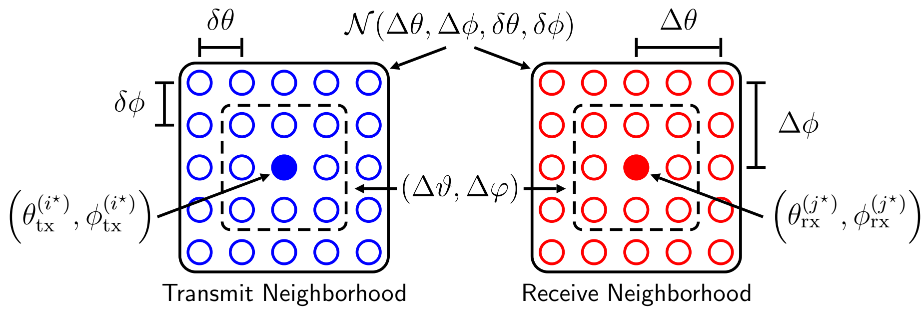

Figure 7.3: The spatial neighborhoods surrounding a given transmit direction and receive direction (shown as filled circles). The size of the neighborhoods is dictated by $(\Delta\theta, \Delta\phi)$ and their resolution by $(\delta\theta, \delta\phi)$. $(\Delta\vartheta, \Delta\varphi)$ will be relevant in Section 7.6.

beam alignment, we explicitly measure self-interference incurred by a number of candidate beams over a small spatial neighborhood. As we will cover shortly, these measurements will not necessarily be taken in real-time (upon each beam selection) but rather may be collected periodically then referenced in real-time, according to the dynamics of self-interference.

If transmitting toward $(\theta_{\text{tx}}, \phi_{\text{tx}})$ and receiving toward $(\theta_{\text{rx}}, \phi_{\text{rx}})$, the IAB node incurs some degree of self-interference, which can be theoretically computed based on (7.8), assuming knowledge of $\mathbf{f}(\theta_{\text{tx}}, \phi_{\text{tx}})$, $\mathbf{w}(\theta_{\text{rx}}, \phi_{\text{rx}})$, G, and $\mathbf{H}$. Practically, however, it is difficult to efficiently and accurately estimate the self-interference channel matrix $\mathbf{H}$ (which is large). Moreover, characterization and modeling of mmWave self-interference is extremely limited, and it is currently impractical for a system to predict what level of self-interference it would incur with a particular transmit beam $\mathbf{f}(\theta_{\text{tx}}, \phi_{\text{tx}})$ and receive beam $\mathbf{w}(\theta_{\text{rx}}, \phi_{\text{rx}})$. All of this—combined with the fact that minor errors in self-interference power can make a significant difference in full-duplex system performance—motivates us to explicitly measure self-interference incurred at the IAB node for particular transmit and receive beams, rather than attempt to estimate it.

To identify attractive steering directions for full-duplex operation, we are interested in measuring the self-interference incurred when transmitting and receiving around the *spatial neighborhoods* surrounding a given transmit direction and receive direction, as illustrated in Figure 7.3 and akin to those introduced in Figure 5.13. Quantifying the size of these spatial neighborhoods, let $\Delta\theta$ and $\Delta\phi$ be maximum absolute azimuthal and elevational deviations from the given transmit direction and receive direction. The spatial neighborhood can be thought of as living within the codebook beam spacing; for instance, $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ when the codebook beams are separated by $(8^\circ, 8^\circ)$. Discretizing these neighborhoods, let $\delta\theta$ and $\delta\phi$ be the measurement resolution in azimuth and elevation, respectively, which should not be larger than $(\Delta\theta, \Delta\phi)$—e.g., $(\delta\theta, \delta\phi) = (1^\circ, 1^\circ)$. The spatial neighborhood $\mathcal{N}$ surrounding a transmit/receive direction can be expressed using the azimuthal neighborhood $\mathcal{N}_\theta$ and elevational neighborhood $\mathcal{N}_\phi$ defined as

$$\mathcal{N}_\theta(\Delta\theta, \delta\theta) = \left\{ m \cdot \delta\theta : m \in \left[-\left\lfloor \frac{\Delta\theta}{\delta\theta} \right\rfloor, \left\lfloor \frac{\Delta\theta}{\delta\theta} \right\rfloor \right] \right\}, \tag{7.18}$$

$$\mathcal{N}_\phi(\Delta\phi, \delta\phi) = \left\{ n \cdot \delta\phi : n \in \left[-\left\lfloor \frac{\Delta\phi}{\delta\phi} \right\rfloor, \left\lfloor \frac{\Delta\phi}{\delta\phi} \right\rfloor \right] \right\} \tag{7.19}$$

where $\lfloor \cdot \rfloor$ is the floor operation and $[a, b] = \{a, a+1, \ldots, b-1, b\}$. The complete neighborhood is the Cartesian product of the azimuthal and elevational neighborhoods as

$$\mathcal{N}(\Delta\theta, \Delta\phi, \delta\theta, \delta\phi) = \mathcal{N}_\theta(\Delta\theta, \delta\theta) \times \mathcal{N}_\phi(\Delta\phi, \delta\phi) \tag{7.20}$$

$$= \left\{(\theta, \phi) : \theta \in \mathcal{N}_\theta(\Delta\theta, \delta\theta),\ \phi \in \mathcal{N}_\phi(\Delta\phi, \delta\phi)\right\}. \tag{7.21}$$

The spatial neighborhoods $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$ surrounding the transmit and receive directions output by conventional beam alignment from the previous section are respectively

$$\mathcal{T}^{(i^\star)} = \left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right) + \mathcal{N}(\Delta\theta, \Delta\phi, \delta\theta, \delta\phi), \tag{7.22}$$

$$\mathcal{R}^{(j^\star)} = \underbrace{\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)}_{\text{initial selection}} + \underbrace{\mathcal{N}(\Delta\theta, \Delta\phi, \delta\theta, \delta\phi)}_{\text{neighborhood}}. \tag{7.23}$$

The size of these sets is

$$\left|\mathcal{T}^{(i^\star)}\right| = \left|\mathcal{R}^{(j^\star)}\right| = |\mathcal{N}(\Delta\theta, \Delta\phi, \delta\theta, \delta\phi)| \tag{7.24}$$

$$= (2 \cdot K_\theta + 1) \cdot (2 \cdot K_\phi + 1) \tag{7.25}$$

where $K_\theta = \left\lfloor \frac{\Delta\theta}{\delta\theta} \right\rfloor$ and $K_\phi = \left\lfloor \frac{\Delta\phi}{\delta\phi} \right\rfloor$, indicating that neighborhoods naturally grow with widened $(\Delta\theta, \Delta\phi)$ or finer resolution $(\delta\theta, \delta\phi)$.

When steering its transmit beam toward $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and receive beam toward $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$, the IAB node incurs an INR of $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$. In DL-DL mode, $\mathsf{INR}_{\mathrm{rx}}$ can be expressed as

$$\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) = \frac{P_{\mathrm{tx}}^{\mathrm{IAB}} \cdot G^2 \cdot \left|\mathbf{w}(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})^* \mathbf{H} \mathbf{f}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})\right|^2}{P_{\mathrm{noise}}^{\mathrm{IAB}}} \tag{7.26}$$

and that during UL-UL mode can be stated analogously. For each potential initial beam selection $(i^\star, j^\star)$, we propose that the IAB node measure and record $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$ for all transmit-receive combinations across the neighborhoods $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$ to populate $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ defined as

$$\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)} = \left\{\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) : (\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}) \in \mathcal{T}^{(i^\star)}, (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \in \mathcal{R}^{(j^\star)}\right\} \quad (7.27)$$

The total number of INR measurements collected in $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ is

$$\left|\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}\right| = \left|\mathcal{T}^{(i^\star)}\right| \cdot \left|\mathcal{R}^{(j^\star)}\right| \quad (7.28)$$

$$= (2 \cdot K_\theta + 1)^2 \cdot (2 \cdot K_\phi + 1)^2 \quad (7.29)$$

which is equal for all $(i^\star, j^\star)$ since we have assumed a fixed neighborhood size. This set of receive link INR measurements $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ will enable the next stage of our proposed design.[2]

Remark 7.1: Measurement overhead and frequency. Naturally, conducting these INR measurements at the full-duplex device may become practically prohibitive if the number of measurements grows too large. This depends on the neighborhood size $(\Delta\theta, \Delta\phi)$ and spatial resolution $(\delta\theta, \delta\phi)$, along with how frequently these measurements need to be collected. System engineers can throttle the neighborhood size and/or spatial resolution to reduce the measurement overhead, though this may reduce the effectiveness of STEER, as we will see. In addition to neighborhood size and spatial resolution, the time-variability of self-interference will heavily dictate the overhead of these measurements. In the extreme case, a nearly static self-interference demands

[2] In Section 7.7, we present an algorithm that can dramatically reduce the number of measurements needed by STEER, requiring only a fraction of $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ to be measured.

infrequent self-interference measurements. Highly dynamic self-interference, on the other hand, will demand more frequent measurements for reliability. Notice that these measurements need not be taken strictly following beam alignment; instead, they can be collected for all (i, j) and referenced for particular $(i^\star, j^\star)$, assuming a sufficiently static self-interference channel. In such a case, the set of all INR measurements can be written as

$$\mathcal{I}_{\mathrm{rx}} = \bigcup_{i=1}^{N_{\mathrm{tx}}} \bigcup_{j=1}^{N_{\mathrm{rx}}} \mathcal{I}_{\mathrm{rx}}^{(i,j)}, \tag{7.30}$$

which has cardinality

$$|\mathcal{I}_{\mathrm{rx}}| = N_{\mathrm{tx}} \cdot N_{\mathrm{rx}} \cdot (2 \cdot K_\theta + 1)^2 \cdot (2 \cdot K_\phi + 1)^2 \tag{7.31}$$

assuming no overlapping transmit neighborhoods or receive neighborhoods. It is important to keep in mind that there is no over-the-air feedback associated with these measurements since they are taken between the transmit and receive panels of the full-duplex IAB node. Reliably measuring $\mathsf{INR}_{\mathrm{rx}}$ is key to the methodology that follows, though small measurement errors would be inherently tolerated; exploring in detail how reliable INR measurements must be would be interesting future work. Note that measuring $\mathsf{INR}_{\mathrm{rx}}$ for some beam pairs may lead to levels of self-interference that saturate the receive chain of the full-duplex transceiver, complicating measurement. It would be valuable future work to develop a means to estimate $\mathsf{INR}_{\mathrm{rx}}$ in such cases (e.g., via transmit power control to avoid saturation), though accuracy would not be especially important, as these beam pairs coupling high self-interference would presumably not be selected by STEER, as we will see.

7.6 STEER: Joint Transmit-Receive Beam Selection

In this section, we present STEER, our methodology for choosing beams that the IAB node uses to serve the transmit link and receive link. STEER incorporates self-interference during beam selection rather than blindly aiming to maximize SNR, as is typically done in conventional beam alignment. To do so, STEER relies on conventional beam alignment from Section 7.4 to identify the general directions in which beams should be steered. Then, based on some design parameters, transmit and receive beams ($\mathbf{f}$ and $\mathbf{w}$) are jointly selected to serve access and backhaul while simultaneously minimizing the degree of self-interference they couple. This is done by leveraging self-interference measurements taken at the full-duplex IAB node as described in Section 7.5.

Using the initial beam selections $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$ from beam alignment, along with the INR measurements $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$, the objective of STEER is to fetch the transmit beam and receive beam from the neighborhoods $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$ that offers full-duplexing gains (in terms of spectral efficiency) over simply using $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$ to serve the transmit and receive links. Naturally, when the SNRs of the transmit and receive links are maximized and the interference on each are simultaneously driven to zero, the achievable spectral efficiencies approach their capacities. In pursuit of appreciable R_{tx} and R_{rx}, we therefore aim to achieve high SNR on each link while reducing interference. However, note that the INR on the transmit link $\mathsf{INR}_{\mathrm{tx}}$ (i.e., cross-link interference) is fixed since we have conditioned on the donor's beamforming weights. Thus, STEER aims to select $\mathbf{f}$ and $\mathbf{w}$—

the transmit and receive beams at the IAB node—so that high SNR can be achieved on each link while simultaneously reducing self-interference. Note that we do not require the final beam selections output by STEER to be from the codebooks $\mathcal{F}$ and $\mathcal{W}$ but rather will be drawn to steer within $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$.

7.6.1 Target Self-Interference Level, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$

Suppose there exists some target receive link INR threshold $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ our system desires. For instance, one may choose $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} \approx 0$ dB to ensure self-interference does not overwhelmingly exceed noise. Our design that follows does not guarantee that $\mathsf{INR}_{\mathrm{rx}} \leq \mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ but rather attempts to meet this target and does not incentivize STEER to provide an $\mathsf{INR}_{\mathrm{rx}}$ further below it. As we will see in our results in Section 7.8, choosing a modest $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ will help STEER yield a more fair and globally optimal solution in terms of sum spectral efficiency since it throttles the sacrifices made in its effort to reduce self-interference. Nonetheless, to force STEER to minimize $\mathsf{INR}_{\mathrm{rx}}$, engineers can use $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -\infty$ dB.

