Robust and Efficient Structure-Based Radar Receive Processing

By

Christian Curtis Jones

Submitted to the graduate degree program in Electrical Engineering and Computer Science and the Graduate Faculty of the University of Kansas in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Committee members

Dr. Shannon Blunt, Chairperson

Dr. Patrick McCormick, Co-chair

Dr. Christopher Allen

Dr. James Stiles

Dr. Suzanne Shontz

Dr. Zsolt Talata

Date defended: December 08, 2023

The Dissertation Committee for Christian Curtis Jones certifies
that this is the approved version of the following dissertation :

Robust and Efficient Structure-Based Radar Receive Processing

Dr. Shannon Blunt, Chairperson

Dr. Patrick McCormick, Co-chair

Date approved: December 08, 2023

Abstract

Legacy radar systems largely rely on repeated emission of a linear frequency modulated (LFM) or chirp waveform to ascertain scattering information from an environment. The prevalence of these chirp waveforms largely stems from their simplicity to generate/process and robustness to physical hardware distortion. However, this traditional design philosophy may lack the flexibility and dimensionality needed to address the dynamic "complexification" of the modern radio frequency (RF) environment or achieve current operational requirements where unprecedented degrees of sensitivity, maneuverability, and adaptability are necessary.

Over the last couple of decades analog-to-digital and digital-to-analog technologies have advanced exponentially, resulting in tremendous design degrees of freedom and arbitrary waveform generation (AWG) capabilities that enable sophisticated design of emissions to better suit operational requirements. However, radar transmitters (TX) typically require high-power amplification (HPA) to contend with two-way propagation loss. Thus, transmitter-amenable waveforms are effectively constrained to be both spectrally contained and constant amplitude, resulting in a non-convex NP-hard design problem. While isolating the global optimal can be intractable for even modest time-bandwidth products (TB), locally optimal TX-amenable solutions that are "good enough" are often readily available. However, traditional matched filter (MF) based estimation may no longer satisfy operational requirements with sub-optimal emissions.

Using knowledge of the radar TX-to-receiver (TX-RX) chain, a discrete linear model can be formed to express the relationship between observed measurements and the complex scattering of the environment. This structured representation enables

more sophisticated least-square (LS) and adaptive estimation techniques that can designed to better satisfy operational requirements, improve estimate fidelity, and extend dynamic range. However, receive dimensionality can be enormous as the aggregate degrees of freedom (DoF) for a coherent processing interval (CPI) is multiplicative with respect to bandwidth, pulse duration, and number of pulses. Brute force implementation of model-based estimation techniques that leverage these DoFs may have unwieldy computational burden on even the most cutting-edge signal processing hardware. Additionally, a discrete linear representation is fundamentally an approximation of a dynamic and continuous physical reality where model errors may induce bias, create false detections, or restrict effective dynamic range. As such, these structure-based approaches must be both computationally efficient and robust to reality.

Here discrete receive models and algorithms for pulsed-radar range-Doppler estimation are introduced. Modifications and alternative solutions are then proposed to improve estimate fidelity, reduce computational complexity, and enhance practical robustness.

Acknowledgements

To the students, graduates, and faculty from the Radar Systems & Remote Sensing Lab, I want to express my gratitude to each of you. Your support and guidance have been invaluable to me, and I wouldn't be the person I am today without you pushing and challenging me. Thank you for being amazing friends and mentors.

To my family and partner, thank you for your endless support and love. I look forward to what life will bring us in the years to come.

This work was supported by the Office of Naval Research under

Contract #N00014-23-C-1053

and the Naval Research Laboratory under

Base Program 53-6B72-Z-1-5.

DISTRIBUTION STATEMENT A.

Approved for Public Release, distribution is unlimited.

Contents

1 Introduction **1**

2 Background **6**

2.1 Radar Range Processing Overview 6

2.1.1 Pulse Echo Model 7

2.1.2 Pulse Compression 10

2.1.3 Discretized Receive Model 12

2.2 Least Squares Mismatched Filtering Overview 18

2.2.1 Template-Based Least Squares MMFs 19

2.2.2 Integrated Sidelobe Level LS MMFs 24

2.2.3 Controlling Mismatch Loss via Quadratic Constraints 29

2.3 Adaptive Range Processing Overview 30

2.4 Adaptive Pulse Compression 31

2.4.1 APC Generic Model Uncertainty 38

2.4.2 APC Covariance Beamspoiling 41

2.4.3 APC Unity Gain Linear Constraint 43

2.5 Doppler Processing Overview 45

2.5.1 Single-Pulse Doppler Processing 46

2.5.2 Multi Pulse Doppler Processing 48

2.5.3 Clutter Cancellation 55

2.6 NIMPC 59

2.7 Select Structured Matrices 61

2.7.1 Sparse and Banded Hermitian Matrices 61

2.7.2 Circulant Matrices 63
2.7.3 Toeplitz Matrices 64
2.7.4 Structure of a Toeplitz Matrix Inverse 66
2.7.5 Toeplitz-Block and Block-Toeplitz Matrices 67
2.8 Iterative Linear System Solvers 70

3 Efficient Straddle-Robust Least Squares Mismatched Filtering 72
3.1 Statistical Straddling Compensation for LS MMFs 72
3.1.1 Low-Rank Approximate Straddling Compensation for Efficient ISL Beam-spoiling 78
3.2 Efficient Solvers for LS MMF QPs 79
3.2.1 Iterative Approximate LS MMFs 81
3.3 Closed-Loop Experimental Assessment 82
3.4 Open-Air Experimental Assessment 89

4 Mismatched Complementary on Receive Filtering 96
4.1 MiCRFt Model 97
4.2 Incorporating Range Straddling Robustness 100
4.3 Efficient Direct Solvers for MiCRFt 107
4.4 MiCRFt Doppler Sensitivity and Aliasing 110
4.5 MiCRFt Closed-Loop Experimental Assessment 114
4.6 MiCRFt Open-Air Experimental Assessment 116

5 Alternate Joint Range-Doppler Clutter Cancellation Representations and Efficient Implementations 120
5.1 Efficient Direct Solvers for NIMPC 120
5.2 NIMPC as a Projection 127
5.2.1 Efficient Solvers for Proj-NIMPC 129
5.3 Open-Air Experimental Assessment 131

6 Robust Adaptive Pulse Compression **138**

6.1 Incorporating Range-Straddling Model Uncertainty 138

6.1.1 Filter Beamspoiling 144

6.2 Full APC representation and relationship to snapshot-form 148

6.2.1 Constructing the Structured-Covariance Matrices via diagonal Convolution 151

6.3 Reduced Complexity Full APC Implementations 153

6.3.1 Reduced Complexity Template-Based Approach 154

6.3.2 Outlier-Based Low-Rank Approximate Implementation 156

6.4 Robust and Efficient APC Experimental Evaluation 158

7 Generalized SINR Loss Constrained LS MMFs **166**

7.1 Generalized SINR Loss Constraint 167

7.2 Generalized Form of the QCQP 171

7.3 The Lagrange Dual Problem 172

7.3.1 Bounds for the Dual 173

7.3.2 Derivatives of the Dual Function 174

7.3.3 Maximizing the Dual 176

7.4 Computational Cost and Efficient Implementations 177

7.4.1 Extending previous efficient solvers to the General SINR QCQP 181

7.4.2 Rapid Estimation of the Template Upper Bound 183

7.4.3 Lower-Upper Preconditioner 184

7.5 A Note on Multi-Waveform Generalizations and Quadratically Constrained MiCRFt 188

7.6 QC LS MMF Closed-Loop Assessment 189

7.6.1 Algorithm Performance 190

7.6.2 Closed-Loop MMF Results 193

7.7 QC LS MMF Open-Air Experimental Results 194

8 Conclusions **204**

A Appendix **213**

A.1 Toeplitz Plus Low Rank Linear Systems 213

A.2 Inverse of the KKT matrix 214

A.3 Algebraic Simplifications for the ISL Dual 215

List of Figures

1.1 Conceptual view of a MTI radar scenario 2

1.2 Conceptual view of a phase code transition 3

1.3 Conceptual view of the instantaneous frequency of random FM waveform compared to an LFM and phase code; Figure modified from [1] 4

2.1 Conceptual Transmission and Receive Measurements for a single scatterer 8

2.2 Theoretical Receiver Operating Characteristics for Scatterer in White Gaussian Noise 9

2.3 Open-Air Measurement pre and post pulse-compression 11

2.4 Matched Filter and Windowed Matched Filter responses 12

2.5 Open-Air Experimental Measurement for the Matched Filter and Windowed Matched Filter 13

2.6 Straddle Loss (left) and Phase (right) for LFM Chirp 16

2.7 Simulated Loopback MF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 22

2.8 Simulated Loopback MF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 22

2.9 Simulated Loopback Akroyd's LS MMF Response for LFM waveform with $\kappa =$ 1.5 (left) and $\kappa = 6$ (right) 22

2.10 Simulated Loopback Akroyd's LS MMF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 23

2.11 Simulated Loopback Template-based LS MMF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 25

2.12 Simulated Loopback Template-based LS MMF PS for LFM waveform with $\kappa =$ 1.5 (left) and $\kappa = 6$ (right) 25

2.13 Simulated Loopback ISL LS MMF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 28
2.14 Simulated Loopback ISL LS MMF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 28
2.15 Simulated APC estimate for idealized measurements with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 36
2.16 Simulated APC noise power for idealized measurements with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 37
2.17 Simulated APC estimate with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 37
2.18 Simulated APC noise power with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 38
2.19 Simulated APC estimate with model uncertainty for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 40
2.20 Simulated APC noise power with model uncertainty for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 40
2.21 Simulated APC estimate with model uncertainty and beamspoiling for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 42
2.22 Simulated APC noise power with model uncertainty and beamspoiling for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 42
2.23 Simulated APC estimate with model uncertainty, beamspoiling, and a unity constraint for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 44
2.24 Simulated APC noise powers with model uncertainty, beamspoiling, and a unity constraint for $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 45
2.25 Ambiguity Function for a LFM Chirp (left) and RFM Waveform (right) 48
2.26 PSF for a 100 repeated LFM Chirps (left) and 100 Unique RFM Waveforms (right) 54
2.27 PSF for a 100 Unique RFM Waveforms (left) and 1000 Unique RFM Waveforms . 54
2.28 Open-Air Experimental Measurement for 1000 repeated LFM Chirps (left) and zoomed in look at intersection (right) 55

2.29 Open-Air Experimental Measurement for 1000 unique RFM waveforms (left) and zoomed in look at intersection (right) 55

3.1 Simulated Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ Uncompensated (left) and Range-Straddle Compensated (right) 77

3.2 Simulated Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ Uncompensated (left) and Range-Straddle Compensated (right) 78

3.3 Eigenvalues of $\mathbf{R}_\gamma$ for LFM waveform with $\kappa = 1.5$ 79

3.4 Residual Error vs Iteration of PCG using Circulant Preconditioners for single LFM (left) and 1000 PROFM (right) Waveforms with $\kappa = 1.5$ 84

3.5 Residual Error vs Iteration of PCG using Circulant Preconditioners for single LFM (left) and 1000 PROFM (right) Waveforms with $\kappa = 6$ 84

3.6 Loopback Template-based LS MMF Response for LFM waveform and $\kappa = 1.5$ . . 85

3.7 Loopback Template-based LS MMF Response for PROFM waveforms and $\kappa = 1.5$ 85

3.8 Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ 86

3.9 Loopback ISL LS MMF Response for PROFM waveforms and $\kappa = 1.5$ 87

3.10 Loopback Straddle Compensated ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ 88

3.11 Loopback Straddle Compensated ISL LS MMF Response for PROFM waveforms and $\kappa = 1.5$ 88

3.12 Open-Air Experimental Test Setup 90

3.13 LS MMF RMS Open-Air Range Profile Estimates $\kappa = 1.5$ 91

3.14 Range-Doppler Estimate for LFM CPI using the MF 92

3.15 Range-Doppler Estimate for LFM CPI using the Straddle Compensated ISL LS MMF 92

3.16 Range-Doppler Estimate for PROFM CPI using the MF 93

3.17 Range-Doppler Estimate for PROFM CPI using the Template-based LS MMF . . . 94

3.18 Range-Doppler Estimate for PROFM CPI using the ISL LS MMF 95

3.19 Range-Doppler Estimate for LFM CPI using the Straddle Compensated ISL LS MMF 95

4.1 Simulated Loopback MF Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 100

4.2 Simulated Loopback MF PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 101

4.3 Simulated Loopback ISL LS MMF Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 101

4.4 Simulated Loopback ISL LS MMF PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 102

4.5 Simulated Loopback ISL MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 102

4.6 Simulated Loopback ISL MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 103

4.7 Simulated Loopback Template-base MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 103

4.8 Simulated Loopback Template-base MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 104

4.9 Simulated Loopback Straddle Compensated ISL MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 106

4.10 Simulated Loopback Straddle Compensated ISL MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 106

4.11 MiCRFt Hessian before permutation (left), after permutation (middle), and zoomed in to the permuted matrix (right) 107

4.12 Preconditioning of MiCRFt QP 108

4.13 Example Toeplitz-like Matrix from (4.17) (left), it's Circulant-approximation preconditioner (middle), and a Truncated Circulant plus Block Preconditioner (right) . 110

4.14 Eigenvalue Distribution for Circulant Preconditioner (left) and Truncated Circulant plus Block Preconditioner (right) 110
4.15 MiCRFt Residual Error vs Iteration of PCG using Circulant Preconditioners for subsets of $Q = 2$ PROFM Waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 111
4.16 MiCRFt Residual Error vs Iteration of PCG using Circulant Preconditioners for subsets of $Q = 4$ PROFM Waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right) 111
4.17 Illustrative $Q = 2$ MiCRFt RSM for 1000 PROFM waveforms for stationary (left) and moving (right) point scatterer 112
4.18 Illustrative $Q = 2$ MiCRFt RSM after presumming for 1000 PROFM waveforms for stationary (left) and moving (right) point scatterer 113
4.19 Peak sidelobe level for MiCRFt subsets as a function of slow-time phase shift . . . 113
4.20 Loopback MiCRFt Response for PROFM CPI and $\kappa = 1.5$ 115
4.21 MiCRFt RMS Open-Air Range Profile Estimates $\kappa = 1.5$ 116
4.22 Range-Doppler Estimate for PROFM CPI using the Template-based MiCRFt with $Q = 2$ and $\kappa = 1.5$ 117
4.23 Range-Doppler Estimate for PROFM CPI using the ISL MiCRFt with $Q = 2$ and $\kappa = 1.5$ 118
4.24 Range-Doppler Estimate for PROFM CPI using the Straddle Compensated ISL MiCRFt with $Q = 2$ and $\kappa = 1.5$ 119

5.1 Scale of complexity to preserve uniform noise floor as a function of Doppler oversampling 126
5.2 Illustrative PCG performance of Proj-NIMPC 132
5.3 Range-Doppler Estimate for PROFM CPI using the NIMPC with unity gain scaling 135
5.4 Range-Doppler Estimate for PROFM CPI using the NIMPC with consistent noise power scaling 136
5.5 Range-Doppler Estimate for PROFM CPI using direct solution of Proj-NIMPC . . 137

5.6 Range-Doppler Estimate for PROFM CPI using 5 iterations of CG to approximate Proj-NIMPC 137
5.7 Range-Doppler Estimate for PROFM CPI using 5 iterations of PCG to approximate Proj-NIMPC 137

6.1 Simulated APC estimate for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right) 142
6.2 Simulated APC noise power for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right) 142
6.3 Simulated APC estimate for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right) 143
6.4 Simulated APC noise power for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right) 143
6.5 Simulated APC estimate for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right) 146
6.6 Simulated APC noise power for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right) 147
6.7 Simulated APC estimate for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right) 147

6.8 Simulated APC noise powers for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right) 148
6.9 Illustration of Extracting the Snapshot Structured-Covariance from the "Full" Structured Covariance 150
6.10 Full Structured Covariance (left) and it's Lower Triangular Cholesky Decomposition (right) 160
6.11 Detected outliers (left) and Structured Covariance (right) for first iteration 161
6.12 Detected outliers (left) and Structured Covariance (right) for second iteration . . . 161
6.13 Detected outliers (left) and Structured Covariance (right) for third/final iteration . . 161
6.14 RMS range estimates (left) and noise powers (right) for the Full APC implementations 163
6.15 Baseline Range-Doppler Estimate for PROFM CPI using APC 163
6.16 Range-Doppler Estimate for PROFM CPI using APC with Straddling Compensation 164
6.17 Range-Doppler Estimate for PROFM CPI using APC with template-based objective 164
6.18 Range-Doppler Estimate for PROFM CPI using Approximate Outlier APC with Straddling Compensation 165

7.1 Iterative estimate of the template-based dual upper bound 185
7.2 Condition Number Bounds for the ULB Preconditioner when $\mu_u > \lambda_b$ 187
7.3 Condition Number Bounds for the ULB Preconditioner when $\mu_u < \lambda_b$ 187
7.4 Resulting MMFs for different tolerances 191
7.5 RMS Filter Response for MF and MMFs with ρ=-2dB; $D = 4N$; k=3 (MMF 1), k=5 (MMF 2), or k=13 (MMF 3) a. - LFM with SNR constraint; b. - LFM with SINR constraint; c. - PROFM with SNR constraint; d. - PROFM with SINR constraint; 195

7.6 Mean Frequency Response for MF and MMFs with ρ=-2dB; $D = 4N$; k=3 (MMF 1), k=5 (MMF 2), or k=13 (MMF 3) a. - LFM with SNR constraint; b. - LFM with SINR constraint; c. - PROFM with SNR constraint; d. - PROFM with SINR constraint; 196
7.7 Open-Air Experimental Test Setup 197
7.8 Toeplitz Model D=4N, ρ=-2dB, k=5 (ISL) or $\kappa_{\mathbf{w}} = 1.3$ (template) with no interference 198
7.9 Matched Filter vs Optimal SINR filters for OFDM interference with 20dB INR . . 199
7.10 Histogram of intersection for PROFM and LFM Waveform sets with (right) and without (left) interference 200
7.11 Toeplitz Model D=4N, ρ=-2dB, k=5 (ISL) or $\kappa_{\mathbf{w}} = 1.3$ (template), OFDM interference with 20dB INR 201
7.12 Approximate Circulant Model D=L_{PRI}, ρ=-2dB, k=3 (ISL) or $\mathbf{p}$-gaussian shape, OFDM interference with 20dB INR 203

List of Tables

3.1 Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms . 83

3.2 Complexity for solving the Toeplitz System of Equations 89

4.1 Table of Experimental Parameters for PROFM CPI of 1000 waveforms 115

4.2 Complexity for constructing MiCRFt Filters $Q = 2$ 115

4.3 Complexity for constructing MiCRFt Filters $Q = 5$ 115

5.1 Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms . 133

5.2 Complexity for solving $\tilde{\mathbf{D}}$ with NIMPC or applying Proj-NIMPC 134

5.3 Complexity for Range-Doppler Processing with NIMPC and Proj-NIMPC 134

6.1 Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms . 159

6.2 Complexity for constructing the Structured Covariance Matrix 159

6.3 Complexity for Cholesky Decompositions 160

6.4 Average Complexity for single APC Iteration 162

7.1 Calls to $EvaluateDual\{\cdot\}$ - ISL MMF D=4N, ρ=-2dB, k=3 190

7.2 Total number of PCG iterations - ISL MMF D=4N, ρ=-2dB, k=3 192

7.3 Order of Complexity - ISL MMF L=4N, ρ=-2dB, k=3 192

7.4 Total number of non-parallelizable Toeplitz multiplies/Levinson recursions - ISL MMF D=4N, ρ=-2dB, k=3 192

7.5 Toeplitz MMFs - median of intersection (dB) 202

7.6 Circulant Approximate MMFs - median of intersection (dB) 203

Chapter 1

Introduction

Traditional radar operation is well established. Most legacy systems actively illuminate a scene with a repeated electromagnetic waveform, usually a linear frequency modulated (LFM) chirp, and collect information by measuring the scattered echoes [2–4]. For a pulsed radar using a repeated waveform, receive measurements are often separated into two different dimensions, a "fast-time" dimension corresponding to measurements within a pulse repetition interval (PRI) and a "slow-time" dimension corresponding to the change from pulse-to-pulse. These fast-time and slow-time components are aggregated into a coherent processing interval (CPI) and, depending on application, subsequent processing is performed to estimate scatterer parameters such as range, velocity, polarimetric reflectivity [5], refractivity [6], changes since last collect [7], etc. See Figure 1.1 for a conceptual view of this work's primary point of emphasis, a moving target indication (MTI) scenario in which the radar seeks to differentiate moving objects from stationary clutter. The prevalence of the repeated chirp framework primary stems from the LFM's simplicity to generate (sweeping the frequency of a tone) and process (convolution with a repeated filter), though it has additional benefits such as Doppler robustness and exceptional spectral containment. However, this transmission scheme may lack the flexibility needed to operate within the dynamic and increasingly congested modern RF environment [8–10].

Extensions beyond use of the LFM can be placed under the broad category of waveform diversity [11–14], which has recently been enabled by advancements in processing power and sophisticated analog-to-digital/digital-to-analog hardware that has rapidly transformed cellular communication over the last few decades. Modern hardware facilitates arbitrary waveform generation (AWG) with tremendous design degrees of freedom. Coupled with advanced digital receive pro-

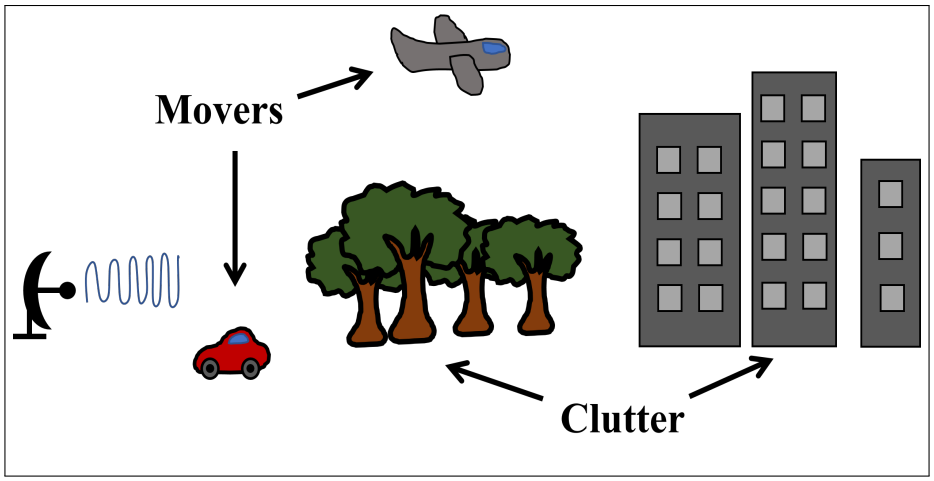

Figure 1.1: Conceptual view of a MTI radar scenario

cessing, system designs can be tailored to better suit operational requirements, e.g. joint radar-communications [15–22], spectral cohabitation [23–35], dual polarization [36], low range sidelobes [37–42], image formation [43], etc., [1, 44]. That said, transitioning these new capabilities to operational systems remains an ongoing effort as theoretically optimal solutions often are not robust, lack real time plausibility, or ignore physical effects of real hardware.

AWGs offers a high degree of flexibility in design, which for MTI applications often seeks to reduce range sidelobes after pulse compression. However, fully unconstrained design faces challenges when implemented on physical hardware. Radar systems often require fast reaction times and inherently contend with a two-way propagation loss from transmission and scattering of spherically spreading electromagnetic waves [2, 5]. Thus, operation with these shorter time scales and additional loss (relative to communication systems) necessitates the use of high-power amplifiers (HPAs) that operate at or near saturation to maintain power efficiency [2]. This requirement effectively restricts the waveform design to a constant modulus (also referred to amplitude or envelope) solution space.

On paper, segmenting a constant envelope waveform into discrete chips and designing the phase, denoted phase coding [45], fits this criterion. These phase values can then be arbitrarily selected via an optimization procedure to reduce pulse compression ambiguities [46], mitigate interference [47], maximize information [48, 49], etc. Unfortunately, most "optimal" phase coding

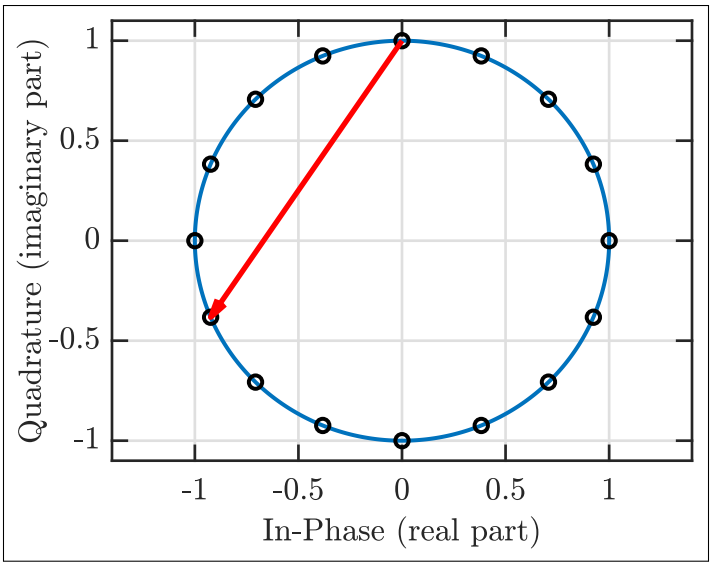

Figure 1.2: Conceptual view of a phase code transition

neglects the transition period between chips, which is not instantaneous in practice [50]. Thus, an unrestricted transition period will induce spectral spreading or amplitude modulation as the shortest path is not restricted to the unit circle (see Figure 1.2). This generally degrades performance of optimal phase coding in practice as they lack robustness towards hardware distortion, Doppler shifts, and range straddling. To improve robustness, it is often necessary to enforce "smooth" transition regions by restricting how far adjacent chips can deviate. From the principle of stationary phase [51], this is (approximately) equivalent to enforcing "spectral containment" outside the waveform's 3-dB bandwidth. However, these constrained solutions often deviate from the theoretically optimal phase coded radar waveforms.

The combination of the constant modulus and spectral containment requirements has led to exploration of higher-order design where the waveform's instantaneous frequency [50, 52] is tuned to produce arbitrary frequency modulated (FM) waveforms. These FM waveforms have been applied to a variety of applications [1] and their robustness towards hardware distortion has been experimentally verified. Figure 1.3 includes a conceptual illustration of the instantaneous frequency for the classical LFM compared with a phase coded and a FM waveform with random instantaneous frequency. The inherent restriction of spectral containment and constant modulus results in an extremely non-convex optimization surface with a plethora of local minima that do not necessary

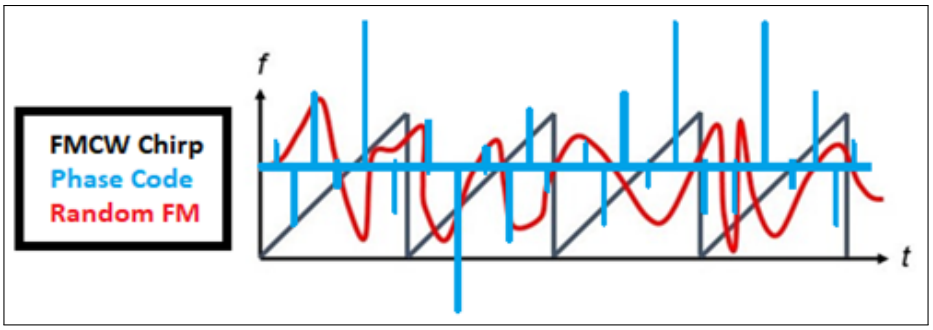

Figure 1.3: Conceptual view of the instantaneous frequency of random FM waveform compared to an LFM and phase code; Figure modified from [1]

perform exceptionally on a per-pulse basis. Exhaustive searches for the globally optimal solution are exceedingly costly, often non-deterministic polynomial-time (NP) hard [53], and rapidly becomes implausible for even a modest number of design variables.

While not immediately apparent, having many local minima can be advantageous. The sheer number of different locally optimal solutions effectively produces entire classes of physically realizable Pseudo–Random FM (RFM) waveforms that can be tuned to address a specific design. Additionally, transmitting multiple unique locally optimal waveforms on a pulse-to-pulse basis provides new information about the scene that can be aggregated by performing subsequent slow-time processing through a coherent integration/combination effect. In some circumstances, this aggregate response may even outperform the repeated transmission of the single globally optimal solution. However, this pulse-to-pulse variation produces a new problem. The variable emission results in a coupling between fast-time (range) and slow-time (Doppler) and induces a phenomena known as range-sidelobe modulation (RSM) [1, 17]. This RSM induces a joint dimensional "sidelobe smearing" in range-Doppler that is unaffected by classic sidelobe mitigation techniques (e.g., Doppler tapering or windowing). With standard matched-filter processing, RSM levels are typically on the order of per-pulse autocorrelation sidelobes divided by the number of unique waveforms. From the central limit theorem, it can be deduced that the per-pulse sidelobe level of RFM waveforms in expectation are roughly proportional to $\frac{1}{TB}$ where TB is the time-bandwidth product. Optimization can theoretically lower this further [37, 46], but most local minima only reduce sidelobe variance and solutions are still bounded by TB.

Alternatively, sophisticated receive processing architectures that leverage the expected structure of the TX-to-RX channel can be employed to reduce pulse-to-pulse sidelobes and thus lower RSM. In particular, structure-based least squares (LS) mismatched filters (MMF) [54–56] can be designed to yield lower sidelobes and employed in lieu of the matched-filter while adaptive techniques such as adaptive pulse compression (APC) [57, 58] construct unique filters for every range bin to that place nulls in the directions of strong scatters. These frameworks can even extend to a joint fast-time slow-time representation where the multiplicative degrees of freedom further enhance RSM suppression capabilities [59–61].

However, all these approaches fundamentally rely on a discretized approximation of a continuous reality and model errors can induce bias, create false detections, or restrict effective dynamic range. As such, these structure-based approaches must be compensated to account for the limitations of physical reality. Additionally, matched-filter based processing only requires knowledge of the transmitted waveforms and can be implemented efficiently using the fast-Fourier transform (FFT) in inverse-FFT (IFFT) with $\sim\mathscr{O}(N\log(N))$ complexity, while LS and adaptive approaches require solving systems of linear equations to effectively utilize the available DoF which has a brute force complexity of $\mathscr{O}(N^3)$[1]. As such, physically meaningful transition of these structure-based approaches must account for both the increased computational burden and discretization based model errors.

[1] Here N is an arbitrary value corresponding to the specific dimensionality of the problem. As such, these complexity orders are provided to illustrate the approximate computational cost as the dimensionality increases.

Chapter 2

Background

The following sections introduce basic radar range estimation theory and its discretized linear representation. Using this framework, more sophisticated range estimation techniques are then examined and their pros/cons discussed. Following the sections on range estimation, the receive models are extended to incorporate Doppler shifts from moving scatterers and range-Doppler estimation techniques are introduced. The chapter is then wrapped up by examining particular matrix structures and linear system solvers to lay the groundwork for subsequent efficient implementations.

2.1 Radar Range Processing Overview

Real radar systems either measure and/or emit real continuous electromagnetic waves in order to ascertain information about an environment. While some legacy systems do operate entirely in the analog domain, the advent of digital signal processing has enabled a plethora of advanced TX/RX architectures and algorithms which offer enhanced capabilities not possible with analog systems. However, approximating the continuous reality as a discrete model inherently produces error and aliasing as fundamentally no time-limited signal can be truly band-limited and no band-limited system can be truly time-limited [62]. Additionally, the TX/RX architectures of a radar system are not linear as RF amplification will induce nonlinear harmonics [63, 64], while analog-to-digital (or digital-to-analog) conversion rely on imperfectly-synchronized reference clocks and requires quantization (potentially with single precision or less), all of which produce non-linear error and bias. That being said, careful design and calibration coupled with high-fidelity hardware can help

mitigate higher order effects, and by performing multiple "looks" a coherent signal gain can be achieved while suppressing non-coherent errors. Thus, approximate discrete linear receive models can be posed if the mapping from continuous-to-discrete is represented with enough fidelity that inherent uncertainties and errors do not significantly degrade dynamic range.

The remainder of this section begins by outlining the mapping of continuous transmission to discrete receive measurements for a pulsed monostatic radar. Leveraging the discretized fast-time receive model, subsequent LS and adaptive range estimation techniques are then introduced to reduce range sidelobes. The subsequent section will then introduce the pulse-to-pulse slow-time model and discuss how fast-time range estimates can be used to perform velocity estimation.

2.1.1 Pulse Echo Model

Conceptually, an active monostatic radar system illuminates a scene with electromagnetic waveform and gathers range information by measuring the time delays of the echoes. For a constant modulus transmit pulse

$$s(t) = \begin{cases} Ae^{j2\pi f_c t} & 0 \leq t < T_p \\ 0 & else \end{cases} \tag{2.1}$$

where f_c is the center frequency, A is the waveform amplitude, and T_p is the pulse duration, the receive measurements can be modeled as the convolution of $s(t)$ with the delay dependent complex scattering of the TX-to-RX propagation path $x(t)$ as

$$y(t) = (s(t)\,(*_C)x(t)) + n(t) \tag{2.2}$$

where $(*_C)$ denotes the continuous convolution operation and $n(t)$ is additive system noise. The measured time delay can then be turned into an approximate baseline range to the scatterer as

$$R = \frac{ct}{2} = \frac{R_{TX} + R_{RX}}{2} \tag{2.3}$$

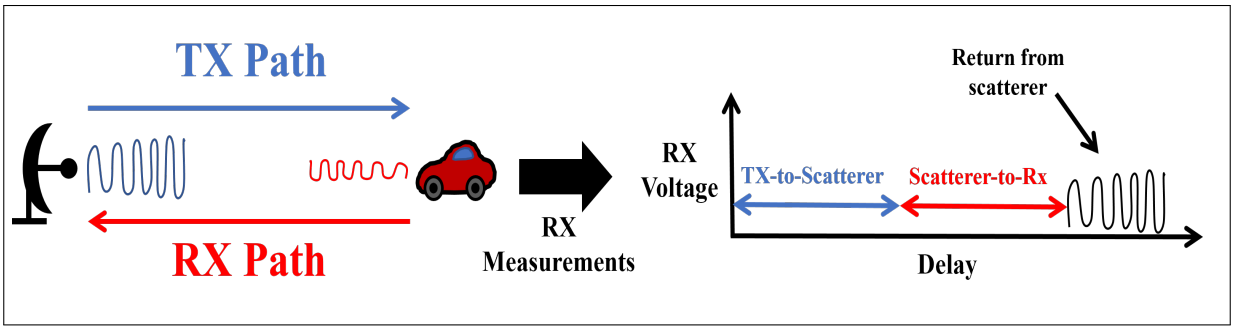

Figure 2.1: Conceptual Transmission and Receive Measurements for a single scatterer

where R_{TX} and R_{RX} are the ranges from TX-to-scatter and scatter-to-RX respectively. See Figure 2.1 for a conceptual visualization for monostatic (co-located TX and RX) radar ranging.

As the propagation path is "two-way" (i.e. TX-to-scatterer-to-RX) the effective range resolution is

$$\Delta R \approx \frac{cT_p}{2} \tag{2.4}$$

where $c \approx 299,792,458 m/s$ is the speed of propagation (for RF) and determines when two closely spaced responses can be separated. The expected received power (P_{RX}) for a scatterer can be approximated using the objects "radar cross section" σ as

$$P_{RX} = \frac{P_{TX} G_{TX} G_{RX} \sigma \lambda_c^2}{(4\pi)^3 L_{RX} R_{TX}^2 R_{RX}^2} \tag{2.5}$$

where $P_{TX} = |A|^2$ is the instantaneous transmit power, G_{TX} and G_{RX} are the TX and RX antenna gains, $\lambda_c = \frac{c}{f_c}$ is the RF wavelength, and L_{RX} is the receiver loss [2], though in reality σ is not constant and P_{RX} will fluctuate as a function of aspect angle, frequency, velocity, etc. For a monostatic configuration $R_{TX} = R_{RX} = R$ simplifying (2.5) to

$$P_{RX} = \frac{P_{TX} G \sigma \lambda_c^2}{(4\pi)^3 L_{RX} R^4} \tag{2.6}$$

where clearly the received power rapidly decreases with range due to the inverse relationship with the R^4 term. On receive, detection with a small chance of a false alarm requires a P_{RX} much greater

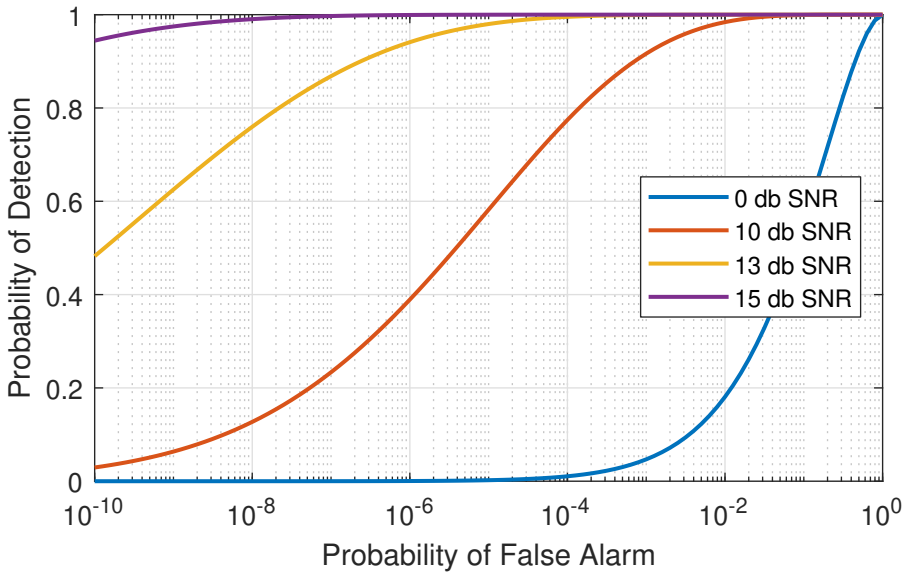

Figure 2.2: Theoretical Receiver Operating Characteristics for Scatterer in White Gaussian Noise

than the receive noise power, which for white-Gaussian system noise is defined as

$$P_N = kT_0B_NF_N \tag{2.7}$$

where $k = 1.38 \times 10^{-23}\ \frac{W}{HzK}$ is Boltzmann's constant, $T_0 = 290\ K$ is the reference temperature of the system noise figure F_N, and B_N is the effective noise bandwidth. Probability of detection can then be characterized by the signal-to-noise ratio (SNR)

$$\frac{P_{RX}}{P_N} = \frac{P_{TX}G\sigma\lambda_c^2}{(4\pi)^3kT_0B_NF_NL_{RX}R^4} \tag{2.8}$$

where a simple SNR threshold can be used as a detector where the probability of detection vs false-alarm for a particular threshold can be computed using the noise and scatterers' probability density functions (pdf) [65, 66]. Figure 2.2 provides a illustration of these receive operation characteristics for an idealized example with a constant scatterer in white Gaussian noise (WGN) for different threshold values. Thus without signal modulation, achieving sufficient SNR to detect and separate close scatters at long range requires both a high instantaneous transmit power and an extremely short pulse.

2.1.2 Pulse Compression

As designing a system that can achieve sufficient energy on target with enough range resolution to resolve close scatterers can be extremely expensive and runs into practical system fidelity limits, modern radars typically take an alternative approach. Rather than a short high power pulse, the transmit signal is instead modulated as

$$s(t) = \begin{cases} Ae^{j2\pi f_c t + j\Theta(t)} & 0 \leq t < T_p \\ 0 & else \end{cases} \tag{2.9}$$

where $\Theta(t)$ is an instantaneous phase-modulation corresponding to the baseband waveform

$$s_{BB}(t) = \begin{cases} Ae^{j\Theta(t)} & 0 \leq t < T_p \\ 0 & else \end{cases} \tag{2.10}$$

which has a 3dB bandwidth of B. This bandwidth enables receive filtering operations that "compresses" the waveform response, aptly denoted pulse-compression. Traditionally, pulse-compression is performed by correlating the receive measurements with the complex conjugate of the transmit waveform via the filtering operation

$$\hat{x}(t) = \left(s^{\beta_C^*}(t) (*_C) y(t) \right) \tag{2.11}$$

where β_C^* denotes time-reversal and complex-conjugation. This specific filter (time-reversed complex-conjugate replica) is denoted the matched filter (MF) as it provides the maximum SNR estimate for white-Gaussian noise. MFing is one of the most common radar receive processing techniques as it is inexpensive to implement, provides an SNR gain proportional to T_p (for constant modulus waveforms), has a range resolution (or mainlobe width) inversely proportional to the bandwidth as

$$\Delta R \approx \frac{c}{2B}, \tag{2.12}$$

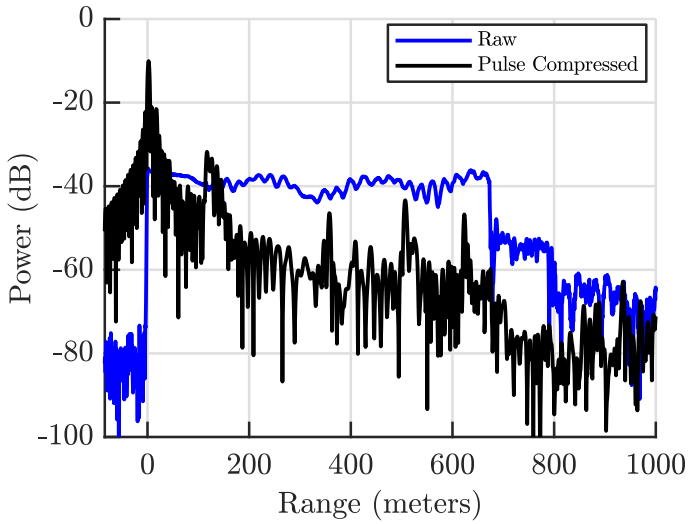

Figure 2.3: Open-Air Measurement pre and post pulse-compression

and is fairly robust to model and system mismatch. See Figure 2.3 for an example of open-air experimental measurements pre and post pulse-compression for a 4.5μs LFM pulse with a bandwidth of 33.3MHz (range resolution of 4.5 meters). An analogous technique which performs range compression in the frequency domain (as opposed to compression in the time domain) can be performed via stretch processing [67] but has traditionally restricted the emissions to LFM chirps, though there have been recent extensions beyond this framework [68, 69]. For brevity, here the scope will be restricted to pulsed-monostatic operation.

Pulse-compression is not without flaw as the filtering process produces a sidelobe structure in range proportional to the waveform's autocorrelation (inverse Fourier transform of the power-spectral-density). These range-sidelobes can in turn mask small scatterers in the vicinity of large scatterers, effectively limiting dynamic range locally and biasing detentions via adjacent scatterer correlation/interference. There are a plethora of techniques aimed at mitigating range-sidelobes, largely focused on using alternative pulse compression filters such as a tapered (sometimes called windowed) version of the transmit waveform or a mismatched filter. However, deviation from the matched filter will always produce some degree of mismatch loss (MML), i.e. less SNR, and may degrade the effective range-resolution through a broadening of the filter response mainlobe. See Figure 2.4 for an example of the filter response for a LFM waveform when using the matched filter

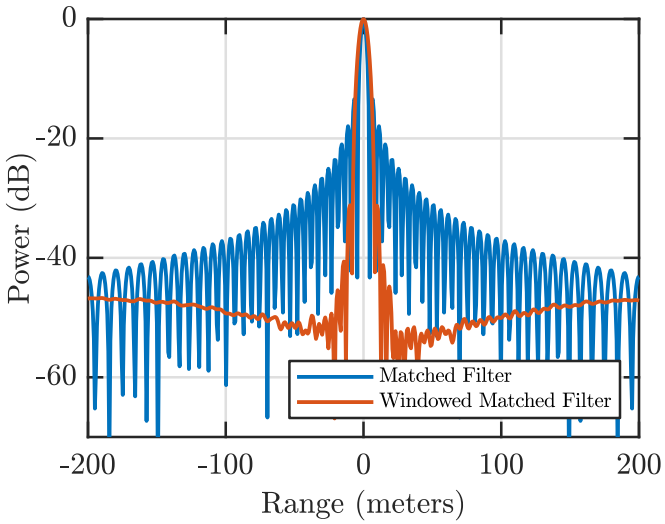

Figure 2.4: Matched Filter and Windowed Matched Filter responses

vs a tapered filter as well as Figure 2.5 for corresponding open-air measurements.

2.1.3 Discretized Receive Model

To offer a greater degree of design flexibility and enable advanced processing, the TX waveform and RX measurements are often represented as discrete variables. From the Nyquist-Shannon sampling theorem a bandlimited signal can be perfectly reconstructed without aliasing when the complex sample rate (both in-phase and quadrature sampling) is equal to the bandwidth [62]. However, the time-frequency Fourier relationships tell us that any “time-limited" signal will inherently have an infinite bandwidth while ”band-limited" signals have infinite duration. In practice, the sample rate is often selected relative to the 3dB bandwidth of the radar waveform and a combination of oversampling and anti-alias filtering are used to mitigate this model mismatch, though some aliasing will always be present.

For complex baseband waveform $s(t)$ with pulse duration T_p and 3dB bandwidth B, a $D \times 1$ collection of contiguous discrete receive measurements can be oriented in a “snapshot” notation

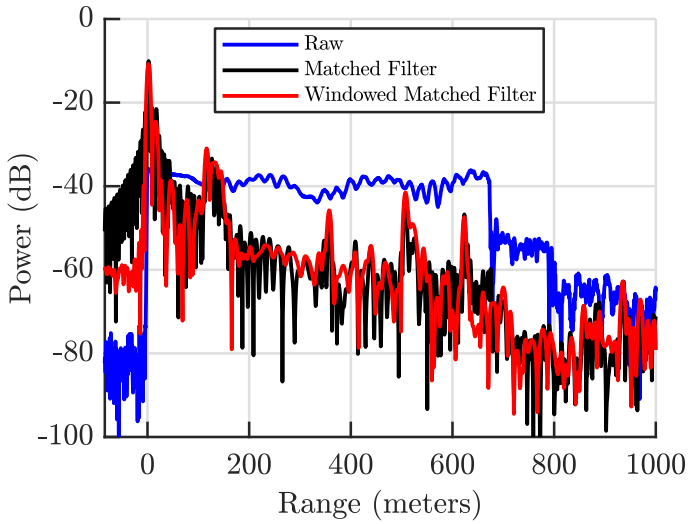

Figure 2.5: Open-Air Experimental Measurement for the Matched Filter and Windowed Matched Filter

and expressed as

$$\begin{aligned}\mathbf{y}(\ell) &= \left[y\left(\tfrac{1}{f_s}(\ell - 1 - \tfrac{D-N}{2})\right) \quad y\left(\tfrac{1}{f_s}(\ell - \tfrac{D-N}{2})\right) \quad \cdots \quad y\left(\tfrac{1}{f_s}(\ell + D - \tfrac{D-N}{2} - 2)\right)\right]^T \\ &= \left[y_{\ell-1-\frac{D-N}{2}} \quad y_{\ell-\frac{D-N}{2}} \quad \cdots \quad y_{\ell+D-\frac{D-N}{2}-2}\right]^T \\ &= \int_{-0.5}^{0.5} \mathbf{S}(\gamma)\,\mathbf{x}(\ell,\gamma)\,d\gamma + \mathbf{n}(\ell)\end{aligned} \tag{2.13}$$

where ℓ is a discrete index corresponding to time delay $\frac{\ell-1}{f_s}$ with f_s being the RX sample rate. Here $\mathbf{x}(\ell,\gamma)$ is the $(D+N-1)\times 1$ vector of complex scattering coefficients

$$\begin{aligned}\mathbf{x}(\ell,\gamma) &= \left[x\left(\tfrac{1}{f_s}(\ell - \tfrac{D}{2} + \gamma)\right) \quad \cdots \quad x\left(\tfrac{1}{f_s}(\ell + \tfrac{D}{2} + N - 3 + \gamma)\right)\right]^T \\ &= \left[x_{\ell-\frac{D}{2}}(\gamma) \quad \cdots \quad x_{\ell+\frac{D}{2}+N-3}(\gamma)\right]^T\end{aligned} \tag{2.14}$$

for sample delay ℓ and sub-sample delay $\gamma \in (-0.5, 0.5]$ corresponding to a time delay of $\frac{\gamma}{f_s}$,

$$\mathbf{S}(\gamma) = \begin{bmatrix} s_N(\gamma) & \cdots & s_1(\gamma) & 0 & \cdots & 0 \\ 0 & s_N(\gamma) & & s_1(\gamma) & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & s_N(\gamma) & \cdots & s_1(\gamma) \end{bmatrix} \tag{2.15}$$

is a $D \times (D+N-1)$ Toeplitz convolution matrix corresponding to discrete samples of the continuous waveform

$$\mathbf{s}(\gamma) = \begin{bmatrix} s(\frac{0+\gamma}{f_s}) & s(\frac{1+\gamma}{f_s}) & \cdots & s(\frac{N-1+\gamma}{f_s}) \end{bmatrix}^T = \begin{bmatrix} s_1(\gamma) & s_2(\gamma) & \cdots & s_N(\gamma) \end{bmatrix}^T \tag{2.16}$$

where $N = T_p f_s$ is the waveform length in samples.