7.6.2 Joint Transmit-Receive Beam Selection

Before beginning our design process, we record the receive link INR when steering along the initial beam selections $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$, which we call the *nominal* receive link INR and express as

$$\mathsf{INR}_{\mathrm{rx}}^{\mathrm{nom}} \triangleq \mathsf{INR}_{\mathrm{rx}}\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}, \theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right). \tag{7.32}$$

This is the receive link INR incurred if our full-duplex system were to use conventional beam selection and will thus be a useful benchmark to compare against. Desirably, our final beam selections will yield $\mathsf{INR}_{\mathrm{rx}} < \mathsf{INR}_{\mathrm{rx}}^{\mathrm{nom}}$ to make our design worthwhile. Note that if $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{nom}} \leq \mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$, we need not proceed with our design for this particular $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}, \theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$ since the target is met inherently by the beams from conventional beam alignment. In such a case, we can simply set the transmit beam and receive beam as $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$. When this is not the case, we proceed with our design as follows.

We begin by denoting the minimum INR over the measured spatial neighborhood as

$$\mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}} \triangleq \min\left(\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}\right). \tag{7.33}$$

Then, we form the beam selection problem (7.34) to retrieve the transmit direction $(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star)$ and receive direction $(\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star)$ that the full-duplex IAB node will steer toward.

$$\begin{aligned}
(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star), (\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star) = \quad & \underset{\substack{(\theta_{\mathrm{tx}},\phi_{\mathrm{tx}}) \\ (\theta_{\mathrm{rx}},\phi_{\mathrm{rx}})}}{\operatorname{argmin}} \ \min_{(\Delta\vartheta,\Delta\varphi)} \ \Delta\vartheta^2 + \Delta\varphi^2 && \text{(7.34a)} \\
\text{s.t.} \quad & \mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \leq \max\left(\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}, \mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}}\right) && \text{(7.34b)} \\
& (\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}) \in \left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right) + \mathcal{N}(\Delta\vartheta, \Delta\varphi, \delta\theta, \delta\phi) && \text{(7.34c)} \\
& (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \in \left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right) + \mathcal{N}(\Delta\vartheta, \Delta\varphi, \delta\theta, \delta\phi) && \text{(7.34d)} \\
& 0 \leq \Delta\vartheta \leq \Delta\theta, \ 0 \leq \Delta\varphi \leq \Delta\phi && \text{(7.34e)}
\end{aligned}$$

The outer maximization aims to find the transmit and receive steering directions $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$ that abide by three constraints. First, the steering directions must satisfy (7.34b), meaning the resulting receive link INR should either be below the desired target $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ or be the minimum INR offered in the surrounding $(\Delta\theta, \Delta\phi)$-neighborhood (i.e., $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}}$). Second, the transmit direction $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ should be within the $(\Delta\vartheta, \Delta\varphi)$-neighborhood surrounding the initial transmit beam selection $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$, as illustrated in Figure 7.3. Third, the receive direction $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$ should be within the $(\Delta\vartheta, \Delta\varphi)$-neighborhood surrounding the initial receive beam selection $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$. To constrain the distance of $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ from $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and of $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$ from $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$, we minimize the neighborhood size $(\Delta\vartheta, \Delta\varphi)$ by minimizing $\Delta\vartheta^2 + \Delta\varphi^2$; other distance measures could also be used. Solving this problem will find the transmit and receive steering directions that meet the INR threshold while minimally deviating from those output by conventional beam alignment. In the next section, we present an algorithm for solving problem (7.34) more efficiently and with fewer INR measurements than an exhaustive search.

Notice that, since $0 \leq \Delta\vartheta \leq \Delta\theta$ and $0 \leq \Delta\varphi \leq \Delta\phi$, we have

$$\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right) + \mathcal{N}(\Delta\vartheta, \Delta\varphi, \delta\theta, \delta\phi) \subseteq \mathcal{T}^{(i^\star)} \tag{7.35}$$

$$\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right) + \mathcal{N}(\Delta\vartheta, \Delta\varphi, \delta\theta, \delta\phi) \subseteq \mathcal{R}^{(j^\star)} \tag{7.36}$$

meaning $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}) \in \mathcal{T}^{(i^\star)}$ and $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \in \mathcal{R}^{(j^\star)}$ for feasible $(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$ and $(\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$, and thus $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \in \mathcal{I}_{\mathrm{rx}}^{(i^\star, j^\star)}$. Hence, solving problem (7.34) simply requires referencing the receive link INR measurements $\mathcal{I}_{\mathrm{rx}}^{(i^\star, j^\star)}$

to find the transmit direction $(\theta_{\mathrm{tx}}^{\star}, \phi_{\mathrm{tx}}^{\star})$ and receive direction $(\theta_{\mathrm{rx}}^{\star}, \phi_{\mathrm{rx}}^{\star})$ that satisfies the INR target while minimizing the deviation from $\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^{\star})}, \phi_{\mathrm{rx}}^{(j^{\star})}\right)$. After solving this problem, our design concludes by setting the beamforming weights as $\mathbf{f}(\theta_{\mathrm{tx}}^{\star}, \phi_{\mathrm{tx}}^{\star})$ and $\mathbf{w}(\theta_{\mathrm{rx}}^{\star}, \phi_{\mathrm{rx}}^{\star})$. When transmitting and receiving with these beams output by STEER, we net $\mathsf{SNR}_{\mathrm{tx}}^{\mathrm{ours}}$ and $\mathsf{SNR}_{\mathrm{rx}}^{\mathrm{ours}}$ analogous to those in (7.16) and (7.17) and a receive link INR of

$$\begin{aligned} \mathsf{INR}_{\mathrm{rx}}^{\mathrm{ours}} &\triangleq \mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}^{\star}, \phi_{\mathrm{tx}}^{\star}, \theta_{\mathrm{rx}}^{\star}, \phi_{\mathrm{rx}}^{\star}) && (7.37) \\ &\leq \mathsf{INR}_{\mathrm{rx}}\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}, \theta_{\mathrm{rx}}^{(j^{\star})}, \phi_{\mathrm{rx}}^{(j^{\star})}\right) = \mathsf{INR}_{\mathrm{rx}}^{\mathrm{nom}}. && (7.38) \end{aligned}$$

The receive link INR achieved by STEER is guaranteed to be no more than that with conventional beam selection. Equality in (7.37) holds when the target $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ is inherently met by the beams output by beam selection. When this target is not inherently met by initial beam selection, STEER will deviate from the initial steering directions $\left(\theta_{\mathrm{tx}}^{(i^{\star})}, \phi_{\mathrm{tx}}^{(i^{\star})}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^{\star})}, \phi_{\mathrm{rx}}^{(j^{\star})}\right)$ only if it leads to lower self-interference.

The steering directions output by STEER relative to those from conventional beam selection are bounded as

$$\begin{aligned} \left|\theta_{\mathrm{tx}}^{\star} - \theta_{\mathrm{tx}}^{(i^{\star})}\right|, \left|\theta_{\mathrm{rx}}^{\star} - \theta_{\mathrm{rx}}^{(j^{\star})}\right| &\leq \Delta\theta, && (7.39) \\ \left|\phi_{\mathrm{tx}}^{\star} - \phi_{\mathrm{tx}}^{(i^{\star})}\right|, \left|\phi_{\mathrm{rx}}^{\star} - \phi_{\mathrm{rx}}^{(j^{\star})}\right| &\leq \Delta\phi. && (7.40) \end{aligned}$$

As such, increasing $(\Delta\theta, \Delta\phi)$ may lead to beams that offer lower SNRs since STEER may use beams that are shifted slightly further away from the transmit and receive devices. However, by throttling $(\Delta\theta, \Delta\phi)$ and courtesy of the INR variability observed over small neighborhoods (on the order of one degree), this potential SNR loss can be constrained and may be greatly outweighed by the reduction in INR, netting it an improved SINR over conventional beam selection. Experimental evaluation of Steer in Section 7.8 confirms this. Note that SNR actually may improve with Steer since it may output beams that are shifted slightly more toward the downlink and uplink devices, compared to those from conventional beam alignment.

Remark 7.2: Choosing a target self-interference level. The optimal choice of $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ (in a sum spectral efficiency sense) cannot be stated analytically since it depends on how much Steer must deviate to meet this target INR, which itself depends on the self-interference channel. In Section 7.8, we find an optimal target heuristically, which shows that $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} \approx -7$ dB is generally near-optimal in maximizing the sum spectral efficiency achieved by Steer. It is difficult to offer commentary on choosing a suitable $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ that generalizes to systems beyond the platform we used to evaluate Steer, since each will have a unique self-interference profile.

Remark 7.3: Design decisions and motivations. We would now like to comment on some of the design decisions and motivations behind Steer. First, we point out that a more attractive beam selection solution would perhaps be one that maximizes the sum spectral efficiency of the transmit and receive links, rather than minimize the receive link INR as is done by Steer. It is practically implausible to reliably maximize sum spectral efficiency since such a problem requires knowing the SNR achieved by a given beam (which would require prohibitive feedback) or downlink and uplink channel knowledge. Minimizing receive link INR by STEER was chosen deliberately, as it relieves the beam selection problem from requiring SNR knowledge. Instead, STEER solely minimizes self-interference through measurements taken at the IAB node and can preserve SNR by reducing its deviation from those output by conventional beam alignment. We also would like to point out that neighborhood size $(\Delta\theta, \Delta\phi)$ and spatial resolution $(\delta\theta, \delta\phi)$ could be uniquely defined for transmit and receive beam selection, rather than having a common one as we have assumed herein. In cases where transmit and receive beam selection are not executed at the same time (e.g., due to a fixed backhaul link), one could simply condition on the beam not being selected (fixing its steering direction) when running STEER. In cases where STEER cannot offer sufficiently low INR to justify full-duplex operation when serving particular transmit and receive users, there is the potential to serve them instead in a half-duplex fashion.

7.7 Efficiently Implementing STEER in Real Systems

Having presented its core components, we now present an algorithm for efficiently executing STEER (by which we mean solving problem (7.34)), along with commentary on key practical considerations. We begin with presentation of Algorithm 7.1. As it was presented in the previous sections, STEER may be executed by first collecting a set of self-interference measurements and then solving problem (7.34), presumably by exhaustive search. While this

Algorithm 7.1 Executing Steer by solving problem (7.34) with a minimal number of measurements.

Input: $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$, $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$, $\mathcal{T}^{(i^\star)}$, $\mathcal{R}^{(j^\star)}$, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$

$\mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}} = \infty$

$\mathcal{V} = \mathcal{T}^{(i^\star)} \times \mathcal{R}^{(j^\star)}$

$\mathcal{D}_{\vartheta} = \left\{\Delta\vartheta = \max\left(\left|\theta_{\mathrm{tx}} - \theta_{\mathrm{tx}}^{(i^\star)}\right|, \left|\theta_{\mathrm{rx}} - \theta_{\mathrm{rx}}^{(j^\star)}\right|\right) : ((\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}), (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})) \in \mathcal{V}\right\}$

$\mathcal{D}_{\varphi} = \left\{\Delta\varphi = \max\left(\left|\phi_{\mathrm{tx}} - \phi_{\mathrm{tx}}^{(i^\star)}\right|, \left|\phi_{\mathrm{rx}} - \phi_{\mathrm{rx}}^{(j^\star)}\right|\right) : ((\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}), (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})) \in \mathcal{V}\right\}$

$\mathcal{D} = \{\Delta\vartheta^2 + \Delta\varphi^2 : \Delta\vartheta \in \mathcal{D}_{\vartheta}, \Delta\varphi \in \mathcal{D}_{\varphi}\}$

$[\sim, \mathcal{J}] = \mathrm{sort}(\mathcal{D}, \mathsf{ascend})$

for $((\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}), (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})) \in [\mathcal{V}]_{\mathcal{J}}$ **do**

Measure (or reference) $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$.

if $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) < \mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}}$ **then**

$\mathsf{INR}_{\mathrm{rx}}^{\mathrm{min}} = \mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$

$(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star) = (\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}})$

$(\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star) = (\theta_{\mathrm{rx}}, \phi_{\mathrm{rx}})$

if $\mathsf{INR}_{\mathrm{rx}}(\theta_{\mathrm{tx}}, \phi_{\mathrm{tx}}, \theta_{\mathrm{rx}}, \phi_{\mathrm{rx}}) \leq \mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ **then**

Break for-loop; target met; no further measurements required.

end if

end if

end for

Output: $(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star)$, $(\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star)$

exhaustive search has fairly low computational complexity, the radio resources consumed to conduct self-interference measurements may be a key bottleneck in practical settings. Motivated by this, we now present an algorithm that executes Steer with a minimal number of measurements.

Algorithm 7.1 begins by sorting all transmit-receive direction pairs $\mathcal{V} = \mathcal{T}^{(i^\star)} \times \mathcal{R}^{(j^\star)}$ based on their deviation from the nominal directions $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$

and $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$ output by beam alignment[3]. The indices of this sorting we denote $\mathcal{J}$, and the sorted set of transmit-receive direction pairs we write as $[\mathcal{V}]_{\mathcal{J}}$, which can be precomputed and is fixed for some neighborhood. In other words, the transmit-receive direction pairs in $[\mathcal{V}]_{\mathcal{J}}$ increase in distance from the nominal direction pair. Starting with the nominal direction pair $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$, the INR of each transmit-receive direction pair in $[\mathcal{V}]_{\mathcal{J}}$ is measured until a transmit-receive pair yields a measured $\mathsf{INR}_{\mathrm{rx}}$ less than the target $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$. Once this target is met, the transmit-receive pair is guaranteed to satisfy all constraints of problem (7.34) and minimizes the distance $\Delta\vartheta^2 + \Delta\varphi^2$, courtesy of our sorting of $\mathcal{V}$. No further measurements are required, having only measured a fraction of the full spatial neighborhood. If no beam pairs meet the threshold, the entire $(\Delta\theta, \Delta\phi)$-neighborhood is measured, with the beam pair offering the lowest INR being selected. Rather than collecting all measurements in $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ and then exhaustively solving problem (7.34), Algorithm 7.1 provides a means to solve problem (7.34) *while* collecting measurements. This reduces its computational overhead since its extremely simple logic can be executed while taking measurements, and more importantly, the overhead consumed to collect INR measurements can be dramatically reduced. We illustrate this reduction in measurement overhead in Section 7.8 using an actual 28 GHz phased array platform.

Remark 7.4: Practical considerations. We highlight that the execution

[3]We slightly abuse convention here and assume sets have ordering for the sake of illustration.

of STEER, along with its associated self-interference measurements, take place solely at the full-duplex IAB node. In fact, the donor and user presumably need not be informed of the beams selected by STEER since only the beams at the IAB node are slightly shifted from those output by conventional beam alignment. This is a practically desirable property of STEER. We would also like to emphasize that a practical implementation of STEER is highly dependent on a number of things. First and foremost, it depends heavily on the time-variability of self-interference, which is currently not well investigated. If self-interference is highly dynamic, measurements will need to be collected more frequently and, with updated measurements, STEER will need to be rerun. However, it may be preferable to first re-measure existing STEER solutions to locate ones that may have become stale, no longer offering sufficiently low INR. Note that, when swapping from the DL-DL operating mode to the UL-UL mode, the self-interference measurements may not be symmetric since the panels presumably swap transmit/receive roles, meaning measurements may need to be collected uniquely for each of the two full-duplexing modes. Characterizing the time-variability and reciprocity of self-interference, along with practical implementations of STEER at mmWave frequencies beyond 28 GHz, are good topics for future work.

Remark 7.5: Precomputing STEER solutions. With collected measurements and known codebooks $\mathcal{F}$ and $\mathcal{W}$, the IAB node can precompute the solution output by STEER for all possible $(i^\star, j^\star)$, keeping a record of $(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star, \theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star)$ for each. Thereafter, the IAB node can directly map initial beam selection indices to the precomputed solution via a lookup table, rather than re-running STEER. This is another practically desirable property of STEER.

$$(i^\star, j^\star) \xrightarrow{\textsf{lookup}} (\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star, \theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star) \tag{7.41}$$

In such a case, the IAB node will need to store $N_{\mathrm{tx}} \cdot N_{\mathrm{rx}} \cdot 4$ values for a given $(\Delta\theta, \Delta\phi)$ and $(\delta\theta, \delta\phi)$. Note that, when the phased arrays are equipped with functions that internally map steering direction to beamforming weights, precomputing STEER solutions only requires storing the steering directions $(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star)$, $(\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star)$ and not the explicit beamforming weights $\mathbf{f}^\star$ and $\mathbf{w}^\star$, reducing storage requirements.

Remark 7.6: Summary of STEER. A summary of our entire beam selection methodology is illustrated in Figure 7.2 and is outlined in Algorithm 7.2. STEER begins by executing conventional beam alignment using codebooks $\mathcal{F}$ and $\mathcal{W}$ to yield initial transmit and receive beam selections that steer toward $\left(\theta_{\mathrm{tx}}^{(i^\star)}, \phi_{\mathrm{tx}}^{(i^\star)}\right)$ and $\left(\theta_{\mathrm{rx}}^{(j^\star)}, \phi_{\mathrm{rx}}^{(j^\star)}\right)$. Then, based on some defined neighborhood size $(\Delta\theta, \Delta\phi)$ and spatial resolution $(\delta\theta, \delta\phi)$, the spatial neighborhoods surrounding these initial transmit and receive beams are constructed as $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$. A receive link INR target $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ is specified by system engineers, likely based on simulation, experimentation, and field trials. Solving problem (7.34) via Algorithm 7.1 will yield the transmit and receive steering directions $(\theta_{\mathrm{tx}}^\star, \phi_{\mathrm{tx}}^\star)$ and $(\theta_{\mathrm{rx}}^\star, \phi_{\mathrm{rx}}^\star)$ that minimally deviate from the initial steering directions while attempting to meet the target INR and minimizing the number of INR measurements collected.

Algorithm 7.2 A summary of our beam selection methodology, STEER.

1. Define transmit and receive coverage regions $\mathcal{A}_{\rm tx}$ and $\mathcal{A}_{\rm rx}$ and corresponding transmit and receive codebooks $\mathcal{F}$ and $\mathcal{W}$ for beam alignment.
2. Conduct beam alignment to yield $\left(\theta_{\rm tx}^{(i^\star)}, \phi_{\rm tx}^{(i^\star)}\right)$ and $\left(\theta_{\rm rx}^{(j^\star)}, \phi_{\rm rx}^{(j^\star)}\right)$.
3. Define the neighborhood spatial resolution $(\delta\theta, \delta\phi)$ and size $(\Delta\theta, \Delta\phi)$.
4. Construct transmit and receive neighborhoods $\mathcal{T}^{(i^\star)}$ and $\mathcal{R}^{(j^\star)}$.
5. Define a desired receive link INR threshold $\mathsf{INR}_{\rm rx}^{\rm tgt}$.
6. Solve problem (7.34) for $(\theta_{\rm tx}^\star, \phi_{\rm tx}^\star)$ and $(\theta_{\rm rx}^\star, \phi_{\rm rx}^\star)$ using Algorithm 7.1, collecting/referencing a minimal number of measurements of self-interference.
7. Set transmit weights as $\mathbf{f}(\theta_{\rm tx}^\star, \phi_{\rm tx}^\star)$ and receive weights as $\mathbf{w}(\theta_{\rm rx}^\star, \phi_{\rm rx}^\star)$.

7.8 Evaluating STEER via Measurement and Simulation

We experimentally evaluate STEER by combining Monte Carlo simulation with INR measurements taken with a 28 GHz phased array platform. A donor and UE are randomly dropped around an IAB node within simulation [104], followed by beam alignment, and then execution of STEER using actual INR measurements. In other words, $\mathsf{INR}_{\rm rx}$ values have been measured, while SNR terms are based on simulation. We consider the case where a full-duplex IAB node transmits and receives using panels on two sides of a sectorized triangular platform, as illustrated in Figure 7.1. The transmit and receive arrays are identical 28 GHz 16×16 half-wavelength UPAs [114]. Using the platform shown previously in Figure 5.3, we measure $\mathsf{INR}_{\rm rx}$ as described in Section 7.5 using a fixed spatial resolution of $(\delta\theta, \delta\phi) = (1^\circ, 1^\circ)$; valuable future work would explore finer resolutions. We take measurements in an anechoic chamber to first explore the impacts of the direct coupling between arrays; investigating the effects of reflections off of realistic environments is also a good topic for dedicated future work. We transmit upconverted Zadoff-Chu sequences with 100 MHz of bandwidth and apply correlation-based processing to reliably estimate self-interference well below the noise floor. Our $\mathsf{INR}_{\mathrm{rx}}$ measurements are typically accurate to within 1 dB, based on validation with high-fidelity test equipment [118] and stepped attenuators [119]. We refer readers to our prior work [44] for more details regarding our measurement methodology, which we also employ herein. We consider identical transmit and receive codebooks comprised of $N_{\mathrm{tx}} = N_{\mathrm{rx}} = 105$ narrow beams that span in azimuth from $-56°$ to $56°$ and in elevation from $-24°$ to $24°$, each with $8°$ spacing. The transmit and receive beams are steered using conjugate beamforming weights (i.e., equal gain/matched filter beamforming), described as $\mathbf{f}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right) = \mathbf{a}_{\mathrm{tx}}\left(\theta_{\mathrm{tx}}^{(i)}, \phi_{\mathrm{tx}}^{(i)}\right)$ and $\mathbf{w}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right) = \mathbf{a}_{\mathrm{rx}}\left(\theta_{\mathrm{rx}}^{(j)}, \phi_{\mathrm{rx}}^{(j)}\right)$, where $\mathbf{a}_{\mathrm{tx}}(\cdot)$ and $\mathbf{a}_{\mathrm{rx}}(\cdot)$ are the transmit and receive array response vectors. The transmit array radiates at an EIRP of 60 dBm and the receive array output has a noise floor of $P_{\mathrm{noise}}^{\mathrm{IAB}} = -68$ dBm over 100 MHz. The transmit and receive beams each have a 3 dB beamwidth of about $7°$. In a Monte Carlo fashion, we randomly drop a donor and UE in the coverage region supplied by our codebooks, from $-60°$ to $60°$ in azimuth and from $-28°$ to $28°$ in elevation. We assume the donor and UE are in LOS of the IAB node for simplicity and to more straightforwardly evaluate STEER against conventional beam selection. We make initial beam selections by choosing transmit and receive beams that maximize their SNRs (e.g., exhaustive beam search), though STEER could be applied atop any beam alignment scheme.

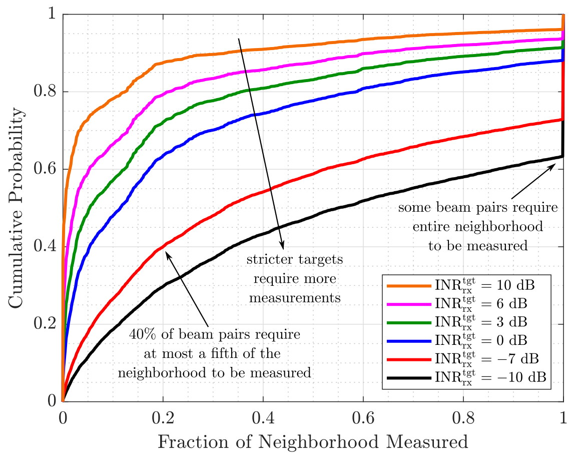

(a) Reducing measurements via Algorithm 7.1.

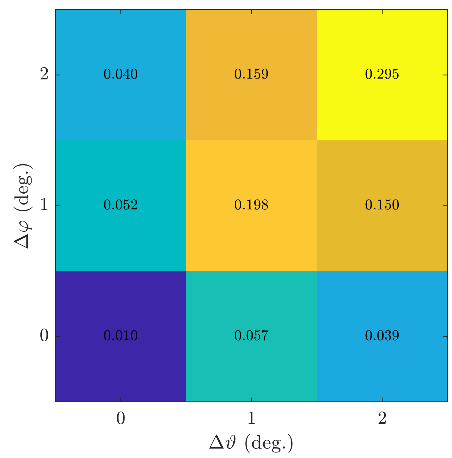

(b) Empirical probability of $(\Delta\vartheta, \Delta\varphi)$.

Figure 7.4: (a) Illustrating the reduction in measurement overhead necessitated by STEER when using Algorithm 7.1. 65% of beam pairs only require 20% of measurements to locate a beam pair that meets a target of $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = 0$ dB. (b) The empirical probability of the resulting $(\Delta\vartheta, \Delta\varphi)$ after executing STEER on the collected INR measurements for $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB.