To avoid redundancy and complications with snapshot centering, the number of samples in a snapshot is typically selected to satisfy

$$D \geq N,$$

and D should have the same parity as N (i.e., both odd or even) with $D = N$ corresponding to the "standard" choice when performing MF processing. Note, (2.13) can be composed of delays that correspond to negative time or times greater than the pulse-repetition interval which after discretizing corresponds to

$$L = \frac{T_{PRI}}{f_s}$$

discrete samples. When receive contributions from pulse-eclipsing (e.g., part of the receive measurement is cut-off due to RX saturation or a front-end TX-RX switch) or multiple-time-around-returns (e.g., scattering from previous transmit intervals arriving in current receive interval) are negligible, the delay indices used for range processing are bounded by

$$1 \leq \ell \leq L$$

and measurement indices outside the range can be treated as or padded with zeros, i.e.,

$$y_\ell = 0$$

for $\ell \leq 0$, $\ell \geq L+1$ and

$$x_\ell(\gamma) = 0$$

$\ell+\gamma \leq 0$, $\ell+\gamma-N-2 \geq L+1$. For certain TX-RX configurations, such as low TX PRF/duty cycle or a well isolated bistatic orientation, these contributions can be below the system noise floor. For the general case, pulse-eclipsing and multiple-time-around-returns are often non-negligible, however, compensation to RX model is relatively straightforward and can be incorporated via an element-wise switch-loss weighting or additional summations for the extended interval respectively. For brevity, such generalizations have been excluded from this work.

The sub-sample shift term γ in (2.13) retains a continuous structure and a range-straddling mismatch will arise if the receive model is truly discrete (i.e., discrete variables with no dependencies on γ). However, by incorporating a frequency oversampling factor $\kappa > 1$ relative to the 3dB bandwidth (i.e. $f_s = \kappa B$) and an anti-aliasing filter, the relative change in waveform structure across these sub-sample shifts can be largely mitigated. For instances, Figure 2.6 illustrate the loss (left) and phase shift (right) respectively for a straddled LFM waveform with different spectral oversampling factors κ. That being said, increasing κ directly corresponds to an increased computational overhead. It is worth noting that the LFM has near idealized spectral containment characteristics and incorporation of an anti-aliasing filter will induce little distortion, while more general FM waveforms may necessitate a higher degree of oversampling.

As the straddle loss

$$\Gamma(\gamma) = \sqrt{\left|\frac{\mathbf{s}(0)^H \mathbf{s}(\gamma)}{\mathbf{s}(0)^H \mathbf{s}(0)}\right|^2}$$

and phase shift

$$\phi(\gamma) \angle \mathbf{s}(0)^H \mathbf{s}(\gamma)$$

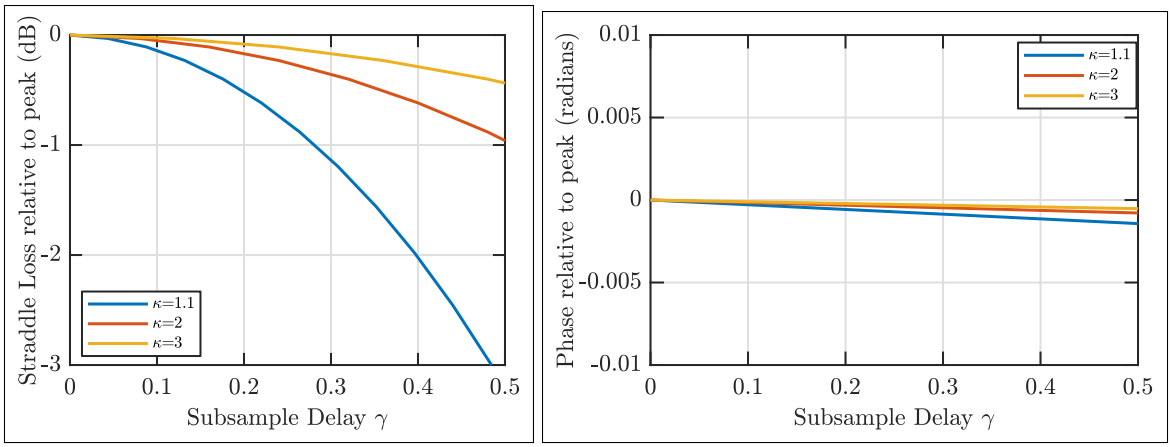

Figure 2.6: Straddle Loss (left) and Phase (right) for LFM Chirp

as a function of γ are relatively small for modest oversampling factors, the approximation

$$\mathbf{s}(\gamma) \approx \Gamma(\gamma)\, e^{j\phi(\gamma)}\mathbf{s}(0) = \Gamma(\gamma)\, e^{j\phi(\gamma)}\mathbf{s} \tag{2.17}$$

is typically incorporated so that (2.13) can be expressed as

$$\mathbf{y}(\ell) = \mathbf{S}\mathbf{x}(\ell) + \mathbf{n}(\ell) \tag{2.18}$$

where

$$\mathbf{S} = \begin{bmatrix} s_N & \cdots & s_1 & 0 & \cdots & 0 \\ 0 & s_N & & s_1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & s_N & \cdots & s_1 \end{bmatrix} \tag{2.19}$$

is the $D \times (D+N-1)$ convolution matrix for the unstraddled discrete waveform $\mathbf{s}$ and

$$\mathbf{x}(\ell) = \int_{-0.5}^{0.5} \Gamma(\gamma)\, e^{j\phi(\gamma)}\mathbf{x}(\ell,\gamma)\, d\gamma \tag{2.20}$$

is the integrated response of the continuous scattering coefficients across a range-bin. This approximation falls into a broad class of grid-based estimation approaches [70] and is often sufficient for spectrally-contained FM waveforms or selected coded waveforms with spectral oversampling, but

may require additional compensation to retain a high dynamic range for near critically sampled ($\kappa \approx 1$) receive architectures.

Rather than the $D \times 1$ snapshot representation in (2.13), it can be convenient to organize the receive measurements into a "full" representation in which all L non-negligible receive samples within a PRI ($T_{PRI} = \frac{L}{f_s}$) are concatenated into the $L \times 1$ receive vector

$$\begin{aligned} \mathbf{u} &= \begin{bmatrix} y_1 & \cdots & y_L \end{bmatrix}^T \\ &= \mathbf{V}\mathbf{x} + \mathbf{n} \end{aligned} \tag{2.21}$$

where

$$\mathbf{V} = \begin{bmatrix} s_1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ s_N & & \ddots & 0 \\ 0 & \ddots & & s_1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & s_N \end{bmatrix} \tag{2.22}$$

is now a $L \times (L-N+1)$ convolution matrix corresponding to a $(L-N+1) \times 1$ vector of non-eclipsed complex scatterers $\mathbf{x}$. Note, here multiple-time-around-returns and pulse-eclipsing effects are again absent from the RX representation and the components previously set to zero in (2.18) are now excised entirely, yielding a slightly different convolution matrix form compared to (2.19).

Maximum SNR estimation for the discretized scattering coefficients can then be acquired using the discrete matched filter

$$\mathbf{w} = \frac{\mathbf{v}}{\mathbf{s}^H \mathbf{s}} \tag{2.23}$$

where $\mathbf{v}$ is the $ceil\{\frac{D+N-1}{2}\}th$ column of (2.19) and is formed via symmetric zero-padded the discretized waveform to yield the $D \times 1$ vector

$$\mathbf{v} = \begin{bmatrix} \mathbf{0}^T & \mathbf{s}^T & \mathbf{0}^T \end{bmatrix}^T \tag{2.24}$$

where the $\mathbf{0}$ components are $\frac{D-N}{2} \times 1$ vectors of zero and a $D = N$ corresponds to $\mathbf{v} = \mathbf{s}$. The MF (or any arbitrary $D \times 1$filter) can then be applied as

$$\hat{x}(\ell) = \mathbf{w}^H \mathbf{y}(\ell) \tag{2.25}$$

to estimate the complex scattering at each delay index. Matched filter estimation can likewise be performed for full PRI representation of (2.21) as

$$\hat{\mathbf{x}} = \frac{1}{\mathbf{s}^H \mathbf{s}} \mathbf{V}^H \mathbf{u}. \tag{2.26}$$

For a single filter applied to estimate all range, this process is equivalent to the discrete convolution between the receive measurements and the time-reversed complex-conjugate of the filter, i.e.,

$$\hat{\mathbf{x}} = \mathbf{u}(*_D)\mathbf{w}^{\beta_d^*} \tag{2.27}$$

where here $(*_D)$ and β_d^* act as the discrete convolution and flip-complex-conjugate operators respectively.

2.2 Least Squares Mismatched Filtering Overview

While MF range processing provides theoretically maximal SNR (assuming the model is ideal and the noise is white), correlation with adjacent bins will produce range-sidelobes proportional to the waveform's autocorrelation. These sidelobes may then obfuscate weak signatures near high powered scatterers, or appear as false detections. Furthermore, the MF assumes all noise is white and may amplify additive interference present in the environment (such as a neighboring communication system), further degrading detection. Alternatively, arbitrary mismatched filters (MMF) can be designed to produce a particular response that better suits to the desired application [54, 55, 71, 72]. These MMFs will always come with a degree of SNR loss and an increased computational burden relative to the MF, but the benefits can greatly outweigh the cost if properly designed.

Within the context of the discrete snapshot model of (2.18) and filter application of (2.25), any $D \times 1$ filter $\mathbf{w}$ will have a $(D+N-1) \times 1$ response $\mathbf{r}$ in range corresponding to

$$\mathbf{r}^T = \mathbf{w}^H \mathbf{S}. \tag{2.28}$$

Thus, as the choice of $\mathbf{w}$ is arbitrary, it can be directly designed to yield a filter response with desirable properties. However, the filter length (D) is strictly less than the extent of the response vector ($D+N-1$), resulting in an overdetermined problem. Since the filter length can be increased arbitrarily, filter-design may asymptotically approach a critically determined case, though longer filters will inherently lead to a extended range response and typically have more computational overhead (larger system of equations and more storage requirements).

MMF design typically falls into two broad categories, template-based approaches which seek to match $\mathbf{r}$ to a particular structure, and sidelobe-minimization approaches which simply attempt to minimize components of $\mathbf{r}$ outside of the near-zero delay mainlobe elements. Both philosophies have functional merit and the choice of approach is often application dependent (e.g., a template approach may provide improved performance when the transmit waveforms are agile while a repeated waveform may benefit more from a sidelobe-minimization approach). In the remainder of this section, select methods to perform MMF design are introduced, and theoretical efficacy with respect to the receive model are discussed.

2.2.1 Template-Based Least Squares MMFs

The original LS MMF approach was posed for phase-coded radar waveforms in [54] and defined an objective function for a $D \times 1$ filter $\mathbf{w}$ that minimizes range sidelobes as

$$\|\mathbf{r}^* - \mathbf{e}\|_2^2 = \left\|\mathbf{S}^H \mathbf{w} - \mathbf{e}\right\|_2^2 \tag{2.29}$$

where $\mathbf{S}$ is defined in (2.19) as the a $(D+N-1)\times D$ Toeplitz convolution matrix for the $N\times 1$ transmit waveform and $\mathbf{e}=\begin{bmatrix}0 & \cdots & 0 & 1 & 0 & \cdots & 0\end{bmatrix}^T$ is a $(D+N-1)\times 1$ elementary vector with a 1 at the ℓth element which acts as the desired filter response (an impulse).

The filter that minimizes (2.29) can be equivalently solved via a standard quadratic program (QP)

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H\mathbf{T}\mathbf{w}-2\,\mathrm{Re}\left\{\mathbf{w}^H\mathbf{b}\right\} \tag{2.30}$$

where $\mathbf{T}=\mathbf{S}\mathbf{S}^H$ is positive-semi-definite (PSD) Toeplitz matrix and

$$\mathbf{b}=\mathbf{S}\mathbf{e}=\left[\mathbf{0}^T,\mathbf{s}^T,\mathbf{0}^T\right]^T=\mathbf{v}.$$

The optimal solution to (2.30) has the well known form

$$\mathbf{w}=\mathbf{T}^{-1}\mathbf{b}=\left(\mathbf{S}\mathbf{S}^H\right)^{-1}\mathbf{S}\mathbf{e} \tag{2.31}$$

when $\mathbf{T}$ is positive-definite (PD), alternatively the pseudo-inverse

$$\mathbf{w}=\left(\mathbf{S}^H\right)^{\dagger}\mathbf{e} \tag{2.32}$$

can be performed when it is PSD.

This formulation was originally designed for critically sampled phase codes, which often face practical limitations as real hardware necessitates some degree of oversampling relative to the waveform's 3dB bandwidth B to minimize distortion induced by anti-alias filtering and avoid amplifier nonlinearities from abrupt phase jumps. This general case with complex sample rate $f_s=\kappa B$ (where $\kappa>1$) does not explicitly change the formulation of (2.29), however, the desired impulse-response $\mathbf{e}$ inherently requires full bandwidth utilization. As such, (2.31) naturally trends the MMF to an inverted spectral shape of $\mathbf{s}$ which can induce significant MML since it is the diametric opposite of the MF. Additionally, the resulting waveform-filter response is extremely sensitive to range straddling as it effectively amplifies the high frequency components that vary the most as a function

of sub-sample shift γ (a time-shift corresponds to a multiplication with a sinusoid in the frequency domain).

To illustrate this effect, a LFM waveform with a pulse duration of 2μs and bandwidth of 33.3MHz was generated at a sample rate of 10GHz. A FIR anti-aliasing filter was then applied followed by decimation to a yield simulated loopback configuration with spectral oversampling factors of $\kappa = 1.5$ and $\kappa = 6$. The matched filter was then generated with the $\gamma = 0$ loopback waveform and applied to generate the filter responses shown in Figure 2.7 with $\kappa = 1.5$ and $\kappa = 6$ corresponding to the left and right plots respectively. Here the black 'x's correspond to the "on-grid" response ($\gamma = 0$) while the blue trace includes the continuum of off-grid shifts to better illustrate the filter performance in practice (range-straddling is unavoidable and will always be present). The respective MF power spectral density (PS) is likewise provided in 2.8 where the enforced spectral containment of the bandpass filter is visible (high frequency components go to zero). Next, Akroyd's LS MMF was generated using the unstraddled waveforms ($\gamma = 0$) with a filter length of $D = 4N$. The resulting solution was then normalization as

$$\mathbf{w}_{\text{norm}} = \frac{1}{\sqrt{\mathbf{w}^H\mathbf{w}}\sqrt{\mathbf{s}^H\mathbf{s}}}\mathbf{w} \tag{2.33}$$

such that the resulting filter response will have a peak corresponding to the filter MML. Note, this filter normalization will always be used for the LS MMF approaches unless otherwise stated. The filter responses and filter PS for Akroyd's LS MMFs are then shown in Figures 2.9 and 2.10 respectively where the inverted spectral shape causes upwards of 45dB MML. Additionally, while the on grid-response of Akroyd's approach provided relatively low sidelobes for the on-grid samples, the off-grid continuum induces a sinc-like ringing and high sidelobes due to the significant contribution of the high-frequency components of the filter PS.

To mitigate this significant MML and ensure invertibility of $\mathbf{T}$, it is common practice to incorporate some form of filter regularization such as

$$\sigma^2\|\mathbf{w}\|_2^2 = \mathbf{w}^H\left(\sigma^2\mathbf{I}\right)\mathbf{w}$$

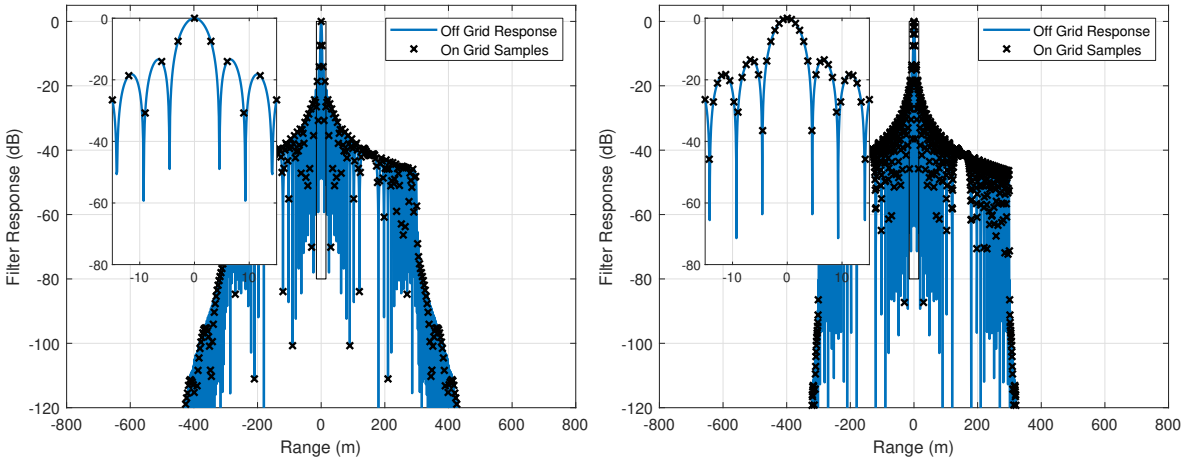

Figure 2.7: Simulated Loopback MF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

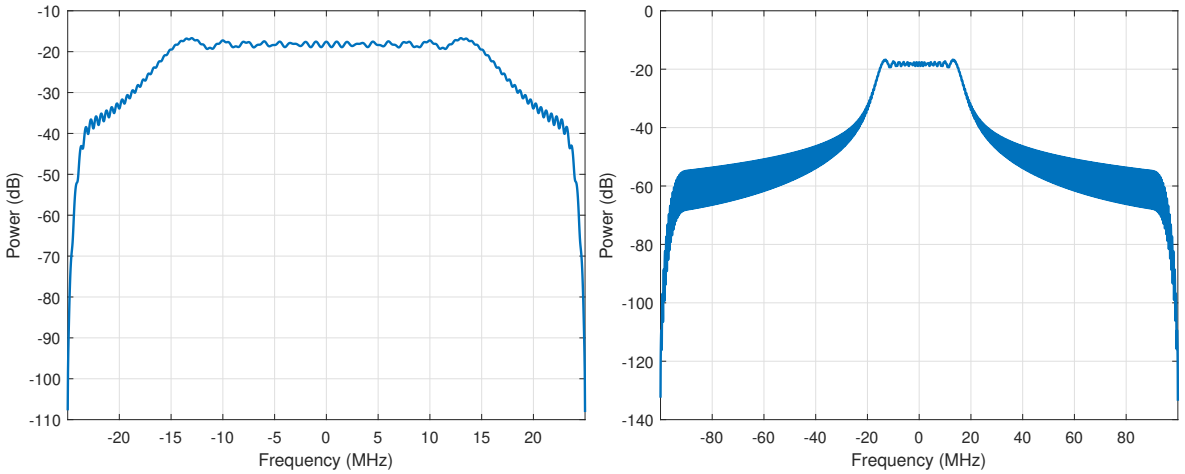

Figure 2.8: Simulated Loopback MF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

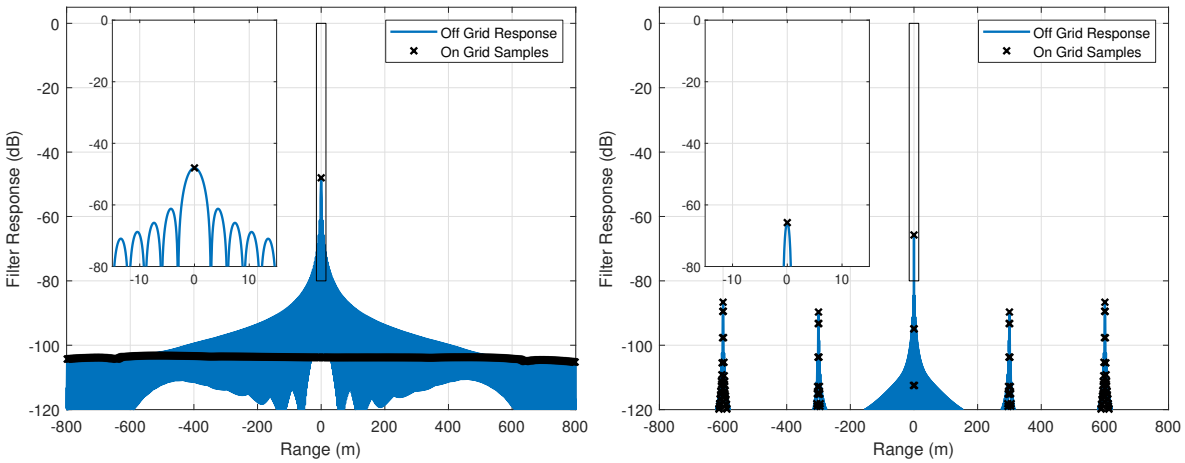

Figure 2.9: Simulated Loopback Akroyd's LS MMF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

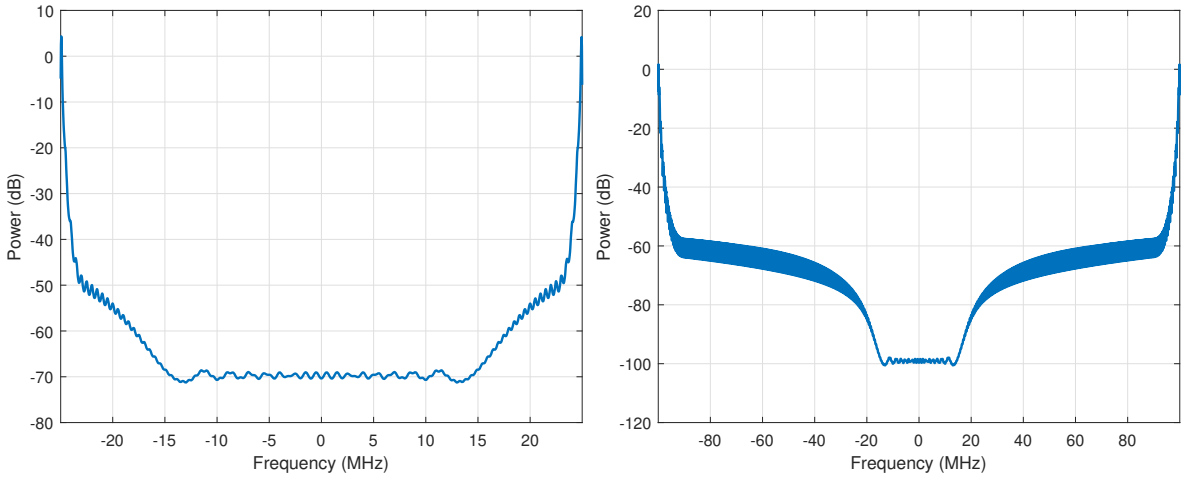

Figure 2.10: Simulated Loopback Akroyd's LS MMF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

(see [56] for alternative regularizers) to (2.30) yielding

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H \left(\mathbf{T} + \sigma^2 \mathbf{I}\right) \mathbf{w} - 2\,\mathrm{Re}\left\{\mathbf{w}^H \mathbf{b}\right\} \tag{2.34}$$

where $\left(\mathbf{T} + \sigma^2 \mathbf{I}\right) = \left(\mathbf{S}\mathbf{S}^H + \sigma^2 \mathbf{I}\right)$ is guaranteed to be PD and thus invertible when $\sigma^2 > 0$ (though it is not necessarily well conditioned or numerically stable). This form of regularization weighting directly controls the filter noise gain as $\sigma^2 \to \infty$ will asymptotically trend the MMF towards a scaled MF since for a large σ^2

$$\mathbf{T} + \sigma^2 \mathbf{I} \approx \sigma^2 \mathbf{I};$$

thus,

$$\left(\mathbf{T} + \sigma^2 \mathbf{I}\right)^{-1} \mathbf{S}\mathbf{e} \approx \frac{1}{\sigma^2} \mathbf{S}\mathbf{e} = \frac{1}{\sigma^2} \mathbf{v}$$

where $\mathbf{v}$ is the zero-padded waveform from (2.24).

While regularization does help with numeric stability, heuristic tuning of σ^2 to yield a desired MML can be a computational cumbersome. Additionally, setting the desired response to an impulse often results in either poor sidelobe suppression or large MML when $\kappa > 1$ as degrees of freedom are expended in an attempt to produce a super-resolved response relative to the nominal resolution. To avoid super-resolution induced MML, [73] proposed replacing the elementary vec-

tor in (2.29) with an alternative desired response $\mathbf{p}$ that has lower (nonzero) range-sidelobes than the MF but retains the nominal resolution (same mainlobe width or 3dB bandwidth). Therefore, the objective function from (2.29) can be modified as

$$\left\|\mathbf{S}^H\mathbf{w}-\mathbf{p}\right\|_2^2. \tag{2.35}$$

The corresponding MMF for an arbitrary $(D+N-1)\times 1$ desired response vector $\mathbf{p}$ can be determined using the same QP form as (2.34) where now $\mathbf{b}=\mathbf{Sp}$. While regularization may still be necessary to ensure invertibility, the resulting MMF can often achieve lower sidelobes for less MML when $\mathbf{p}$ is appropriately selected. However, as the desired response is no longer an impulse, increasing $\sigma^2\to\infty$ will not trend the MMF to the MF as $\mathbf{b}\neq\mathbf{v}$ and the best case MML$\neq 0$ in general.

Using the same simulated setup as before, template-based LS MMFs were generated for the unstraddled LFM waveform using $D=4N$, $\sigma=\frac{10^{-3}}{D}\mathrm{Tr}\{\mathbf{SS}^H\}$ (where $\mathrm{Tr}\{\cdot\}$ denotes the trace of the matrix), and $\mathbf{p}$ corresponding to a response with a super-Gaussian shape of order 6 in the frequency domain which has the same 3dB bandwidth as the LFM waveform but lower sidelobes. The corresponding filter responses and PSs are then provided in Figures 2.11 and 2.12 respectively. Here the LS MMFs are fairly successful at lowering range sidelobes for a much improved $\sim 1dB$ MML. Additionally, the choice of template provides spectral containment to LS MMF PS and as such the off-grid response is fairly consistent with the on-grid samples.

2.2.2 Integrated Sidelobe Level LS MMFs

Another common MMF approach designs QPs with various weightings of the sidelobe power and the mainlobe power [55, 56, 74, 75] via an integrated sidelobe level (ISL) metric that is fundamentally analogous to Akroyd's philosophy when $\kappa=1$. A regularized ISL QP can be posed

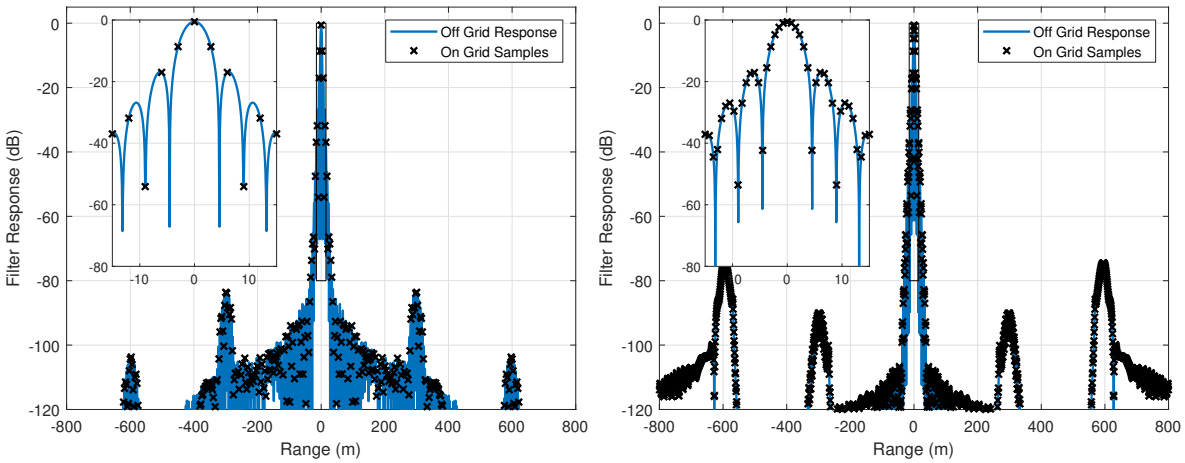

Figure 2.11: Simulated Loopback Template-based LS MMF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

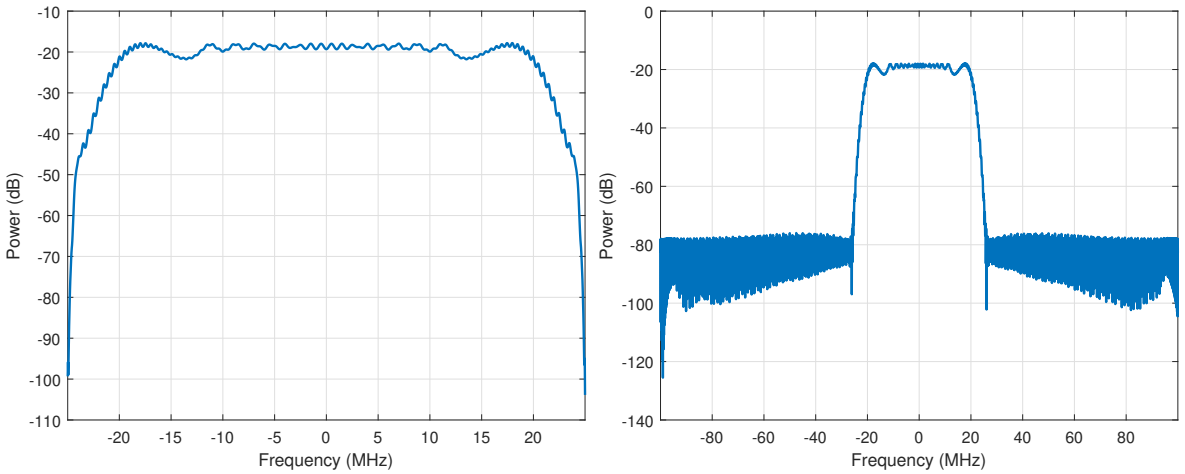

Figure 2.12: Simulated Loopback Template-based LS MMF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

as

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H \left(\mathbf{T} + \sigma^2 \mathbf{I}\right) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{v} = 1 \end{aligned} \tag{2.36}$$

where the linear constraint $\mathbf{w}^H \mathbf{v} = \mathbf{w}^H \mathbf{S}^H \mathbf{e}$ enforces that the zero-delay filter response is unity and prevents the trivial $\mathbf{w} = \mathbf{0}$ solution.

In this base form, the QP is functionally equivalent to (2.34) and the resulting MMF will again suffer from the super-resolution drawbacks. To alleviate MML and provide filter nominal resolution flexibility, [55] suggests a mainlobe "beamspoiling" approach where rows were excised from the convolution matrix $\mathbf{S}$ corresponding to the desired mainlobe width (though there are other approaches to mitigate super resolution such as the mainlobe broadening factor in [75]).

For fixed desired resolution of k samples, beamspoiling is incorporated into (2.36) via a rank-k downdate of the Toeplitz matrix as

$$\tilde{\mathbf{T}} = \mathbf{T} - \tilde{\mathbf{S}} \tilde{\mathbf{S}}^H \tag{2.37}$$

where $\tilde{\mathbf{T}}$ retains a PSD Hermitian structure as $\tilde{\mathbf{S}}$ contains the k columns of $\mathbf{S}$ that correspond to the desired mainlobe location and width. Note, k does not necessarily need to retain the nominal resolution of the MF and, much like tapering of an LFM MF, a broadened filter responds can greatly improve sidelobe reduction at the cost of a degraded range resolution [45, p. 61]. Additionally, the beamspoiling indices can include the zero delay index, though doing so should not significantly alter the optimal filter as the constraint vector will effectively negate the matrix contribution from that direction regardless. Thus, the beamspoiled ISL QP becomes

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H \left(\tilde{\mathbf{T}} + \sigma^2 \mathbf{I}\right) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{v} = 1 \end{aligned} \tag{2.38}$$

where the Lagrangian

$$\mathscr{L}\{\mathbf{w}, \mu\} = \mathbf{w}^H \left(\tilde{\mathbf{T}} + \sigma^2 \mathbf{I}\right) \mathbf{w} + 2\,\mathrm{Re}\left\{\mu \left[\mathbf{w}^H \mathbf{v} - 1\right]\right\} \tag{2.39}$$

can be equivalently expressed as

$$\mathscr{L}\{\mathbf{z}\} = \mathbf{z}^H \mathbf{D}\mathbf{z} - 2\,\mathrm{Re}\left\{\mathbf{z}^H \tilde{\mathbf{b}}\right\} \tag{2.40}$$

where

$$\mathbf{z} = \begin{bmatrix} \mathbf{w} \\ \mu \end{bmatrix},$$

$$\mathbf{D} = \begin{bmatrix} \tilde{\mathbf{T}} + \sigma^2 \mathbf{I} & \mathbf{v} \\ \mathbf{v}^H & 0 \end{bmatrix},$$

and

$$\tilde{\mathbf{b}} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix},$$

are formed by concatenating the MMF $\mathbf{w}$ and Lagrange multiplier μ into a single vector.

The optimal $\mathbf{w}$ and μ correspond to the stationary point of the Lagrangian, which, for the form of (2.40), can be solved directly by taking the gradient w.r.t. $\mathbf{z}^*$ and solving the linear system of equations

$$\mathbf{D}\mathbf{z}_\star = \tilde{\mathbf{b}} \tag{2.41}$$

where $\mathbf{z}_\star$ satisfies the Karush–Kuhn–Tucker (KKT) conditions [76, p. 454]. Note, $\mathbf{D}$ is strictly not PD, and will in fact has a single negative eigenvalue. Regardless, a solution will exists as the system is not singular when $\tilde{\mathbf{T}} + \sigma^2 \mathbf{I}$ is PD and can be solved using the identity in Appendix A.2 and is expressed as

$$\mathbf{w}_\star = \frac{1}{\mathbf{v}^H \left(\tilde{\mathbf{T}} + \sigma^2 \mathbf{I}\right)^{-1} \mathbf{v}} \left(\tilde{\mathbf{T}} + \sigma^2 \mathbf{I}\right)^{-1} \mathbf{v}. \tag{2.42}$$

Again using the simulated loopback setup, ISL LS MMFs were generated for the unstraddled LFM waveform using $D = 4N$, $\sigma^2 = \frac{10^{-3}}{D}\mathrm{Tr}\{\mathbf{S}\mathbf{S}^H\}$ (where $\mathrm{Tr}\{\cdot\}$ denotes the trace of the matrix), and $k = 2ceil\{\kappa\}$ corresponds to the off-zero delays corresponding to the MF nominal resolution broadened by a sample. The respective filter responses and PSs are then provided in Figures

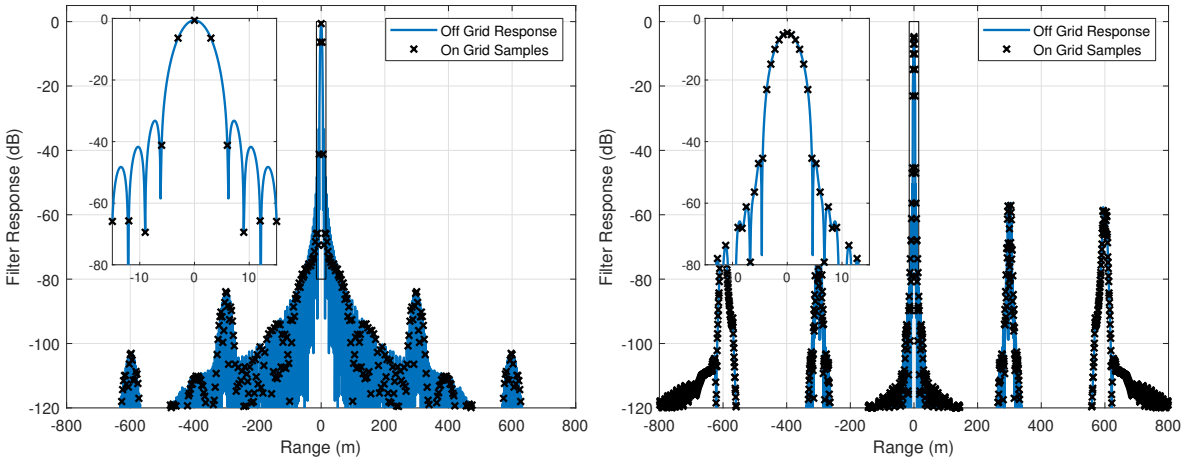

Figure 2.13: Simulated Loopback ISL LS MMF Response for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

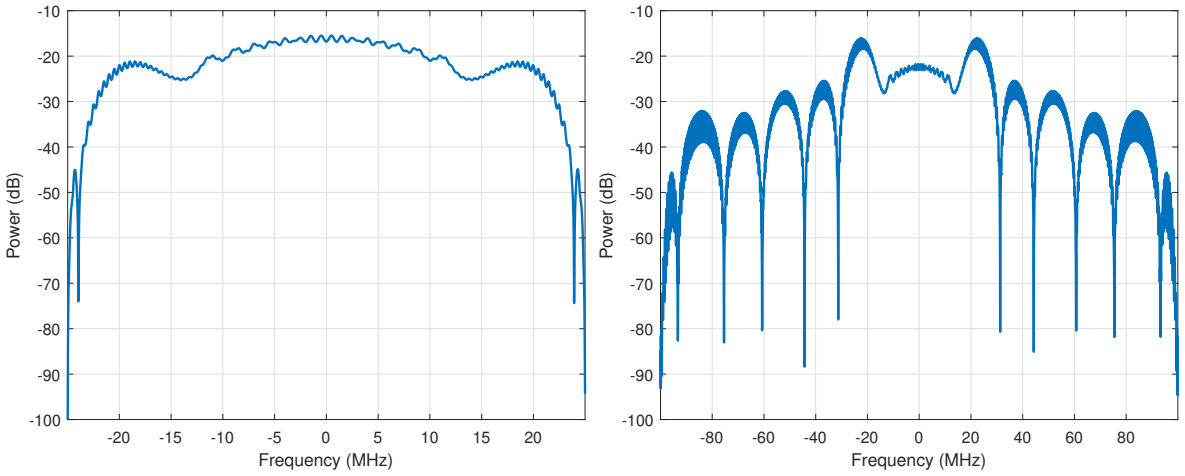

Figure 2.14: Simulated Loopback ISL LS MMF PS for LFM waveform with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

2.13 and 2.14. Here the both the LS MMFs successful successfully suppress sidelobes and have significantly less MML than Akroyd's approach. However, the $\kappa = 1.5$ case has effectively placed the on-grid response on the nulls of the continuous, meaning any straddling will yield significant degradation to sidelobes suppression for close in ranges. On the other hand, the $\kappa = 6$ case is more robust to straddling but has significantly more mismatch loss for the same relative contribution of the regularization (both scaled by the matrix trace and filter length). Examining the filter PSs we see that this form of beamspoiling induces a rippling effect to the spectrum, likely in relation to the Fourier pair between the quasi-rectangular beamspoiling operation and a sinc-function.

This straddling sensitivity was examined in [77] and the author showed consistent results with those illustrated here where template-based MMFs (with an appropriate desired response) are generally more robust to straddling than the ISL MMFs. In [55], the authors propose some heuristic approaches to enhance robustness towards straddling. In particular, they suggest an averaging approach using multiple copies of the sub-sample shifted waveform. However, doing so would require collection of multiple loopback waveform replicas and solving additional linear systems. This may be a reasonable approach for a repeated waveform CPI, but the increased computational burden of collecting the multiple loopback collects and solving many additional linear systems approaches infeasiblilty when the CPI is non-repeating and agile. Additionally, this averaging approach is no longer optimal with respect to the recieve model/ISL QP. In Section 3.1, an alternative range straddling compensation approach is proposed and assessed for the ISL LS MMF which does not require multiple loopback collects and retains optimality with respect to the recieve model and ISL QP.

2.2.3 Controlling Mismatch Loss via Quadratic Constraints

While regularization may be incorporated to both the Template and ISL optimization to adjust MML and ensure the QP is convex, tuning of σ^2 is often heuristic, computationally expensive, and the resulting MMF may lack explicit optimality with respect to the original objective function when using the template-based form. Alternatively, [72, 74, 75, 78] examined incorporating a quadratic inequality constraint (QC) of the form

$$\mathbf{w}^H \mathbf{C} \mathbf{w} + a \leq 0 \tag{2.43}$$

to control MML, while [79] proposed a similar approach to constrain MMF ISL or peak-sidelobe level. In general, a single QC in a QP can yield a tractable optimization problem with a global optimal solution, as long as there exists some $\mathbf{w}$ that satisfies the constraints (even if $\mathbf{T}$ and $\mathbf{C}$ are not PD [80]). Thus, a QCQP framework provides MMF design flexibility while ensuring

that the resulting filter is globally optimal. However, brute force solutions to the QCQP can be computationally intensive for even modest D. Both previous works proposed techniques to combat computational burden. [74] elected to perform a costly initial calculation to reduce the cost of subsequent operations. [75] examined a Circulant MMF formulation which greatly reduced the computational cost, but required an infinite impulse response (IIR) approximation to an FIR filter and is not applicable all modes of operation [81].

In Section 7, the work of [74, 75] is generalized for the template-based objective function and further extensions are developed to the QC such that arbitrary signal-to-interference-plus-noise (SINR) loss constraints can be incorporated to both MMF forms. Furthermore, the structure of $\mathbf{T}$ and $\mathbf{C}$ is then leveraged to develop an efficient solver for both the traditional Hermitian Toeplitz QP [74] and the approximate form of [75]. Hardware-in-the-loop and open-air experimental evaluations are then provided.

2.3 Adaptive Range Processing Overview

While LS MMFs can help mitigate range-sidelobes for a modest MML if designed appropriately, using a single filter for all delays effectively weights contributions from all scatterers uniformly and therefore wastes DoF over-suppressing sidelobes from directions where there is low-power scattering. Alternatively, adaptive signal processing techniques attempt to dynamically allocate DoFs to optimally estimate the scene. This results in a unique filter for each delay which improves sidelobe-suppression for the less MML, but comes with an increased computational burden. This process can be performed by solving for the optimal filters that satisfy a minimum mean square error (MMSE) objective function [82, p. 96], though direct "filterless" adaptive estimators can likewise be posed [83, p. 259].

However, classic adaptive approaches require characterization of the receive statistics, which is not generally known a priori. In regimes where the receive model is approximately linear time invariant (LTI) and wide sense stationary (WSS), measured snapshots can be used to form a sample covariance matrix (SCM) that approximates the receive statistics and can be utilized for MMSE

processing [82, p. 359]. Unfortunately, forming a SCM for the range domain proves challenging as scattering statistics will not generally be LTI or WSS with respect to fast-time delay, while movers and waveform agility degrade SCM estimation across slow-time. Though, it is possible to form fast-time sample covariance matrices in some circumstances [84].

Alternatively, the structure of receive model can be leveraged to approximated the idealized SCM with a structured covariance matrix. This approximate form significantly reduces the required number of snapshots (typically only a single snapshot is required). However, as with the LS MMF leveraging a structured form fundamentally relies on the accuracy of a discrete representation of a continuous reality. As such, robustness measures must be incorporated to reduce sensitivity towards practical fidelity limitations.

2.4 Adaptive Pulse Compression

Standard adaptive processing is seeks to determine an optimal filter that minimizes the mean square value of some statistical error metric e as

$$J = \mathscr{E}\left\{\|e\|_2^2\right\} \tag{2.44}$$

where $\mathscr{E}\{\cdot\}$ denotes the expectation operator. The adaptive pulse compression (APC) [57, 58] algorithm leverages this framework to pose an MMSE objective function to solved for a $D \times 1$filter $\mathbf{w}(\ell)$ at each delay ℓ in the linear model of (2.18) as

$$\begin{aligned} J &= \mathscr{E}\left\{\left\|x(\ell) - \mathbf{w}^H(\ell)\,\mathbf{y}(\ell)\right\|_2^2\right\} \\ &= \rho(\ell) - 2\,\mathrm{Re}\left(\mathbf{w}^H(\ell)\,\mathbf{r}_{\mathbf{y}x}(\ell)\right) + \mathbf{w}^H(\ell)\,\mathbf{R}_{\mathbf{yy}}(\ell)\,\mathbf{w}(\ell) \end{aligned} \tag{2.45}$$

with

$$\mathbf{r}_{\mathbf{y}x}(\ell) = \mathscr{E}\{\mathbf{y}(\ell)x^*(\ell)\} \tag{2.46}$$

corresponding to the $D \times 1$ measurement-scatterer cross-correlation vector,

$$\rho(\ell) = \mathscr{E}\{x(\ell)x^*(\ell)\} \tag{2.47}$$

the scatterer-scatterer correlation, and

$$\mathbf{R}_{\mathbf{yy}}(\ell) = \mathscr{E}\left\{\mathbf{y}(\ell)\mathbf{y}^H(\ell)\right\} \tag{2.48}$$

the $D \times D$ measurement-measurement correlation matrix respectively[1].

Similar to (2.30), the filter that minimizes (2.45) can be solved with a QP as

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H\mathbf{R}_{\mathbf{yy}}\mathbf{w} - 2\,\mathrm{Re}\left\{\mathbf{w}^H\mathbf{r}_{\mathbf{y}x}\right\} \tag{2.49}$$

where the dependency on ℓ has been dropped with vectors and matrices for notional convenience throughout the remainder of this chapter unless explicitly stated. When $\mathbf{R}_{\mathbf{yy}}$ is PD, (2.49) has the optimal solution of

$$\mathbf{w} = \mathbf{R}_{\mathbf{yy}}^{-1}\mathbf{r}_{\mathbf{y}x}, \tag{2.50}$$

however, unlike with the LS MMFs, $\mathbf{R}_{\mathbf{yy}}$ and $\mathbf{r}_{\mathbf{y}x}$ are not deterministic or known a priori.

Where classical adaptive techniques would attempt to estimate $\mathbf{R}_{\mathbf{yy}}$ and $\mathbf{r}_{\mathbf{y}x}$ through the accumulation of measurement snapshots, APC leverages the structure of the linear receive model (2.18)

$$\mathbf{y} = \mathbf{S}\mathbf{x} + \mathbf{n}$$

to form the structured-covariance matrix

$$\mathbf{R}_{\mathbf{yy}} = \mathbf{S}\mathscr{E}\left\{\mathbf{x}\mathbf{x}^H\right\}\mathbf{S}^H + \mathscr{E}\left\{\mathbf{n}\mathbf{n}^H\right\}$$

[1]Note, the original derivation of the algorithm strictly used $D = N$. Here for generality, an arbitrary filter length is used where $D > N$ increases the available DoF.

and cross-correlation vector

$$\mathbf{r}_{\mathbf{y}x} = \mathbf{S}\mathscr{E}\left\{\mathbf{x}x^*\left(\ell\right)\right\}$$

where the noise and scatterers are treated as zero-mean random variables independent from one another. As the noise statistics can often be estimated a priori using the system noise figure or measured with "quiet" data (i.e. no active TX), the noise covariance matrix is assumed to be known as

$$\mathscr{E}\left\{\mathbf{n}\mathbf{n}^H\right\} = \mathbf{R}_\mathbf{n} \tag{2.51}$$

which for zero-mean white Gaussian system noise with a variance of $\sigma_\mathbf{n}^2$ is delay invariant and defined as

$$\mathbf{R}_\mathbf{n} = \sigma_\mathbf{n}^2\mathbf{I}.$$

Next, APC imposes a structure treats the scattering coefficients as independent variables to yield a $(D+N-1)\times(D+N-1)$ diagonal source correlation matrix

$$\mathscr{E}\left\{\mathbf{x}\mathbf{x}^H\right\} = \left(\mathbf{x}\mathbf{x}^H\right)\odot\mathbf{I} = \mathbf{P} \tag{2.52}$$

where

$$\mathbf{P}_{i,i} = \left|x\left(\ell-D+i\right)\right|^2 = \rho\left(\ell-D+i\right), \tag{2.53}$$

and

$$\mathscr{E}\left\{\mathbf{x}x^*\left(\ell\right)\right\} = \rho\left(\ell\right)\mathbf{e}, \tag{2.54}$$

thus, the structured covariance is expressed as

$$\begin{aligned}\mathbf{R}_{\mathbf{yy}} &= \mathbf{R}_\mathbf{s} + \mathbf{R}_\mathbf{n} \\ &= \mathbf{S}\mathbf{P}\mathbf{S}^H + \sigma_\mathbf{n}^2\mathbf{I}\end{aligned} \tag{2.55}$$

and structured cross-correlation vector

$$\mathbf{r}_{\mathbf{y}x} = \mathbf{S}\mathscr{E}\left\{\mathbf{x}x^{*}\left(\ell\right)\right\} = \rho\left(\ell\right)\mathbf{S}\mathbf{e} = \rho\left(\ell\right)\mathbf{v} \tag{2.56}$$

where $\mathbf{e}$ is the $D+N-1\times 1$ elementary vector of all zeros except for unity in the center element[2].

These derivations provide a framework for the structured covariance, but as $\mathbf{P}$ corresponds to the scattering coefficient powers APC is trying to estimate, it is inherently unknown. To draw a quick parallel, the LS MMF approaches can be derived from this perspective by approximating the scattering coefficients as constant and delay invariant, i.e.

$$\mathbf{P} = \sigma_{\mathbf{x}}^{2}\mathbf{I}$$

such that

$$\mathbf{R}_{\mathbf{yy}} = \mathbf{S}\mathbf{P}\mathbf{S}^{H} + \sigma_{\mathbf{n}}^{2}\mathbf{I} = \sigma_{\mathbf{x}}^{2}\mathbf{S}\mathbf{S}^{H} + \sigma_{\mathbf{n}}^{2}\mathbf{I}$$

where the relative scaling of $\sigma_{\mathbf{x}}^{2}$ and $\sigma_{\mathbf{n}}^{2}$ corresponds to the degree of regularization. The MMF approaches then attempt control the contribution from adjacent scattering via design of the filter response using (2.28) ($\mathbf{r}^{T} = \mathbf{w}^{H}\mathbf{S}$). This allows a single filter to be applied for all range, but performance will be strictly worse than a true MMSE estimate (unless the scattering statistics are truly range invariant). Rather than a delay invariance assumption, APC leverages a reiterative minimum mean square error (RMMSE) framework [57, 88] in which estimates for the scattering coefficients and structured covariance matrix are iteratively refined. From an algorithmic perspective, RMMSE effectively takes an adaptive nulling approach where the current source estimates are used to place nulls in the presumed direction high powered scattering.

[2]Note, for the previously established APC formulation $D = N$ and thus $\mathbf{e}$ is a $2N-1\times 1$ vector. Likewise, $\mathbf{v} = \mathbf{S}\mathbf{e} = \mathbf{s}$ is the form presented in [57, 58, 85, 86]. As such, forming the APC filters near the delay boundaries will necessitate truncation of $\mathbf{v}$ and $\mathbf{R}_{\mathbf{yy}}$ in a manner similar to [87]. Here the mathematically equivalent notation from (2.24) is used for consistency to retain the generalized snapshot length representation of (2.13).

For the initial iteration, APC solves a QP

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H \left(\mathbf{S}\mathbf{P}^{(0)}\mathbf{S}^H + \sigma_{\mathbf{n}}^2 \mathbf{I} \right) \mathbf{w} - 2\,\mathrm{Re}\left\{ \rho^{(0)}(\ell)\, \mathbf{w}^H \mathbf{v} \right\} \tag{2.57}$$

for every delay index ℓ where

$$\rho^{(0)}(\ell) = \left| \hat{x}^{(0)}(\ell) \right|^2$$

is formed via the MF (2.23) estimate in (2.25) and is likewise used to form $\mathbf{P}^{(0)}$ via (2.53). As $\mathbf{P}^{(0)}$ is strictly positive and $\sigma_{\mathbf{n}}^2 > 0$, the hessian of (2.57) is strictly PD and thus

$$\mathbf{w} = \rho^{(0)}(\ell) \left(\mathbf{S}\mathbf{P}^{(0)}\mathbf{S}^H + \sigma_{\mathbf{n}}^2 \mathbf{I} \right)^{-1} \mathbf{v}. \tag{2.58}$$

Solving and applying (2.58) for every delay index is then used to update the estimate scattering coefficients[3]

$$\hat{x}^{(1)}(\ell) = \mathbf{w}(\ell)^H \mathbf{y}(\ell),$$

scattering powers

$$\rho^{(1)}(\ell) = \left| \hat{x}^{(1)}(\ell) \right|^2,$$

and source correlation matrix $\mathbf{P}^{(1)}$. This process is then repeated for I RMMSE iterations or until convergence where the ith RMMSE QP is posed as

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H \left(\mathbf{S}\mathbf{P}^{(i)}\mathbf{S}^H + \sigma_{\mathbf{n}}^2 \mathbf{I} \right) \mathbf{w} - 2\,\mathrm{Re}\left\{ \rho^{(i)}(\ell)\, \mathbf{w}^H \mathbf{v} \right\} \tag{2.59}$$

which is minimized by

$$\mathbf{w} = \rho^{(i)}(\ell) \left(\mathbf{S}\mathbf{P}^{(i)}\mathbf{S}^H + \sigma_{\mathbf{n}}^2 \mathbf{I} \right)^{-1} \mathbf{v}. \tag{2.60}$$

The original RMMSE framework performed well on idealized (simulated) measurements, however, the explicit fitting to a presumed structure often lacks practical robustness as it ignores the

[3]Note, here the snapshot notation for the filter has temporarily returned to emphasize that a unique RMMSE filter must be solved for every delay index.

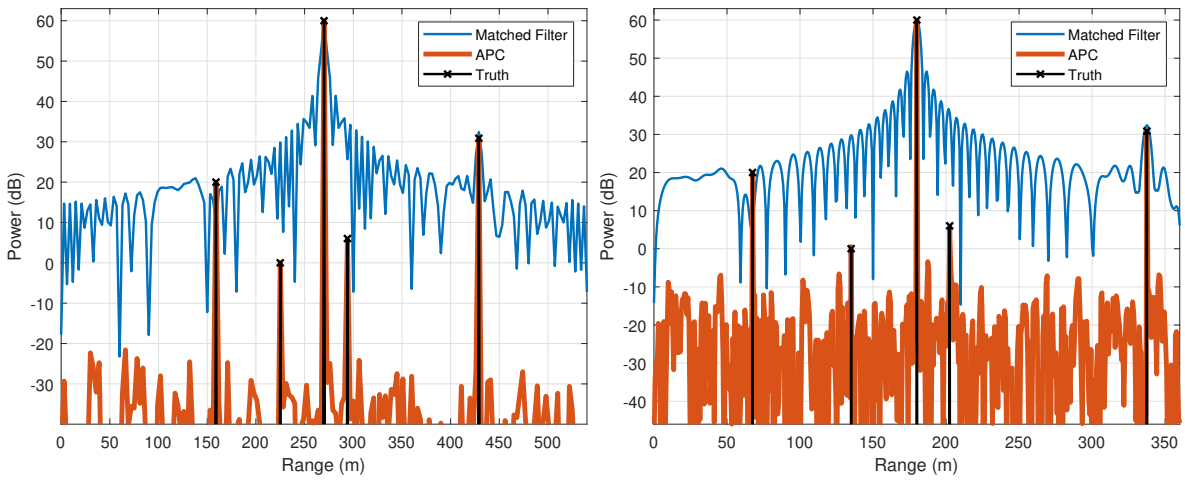

Figure 2.15: Simulated APC estimate for idealized measurements with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

continuous reality where hardware and model errors (e.g. phase noise, amplifier nonlinearities, range straddling, etc.) can lead to over suppression or false peaks. Like to Akroyd's LS MMF, the original form of APC also attempts to provide a super-resolved estimate which induces significant MML and is extremely sensitive to straddling effects. To illustrate these drawbacks, the simulated loopback setup from Section 2.2 was used to generate synthetic receive measurements with known truth values. Five scatterers were placed in the receive measurements and 3 iterations of APC was used to estimate the scene. For the first case, all the scatterers were placed on-grid (meaning no straddling is present) to yield an idealized receive model. The respective APC estimates for $\kappa = 1.5$ and $\kappa = 6$ waveforms are then shown on the left and right side of Figure 2.15 respectively, where the orange trace is the APC estimate, the blue trace is the MF estimate and the black trace is the ground truth. Both approaches were able to isolate all five sources, however, the $\kappa = 6$ case created a significantly higher background floor. Figure 2.16 then illustrates the expected noise power for each range bin (since APC constructs a unique filter for every delay), where clearly the $\kappa = 1.5$ case yields noise power fairly close to that of the MF, while the $\kappa = 6$ case yields noise powers $\sim$15dB higher than the MF due to the super-resolution.

Next, the same scenario is performed, but this time each of the five sources was given a sub-sample shift and the dominant scatter in the middle was selected to be maximally straddled. The

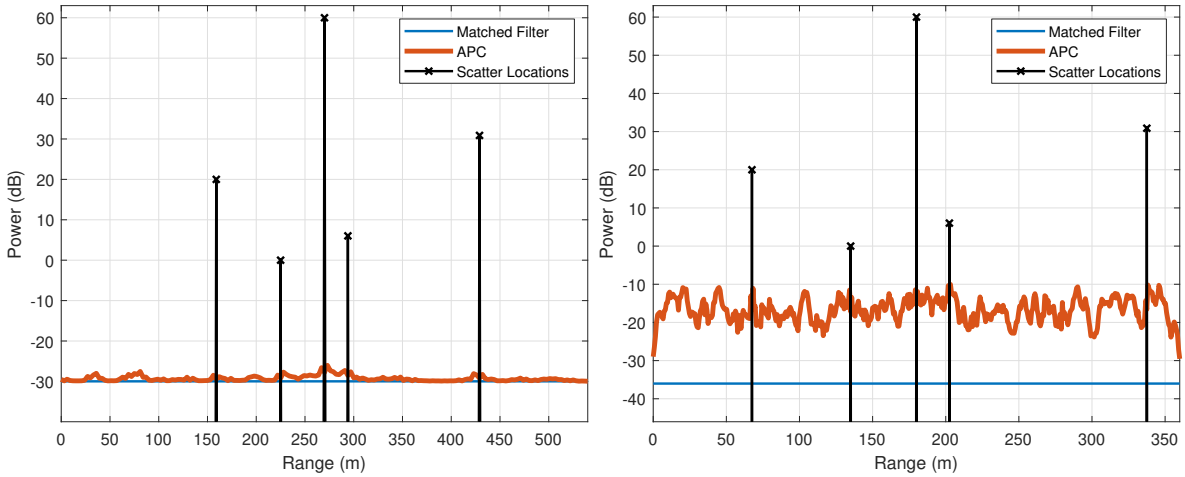

Figure 2.16: Simulated APC noise power for idealized measurements with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

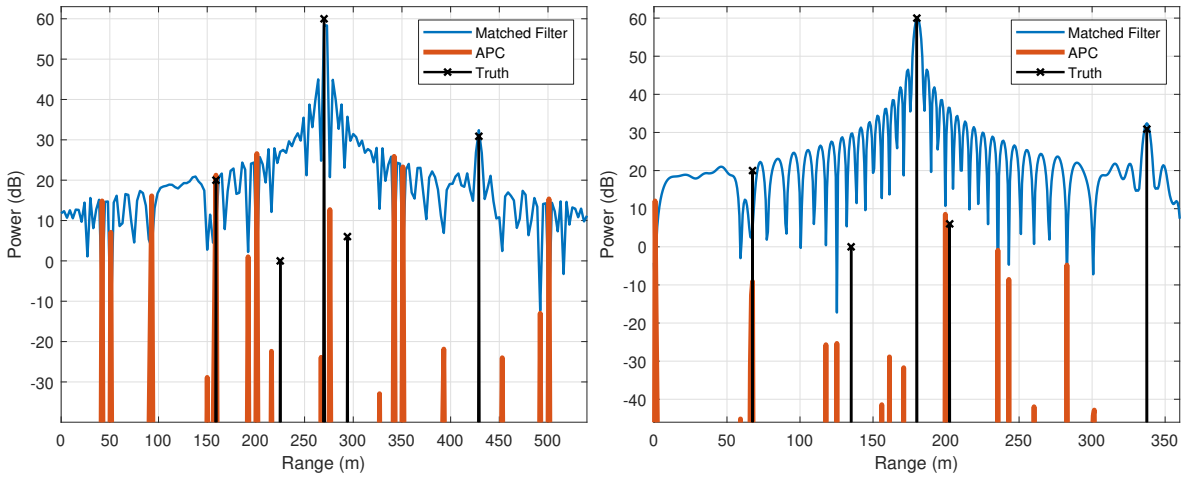

Figure 2.17: Simulated APC estimate with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

APC estimates for this case are shown in Figure 2.17 where for both cases APC performs erratically and is not able to converge to a physically meaningful estimate. Looking at corresponding respective noise powers in Figure 2.18, the noise amplification is significant and the respective estimates likely just correspond to amplified noise.

A myriad of compensation techniques/extensions for RMMSE have been proposed to enhance performance or enable estimation for different domains/applications [55, 89–97], where a robust range estimate typically necessitates incorporation of a gain constraint to avoid oversuppression, a form of beamspoiling to retain the filter nominal resolution, and statistical model uncertainties to

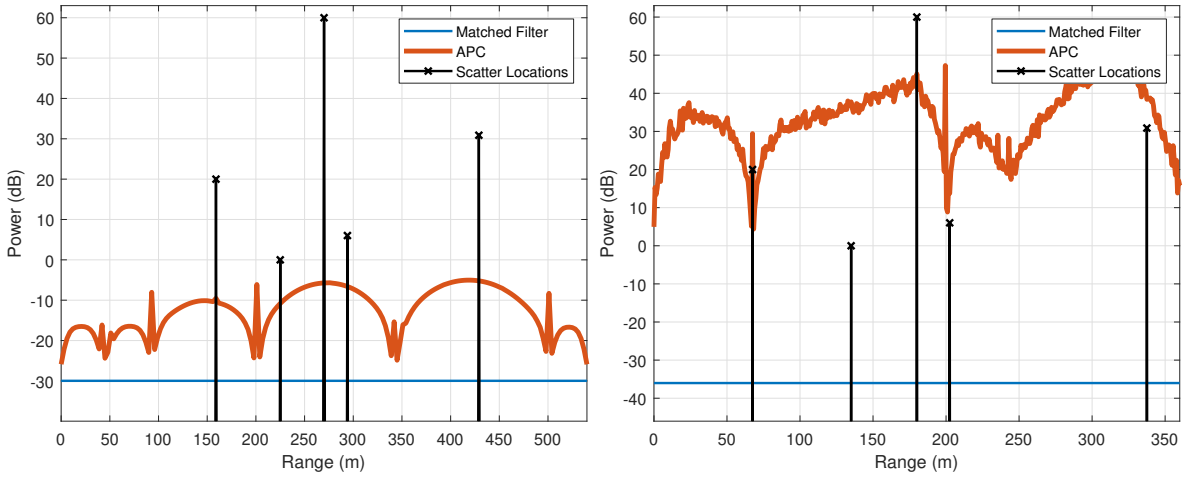

Figure 2.18: Simulated APC noise power with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

avoid amplification of model errors. The remainder of this section will outline these modifications and discuss theoretical efficacy using this simulated receive measurements. Further developments and a computationally efficient approximate implementation will later be presented in Section 6 as well as an open-air experimental demonstration.

2.4.1 APC Generic Model Uncertainty

The idealized model of (2.18) neglects the physical reality where nonlinearities, phase noise, range straddling, fast-time Doppler, etc. are always present. Many of these effects can be reduced with more sophisticated (and expensive) hardware and or careful waveform design, but they will always act as a limiting factor for adaptive processing. Fortunately, the arbitrary nature of the RMMSE objective function allows for the incorporation of model errors to enhance robustness towards these effects [89, 91, 92], though, doing so requires both a statistical model for the error source and an estimate mean/variance.

While individually modeling and incorporating all possible error sources would be ideal, the aggregate model error imposed on the waveform from independent sources (e.g. straddling, amplifier distortion, phase noise) can be coarsely approximated as Gaussian via the central limit theorem [98, p. 241]. Additionally, this sum of model errors can be conveniently measured using hardware-in-the-loop collects, though, distortion due to the antenna and open-air channel would

not be incorporated. This approach suggested in [90] (an extension of [99]), where the linear model of (2.18) can be approximately compensated as

$$\mathbf{y} = [\mathbf{1} + \mathbf{z}] \odot [\mathbf{S}\mathbf{x}]\,\mathbf{n} \tag{2.61}$$

where $\mathbf{z}$ is an additional zero mean complex i.i.d. random variable with variance $\sigma_{\mathbf{z}}^2$ that subsumes the aggregate model error that was measured in loopback. This modified linear model assumes that each receive sample has a multiplicative error term dependent on the measurement strength and effectively places a delay variant floor on the effective dynamic range. The structured covariance is found as

$$\begin{aligned} \mathbf{R}_{\mathbf{yy}} &= \mathbf{R}_{\mathbf{s}} + \mathbf{R}_{\mathbf{z}} + \mathbf{R}_{\mathbf{n}} \\ &= \mathbf{S}\mathbf{P}\mathbf{S}^H + \left(\sigma_{\mathbf{z}}^2\mathbf{I}\right) \odot \left[\mathbf{S}\mathbf{P}\mathbf{S}^H\right] + \sigma_{\mathbf{n}}^2\mathbf{I} \end{aligned} \tag{2.62}$$

where $\mathbf{R}_{\mathbf{z}} = \left(\sigma_{\mathbf{z}}^2\mathbf{I}\right) \odot \left[\mathbf{S}\mathbf{P}\mathbf{S}^H\right]$ is the model uncertainty covariance matrix. The RMMSE QP for each iteration will still have the form

$$\min_{\mathbf{w}^*} \quad \mathbf{w}^H\mathbf{R}_{\mathbf{yy}}\mathbf{w} - 2\,\mathrm{Re}\left\{\rho\left(\ell\right)\mathbf{w}^H\mathbf{v}\right\} \tag{2.63}$$

where the iteration index has been excluded for brevity. As $\mathbf{R}_{\mathbf{s}}$ is PSD and has an all positive diagonal, hessian of this QP is still PD and each filter can be solved via

$$\begin{aligned} \mathbf{w} &= \rho\left(\ell\right)\mathbf{R}_{\mathbf{yy}}^{-1}\mathbf{v} \\ &= \rho\left(\ell\right)\left(\mathbf{S}\mathbf{P}\mathbf{S}^H + \left(\sigma_{\mathbf{z}}^2\mathbf{I}\right) \odot \left[\mathbf{S}\mathbf{P}\mathbf{S}^H\right] + \sigma_{\mathbf{n}}^2\mathbf{I}\right)^{-1}\mathbf{v} \end{aligned} \tag{2.64}$$

where the same RMMSE initialization and refinement procedure is applied as before.

For the simulated setup, $\sigma_{\mathbf{z}}^2$ was set to the variance across waveforms as a function of the sub-sample receive shifts $0.5 \leq \gamma \leq 0.5$. While in practice, measuring the loopback response for different sub-samples shifts can be somewhat impractical, here we do so to provide an illustration of performance when $\sigma_{\mathbf{z}}^2$ is configured perfectly. APC with model uncertainty was then applied to

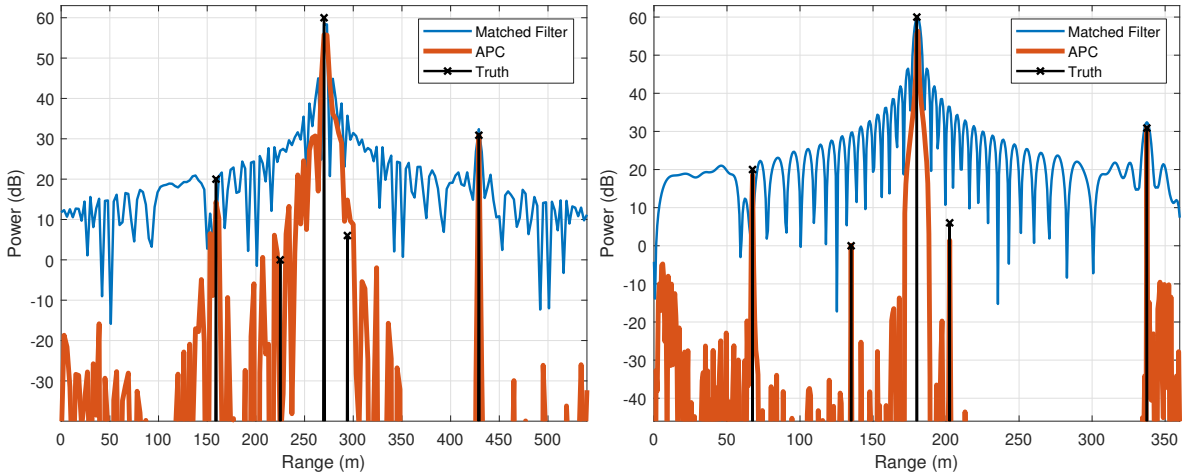

Figure 2.19: Simulated APC estimate with model uncertainty for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

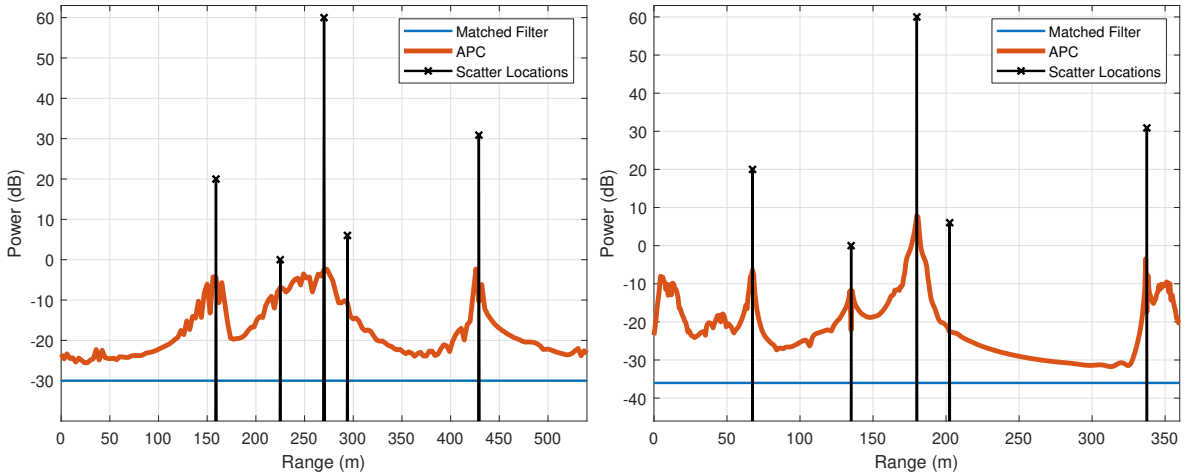

Figure 2.20: Simulated APC noise power with model uncertainty for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

the simulated measurements with range straddling and the corresponding estimates are shown in Figure 2.19, while Figure 2.20 shows the respective range varying noise power. Here APC was able to resolve three of the four signatures for both cases, however, a range smearing effect is visible with the center source due to straddling, which obfuscates the low powered adjacent sources. The noise power for both cases is lower than when model uncertainty was excluded, but the noise gain is still high relative to the MF and a distinct peaking behavior is visible near the scatterer true locations.

2.4.2 APC Covariance Beamspoiling

While incorporating generic model uncertainty enhanced algorithm stability and robustness towards false peaks, it does not directly address super-resolution induced noise gain or error from straddling. The super-resolved estimates in turn induce significant variation of the noise power across range, and limits the local dynamic range as a straddled scatterer will deviate from the expected model. To combat super-resolution, [55] proposed a covariance "beamspoiling" method in an attempt to preserve the nominal resolution of the MF. This was achieved in a similar manner as the LS MMF ISL beamspoiling (2.37) by zeroing out Z elements of the source covariance $\mathbf{P}$ corresponding to the matched filter mainlobe. Note, unlike (2.37) to retain stability the center most element corresponding the ℓth direction and can not be excised[4], e.g. $\mathbf{P}_{z,z} = 0$ for

$$z \in \left(\frac{D-N}{2} - Z, \frac{D-N}{2} + Z \right), \; z \neq N$$

where $Z = ceil\{\kappa\} - 1$. This can be implemented as a rank-$2Z$ downdate if the structured covariance as

$$\mathbf{R_s} = \mathbf{SPS}^H - \tilde{\mathbf{S}}\tilde{\mathbf{P}}\tilde{\mathbf{S}}^H \tag{2.65}$$

where $\tilde{\mathbf{S}}$ and $\tilde{\mathbf{P}}$ corresponds to the respective mainlobe contributions of $\mathbf{S}$ and $\mathbf{P}$ respectively. Outside the zeroed elements in $\mathbf{P}$, the RMMSE framework remains unchanged and thus beamspoiling does not prevent incorporation of generic model uncertainty from (2.63).

Beamspoiling and model uncertainty were incorporated into APC simulation and applied to yield the estimates in Figure 2.19 and noise power in Figure 2.20. Here APC was again able to resolve three of the five signatures for the $\kappa = 1.5$ case, but the center most source sill sees a range smearing from straddling. The $\kappa = 6$ case now able to clearly resolve all 5 sources and does not see the same spiking behavior as Figure 2.19. Additionally, the $\kappa = 6$ case see a modest improvement in overall the noise power.

[4] To be more precise, the center most element cannot be removed unless some form of unity constraint is enforced similar to that of the ISL MMF. Such constraints will be discussed in the subsequent section.

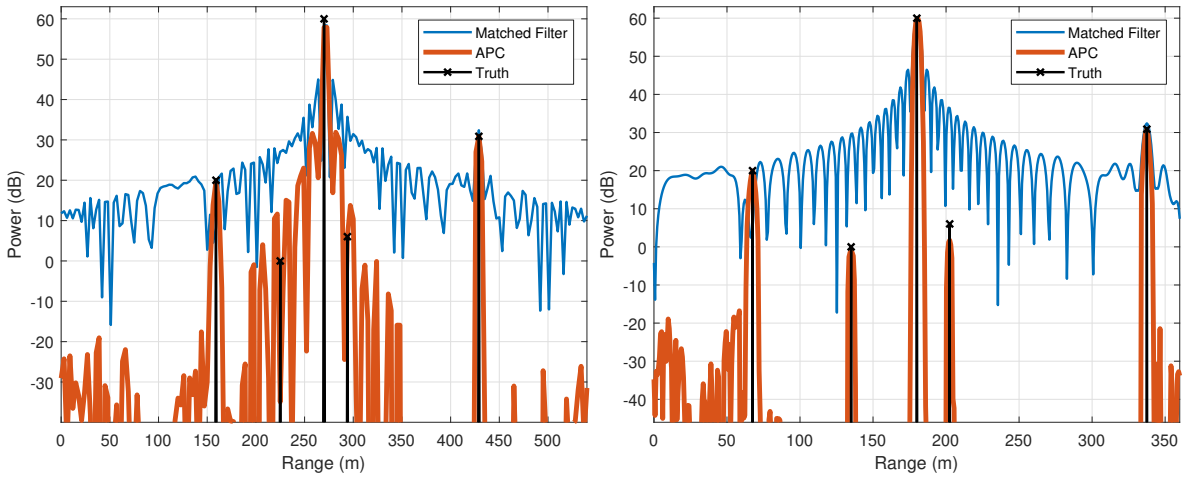

Figure 2.21: Simulated APC estimate with model uncertainty and beamspoiling for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

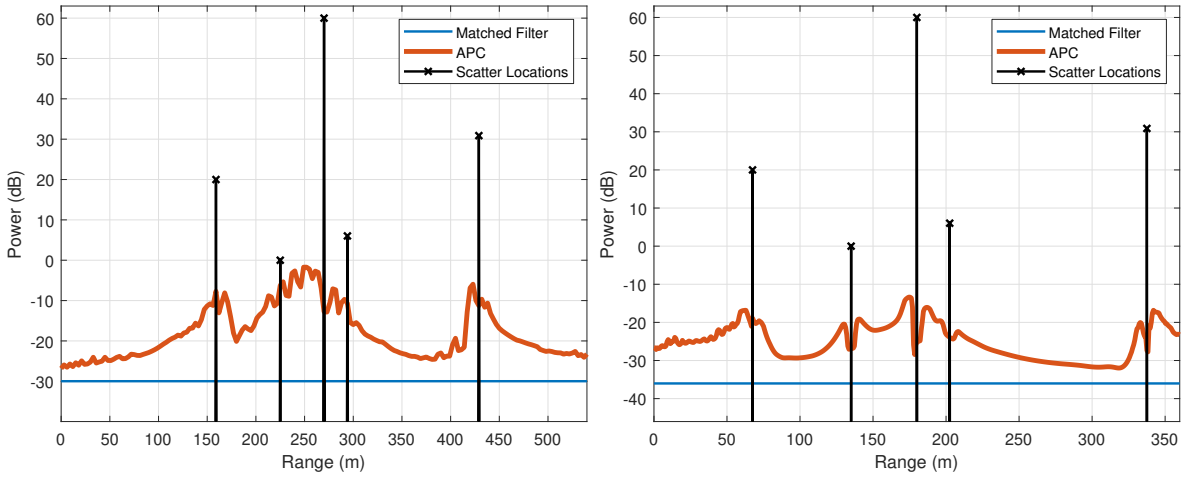

Figure 2.22: Simulated APC noise power with model uncertainty and beamspoiling for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

2.4.3 APC Unity Gain Linear Constraint

While incorporating beamspoiling and model uncertainty enhances robustness towards super-resolution and model errors, RMMSE in general has a tendency to over-suppress low-power scattering and produces erratic peaks below the MF noise floor. For range only detection, this may actually be beneficial as low SNR detections are often thrown out to avoid false alarms. However, subsequent velocity estimation processing will provide a coherent gain and low-SNR signatures may rise above the detection threshold. As such, it is typically desirable to prevent this oversuppression.

To do so, [86] proposed incorporating a unity gain constraint into the RMMSE QP formulation as

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H \mathbf{R}_{\mathbf{yy}} \mathbf{w} - 2\,\mathrm{Re}\left\{\rho\left(\ell\right) \mathbf{w}^H \mathbf{v}\right\} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{v} = 1 \end{aligned} \tag{2.66}$$

where $\mathbf{R}_{\mathbf{yy}}$ can incorporate beamspoiling or model uncertainty if desired. As the only feasible solutions to the QP must satisfy the linear constraint (i.e., $\mathbf{w}^H \mathbf{v} = 1$), (2.66) can be simplified to a CAPON esc form [100] as

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H \mathbf{R}_{\mathbf{yy}} \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{v} = 1 \end{aligned} \tag{2.67}$$

which can again be solved using the KKT identity from Appendix A.2 and yields the optimal solution

$$\mathbf{w} = \frac{1}{\mathbf{v}^H \mathbf{R}_{\mathbf{yy}}^{-1} \mathbf{v}} \mathbf{R}_{\mathbf{yy}}^{-1} \mathbf{v}. \tag{2.68}$$

As this linear constraint is enforced every iteration, gain constrained APC preserves a higher but less erratic estimate floor and will not over-suppress signatures near the noise floor. There exists hybrid constrained RMMSE formulation that leverages a geometric averaging of the unity constrained filter and unconstrained filters and has been empirically shown to enhanced estimation of low SNR signatures in a frequency domain RMMSE context [101]. However, this hybrid or partially constrained approach still may produce erratic behavior below the noise floor which may affect subsequent Doppler estimation.

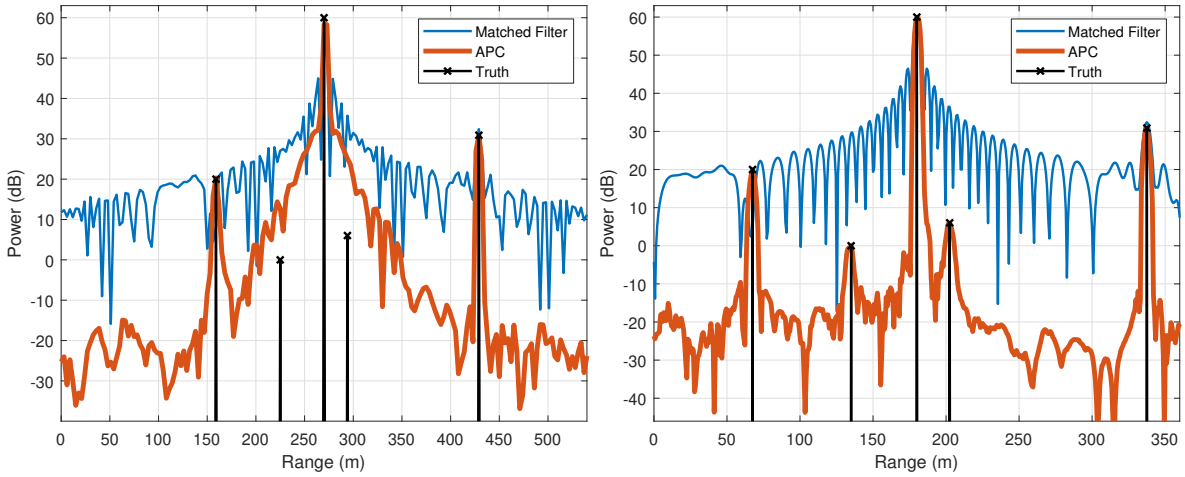

Figure 2.23: Simulated APC estimate with model uncertainty, beamspoiling, and a unity constraint for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

As a final simulated example, the unity gain constraint was incorporated into APC where the beamspoiling and model uncertainty parameters were configured as before. This formulation then was applied to the simulated measurements and produced the estimates in Figure 2.23 with noise powers shown in Figure 2.24. This form of APC performed on par with the beamspoiling and no unity constraint approach, but has a less erratic estimate floor that is proportional to the range varying noise power.

While APC with these incorporated robustness considerations performs well for larger κ, range straddling leads to poor estimation for near critically sampled architectures. In [102], the author explored range-straddling robustness measures for an LFM waveforms via covariance matrix tapering. However, this approach required explicit knowledge of the continuous waveform phase function which is not always possible for arbitrary waveforms, and does not account for RX anti-alias filter induced distortion. In Section 6 general straddling robustness measures for arbitrary waveforms will be developed leveraging hardware-in-the-loop measurements and a structured straddle interpolated signal model.

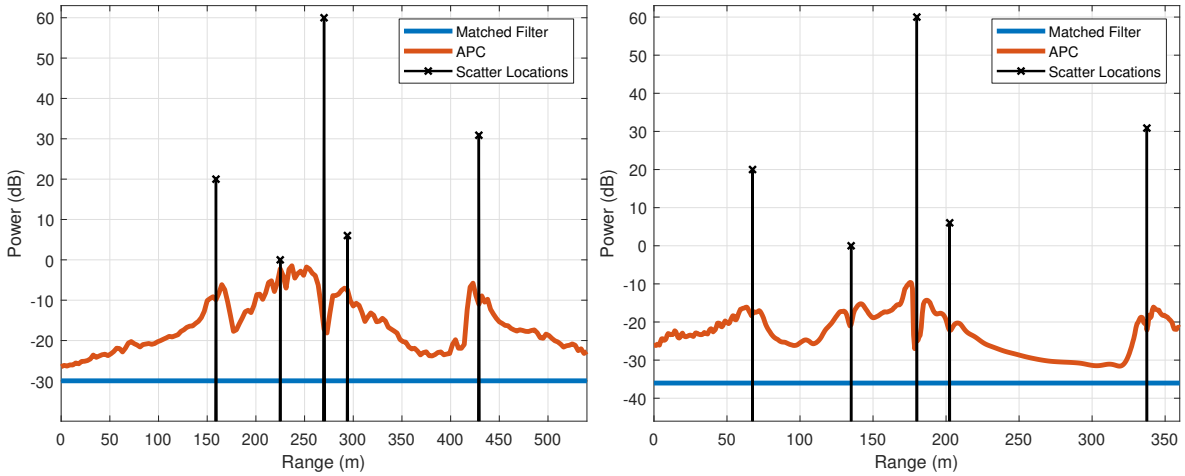

Figure 2.24: Simulated APC noise powers with model uncertainty, beamspoiling, and a unity constraint for $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

2.5 Doppler Processing Overview

For MTI applications, it is necessary to separate moving objects from stationary scatterers, denoted clutter. Radial movement relative to the TX/RX baseline will induce a Doppler shift, thus, knowledge of the TX/RX position and speed can be combined with Doppler measurements to estimate velocity via

$$\omega = 2\pi \frac{2\,(v_o)}{c} f_c \tag{2.69}$$

where ω is the radial Doppler frequency, v_o is a scatterer's radial-velocity relative to the radar platform baseline, f_c is the waveform center frequency, and c is the speed of propagation for the plane wave (i.e. the speed of light). Here (2.69) assumes v_o is small relative to the speed of propagation and that the waveform is "narrowband" so that relativistic and pulse dilation effects can be ignored, which is often (but not always) the case for MTI applications. Using this knowledge, Doppler processing seek to estimate the frequency shifts present in the receive measurements so that the velocity of each scatterer can be determined.

2.5.1 Single-Pulse Doppler Processing

To perform velocity estimation, the receive model must be generalized to incorporate Doppler. As such, the idealized single-pulse continuous receive model for range and Doppler shifted scatters can be expressed as

$$y(t) = \int s(t,\omega) * x(t,\omega)\, d\omega + n(t) \tag{2.70}$$

where

$$s(t,\omega) = s(t)\, e^{j\omega t} \tag{2.71}$$

is the Doppler shifted return from complex scattering coefficient $x(t,\omega)$. Similar to (2.13), the receive measurements can be discretized in snapshot form as

$$\mathbf{y}(\ell) = \int \mathbf{S}(\omega)\, \mathbf{x}(\ell,\omega)\, d\omega + \mathbf{n}(\ell) \tag{2.72}$$

where

$$\mathbf{S}(\omega) = \begin{bmatrix} s_N(\omega) & \cdots & s_1(\omega) & 0 & \cdots & 0 \\ 0 & s_N(\omega) & & s_1(\omega) & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & s_N(\omega) & \cdots & s_1(\omega) \end{bmatrix} \tag{2.73}$$

is the convolution matrix for the Doppler shifted waveform

$$\mathbf{s}(\omega) = \begin{bmatrix} s_1 e^{j\omega\frac{0}{f_s}} & s_2 e^{j\omega\frac{1}{f_s}} & \cdots & s_N e^{j\omega\frac{N-1}{f_s}} \end{bmatrix}^T. \tag{2.74}$$

As ω contains a continuum of possible Doppler shifts which can be computationally impractical to analyze in practice, once again the model is approximated with a discrete grid of K points as

$$\Omega = \begin{bmatrix} \omega_1 & \omega_2 & \cdots & \omega_K \end{bmatrix} \tag{2.75}$$

which assumes there are not scatterers moving faster than the highest/lowest Doppler frequency in Ω and there is not significant amplitude and phase variation between grid points. This approximation again induces "off-grid" or straddling mismatch. However, unlike range straddling where the 3dB oversampling factor κ dictates the fidelity, the sampling of Ω is arbitrary and in turn can be resampled as needed to refine the estimate. Thus, the single-pulse discrete linear receive model is expressed as

$$\mathbf{y}(\ell) = \sum_{\omega \in \Omega} \mathbf{S}(\omega)\mathbf{x}(\ell,\omega) + \mathbf{n}(\ell). \tag{2.76}$$

Matched filtering can be performed in a similar manner as (2.23) and (2.25), where now each of the K waveforms will be used to form a unique filter for each respective Doppler shift, i.e.

$$\hat{x}(\ell,\omega) = \frac{1}{\mathbf{s}^H(\omega)\mathbf{s}(\omega)}\mathbf{s}^H(\omega)\mathbf{y}(\ell). \tag{2.77}$$

Here resolution or Doppler mainlobe width (in radians) can be determined by the pulse duration T_p as

$$\Delta\omega = \frac{2\pi}{T_p}. \tag{2.78}$$

Similar to range sidelobes induced by the autocorrelation, this bank of matched filters will have a degree of correlation across Doppler, i.e.

$$\mathbf{s}^{B^*}(\omega_i) * \mathbf{s}(\omega_j) \neq \mathbf{0} \quad i,j \in K$$

which now induces sidelobe ambiguities in both range and Doppler. This sidelobe structure is typically characterized by waveform's ambiguity function defined as

$$AF(\ell,\omega) = \left[\mathbf{s}^{B^*}(\omega)(*_D)\mathbf{s}(0)\right]_\ell \tag{2.79}$$

where $[\cdot]_\ell$ extracts the ℓth element of the cross correlation [45]. Note, the ambiguity function is centered with respect to the zero-Doppler signal ($\mathbf{s}(0)$) as other Doppler shifts will follow the same

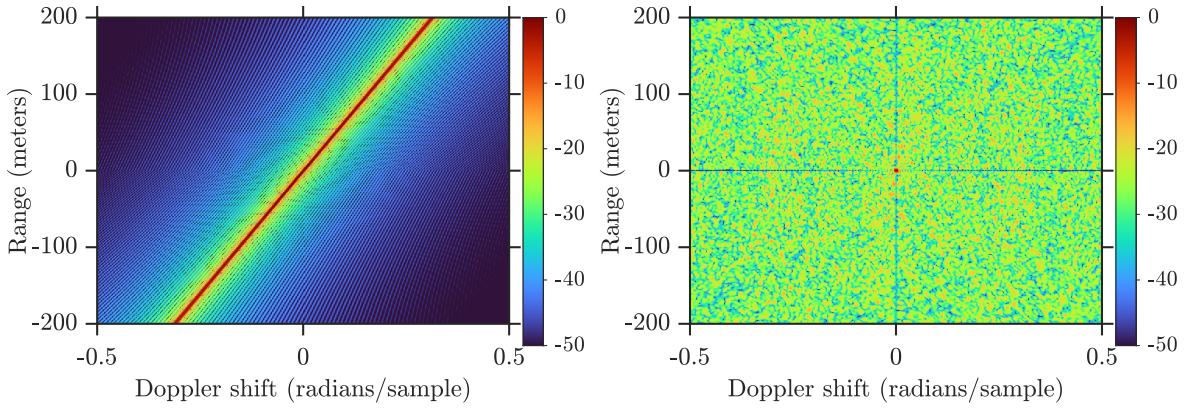

Figure 2.25: Ambiguity Function for a LFM Chirp (left) and RFM Waveform (right)

sidelobe structure via linearity, albeit with a corresponding shift in Doppler and potential effects from the anti-aliasing filter.

Generally waveforms are placed into one of two classes based upon their ambiguity characteristics [45, 103]. The class of Doppler-tolerant waveforms (such as the LFM) typically have a relatively flat ridge extending in range and Doppler (see left side of Figure 2.25), meaning a Doppler shift corresponds to an apparent shifting in range when processing with the zero-Doppler matched filter. Doppler "intolerant" or Doppler-selective waveforms on the other hand, typically have a thumbtack like ambiguity function were the range and Doppler of a single source can be easily resolved, but the sidelobe floor is higher and is spread in range and Doppler (see right side of Figure 2.25). Regardless, both classes of waveform have functional merit depending on application [93].

2.5.2 Multi Pulse Doppler Processing

From (2.78) and (2.69) we can derive an equality that determines when a platform can nominally separate a moving object with radial velocity (v_o) from stationary clutter as

$$1 < T_p \frac{2v_r}{c} f_c \tag{2.80}$$

for pulse duration T_p and center frequency f_c. For acoustical/sonar applications, the relatively slow speed of propagation enables Doppler separability for modest T_p and f_c. For radio frequencies however, the high speed of propagation necessitates higher frequencies and long pulses, which historically has been both monetarily expensive to construct and computationally prohibitive to process.

To enhance Doppler resolution, pulsed radar systems instead aggregate the response of multiple pulses across a longer coherent processing interval (CPI). For a PRI duration T_{pri} corresponding to L samples, M transmit pulses are concatenated to yield the $LM \times 1$ discrete transmission

$$\bar{\bar{\mathbf{s}}} = \begin{bmatrix} \mathbf{s}_1^T & \mathbf{0}_{(L-N)\times 1}^T & \mathbf{s}_2^T & \cdots & \mathbf{s}_M^T & \mathbf{0}_{(L-N)\times 1}^T \end{bmatrix}^T \tag{2.81}$$

for the CPI. The range-Doppler receive snapshot model can then be derived in the same manner as (2.18) by treating (2.81) as the waveform and the ambiguity function (2.79) can be used to characterize the range-Doppler sidelobe structure. Here the Doppler resolution is now proportional to $T_{cpi} = MT_{pri}$ rather than T_p, however the computational complexity of matched filtering now scales with LM rather than N where in general $LM \gg N$.

To reduce computational burden, the intra-pulse Doppler is typically approximated as negligible relative to pulse-to-pulse Doppler shift, i.e. for the mth transmit waveform $\mathbf{s}_m = [s_{1,m}\ s_{2,m}\ \cdots\ s_{N,m}]^T$ the Doppler shift for ω is incorporated as

$$\begin{bmatrix} s_{1,m} \\ s_{2,m} \\ \vdots \\ s_{N,m} \end{bmatrix} \odot \begin{bmatrix} e^{j\omega\frac{L(m-1)+0}{f_s}} \\ e^{j\omega\frac{L(m-1)+1}{f_s}} \\ \vdots \\ e^{j\omega\frac{L(m-1)+N-1}{f_s}} \end{bmatrix} \approx e^{j\omega\frac{L(m-1)}{f_s}}\mathbf{s}_m = e^{j\omega\frac{m-1}{f_{prf}}}\mathbf{s}_m \tag{2.82}$$

where $f_{prf} = \frac{1}{T_{pri}}$ is the pulse repetition frequency (PRF) of the CPI. Thus, a discretized $D \times 1$

collection of receive samples for the mth pulse is denoted

$$\mathbf{y}_m(\ell) = \sum_{\omega\in\Omega} e^{j\omega\frac{m-1}{f_{prf}}}\mathbf{S}_m\mathbf{x}(\ell,\omega) + \mathbf{n}_m(\ell) \tag{2.83}$$

where

$$\mathbf{S}_m = \begin{bmatrix} s_{N,m} & \cdots & s_{1,m} & 0 & \cdots & 0 \\ 0 & s_{N,m} & & s_{1,m} & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & s_{N,m} & \cdots & s_{1,m} \end{bmatrix} \tag{2.84}$$

is the $D\times(D+N-1)$ convolution matrix for $\mathbf{s}_m$ and ℓ is the sample delay relative to the start of the mth PRI. This organization of receive measurements effectively splits the time domain into two dimensions, sample delay from the start of the PRI (denoted fast-time) corresponding to range via index ℓ, and relative time from pulse-to-pulse (denoted slow-time) corresponding to Doppler via the pulse index m. This fast/slow-time notation is a convenient approximation that allows for efficient separate/sequential processing of each respective domain when the fast-time Doppler shift is negligible.

The response in (2.83) for each pulse is then collected into the matrix

$$\mathbf{Y}(\ell) = \begin{bmatrix} \mathbf{y}_1(\ell) & \mathbf{y}_2(\ell) & \cdots & \mathbf{y}_M(\ell) \end{bmatrix} \tag{2.85}$$

which is organized in terms of fast-time along each column and slow-time along each row. A repeated transmit waveform $\mathbf{s}_m = \mathbf{s}$ results in a decoupling of range and Doppler via linearity and superposition. In this repeated framework matched filtering in range and Doppler can be performed independently along the respective fast-time and slow-time axis as

$$\hat{x}(\ell,\omega) = \frac{1}{\mathbf{s}^H\mathbf{s}}\mathbf{s}^H\mathbf{Y}(\ell)\mathbf{a}^*(\omega)\frac{1}{\mathbf{a}^H(\omega)\mathbf{a}(\omega)} \tag{2.86}$$

where

$$\mathbf{a}(\omega) = \begin{bmatrix} e^{j\omega\frac{0}{f_{prf}}} & e^{j\omega\frac{1}{f_{prf}}} & \cdots & e^{j\omega\frac{M-1}{f_{prf}}} \end{bmatrix}^T \tag{2.87}$$

is a $M \times 1$ frequency steering vector for Doppler frequency ω that acts as a matched filter. As with pulse-compression, windowing can likewise be applied to the Doppler matched filter to reduce sidelobes.