7.8.1 Reducing Measurement Overhead via Algorithm 7.1

To begin our evaluation of STEER, we first consider Figure 7.4a, which highlights the reduction in measurements needed when using Algorithm 7.1 to solve problem (7.34), compared to measuring the entire spatial neighborhood and then applying exhaustive search. Here, we consider a neighborhood of size $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ with resolution $(\delta\theta, \delta\phi) = (1^\circ, 1^\circ)$. Figure 7.4a shows the empirical CDF of the fraction of $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ measured after running Algorithm 7.1 across all possible $(i^\star, j^\star)$. With $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = 0$ dB, for instance, nearly 65% of all beam pairs $(i^\star, j^\star)$ require at most 20% of the neighborhood to be measured. This highlights the impressive savings Algorithm 7.1 can offer in terms of measurement overhead, a key practical consideration. Around 12% of beam pairs $(i^\star, j^\star)$ require the entire neighborhood $\mathcal{I}_{\mathrm{rx}}^{(i^\star,j^\star)}$ to be measured for $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = 0$ dB. With stricter $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$, more measurements are required in order to locate a transmit-receive beam pair that can meet the target. Notice that some fraction of beam pairs require the entire neighborhood to be measured, which is almost exclusively due to the fact that $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ cannot be met within the neighborhood. In Figure 7.4b, we show the fraction of beam pairs $(i^\star, j^\star)$ that yield each possible $(\Delta\vartheta, \Delta\varphi)$ when $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ and $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB. For example, around 29.5% of beam pairs make use of the full $(2^\circ, 2^\circ)$ tolerance allowed. Around 20% of beam pairs reach $\mathsf{INR}_{\mathrm{rx}} \leq -7$ dB with $(1^\circ, 1^\circ)$ of shifting.

7.8.2 Performance Metrics

We now introduce performance metrics used to evaluate Steer. Recall, the Shannon capacities of the links we denote C_{tx} and C_{rx} are based on the inherent link qualities $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$, respectively. Perhaps more practically meaningful, the capacities of the links after conventional beam alignment we refer to as the *codebook capacities*, which are

$$C_{\mathrm{tx}}^{\mathrm{cb}} = \log_2(1 + \mathsf{SNR}_{\mathrm{tx}}^{\mathrm{nom}}) \leq C_{\mathrm{tx}}, \tag{7.42}$$

$$C_{\mathrm{rx}}^{\mathrm{cb}} = \log_2(1 + \mathsf{SNR}_{\mathrm{rx}}^{\mathrm{nom}}) \leq C_{\mathrm{rx}} \tag{7.43}$$

with these nominal SNRs defined in (7.16)–(7.17) and the sum codebook capacity as $C_{\mathrm{sum}}^{\mathrm{cb}} = C_{\mathrm{tx}}^{\mathrm{cb}} + C_{\mathrm{rx}}^{\mathrm{cb}}$. Using conventional beam alignment, the achievable sum spectral efficiency of the transmit and receive links under equal TDD with fixed power control (i.e., an instantaneous transmit power constraint) are

$$R_{\mathrm{sum}}^{\mathrm{TDD}} = 0.5 \cdot \log_2(1 + \mathsf{SNR}_{\mathrm{tx}}^{\mathrm{nom}}) + 0.5 \cdot \log_2(1 + \mathsf{SNR}_{\mathrm{rx}}^{\mathrm{nom}}). \tag{7.44}$$

Under an average power constraint, power control can be used during equal TDD operation to boost transmit power inversely proportional to the transmit duration. In such a case, the sum spectral efficiency becomes

$$R_{\mathrm{sum}}^{\mathrm{GTDD}} = 0.5 \cdot \log_2(1 + 2 \cdot \mathsf{SNR}_{\mathrm{tx}}^{\mathrm{nom}}) + 0.5 \cdot \log_2(1 + 2 \cdot \mathsf{SNR}_{\mathrm{rx}}^{\mathrm{nom}}). \tag{7.45}$$

Note that this achievable sum spectral efficiency coincides with that of equal FDD. While instantaneous transmit power constraints are more practical, it is still useful for us to compare against an average power constraint to better examine the gains offered by full-duplex. The achievable sum spectral efficiency under Steer we denote as $R_{\mathrm{sum}}^{\mathrm{ours}} = R_{\mathrm{tx}}^{\mathrm{ours}} + R_{\mathrm{rx}}^{\mathrm{ours}}$, which can be computed using (7.10) with the SINRs achieved by Steer. The achievable sum spectral efficiency under conventional beam selection is defined analogously.

To normalize these achievable sum spectral efficiencies to the sum codebook capacity, we translate them to quantities denoted by γ_{sum} by dividing by $C_{\mathrm{tx}}^{\mathrm{cb}} + C_{\mathrm{rx}}^{\mathrm{cb}}$; for instance, the fraction of the codebook capacity achieved when full-duplexing with Steer is

$$\gamma_{\mathrm{sum}}^{\mathrm{ours}} = \frac{R_{\mathrm{sum}}^{\mathrm{ours}}}{C_{\mathrm{sum}}^{\mathrm{cb}}} = \frac{R_{\mathrm{tx}}^{\mathrm{ours}} + R_{\mathrm{rx}}^{\mathrm{ours}}}{C_{\mathrm{tx}}^{\mathrm{cb}} + C_{\mathrm{rx}}^{\mathrm{cb}}}. \tag{7.46}$$

Note that γ_{sum} is typically less than 1 but is not truly bounded since the codebook capacity is not a true upper bound on achievable spectral efficiency.

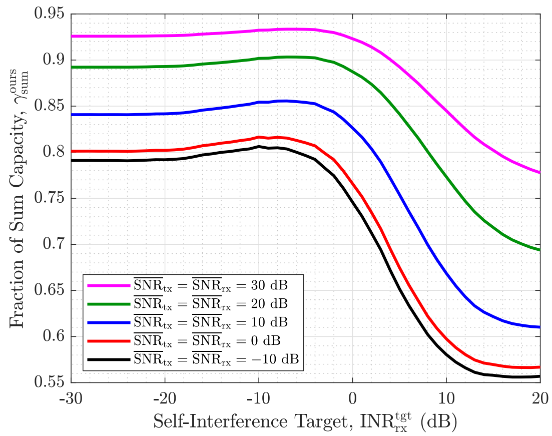

Figure 7.5: The fraction $\gamma_{\mathrm{sum}}^{\mathrm{ours}}$ of the sum capacity $C_{\mathrm{sum}}^{\mathrm{cb}}$ achieved by STEER as a function of the design parameter $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ for various $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. Heuristically, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB proves to be broadly optimal or near-optimal.

Nonetheless, codebook capacity is a useful metric since it provides insight on best-case full-duplex performance with a conventional beam codebook (i.e., in the presence of no cross-link or self-interference).

7.8.3 Choosing the Self-Interference Target, $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$.

In Figure 7.5, we plot the fraction of the sum capacity $\gamma_{\mathrm{sum}}^{\mathrm{ours}}$ achieved by STEER as a function of the design parameter $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ for various $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$, where we have used $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$ and $\mathsf{INR}_{\mathrm{tx}} = 0$ dB for this illustration. Heuristically, we observe that the $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ which maximizes sum spectral efficiency ranges from -10 dB to 0 dB across this broad range of $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. While choosing $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -\infty$ dB is near-optimal (especially at high SNR), the optimal $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ is finite. This is thanks to the fact that choosing a modest $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}}$ can throttle the deviation that STEER makes from the nominal

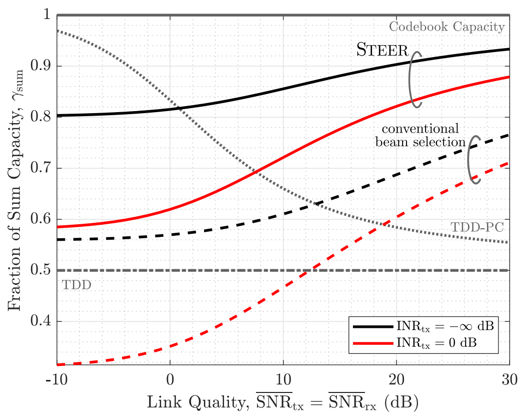

(a) As a function of link quality.

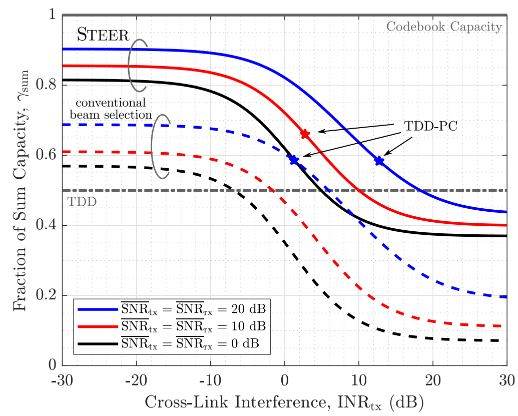

(b) As a function of cross-link interference.

Figure 7.6: (a) The fraction of the sum capacity γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$ for various INR_{tx}. (b) The fraction of the sum capacity γ_{sum} as a function of INR_{tx} for various $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$. $\star$ markers indicate intersections with TDD-PC.

beams, reducing the chance the donor and/or UE sees low beamforming gain. Henceforth, these numerical results use $\mathsf{INR}_{\text{rx}}^{\text{tgt}} = -7$ dB since it is observed to be optimal or near-optimal broadly across $\overline{\mathsf{SNR}}_{\text{tx}}$ and $\overline{\mathsf{SNR}}_{\text{rx}}$.

7.8.4 Full-Duplex with Steer vs. Other Multiplexing Strategies

In Figure 7.6, we compare full-duplexing with Steer to both half-duplexing and full-duplexing with beams from conventional beam selection. We let $(\Delta\theta, \Delta\phi) = (2°, 2°)$ with a spatial resolution of $(\delta\theta, \delta\phi) = (1°, 1°)$ and $\mathsf{INR}_{\text{rx}}^{\text{tgt}} = -7$ dB when running Steer. First, let us examine Figure 7.6a, which shows the fraction of the sum capacity γ_{sum} achieved by various multiplexing strategies as a function of $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$. (For now, we let $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$ for simplicity and examine $\overline{\mathsf{SNR}}_{\text{tx}} \neq \overline{\mathsf{SNR}}_{\text{rx}}$ shortly.) We aim for the codebook capacity $\gamma_{\text{sum}} = 1$ during full-duplex operation, whereas half of this, $\gamma_{\text{sum}} = 0.5$, can be achieved via half-duplexing with equal TDD. With TDD-PC, high γ_{sum} can be had at low SNR thanks to $\log(1+x) \approx x$ at low x, though these gains diminish toward 0.5 as SNR increases. Without cross-link interference (when $\mathsf{INR}_{\text{tx}} = -\infty$ dB; shown in black), STEER vastly outperforms TDD across SNRs, achieving 80% of the codebook capacity at low SNR and over 90% at high SNR. Albeit less practical, TDD-PC can outperform STEER at low SNR, where doubling SNR approximately doubles spectral efficiency, but falls short for $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} \geq 1$ dB. Conventional beam selection also broadly outperforms TDD but only outperforms TDD-PC at $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} \geq 13$ dB. The sizable gap between STEER and conventional beam alignment of around 20% of the codebook capacity is attributed to better self-interference mitigation of STEER. With cross-link interference that is equal to noise (when $\mathsf{INR}_{\text{tx}} = 0$ dB; shown in red), we naturally see a drop in performance of both STEER and conventional beam selection. In this case, STEER still broadly outperforms half-duplexing with TDD, while the same cannot be said about conventional beam selection. Rather, attempting to full-duplex with beams from conventional beam selection falls short of TDD for $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}} \leq 12$ dB. This is due to higher self-interference with conventional beams, which plagues the receive link, and the presence of cross-link interference, which plagues the transmit link. STEER can reduce self-interference to levels that make full-duplex operation worthwhile, even in the presence of cross-link interference.

In Figure 7.6b, we plot the fraction of the sum capacity γ_{sum} achieved by

STEER as a function of cross-link interference $\mathsf{INR}_{\mathrm{tx}}$ for various $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}}$. The starred markers indicate the intersection with performance of TDD-PC. At low cross-link interference, full-duplexing with conventional beam selection outperforms half-duplexing with TDD across SNRs. With STEER, significantly higher spectral efficiencies are obtained largely thanks to its ability to better mitigate self-interference on the receive link. As cross-link interference increases, conventional beam selection degrades as its marginal full-duplexing gains are negated, eventually falling below TDD. At $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} = 10$ dB, for instance, conventional beam selection can only tolerate $\mathsf{INR}_{\mathrm{tx}} \leq -2$ dB, whereas STEER can tolerate $\mathsf{INR}_{\mathrm{tx}} \leq 10$ dB—a gain of about 12 dB in robustness to cross-link interference. Here, the rightward and upward shift of STEER compared to conventional beam selection captures its increased robustness to cross-link interference and its improved sum spectral efficiency. Nonetheless, these results emphasize that justifying full-duplex operation with STEER over half-duplexing strategies depends on cross-link interference levels and the inherent transmit and receive link qualities. This motivates the need to study and measure practical cross-link interference levels and routes to mitigate it as needed, potentially via user selection, which may also be used to meet SNR requirements.