For a generalized arbitrary/non-repeating set of pulses, range and Doppler are coupled and require a specific order of operations to ensure a scattering response coheres across range and Doppler. This can be achieved by first performing range estimation for each column of (2.85) as

$$\hat{x}_m(\ell) = \mathbf{w}_m^H \mathbf{y}_m(\ell) \tag{2.88}$$

where $\mathbf{w}_m$ can be the MF, LS MMF, or the APC[5] filters discussed previously. These estimates are then accumulated into the $M \times 1$ range-compressed slow-time snapshot

$$\hat{\mathbf{y}}(\ell) = \begin{bmatrix} \hat{x}_1(\ell) & \hat{x}_2(\ell) & \cdots & \hat{x}_M(\ell) \end{bmatrix}^T \tag{2.89}$$

where subsequent Doppler processing is performed as

$$\hat{x}(\ell, \omega) = \frac{1}{\mathbf{a}^H(\omega)\mathbf{a}(\omega)} \mathbf{a}^H(\omega)\hat{\mathbf{y}}(\ell) \tag{2.90}$$

Equivalently, joint range-Doppler estimation [59] can be performed by expressing (2.85) as a vector through a concatenation (or vectorization) of the columns in the snapshot matrix as

$$\bar{\mathbf{y}}(\ell) = \begin{bmatrix} \mathbf{y}_1^T(\ell) & \mathbf{y}_2^T(\ell) & \cdots & \mathbf{y}_M^T(\ell) \end{bmatrix}^T \tag{2.91}$$

[5]Note, for the APC filters $\mathbf{w}_m$ would also be a function of ℓ, i.e. $\mathbf{w}_m(\ell)$

where the joint range-Doppler estimation is expressed as

$$\hat{x}(\ell,\omega) = \bar{\mathbf{w}}^H(\omega)\,\bar{\mathbf{y}}(\ell) \tag{2.92}$$

where

$$\bar{\mathbf{w}}(\omega) = \begin{bmatrix} e^{j\omega\frac{0}{f_{prf}}}\mathbf{w}_1^T & e^{j\omega\frac{1}{f_{prf}}}\mathbf{w}_2^T & \cdots & e^{j\omega\frac{M-1}{f_{prf}}}\mathbf{w}_M^T \end{bmatrix}^T \tag{2.93}$$

is the Doppler-shifted waveform sequence, and normalization in (2.92) is included for consistent scaling. The linear receive model for (2.91) can likewise be expressed in a form akin to (2.83) as the $DM \times 1$ snapshot

$$\bar{\mathbf{y}}(\ell) = \sum_{\omega\in\Omega} \bar{\mathbf{S}}(\omega)\,\mathbf{x}(\ell,\omega) + \bar{\mathbf{n}}(\ell) \tag{2.94}$$

where the concatenated set of convolution matrices are

$$\bar{\mathbf{S}}(\omega) = \begin{bmatrix} e^{j\omega\frac{0}{f_{prf}}}\mathbf{S}_1^T & e^{j\omega\frac{1}{f_{prf}}}\mathbf{S}_2^T & \cdots & e^{j\omega\frac{M-1}{f_{prf}}}\mathbf{S}_M^T \end{bmatrix}^T \tag{2.95}$$

and the DM noise samples are subsumed in

$$\bar{\mathbf{n}}(\ell) = \begin{bmatrix} \mathbf{n}_1(\ell)^T & \mathbf{n}_2(\ell)^T & \cdots & \mathbf{n}_M(\ell)^T \end{bmatrix}^T. \tag{2.96}$$

Note, the approximation procedure in (2.82) effectively reduces the sample rate for the Doppler dimension from the system (or fast-time) sample rate of f_s to the PRF (or slow-time sample rate) f_{prf}. As idealized Doppler shifts correspond to complex sinusoids, this reduced sample rate will produce aliasing when

$$f_{prf} < \left|\frac{2v_r}{c} f_c\right|. \tag{2.97}$$

Thus, the maximum unambiguous velocity v_{unamb} is bounded by f_{prf} as

$$|v_{unamb}| < \frac{c f_{prf}}{4 f_c} \tag{2.98}$$

where the additional factor of 2 comes from the ability to differentiate forward and backward radial velocities. Additionally, for a fast-time Doppler selective CPI, movers with significant radial velocity may require additional compensation [89, 92, 93, 104, 105].

The range-Doppler sidelobe structure for a pulsed emission can then be characterized in terms of a point-spread function (PSF) which approximates the CPI's ambiguity function within the unambiguous velocity bounds. The PSF is defined as

$$PSF\left(\ell,\omega\right)=\frac{\bar{\mathbf{s}}^{H}\left(\omega\right)\tilde{\mathbf{y}}\left(\ell\right)}{\bar{\mathbf{s}}^{H}\left(\omega\right)\bar{\mathbf{s}}\left(\omega\right)} \tag{2.99}$$

where $\tilde{\mathbf{y}}\left(\ell\right)$ is the idealized response for a point scatter at zero-delay and zero-Doppler.

Similar to the classes of Doppler tolerant and intolerant waveforms mentioned in Section 2.5.1 the choice of transmit waveforms across a CPI can be broadly grouped in two categories, repeated and agile. Most legacy systems use a repeated framework in which the transmit waveform is kept consistent across the CPI so that the range-sidelobes are homogeneous in slow-time. As shown in (2.86), homogenization effectively decouples fast-time from Doppler and slow-time from range, resulting in a sinc-sidelobe spreading in the Doppler domain. Alternatively, an agile emission changes the transmit waveform within the CPI, resulting in a heterogeneous range-sidelobe structure across slow-time.

This heterogeneity effectively couples range and Doppler estimation as every range-delay produces a unique sidelobe modulation across slow-time i.e. RSM. This RSM creates a delay dependent smearing in the Doppler domain as the slow-time response is no longer uniform and sinusoidal. An example PSF for 100 repeated LFM waveforms (left) and 100 unique RFM waveforms (right) is shown in Figure 2.26 where the sinc-sidelobes and RSM are clearly visible for the repeated and agile cases respectively. Unlike the repeated waveform, the psuedo-random nature of the RFM sidelobes creates a decoherence outside the range-Doppler mainlobe and RSM is coherently suppress like noise. This results in lower peak range-Doppler sidelobes (relative to the repeated case) but a spreading across all Doppler. The coherent suppression effect is demonstrated

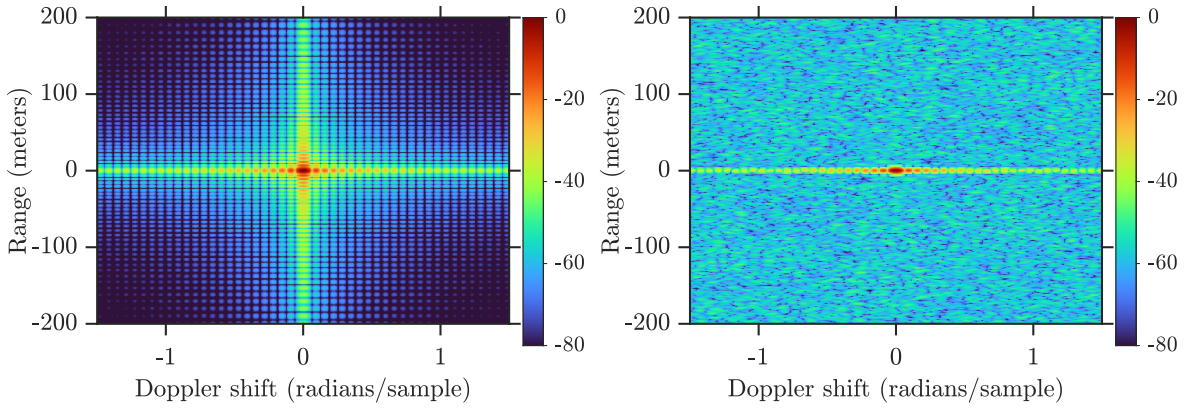

Figure 2.26: PSF for a 100 repeated LFM Chirps (left) and 100 Unique RFM Waveforms (right)

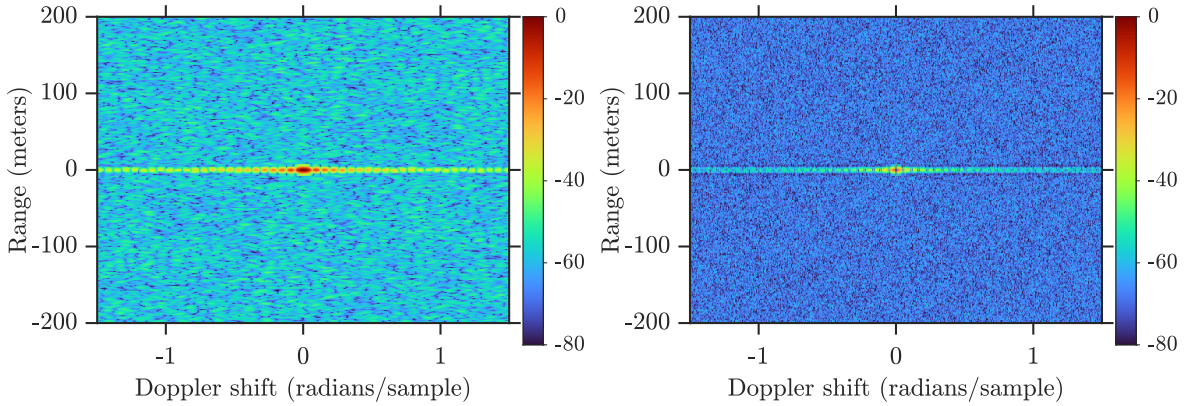

Figure 2.27: PSF for a 100 Unique RFM Waveforms (left) and 1000 Unique RFM Waveforms

in the PSFs of Figure 2.27 where CPIs of 100 and 1000 unique RFM waveforms shown that increasing the number of unique waveforms by a factor of 10 corresponds to $\sim$10dB of additional RSM suppression.

Supplemental (these results will be examined with further detail in Chapters 3-6) open-air experimental measurements for 1000 repeated LFM and 1000 unique RFM waveforms were collected and pulsed range-Doppler processing was performed with a Taylor window applied in Doppler to reduce sidelobes (in Doppler only). Here the particular RFM waveforms were designed with an alternating projection routine such that the frequency content has a Gaussian shape and are denoted as PROFM waveforms [73]. Figures 2.28 and 2.29 show the LFM and RFM results respectively with a focused look at the intersection on the right. The vehicles are clearly visible for both the LFM and RFM case, however RSM from the proportionally high-power clutter creates added back-

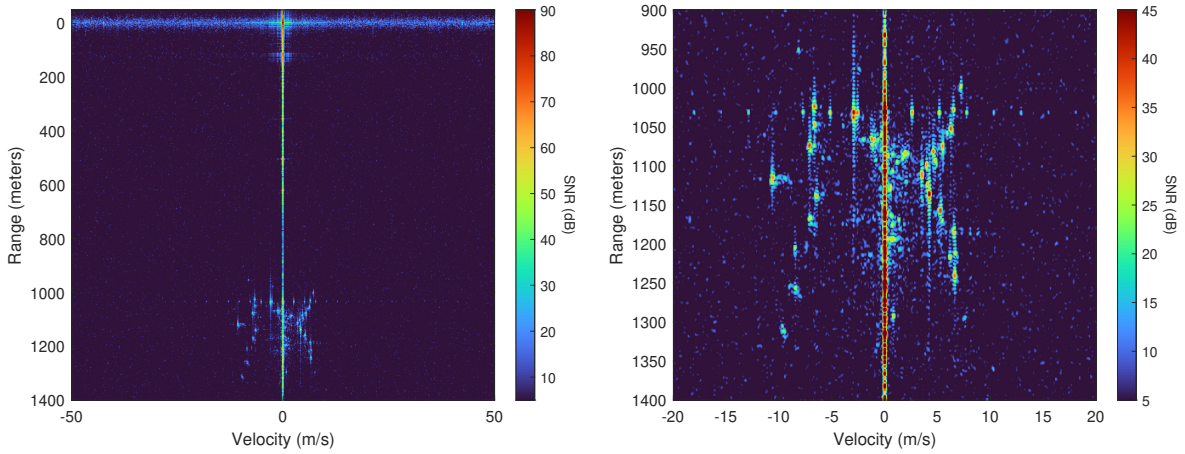

Figure 2.28: Open-Air Experimental Measurement for 1000 repeated LFM Chirps (left) and zoomed in look at intersection (right)

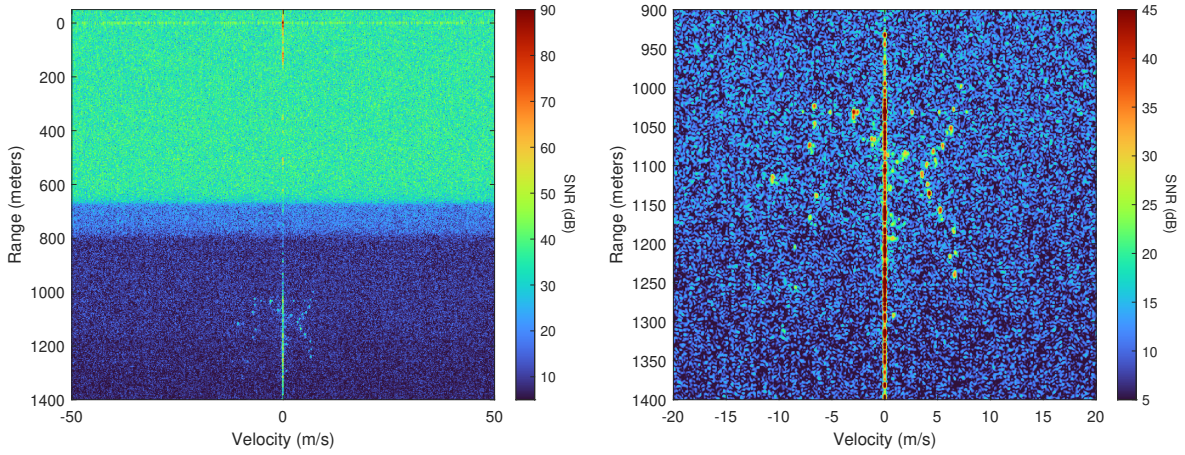

Figure 2.29: Open-Air Experimental Measurement for 1000 unique RFM waveforms (left) and zoomed in look at intersection (right)

ground speckle for the RFM case. As the RSM is uncorrelated, this speckle effectively acts as an increased noise power and will degrade detection performance.

2.5.3 Clutter Cancellation

Range-Doppler processing enables separation of moving targets from stationary clutter, but clutter generally has a much higher power than moving objects [67, 84] and post pulse-compression scatterer peaks will have a sinc-sidelobe structure in Doppler. Slow-time windowing can be applied to reduce these sidelobes, but will broaden the Doppler mainlobe (degrading Doppler resolution) and

will create a mismatch loss. Rather than attempting to minimize Doppler sidelobes from clutter, performance can often be improved by applying clutter cancellation.

While there are a wide variety of clutter cancellation techniques [4, 106], all typically involve some form of filtering or a projection operation across the slow-time domain. Consider the $M \times 1$ slow-time response for a idealized range compressed scatter at the ℓth delay

$$\hat{\mathbf{y}}(\ell) = \sum_{\omega \in \Omega} \mathbf{a}(\omega) x(\ell, \omega) + \mathbf{n}(\ell) = \mathbf{A}\mathbf{x}_{\Omega}(\ell) + \mathbf{n}(\ell) \tag{2.100}$$

where

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}(\omega_1) & \mathbf{a}(\omega_2) & \cdots & \mathbf{a}(\omega_K) \end{bmatrix} \tag{2.101}$$

is the bank of steering vectors and

$$\mathbf{x}_{\Omega}(\ell) = \begin{bmatrix} x(\ell, \omega_1) & x(\ell, \omega_2) & \cdots & x(\ell, \omega_K) \end{bmatrix} \tag{2.102}$$

are the complex scattering coefficients for the Doppler shifts in Ω at delay ℓ respectively. To suppress clutter while retaining our desired signals, the measurement vector is separated into a clutter and target components as

$$\hat{\mathbf{y}}(\ell) = \hat{\mathbf{y}}_c(\ell) + \hat{\mathbf{y}}_t(\ell) + \mathbf{n}(\ell) \tag{2.103}$$

where

$$\hat{\mathbf{y}}_c(\ell) = \sum_{\omega \in \Omega} \mathbf{a}(\omega) x_c(\ell, \omega) \tag{2.104}$$

are the measurements due to the complex scattering from clutter $x_c(\ell, \omega)$ and

$$\hat{\mathbf{y}}_t(\ell) = \sum_{\omega \in \Omega} \mathbf{a}(\omega) x_t(\ell, \omega) \tag{2.105}$$

is a measurement vector due to the remaining complex scatterers $x_t(\ell, \omega)$ (what we wish to estimate). Statistical clutter cancellation can be performed by treating $\hat{\mathbf{y}}_c(\ell)$ and $\mathbf{n}(\ell)$ as independent

random variables to form a correlation matrix

$$\mathbf{R}_{cc} = \mathbf{R}_c + \mathbf{R}_n = \mathbb{E}\{\mathbf{y}_c(\ell)\,\mathbf{y}_c^H(\ell)\} + \mathbb{E}\{\mathbf{n}(\ell)\,\mathbf{n}^H(\ell)\} \tag{2.106}$$

which can be used to from a Karhunen–Loève transform (KLT) [66] that will whiten the data of the clutter contribution as

$$\hat{\mathbf{y}}_{cc}(\ell) = (\mathbf{R}_{cc})^{-\frac{1}{2}}\,\hat{\mathbf{y}}(\ell). \tag{2.107}$$

Instead of KLT, in radar it is more common to perform maximum signal-to-clutter-plus-noise ratio SCNR filtering which can be derived similar to the max SINR filter by treating clutter as interference [84] and transforming the measurements as

$$\hat{\mathbf{y}}_{cc}(\ell) = (\mathbf{R}_{cc})^{-1}\,\hat{\mathbf{y}}(\ell) \tag{2.108}$$

followed by standard matched filtering in Doppler. However, the matrix inversion will induce some distortion, scaling, and noise gain with effects most prominent on sources close to the clutter ridge. This procedure can be equivalently performed using a Doppler filter

$$\mathbf{w}_{cc}(\omega) = \frac{1}{\mathbf{a}^H(\omega)\,\mathbf{a}(\omega)}\,(\mathbf{R}_{cc})^{-1}\,\mathbf{a}(\omega) \tag{2.109}$$

and performing max SCNR estimation via

$$\hat{x}_t(\ell,\omega) = \mathbf{w}_{cc}^H(\omega)\,\hat{\mathbf{y}}(\ell). \tag{2.110}$$

With the filter formulation, normalization can also be incorporated as

$$\mathbf{w}_{cc}(\omega) = \frac{1}{\mathbf{a}^H(\omega)\,(\mathbf{R}_{cc})^{-1}\,\mathbf{a}(\omega)}\,(\mathbf{R}_{cc})^{-1}\,\mathbf{a}(\omega) \tag{2.111}$$

to remove the scale factor due to the matrix inversion, but not the distortion or noise gain.

This filtering is optimal in terms of a SCNR [84] but in practice clairvoyant knowledge of the clutter covariance matrix will not be known. There are a multitude of techniques for estimating the clutter covariance matrix, but for simplicity here a stationary platform and uniform clutter power is used to construct a structured clutter covariance. When stationary, clutter will largely be focused about zero-Doppler (though there may be some spread due to motion from wind or ocean waves) and a subset of K_c Doppler frequencies

$$\tilde{\Omega} = \begin{bmatrix} \tilde{\omega}_1 & \tilde{\omega}_2 & \cdots & \tilde{\omega}_{K_c} \end{bmatrix} \tag{2.112}$$

corresponding to Doppler steering vectors

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{a}(\tilde{\omega}_1) & \mathbf{a}(\tilde{\omega}_2) & \cdots & \mathbf{a}(\tilde{\omega}_{K_c}) \end{bmatrix} \tag{2.113}$$

can be used to approximately characterize the clutter response and extent. Assuming a uniform spread of clutter with power σ_c^2 and white Gaussian noise (WGN) with variance σ_n^2 , the optimal clutter-plus-noise correlation matrix is

$$\mathbf{R}_{cc} = \sigma_c^2 \tilde{\mathbf{A}} \tilde{\mathbf{A}}^H + \sigma_n^2 \mathbf{I}_M \tag{2.114}$$

which can be used to cancel the clutter as before.

Alternatively, if the clutter power is not known but the Doppler extent is, the measurement vector can be projected onto the nullspace of the clutter by minimizing

$$\|\hat{\mathbf{y}}(\ell) - \tilde{\mathbf{A}}\boldsymbol{\gamma}\|^2 \tag{2.115}$$

with respect to the $K_c \times 1$ vector $\boldsymbol{\gamma}$, which has the closed form solution:

$$\boldsymbol{\gamma} = \left(\tilde{\mathbf{A}}^H \tilde{\mathbf{A}}\right)^{-1} \tilde{\mathbf{A}}^H \hat{\mathbf{y}}(\ell). \tag{2.116}$$

Thus, the projection matrix

$$\Gamma_{cc} = \mathbf{I} - \tilde{\mathbf{A}} \left(\tilde{\mathbf{A}}^H \tilde{\mathbf{A}} \right)^{-1} \tilde{\mathbf{A}}^H \tag{2.117}$$

can be used in lieu of $(\mathbf{R}_{cc})^{-1}$.

These clutter cancellation techniques can be very effective at suppressing signatures that are sinusoidal across slow-time. For the repeated waveform case, both the range-mainlobe and range-sidelobes of clutter have this sinusoidal response due to the decoupling of fast-time and slow-time. For an agile CPI, the range-mainlobe will fit this structure but the RSM is a joint range-Doppler effect and can not be suppressed via slow-time processing alone.

2.6 NIMPC

Rather than just utilizing the dimensionality provided by the slow-time dimension (i.e. number of pulses), the Non-Identical Multiple Pulse Compression (NIMPC) framework was proposed to perform clutter cancellation and pulse compression using the joint slow-time-fast-time dimensionality of an agile CPI. This joint representation can then perform simultaneous cancellation of both the sinusoidal mainlobe Doppler sidelobes and non-sinusoidal RSM clutter contributions. NIMPC utilizes the form of (2.94) with $D = N$ (though without loss of generality any $D \geq N$ can be used) where a receive snapshot with fast-time delay ℓ for CPI with M pulses is expressed as

$$\bar{\mathbf{y}}(\ell) = \sum_{\omega \in \Omega} \bar{\mathbf{S}}(\omega) \mathbf{x}(\ell, \omega) + \bar{\mathbf{n}}(\ell) \tag{2.118}$$

where Ω is a collection of K radial Doppler frequencies and

$$\bar{\mathbf{S}}(\omega) = \left[e^{j\omega \frac{0}{f_{prf}}} \mathbf{S}_1^T \quad e^{j\omega \frac{1}{f_{prf}}} \mathbf{S}_2^T \quad \cdots \quad e^{j\omega \frac{M-1}{f_{prf}}} \mathbf{S}_M^T \right]^T \tag{2.119}$$

is formed by concatenating the Doppler shifted convolution matrices for each waveform. Using a similar clutter cancellation methodology as (2.5.3), a subset of K_c Doppler frequencies

$$\tilde{\Omega} = \begin{bmatrix} \tilde{\omega}_1 & \tilde{\omega}_2 & \cdots & \tilde{\omega}_{K_c} \end{bmatrix} \tag{2.120}$$

corresponding to the clutter can be used to form the joint Range-Doppler response matrix

$$\tilde{\bar{\mathbf{S}}} = \begin{bmatrix} \bar{\mathbf{S}}(\tilde{\omega}_1) & \bar{\mathbf{S}}(\tilde{\omega}_2) & \cdots & \bar{\mathbf{S}}(\tilde{\omega}_{K_c}) \end{bmatrix} \tag{2.121}$$

which is used to form an $MD \times MD$ estimate clutter correlation matrix as

$$\bar{\mathbf{R}}_{cc} = \sigma_c^2 \tilde{\bar{\mathbf{S}}} \tilde{\bar{\mathbf{S}}}^H + \bar{\mathbf{R}}_n \tag{2.122}$$

where $\bar{\mathbf{R}}_n = \sigma_n^2 \mathbf{I}$ is the noise correlation matrix for WGN with noise power σ_n^2 and σ_c^2 is the expected clutter power. This joint range-Doppler correlation matrix can then be used to perform maximum SINR processing by transforming the receive snapshot as

$$\bar{\mathbf{y}}_{cc}(\ell) = \left(\bar{\mathbf{R}}_{cc}\right)^{-1} \bar{\mathbf{y}}(\ell) \tag{2.123}$$

or by equivalently forming a normalized clutter-canceling filter as

$$\bar{\mathbf{w}}_{cc}(\omega) = \frac{1}{\bar{\mathbf{s}}^H(\omega)\left(\bar{\mathbf{R}}_{cc}\right)^{-1}\bar{\mathbf{s}}(\omega)} \left(\bar{\mathbf{R}}_{cc}\right)^{-1} \bar{\mathbf{s}}(\omega) \tag{2.124}$$

which is then applied to the measurements to yield the clutter-canceled range-Doppler estimate as

$$\hat{x}_{cc}(\ell, \omega) = \bar{\mathbf{w}}_{cc}^H(\omega)\, \bar{\mathbf{y}}(\ell). \tag{2.125}$$

Alternatively, the filter with scaling of

$$\bar{\mathbf{w}}_{cc}(\omega) = \frac{1}{\sqrt{\bar{\mathbf{s}}^H(\omega)\left(\bar{\mathbf{R}}_{cc}\right)^{-2}\bar{\mathbf{s}}(\omega)}}\left(\bar{\mathbf{R}}_{cc}\right)^{-1}\bar{\mathbf{s}}(\omega)$$

will retain a uniform noise floor which can be useful when performing detection.

The complexity of NIMPC clutter cancellation scales with entire slow time dimension and storage/inversion of the $MD \times MD$ correlation matrix rapidly becomes infeasible for even a modest M and D. For the supplemental measurements in Figure 2.29, this would correspond to a $300,000 \times 300,000$ complex matrix which would require $\sim$1.5 terabytes of memory just to store (for complex double precision). [107] exploited redundancies and structure of the NIMPC framework to decrease computational burden and storage by a factor of M, but complexity and storage were still a practical limitation. In Chapter 5 alternative representations and solutions will be developed to drastically improve this computational burden.

2.7 Select Structured Matrices

The QPs for the LS MMFs and APC all require construction and solution to a $D \times D$ linear system, which through brute force have a complexity $\mathcal{O}(D^3)$ for the solve alone. This computational burden can significantly degrade the practical viability of these approaches as D increases. As such, it is necessary to develop approaches and solvers that can realistically be implemented on modern hardware. The following subsection outlines properties of select structured matrices and techniques that enable efficient operations. This matrix structure will later be used to reduce computational complexity and push these algorithms closer to real-time viability.

2.7.1 Sparse and Banded Hermitian Matrices

Sparse hermitian matrices correspond to a class of matrices in which the majority of elements are zero. For a generic $D \times D$ sparse matrix $\mathbf{H}$ with a full main diagonal and h non-zero off-diagonal

entries, matrix-vector multiplication with an arbitrary vector, i.e.,

$$\mathbf{a} = \mathbf{H}\mathbf{b},$$

can ignore operations corresponding to the zero-entries to yield a complexity of $\mathscr{O}(hD+D)$. Additionally, when $\mathbf{H}$ is positive definite, the Cholesky factorization

$$\mathbf{H} = \mathbf{L}\mathbf{L}^H \tag{2.126}$$

can be computed with a complexity of $\sim\mathscr{O}(h^2D)$ and the $D \times D$ lower-triangular matrix preserves the same sparsity structure as $\mathbf{H}$. However, this cost is inexact as it is a function of the sparsity structure and can potentially reduced further depending on where the non-zero elements are. Using the sparse Cholesky factorization, linear systems involving $\mathbf{H}$ can be solved via forward and backward substitution with $\mathbf{L}$ with $\sim\mathscr{O}(hD)$ complexity. Thus the aggregate complexity to solve a QP with a sparse matrix will be $\sim\mathscr{O}(h^2D+hD)$ [108, p. 670].

Banded hermitian matrices are of the same vein in which the matrix has a dense "band" of non-zero diagonals. A $D \times D$ banded hermitian matrix $\mathbf{B}$ with a matrix-bandwidth of h when the $-(h-1)th$ to $(h-1)th$ diagonals are the only components that have non-zero entries, i.e.,

$$\mathbf{B} = \begin{bmatrix} b_{1,1} & \cdots & b_{1,h} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \ddots & \ddots & \vdots \\ b_{h,1} & \cdots & b_{h,h} & & \ddots & 0 \\ 0 & \ddots & & \ddots & & b_{D,D-h+1} \\ \vdots & \ddots & \ddots & & \ddots & \vdots \\ 0 & \cdots & 0 & b_{D,D-h+1} & \cdots & b_{D,D} \end{bmatrix}. \tag{2.127}$$

Matrix vector multiplication with $\mathbf{B}$ can be computed with $\mathscr{O}(hD)$ and the banded cholesky decomposition

$$\mathbf{B} = \mathbf{L}\mathbf{L}^H \tag{2.128}$$

can be computed with $\mathcal{O}(h^2D)$ which preserves the banded structure and

$$\mathbf{L} = \begin{bmatrix} L_{1,1} & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots & & \vdots \\ L_{h,1} & \cdots & L_{h,h} & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & L_{h,1} & \cdots & L_{D,h} \end{bmatrix} \tag{2.129}$$

can be leveraged to solve QPs with $\mathbf{B}$ as the hessian with $\mathcal{O}(h^2D + hD)$ complexity [108, p. 670].

2.7.2 Circulant Matrices

Circulant matrices [109, p. 61] are a special class of square matrices in which the columns of the matrix are expressed as shifted and wrapped replicas of the $D \times 1$ vector $\mathbf{c} = \begin{bmatrix} c_0 & \cdots & c_{D-1} \end{bmatrix}^T$ with the form

$$\mathbf{C} = \begin{bmatrix} c_0 & c_{D-1} & \cdots & c_1 \\ c_1 & c_0 & & \vdots \\ \vdots & & \ddots & c_{D-1} \\ c_{D-1} & \cdots & c_1 & c_0 \end{bmatrix} \tag{2.130}$$

for a $D \times D$ Circulant matrix. This shifting behavior enforces the property that $\mathbf{C}$ is has eigenvectors corresponding to the discrete Fourier transform (DFT), thus it can be diagonalized as

$$\mathbf{C} = \mathbf{F}^{-1}\mathbf{C}_f\mathbf{F} \tag{2.131}$$

where $\mathbf{F}$ is defined as a $D \times D$ discrete Fourier transform (DFT) matrix with elements

$$f_{i,k} = \exp\{-j2\pi(i-1)(k-1)\},$$

and

$$\mathbf{C}_f = diag\{\mathbf{Fc}\}$$

where $diag\{\cdot\}$ is a diagonalization operation that places the input vector onto the main diagonal of the output matrix.

Thus, matrix-vector multiplication with $\mathbf{C}$ and an arbitrary $D \times 1$ vector $\mathbf{b}$ can be performed efficiently by utilizing the FFT and IFFT to perform element-wise frequency domain multiplication via

$$\begin{aligned} \mathbf{a} = \mathbf{Cb} &= \mathbf{F}^{-1}\mathbf{C}_f\mathbf{Fb} \\ &= \mathbf{F}^{-1}\left[(\mathbf{Fc}) \odot (\mathbf{Fb})\right] \\ &= \text{IFFT}\{\,\text{FFT}\{\mathbf{c}\} \odot \text{FFT}\{\mathbf{b}\}\,\} \end{aligned} \tag{2.132}$$

where $\odot$ is the Hadamard product and $\mathbf{a}$ equivalent to the circular convolution of $\mathbf{b}$ and $\mathbf{c}$ [110].

2.7.3 Toeplitz Matrices

Circulant matrices are a special case of a more general class of matrices that frequently show up in linear systems involving convolution. The parent class where the main diagonal and off-diagonals are constant are denoted Toeplitz matrices. The general form of a non-square $D \times K$ Toeplitz matrix can be expressed as

$$\mathbf{T} = \begin{bmatrix} t_0 & t_{-1} & \cdots & t_{-K+1} \\ t_1 & t_0 & & \vdots \\ \vdots & & \ddots & t_{-1} \\ t_{D-1} & \cdots & t_1 & t_0 \end{bmatrix}. \tag{2.133}$$

Toeplitz matrix-vector multiplication can again be performed efficiently in the frequency domain by first embedding the Toeplitz matrix into a larger Circulant matrix [110]. For instance, consider

the $(D+K-1)\times(D+K-1)$ Circulant embedding constructed as

$$\mathbf{C}\{\mathbf{T}\} = \begin{bmatrix} \mathbf{T} & \mathbf{U}^T \\ \mathbf{U} & \mathbf{T} \end{bmatrix} \tag{2.134}$$

where

$$\mathbf{U} = \begin{bmatrix} t_{-K+1} & t_{D-1} & t_{D-2} & \cdots & t_1 \\ t_{-K+2} & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & t_{D-2} \\ t_{-1} & \cdots & t_{-K+2} & t_{-K+1} & t_{D-1} \end{bmatrix}.$$

Note, this particular embedded matrix requires the least amount of additional dimensionality, but it is often preferable to pad additional zeros to $\mathbf{U}$ so that the corresponding FFT/IFFT pairs are a power of two. Additionally, when $\mathbf{T}$ is a banded matrix, then the embedding dimensionality can be further reduced. Using (2.134) Toeplitz matrix-vector multiplication with an arbitrary $K\times 1$ vector $\mathbf{b}$ is performed as

$$\begin{aligned} \mathbf{T}\mathbf{b} &= \begin{bmatrix} \mathbf{I}_D & \mathbf{0}_{D\times(K-1)} \end{bmatrix} \mathbf{C}\{\mathbf{T}\} \begin{bmatrix} \mathbf{b} \\ \mathbf{0}_{(D-1)\times 1} \end{bmatrix} \\ &= \left[\text{IFFT}\left\{\text{FFT}\left\{\hat{\mathbf{t}}\right\} \odot \text{FFT}\left\{\hat{\mathbf{b}}\right\}\right\}\right]_{1:D} \end{aligned} \tag{2.135}$$

where

$$\hat{\mathbf{t}} = \begin{bmatrix} t_0 & t_1 & \cdots & t_{D-1} & t_{-K+1} & t_{-K+2} & \cdots & t_{-1} \end{bmatrix}^T$$

and $\hat{\mathbf{b}}$ is a zero padded replica of $\mathbf{b}$ shown in (2.135).

Another special case that shows up frequently in least squares problems (LS) is the Hermitian Toeplitz matrix expressed as

$$\mathbf{H} = \begin{bmatrix} t_0 & t_1^* & \cdots & t_{D-1}^* \\ t_1 & t_0 & & \vdots \\ \vdots & & \ddots & t_1^* \\ t_{D-1} & \cdots & t_1 & t_0 \end{bmatrix}. \tag{2.136}$$

Hermitian Toeplitz matrices have some nice properties, such as positive semi-definiteness (when t_0 is positive), a purely real frequency domain diagonal when embedding in an appropriate Circulant matrix, and an inverse structure that can be fully determined by its first column [109, p. 62].

2.7.4 Structure of a Toeplitz Matrix Inverse

Square non-singular Toeplitz matrices have some convenient properties that simplify computation of their inverse. Here we will focus on the structure of the Circulant inverse and the Hermitian Toeplitz inverse for simplicity, though there are analogous forms for the more general case.

From (2.131) we see that the Circulant matrix form can be decomposed into three matrices with an inverse

$$\begin{aligned}\mathbf{C}^{-1} &= \left(\mathbf{F}^{-1}\mathbf{C}_f^{-1}\mathbf{F}\right)^{-1} \\ &= \mathbf{F}^{-1}\mathbf{C}_f^{-1}\mathbf{F}\end{aligned} \tag{2.137}$$

with

$$\mathbf{C}_f^{-1} = diag\left\{\mathbf{1} \oslash (\mathbf{Fc})\right\}$$

where $\oslash$ represents element-wise division. Thus, the structure of the inverse of a Circulant matrix is also Circulant and inversion can be done element-wise in the frequency domain.

Diagonalizing the Hermitian Toeplitz case is not as straightforward, but can be determined by solving for the first column of the inverse matrix which can be solved directly using levinson-recursion in with $\mathcal{O}\left(D^2\right)$ complexity [109, p. 85] or though for larger and well conditioned matrices this can be lowered even further [111, 112]. In general, the first column of the inverse $\mathbf{a}$ can be found by solving

$$\mathbf{Ta} = \mathbf{e}_1 \tag{2.138}$$

where $\mathbf{e}_1$ is an $D \times 1$ elementary matrix will all zeros except a 1 in the first element. With the first column, $\mathbf{T}^{-1}$ can be reconstructed as

$$\mathbf{T}^{-1} = \frac{1}{a_1}\left(\mathbf{T}_f{\mathbf{T}_f}^H - \mathbf{T}_b{\mathbf{T}_b}^H\right) \tag{2.139}$$

where

$$\mathbf{T}_f = \begin{bmatrix} a_1 & 0 & \cdots & 0 \\ a_2 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_D & \cdots & a_2 & a_1 \end{bmatrix}$$

and

$$\mathbf{T}_b = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ a_D^* & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_2^* & \cdots & a_D^* & 0 \end{bmatrix}$$

are both Toeplitz matrices with structure than can be used to perform efficient matrix vector multiplication. Both the Circulant and Hermitian Toeplitz matrix inverse model structure will later be used to efficiently compute our optimal mismatch filters.

2.7.5 Toeplitz-Block and Block-Toeplitz Matrices

Another special class of matrix is denoted block-Toeplitz [109, p. 344] in which the $DQ \times DQ$ matrix $\bar{\mathbf{T}}$ is composed of $2D-1$ matrices $\mathbf{R}_d$ of size $Q \times Q$. The block-Toepltiz matrix then has consistent block diagonals as

$$\bar{\mathbf{T}} = \begin{bmatrix} \mathbf{R}_0 & \mathbf{R}_{-1} & \cdots & \mathbf{R}_{-D+1} \\ \mathbf{R}_1 & \mathbf{R}_0 & & \vdots \\ \vdots & & \ddots & \mathbf{R}_{-1} \\ \mathbf{R}_{D-1} & \cdots & \mathbf{R}_1 & \mathbf{R}_0 \end{bmatrix}. \tag{2.140}$$

Similarly Toeplitz-block matrices are defined as a $DQ \times DQ$ matrix $\bar{\bar{\mathbf{T}}}$ that is composed of Q^2 Toeplitz matrices $\mathbf{T}_{q_1,q_2}$ oriented in a $Q \times Q$ block matrix form as

$$\bar{\bar{\mathbf{T}}} = \begin{bmatrix} \mathbf{T}_{1,1} & \mathbf{T}_{1,2} & \cdots & \mathbf{T}_{1,Q} \\ \mathbf{T}_{2,1} & \mathbf{T}_{2,2} & & \vdots \\ \vdots & & \ddots & \mathbf{T}_{Q-1,Q} \\ \mathbf{T}_{Q,1} & \cdots & \mathbf{T}_{Q,Q-1} & \mathbf{T}_{Q,Q} \end{bmatrix}. \tag{2.141}$$

Any Toeplitz-block matrix can be transformed into a block-Toeplitz matrix by applying permutations as

$$\bar{\mathbf{T}} = \bar{\mathbf{J}}\bar{\bar{\mathbf{T}}}\bar{\mathbf{J}}^T \tag{2.142}$$

with the permutation matrix

$$\bar{\mathbf{J}} = \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_{D+1} & \mathbf{e}_{2D+1} & \cdots & \mathbf{e}_{(Q-1)D+1} & \mathbf{e}_2 & \mathbf{e}_{D+2} & \cdots & \mathbf{e}_{(Q-1)D+Q} \end{bmatrix} \tag{2.143}$$

where $\mathbf{e}_i$ is an elementary vector with all zeros except for unity at the ith element. The reverse is likewise true as

$$\bar{\bar{\mathbf{T}}} = \bar{\mathbf{J}}^T\bar{\mathbf{T}}\bar{\mathbf{J}} \tag{2.144}$$

since permutations are unitary, i.e.,

$$\mathbf{I} = \bar{\mathbf{J}}^T\bar{\mathbf{J}}.$$

Using the Toeplitz structure of each block of (2.141), (2.134) can again be leveraged to perform matrix-vector multiplication in the frequency domain with a cost of $\mathscr{O}(2Q^2D\log(2D))$ if each $\mathbf{T}_{q_1,q_2}$ is fully filled.

Additionally, if a block-Toeplitz matrix is PD, its inverse will inherit an structure similar to

(2.139) and can be characterized by

$$\bar{\mathbf{T}} \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_D \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \tag{2.145}$$

and

$$\bar{\mathbf{T}} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \\ \vdots \\ \mathbf{B}_D \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{I} \end{bmatrix} \tag{2.146}$$

where $\mathbf{A}_i$, $\mathbf{B}_i$, and $\mathbf{I}$ are $D \times D$ matrices [109, p. 348] and can be solved efficiently via multi-channel Levinson recursion with $\mathcal{O}(Q^3D^2)$ [109, p. 345] complexity. Note, when the blocks of $\bar{\mathbf{T}}$ are also Toeplitz and the main diagonal block is hermitian, then $\mathbf{B}_i = \mathbf{A}_i^H$. This yields the block inverse structure

$$\bar{\mathbf{T}}^{-1} = \bar{\mathbf{T}}_f \bar{\mathbf{E}}_f^{-1} \bar{\mathbf{T}}_f^H - \bar{\mathbf{T}}_b \bar{\mathbf{E}}_b^{-1} \bar{\mathbf{T}}_b^H \tag{2.147}$$

where

$$\mathbf{T}_f = \begin{bmatrix} \mathbf{A}_1 & 0 & \cdots & 0 \\ \mathbf{A}_2 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \mathbf{A}_D & \cdots & \mathbf{A}_2 & \mathbf{A}_1 \end{bmatrix}$$

and

$$\mathbf{T}_b = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ \mathbf{B}_D & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \mathbf{B}_2 & \cdots & \mathbf{B}_D & 0 \end{bmatrix}$$

are both block-Toeplitz matrices, while

$$\bar{\mathbf{E}}_f = \begin{bmatrix} \mathbf{A}_1 & \mathbf{0} & \cdots & 0 \\ \mathbf{0} & \mathbf{A}_1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{A}_1 \end{bmatrix}$$

and

$$\bar{\mathbf{E}}_f = \begin{bmatrix} \mathbf{B}_1 & \mathbf{0} & \cdots & 0 \\ \mathbf{0} & \mathbf{B}_1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{B}_1 \end{bmatrix}$$

are block diagonal matrices and matrix-vector multiplication with (2.147) can be performed with $\mathscr{O}(Q^2 D \log(2D))$.

2.8 Iterative Linear System Solvers

The solution to a $D \times D$ linear system of equations

$$\mathbf{A}\mathbf{x} = \mathbf{b} \tag{2.148}$$

can be equivalently posed as a convex QP when $\mathbf{A}$ is Hermitian PD as

$$\min_{\mathbf{x}^*} \quad \mathbf{x}^H \mathbf{A}\mathbf{x} - 2\,\mathrm{Re}\left\{\mathbf{x}^H \mathbf{b}\right\} \tag{2.149}$$

as taking the derivative w.r.t. $\mathbf{x}^*$ and setting equal to zero reveals that the optimal solution

$$\mathbf{x}_\star = \mathbf{A}^{-1}\mathbf{b}$$

satisfies (2.148). Thus, rather than performing generally expensive matrix decomposition/inverse to directly solve $\mathbf{A}^{-1}\mathbf{b}$, (2.149) can instead be minimized via convex optimization techniques. Additionally, stopping the optimization prior to convergence may provide a "good enough" solution that while not optimal, can be achieved at a much lower cost than a direct solver.

Steepest descent would be simple to implement, yet it only has linear convergence. Newton/quasi-Newton methods are faster, but have complexity/memory requirements that can be comparable to direct inversion. Alternatively, the linear conjugate gradient (CG) algorithm can provide superlinear convergence to the optimum [76, p. 101-132], with each iteration moving the optimal step in a direction conjugate to all previous directions. Unlike direct solutions, the per iteration complexity of PCG is dictated by matrix-vector multiplication with $\mathbf{A}$. Thus, if $\mathbf{A}$ has structure that enables fast matrix-vector multiplication (such as a Toeplitz structure), PCG can offer a significant reduction in computational cost. However, PCG may not be suitable to every linear system and can be unstable when numerical errors yield steps that lack conjugacy. Alternatively, the generalized minimal residual method (GMRES) [113] can offer improved stability for a higher per iteration cost as each step performs an additional decomposition to ensure conjugacy but requires a complexity increase of $\sim\mathscr{O}\left\{i^2\right\}$ where i is the iteration index.

Convergence for both approaches would otherwise be dictated by the condition number and eigenvalue clustering of $\mathbf{A}$, thus, a hermitian positive definite preconditioner matrix can $\mathbf{M}$ is typically introduced such that

$$\mathbf{M}^{-1}\mathbf{A}$$

has an improved condition number and the eigenvalues are clustered about unity. However, each iteration preconditioned CG (PCG) or preconditioned GMRES (PGMRES) then requires linear system solves involving $\mathbf{M}$, thus, $\mathbf{M}$ must be carefully selected to both improve convergence while not significantly degrading the overall computational cost.

Chapter 3

Efficient Straddle-Robust Least Squares Mismatched Filtering

While the LS MMFs from Section 2.2 offered improved sidelobe suppression for a minor MML when properly designed, their efficacy tends to degrade for lower oversampling factors (κ). Additionally, solving each LS MMF required a matrix inverse with a brute force complexity of $\mathcal{O}(D^3)$ which can hinder real time viability. Here we seek to incorporate further robustness measures to support near critically sampled architectures ($\kappa \approx 1$) so that RX processing can support (near) full bandwidth utilization. Computationally efficient solutions and alternatives are then proposed and assessed using closed-loop and open-air experimental measurements.

3.1 Statistical Straddling Compensation for LS MMFs

While incorporating beamspoiling into the ISL MMF helps retain the filter nominal resolution, the lack of an explicit constraint on high frequency components results in a sensitivity to range straddling more so than the template-based LS MMFs. As mentioned in Section 2.2.2, [55] proposed averaging of multiple MMFs for different straddled replicas of the waveform, however, doing so lacks optimality in terms of the original MMF QPs and necessitates multiple loopback collects. [114] examined straddle compensation for the particular case of an LFM waveform via statistical phase uncertainties which has a convenient closed form solution. For an arbitrary waveform, this same framework is not necessary extensible unless the waveform can be characterized by a specific phase-basis, i.e., [50, 115–118]. Additionally, this phase-based approach does not incorporate distortion induced by the RX anti-alias filter, which may be significant for waveforms less spectrally contained than chirps. Alternatively, using hardware-in-the-loop measurements that in-

corporate the RX anti-aliasing filter or waveform modeling the incorporates TX DAC modeling [119], a structure-based interpolation can be posed to express the waveforms as a function of the sub-sample delay γ to account for the continuous reality. This generic framework can then be used to evaluate the straddling uncertainty, which is then directly incorporated into a compensated QP.

To incorporate straddling uncertainty in the LS MMF framework, the hessian of the QPs prior to beamspoiling ($\mathbf{T} = \mathbf{S}\mathbf{S}^H$) is expanded using the continuous representation of the convolution matrix from (2.15) and rearranged using the distributive property of the expectation operator as

$$\begin{aligned}
\mathbf{T} &= \mathscr{E}\{\mathbf{S}(\gamma)\mathbf{S}^H(\gamma)\} \\
&= \mathscr{E}\left\{\sum_{i=-N+1}^{D-1} \mathbf{J}_i \mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\mathbf{J}_i^H\right\} \\
&= \sum_{i=-N+1}^{D} \mathbf{J}_i \mathscr{E}\{\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\}\mathbf{J}_i^T \\
&= \sum_{i=-N+1}^{D} \mathbf{J}_i \left(\int_{-0.5}^{0.5} \mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\,d\gamma\right)\mathbf{J}_i^T \\
&= \sum_{i=-N+1}^{D} \mathbf{J}_i \mathbf{R}_\gamma \mathbf{J}_i^T
\end{aligned} \tag{3.1}$$

where is $\mathbf{J}_i$ a reversing and shifting matrix such that matrix-vector multiplication with an arbitrary $N \times 1$ vector $\mathbf{b}$ yields a $D \times 1$ vector

$$\mathbf{x} = \mathbf{J}_i \mathbf{b}$$

as

$$x_d = \begin{cases} b_{i-d} & 1 \le i-d \le N \\ 0 & else \end{cases} \tag{3.2}$$

for $d = 1, 2, \cdots, D$ which has most N non-zero elements.