7.8.5 Impact of Neighborhood Size on STEER's Performance

Having examined the performance of STEER versus other multiplexing strategies for a fixed neighborhood size $(\Delta\theta, \Delta\phi) = (2^\circ, 2^\circ)$, we now inspect its

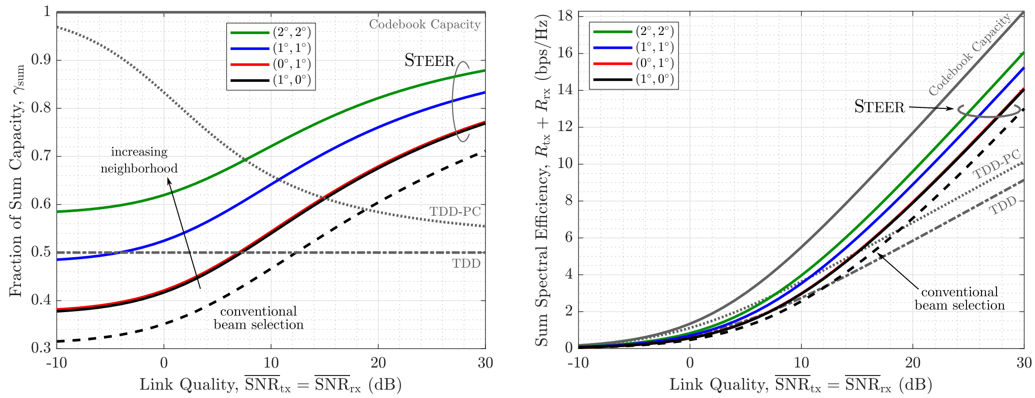

(a) Fraction of the sum capacity, γ_{sum}.

(b) Sum spectral efficiency, $R_{\text{tx}} + R_{\text{rx}}$.

Figure 7.7: (a) The fraction of the sum capacity γ_{sum} as a function of $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$ for various $(\Delta\theta, \Delta\phi)$, where cross-link interference $\mathsf{INR}_{\text{tx}} = 0$ dB. (b) The unnormalized counterpart of (a). Enlarging $(\Delta\theta, \Delta\phi)$ offers noteworthy spectral efficiency gains, especially at high SNR.

performance for various neighborhood sizes while fixing the spatial resolution to $(\delta\theta, \delta\phi) = (1^\circ, 1^\circ)$. A larger neighborhood size $(\Delta\theta, \Delta\phi)$ improves STEER's ability to reduce self-interference since it widens the search space, though this comes at the cost of additional measurement overhead and the potential for gaps in coverage (i.e., reduced SNR). In Figure 7.7a, we plot the fraction of the sum capacity γ_{sum} achieved by STEER for various $(\Delta\theta, \Delta\phi)$ as a function of $\overline{\mathsf{SNR}}_{\text{tx}} = \overline{\mathsf{SNR}}_{\text{rx}}$, where $\mathsf{INR}_{\text{tx}} = 0$ dB. The dashed line shows the performance with conventional beam selection. In Figure 7.7b, we plot the unnormalized counterpart of Figure 7.7a (i.e., absolute sum spectral efficiency $R_{\text{tx}} + R_{\text{rx}}$).

Full-duplexing with conventional beam selection offers gains over TDD only beyond SNRs of 12 dB, which are fairly modest until high SNR. By allowing STEER to shift beams by at most 1° in azimuth and elevation (i.e.,

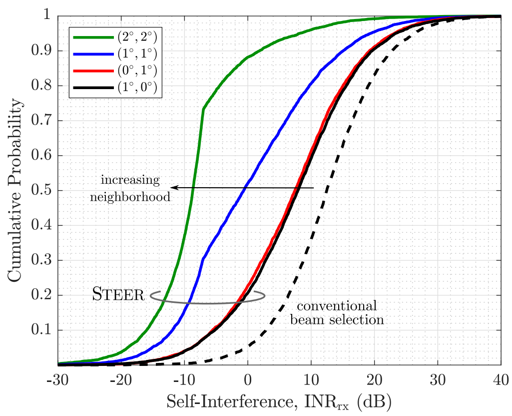

(a) Reducing self-interference.

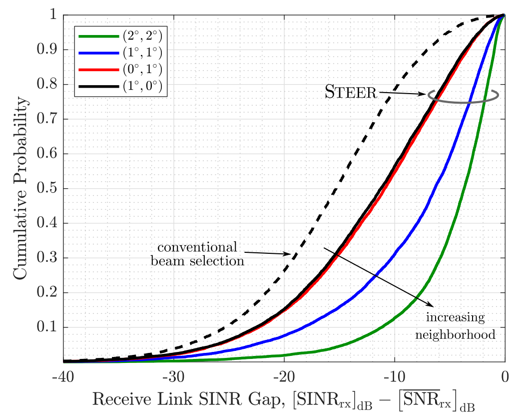

(b) Improving uplink SINR.

Figure 7.8: (a) The CDF of $\mathsf{INR}_{\mathrm{rx}}$ for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. (b) The CDF of the gap between $\mathsf{SINR}_{\mathrm{rx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. STEER reliably reduces $\mathsf{INR}_{\mathrm{rx}}$, as evident in (a), while maintaining high beamforming gain, shifting $\mathsf{SINR}_{\mathrm{rx}}$ closer to $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ (its upper bound), as shown in (b).

$(\Delta\theta, \Delta\phi) = (1^\circ, 1^\circ)$), it can choose beams that greatly reduce self-interference to levels such that full-duplex operation matches or significantly outperforms TDD. Notice, with $(1^\circ, 1^\circ)$, STEER only requires $\overline{\mathsf{SNR}}_{\mathrm{tx}} = \overline{\mathsf{SNR}}_{\mathrm{rx}} \geq -5$ dB to justify full-duplex operation over TDD. Compared to full-duplexing with conventional beam alignment, this is an SNR gain of over 15 dB. It is important to realize that this $(1^\circ, 1^\circ)$-neighborhood lives well within the 7° beamwidth of our beams. A gain of about 0.17 in γ_{sum} is observed with $(1^\circ, 1^\circ)$ and this jumps to around 0.27 with $(2^\circ, 2^\circ)$.

In Figure 7.8a, we plot the CDF of receive link INR offered by STEER for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. Under conventional beam selection, the INR distribution undesirably lay largely above 0 dB, where self-interference is stronger than noise. Notice that shifting the transmit and receive beams by at most 1° in azimuth and elevation, the $\mathsf{INR}_{\mathrm{rx}}$ distribution shifts notably leftward—by about 13 dB in median. Half of all possible transmit-receive user pairs enjoy an $\mathsf{INR}_{\mathrm{rx}} \leq 0$ dB with STEER when $(\Delta\theta, \Delta\phi) = (1^\circ, 1^\circ)$. With $(2^\circ, 2^\circ)$, almost 90% of user pairs enjoy $\mathsf{INR}_{\mathrm{rx}} \leq 0$ dB. As a result of using $\mathsf{INR}_{\mathrm{rx}}^{\mathrm{tgt}} = -7$ dB, we see a sharp bend around $\mathsf{INR}_{\mathrm{rx}} = -7$ dB since STEER is not incentivized to reduce the INR below such. Note that there exist select user pairs that require further deviation beyond $(2^\circ, 2^\circ)$ in order to deliver low INR. Around 4% of user pairs see $\mathsf{INR}_{\mathrm{rx}} \geq 10$ dB even with searching for beams across a $(2^\circ, 2^\circ)$-neighborhood. Enlarging the neighborhood $(\Delta\theta, \Delta\phi)$ or using a finer spatial resolution $(\delta\theta, \delta\phi)$ may facilitate full-duplexing these user pairs or perhaps they are better off served in a half-duplex fashion.

In Figure 7.8b, we plot the CDF of the difference in $[\mathsf{SINR}_{\mathrm{rx}}]_{\mathrm{dB}}$ and $\left[\overline{\mathsf{SNR}}_{\mathrm{rx}}\right]_{\mathrm{dB}}$ (its upper bound) of STEER for various neighborhood sizes $(\Delta\theta, \Delta\phi)$. This difference is useful in capturing two artifacts of STEER: (i) its reduction in $\mathsf{INR}_{\mathrm{rx}}$ and (ii) its effectiveness in receive beamforming. Recall that $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ is the maximum achievable SNR on the receive link and is only achieved by beamforming directly toward the receive device; a conventional codebook would only achieve this SNR if the receive device was precisely in the direction of one of its receive beams. With conventional beam alignment, $\mathsf{SINR}_{\mathrm{rx}}$ is typically well over 10 dB short of $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ and is not unlikely to fall over 20 dB short. When STEER is supplied a $(1^\circ, 1^\circ)$-neighborhood, the distribution greatly shifts rightward. Around 40% of the time, STEER delivers an $\mathsf{SINR}_{\mathrm{rx}}$ that is within 5 dB of $\overline{\mathsf{SNR}}_{\mathrm{rx}}$. This is thanks to STEER's ability to reduce $\mathsf{INR}_{\mathrm{rx}}$ and simultaneously deliver high beamforming gain (i.e., high $\mathsf{SNR}_{\mathrm{rx}}$). With $(2^\circ, 2^\circ)$, even better performance is delivered by STEER: around 70% of the time $\mathsf{SINR}_{\mathrm{rx}}$ is within 6 dB and around 40% of the time, within 3 dB. As highlighted before, a tail exists even with a $(2^\circ, 2^\circ)$-neighborhood since some user pairs simply cannot be offered low $\mathsf{INR}_{\mathrm{rx}}$ without further deviation (or potentially finer spatial resolution)—a characteristic of the self-interference channel. Due to space constraints, we have omitted an examination of the transmit link. Similar conclusions are drawn: SNR loss is throttled by neighborhood size, as was the motivation behind the design of STEER. It would be valuable future work to investigate how the gains of STEER saturate as the neighborhood $(\Delta\theta, \Delta\phi)$ is widened beyond $(2^\circ, 2^\circ)$ or as the resolution $(\delta\theta, \delta\phi)$ is reduced below $(1^\circ, 1^\circ)$.

7.8.6 For Disparate Transmit and Receive Links

For our final set of results, we investigate how STEER performs for $\overline{\mathsf{SNR}}_{\mathrm{tx}} \neq \overline{\mathsf{SNR}}_{\mathrm{rx}}$. This is especially important in IAB applications since network-side infrastructure, like fiber-connected donor base stations and IAB nodes, will likely have higher transmit powers, more antennas, and lower noise figures compared to user equipment. In Figure 7.9a, as a function of $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$, we show the fraction of the sum capacity γ_{sum} achieved when full-duplexing with beams from conventional beam selection, where $\mathsf{INR}_{\mathrm{tx}} = 0$ dB. In Figure 7.9b, we plot that with beams from STEER. Recall that $\gamma_{\mathrm{sum}} = 0.5$ can always be achieved by half-duplexing with TDD, meaning we desire full-duplex

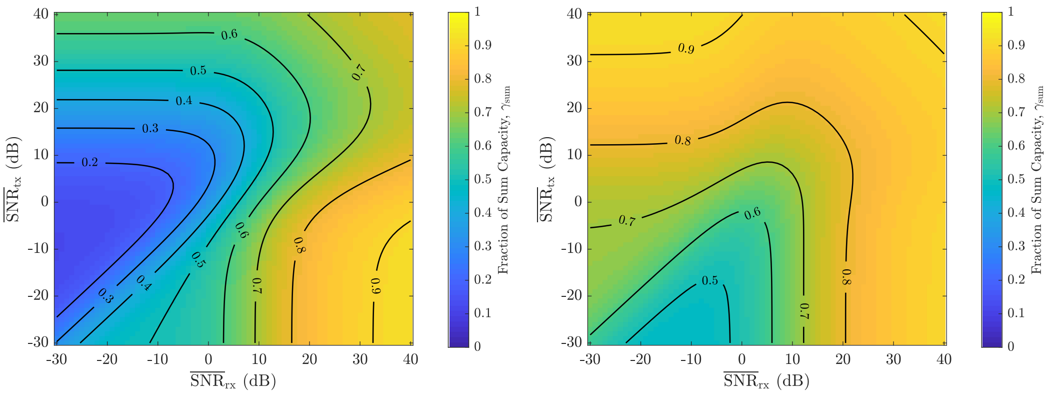

(a) Achieved by conventional beam selection.

(b) Achieved by STEER.

Figure 7.9: The fraction of the sum capacity achieved by (a) conventional beam selection and (b) STEER as a function of $(\overline{\mathsf{SNR}}_{\mathrm{rx}}, \overline{\mathsf{SNR}}_{\mathrm{tx}})$, where $\mathsf{INR}_{\mathrm{tx}} = 0$ dB. With STEER, greater sum spectral efficiency is achieved broadly across SNRs, and the region that nets $\gamma \leq 0.5$ (where half-duplexing is preferred) is greatly reduced with STEER.

operation that exceeds this. With conventional beam selection, high γ_{sum} is seen only at high $\overline{\mathsf{SNR}}_{\mathrm{rx}}$, where self-interference is less impactful due to the diminishing gains of $\log_2(1+x)$. Notice that, when $\overline{\mathsf{SNR}}_{\mathrm{rx}} \leq 0$ dB, conventional beam selection largely yields γ_{sum} that is worse than TDD for practical $\overline{\mathsf{SNR}}_{\mathrm{tx}}$. When STEER is employed, on the other hand, an obvious improvement can be seen, as higher fractions of the sum capacity are achieved broadly across combinations of $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and $\overline{\mathsf{SNR}}_{\mathrm{rx}}$. Only a small region at low $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ and low $\overline{\mathsf{SNR}}_{\mathrm{rx}}$ yields $\gamma_{\mathrm{sum}} \leq 0.5$, in which case the system is better off operating using TDD in terms of sum spectral efficiency. Recall these results are with $\mathsf{INR}_{\mathrm{tx}} = 0$ dB, meaning they would only improve with reduced cross-link interference.

Finally, in Figure 7.10, we compare the $(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}})$-regions where

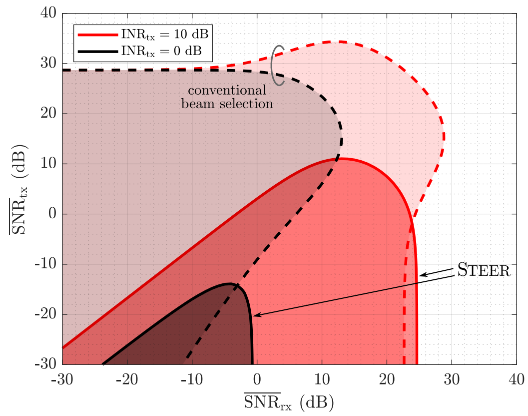

Figure 7.10: For various cross-link interference levels, the $\left(\overline{\mathsf{SNR}}_{\mathrm{tx}}, \overline{\mathsf{SNR}}_{\mathrm{rx}}\right)$-region where $\gamma_{\mathrm{sum}} \leq 0.5$ (achieved by equal TDD). This shaded region shrinks dramatically with STEER, outside of which full-duplexing is justified. STEER generally demands lower SNRs to offer gains over half-duplex and can tolerate higher cross-link interference .

$\gamma_{\mathrm{sum}} \leq 0.5$ for STEER and for conventional beam selection at various cross-link interference levels. Within the shaded regions, it is advantageous to operate using TDD; outside of them, full-duplexing is worthwhile (i.e., $\gamma_{\mathrm{sum}} \geq 0.5$). When $\mathsf{INR}_{\mathrm{tx}} = 0$ dB, the dashed and solid black lines correspond to the $\gamma_{\mathrm{sum}} = 0.5$ contours in Figure 7.9a and Figure 7.9b, respectively. The region bounded by the solid black line is notably smaller than that bounded by the dashed black line, highlighting the dramatic SNR improvement offered by STEER. At $\overline{\mathsf{SNR}}_{\mathrm{rx}} = -10$ dB, for instance, STEER offers an $\overline{\mathsf{SNR}}_{\mathrm{tx}}$ gain of over 30 dB. With higher $\mathsf{INR}_{\mathrm{tx}}$, the regions naturally grow as cross-link interference erodes some of the full-duplexing gains. Still, STEER proves to be more robust to cross-link interference as it generally demands lower SNRs to outperform half-duplex.

7.9 Conclusion

In this work, we present Steer, a measurement-driven beam selection methodology for full-duplex mmWave systems that leverages small shifts of the steering directions of the transmit and receive beams to significantly reduce self-interference and deliver high beamforming gain. Evaluation of Steer through measurements with a 28 GHz phased array platform along with further simulation highlights its ability to reduce self-interference to levels near or below the noise floor, offering noteworthy spectral efficiency gains over half-duplex and full-duplex operation that uses conventional beam selection. Steer can facilitate the deployment of full-duplex mmWave systems to deliver high-throughput, low-latency wireless connectivity, while importantly supporting existing beam alignment schemes in 5G. Valuable future work would include further measuring the time dynamics and small-scale spatial variability of mmWave self-interference, which would ultimately drive design decisions of Steer at deployment. The design of beamforming codebooks that are inherently robust to self-interference and the integration of full-duplex beam selection into mmWave network standards would also be useful contributions. Extending Steer to multi-user/multi-beam systems, along with evaluating Steer using mmWave platforms at other carrier frequencies, in different configurations, and in a variety of settings, would be necessary future work.

Chapter 8

Concluding Remarks

8.1 Summary of Contributions

This dissertation has investigated full-duplex mmWave communication systems from multiple directions, spanning theory and practice. The contributions herein have proposed beamforming-based solutions to allow mmWave communication systems to operate in a full-duplex fashion—simultaneously transmitting and receiving over the same frequency spectrum.

- Contribution I presents a hybrid beamforming design for full-duplex mmWave systems that holistically accounts for a number of key practical considerations. The design uses convex optimization to design beams that reduce self-interference spatially to levels sufficiently low to prevent saturation of receive chain components, such as LNAs and ADCs.
- Contribution II introduces LoneSTAR, the first transmit and receive beamforming codebooks designed specifically for full-duplex mmWave systems. LoneSTAR allows a full-duplex mmWave system to conduct codebook-based beam alignment while simultaneously reducing self-interference through beamforming.
- Contribution V presents Steer, a measurement-driven beam selection

methodology for full-duplex mmWave systems. By strategically—yet slightly—shifting the steering directions of its transmit and receive beams, a full-duplex mmWave transceiver can significantly reduce self-interference while maintaining high beamforming gain in desired directions.

In addition to these three solutions, this dissertation has extensively measured and modeled self-interference in full-duplex mmWave communication systems.

- Contribution III collected nearly 6.5 million measurements of self-interference across a dense spatial profile using 28 GHz phased arrays. Analysis of these measurements sheds light on the levels of real-world self-interference, as well as its high-level spatial trends. Notable small-scale spatial variability was also discovered and characterized; it was found that slight shifts of the transmit and receive steering directions (on the order of one degree) can greatly reduce self-interference.
- Contribution IV proposed a measurement-backed spatial and statistical model of mmWave self-interference, based on the measurements in Contribution III. To construct this model, it was uncovered that the dominant source of self-interference couples from transmit array to the receive array along clusters of rays in a far-field fashion, a stark difference from the idealized near-field model that has been adopted widely thus far in the related literature. Unlike this idealized near-field model, the model proposed in Contribution IV aligns both spatially and statistically with collected measurements.

8.2 Topics for Future Work

This dissertation concludes by detailing directions for future research that can capitalize on the progress made herein.

8.2.1 Extensions of Steer and LoneSTAR

There are some natural extensions to both Steer and LoneSTAR. First, Steer could be generalized to multi-beam/multi-user systems. For instance, suppose a full-duplex mmWave system juggles multiple transmit beams and multiple receive beams. It is unclear how to best extend Steer to such systems, since each transmit beam couples with multiple receive beams. This also applies to the case where more than one sector is operating in a full-duplex fashion at a sectorized mmWave base station. For LoneSTAR, perhaps the most natural extension is to create a higher-dimensional codebook that designs transmit-receive beam *pairs* $\{(\mathbf{f}, \mathbf{w})\}$ that are jointly selected via beam alignment, as opposed to two independent transmit and receive codebooks; this is similar to how Steer operates. This presumably would yield better performance but may be computationally costly or may not be desirable practically since the transmitter and receiver would not be able to operate independently. For both Steer and LoneSTAR, it would be valuable to extend them to be more adaptive, to reduce their overhead, and to apply machine learning.

8.2.2 Shared-Antenna Full-Duplex mmWave Systems

Throughout this dissertation, it was often assumed that separate arrays are used at a full-duplex transceiver for transmission and reception. This is a fitting assumption for inter-sector full-duplex contexts (see Figure 1.5b). For other contexts, such as intra-sector full-duplex, it would be ideal to have a single array at a full-duplex transceiver, where each antenna is shared for transmission and reception. Enabling such shared-antenna full-duplex mmWave systems would be attractive future work. It would be interesting to characterize self-interference in such systems and then explore the feasibility of beamforming as a means to mitigate self-interference.

8.2.3 Real-Time, Adaptive Full-Duplex Solutions

Most full-duplex solutions have been validated in simulation, lab settings, or static/controlled environments. Moreover, most evaluations ignore the overhead associated with configuring a full-duplex solution; in practice, the radio resources consumed to configure a full-duplex solution must not outweigh the gains it offers. It is essential that full-duplex solutions be designed and evaluated with real-time deployments in mind—with strict overhead requirements, with the ability to adapt to dynamic environments, and with minimal form-factors and power consumption. It would be useful to study the time-variability of self-interference in practical applications and environments, especially in mmWave settings. Then, based on these findings, full-duplex solutions can be created that adapt to self-interference with low overhead.

8.2.4 Network-Level Studies of Full-Duplex

To justify the deployment of wireless networks equipped with full-duplex, it is paramount that researchers conduct studies that prove its network-level gains over traditional multiplexing strategies, such as TDD, FDD, and space-division multiple access (SDMA), as well as the recently-proposed cross-division duplexing (XDD) [121]. This has not been examined very extensively for mmWave networks, especially in the context of IAB networks and sectorized base stations.

8.2.5 Capitalizing on Advances in Machine Learning

To supplement existing work [26–29, 50–52], the prospects of using machine learning for full-duplex are still quite open-ended, especially beyond digital SIC. Using machine learning to configure and adaptively update analog SIC filters, for instance, or to configure beamformers that mitigate self-interference in full-duplex mmWave systems are topics that have yet to be fully explored. Perhaps machine learning can be leveraged to create solutions that do not rely on explicit self-interference MIMO channel knowledge (i.e., $\mathbf{H}$) but rather measurements across the channel. In addition, machine learning may be able to overcome transceiver nonidealities when using beamforming to cancel self-interference, such as phased array calibration errors and phase shifter quantization. There are also the prospects of using machine learning to intelligently schedule users and proactively manage cross-link interference within full-duplex networks.

Bibliography

[1] R. K. Mishra, Y. Chen, and I. P. Roberts, "Collision detection in dense Wi-Fi networks using self-interference cancellation," in *Proc. IEEE ICC Wkshp.*, Jun. 2020, pp. 1–6.

[2] J.-S. Jiang and M. A. Ingram, "Spherical-wave model for short-range MIMO," *IEEE Trans. Commun.*, vol. 53, no. 9, pp. 1534–1541, Sep. 2005.

[3] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "Steer: Beam selection for full-duplex millimeter wave communication systems," *IEEE Trans. Commun.*, vol. 70, no. 10, pp. 6902–6917, Oct. 2022.

[4] A. Goldsmith, *Wireless Communications.* Cambridge, UK: Cambridge University Press, 2005.

[5] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication.* Cambridge, UK: Cambridge University Press, 2005.

[6] A. Sabharwal, P. Schniter, D. Guo, D. W. Bliss, S. Rangarajan, and R. Wichman, "In-band full-duplex wireless: Challenges and opportunities," *IEEE JSAC*, vol. 32, no. 9, pp. 1637–1652, Sep. 2014.

[7] K. E. Kolodziej, B. T. Perry, and J. S. Herd, "In-band full-duplex technology: Techniques and systems survey," *IEEE Trans. Microw. Theory Techn.*, vol. 67, no. 7, pp. 3025–3041, Feb. 2019.

[8] S. Chen, M. A. Beach, and J. P. McGeehan, "Division-free duplex for wireless applications," *Electronics Letters*, vol. 34, pp. 147–148, Jan. 1998.

[9] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *ACM MobiCom*, Sep. 2010, pp. 1–12.

[10] D. Bharadia, E. McMilin, and S. Katti, "Full duplex radios," in *ACM SIGCOMM Conf.*, Aug. 2013, pp. 375–386.

[11] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *ACM MobiCom*, Sep. 2011, pp. 301–312.

[12] D. Bharadia and S. Katti, "Full duplex MIMO radios," in *Proc. USENIX NSDI*, Apr. 2014, pp. 359–372.

[13] S. Hong, J. Brand, J. I. Choi, M. Jain, J. Mehlman, S. Katti, and P. Levis, "Applications of self-interference cancellation in 5G and beyond," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 114–121, Feb. 2014.

[14] H. Alves, T. Riihonen, and H. A. Suraweera, *Full-Duplex Communications for Future Wireless Networks.* Singapore: Springer, 2020.

[15] M. Chung, M. S. Sim, J. Kim, D. K. Kim, and C.-B. Chae, "Prototyping real-time full duplex radios," *IEEE Commun. Mag.*, vol. 53, no. 9, pp. 56–63, Sep. 2015.

[16] D. Korpi, L. Anttila, V. Syrjälä, and M. Valkama, "Widely linear digital self-interference cancellation in direct-conversion full-duplex transceiver," *IEEE JSAC*, vol. 32, no. 9, pp. 1674–1687, Sep. 2014.