With this representation, incorporating range straddling uncertainty necessitates solving for the straddled-signal correlation matrix

$$\mathscr{E}\{\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\} = \int_{-0.5}^{0.5} \mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\,d\gamma = \mathbf{R}_\gamma \tag{3.3}$$

to construct the compensated $\mathbf{T}$. Likewise, it may be desirable (but not necessary) to compensate the zero-delay filter $\mathbf{v}$ by constructing it with the expected value of the straddled waveform

$$\mathbf{s}_\gamma = \mathscr{E}\{\mathbf{s}(\gamma)\}. \tag{3.4}$$

as

$$\mathbf{v} = \begin{bmatrix} \mathbf{0}^T & \mathbf{s}_\gamma^T & \mathbf{0}^T \end{bmatrix}^T \tag{3.5}$$

in the same manner as (2.24). Note, $\mathbf{R}_\gamma$ should be a Hermitian matrix resulting in a $\mathbf{T}$ that retains a Banded-Hermitian Toeplitz structure where the kth diagonal is found as

$$t_k = \begin{cases} \sum_{i=1}^{N-k} \left[\mathbf{R}_\gamma\right]_{i,i+k} & 0 \le k \le N-1 \\ \sum_{i=1}^{N+k} \left[\mathbf{R}_\gamma\right]_{i-k,i} & -N+1 \le k \le -1 \\ 0 & else \end{cases} \tag{3.6}$$

which corresponds to the sum of the diagonals of $\mathbf{R}_\gamma$. The beamspoiling downdate in the ISL objective will now correspond to the sum of k shifted $\mathbf{R}_\gamma$, which is not necessarily rank-k as in (2.37).

Estimating the waveform expected value and covariance matrix can be performed in a variety of different manners. As modern AWGs generally support sample rates much larger than the RX bandwidth, sub-sample shifts can be incorporated into the TX emission and their response can be measured in a hardware-in-the-loop configuration to accumulate straddled snapshots and form a sample correlation matrix. Doing so would simultaneously characterize other sources of distortions like amplifier nonlinearities and may further enhance filter robustness. However, this process requires collection of many loopback measurements (on the order of N to ensure a full rank estimate) for every single TX waveform which may not always be practical.

Rather than a measurement based approach, straddling can be approximated via interpolation models so that a structure based estimate of $\mathbf{R}_\gamma$ can be formed without needing multiple loopback

measurements, though the expectation fidelity will be bound by the interpolation accuracy. Using a loopback collection of $\mathbf{s}$ or a sufficient model of the TX upsampling and RX anti-aliasing filter, a sub-sample shift will correspond to multiplication by a sinusoid in the frequency domain. Thus, a truncated sinc-interpolation approach can be used to approximate straddling via

$$\begin{aligned}\mathbf{s}(\gamma) &= \frac{1}{Z}\begin{bmatrix} \mathbf{I}_N & \mathbf{0}_{N\times(Z-N)} \end{bmatrix}\mathbf{F}^H\left(\mathbf{d}(\gamma)\odot\mathbf{F}\begin{bmatrix} \mathbf{s} \\ \mathbf{0}_{(Z-N)\times 1} \end{bmatrix}\right) \\ &= \frac{1}{Z}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\mathbf{d}(\gamma)\odot\mathbf{s}_f\right)\end{aligned} \tag{3.7}$$

where $\mathbf{F}$ is a $Z\times Z$ DFT matrix where Z corresponds to the nearest power of two,

$$\mathbf{d}(\gamma) = \exp(-j2\pi\gamma\tilde{\mathbf{z}}) \tag{3.8}$$

is a complex exponential vector corresponding to a sub-sample temporal shift γ where

$$\tilde{\mathbf{z}} = \begin{bmatrix} 0 & 1 & \cdots & \frac{Z}{2}-1 & -\frac{Z}{2} & -\frac{Z}{2}+1 & \cdots & -1 \end{bmatrix}^T$$

is a $Z\times 1$ vector corresponding to the normalized frequency index, $\mathbf{s}_f$ is the discrete Fourier transform of the padded waveform, $\mathbf{I}_{N\times Z}$ is a selection matrix composed of the first N rows of a $Z\times Z$ identity matrix. This form of interpolation assumes that the waveform is truly band-limited and aliased energy will propagate error. Thus, it is important that the $\mathbf{s}$ used for estimation sufficiently models the upsampling and downsampling from TX-to-RX or is measured via a hardware-in-the-loop collect.

As $\mathbf{d}(\gamma)$ is a complex exponential, the integral in (3.3) results a sinc-structure since

$$\exp(a)\exp(b) = \exp(a+b) \tag{3.9}$$

and

$$\begin{aligned}\int_{-0.5}^{0.5}\exp\{j2\pi bx\}dx &= \int_{-0.5}^{0.5}\exp\{-j2\pi bx\}dx \\ &= \int_{-0.5}^{0.5}\cos(2\pi bx)dx + \int_{-0.5}^{0.5}\sin(2\pi bx)dx \\ &= \frac{\sin(2\pi bx)}{2\pi b}\Big|_{-0.5}^{0.5} + 0 \\ &= \frac{\sin(\pi b)}{\pi b} \\ &= \mathrm{sinc}(b)\end{aligned} \tag{3.10}$$

can be computed from the symmetric bounds of the integral and Euler's identity where

$$\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$$

is the "so called" sinc-function.

Thus, (3.3) can be directly evaluated as

$$\mathbf{R}_\gamma = \frac{1}{Z^2}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\mathbf{T}_\gamma \odot \mathbf{s}_f\mathbf{s}_f^H\right)\mathbf{F}\mathbf{I}_{N\times Z}^T \tag{3.11}$$

where $\mathbf{T}_\gamma$ is a symmetric sinc-function covariance matrix taper (CMT) corresponding to

$$\mathbf{T}_\gamma = \mathrm{sinc}\left(\frac{1}{Z}\tilde{\mathbf{z}}\mathbf{1}^T - \frac{1}{Z}\mathbf{1}\tilde{\mathbf{z}}^T\right) \tag{3.12}$$

and $\mathbf{R}_\gamma$ is again a hermitian matrix. Likewise, the expected waveform can be computed as

$$\begin{aligned}\mathbf{s}_\gamma &= \mathscr{E}\{\mathbf{s}(\gamma)\} \\ &= \int_{-0.5}^{0.5}\frac{1}{Z}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\mathbf{d}(\gamma)\odot\mathbf{s}_f\right)d\gamma \\ &= \frac{1}{Z}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\tilde{\mathbf{d}}\odot\mathbf{s}_f\right)\end{aligned} \tag{3.13}$$

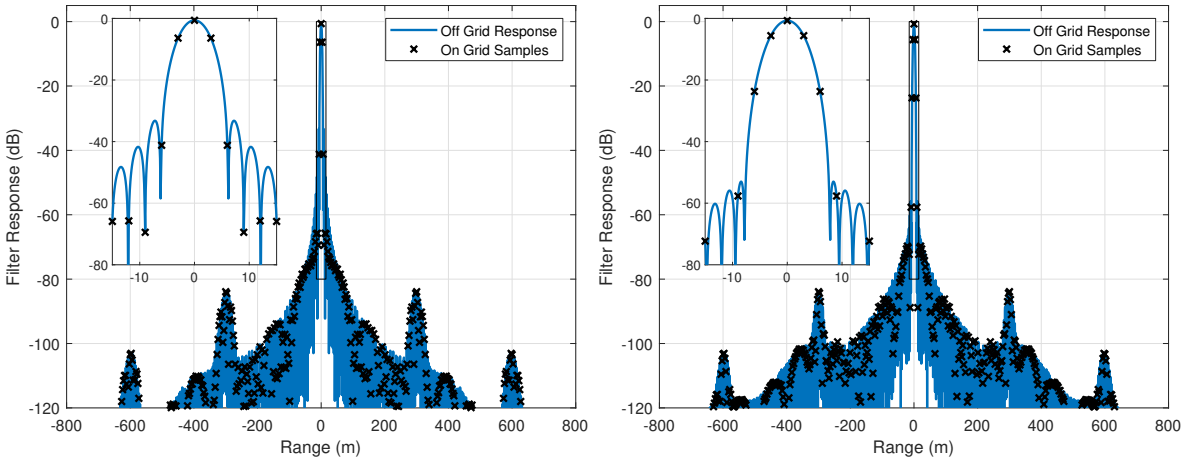

Figure 3.1: Simulated Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ Uncompensated (left) and Range-Straddle Compensated (right)

where

$$\tilde{\mathbf{d}} = \text{sinc}\left(\frac{1}{Z}\tilde{\mathbf{z}}\right)$$

acts as a frequency domain sinc taper.

Using the straddle compensated model, the $\kappa = 1.5$ ISL LS MMF from Figures 2.13-2.14 was regenerated using the same beamspoiling, filter length, and regularization factors as before. Figure 3.1 illustrates this LS MMF response on the left and compares it to the baseline uncompensated case on the right. Here we see that the straddle compensated formulation has slightly broadened the filter-response mainlobe, but now the on-grid samples are much more commensurate with the off-grid continuum (recall the uncompensated case effectively places the on-grid sample points on the nulls of the continuum). Looking at the respective filter PSs in Figure 3.2, we see that the compensated model has a more gradual spectral roll-off from $\sim 10 - 23$MHz and the edge of the band now drops straight down as opposed to the rippling behavior present at $\sim$24.5MHz with the uncompensated approach.

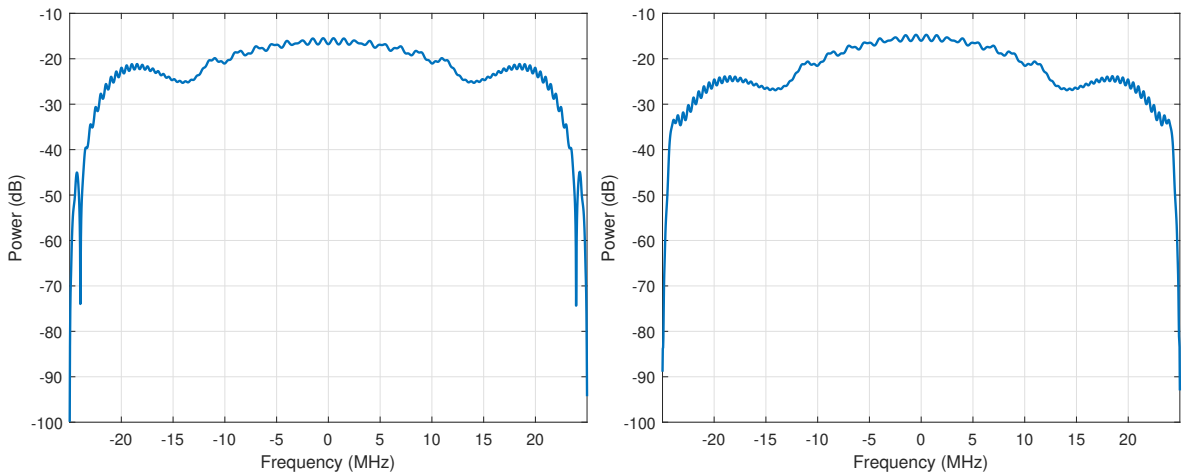

Figure 3.2: Simulated Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$ Uncompensated (left) and Range-Straddle Compensated (right)

3.1.1 Low-Rank Approximate Straddling Compensation for Efficient ISL Beamspoiling

While beamspoiling can be incorporated by subtracting the mainlobe indices from (3.1) as

$$\begin{aligned} \tilde{\mathbf{T}} &= \sum_{i=-N+1}^{D} \mathbf{J}_i \mathbf{R}_\gamma \mathbf{J}_i^T - \sum_k \mathbf{J}_k \mathbf{R}_\gamma \mathbf{J}_k^T \\ &= \mathbf{T} - \mathbf{T}_{ML} \end{aligned}, \tag{3.14}$$

it can no longer be expressed as a rank-k downdate of $\mathbf{T}$, which will prevent use of the efficient Toeplitz solvers discussed in the next section. However, examining the eigenvalues of $\mathbf{R}_\gamma$ reveals a steep drop-off (see Figure 3.3 for an example) and while it is no longer rank-1 (e.g., $\mathbf{s}\mathbf{s}^H$ is rank-1 by definition) it can be well approximated by a low-rank matrix. Thus, in practice a truncated $\mathbf{R}_\gamma$ is expressed as

$$\mathbf{R}_\gamma = \sum_{a=1}^{A} \mathbf{u}_a \lambda_a \mathbf{u}_a^H \tag{3.15}$$

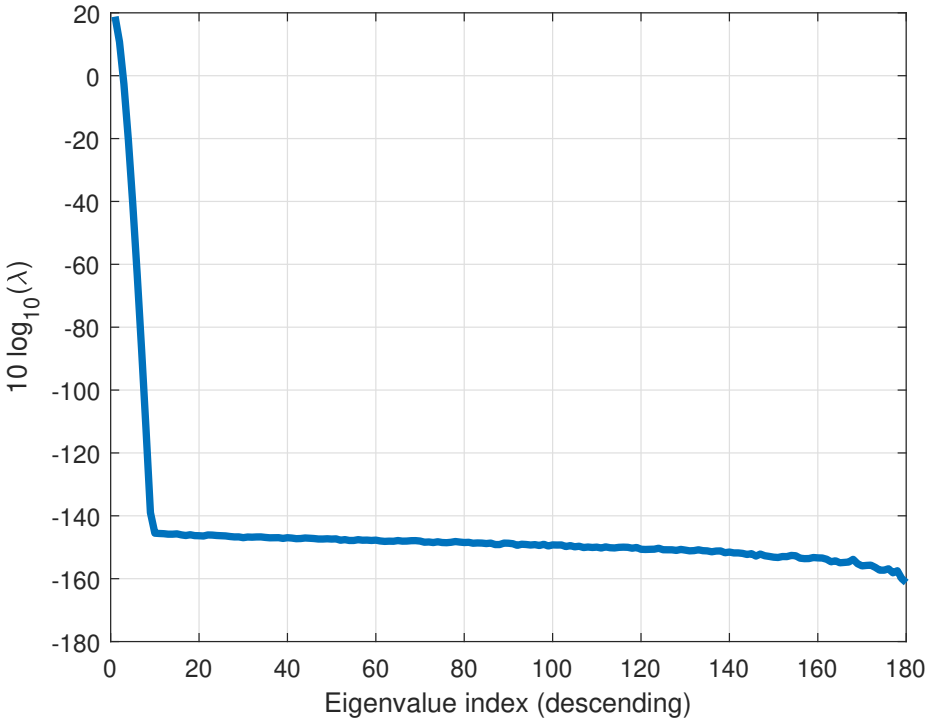

Figure 3.3: Eigenvalues of $\mathbf{R}_\gamma$ for LFM waveform with $\kappa = 1.5$

using the dominant A eigenvalues (λ_a) and eigenvectors $\mathbf{u}_a$ such that beamspoiling can be expressed as a rank-(Ak) downdate via

$$\begin{aligned}\tilde{\mathbf{T}} &= \sum_{i=-N+1}^{D} \mathbf{J}_i \mathbf{R}_\gamma \mathbf{J}_i^T - \sum_k \mathbf{J}_k \mathbf{R}_\gamma \mathbf{J}_k^T \\ &= \mathbf{T} - \sum_a \sum_k \mathbf{J}_k \mathbf{u}_a \lambda_a \mathbf{u}_a^H \mathbf{J}_k^T\end{aligned} . \tag{3.16}$$

For the rest of this work, this rank-(Ak) downdate scaling is treated as implicit when straddling compensation is used and just referenced as a rank-k operation.

3.2 Efficient Solvers for LS MMF QPs

As the LS MMFs problem require a the solution to a system of linear equations proportional to the filter length D, it can be a computationally expensive procedure for longer filters or larger TB. However, as both the template-based and ISL MMFs are driven by the banded hermitian Toeplitz

matrix $\mathbf{T}$ (plus a rank-k downdate for the ISL case), the matrix structure can be leveraged to solve the filter efficiently. Here these efficient direct solutions are split into two cases depending on the relative size of D and N. While there are other potentially faster solvers than the methods shown here (such as a banded levinson approach [120] or general superfast Toeplitz solvers [111, 112], these approaches can be numerically unstable or only provide benefits for very longer filters. Thus, the solvers in this section are focused on empirically stable techniques that provide complexity reduction for filter lengths $D \sim 100-10,000$.

The first approach leverages the hermitian banded structure of $\mathbf{T}$ to perform a banded Cholesky decomposition of the regularized Hessian (to ensure positive definiteness) as

$$\mathbf{T} + \sigma^2 \mathbf{I} = \mathbf{L}\mathbf{L}^H \tag{3.17}$$

where the $D \times D$ lower-triangular matrix

$$\mathbf{L} = \begin{bmatrix} L_{1,1} & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots & & \vdots \\ L_{N,1} & \cdots & L_{N,N} & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & L_{D,1} & \cdots & L_{D,N} \end{bmatrix} \tag{3.18}$$

retains a lower-matrix bandwidth of N and can be computed with $\mathscr{O}(N^2 D)$ by ignoring the zero entries in $\mathbf{T}$. For the ISL case, beamspoiling can then be incorporated via a Cholesky downdate [121], which again by ignoring the zero entries can be performed with $\mathscr{O}(kND)$ complexity and retains the banded structure. The LS MMF can the be solved via forward backward substitution using $\mathbf{L}$ which has $\mathscr{O}(ND)$ complexity. Alternatively, when $N^2 > D$ levinson-durbin recursion [109, p. 85] can be used to solve for the first column of the inverse hessian $\mathbf{a}$ via (2.138) (again including regularization) with $\mathscr{O}(D^2)$ complexity. The Toeplitz structure of (2.139) can then be used

to perform frequency domain matrix matrix-vector multiplication with $\mathcal{O}(D\log(2D))$ complexity. The beamspoiling downdates for the ISL case can then be performed via Woodbury's inverse identity in Appendix A.1 with $\mathcal{O}(k^3 + kD\log(2D))$ complexity or the iterative update through PCG or GMRES with $\mathcal{O}(k^2 + k^2 D\log(2D))$ can be applied if the inverse update is not numerically stable.

3.2.1 Iterative Approximate LS MMFs

While direct solutions to the LS MMF problem provide a solution within numerical precision, the imperfect nature of the RX model means the optimal filter will always be "over-fit" and may have a theoretical response beyond the achievable fidelity. Alternatively, an iterative approximate filter can be solved rapidly and with appropriate stopping conditions, yields little-to-no degradation in practice. As the regularized Hessian is a banded hermitian Toeplitz matrix, the matrix-vector operations in CG or GMRES can be computed in the frequency domain with $\mathcal{O}((D+N-1)\log(D+N-1))$ complexity (or more accurately the next power of 2 from $D+N-1$). Additionally, Circulant preconditioners can significantly improve convergence and only require an additional $\mathcal{O}((D)\log(D))$ operation each iteration.

The first Circulant preconditioner (denoted $\mathbf{C}_1$) has a frequency domain diagonal formed using the PS of the waveform zero-padded to D samples plus regularization as

$$\mathbf{c}_1 = |\mathbf{F}\mathbf{v}|^2 + \sigma^2\mathbf{1} \tag{3.19}$$

where the inverse preconditioner matrix-vector multiplication is applied as

$$\mathbf{C}_1^{-1}\mathbf{x} = ifft\{\, fft\{\mathbf{x}\} \oslash \mathbf{c}_1 \,\}. \tag{3.20}$$

The second Circulant preconditioner is similar, but leverages a form that minimizes the Frobenius norm between the preconditioner and hessian [122] via

$$\mathbf{c}_2 = \mathbf{F}\tilde{\mathbf{t}} + \sigma^2\mathbf{1} \tag{3.21}$$

where

$$\tilde{\mathbf{t}} = \begin{bmatrix} t_0 & \frac{D-1}{D}t_1 & \cdots & \frac{D-N+1}{D}t_{N-1} & 0 & \cdots & 0 & \frac{D-N+1}{D}t_{-N+1} & \cdots & \frac{D-1}{D}t_{-1} \end{bmatrix}^T \tag{3.22}$$

can be applied when $D > 2N$.

For the template objective, PCG or PGMRES can be used directly on (2.34) as the QP is already in a form that is amenable to iterative solvers. However, the linear constraint in the ISL objective complicates things as the Lagrangian will have a negative eigenvalue. To yield a tractable iterative implementation one may define confine the optimization to the a null-space of the constraint in a manner similar to [76, p. 461-462], however, forming orthogonal steps has a tendency to be unstable and may require "reset" operations internal that slow convergence. Alternatively, the first column of $\left(\mathbf{T}+\sigma^2\mathbf{I}\right)^{-1}$ can be approximated via the iterative method and (2.139) is then used to apply the beamspoiling and constraint in the same manner as in Section 3.2.

3.3 Closed-Loop Experimental Assessment

To evaluate the efficiency of the iterative approximate LS MMFs and evaluate the efficacy of the range straddling compensation, two coherent processing intervals (CPIs) of FM waveforms were generated on a Tektronix arbitrary waveform generator (AWG). Both CPIs consisted of 1000 pulses with a center frequency of 3.5 GHz, 4.5 micro-second pulse duration, a 3dB bandwidth of 33 MHz (time-bandwidth product of 330), and a 5 kHz PRF (2.25% duty cycle). The first CPI followed a standard repeated LFM chirp waveform framework, while the second used a waveform agile set of non-repeating optimized random FM (RFM) waveforms with a Gaussian power spectral density [73]. The AWG was connected to a bandpass filter and class A amplifier to model TX amplification, followed by an attenuator, bandpass filter, low-noise amplifier (LNA), and then a Rhode & Schwarz spectrum analyzer with a sample rate of 200 MHz (κ=6) so that closed-loop measurements of the waveforms on real hardware could be collected. Additionally, a kaiser window based FIR bandpass filter was applied to these loopback measurements followed by decimation by a factor of 4 to

Measurement Parameters		
Parameter	Variable	Value
Sample Rate	f_s	50MHz
3dB Bandwidth	B	33.3MHz
Pulse Duration	T	4.5μseconds
Samples in Waveform	N	276
Fast-Time Filter Samples	D	$4N$

Table 3.1: Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms

yield a sample rate of 50 MHz ($\kappa = 1.5$) corresponding to measurements more representative of a system with near full bandwidth utilization. A tabularized list of variables and parameters used in this experimental assessment is provided in Table 3.1. Note, the number of waveform samples corresponds to a slightly longer duration than the pulse to account for effects from the anti-aliasing filter.

Using these loopback measurements, PCG was applied to solve for the Hermitian-Toeplitz inverse representation from (2.138) for a filter length $D = 4N$, a regularization $\sigma = \frac{10^{-5}}{D}\text{Tr}\{\mathbf{S}\mathbf{S}^H\}$, and the range straddling compensation from Section 3.1 to give a sense of convergence for PCG. Figures 3.4-3.5 illustrate the residual error of PCG for $\kappa = 1.5$ and $\kappa = 6$ respectively when using no preconditioner, (3.19) (denoted Precond 1), and (3.21) (denoted Precond 2) for the LFM and 1000 PROFM waveforms respectively. Both preconditioners offer improved convergence over CG where (3.19) typically converges faster for the $\kappa = 1.5$ case while (3.21) offers a modest improvement for the $\kappa = 6$ case.

Next, to asses the efficacy the approximate LS MMF solutions, PCG with (3.19) as a preconditioner and a fixed number of iterations was used to generate LS MMFs for the $\kappa = 1.5$ case. First, PCG was used to directly solve the template-based LS MMF QP with no statistical straddle compensation. The selected templates were a super-gaussian of order 6 for the LFM and a gaussian for the PROFM waveforms. Both templates were selected to have the same 3dB bandwidth of 33MHz and an additional tapering was placed on the high frequency components of the template to retain spectral containment and provide a response robust to range-straddling. PCG was then ran for 10, 20, and 100 iterations to solve the respective QPs. The left side of Figures 3.6-3.7 illustrates the

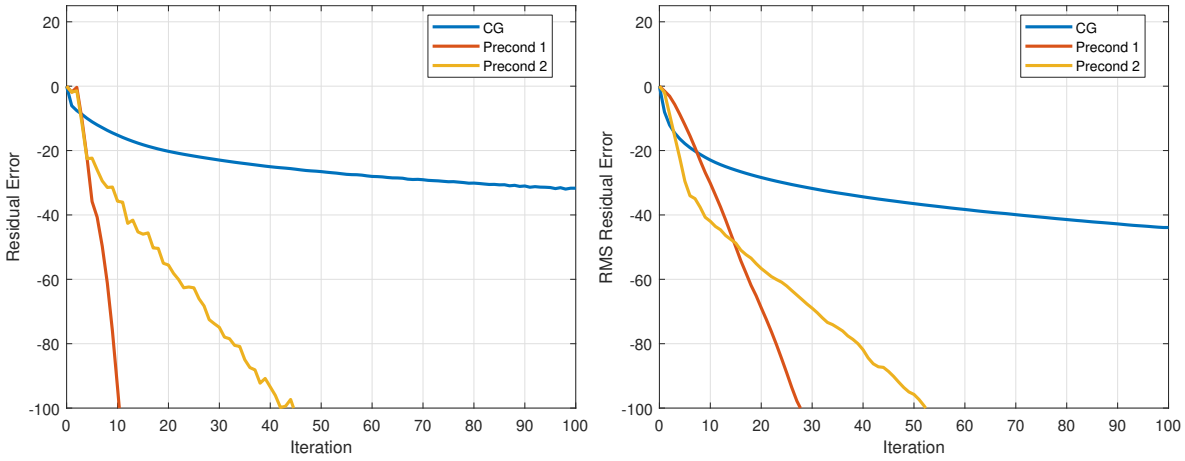

Figure 3.4: Residual Error vs Iteration of PCG using Circulant Preconditioners for single LFM (left) and 1000 PROFM (right) Waveforms with $\kappa = 1.5$

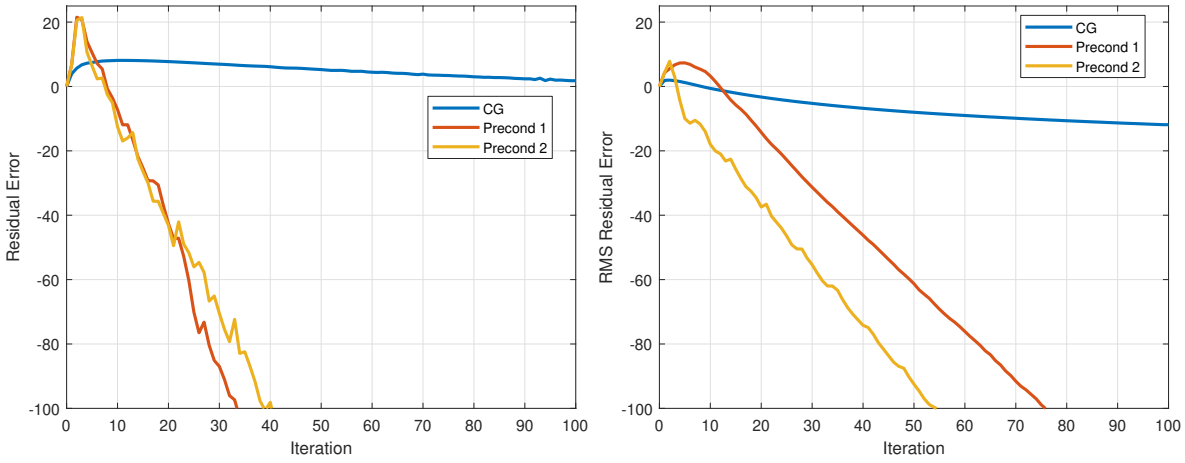

Figure 3.5: Residual Error vs Iteration of PCG using Circulant Preconditioners for single LFM (left) and 1000 PROFM (right) Waveforms with $\kappa = 6$

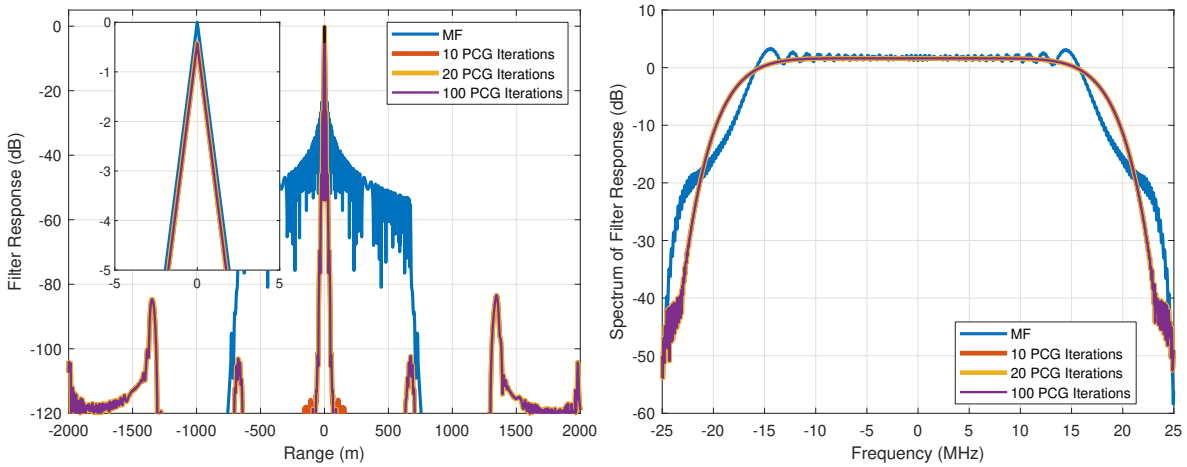

Figure 3.6: Loopback Template-based LS MMF Response for LFM waveform and $\kappa = 1.5$

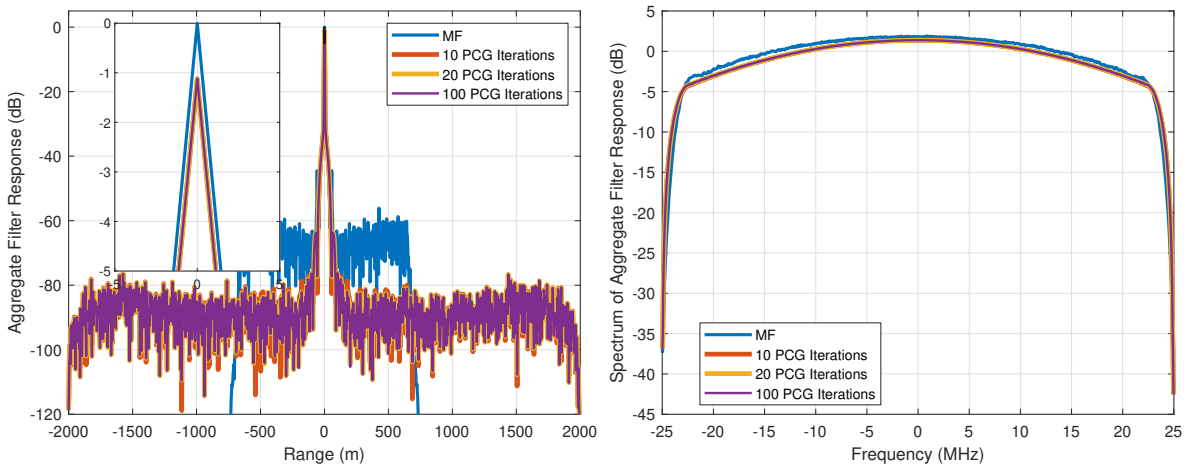

Figure 3.7: Loopback Template-based LS MMF Response for PROFM waveforms and $\kappa = 1.5$

coherently combined aggregate filter response for the LFM waveform and PROFM waveforms respectively. Additionally, the corresponding waveform-filter frequency response are shown on the left of Figures 3.6-3.7. For this first test case, solving the QP directly yields rapid convergence for both waveform sets and even after only 10 iterations, the filter response has largely converged outsize of some minor ripples near delays $\sim$100 meters. The frequency response likewise shows little to no errors.

A similar procedure was then applied to the uncompensated ISL LS MMF. Here, PCG was used to approximate the Toeplitz inverse structure via (2.138). The resulting approximate inverse

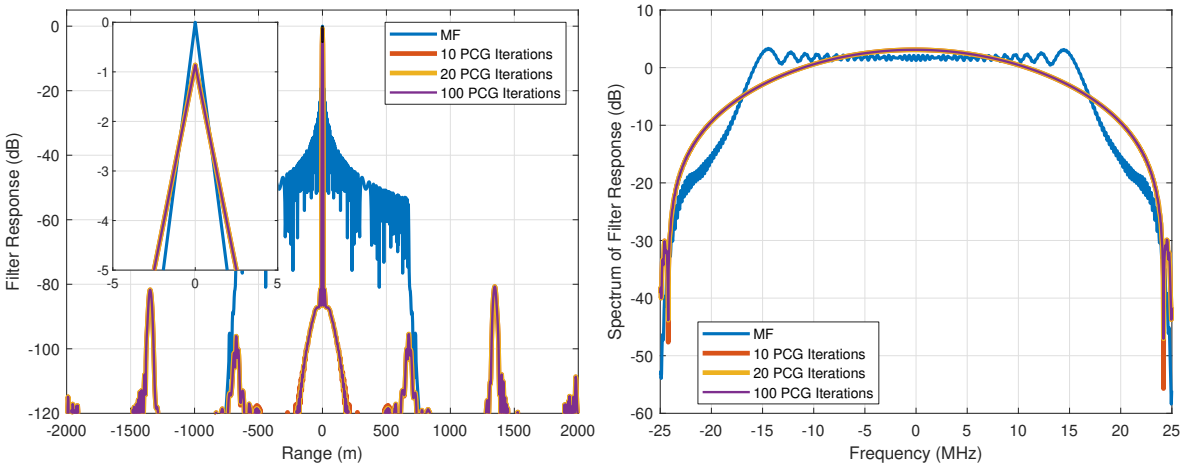

Figure 3.8: Loopback ISL LS MMF Response for LFM waveform and $\kappa = 1.5$

was then used to efficiently form

$$\left(\mathbf{T} - \tilde{\mathbf{S}}\tilde{\mathbf{S}}^H\right)^{-1}\mathbf{v}$$

via the beamspoiled downdate via fast inverse Teopltiz matrix-vector multiplication and Woodbury's identity from (A.2) for $k = 3$. The resulting vector was then used to form the approximate ISL LS MMF via (2.42). Figures 3.8-3.9 illustrates the resulting filter responses and frequency response for the LFM and PROFM waveforms respectively with the same orientation as the template case. Here for the LFM case, we see that 10 iterations was once again enough to converge sufficiently yield virtually no difference between the resulting filters. As expected, the lack of range straddling compensation does produces a ripple near the edge of the band ($\pm$25MHz) which will result in raised sidelobes when a scatterer is straddled in range. For the PROFM case, we see that 10 iterations was not sufficient to achieve convergences and subsequent application of Woodbury's identity results in an unstable update (for at least one of the MMFs) leading to significant MML and an inaccurate filter. After 20 iterations, the respective filters have once again converged resulting in a shape commensurate with what is expected for the LS MMFs without straddle compensation.

Lastly, the straddle compensated formulation LS MMFs were formed and $\mathbf{R}_\gamma$ was approximated using $A = 3$ dominant eigenvalues for each waveform. Here, PCG was again used to approximate the Toeplitz inverse structure via (2.138), Woodbury's (A.1) for $k = 9$ (update size has increased

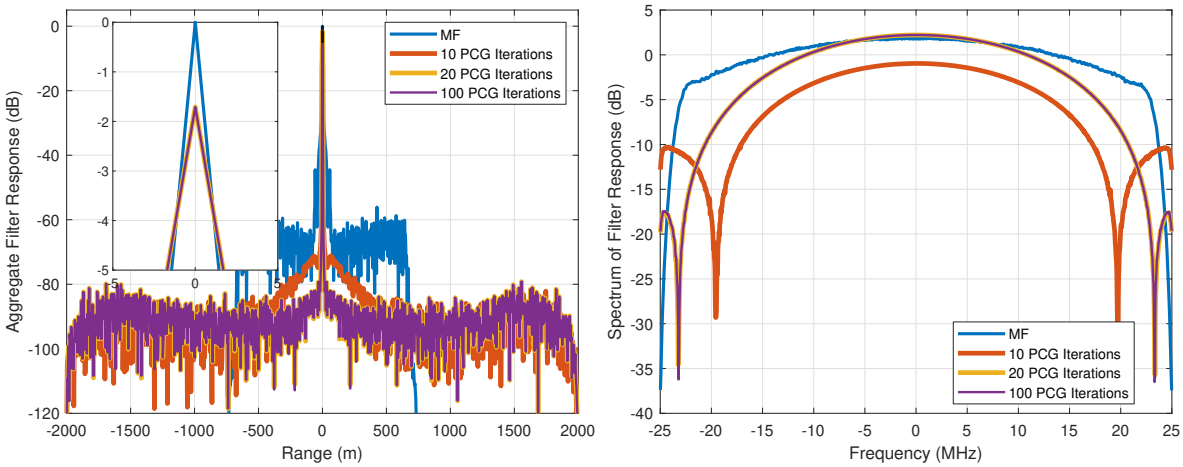

Figure 3.9: Loopback ISL LS MMF Response for PROFM waveforms and $\kappa = 1.5$

by a size of A relative to the uncompensated case), and (2.42) were then used to solve the straddle compensated ISL LS MMFs. Figures 3.10-3.11 illustrates the resulting filter responses and frequency response for the LFM and PROFM waveforms respectively with the same orientation as the template case. Again for the LFM case, we see that 10 iterations was enough converge sufficiently and there is only minor difference between the filters. Range straddling compensation has removed the high frequency ripple effect and the resulting filter has a smooth frequency roll-off and sharp null at the edge of the band. For the PROFM case, surprising 10 iterations was enough to effectively converge for all cases, which is unexpected since the increase in downdate dimensionality suggests a higher likelihood for instability with Woodbury's matrix identity. It is possible that incorporation of straddle compensation yields a Toeplitz structure that is better suited for PCG, but the exact reasoning of this phenomena is unclear. As with the LFM case, the frequency response of the PROFM waveforms is has a smooth roll-off and does not have much high frequency contribution. That being said, the notch on the edge of the band is raised relative to the LFM case likely due to the variable nature of the frequency response pulse-to-pulse, e.g. some PROFM LS MMFs will perform better than others since each waveform is only locally-optimal and unique.

Table 3.2 provides the order of complexity for each of the respective solvers where PCG with a Circulant preconditioner offers and order of magnitude improvement in cost relative to a Levinson approach for 10 iterations. It should be noted, that the operations involved in PCG are well opti-

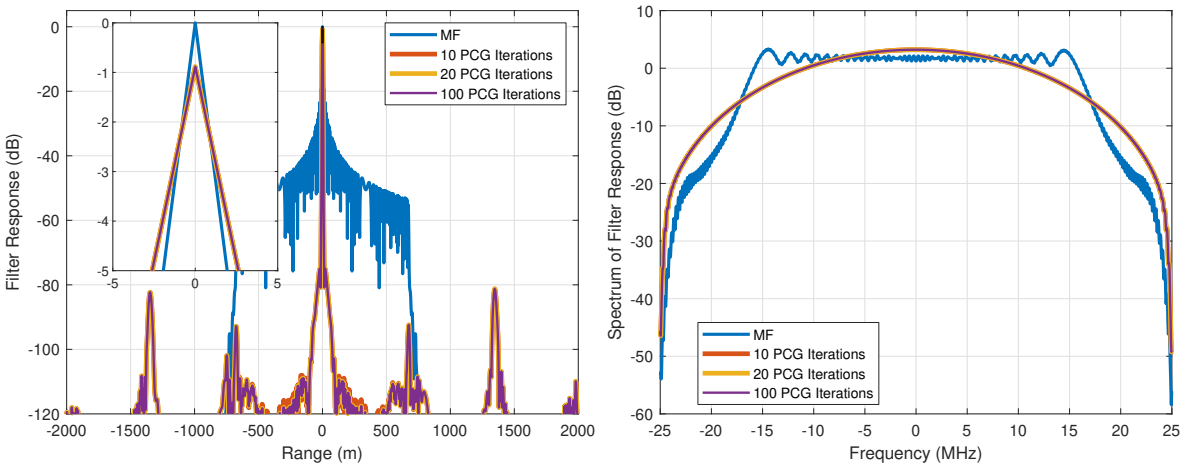

Figure 3.10: Loopback Straddle Compensated ISL LS MMF Response for LFM waveform and $\kappa = 1.5$

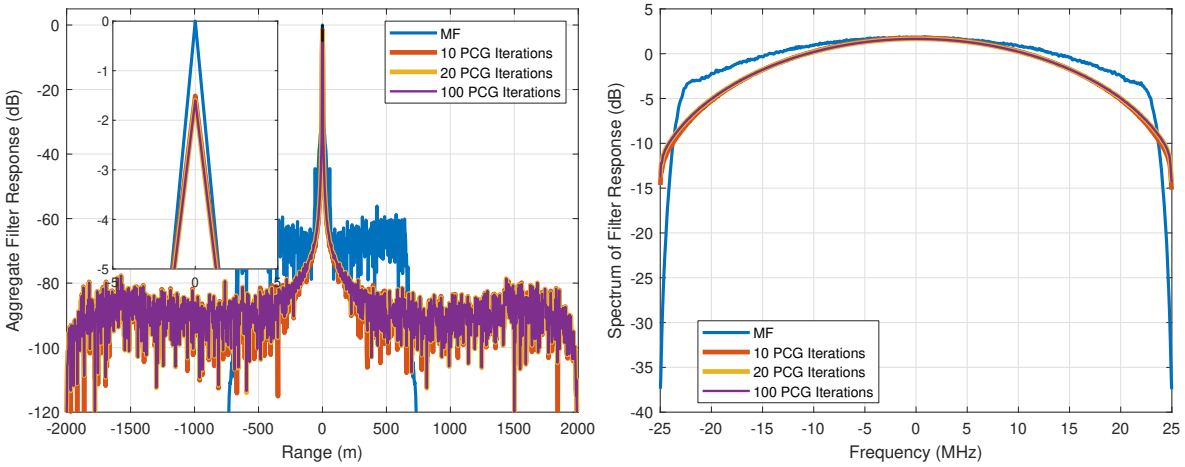

Figure 3.11: Loopback Straddle Compensated ISL LS MMF Response for PROFM waveforms and $\kappa = 1.5$

Toeplitz Solve Order of Complexity	
Direct	1.35E+09
Levinson	1.22E+06
Banded	8.44E+07
PCG (10 iter)	2.55E+05
PCG (20 iter)	5.11E+05
PCG (100 iter)	2.55E+06

Table 3.2: Complexity for solving the Toeplitz System of Equations

mized and amenable to hardware acceleration (e.g., the FFTs) and the actual runtime will likely see improvements much greater than a single order of magnitude [123].

3.4 Open-Air Experimental Assessment

Next the attenuator from the closed-loop setup was removed and the respective transmitter and receiver chains were connected to parabolic antennas. This test setup was taken to the roof of Nichols hall at the University of Kansas and placed in a pseudo monostatic configuration where TX and RX happens simultaneously (see Fig. 3.12). The antennas were pointed at an intersection with moving vehicles $\sim$1000 meters away and the two CPIs from the previous section were sequentially transmitted to collect open-air experimental measurements of a scene with moving vehicles and stationary clutter. Note, these the selected pulse duration and test setup resulted in direct path sidelobes that extend $\sim$700 meters.

The receive measurements were then orientated into fast-time slow-time measurement matrices, and the MF each of the respective LS MMFs from the previous sections were applied to estimate the scattering profile of the scene. Figure 3.13 illustrates the RMS average across pulses for the range profile for the LFM (left) and PROFM (right) CPIs. For the LFM CPI, we see that each of the LS MMF approaches were successful at suppressing the direct path sidelobes. Examining the subfigure in the top right, we can clearly see that while the template-based yields a lower but non-negligible sidelobe roll-off while the beamspoiled approaches have a sharper sidelobe roll-off. While somewhat minor for this case, the direct path is straddled enough in range such

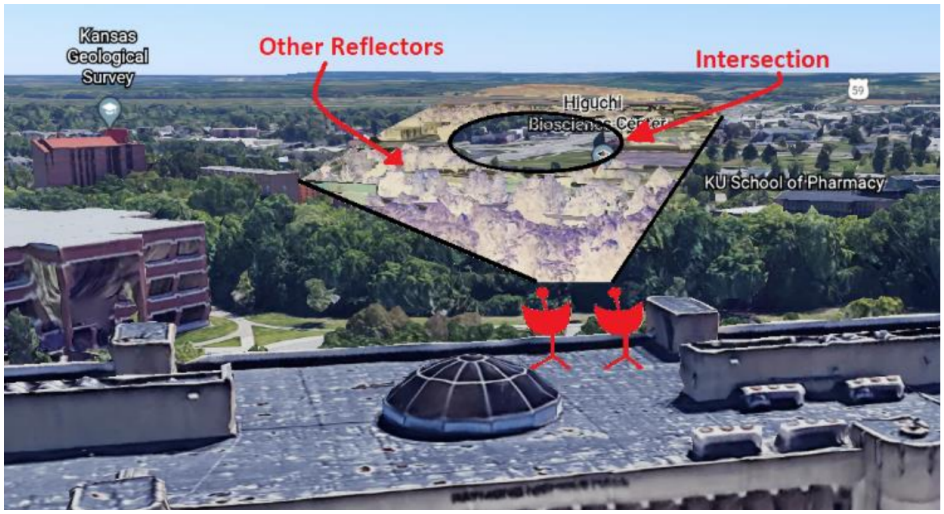

Figure 3.12: Open-Air Experimental Test Setup

that the uncompensated LS MMF (yellow) a roll-off ∼5 dB higher than the straddle compensated LS MMF.

For the PROFM CPI, we see that the LS MMFs are able to suppress the direct path sidelobes by ∼25dB further than the matched filter, though we see that the longer filter length has extended these sidelobes in range and the sidelobe level at the intersection (∼1000 meters) is actually higher with the LS MMFs than the MF. Looking at the subfigure, the effect from range straddling is clearly visible for the uncompensated ISL case, where the near in roll-off is up to ∼20dB higher than the compensated ISL MMF. However, the template-based and straddle compensated ISL MMFs see $\sim 3dB$ and $\sim 1.5dB$ higher sidelobes at the intersection that the uncompensated ISL MMF. For the template MMF, this result follows intuition as the LS MMFs expend DoF attempting to match the template, rather than explicitly minimizing sidelobes. For the compensated ISL case, the slightly higher sidelobes naturally arise as the QP and $\mathbf{R}_\gamma$ now must consider a continuum of delay samples as opposed to the $D+N-1$ on-grid delays.

Next, Doppler processing was performed across each range profile and a Taylor window was applied along slow-time to reduce Doppler sidelobes (but not RSM). Corresponding range-Doppler estimates are then shown in Figures 3.14-3.19 where each response was normalized so that the color-axis corresponds to SNR. The left side of each figure contains a "zoomed out" look such

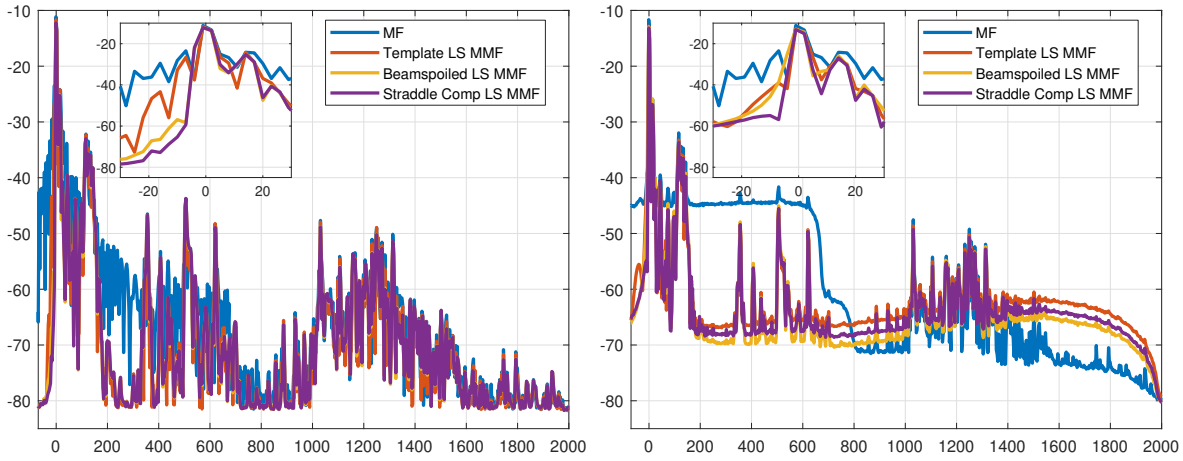

Figure 3.13: LS MMF RMS Open-Air Range Profile Estimates $\kappa = 1.5$

that the direct-path response is visible, while the right side corresponds to a focused look of the intersection with moving vehicles.

First examining the LFM CPI using the MF for range estimation in Figure 3.14, we see that use of a stationary platform and repeated waveform has localized the clutter response to zero-Doppler. While the direct path (0 meters) Doppler sidelobes have been largely suppressed due to the Taylor window, there remains residual energy smearing 20dB above the noise floor and to the model uncertainty floor as even using loop-back measurements to form the MF can not perfectly replicate reality as phase-noise an nonlinearities will still be present. This residual is on average $\sim$85dB down from the direct path and effectively set a local floor on the dynamic range. Looking at the intersection, the moving vehicles are clearly visible and distinctly separate from the 0-velocity clutter ridge. The SNR of the cars is strong enough such that their range sidelobes are clearly visible and they may be potentially obfuscating some of the scene characteristics. The $\sim$1040 meter Doppler slice has a distinct periodic return spaced every $\sim$2.5m/s, likely corresponding to the Doppler signature of a spinning fan-blade from an air-conditioning unit on one of the buildings. There is also a notable non-zero Doppler extent of the clutter spanning $1050 - 1250$ which may correspond to either slow moving vehicles or gradual sway of the light poles at the intersection. This MF LFM case will act as a baseline for the other approaches as it is representative of a traditional radar architecture.

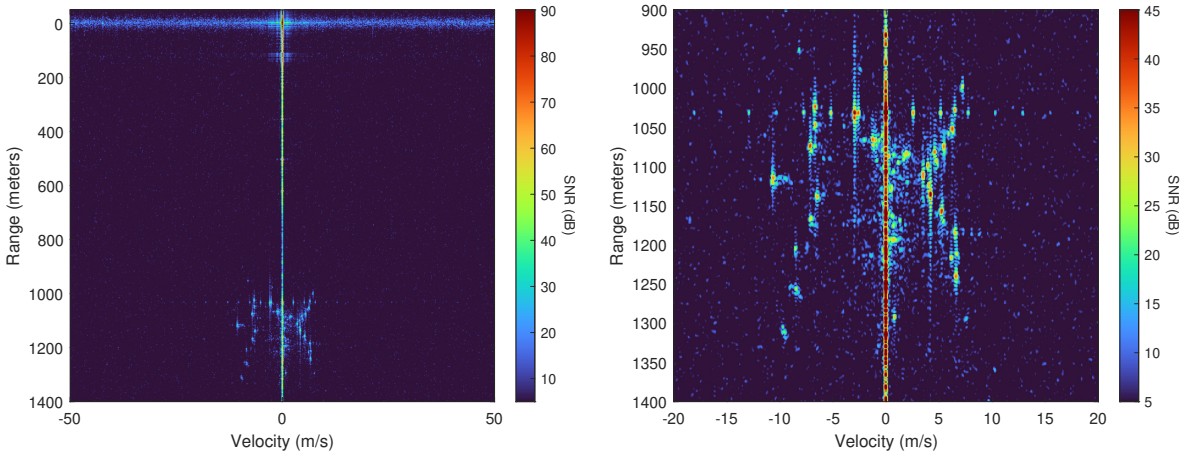

Figure 3.14: Range-Doppler Estimate for LFM CPI using the MF

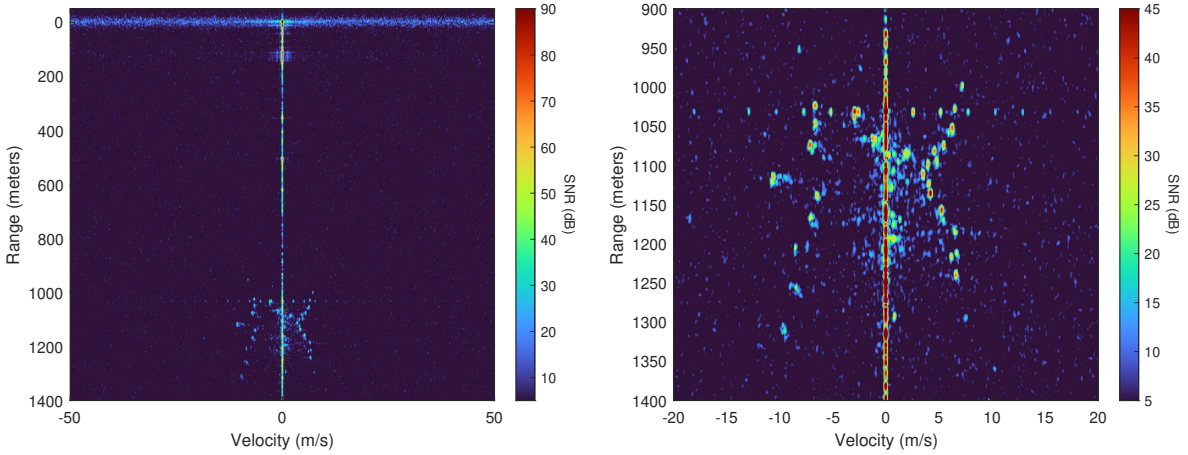

Figure 3.15: Range-Doppler Estimate for LFM CPI using the Straddle Compensated ISL LS MMF

Next, we observe the Straddle Compensated LS MMF for the LFM CPI shown in Figure 3.15. Here the zoomed-out look appears largely the same, though the direct-path range sidelobes are now further down. The residual from model uncertainty remains unchanged and still acts as a fundamental bound. Looking at the intersection, the benefit of applying the LS MMF is now clearly visible and the cars' range sidelobes have effectively been mitigated and the Doppler signature are more apparent, though the range mainlobe has been very slightly broadened. Overall, these filter provide an enhanced view of the intersection, for a modest MML of 0.9dB.

As the other LS MMFs for the LFM CPI perform fairly similar response, we now elect to move on and examine the PROFM CPI. First, the MF estimate is provided in Figure 3.16 where we

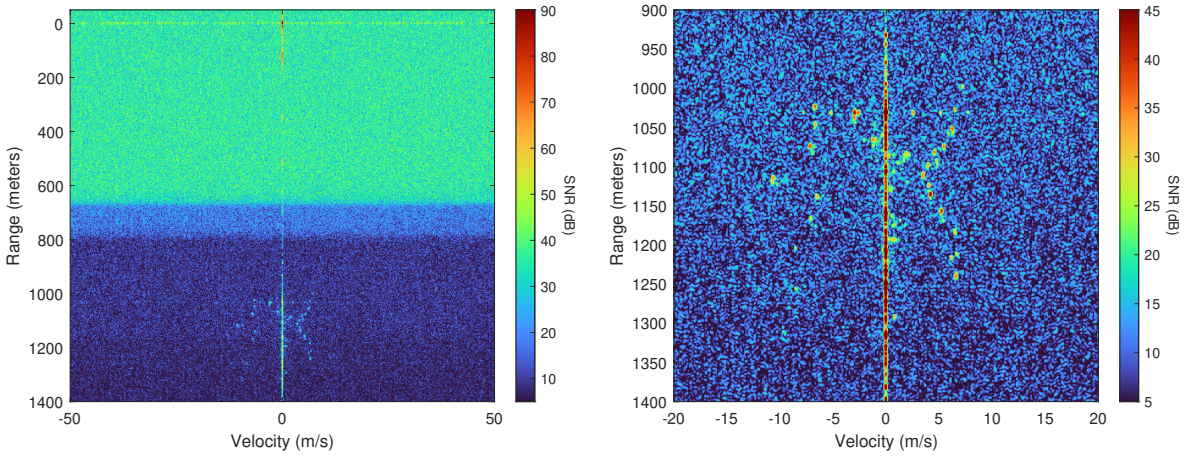

Figure 3.16: Range-Doppler Estimate for PROFM CPI using the MF

see that the pulse-to-pulse variation of the direct path range-sidelobes has produced RSM 40dB above the noise floor. This RSM in turn will effectively act as range varying noise floor which will naturally degrade detection. The model uncertainty floor is still likewise present, though now the RSM makes it less apparent that the LFM CPI. Examining the intersection, we see that while the SNR of the vehicles is largely the same, the nearby clutter has produced RSM and the "effective" SNR has been degraded by $\sim$10dB compared the the LFM baseline. However, unlike the LFM MF, the car's do not have any apparent range sidelobes as they are now lower ($\sim$60dB down from the mainlobe return) but spread in both range and Doppler.

Next, the range-Doppler response for the template-based LS MMFs and PROFM CPI is shown in Figure 3.17. Here the MMF has lowered the pulse-to-pulse range sidelobes by $\sim 20dB$, which proportionally reduced the direct path RSM. However, as the filter length has increase by a factor of 4, the direct path response is now extended in range to the point where the intersection RSM is now higher that the MF case.

Examining the ISL LS MMF (without compensation) response in Figure 3.18, we see a similar result where the MMF has reduced the direct path RSM and by directly minimizing sidelobes (as opposed to matching a desired response) offers a slight improvement to RSM levels. However, for this case a distinct Doppler-streaking behavior is now present corresponding to large scatters which is due to the nature of how the beamspoiling was incorporated into the ISL objective. As

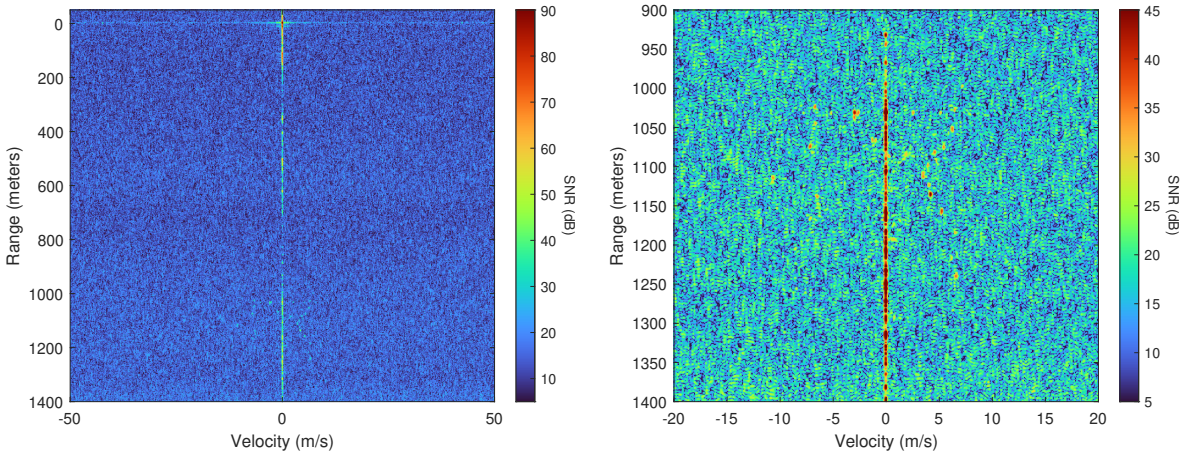

Figure 3.17: Range-Doppler Estimate for PROFM CPI using the Template-based LS MMF

the ISL objective only constrains the zero-delay response to unity and "does not care" about the beamspoiled delay indices, the resulting filter may induce a mainlobe response that varies pulse-to-pulse. From the Fourier relationships, this equivalently corresponds to variation in the spectral shape of the filter. This subsequently induces a mainlobe-modulation effect similar to that of RSM and can not be compensated for using traditional Doppler-tapering or clutter cancellation. As the template-based LS MMFs explicitly desires a specific response, this effect is implicitly avoided. Additionally, since the mainlobe is higher energy than the sidelobes, mainlobe modulation results in fairly significant (but localized) Doppler smearing. An extreme case of this phenomena was observed in [124] where the transmit waveform spectral shape varied pulse-to-pulse to enable spectral cohabitation, and this mainlobe modulation was present for both the LS MMF and the MF. Post-processing techniques can be applied to help mitigate this effect [125, 126] but are outside the scope of this work. Examining the intersection, we see that again the RSM has been further suppressed relative to the template-based MMF, but the strong scatterers induce visible mainlobe modulation.

The final range-Doppler map in Figure 3.19 uses the Straddle compensated ISL LS MMF where, as expected from 3.13, this approach sees a slightly higher direct path RSM ($\sim$1.5dB) compared to the uncompensated ISL LS MMFs. However, incorporating straddle compensation significantly reduces the mainlobe modulation, even though the degree of beamspoiling is identical

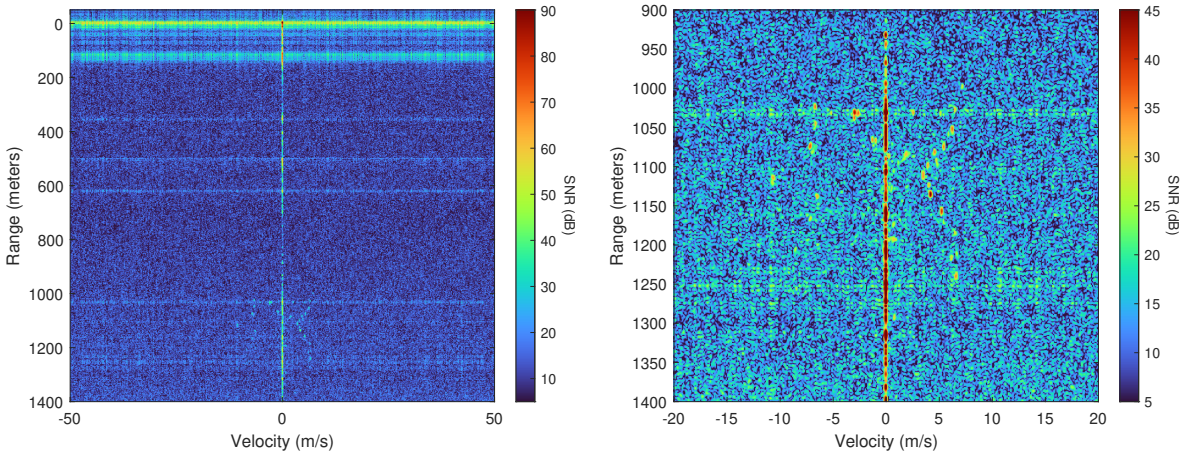

Figure 3.18: Range-Doppler Estimate for PROFM CPI using the ISL LS MMF

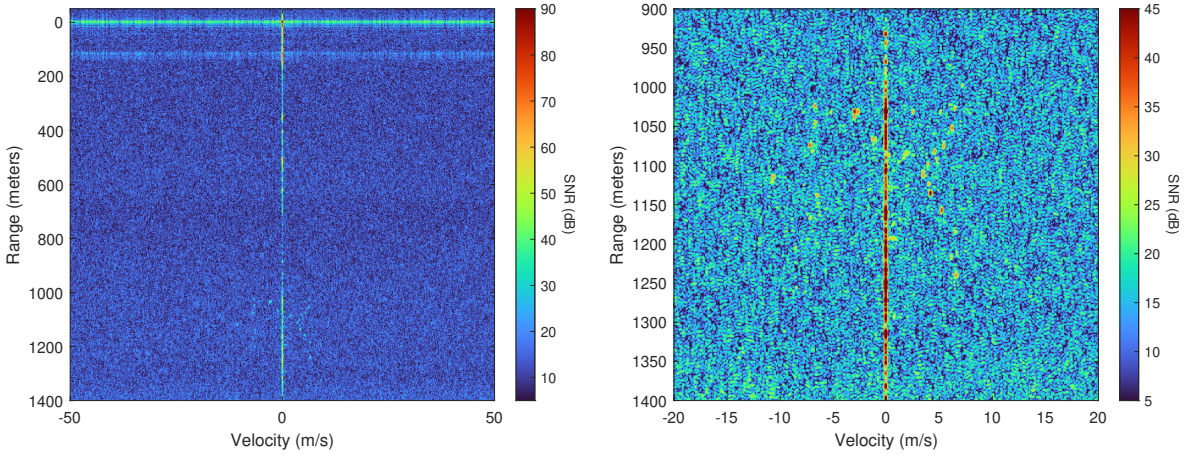

Figure 3.19: Range-Doppler Estimate for LFM CPI using the Straddle Compensated ISL LS MMF

that of the uncompensated case. As the straddle compensation effectively prevents high frequency components in the filter-response, the possible MMF solution space has been slightly restricted and thus, the resulting MMFs are more likely to share a consistent spectral shape. Thus, straddle compensation implicitly reduces mainlobe modulation effects.

Chapter 4

Mismatched Complementary on Receive Filtering

While the LS MMF approaches can provide an improvement in sidelobe level relative to the MF, the objective functions are inherently overdetermined as discrete non-circular convolution results in a response strictly longer than the filter length, which limits how well as single filter can perform. As the filter dimensionality D increases sidelobe levels can be decreased further, however, longer filters spread energy in range as seen in Figures 3.17-3.19. Additionally, longer filters inherently increase computational burden and may face additional challenges when model deviations like pulse-eclipsing or multiple time around clutter are present [81].

From a waveform design perspective, Golay [127] first proposed a method to address this lack of dimensionality in which complementary sets of codes were designed so coherent combination of their matched filter responses would (theoretically) achieve complete sidelobe removal. However, these codes relied on instantaneous phase transition and were sensitive to hardware distortion and straddling effects. In [128] it was experimentally demonstrated that complementary subsets of FM waveform can instead be designed to prove robustness towards these limitations, but could not achieve "perfect" sidelobe removal.

[129] extended the LS MMF work of Akroyd for binary phase codes to design sets of complementary LS MMFs to minimize sidelobes for phase coded waveforms. Here this same philosophy is taken on receive in which the degrees of freedom provided by agile CPI of random FM waveforms are used to jointly design subsets of LS MMFs for FM waveforms that achieve a complementary sidelobe cancellation and are robust to hardware and model mismatch effects. Additionally, the incorporation of the range straddling uncertainty from Section 3.1, the resulting filters yield robust performance for near critically sampled architectures ($\kappa \approx 1$) and help prevent

mainlobe modulation. However, there is no free lunch, as the complementary sidelobe cancellation necessitates a reduction in unambiguous velocity and becomes less efficacious for scatterers when there are Doppler shifts present within a complementary subset. This approach is denoted mismatched complementary-on receive filtering, or MiCRFt (pronounced like "Mycroft" Holmes, fictional elder brother of Sherlock Holmes) [130, 131].

4.1 MiCRFt Model

Consider the design of Q mismatched filters for a corresponding set of Q unique waveforms such that, when the respective filter responses are combined, a complementary condition is achieved. Here, the template-based LS MMF model of (2.34) is extended as

$$J = \left\| \mathbf{p} - \sum_{q=1}^{Q} \mathbf{S}_q^H \mathbf{w}_q \right\|_2^2 \tag{4.1}$$

where $\mathbf{S}_q$ and $\mathbf{w}_q$ are the convolution matrix and complementary filter for the qth waveform respectively. Note, MiCRFt inherently relies on the use of unique waveforms within a complementary subset as the repeated case will effectively simplify to a single LS MMF via linearity, i.e.,

$$\begin{aligned}
\bar{\bar{\mathbf{r}}}^T &= \sum_{q=1}^{Q} \mathbf{r}_q^T \\
&= \sum_{q=1}^{Q} \mathbf{w}_q^H \mathbf{S}_q \\
&= \sum_{q=1}^{Q} \mathbf{w}_q^H \mathbf{S} \\
&= \left(\sum_{q=1}^{Q} \mathbf{w}_q^H \right) \mathbf{S} \\
&= Q \mathbf{w}^H \mathbf{S}.
\end{aligned}$$

The summations can be posed as a single matrix multiplication by concatenating the convolution matrices and filters as $\bar{\bar{\mathbf{S}}}^H = \begin{bmatrix} \mathbf{S}_1^H & \mathbf{S}_2^H & \cdots & \mathbf{S}_Q^H \end{bmatrix}$ and $\bar{\bar{\mathbf{w}}} = \begin{bmatrix} \mathbf{w}_1^T & \mathbf{w}_2^T & \cdots & \mathbf{w}_Q^T \end{bmatrix}^T$ which, after incorpo-

ration of regularization, can be solved via the equivalent QP

$$\min_{\bar{\tilde{\mathbf{w}}}^*} \quad \bar{\tilde{\mathbf{w}}}^H \left(\bar{\mathbf{T}} + \sigma^2 \mathbf{I}\right) \bar{\tilde{\mathbf{w}}} - 2\,\mathrm{Re}\left\{\bar{\tilde{\mathbf{w}}}^H \bar{\mathbf{b}}\right\} \tag{4.2}$$

where

$$\bar{\mathbf{T}} = \bar{\tilde{\mathbf{S}}}\bar{\tilde{\mathbf{S}}}^H = \begin{bmatrix} \mathbf{S}_1\mathbf{S}_1^H & \mathbf{S}_1\mathbf{S}_2^H & \cdots & \mathbf{S}_1\mathbf{S}_Q^H \\ \mathbf{S}_2\mathbf{S}_1^H & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ \mathbf{S}_Q\mathbf{S}_2^H & \cdots & \cdots & \mathbf{S}_Q\mathbf{S}_Q^H \end{bmatrix} \tag{4.3}$$

is now a $DQ \times DQ$ Toeplitz-block matrix and

$$\bar{\mathbf{b}} = \bar{\tilde{\mathbf{S}}}\mathbf{p} = \begin{bmatrix} \mathbf{S}_1\mathbf{p} \\ \mathbf{S}_2\mathbf{p} \\ \vdots \\ \mathbf{S}_Q\mathbf{p} \end{bmatrix} \tag{4.4}$$

with

$$\bar{\tilde{\mathbf{w}}} = \left(\bar{\mathbf{T}} + \sigma^2 \mathbf{I}\right)^{-1} \bar{\tilde{\mathbf{S}}}\mathbf{p} \tag{4.5}$$

corresponding to the optimal solution. Likewise, the ISL MiCRFt QP will be of the form

$$\begin{aligned} \min_{\bar{\tilde{\mathbf{w}}}^*} \quad & \bar{\tilde{\mathbf{w}}}^H \left(\bar{\mathbf{T}} - \bar{\tilde{\mathbf{T}}} + \sigma^2 \mathbf{I}\right) \bar{\tilde{\mathbf{w}}} \\ \text{s.t.} \quad & \bar{\tilde{\mathbf{w}}}^H \bar{\tilde{\mathbf{v}}} = Q \end{aligned} \tag{4.6}$$

where

$$\bar{\tilde{\mathbf{T}}} = \begin{bmatrix} \tilde{\mathbf{S}}_1\tilde{\mathbf{S}}_1^H & \tilde{\mathbf{S}}_1\tilde{\mathbf{S}}_2^H & \cdots & \tilde{\mathbf{S}}_1\tilde{\mathbf{S}}_Q^H \\ \tilde{\mathbf{S}}_2\tilde{\mathbf{S}}_1^H & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ \tilde{\mathbf{S}}_Q\tilde{\mathbf{S}}_2^H & \cdots & \cdots & \tilde{\mathbf{S}}_Q\tilde{\mathbf{S}}_Q^H \end{bmatrix} \tag{4.7}$$

is the rank-k beamspoiling downdate and $\bar{\bar{\mathbf{v}}} = \left[\mathbf{v}_1^T\ \mathbf{v}_2^T\ \cdots\ \mathbf{v}_Q^T\right]^T$ is the $DQ \times 1$ vector of concatenated zero-padded waveforms. (4.6) can again be solved using the KKT identity (Appendix $A.2$) and has the solution

$$\bar{\bar{\mathbf{w}}} = \frac{Q}{\bar{\bar{\mathbf{v}}}^H \left(\bar{\mathbf{T}} - \tilde{\bar{\mathbf{T}}} + \sigma^2 \mathbf{I}\right)^{-1} \bar{\bar{\mathbf{v}}}} \left(\bar{\mathbf{T}} - \tilde{\bar{\mathbf{T}}} + \sigma^2 \mathbf{I}\right)^{-1} \bar{\bar{\mathbf{v}}}. \tag{4.8}$$

Prior to incorporating regularization or a QC, the objective function attempts to minimize the error in $D+N-1$ equations using QD degrees of freedom, resulting in an underdetermined system with infinite optimal solutions. With regularization, the solution space collapses to a single set of filters that can suppress sidelobes (after coherent combination) to a much greater degree than a set of individual LS MMFs (even without diagonal loading) due to the complementary property of the filter response.

To illustrate the benefit of complementary design, two PROFM waveforms with a pulse duration of 2μs and bandwidth of 33.3MHz was generated at a sample rate of 10GHz. A FIR anti-aliasing filter was then applied followed by decimation to a yield simulated loopback configuration with spectral oversampling factors of $\kappa = 1.5$ and $\kappa = 6$. The MFs were then generated with the $\gamma = 0$ loopback waveform and applied to generate the aggregate filter response (coherent combination of both responses) shown in Figure 4.1 with $\kappa = 1.5$ and $\kappa = 6$ corresponding to the left and right plots respectively. The black xs correspond to the "on-grid" response ($\gamma = 0$) while the blue trace includes the continuum of off-grid shifts to better illustrate the filter performance in practice. The respective MF PS is likewise provided in 4.2 where the enforced spectral containment of the bandpass filter is again visible (high frequency components go to zero). Next, ISL LS MMF were generated for each waveform with $k = 2ceil\{\kappa\}$, $\sigma = \frac{10^{-3}}{D}\mathrm{Tr}\{\mathbf{S}\mathbf{S}^H\}$ and a length of $D = 4N$. The resulting solutions oriented into a subset and then normalization in terms of their aggregate MML as

$$\bar{\bar{\mathbf{w}}}_{\text{norm}} = \frac{Q}{\sqrt{\bar{\bar{\mathbf{w}}}^H \bar{\bar{\mathbf{w}}}} \sqrt{\bar{\bar{\mathbf{v}}}^H \bar{\bar{\mathbf{v}}}}} \bar{\bar{\mathbf{w}}} \tag{4.9}$$

such that the response will have a peak corresponding to the subsets MML. The aggregate filter response and filter PSs shown in Figures 4.3 and 4.3 respectively, where the LS MMF yields range

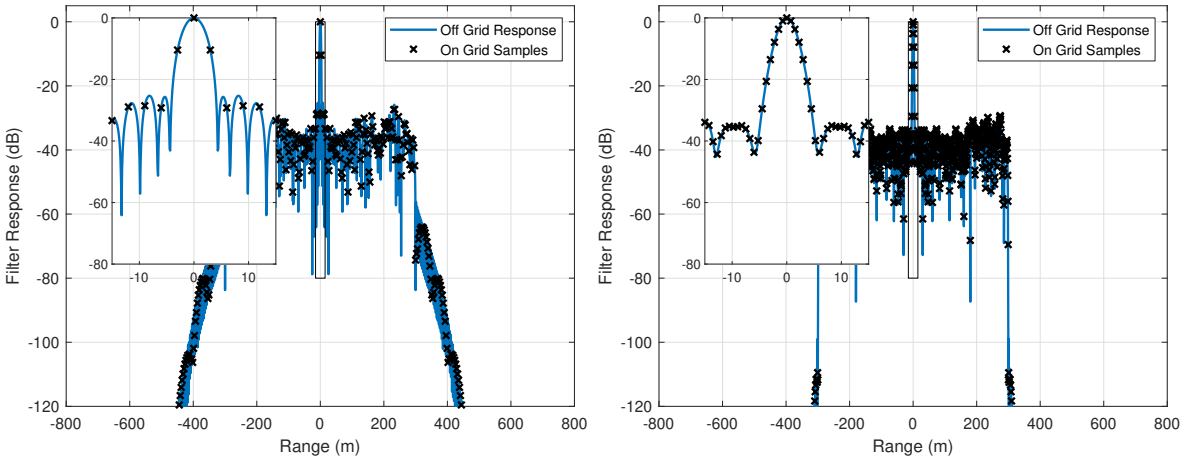

Figure 4.1: Simulated Loopback MF Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

sidelobes are $\sim$20dB lower than the MF. Next, a complementary subset of ISL objective MiCRFt filters were designed using the same parameters and applied in the same manner to generate the aggregate filter response and filter PSs in Figures 4.5 and 4.6 respectively. Here the complementary combination offers a significant reduction in sidelobe levels (40dB down from the MF) for no additional MML. However, the lack of straddling compensation results a fairly significant rippling/raising of near-in sidelobes for the $\kappa = 1.5$ case, though the off-grid response still generally on par or better than the non-complementary ISL LS MMFs. Leveraging a template-based MiCRFt formulation can provide robustness to straddling as illustrated in Figures 4.7 and 4.7, but it is once again desirable to incorporate straddling robustness into the MiCRFt ISL formulation.

4.2 Incorporating Range Straddling Robustness

Range straddling robustness for MiCRFt can be incorporated in a similar manner to the LS MMF, but now within a block matrix context. For the straddling model, the main-diagonal blocks (4.7) of retain the same form of (3.1), but the off diagonal blocks now have a cross-correlation component across pulses, (i.e. $\mathbf{T}_{a,b} = \mathbf{S}_a\mathbf{S}_b^H$ for $i \neq j$. Fortunately, $\mathbf{T}_{i,k}$ still retains a banded-Toeplitz structure (though it is not longer Hermitian) and the procedure of estimating the straddling cross-correlation

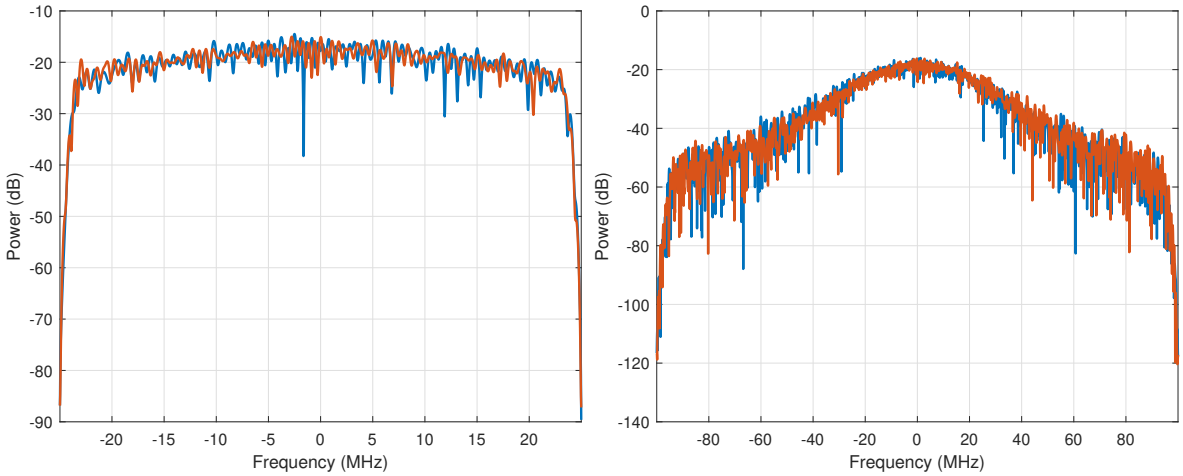

Figure 4.2: Simulated Loopback MF PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

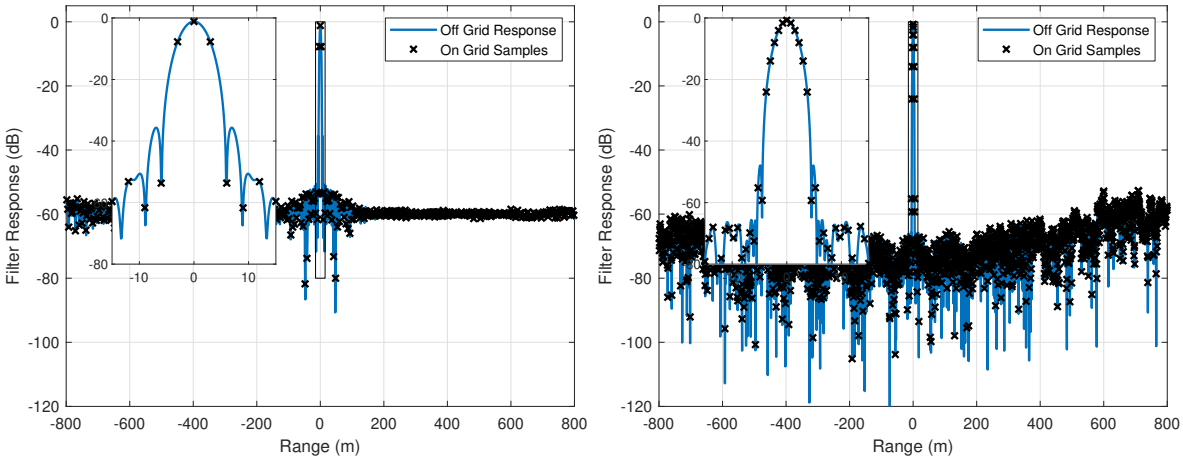

Figure 4.3: Simulated Loopback ISL LS MMF Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

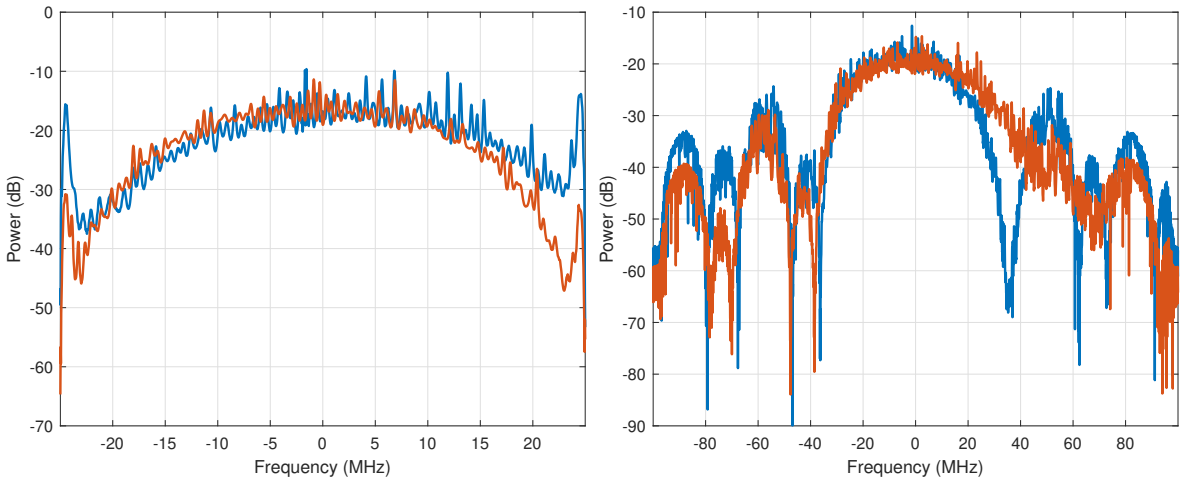

Figure 4.4: Simulated Loopback ISL LS MMF PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

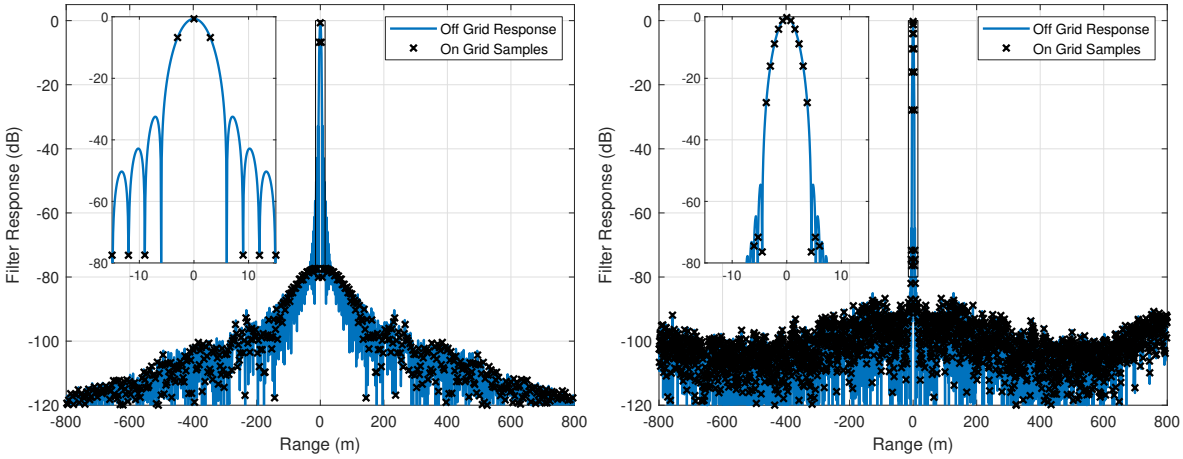

Figure 4.5: Simulated Loopback ISL MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

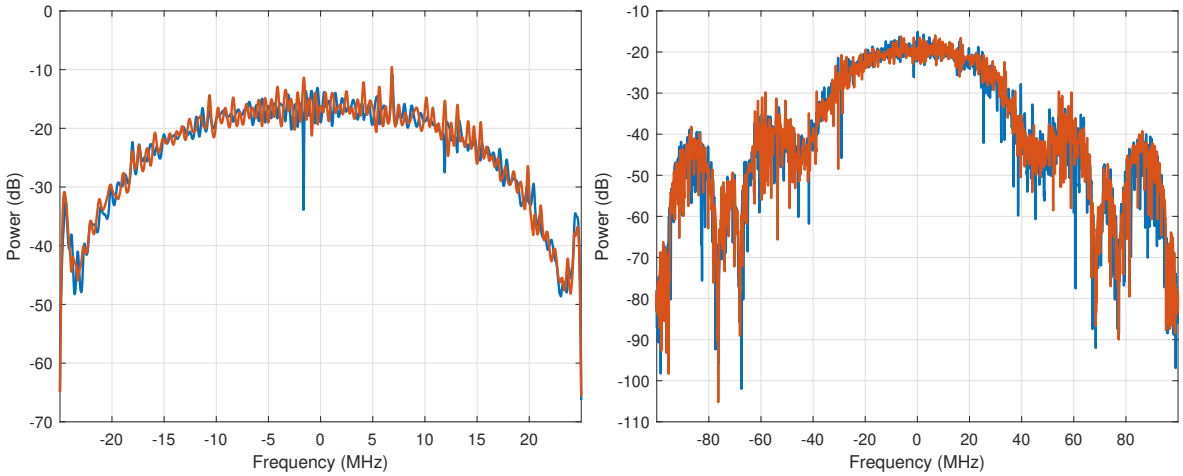

Figure 4.6: Simulated Loopback ISL MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

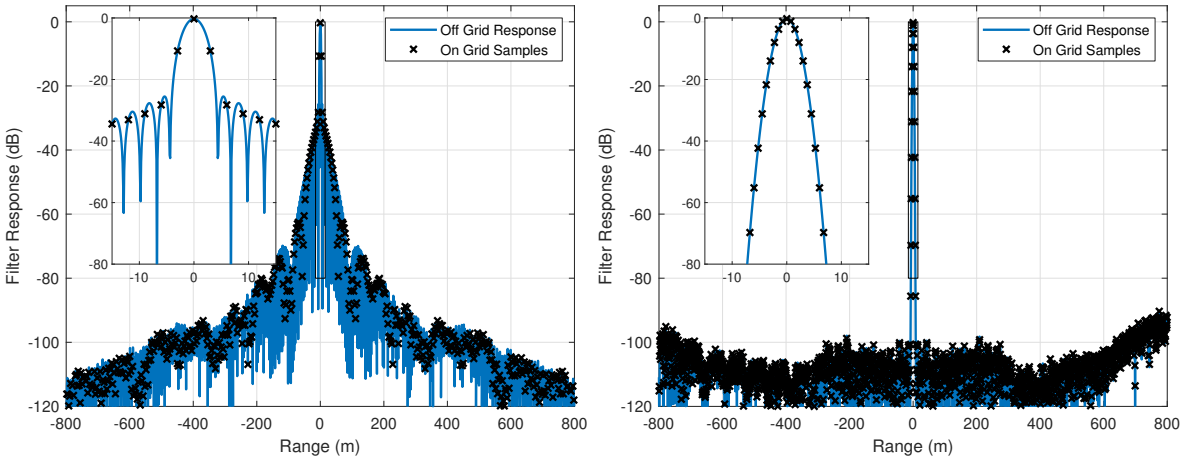

Figure 4.7: Simulated Loopback Template-base MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

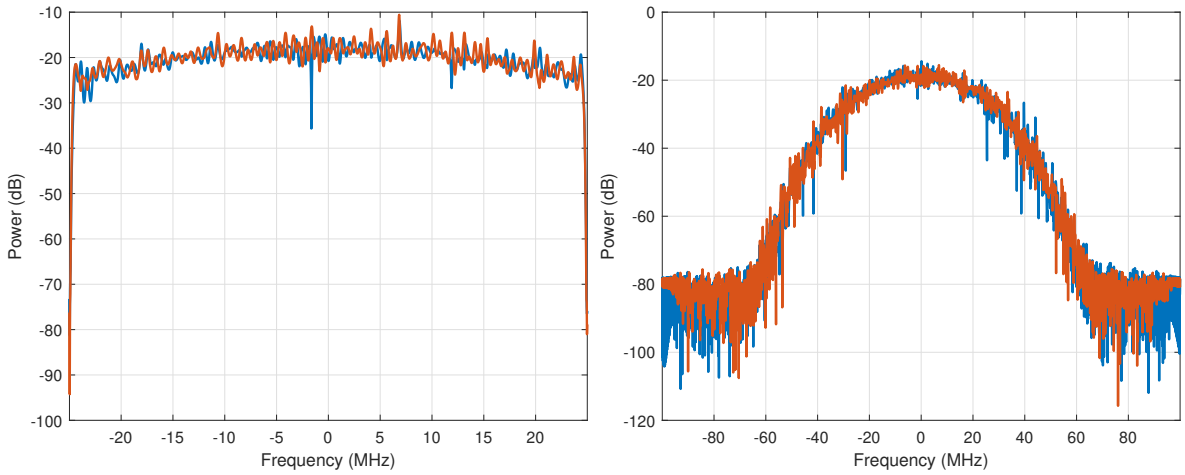

Figure 4.8: Simulated Loopback Template-base MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

matrix via

$$\begin{aligned}\mathbf{T}_{a,b} &= \mathscr{E}\{\mathbf{S}_a(\gamma)\,\mathbf{S}_b^H(\gamma)\} \\ &= \mathscr{E}\left\{\sum_{i=0}^{D-1}\mathbf{J}_i\mathbf{s}_a(\gamma)\,\mathbf{s}_b^H(\gamma)\,\mathbf{J}_i^H\right\} \\ &= \sum_{i=-N+1}^{D}\mathbf{J}_i\mathscr{E}\{\mathbf{s}_a(\gamma)\,\mathbf{s}_b^H(\gamma)\}\mathbf{J}_i^T \\ &= \sum_{i=-N+1}^{D}\mathbf{J}_i\left(\int_{-0.5}^{0.5}\mathbf{s}_a(\gamma)\,\mathbf{s}_b^H(\gamma)\,d\gamma\right)\mathbf{J}_i^T\end{aligned} \tag{4.10}$$

can again be computed via a structured-straddling models of (3.11) as this approach independent of the respective waveforms (e.g. γ is the only random variable) and the

$$\mathbf{R}_{\gamma,a,b} = \frac{1}{Z^2}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\mathbf{T}_\gamma \odot \mathbf{s}_{f,a}\mathbf{s}_{f,b}^H\right)\mathbf{F}\mathbf{I}_{N\times Z}^T \tag{4.11}$$

where $\mathbf{T}_\gamma$ is the same symmetric sinc-function covariance matrix taper (CMT) from (3.12). Additionally, all $\mathbf{T}_{a,b}$ sill correspond to $D\times D$ banded Toeplitz matrices with a matrix bandwidth of N

that have the form

$$\mathbf{T}_{a,b} = \begin{bmatrix} t_{0,a,b} & \cdots & t_{-N+1,a,b} & 0 & \cdots & 0 \\ \vdots & \ddots & & \ddots & \ddots & \vdots \\ t_{N-1,a,b} & & \ddots & & \ddots & 0 \\ 0 & \ddots & & \ddots & & t_{-N+1,a,b} \\ \vdots & \ddots & \ddots & & \ddots & \vdots \\ 0 & \cdots & 0 & t_{N-1,a,b} & \cdots & t_{0,a,b} \end{bmatrix} \tag{4.12}$$

that can be constructed efficiently via a generalization of (3.6) as

$$t_{k,a,b} = \begin{cases} \sum_{i=1}^{N-k} \left[\mathbf{R}_{\gamma,a,b}\right]_{i,i+k} & 0 \leq k \leq N-1 \\ \sum_{i=1}^{N+k} \left[\mathbf{R}_{\gamma,a,b}\right]_{i-k,i} & -N+1 \leq k \leq -1 \cdot \\ 0 & else \end{cases} \tag{4.13}$$

This straddle compensated representation can likewise be formed via truncated low-rank approximations of $\mathbf{R}_{\gamma,\mathbf{a},\mathbf{b}}$ in the similar manner as Section 3.1.1, though here the singular value decomposition is necessary as opposed to eigen decomposition. This low rank representation can then be leveraged to enable efficient beamspoiling via low-rank Woodbury updates.

Straddle compensated ISL MiCRFt filters were then generated using the same parameters as previous and used to generate the aggregate filter response and PSs in Figures 4.9 and 4.10 respectively. Here we see that straddle compensated MiCRFt is not able to suppress the on-grid sidelobes as far as, but the the on-grid samples have are now more commensurate with the continuum of straddles. Interestingly, the PSs do not yield the same degree of high frequency nulling behavior present with the individually designed filters on a per filter basis.