[17] D. Korpi, T. Riihonen, V. Syrjälä, L. Anttila, M. Valkama, and R. Wichman, "Full-duplex transceiver system calculations: Analysis of ADC and linearity challenges," *IEEE Trans. Wireless Commun.*, vol. 13, no. 7, pp. 3821–3836, Jul. 2014.

[18] D. Korpi, "Full-duplex wireless: Self-interference modeling, digital cancellation, and system studies," Ph.D. dissertation, Tampere University of Technology, Dec. 2017.

[19] Z. Xiao, P. Xia, and X. Xia, "Full-duplex millimeter-wave communication," *IEEE Wireless Commun.*, vol. 24, no. 6, pp. 136–143, Dec. 2017.

[20] I. P. Roberts, J. G. Andrews, H. B. Jain, and S. Vishwanath, "Millimeter-wave full duplex radios: New challenges and techniques," *IEEE Wireless Commun.*, vol. 28, no. 1, pp. 36–43, Feb. 2021.

[21] G. F. G. A. H. A. S. C.-B. C. Besma Smida, Ashutosh Sabharwal, "Full-duplex wireless for 6G: Progress brings new opportunities

and challenges," Apr. 2023. [Online]. Available: https://arxiv.org/abs/2304.08789

[22] C. B. Barneto, S. D. Liyanaarachchi, M. Heino, T. Riihonen, and M. Valkama, "Full duplex radio/radar technology: The enabler for advanced joint communication and sensing," *IEEE Wireless Commun.*, vol. 28, no. 1, pp. 82–88, Feb. 2021.

[23] Y. Liao, T. Wang, L. Song, and Z. Han, "Listen-and-talk: Full-duplex cognitive radio networks," in *Proc. IEEE GLOBECOM*, Dec. 2014, pp. 3068–3073.

[24] Y. Liao, L. Song, Z. Han, and Y. Li, "Full duplex cognitive radio: a new design paradigm for enhancing spectrum usage," *IEEE Commun. Mag.*, vol. 53, no. 5, pp. 138–145, May 2015.

[25] Y. Liao, K. Bian, L. Song, and Z. Han, "Full-duplex MAC protocol design and analysis," *IEEE Commun. Lett.*, vol. 19, no. 7, pp. 1185–1188, Jul. 2015.

[26] Y. Kurzo, A. Burg, and A. Balatsoukas-Stimming, "Design and implementation of a neural network aided self-interference cancellation scheme for full-duplex radios," in *Proc. Asilomar Conf. Signals, Sys., and Comput.*, Oct. 2018, pp. 589–593.

[27] Y. Kurzo, A. T. Kristensen, A. Burg, and A. Balatsoukas-Stimming, "Hardware implementation of neural self-interference cancellation," *IEEE*

Trans. Emerg. Sel. Topics Circuits Syst., vol. 10, no. 2, pp. 204–216, Jun. 2020.

[28] A. Balatsoukas-Stimming, "Joint detection and self-interference cancellation in full-duplex systems using machine learning," in *Proc. Asilomar Conf. Signals, Sys., and Comput.*, Oct. 2021, pp. 989–992.

[29] A. Balatsoukas-Stimming, "Non-linear digital self-interference cancellation for in-band full-duplex radios using neural networks," in *Proc. IEEE SPAWC*, Jun. 2018, pp. 1–5.

[30] L. Song, Y. Liao, K. Bian, L. Song, and Z. Han, "Cross-layer protocol design for CSMA/CD in full-duplex WiFi networks," *IEEE Commun. Lett.*, vol. 20, no. 4, pp. 792–795, Apr. 2016.

[31] K. M. Thilina, H. Tabassum, E. Hossain, and D. I. Kim, "Medium access control design for full duplex wireless systems: challenges and approaches," *IEEE Commun. Mag.*, vol. 53, no. 5, pp. 112–120, May 2015.

[32] M. Gupta, I. P. Roberts, and J. G. Andrews, "System-level analysis of full-duplex self-backhauled millimeter wave networks," *IEEE Trans. Wireless Commun.*, vol. 22, no. 2, pp. 1130–1144, Feb. 2022.

[33] S. Haddad, A. Özgür, and E. Telatar, "Can full-duplex more than double the capacity of wireless networks?" in *Proc. IEEE ISIT*, Jun. 2017, pp. 963–967.

[34] Z. Xiao and Y. Zeng, “Waveform design and performance analysis for full-duplex integrated sensing and communication,” *IEEE JSAC*, vol. 40, no. 6, pp. 1823–1837, Mar. 2022.

[35] X. Wang, Z. Fei, J. A. Zhang, and J. Huang, “Sensing-assisted secure uplink communications with full-duplex base station,” *IEEE Commun. Lett.*, vol. 26, no. 2, pp. 249–253, Dec. 2022.

[36] G. Zheng, I. Krikidis, J. Li, A. P. Petropulu, and B. Ottersten, “Improving physical layer secrecy using full-duplex jamming receivers,” *IEEE Trans. Signal Process.*, vol. 61, no. 20, pp. 4962–4974, Oct. 2013.

[37] T. Riihonen, D. Korpi, O. Rantula, H. Rantanen, T. Saarelainen, and M. Valkama, “Inband full-duplex radio transceivers: A paradigm shift in tactical communications and electronic warfare?” *IEEE Commun. Mag.*, vol. 55, no. 10, pp. 30–36, Oct. 2017.

[38] I. P. Roberts and H. A. Suraweera, “Full-duplex transceivers for next-generation wireless communication systems,” in *Fundamentals of 6G Communications and Networking*, Y. Liu, J. Zhang, X. Lin, and J. Kim, Eds. Berlin: Springer Nature, 2023, (to be published). [Online]. Available: https://arxiv.org/abs/2210.08094

[39] R. W. Heath, N. González-Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, “An overview of signal processing techniques for millimeter wave MIMO systems,” *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 3, pp. 436–453, Apr. 2016.

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

[41] I. P. Roberts, J. G. Andrews, and S. Vishwanath, "Hybrid beamforming for millimeter wave full-duplex under limited receive dynamic range," *IEEE Trans. Wireless Commun.*, vol. 20, no. 12, pp. 7758–7772, Dec. 2021.

[42] I. P. Roberts, S. Vishwanath, and J. G. Andrews, "LoneSTAR: Analog beamforming codebooks for full-duplex millimeter wave systems," *IEEE Trans. Wireless Commun.*, Jan. 2023, (early access).

[43] I. P. Roberts, H. B. Jain, S. Vishwanath, and J. G. Andrews, "Millimeter wave analog beamforming codebooks robust to self-interference," in *Proc. IEEE GLOBECOM*, Dec. 2021, pp. 1–6.

[44] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "Beamformed self-interference measurements at 28 GHz: Spatial insights and angular spread," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9744–9760, Nov. 2022.

[45] I. P. Roberts, A. Chopra, T. Novlan, S. Vishwanath, and J. G. Andrews, "Spatial and statistical modeling of multi-panel millimeter wave self-interference," 2023, (to appear).

[46] J. Suarez, K. Kravtsov, and P. R. Prucnal, "Incoherent method of optical interference cancellation for radio-frequency communications," *IEEE J. Quantum Electron.*, vol. 45, no. 4, pp. 402–408, Mar. 2009.

[47] X. Han, B. Huo, Y. Shao, and M. Zhao, "Optical RF self-interference cancellation by using an integrated dual-parallel MZM," *IEEE Photon. J.*, vol. 9, no. 2, pp. 1–8, Apr. 2017.

[48] X. Su, X. Han, S. Fu, S. Wang, C. Li, Q. Tan, G. Zhu, C. Wang, Z. Wu, Y. Gu, and M. Zhao, "Optical multipath RF self-interference cancellation based on phase modulation for full-duplex communication," *IEEE Photon. J.*, vol. 12, no. 4, pp. 1–14, Jun. 2020.

[49] M. S. Sim, M. Chung, D. Kim, J. Chung, D. K. Kim, and C.-B. Chae, "Nonlinear self-interference cancellation for full-duplex radios: From link-level and system-level performance perspectives," *IEEE Commun. Mag.*, vol. 55, no. 9, pp. 158–167, Sep. 2017.

[50] H. Guo, J. Xu, S. Zhu, and S. Wu, "Realtime software defined self-interference cancellation based on machine learning for in-band full duplex wireless communications," Mar. 2018, pp. 779–783.

[51] H. Guo, S. Wu, H. Wang, and M. Daneshmand, "DSIC: Deep learning based self-interference cancellation for in-band full duplex wireless," in *Proc. IEEE GLOBECOM*, Dec. 2019, pp. 1–6.

[52] Y. Chen, R. K. Mishra, D. Schwartz, and S. Vishwanath, "MIMO full duplex radios with deep learning," in *Proc. IEEE ICC Wkshp.*, Jun. 2020, pp. 1–6.

[53] L. Laughlin, M. A. Beach, K. A. Morris, and J. L. Haine, "Optimum single antenna full duplex using hybrid junctions," *IEEE JSAC*, vol. 32, no. 9, pp. 1653–1661, Jun. 2014.

[54] T. Dinc, M. Tymchenko, A. Nagulu, D. Sounas, A. Alu, and H. Krishnaswamy, "Synchronized conductivity modulation to realize broadband lossless magnetic-free non-reciprocity," *Nature Commun.*, vol. 8, no. 11, pp. 1–9, Oct. 2017.

[55] T. Dinc, A. Chakrabarti, and H. Krishnaswamy, "A 60 GHz CMOS full-duplex transceiver and link with polarization-based antenna and RF cancellation," *IEEE J. Solid-State Circuits*, vol. 51, no. 5, pp. 1125–1140, May 2016.

[56] A. Alkhateeb, O. E. Ayach, G. Leus, and R. W. Heath, "Channel estimation and hybrid precoding for millimeter wave cellular systems," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 831–846, Oct. 2014.

[57] A. Natarajan, S. K. Reynolds, M.-D. Tsai, S. T. Nicolson, J.-H. C. Zhan, D. G. Kam, D. Liu, Y.-L. O. Huang, A. Valdes-Garcia, and B. A. Floyd, "A fully-integrated 16-element phased-array receiver in SiGe BiCMOS for 60-GHz communications," *IEEE J. Solid-State Circuits*, vol. 46, no. 5, pp. 1059–1075, May 2011.

[58] B. Sadhu, Y. Tousi, J. Hallin, S. Sahl, S. K. Reynolds, Ö. Renström, K. Sjögren, O. Haapalahti, N. Mazor, B. Bokinge, G. Weibull, H. Bengtsson, A. Carlinger, E. Westesson, J.-E. Thillberg, L. Rexberg, M. Yeck, X. Gu, M. Ferriss, D. Liu, D. Friedman, and A. Valdes-Garcia, "A 28-GHz 32-element TRX phased-array IC with concurrent dual-polarized operation and orthogonal phase and gain control for 5G communications," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3373–3391, Dec. 2017.

[59] "5G Products | Anokiwave," 2021. [Online]. Available: https://www.anokiwave.com/5g/

[60] Y. Heng, J. G. Andrews, J. Mo, V. Va, A. Ali, B. L. Ng, and J. C. Zhang, "Six key challenges for beam management in 5.5G and 6G systems," *IEEE Commun. Mag.*, vol. 59, no. 7, pp. 74–79, Jul. 2021.

[61] J. Wang, Z. Lan, C. woo Pyo, T. Baykas, C. sean Sum, M. Rahman, J. Gao, R. Funada, F. Kojima, H. Harada, and S. Kato, "Beam codebook based beamforming protocol for multi-Gbps millimeter-wave WPAN systems," *IEEE JSAC*, vol. 27, no. 8, pp. 1390–1399, Oct. 2009.

[62] S. Huberman and T. Le-Ngoc, "MIMO full-duplex precoding: A joint beamforming and self-interference cancellation structure," *IEEE Trans. Wireless Commun.*, vol. 14, no. 4, pp. 2205–2217, Apr. 2015.

[63] T. Riihonen, S. Werner, and R. Wichman, "Mitigation of loopback self-interference in full-duplex MIMO relays," *IEEE Trans. Signal Process.*,

vol. 59, no. 12, pp. 5983–5993, Aug. 2011.

[64] H. A. Suraweera, I. Krikidis, G. Zheng, C. Yuen, and P. J. Smith, "Low-complexity end-to-end performance optimization in MIMO full-duplex relay systems," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, pp. 913–927, Jan. 2014.

[65] H. Q. Ngo, H. A. Suraweera, M. Matthaiou, and E. G. Larsson, "Multipair full-duplex relaying with massive arrays and linear processing," *IEEE JSAC*, vol. 32, no. 9, pp. 1721–1737, Jun. 2014.

[66] B. P. Day, A. R. Margetts, D. W. Bliss, and P. Schniter, "Full-duplex MIMO relaying: Achievable rates under limited dynamic range," *IEEE JSAC*, vol. 30, no. 8, pp. 1541–1553, Sep 2012.