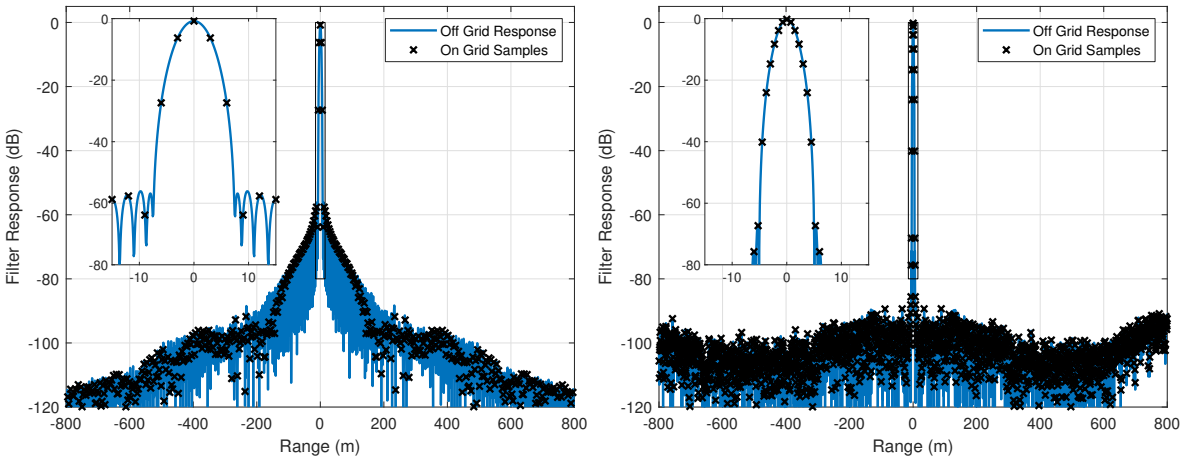

Figure 4.9: Simulated Loopback Straddle Compensated ISL MiCRFt Combined Response for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

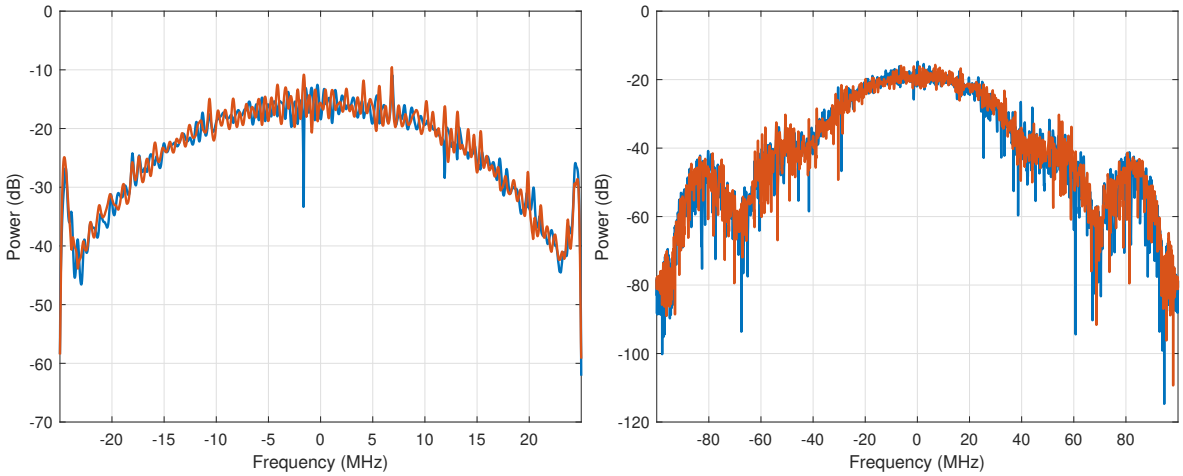

Figure 4.10: Simulated Loopback Straddle Compensated ISL MiCRFt PSs for Q=2 PROFM waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

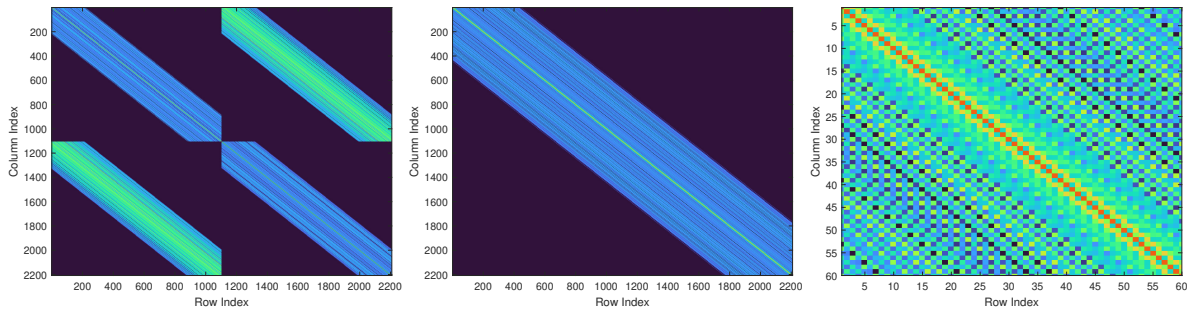

Figure 4.11: MiCRFt Hessian before permutation (left), after permutation (middle), and zoomed in to the permuted matrix (right)

4.3 Efficient Direct Solvers for MiCRFt

For the MiCRFt QP models, the hessian corresponds to a Toeplitz-block matrix. As outlined in Section 2.7.5, a permutation matrix can be applied to yield a hermitian block-Toeplitz structure where the resulting matrix is likewise banded with a matrix bandwidth of QN as illustrated in Figure 4.11. Similar to the direct LS MMF solvers in Section 3.2, the MiCRFt filters can be solved efficiently through a banded Cholesky solver with $\mathcal{O}(Q^2N^2D)$ for the decomposition, $\mathcal{O}(kQND)$ for the beamspoiling updates and $\mathcal{O}(QND)$ for the forward/backward substitution . Alternatively, multi-channel levinson recursion can be used to solve for the structured inverse matrix with $\mathcal{O}(Q^3D^2)$ complexity, and matrix-vector multiplication can be applied efficiently in the frequency domain for a cost of $\mathcal{O}(Q^2D\log(2D))$.

Unfortunately, MiCRFt does not yield an immediate structure that is as amenable to an iterative solvers like PCG. For certain Toeplitz-block matrices, a Circulant-block approximation to $\bar{\mathbf{T}}$ can be constructed by forming Circulant approximations for each block

$$\mathbf{C}_{i,j} \approx \mathbf{S}_i\mathbf{S}_j^H \tag{4.14}$$

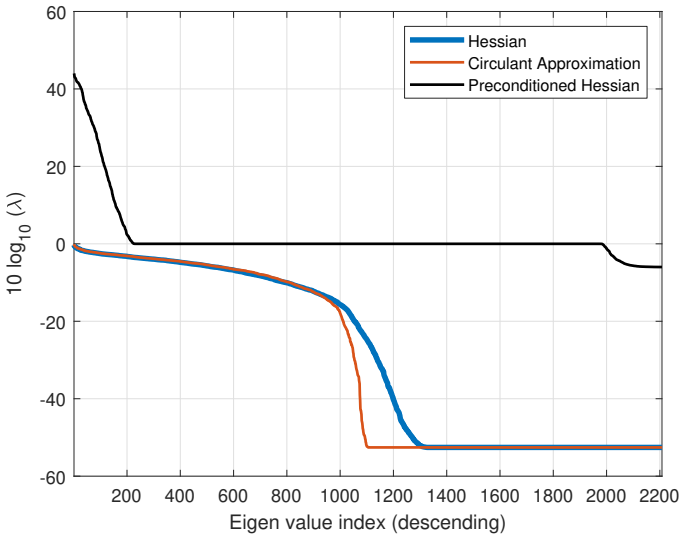

Figure 4.12: Preconditioning of MiCRFt QP

followed by concatenation into the block matrix

$$\bar{\mathbf{T}} \approx \bar{\mathbf{C}} = \begin{bmatrix} \mathbf{C}_{1,1} & \mathbf{C}_{1,2} & \cdots & \mathbf{C}_{1,Q} \\ \mathbf{C}_{2,1} & \mathbf{C}_{2,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{C}_{Q-1,Q} \\ \mathbf{C}_{Q,1} & \cdots & \mathbf{C}_{Q,Q-1} & \mathbf{C}_{Q,Q} \end{bmatrix} + \sigma^2 \mathbf{I}. \tag{4.15}$$

which can be implemented efficiently using block-Fourier transforms and sparse Cholesky decomposition [132]. However, the MiCRFt off-diagonal blocks (i.e., $\mathbf{S}_i \mathbf{S}_j^H$ for $i \neq j$) induce a transition between the regularization and signal eigen subspace that is poorly approximated by Circulant matrices. Consider the $Q = 2$ case from Figure 4.11 where a complementary subset was defined for PROFM waveforms with $\kappa = 1.5$, $D = 4N$, and $\sigma^2 = \frac{10^{-3}}{D}\text{Tr}\{\bar{\bar{\mathbf{S}}}\bar{\bar{\mathbf{S}}}^H\}$. Figure 4.12 illustrates the eigen value distribution of QP Hessian, it's Circulant-block approximation, and the resulting Hessian after preconditioning. While the Circulant-block system matches the dominant eigenvalues of the Hessian well, the transition to the regularization floor is poorly fit. Preconditioning approximately yields a division between these distributions, resulting in a large unclustered group and a poor condition number.

However, an alternative representation of the MiCRFt solution via Woodbury's identity can

yield a reduced complexity implementation. However, doing so is restricted to the template-based MiCRFt objective for $\kappa \approx 1$ as incorporation of straddling uncertainty would increase the computational complexity (low-rank approximate or otherwise). This implementation applies an identity to (4.5) as

$$\begin{aligned}\bar{\bar{\mathbf{w}}} &= \left(\bar{\bar{\mathbf{S}}}\bar{\bar{\mathbf{S}}}^H + \sigma^2\mathbf{I}\right)^{-1}\bar{\bar{\mathbf{S}}}\mathbf{p} \\ &= \frac{1}{\sigma^2}\bar{\bar{\mathbf{S}}}\left(\frac{1}{\sigma^2}\bar{\bar{\mathbf{S}}}^H\bar{\bar{\mathbf{S}}} + \mathbf{I}\right)^{-1}\mathbf{p} \\ &= \frac{1}{\sigma^2}\bar{\bar{\mathbf{S}}}\left(\frac{1}{\sigma^2}\sum_{q=1}^{Q}\left(\mathbf{S}_q^H\mathbf{S}_q\right) + \mathbf{I}\right)^{-1}\mathbf{p}\end{aligned} \tag{4.16}$$

which alters the linear system to the $(D+N-1)\times(D+N-1)$ matrix

$$\frac{1}{\sigma^2}\sum_{q=1}^{Q}\left(\mathbf{S}_q^H\mathbf{S}_q\right) + \mathbf{I} \tag{4.17}$$

which is no longer a Toeplitz matrix, but the Toeplitz and banded properties of $\mathbf{S}_q^H$ and $\mathbf{S}_q$ can still be leveraged for efficient construction. (4.17) retains a matrix-bandwidth of N, thus (4.16) can be solved via banded-Cholesky decomposition and efficient Toeplitz matrix-vector multiplication for an aggregate complexity of $\mathscr{O}(Q(D+N-1)\log(D+N-1)+N^2(D+N-1)+N(D+N-1))$ which will generally be faster than multi-channel Levinson. However, as (4.16) leverages Woodbury's identity, numerical precision errors may arise when using single precision and small values of σ^2.

The eigenvalue distribution of (4.17) does not have the same sharp transition and is more amenable to PCG approaches. As the linear system is no longer Toeplitz but rather Toeplitz-like, the Circulant preconditioners from Section 3.2.1 still improve convergence, but not to the same degree as the LS MMFs as shown. Alternatively, a Truncated Circulant plus Block Preconditioner can be formed via a Circulant approximation for the center-most $D\times D$ elements of (4.17) while the first and last $\frac{N}{2}-1\times\frac{N}{2}-1$ are left unchanged as illustrated in Figures 4.13. This preconditioner can be generated with a $\mathscr{O}(\left(\frac{N}{2}-1\right)^3)$ complexity, applied with $\mathscr{O}(\left(\frac{N}{2}-1\right)^2+D\log(D))$, and clusters the Hessian eigenvalues well, outside of a small number of low rank components as

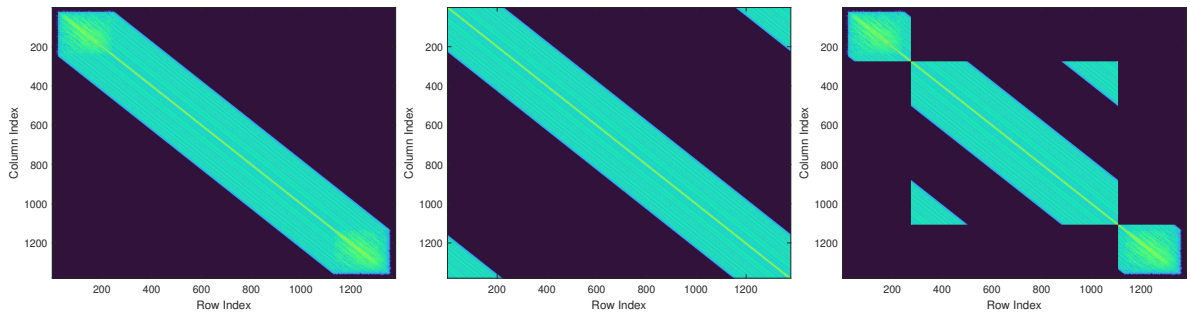

Figure 4.13: Example Toeplitz-like Matrix from (4.17) (left), it's Circulant-approximation preconditioner (middle), and a Truncated Circulant plus Block Preconditioner (right)

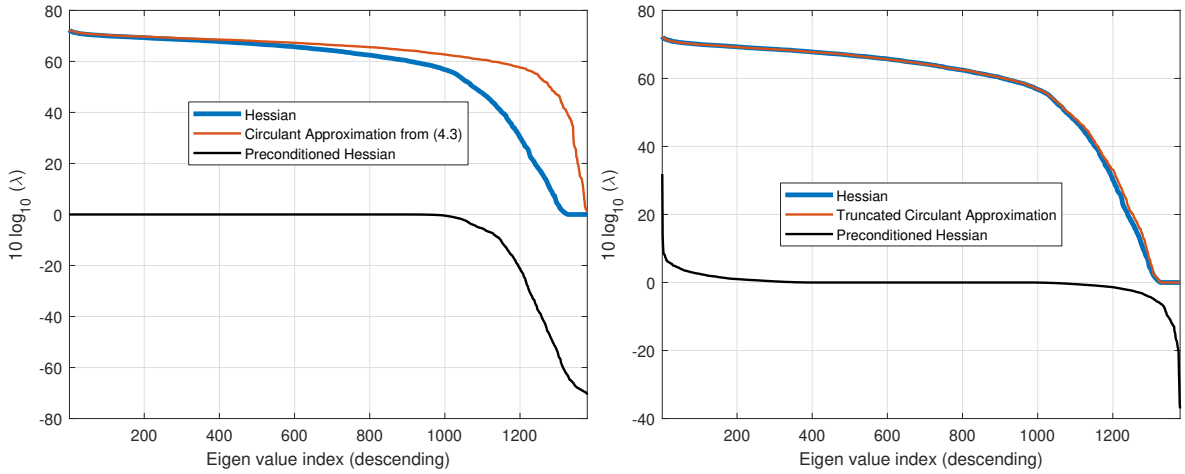

Figure 4.14: Eigenvalue Distribution for Circulant Preconditioner (left) and Truncated Circulant plus Block Preconditioner (right)

illustrated in Figure 4.14. PCG with this approach (denoted Precond 2) converges slower at first relative to the Circulant preconditioner (denoted Precond 1), but after a certain point accelerates and converges rapidly to the optimal solution as illustrated in Figures 4.15-4.16. That being said, alternative preconditioners likely have the potential to further improve convergence with less computational overhead.

4.4 MiCRFt Doppler Sensitivity and Aliasing

Complementary design across a subset effectively ignores the per-pulse sidelobe levels and instead minimizes the aggregate response by inducing a coherent cancellation effect when summing across the set. Thus, to achieve cancellation the slow-time phase progression must match the expected

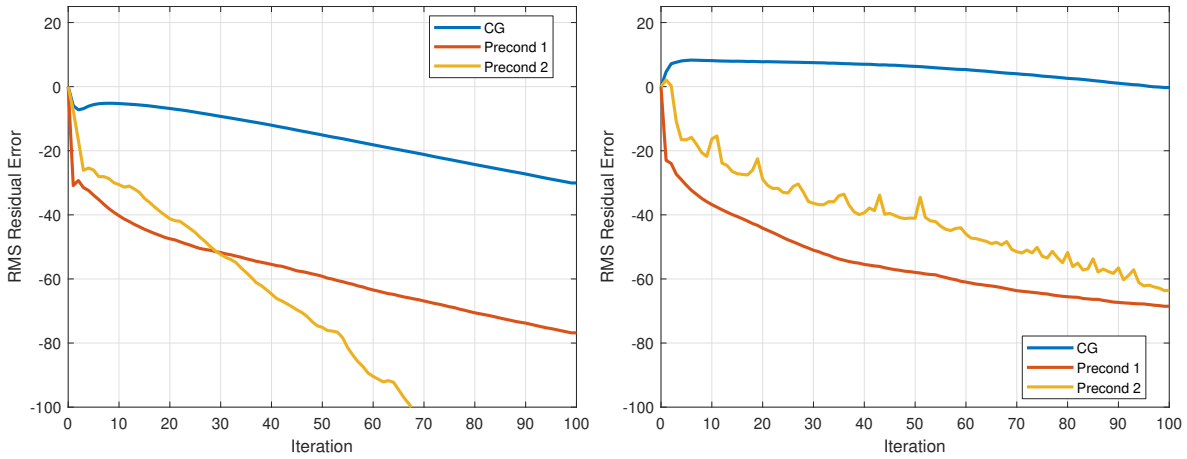

Figure 4.15: MiCRFt Residual Error vs Iteration of PCG using Circulant Preconditioners for subsets of $Q = 2$ PROFM Waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

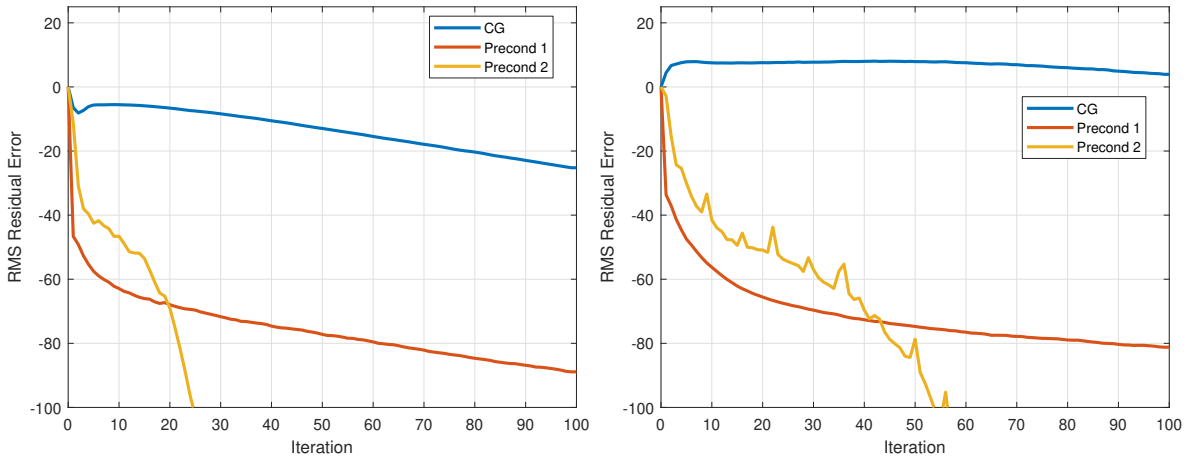

Figure 4.16: MiCRFt Residual Error vs Iteration of PCG using Circulant Preconditioners for subsets of $Q = 4$ PROFM Waveforms with $\kappa = 1.5$ (left) and $\kappa = 6$ (right)

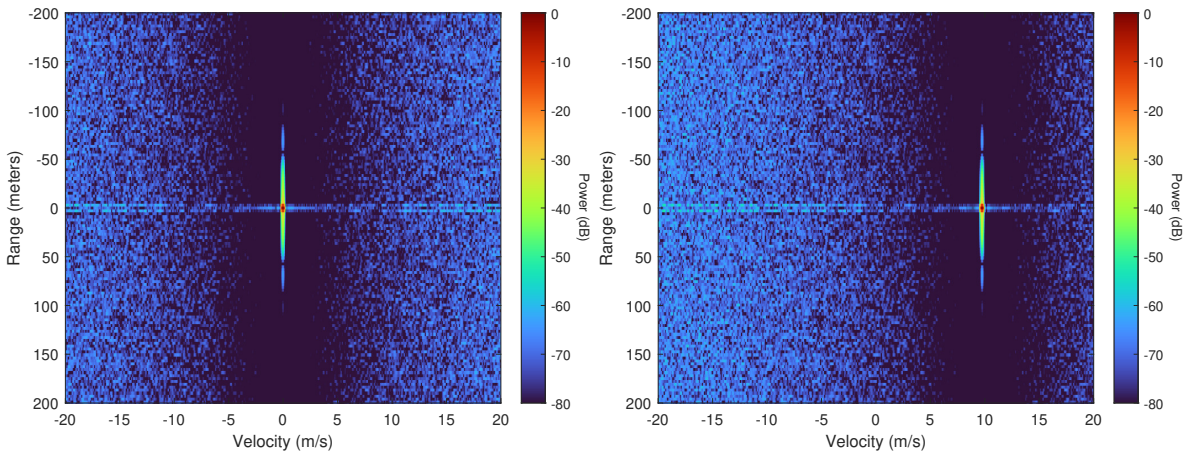

Figure 4.17: Illustrative $Q = 2$ MiCRFt RSM for 1000 PROFM waveforms for stationary (left) and moving (right) point scatterer

complementary design. However, when Doppler processing is performed (e.g, (2.90)), pulses within a set will be phase shifted by the corresponding Doppler steering vector, resulting in a deviation from the model in (4.1). Figure 4.17 illustrates this effect where complementary subsets of $Q = 2$ MiCRFt filters are generated for 1000 PROFM waveforms and applied to a simulated point scatter that is both stationary (zero-Doppler) on the left subfigure and moving at 10m/s on the right. Here the RSM near zero Doppler has been suppressed as the slow-time phase progression closely matches the complementary design, larger Doppler shifts correspond to RSM levels closer to that of the MF.

This effect is not exclusive to MiCRFt and is a general artifact of complementary design. Traditional complementary processing contends with this effect by approximating the expected Doppler shift within a subset as small (i.e. the unambiguous velocity is much larger than the fastest moving scatterer). When this is the case, presumming can be applied across the subsets (in slow-time) prior to Doppler processing. Note, this operation does not mitigate the complementary "decoherence" and sidelobe levels will increase when a nonzero Doppler shift is present within a subset as shown in Figure 4.17. Rather, presumming effectively induces a fixed range-sidelobe level unique to each respective Doppler shift. As such, subsequent Doppler processing produces a unique PSF for every respective Doppler slice, where the RSM levels are proportional to the range-sidelobes of each Doppler shifted complementary subset. Since clutter typically contains the dominant scatterers,

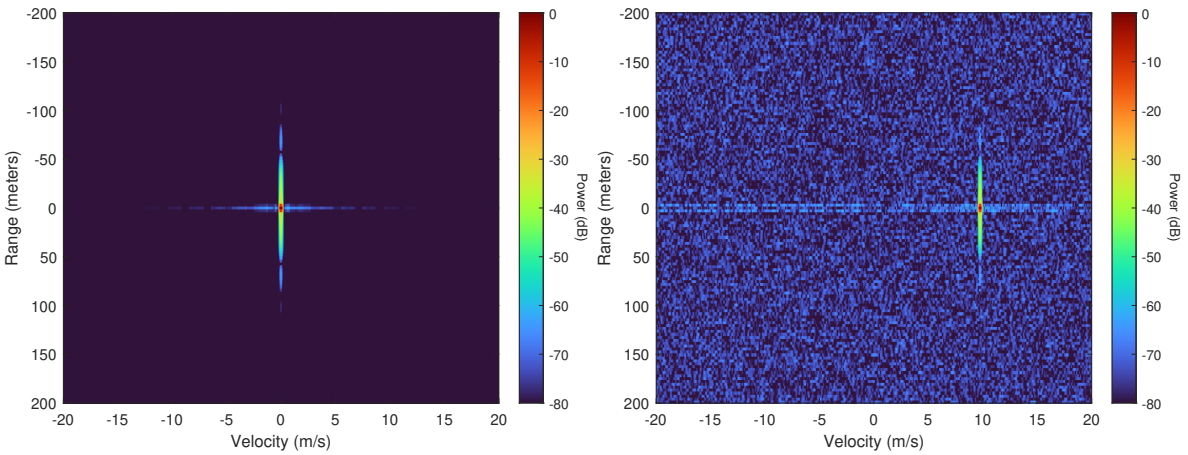

Figure 4.18: Illustrative $Q = 2$ MiCRFt RSM after presumming for 1000 PROFM waveforms for stationary (left) and moving (right) point scatterer

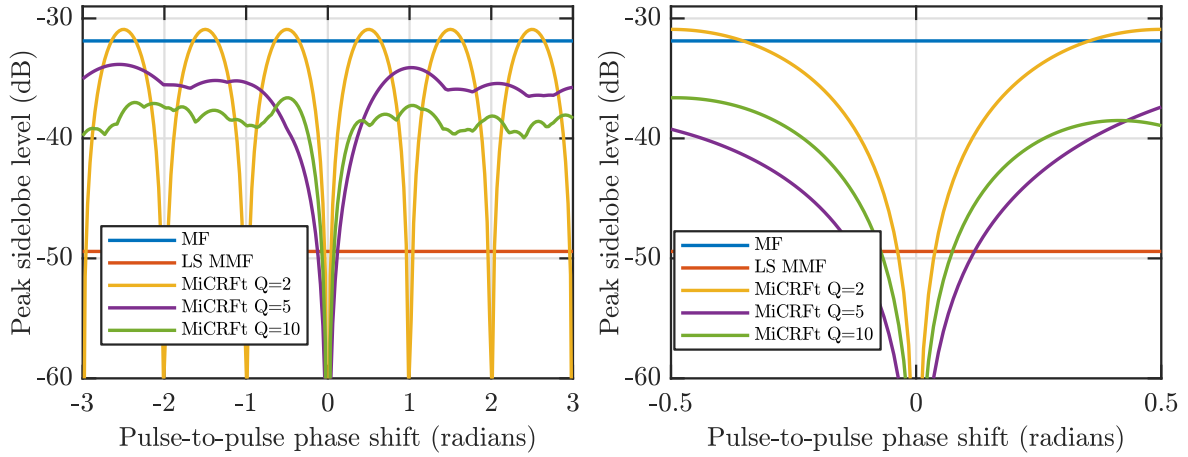

Figure 4.19: Peak sidelobe level for MiCRFt subsets as a function of slow-time phase shift

the presumming operation can be tuned such that the complementary effect is maximized at the expected clutter Doppler, resulting in suppression of RSM from clutter but a potential to increase RSM from other moving scatterers (relative to the LS MMF). Additionally, presumming reduces the effective PRF by a factor of Q, while likewise proportionally reduces the maximum unambiguous velocity. See Figure 4.18 for the same illustrative example, but now with presumming by a factor of Q applied as well as Figure 4.19 for an example of MiCRFt peak-sidelobe-level as a function of the-pulse-to-pulse phase change within a subset.

4.5 MiCRFt Closed-Loop Experimental Assessment

To assess the efficacy of MiCRFt, the closed-loop experimental measurements from Section 3.3 were used to generate 500 $Q = 2$ complementary subsets of MiCRFt filters for the PROFM CPI of 1000 waveforms with the same filter length, regularization, and beamspoiling/template as before. MiCRFt filters were generated using the template-based, the uncompensated ISL, and the straddle compensated ISL objectives and applied to the loopback measurements to yield the aggregate filter response and frequency spectrum illustrated in the left and right sides respectively of Figure (4.20). The template-based MiCRFt yields a high fidelity spectral match with the desired response, and in range yields a near perfect match down to -110dB. However, as this desired response accounts for the anti-aliasing filter it has non-zero sidelobes that are higher than the other approaches. For the ISL case without straddling compensation the range response has no sidelobes down to -100dB and the aggregate spectrum has a smooth roll-off. However, as with the parallel LS MMF formulation, the lack of straddling compensation may yield a broadening of this response and it will be susceptible to mainlobe-modulation from variations in spectral shape. Lastly, the straddle compensated ISL MiCRFt formulation yields a a broadened mainlobe relative the other approaches as well as 0.5dB more MML. However, it has lower near in sidelobes than the Template-based formulation should be more robust than the uncompensated formulation.

Table 4.1 provides a list of the parameters using in this closed-loop assessment while Table 4.2 provides the corresponding order of complete for solving these filters. Note, for the $Q = 2$ case multichannel levinson offers a rather efficient solution, though again levinson based approaches are not as amenable to parallelization or hardware acceleration like PCG. Additionally, if the complementary subset was increased to $Q = 5$ (shown in Table 4.3), levinson sees a large increase in cost while the reformulated banded approach and PCG largely stay the same.

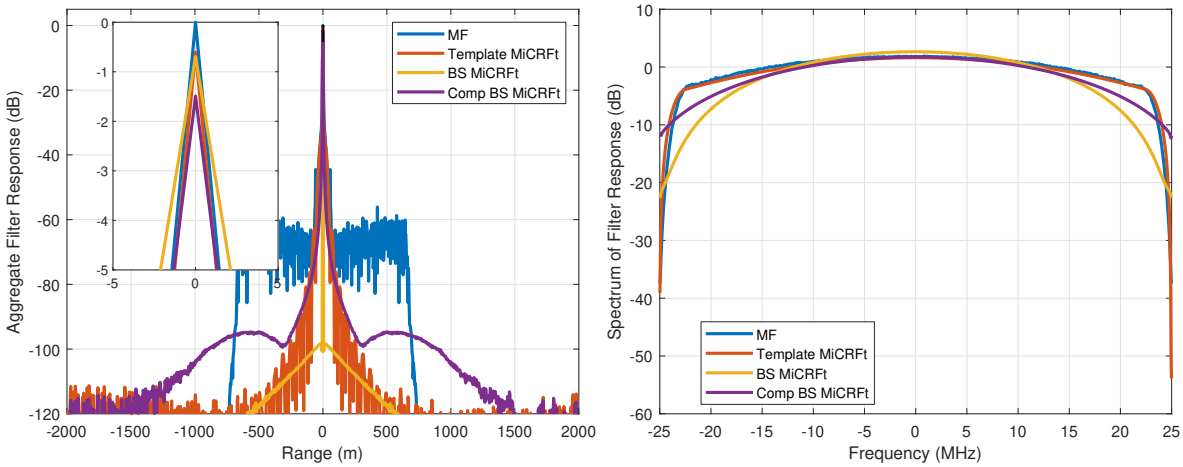

Figure 4.20: Loopback MiCRFt Response for PROFM CPI and $\kappa = 1.5$

Measurement Parameters		
Parameter	Variable	Value
Sample Rate	f_s	50MHz
3dB Bandwidth	B	33.3MHz
Pulse Duration	T	4.5μseconds
Samples in Waveform	N	276
Fast-Time Filter Samples	D	$4N$
Size of Complementary Subset	Q	2

Table 4.1: Table of Experimental Parameters for PROFM CPI of 1000 waveforms

MiCRFt Order of Complexity	
Direct	1.08E+10
Multi-channel Levinson	9.75E+06
Reformulated Banded	1.05E+08
Reformulated CG (10 iter)	1.73E+05
Reformulated CG (20 iter)	3.16E+05
Reformulated CG (100 iter)	1.47E+06

Table 4.2: Complexity for constructing MiCRFt Filters $Q = 2$

MiCRFt Order of Complexity	
Direct	1.68E+11
Multi-channel Levinson	1.52E+08
Reformulated Banded	1.05E+08
Reformulated CG (10 iter)	2.16E+05
Reformulated CG (20 iter)	3.60E+05
Reformulated CG (100 iter)	1.51E+06

Table 4.3: Complexity for constructing MiCRFt Filters $Q = 5$

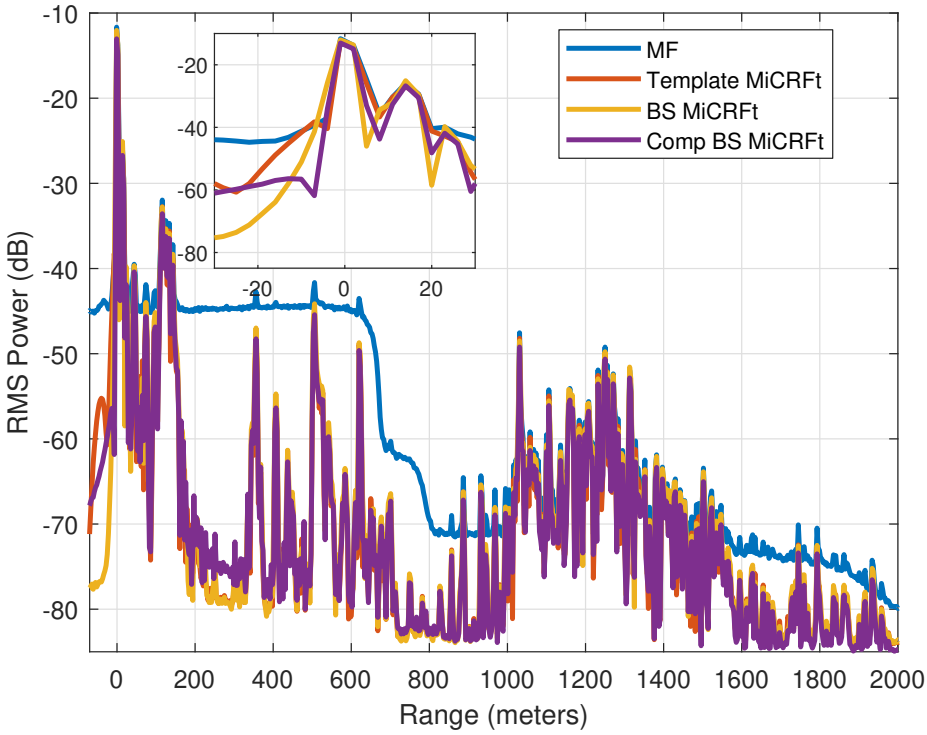

Figure 4.21: MiCRFt RMS Open-Air Range Profile Estimates $\kappa = 1.5$

4.6 MiCRFt Open-Air Experimental Assessment

These respective MiCRFt approaches were then applied to the open-air experimental measurements from (3.4) and presumming was performed to achieve the complementary effect, which likewise halved maximum unambiguous velocity $Q = 2$. The RMS range response for respective approach illustrated in Figure 4.21 where each of the MiCRFt approaches yields sidelobes 15-20dB lower than their respective LS MMF counterpart. Additionally, even with the extended filter length ($D = 4N$) the direct path sidelobes are not visible in the intersection the sidelobe floor is now than the PROFM MF and has a level commensurate with the LS MMFs applied to the LFM CPI. In the subfigure, we see that the template-based MiCRFt has the same high near-in sidelobes as the LS MMF case, ISL MiCRFt has a broadened response relative to the expected loop-back filter response due to straddling, while the straddle compensated ISL MiCRFt does not suffer from this broadening but does not roll-off as fast as uncompensated case.

Next, Doppler processing was performed on the respective range profiles were again a Taylor

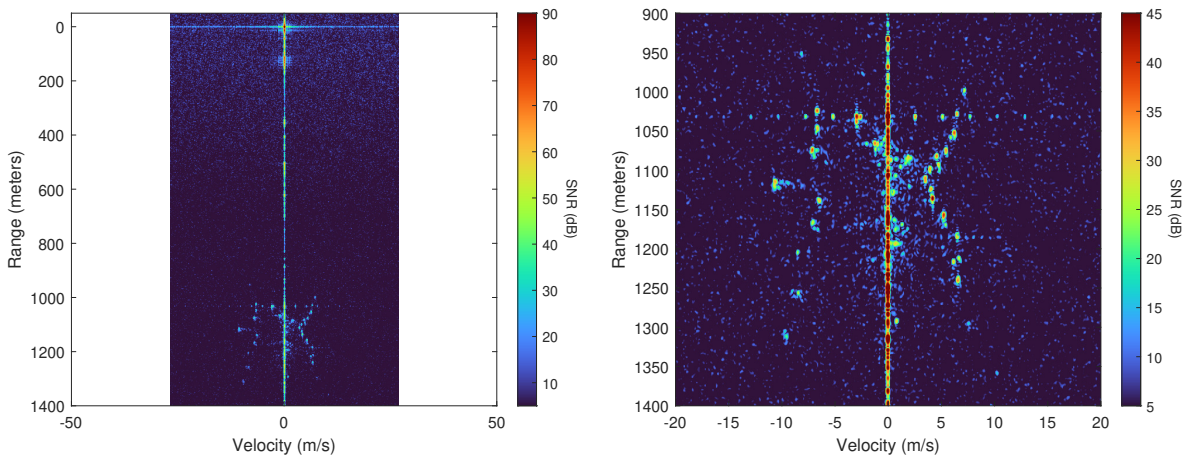

Figure 4.22: Range-Doppler Estimate for PROFM CPI using the Template-based MiCRFt with $Q = 2$ and $\kappa = 1.5$

window was applied to suppress Doppler sidelobes. Figure (4.22) first illustrates the template-based MiCRFt approach where the reduction in unambiguous velocity is clear in the zoomed out look on the left. Here the direct path RSM has been significantly reduced ($\sim$40dB lower than the MF) and the residual speckle corresponds to the model uncertainty floor as the expected sidelobe level (based on the loopback measurements in Figure 4.20) would yield RSM 20dB below the noise floor. The zoomed in look of the intersection reveals that MiCRFt provides an estimate of the cars and their respective Doppler signatures as good as the non-agile LFM LS MMF case, though there is a slight increase (2dB) in background speckle corresponding to either model errors or raise sidelobes from MiCRFt's Doppler sensitivity.

Next the ISL MiCRFt Range-Doppler estimate is illustrated in Figure 4.23. Here the direct path RSM levels are identical to the template based approach since both cases are effectively model-error limited. However, there is enough subset-to-subset spectral variation that the ISL case induces notable mainlobe modulation. Relative to the ISL LS MMF, this Doppler spreading is $\sim$6dB lower and is roughly commensurate with the straddle compensate ISL LS MMF. This drawback makes use of ISL MiCRFt nonviable for fully agile radar in practice, even if it can provide theoretically lower pulse-to-pulse sidelobes than the template-based approach.

Finally, the straddle compensated ISL MiCRFt range-profiles were Doppler processed and the corresponding figures are shown in Figure 4.24. The direct path RSM suppression is on par with

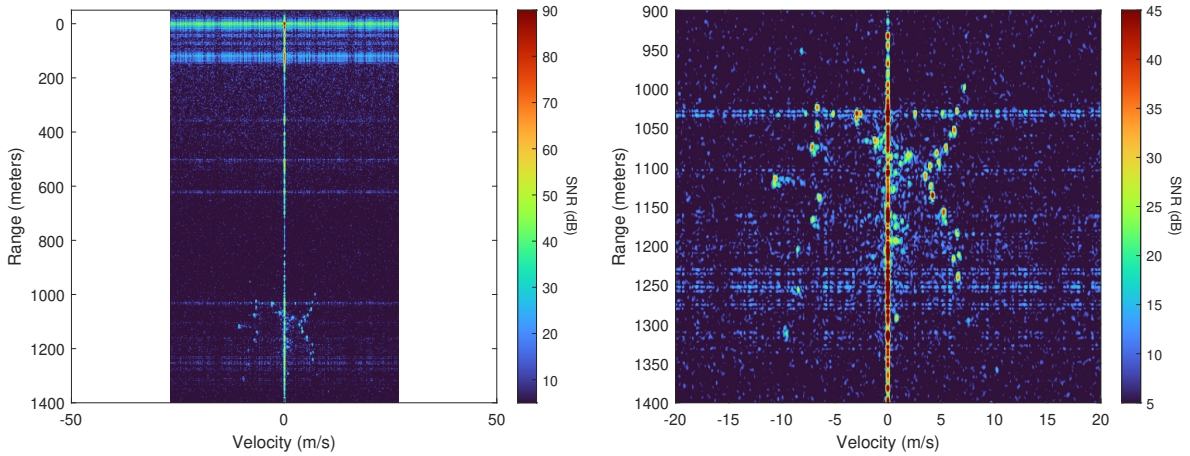

Figure 4.23: Range-Doppler Estimate for PROFM CPI using the ISL MiCRFt with $Q = 2$ and $\kappa = 1.5$

the other two approaches and the background speckle corresponds to the model fidelity limitation. As with the LS MMFs, incorporation of straddle uncertainty effectively restricts the spectral solution space, and the resulting range-Doppler map also does not see any mainlobe modulation, though it incurs and additional 0.5dB of MML. Additionally, this MiCRFt does not see high-power near in sidelobes like the template case, though the mainlobe roll-off has slightly broadened. Unfortunately, this dataset being model fidelity limited does not allow for full characterization of the Doppler sensitivity of these respective approaches as RSM from the vehicles are likely below the noise floor. As future work, it would be interesting to explore MiCRFt analysis on higher fidelity measurements with a larger dynamic range.

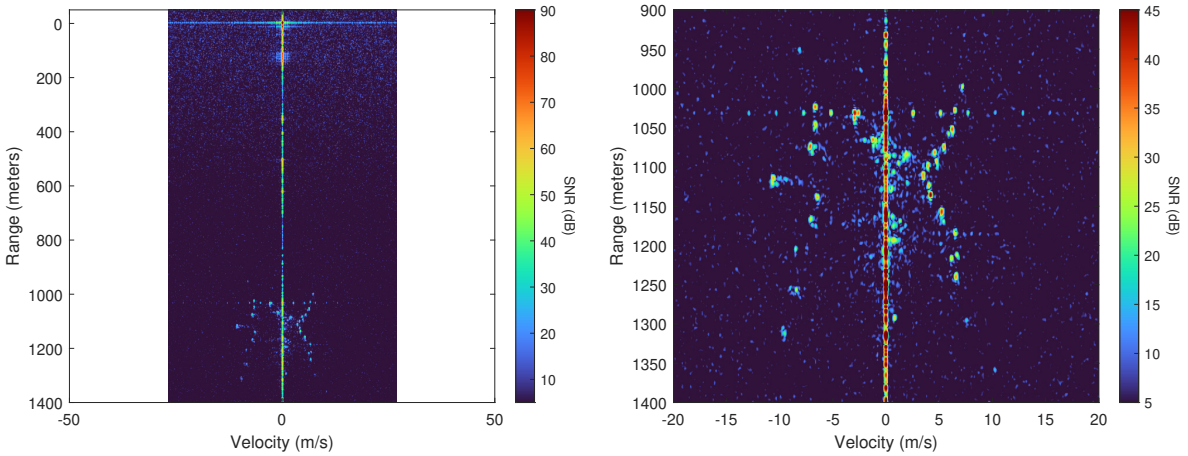

Figure 4.24: Range-Doppler Estimate for PROFM CPI using the Straddle Compensated ISL Mi-CRFt with $Q = 2$ and $\kappa = 1.5$

Chapter 5

Alternate Joint Range-Doppler Clutter Cancellation Representations and Efficient Implementations

The joint representation of NIMPC can be extremely effective at suppressing clutter RSM, however, as the filters requires storage and inversion of a $DM \times DM$ matrix, computational complexity acts as a bottleneck in practice. To enable solutions with practical merit, here we examine computationally efficient implementations of NIMPC as well as alternative join-Dimension clutter cancellation frameworks that offer significant complexity reductions.

5.1 Efficient Direct Solvers for NIMPC

The NIMPC filter model likewise follows a Toeplitz-Block form which allows for incorporation of the multi-channel Levinson approach or a banded Cholesky approach when $D > N$. Additionally, the Fourier relationship with respect to Doppler allows for some additional computational benefits that enable efficient calculation of

$$\begin{aligned} \hat{x}_{cc}(\ell, \omega) &= \bar{\mathbf{w}}_{cc}^{H}(\omega)\, \bar{\mathbf{y}}(\ell) \\ &= \frac{1}{\bar{\mathbf{s}}^{H}(\omega) \left(\bar{\mathbf{R}}_{cc}\right)^{-1} \bar{\mathbf{s}}(\omega)} \bar{\mathbf{s}}^{H}(\omega) \left(\mathbf{R}_{cc}\right)^{-1} \bar{\mathbf{y}}(\ell) \end{aligned} \tag{5.1}$$

without explicitly performing

$$\left(\bar{\mathbf{R}}_{cc}\right)^{-1} \bar{\mathbf{s}}(\omega)$$

for every Doppler bin.

First, by permuting to a Block-Toeplitz form and solving for inverse form via multi-channel

levinson we have

$$\begin{aligned}\bar{\mathbf{R}}_{cc}^{-1} &= \left(\bar{\mathbf{J}}^T\bar{\mathbf{J}}\bar{\mathbf{R}}_{cc}\bar{\mathbf{J}}^T\bar{\mathbf{J}}\right)^{-1} \\ &= \bar{\mathbf{J}}^T\left(\bar{\mathbf{J}}\bar{\mathbf{R}}_{cc}\bar{\mathbf{J}}^T\right)^{-1}\bar{\mathbf{J}}^T \\ &= \bar{\mathbf{J}}\left(\bar{\mathbf{T}}_f\bar{\mathbf{E}}_f^{-1}\bar{\mathbf{T}}_f^H - \bar{\mathbf{T}}_b\bar{\mathbf{E}}_b^{-1}\bar{\mathbf{T}}_b^H\right)\bar{\mathbf{J}}^T\end{aligned} \tag{5.2}$$

which will have $\mathscr{O}(M^3D^2)$ complexity. Next [107], noted that $\bar{\mathbf{s}}(\omega)$ could be decomposed as

$$\bar{\mathbf{s}}(\omega) = \mathbf{D}\mathbf{a}(\omega) \tag{5.3}$$

where $\mathbf{a}(\omega)$ is the $M \times 1$ Doppler steering vector from (2.87) and

$$\mathbf{D} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{v}_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{v}_M \end{bmatrix} \tag{5.4}$$

is a $DM \times M$ matrix composed of the zero-padded waveform for every direction. Thus, only M systems of equations need to be solved as

$$\tilde{\mathbf{D}} = \left(\bar{\mathbf{R}}_{cc}\right)^{-1}\mathbf{D} \tag{5.5}$$

to construct the NIMPC filter for every Doppler. Using the block-Toeplitz inverse form, $\tilde{\mathbf{D}}$ can be solved with $\mathscr{O}(M^3(D+N-1)\log(D+N-1))$ (though parallel processing for each of the M vectors in $\tilde{\mathbf{D}}$ can yield a $\mathscr{O}(M^2(D+N-1)\log(D+N-1))$ cost). In total, solving for the NIMPC filters via the multi-channel levinson approach will have a cost of

$$\mathscr{O}(M^3D^2 + M^3(D+N-1)\log(D+N-1))$$

where the dependency on M^3 can still be a practical hindrance.

As an alternative, when the number of "cancellation" directions $K_c \ll M$ it is beneficial to

repose (5.5) in a similar manner to (4.16) by expanding

$$\begin{aligned}(\mathbf{R}_{cc})^{-1}\mathbf{D} &= \left(\sigma_c^2\tilde{\tilde{\mathbf{S}}}\tilde{\tilde{\mathbf{S}}}^H+\sigma_n^2\mathbf{I}\right)^{-1}\mathbf{D}\\ &= \frac{1}{\sigma_n^2}\mathbf{D}-\frac{1}{\sigma_n^4}\tilde{\tilde{\mathbf{S}}}\left(\frac{1}{\sigma_n^2}\tilde{\tilde{\mathbf{S}}}^H\tilde{\tilde{\mathbf{S}}}+\frac{1}{\sigma_c^2}\mathbf{I}\right)^{-1}\tilde{\tilde{\mathbf{S}}}^H\mathbf{D}\end{aligned} \tag{5.6}$$

using woodbury's identity where the $K_c(D+N-1)\times K_c(D+N-1)$ matrix

$$\tilde{\tilde{\mathbf{S}}}^H\tilde{\tilde{\mathbf{S}}}=\sum_{m=1}^{M}\begin{bmatrix}\mathbf{S}_m^H\mathbf{S}_m & e^{j(\omega_2-\omega_1)\frac{m-1}{f_{prf}}}\mathbf{S}_m^H\mathbf{S}_m & \cdots & e^{j(\omega_{K_c}-\omega_1)\frac{m-1}{f_{prf}}}\mathbf{S}_m^H\mathbf{S}_m\\ e^{j(\omega_1-\omega_2)\frac{m-1}{f_{prf}}}\mathbf{S}_m^H\mathbf{S}_m & \mathbf{S}_m^H\mathbf{S}_m & & \vdots\\ \vdots & & \ddots & \vdots\\ e^{j(\omega_1-\omega_{K_c})\frac{m-1}{f_{prf}}}\mathbf{S}_m^H\mathbf{S}_m & \cdots & \cdots & \mathbf{S}_m^H\mathbf{S}_m\end{bmatrix} \tag{5.7}$$

is a Block-Toeplitz matrix with K_c^2 blocks of Toeplitz-like matrices

$$\sum_{m=1}^{M}e^{j(\omega_a-\omega_b)\frac{m-1}{f_{prf}}}\mathbf{S}_m^H\mathbf{S}_m$$

that have a matrix-bandwidth of N. (5.6) is then broken into four significant operations as

$$\begin{aligned}(\mathbf{R}_{cc})^{-1}\mathbf{D} &= \frac{1}{\sigma_n^2}\mathbf{D}-\frac{1}{\sigma_n^4}\tilde{\tilde{\mathbf{S}}}\left(\frac{1}{\sigma_n^2}\tilde{\tilde{\mathbf{S}}}^H\tilde{\tilde{\mathbf{S}}}+\frac{1}{\sigma_c^2}\mathbf{I}\right)^{-1}\tilde{\tilde{\mathbf{S}}}^H\mathbf{D}\\ &= \frac{1}{\sigma_n^2}\mathbf{D}-\frac{1}{\sigma_n^4}\tilde{\tilde{\mathbf{S}}}\left(\frac{1}{\sigma_n^2}\tilde{\tilde{\mathbf{S}}}^H\tilde{\tilde{\mathbf{S}}}+\frac{1}{\sigma_c^2}\mathbf{I}\right)^{-1}\mathbf{U}_1\\ &= \frac{1}{\sigma_n^2}\mathbf{D}-\frac{1}{\sigma_n^4}\tilde{\tilde{\mathbf{S}}}\mathbf{U}_2\\ &= \frac{1}{\sigma_n^2}\mathbf{D}-\mathbf{U}_3\\ &= \mathbf{U}_4\end{aligned} \tag{5.8}$$

where $\mathbf{U}_1$ and $\mathbf{U}_3$ are implemented via efficient Toeplitz-block matrix-vector multiplication for $\mathcal{O}(K_cM^2(D+N-1)\log(D+N-1))$ and $\mathcal{O}(K_c^2M(D+N-1)\log(D+N-1))$ respectively, while solving $\mathbf{U}_4$ requires only negligible elementwise operations. The matrix inversion can be solved

via a banded-block Cholesky decomposition or multichannel levinson recursion on

$$\frac{1}{\sigma_n^2}\tilde{\tilde{\mathbf{S}}}^H\tilde{\tilde{\mathbf{S}}}+\frac{1}{\sigma_c^2}\mathbf{I}$$

where Cholesky decomposition will cost $\mathscr{O}(Kc^2N^2(D+N-1))$ and forward/backward substitution for all M vectors of $\mathbf{U}_1=\tilde{\tilde{\mathbf{S}}}^H\mathbf{D}$ will have complexity of $\mathscr{O}(KcNM(D+N-1))$, while multichannel levinson will cost $\mathscr{O}(K_c^3(D+N-1)^2)$ to yield the inverse matrix and $\mathscr{O}(MK_c^2(D+N-1)\log(D+N-1))$ will be the cost the subsequent matrix-vector multiplication. Thus, the total complexity for this formulation with be on the order of

$$\begin{aligned}&\mathscr{O}(Kc^2N^2(D+N-1)+KcNM(D+N-1)\\&\quad+K_cM^2(D+N-1)\log(D+N-1)+K_c^2M(D+N-1)\log(D+N-1))\end{aligned}$$

for a Cholesky approach and

$$\mathscr{O}(K_c^3(D+N-1)^2+MK_c^2(D+N-1)\log(D+N-1)+K_cM^2(D+N-1)\log(D+N-1))$$

for the multichannel Levinson approach.