[67] E. Everett, C. Shepard, L. Zhong, and A. Sabharwal, "SoftNull: Many-antenna full-duplex wireless via digital beamforming," *IEEE Trans. Wireless Commun.*, vol. 15, no. 12, pp. 8077–8092, Dec. 2016.

[68] X. Liu, Z. Xiao, L. Bai, J. Choi, P. Xia, and X.-G. Xia, "Beamforming based full-duplex for millimeter-wave communication," *Sensors*, vol. 16, no. 7, p. 1130, Jul. 2016.

[69] R. López-Valcarce and N. González-Prelcic, "Beamformer design for full-duplex amplify-and-forward millimeter wave relays," in *Proc. ISWCS*, Aug. 2019, pp. 86–90.

[70] R. López-Valcarce and N. González-Prelcic, "Analog beamforming for full-duplex millimeter wave communication," in *Proc. ISWCS*, Aug. 2019, pp. 687–691.

[71] J. Palacios, J. Rodríguez-Fernández, and N. González-Prelcic, "Hybrid precoding and combining for full-duplex millimeter wave communication," in *Proc. IEEE GLOBECOM*, Dec. 2019, pp. 1–6.

[72] K. Satyanarayana, M. El-Hajjar, P. Kuo, A. Mourad, and L. Hanzo, "Hybrid beamforming design for full-duplex millimeter wave communication," *IEEE Trans. Veh. Technol.*, vol. 68, no. 2, pp. 1394–1404, Feb. 2019.

[73] L. Zhu, J. Zhang, Z. Xiao, X. Cao, X. Xia, and R. Schober, "Millimeter-wave full-duplex UAV relay: Joint positioning, beamforming, and power control," *IEEE JSAC*, vol. 38, no. 9, pp. 2057–2073, Sep. 2020.

[74] J. M. B. da Silva, A. Sabharwal, G. Fodor, and C. Fischione, "1-bit phase shifters for large-antenna full-duplex mmWave communications," *IEEE Trans. Wireless Commun.*, vol. 19, no. 10, pp. 6916–6931, Oct. 2020.

[75] R. López-Valcarce and M. Martínez-Cotelo, "Full-duplex mmWave MIMO with finite-resolution phase shifters," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 8979–8994, Nov. 2022.

[76] A. Koc and T. Le-Ngoc, "Full-duplex mmWave massive MIMO systems: A joint hybrid precoding/combining and self-interference cancellation

design," vol. 2, pp. 754–774, Mar. 2021.

[77] Y. Cai, Y. Xu, Q. Shi, B. Champagne, and L. Hanzo, "Robust joint hybrid transceiver design for millimeter wave full-duplex MIMO relay systems," *IEEE Trans. Wireless Commun.*, vol. 18, no. 2, pp. 1199–1215, Feb. 2019.

[78] I. P. Roberts, H. B. Jain, and S. Vishwanath, "Equipping millimeter-wave full-duplex with analog self-interference cancellation," in *Proc. IEEE ICC Wkshp.*, Jun. 2020, pp. 1–6.

[79] G. C. Alexandropoulos, M. A. Islam, and B. Smida, "Full duplex hybrid A/D beamforming with reduced complexity multi-tap analog cancellation," in *Proc. IEEE SPAWC*, May 2020, pp. 1–5.

[80] G. C. Alexandropoulos and M. Duarte, "Joint design of multi-tap analog cancellation and digital beamforming for reduced complexity full duplex MIMO systems," in *Proc. IEEE ICC*, May 2017, pp. 1–7.

[81] V. Singh, S. Mondal, A. Gadre, M. Srivastava, J. Paramesh, and S. Kumar, "Millimeter-wave full duplex radios," in *ACM MobiCom*, Apr. 2020.

[82] C. Zhang and X. Luo, "Adaptive digital self-interference cancellation for millimeter-wave full-duplex backhaul systems," *IEEE Access*, vol. 7, pp. 175 542–175 553, Dec. 2019.

[83] G. Y. Suk, S.-M. Kim, J. Kwak, S. Hur, E. Kim, and C.-B. Chae, "Full duplex integrated access and backhaul for 5G NR: Analyses and proto-

type measurements," *IEEE Wireless Commun.*, vol. 29, no. 4, pp. 40–46, Aug. 2022.

[84] A. Sabharwal, P. Schniter, D. Guo, D. W. Bliss, S. Rangarajan, and R. Wichman, "In-band full-duplex wireless: Challenges and opportunities," *IEEE JSAC*, vol. 32, no. 9, pp. 1637–1652, Sep. 2014.

[85] I. P. Roberts and S. Vishwanath, "Beamforming cancellation design for millimeter-wave full-duplex," in *Proc. IEEE GLOBECOM*, Dec. 2019, pp. 1–6.

[86] I. P. Roberts, H. B. Jain, and S. Vishwanath, "Frequency-selective beamforming cancellation design for millimeter-wave full-duplex," in *Proc. IEEE ICC*, Jun. 2020, pp. 1–6.

[87] Q. Gu, *RF System Design of Transceivers for Wireless Communications.* New York, NY, USA: Springer, 2005.

[88] R. W. Heath Jr. and A. Lozano, *Foundations of MIMO Communication.* Cambridge, UK: Cambridge University Press, 2018.

[89] M. Grant and S. Boyd, "CVX: Software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[90] R. J. Mailloux, *Phased Array Antenna Handbook.* Norwood, MA, USA: Artech House, 2005.

[91] T. Taylor, "Design of line-source antennas for narrow beamwidth and low side lobes," *Transactions of the IRE Professional Group on Antennas and Propagation*, vol. 3, no. 1, pp. 16–28, Jan. 1955.

[92] W. G. Carrara, R. S. Goodman, and R. M. Majewski, *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms.* Norwood, MA, USA: Artech House, 1995.

[93] S. Huberman and T. Le-Ngoc, "MIMO full-duplex precoding: A joint beamforming and self-interference cancellation structure," *IEEE Trans. Wireless Commun.*, vol. 14, no. 4, pp. 2205–2217, Apr. 2015.

[94] C. Dehos, J. L. González, A. D. Domenico, D. Kténas, and L. Dussopt, "Millimeter-wave access and backhauling: the solution to the exponential data traffic increase in 5G mobile communications systems?" *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 88–95, Sep. 2014.

[95] 3GPP, "3GPP TS 38.174: New WID on IAB enhancements," 2021. [Online]. Available: https://www.3gpp.org/dynareport/38174.htm

[96] Z. Wei, X. Zhu, S. Sun, Y. Huang, L. Dong, and Y. Jiang, "Full-duplex versus half-duplex amplify-and-forward relaying: Which is more energy efficient in 60-GHz dual-hop indoor wireless systems?" vol. 33, no. 12, pp. 2936–2947, Dec. 2015.

[97] C. Balanis, *Antenna Theory: Analysis and Design.* Hoboken, New Jersey: John Wiley & Sons, 2016.

[98] M. Cui and L. Dai, "Channel estimation for extremely large-scale MIMO: Far-field or near-field?" *IEEE Trans. Wireless Commun.*, vol. 70, no. 4, pp. 2663–2677, Apr. 2022.

[99] S. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge, UK: Cambridge University Press, 2004.

[100] A. B. Gershman, N. D. Sidiropoulos, S. Shahbazpanahi, M. Bengtsson, and B. Ottersten, "Convex optimization-based beamforming," *IEEE Signal Process. Mag.*, vol. 27, no. 3, pp. 62–75, May 2010.

[101] X. Zhang, D. P. Palomar, and B. Ottersten, "Statistically robust design of linear MIMO transceivers," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3678–3689, Aug. 2008.

[102] J. Wang and D. P. Palomar, "Worst-case robust MIMO transmission with imperfect channel knowledge," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 3086–3100, Aug. 2009.

[103] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions.* Boca Raton, FL, USA: Chapman & Hall/CRC, 2000.

[104] I. P. Roberts, "MIMO for MATLAB: A toolbox for simulating MIMO communication systems," Nov. 2021. [Online]. Available: https://arxiv.org/abs/2111.05273

[105] H. Ronkainen, J. Edstam, A. Ericsson, and C. Östberg, “Integrated access and backhaul: A new type of wireless backhaul in 5G,” *Frontiers in Communications and Networks*, vol. 2, Apr. 2021.

[106] S. Rajagopal, R. Taori, and S. Abu-Surra, “Self-interference mitigation for in-band mmWave wireless backhaul,” in *Proc. IEEE CCNC*, Jan. 2014, pp. 551–556.

[107] Y. Kohda, K. Takano, D. Nakano, N. Ohba, T. Yamane, and Y. Katayama, “Single-channel full-duplex mmWave link using phased-array for Ethernet,” in *Proc. IEEE CCNC*, Jan. 2015, pp. 400–405.

[108] B. Lee, J. Lim, C. Lim, B. Kim, and J. Seol, “Reflected self-interference channel measurement for mmWave beamformed full-duplex system,” in *Proc. IEEE GLOBECOM Wkshp.*, Dec. 2015, pp. 1–6.

[109] H. Yang, Y. He, C. Jen, C. Liu, S. Jou, X. Yin, M. Ma, and B. Jiao, “Interference measurement and analysis of full-duplex wireless system in 60 GHz band,” in *Proc. IEEE APCCAS*, Oct. 2016, pp. 273–276.

[110] Y. He, X. Yin, and H. Chen, “Spatiotemporal characterization of self-interference channels for 60-GHz full-duplex communication,” *IEEE Antennas Wireless Propag. Lett.*, vol. 16, pp. 2220–2223, May 2017.

[111] K. Haneda, J. JÃďrvelÃďinen, A. Karttunen, and J. Putkonen, “Self-interference channel measurements for in-band full-duplex street-level

backhaul relays at 70 GHz," in *Proc. IEEE PIMRC*, Sep. 2018, pp. 199–204.

[112] L. Li, K. Josiam, and R. Taori, "Feasibility study on full-duplex wireless millimeter-wave systems," in *Proc. IEEE ICASSP*, May 2014, pp. 2769–2773.

[113] A. Chopra, I. P. Roberts, T. Novlan, and J. G. Andrews, "28 GHz phased array-based self-interference measurements for millimeter wave full-duplex," in *Proc. IEEE WCNC*, Apr. 2022, pp. 2583–2588.

[114] "Anokiwave AWA-0134 5G active antenna innovator kit," 2021. [Online]. Available: https://www.anokiwave.com/products/awa-0134/

[115] Qualcomm, "3GPP RP-193251: New WID on IAB enhancements," Dec. 2019. [Online]. Available: https://portal.3gpp.org/ngppapp/CreateTDoc.aspx?mode=view&contributionUid=RP-193251

[116] "Getting Started Guide USRP-2950/2952/2953/2954/2955 - National Instruments," 2021. [Online]. Available: https://www.ni.com/pdf/manuals/376355c.pdf

[117] "Millimeter wave products | components | 7GHz to 320GHz," 2021. [Online]. Available: https://www.miwv.com

[118] "NRP40SN three path diode power sensor - Rohde & Schwarz," 2021. [Online]. Available: https://www.rohde-schwarz.com/my/product/nrp40sn-options_63490-160769.html

[119] "RSC-Z405 RSC step attenuator - Rohde & Schwarz," 2021. [Online]. Available: https://www.rohde-schwarz.com/us/product/rsc-productstartpage_63493-11395.html

[120] M. Cudak, A. Ghosh, A. Ghosh, and J. G. Andrews, "Integrated access and backhaul: A key enabler for 5G millimeter-wave deployments," *IEEE Commun. Mag.*, vol. 59, no. 4, pp. 88–94, Apr. 2021.

[121] H. Ji, Y. Kim, K. Muhammad, C. Tarver, M. Tonnemacher, T. Kim, J. Oh, B. Yu, G. Xu, and J. Lee, "Extending 5G TDD coverage with XDD: Cross division duplex," *IEEE Access*, vol. 9, pp. 51 380–51 392, Mar. 2021.

Vita

Ian Parker Roberts was raised in Gordonville, Missouri. He attended Jackson R-2 Public Schools from kindergarten through his high school graduation in 2014. He received the Bachelor of Science degree in electrical engineering from Missouri University of Science and Technology in 2018. He received the Master of Science in Engineering degree in electrical and computer engineering from The University of Texas at Austin in 2020. He held internships at AT&T Labs, GenXComm, Amazon, Sandia National Labs, and Dynetics. During the summer of 2022, he was a visiting student at Arizona State University and at Yonsei University. In 2023, he received the Andrea Goldsmith Young Scholars Award from the Communication Theory Technical Committee of the IEEE Communications Society. Also in 2023, he received the Student Leadership Award from the Wireless Networking and Communications Group at The University of Texas at Austin. His graduate studies were supported by the National Science Foundation Graduate Research Fellowship Program.

Permanent address: `ipr@utexas.edu`

This dissertation was typeset with LaTeX† by the author.

† LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TeX Program.