[107] also noted that $\tilde{\mathbf{D}}^H\bar{\mathbf{y}}(\ell)$ follows a convolution relationship with range, and thus

$$\mathbf{a}^H(\omega)\tilde{\mathbf{D}}^H\bar{\mathbf{y}}(\ell)$$

could be applied for all range and Doppler via frequency domain convolution and FFT-based Doppler processing. However, [107] did not examine speedups for computing the unity scaling

$$\frac{1}{\bar{\mathbf{s}}^H(\omega)\left(\bar{\mathbf{R}}_{cc}\right)^{-1}\bar{\mathbf{s}}(\omega)}$$

or the noise scaling

$$\frac{1}{\left\|\left(\bar{\mathbf{R}}_{cc}\right)^{-1}\bar{\mathbf{s}}(\omega)\right\|}=\frac{1}{\sqrt{nc(\omega)}}$$

which is typically necessary in practice to avoid significant distortion across Doppler. Without speedups, forming these denominators requires computing the filter for each of the K Doppler slices (where $K \geq M$) that are to be estimated, yielding a cost of $\mathscr{O}(KMD + MDK\log(K))$ using (5.3) and (5.5), which at best $K = M$ yielding a cost of $\mathscr{O}(M^2D + M^2D\log(M))$ but in practice $K > M$ to minimize the loss from straddling in Doppler.

To compute the scale factors efficiently, evaluating

$$\begin{aligned} gc(\omega) &= \bar{\mathbf{s}}^H(\omega)\left(\bar{\mathbf{R}}_{cc}\right)^{-1}\bar{\mathbf{s}}(\omega) \\ &= \mathbf{a}^H(\omega)\mathbf{D}^H\tilde{\mathbf{D}}\mathbf{a}(\omega) \\ &= \mathbf{a}^H(\omega)\mathbf{G}\mathbf{a}(\omega) \end{aligned} \tag{5.9}$$

yields a $M \times M$ hermitian matrix $\mathbf{G}$ that can be used to perform our efficient implementation. Note, to incorporating Doppler windowing one can simply weight the respective columns of $\mathbf{G}$ by the respective window at this stage. Next, as $\mathbf{a}(\omega)$ corresponds to complex exponential, (5.9) can be rearranged and expressed as the frequency transform [109, 133] as

$$gc(\omega) = \sum_{m=-M+1}^{M+1} e^{-j\omega\frac{m}{f_{prf}}} g_m \tag{5.10}$$

where

$$g_m = \begin{cases} \sum_{i=1}^{M-m}[\mathbf{G}]_{i,i+m} & 0 \leq k \leq M-1 \\ \sum_{i=1}^{M+m}[\mathbf{G}]_{i-m,i} & -M+1 \leq k \leq -1 \end{cases} \tag{5.11}$$

is the sum of mth the diagonal of $\mathbf{G}$. As (5.10) is simply the two-sided Fourier transform of the g_m coefficients, the scaling for each Doppler bin can be computed efficiently via FFTs with $\mathscr{O}(K\log(K))$ cost after computing $\mathbf{D}^H\tilde{\mathbf{D}}$ which, using the sparsity of $\mathbf{D}$, has a $\mathscr{O}(M^2D)$ complexity. Thus, the total complexity is $\mathscr{O}(M^2D + K\log(K))$ which laregly removes scaling with K as the $K\log(K)$ operation will typically be negligible.

The filter scaling of

$$\frac{1}{\left\| \left(\bar{\mathbf{R}}_{cc}\right)^{-1} \bar{\mathbf{s}}\left(\omega\right) \right\|} = \frac{1}{\sqrt{nc\left(\omega\right)}}$$

which retains a uniform noise floor is a more challenging case. It can likewise be expressed as

$$\begin{aligned} nc\left(\omega\right) &= \bar{\mathbf{s}}^{H}\left(\omega\right)\left(\bar{\mathbf{R}}_{cc}\right)^{-2}\bar{\mathbf{s}}\left(\omega\right) \\ &= \mathbf{a}^{H}\left(\omega\right)\tilde{\mathbf{D}}^{H}\tilde{\mathbf{D}}\mathbf{a}\left(\omega\right) \\ &= \mathbf{a}^{H}\left(\omega\right)\mathbf{G}\mathbf{a}\left(\omega\right) \end{aligned} \tag{5.12}$$

which requires evaluating the matrix-multiplication $\tilde{\mathbf{D}}^{H}\tilde{\mathbf{D}}$ that does not have a sparse structure, yielding a cost of $\mathscr{O}(M^3 D)$. The aggregate cost is then $\mathscr{O}(M^3 D + K\log(K))$ which has an increase in order by M, relative to the unity scaling, and may not always be faster than direct evaluation using (5.3) and (5.5) ($\mathscr{O}(KMD + MDK\log(K))$). Ignoring the standalone $K\log(K)$ term in the efficient implementation, removing the common scaling with D, and expressing $K = aM$ as a multiplicative size increase by the scalar $a \geq 1$, we can define the relationship

$$M^3 < aM^2 + aM^2\left(\log(M) + \log(a)\right)$$

to determine if this implementation will have be more computationally efficient for a particular value of a and M. This trend is illustrated in Figure 5.1 where the dashed lines correspond to the complexity scaling with the proposed approach (M^3) and full lines correspond to a direct approach ($aM^2 + aM^2\left(\log(M)\right)^a$). Here, for small M this alternative implementation but the scaling with M reveals that generally the other approach is more efficacious. However, as $\left(\bar{\mathbf{R}}_{cc}\right)^{-1}$ is hermitian and has already been decomposed to yield $\tilde{\mathbf{D}}$, a second application of

$$\tilde{\tilde{\mathbf{D}}} = \left(\bar{\mathbf{R}}_{cc}\right)^{-1}\tilde{\mathbf{D}} \tag{5.13}$$

can be applied without a second decomposition of $\bar{\mathbf{R}}_{cc}$, such that the sparsity of $\mathbf{D}$ can be leveraged

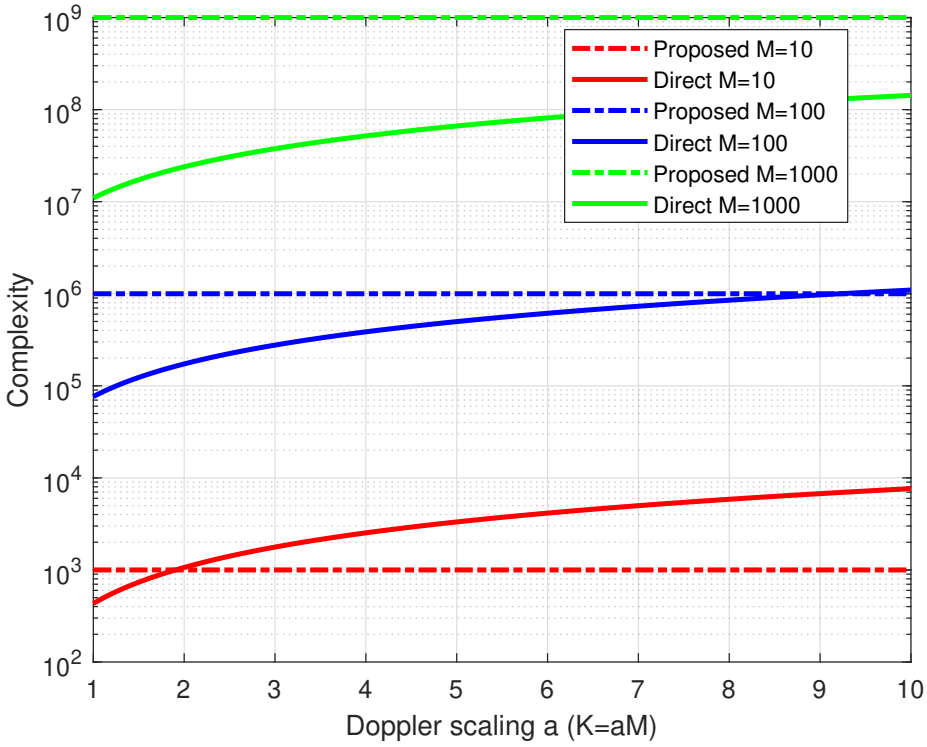

Figure 5.1: Scale of complexity to preserve uniform noise floor as a function of Doppler oversampling

to compute

$$\begin{aligned} \mathbf{G} &= \mathbf{D}^H \left(\bar{\mathbf{R}}_{cc}\right)^{-2} \mathbf{D} \\ &= \mathbf{D}^H \left(\bar{\mathbf{R}}_{cc}\right)^{-1} \tilde{\mathbf{D}} \quad = \mathbf{D}^H \tilde{\tilde{\mathbf{D}}} \end{aligned} \tag{5.14}$$

with $\mathscr{O}(M^2 D)$ complexity and thus the noise-scale factor will have a cost of $\mathscr{O}(M^2 D + K\log(K))$. Though again, the reduction in complexity for this approach will be dependent on what initial solver was used for $\bar{\mathbf{R}}_{cc}$ as forming (5.13) will require an additional forward-backward substitution or inverse Block-Toeplitz multiply. As such, there remains further room for complexity reduction.

Similar to MiCRFt, the Toepltiz-block structure may be leveraged to iteratively approximate the NIMPC filters using PCG with block-Circulant preconditioners via (4.15), but again these approaches tend to leave an unclustered group of eigenvalues that slow convergence. That being said, PCG does not require explicit storage of $\bar{\mathbf{R}}_{cc}$, which provides some practical benefits. Even without a preconditioner, CG should converge in at most $K_c(D+N)$ iterations (corresponding to how many unclustered eigenvalues there are) and has a per iteration cost of $\mathscr{O}(K_c M(D+N-$

$1)\log(D+N-1))$. Meaning for complete convergence for all M vectors in $\mathbf{D}$, the complexity would be $\mathcal{O}(K_c^2 M^2 (D+N)^2 \log(D+N-1))$, though stopping early can often yield filters just as efficacious. I did explore said implementation further in [134], but I have since concluded that for this NIMPC framework the direct approaches proposed in the previous section typically offer more efficacious solutions and a manageable computational burden. That being said, a reposing of the NIMPC clutter cancellation approach yields a model that is extremely amenable to PCG and can provide a reduction in complexity by several orders of magnitude.

5.2 NIMPC as a Projection

Alternatively, a nullspace projection away from K_c Doppler frequencies similar to (2.115) can likewise be formulated for the NIMPC model. Unlike (2.115), this NIMPC nullspace projection (Proj-NIMPC) requires incorporation of the fast-time dimension to retain the joint range-Doppler structure. As such, the projection objective must utilize the "full" receive model (akin to 2.21) to subsume all fast-time snapshots. This approach is inspired by the extensive cancellation algorithm [135, 136] used in passive bistatic radar [137]. From a high level, ECA leverages the fast-time Doppler model of (2.76) to form a projection onto the nullspace of clutter. Here, this philosophy is generalized for a pulsed-agile-radars, where by the "full" range model of (2.21) and the respective slow-time approximations from (2.83) are combined to yield a projection that can be implemented extremely efficiently.

Again, the absence of multiple-time-around-clutter and pulse eclipsing is assumed for brevity, though, they can be incorporated into Proj-NIMPC without loss of generality. Thus, the "full" $LM \times 1$ fast-time slow-time receive vector is expressed as

$$\dot{\mathbf{u}} = \sum_{\omega \in \Omega} \dot{\mathbf{V}}(\omega)\, \dot{\mathbf{x}}(\omega) + \dot{\mathbf{n}} \tag{5.15}$$

where

$$\dot{\mathbf{u}} = \begin{bmatrix} \mathbf{u}_1^T & \mathbf{u}_2^T & \cdots & \mathbf{u}_M^T \end{bmatrix}^T \tag{5.16}$$

is the concatenated $L \times 1$ measurement vectors for each pulse,

$$\dot{\mathbf{V}}(\omega) = \begin{bmatrix} e^{j\omega\frac{0}{f_{prf}}}\mathbf{V}_1^T & e^{j\omega\frac{1}{f_{prf}}}\mathbf{V}_2^T & \cdots & e^{j\omega\frac{M-1}{f_{prf}}}\mathbf{V}_M^T \end{bmatrix}^T \tag{5.17}$$

is the $ML \times (L-N+1)$ collection of concatenated $L \times (L-N+1)$ convolution matrices

$$\mathbf{V}_m = \begin{bmatrix} s_{1,m} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ s_{N,m} & & \ddots & 0 \\ 0 & \ddots & & s_{1,m} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & s_{N,m} \end{bmatrix}, \tag{5.18}$$

$\dot{\mathbf{x}}(\omega)$ is a $(L-N+1) \times 1$ vector of fast-time complex scattering coefficients for Doppler frequency ω and $\dot{\mathbf{n}}$ is the $ML \times 1$ collection of noise samples.

A projection onto the nullspace of all range samples spanning the $\tilde{\Omega}$ Doppler swath is then formed by minimizing

$$J = \|\dot{\mathbf{u}} - \tilde{\dot{\mathbf{V}}}\gamma\|^2 \tag{5.19}$$

where

$$\tilde{\dot{\mathbf{V}}} = \begin{bmatrix} \dot{\mathbf{V}}(\tilde{\omega}_1) & \dot{\mathbf{V}}(\tilde{\omega}_2) & \cdots & \dot{\mathbf{V}}(\tilde{\omega}_{K_c}) \end{bmatrix} \tag{5.20}$$

which yields the $LM \times LM$ projection matrix

$$\dot{\Gamma}_{cc} = \mathbf{I} - \tilde{\dot{\mathbf{V}}}\left(\tilde{\dot{\mathbf{V}}}^H\tilde{\dot{\mathbf{V}}}\right)^{-1}\tilde{\dot{\mathbf{V}}}^H . \tag{5.21}$$

Unlike the original form of NIMPC, Proj-NIMPC's storage and complexity are dominated by the

$K_c\,(L-N+1)\times K_c\,(L-N+1)$ Hermitian Toeplitz-Block-Toeplitz matrix

$$\tilde{\dot{\mathbf{V}}}^H\tilde{\dot{\mathbf{V}}}=\sum_{m=1}^{M}\begin{bmatrix}\mathbf{V}_m^H\mathbf{V}_m & e^{j(\omega_2-\omega_1)\frac{m-1}{f_{prf}}}\mathbf{V}_m^H\mathbf{V}_m & \cdots & e^{j(\omega_{K_c}-\omega_1)\frac{m-1}{f_{prf}}}\mathbf{V}_m^H\mathbf{V}_m\\ e^{j(\omega_1-\omega_2)\frac{m-1}{f_{prf}}}\mathbf{V}_m^H\mathbf{V}_m & \mathbf{V}_m^H\mathbf{V}_m & & \vdots\\ \vdots & & \ddots & \vdots\\ e^{j(\omega_1-\omega_{K_c})\frac{m-1}{f_{prf}}}\mathbf{V}_m^H\mathbf{V}_m & \cdots & \cdots & \mathbf{V}_m^H\mathbf{V}_m\end{bmatrix} \tag{5.22}$$

that does not scale with the slow-time dimension (M) and can be much cheaper than NIMPC when $K_c \ll M$. The clutter-canceled range-Doppler estimate is then formed by applying (5.21) to the measurements as

$$\dot{\mathbf{u}}_{cc}=\dot{\Gamma}_{cc}\dot{\mathbf{u}} \tag{5.23}$$

followed by standard range-Doppler matched filtering, though LS MMFs, MiCRFt, or APC can likewise be applied.

5.2.1 Efficient Solvers for Proj-NIMPC

To develop efficient solvers for Proj-NIMPC, the projection is separated as

$$\begin{aligned}\dot{\mathbf{u}}_{cc}&=\dot{\Gamma}_{cc}\dot{\mathbf{u}}\\&=\dot{\mathbf{u}}-\tilde{\dot{\mathbf{V}}}\left(\tilde{\dot{\mathbf{V}}}^H\tilde{\dot{\mathbf{V}}}\right)^{-1}\tilde{\dot{\mathbf{V}}}^H\dot{\mathbf{u}}\\&=\dot{\mathbf{u}}-\tilde{\dot{\mathbf{V}}}\left(\tilde{\dot{\mathbf{V}}}^H\tilde{\dot{\mathbf{V}}}\right)^{-1}\dot{\mathbf{u}}_1\\&=\dot{\mathbf{u}}-\tilde{\dot{\mathbf{V}}}\dot{\mathbf{u}}_2\\&=\dot{\mathbf{u}}-\dot{\mathbf{u}}_3\end{aligned} \tag{5.24}$$

which has four major operations. First, there is a Toeplitz-block matrix-vector multiplication

$$\dot{\mathbf{u}}_1=\tilde{\dot{\mathbf{V}}}^H\dot{\mathbf{u}} \tag{5.25}$$

followed by a Toeplitz-Block-Toeplitz linear system solve

$$\left(\tilde{\mathbf{V}}^H\tilde{\mathbf{V}}\right)\dot{\mathbf{u}}_2 = \dot{\mathbf{u}}_1 \tag{5.26}$$

another Toeplitz-block matrix-vector multiplication

$$\dot{\mathbf{u}}_3 = \tilde{\mathbf{V}}\dot{\mathbf{u}}_2 \tag{5.27}$$

and a final vector-vector subtraction

$$\dot{\mathbf{u}}_{cc} = \dot{\mathbf{u}} - \dot{\mathbf{u}}_3. \tag{5.28}$$

Operations (5.25) and (5.27) can be computed efficiently with frequency domain multiplication with a complexity of $\mathscr{O}(K_cML\log(L))$ and is extremely parallelizable ($\mathscr{O}(L\log(L))$), while (5.28) has a negligible cost as it corresponds to trivial elementwise subtraction ($\mathscr{O}(MD)$). As such, the computational complexity is largely dominated by the $K_cL \times K_cL$ Toeplitz-block-Toeplitz linear system solve in (5.26), which can be permuted to yield a multi-channel levinson complexity of either $\mathscr{O}(K_c^3L^2)$ or $\mathscr{O}(L^3K_c^2)$, though $K_c \ll L$ so the former is more desirable. Banded-Cholesky will typically offer even lower complexity as the decomposition will cost $\mathscr{O}(K_c^2N^2L)$ and forward backward substitution will be $\mathscr{O}(K_cNL)$. The aggregate order of complexity (without parallelization) for a direct solution for Proj-NIMPC would then be

$$\mathscr{O}(K_cML\log(L) + K_c^3L^2)$$

using a multi-channel levinson approach and

$$\mathscr{O}(K_cML\log(L) + K_c^2N^2L + K_cNL)$$

for a Cholesky approach.

Proj-NIMPC is also very amenable to a PCG implementations. As the Hessian is now Toeplitz-

block-Toeplitz, an additional computational speedup is possible where and the per iteration matrix-vector multiplication can be performed with $\mathscr{O}(LK_c\log(2K_c)\log(L))$ by leveraging the Toeplitz properties in both domains.

While one would expect Circulant-block preconditioners (4.15) (or even Circulant-Block-Circulant) to offer adequate performance as each block is Toeplitz, empirically diagonal-block Circulant of the form

$$\bar{\mathbf{C}} = \begin{bmatrix} \mathbf{C} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{C} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{C} \end{bmatrix} + \sigma^2\mathbf{I} \tag{5.29}$$

where $\mathbf{C}$ is the Circulant approximation to $\sum_{m=1}^{M}\mathbf{V}_m^H\mathbf{V}_m$ provides exceptional convergence and multiplication with it's inverse can be applied in $\mathscr{O}(K_cL\log(L))$ per iteration (or parallelized to $\mathscr{O}(L\log(L))$). The more sophisticated Circulant-block preconditioner generally converges slightly slower and has a cost of $\mathscr{O}(K_c^2L+K_cL\log(L))$ for inverse construction (via Cholesky decomposition) and a per iteration cost of $\mathscr{O}(K_cL+K_cL\log(L))$ (though, again the FFTs can be parallelized for an alternative per iteration complexity of $\mathscr{O}(K_cL+L\log(L))$). Figure provides an example PCG performance for a somewhat extreme case where $M = 1000$ PROFM waveforms with $\kappa = 6$ and $N = 950$ were used to define a projection applied to a receive interval of $L = 10,000$ samples and $K_c = 20$ Doppler steering vectors spanning $\pm$20m/s defined the projection span. Even though the diagonal-block Circulant preconditioner does not account for the off-diagonal contributions of the hessian, it out-paces CG and the Circulant-block preconditioner and has practically converged after only 10 iterations.

5.3 Open-Air Experimental Assessment

We now seek to apply these efficient NIMPC implementation to the open-air experimental PROFM CPIs used in Sections 3.4 and 4.6. To give a sense of complexity improvement, the corresponding parameters and dimensionality are tabularized in Table 5.1. Note, that while the 4.5μsecond pulse

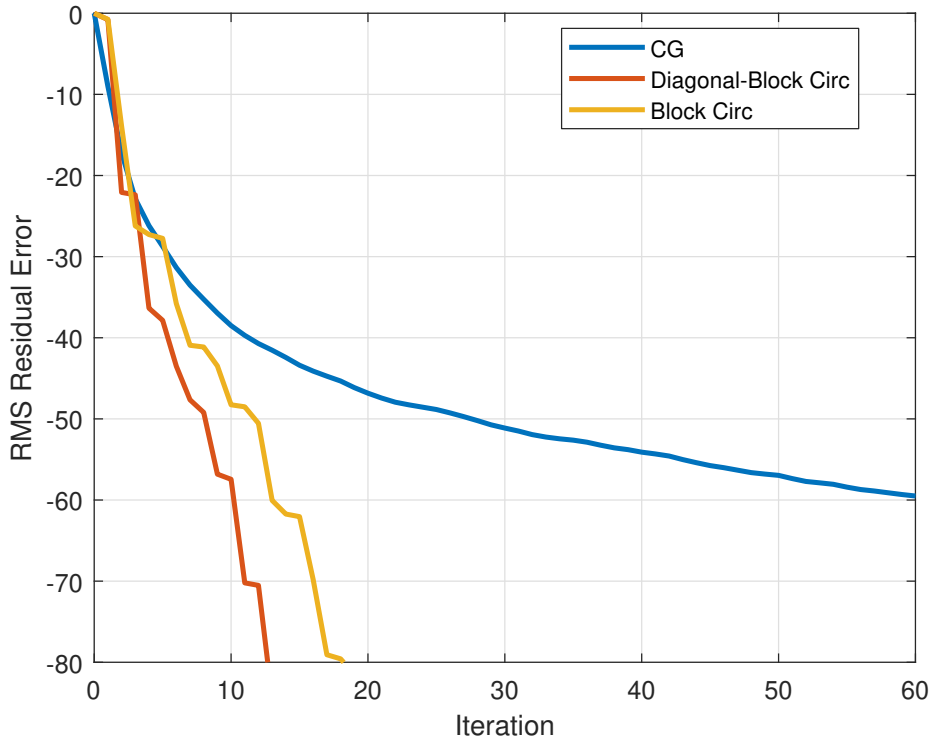

Figure 5.2: Illustrative PCG performance of Proj-NIMPC

duration would normally correspond to $N = 225$ samples, here the value of $N = 276$ is used as the corresponding MFs and linear models are constructed using hardware-in-the-loop measurements of $\mathbf{s}_m$ where the respective anti-aliasing filter has led to an extended response. As the clutter ride is relatively narrow for our case, the clutter cancellation dimensionality was set to $K_c = 3$ where the corresponding clutter Doppler shifts were set to -5, 0, and 5Hz.

For these parameters were used to evaluate the order of complexity for solving $\tilde{\mathbf{D}}$ with NIMPC or applying Proj-NIMPC for various implementations as provided in Table 5.2. Likewise, Table 5.3 provides a tabularized look at order of complexity when using the respective approaches perform range-Doppler estimation. Here the text in parenthesis corresponds to the respective solver used with chol corresponding to Cholesky decomposition, lev corresponding to multi-channel Levinson recursion, and PCG for a fixed number of iterations. The Reformulated NIMPC values correspond to the Woodbury form from (5.6) and leverage the efficient unity gain formulation from (5.10) when performing range-Doppler processing. Note, these values are not true flop counts as complexity notation drops constants and redundant scalars, rather, these numbers are here to just provide a

Measurement Parameters		
Parameter	Variable	Value
Sample Rate	f_s	50MHz
3dB Bandwidth	B	33.3MHz
Pulse Duration	T	4.5μseconds
Samples in Waveform	N	276
Number of Pulses	M	1000
Fast-time Samples processed	L	2500
Fast-Time Filter Samples	D	N
Number of Doppler Bins	K	3000
Number of Clutter Bins	K_c	3
NIMPC CNR	$\sigma_c^2/\sigma_{\mathbf{n}}^2$	10^5

Table 5.1: Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms

sense of scalability and performance.

The reformulated NIMPC approach solvers see a $\sim$3 orders of magnitude improvement in complexity order over that of the approach from [107] as K_c is fairly small. As K_c approaches M, these orders of complexity will start to converge, however, the efficient calculation of the unity scale factor will always provide a reduction in complexity. For this particular CPI, the multi-channel levinson approaches were the most efficient direct solvers which likely relates to the fact that $D = N$ and may not always be true in general. The direct implementations of Proj-NIMPC sees another complexity reduction of $\sim 1-2$ orders of magnitude over their NIMPC counterparts. The iterative approximation of Proj-NIMPC yields further complexity reduction over their respective direct solvers and also provide additional parallelization storage benefits not illustrated by the order of complexity. However, unlike the NIMPC filter approaches, Proj-NIMPC cannot provide a unity gain response and some distortion will be present in the Doppler-bins near the respective clutter cancellation swath.

Next the direct NIMPC filters were solved using the reformulated framework and applied to the open-air measurements. First, we examine the a baseline in Figure 5.3 where the NIMPC filters leveraged the unity scale factor. Again the left range-Doppler image corresponds to a zoomed out look that includes the direct path contribution while the right image focuses on the intersection with cars. Here NIMPC is able to suppress the direct path RSM to the same floor the MiCRFt achieved

Order of complexity to form filters/apply projection	
NIMPC (Brute Force)	2.10E+19
Efficient NIMPC [107]	7.62E+13
Reformulated NIMPC (Lev)	1.51E+10
Reformulated NIMPC (Chol)	1.59E+10
Proj-NIMPC (Lev)	2.53E+08
Proj-NIMPC (Chol)	1.80E+09
Proj-NIMPC (PCG 5 iter)	8.62E+07
Proj-NIMPC (CG 5 iter)	8.56E+07

Table 5.2: Complexity for solving $\tilde{\mathbf{D}}$ with NIMPC or applying Proj-NIMPC

Aggregate order of complexity for range-Doppler processing	
NIMPC (Brute Force)	2.10E+19
Efficient NIMPC [107]	7.62E+13
Reformulated NIMPC (Lev)	4.72E+10
Reformulated NIMPC (Chol)	4.80E+10
Proj-NIMPC (Lev)	3.72E+08
Proj-NIMPC (Chol)	1.92E+09
Proj-NIMPC (PCG 5 iter)	2.05E+08
Proj-NIMPC (CG 5 iter)	2.04E+08

Table 5.3: Complexity for Range-Doppler Processing with NIMPC and Proj-NIMPC

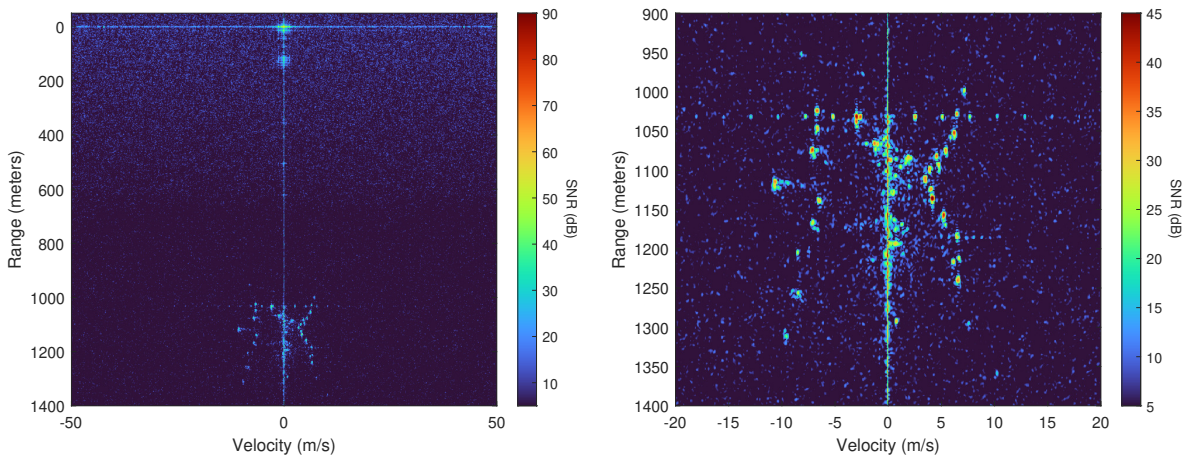

Figure 5.3: Range-Doppler Estimate for PROFM CPI using the NIMPC with unity gain scaling

but without the reduction in unambiguous velocity. The Doppler profiles of the intersection is likewise fairly clean and no clear sidelobes are visible since the transmit waveforms are unique pulse-to-pulse. As with MiCRFt, there is still residual background speckle/noise slightly higher than that of the LFM case, possibly due to RSM from the non-zero Doppler cars or from signal model-error. Additionally, the clutter ridge itself appear super-resolved in range and Doppler. This effect stems from the relationship between the LS MMF and NIMPC, as the trivial $M = 1$ case would yield a NIMPC filter equivalent to the ISL LS MMF without any beamspoiling. As this super-resolution effect will likely induce significant noise gain (especially since $D = N$), the clutter ridge estimates are prone to errors and would likely be excised from any detector outputs.

To illustrate this noise sensitivity, next these same NIMPC filters were used to generate the uniform noise scaling denominator, and this respective range-Doppler provided in Figure 5.4. Here the clutter ridge and some of the adjacent scattering has now been suppressed and the remaining signatures all correspond to presumed moving non-zero Doppler scattering. The consistent noise power across Doppler acts as a convenient surface for subsequent detection, however, it will increase the computational burden relative to the unity scale factor.

Next, Proj-NIMPC was applied to the measurements and subsequent range-Doppler estimation was performed using the MF as illustrated in Figure 5.5. As NIMPC is a projection, there is now a distinct null near zero-Doppler and adjacent bins are either slightly suppressed or distorted. The

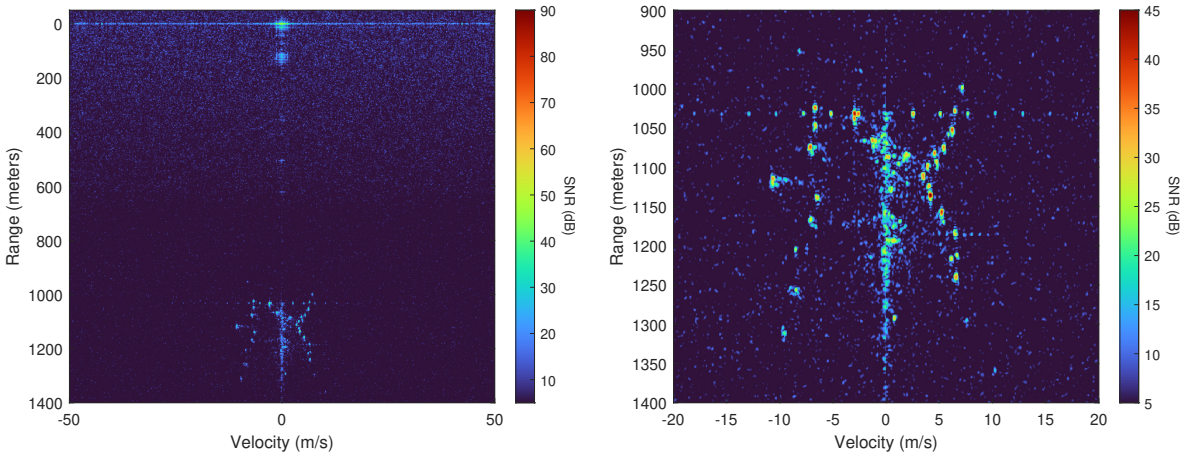

Figure 5.4: Range-Doppler Estimate for PROFM CPI using the NIMPC with consistent noise power scaling

noise scaling parameter for Proj-NIMPC can likewise be computed, however doing so is computationally intensive. Thus, there will generally be some small but unknown noise distortion with Proj-NIMPC. That being said, the RSM suppression is on par with that of NIMPC for a much lower computational cost and the intersection estimate is fairly consistent with all previous results.

Using the same parameters, 5 iterations of CG were performed in order to approximate Proj-NIMPC. The resulting range Doppler map is shown in Figure 5.6 where the direct path and intersection RSM has been suppressed by $\sim$20dB relative to the matched filter. That being said, this approximation with only 5 iterations and no preconditioner could not achieve the same estimate fidelity as the direct approaches. Next PCG is used to approximate Proj-NIMPC with the diagonal-block Circulant preconditioner as shown in Figure 5.7 where after only 5 iterations, PCG is able to yield effectively same the RSM suppression as the direct solver.

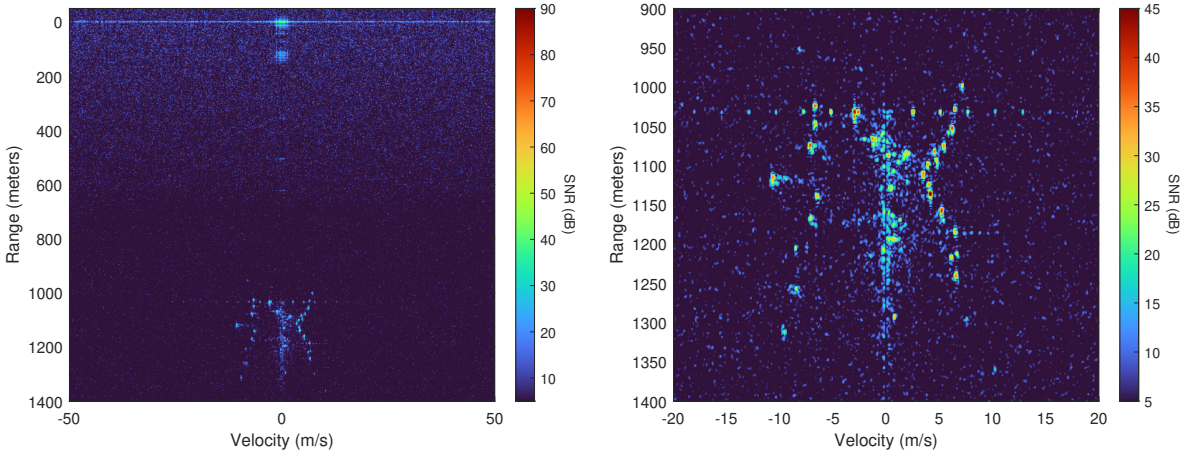

Figure 5.5: Range-Doppler Estimate for PROFM CPI using direct solution of Proj-NIMPC

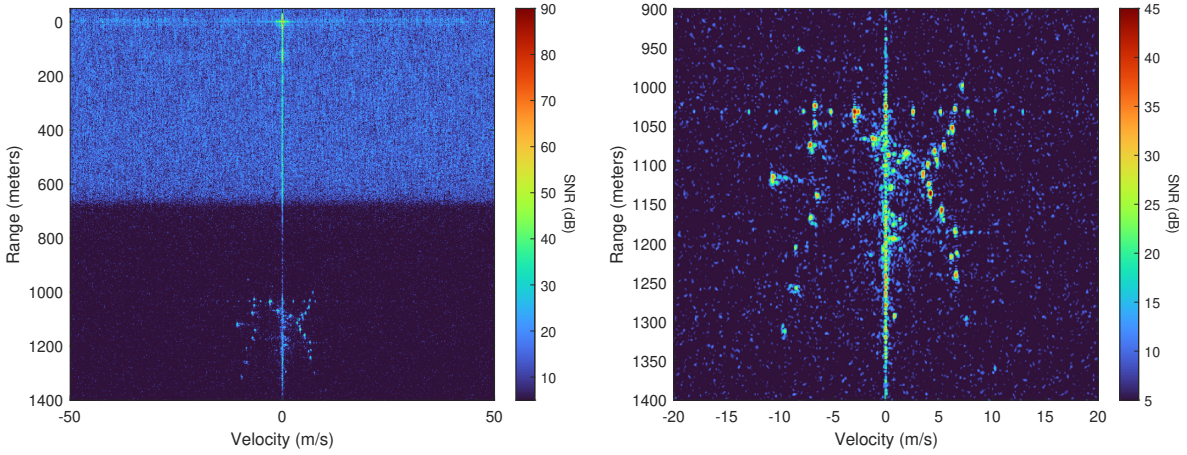

Figure 5.6: Range-Doppler Estimate for PROFM CPI using 5 iterations of CG to approximate Proj-NIMPC

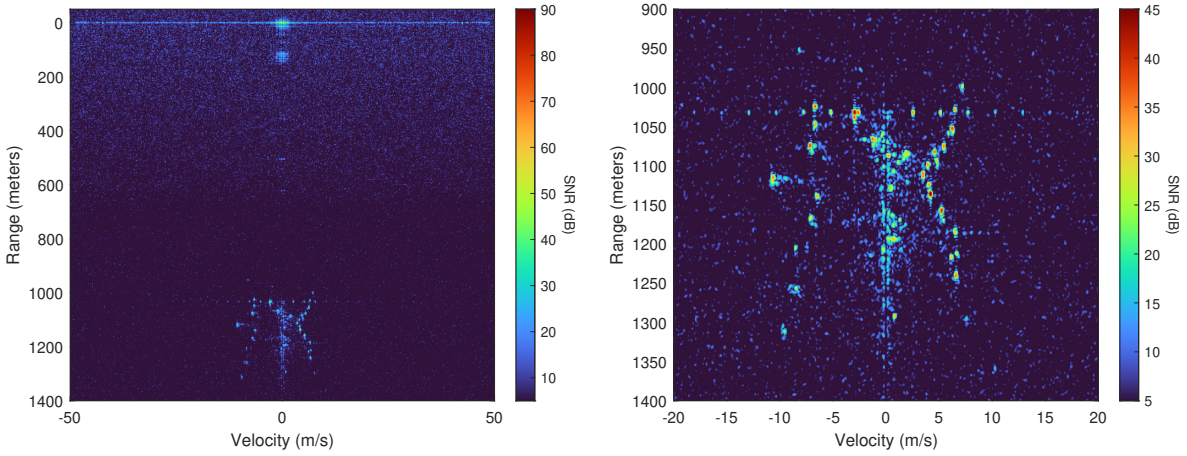

Figure 5.7: Range-Doppler Estimate for PROFM CPI using 5 iterations of PCG to approximate Proj-NIMPC

Chapter 6

Robust Adaptive Pulse Compression

While the LS approaches can offer improvements over the MF, they often necessitate some form of tradespace, (e.g., the complementary subset combination of MiCRFt, or canceling certain Doppler slices with NIMPC) to yield enough DoF to sufficiently suppress RSM (to the model error/noise floor). The APC approaches on the other hand, offer improved sidelobe suppression, but are more computationally intensive and are not robust to straddling unless the 3dB spectral oversampling is large. Here we seek to extend the straddling compensation framework to the APC and provide alternative techniques and representations to offer a more manageable computational load.

6.1 Incorporating Range-Straddling Model Uncertainty

Similar to the approach in Section 3.1, range straddling in APC can be addressed by incorporating model based uncertainty. In [102], the effects of sub-sample straddling with LFM chirp waveforms were compensated via a CMT weighting on the structured covariance. For an arbitrary waveform, this same framework is not necessary extensible, though, if the passband-waveform is characterized by a specific phase-basis, i.e., [50, 115–118], a similar CMT methodology can be derived. [138, 139] attempted to directly estimate the straddling offset in a manner akin to [92], but again required a phase-based waveform model and necessitated additional refinement stage that increased computational burden or required a dimensionality reduction (less adaptive DoF). However, these phase-based models do not typically incorporate distortion induced by the RX anti-alias filter, which may be significant for waveforms less spectrally contained than chirps. Thus, we again leverage loopback captures of the waveform to formulate our recieve model and incorporate

straddling uncertainty via the interpolation-based approach from Section 3.1.

Under the assumption that range-straddling is independent from the scatterer-variance, the straddle-uncertain structured covariance will have the form

$$\begin{aligned}\mathbf{R}_{\mathbf{s}}(\ell) &= \mathscr{E}\{\mathbf{S}(\gamma)\mathbf{P}(\ell)\mathbf{S}^H(\gamma)\} \\ &= \mathscr{E}\left\{\sum_{i=-N+1}^{D-1} p_{i+N}(\ell)\mathbf{J}_i\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\mathbf{J}_i^H\right\} \\ &= \sum_{i=-N+1}^{D} p_{i+N}(\ell)\mathbf{J}_i\mathscr{E}\{\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\}\mathbf{J}_i^T \\ &= \sum_{i=-N+1}^{D} p_{i+N}(\ell)\mathbf{J}_i\left(\int_{-0.5}^{0.5}\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\,d\gamma\right)\mathbf{J}_i^T\end{aligned} \tag{6.1}$$

where

$$p_i(\ell) = \mathbf{P}_{i,i}(\ell)$$

corresponds to the diagonal elements of the $D+N-1 \times D+N-1$ estimate source covariance $\mathbf{P}(\ell)$ and $\mathbf{J}_i$ is a reversing and shifting matrix such that matrix-vector multiplication with an arbitrary $N \times 1$ vector $\mathbf{b}$ yields a $D \times 1$ vector

$$\mathbf{x} = \mathbf{J}_i\mathbf{b}$$

with at most N nonzero elements

$$x_d = \begin{cases} b_{i-d} & 1 \le i-d \le N \\ 0 & else \end{cases} \tag{6.2}$$

for $d = 1, 2, \cdots, D$.

As with the LS MMFs, the only unknown in (6.1) is the straddled-signal correlation matrix

$$\mathscr{E}\{\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\} = \int_{-0.5}^{0.5}\mathbf{s}(\gamma)\mathbf{s}^H(\gamma)\,d\gamma = \mathbf{R}_\gamma \tag{6.3}$$

which can be approximated via a sample correlation matrix or the sinc interpolation model

$$\mathbf{R}_{\gamma} = \frac{1}{Z^2}\mathbf{I}_{N\times Z}\mathbf{F}^H\left(\mathbf{T}_{\gamma}\odot\mathbf{s}_f\mathbf{s}_f^H\right)\mathbf{F}\mathbf{I}_{N\times Z}^T \tag{6.4}$$

from Section 3.1 where $\mathbf{T}_{\gamma}$ is the respective sinc-CMT. (6.1) reveals that $\mathbf{R}_{\mathbf{s}}(\ell)$ retains a matrix-bandwidth of N and hermitian structure (since $\mathbf{R}_{\gamma}$ is a hermitian matrix). Note, $\mathbf{R}_{\gamma}$ can again be approximated as low-rank to enable efficient beamspoiling downdates.

As general model uncertainty covariance is corresponds to the diagonal of (6.1),

$$\mathbf{R}_{\mathbf{z}}(\ell) = \sigma_{\mathbf{z}}^2\mathbf{I}\odot\mathbf{R}_{\mathbf{s}}(\ell)$$

too can be written in terms of the straddled-signal correlation matrix and can the be constructed via

$$\begin{aligned}\mathbf{R}_{\mathbf{sz}}(\ell) &= \mathbf{R}_{\mathbf{s}}(\ell)+\mathbf{R}_{\mathbf{z}}(\ell)\\ &= \sum_{i=-N+1}^{D} p_{i+N}(\ell)\mathbf{J}_i\mathbf{R}_{\gamma}\mathbf{J}_i^T + \sigma_{\mathbf{z}}^2\mathbf{I}\odot\left(p_{i+N}(\ell)\mathbf{J}_i\mathbf{R}_{\gamma}\mathbf{J}_i^T\right)\\ &= \sum_{i=-N+1}^{D} p_{i+N}(\ell)\mathbf{J}_i\mathbf{R}_{\gamma,\mathbf{z}}\mathbf{J}_i^T\end{aligned} \tag{6.5}$$

where

$$\mathbf{R}_{\gamma,\mathbf{z}} = \mathbf{R}_{\gamma}+\sigma_{\mathbf{z}}^2\mathbf{I}\odot\mathbf{R}_{\gamma} \tag{6.6}$$

is the straddled-signal plus model uncertainty correlation matrix.

Next, straddling uncertainty is incorporated into the cross correlation component as

$$\begin{aligned}\mathbf{r}_{\mathbf{y}x}(\ell) &= \mathscr{E}\{\mathbf{S}(\gamma)\}\mathscr{E}\{\mathbf{x}(\ell)x^*(\ell)\}\\ &= \rho(\ell)\mathscr{E}\{\mathbf{S}(\gamma)\}\mathbf{e}\\ &= \rho(\ell)\mathscr{E}\{\mathbf{v}(\gamma)\}\\ &= \rho(\ell)\mathbf{v}_{\gamma}\end{aligned} \tag{6.7}$$

which again follows the derivation in (3.5) and (3.13) as

$$\mathbf{v}_\gamma = \begin{bmatrix} \mathbf{0}^T & \mathbf{s}_\gamma^T & \mathbf{0}^T \end{bmatrix}^T \tag{6.8}$$

where

$$\begin{aligned} \mathbf{s}_\gamma =& \mathscr{E}\{\mathbf{s}(\gamma)\} \\ =& \int_{-0.5}^{0.5} \frac{1}{Z} \mathbf{I}_{N\times Z} \mathbf{F}^H \left(\mathbf{d}(\gamma) \odot \mathbf{s}_f\right) d\gamma \\ =& \frac{1}{Z} \mathbf{I}_{N\times Z} \mathbf{F}^H \left(\tilde{\mathbf{d}} \odot \mathbf{s}_f\right) \end{aligned} \tag{6.9}$$

and

$$\tilde{\mathbf{d}} = \text{sinc}\left(\frac{1}{Z}\tilde{\mathbf{z}}\right)$$

acts as a frequency domain sinc taper.

For illustration purposes, this straddle-uncertainty was incorporated into the simulated APC configuration from Section 2.3. The left side of Figure 6.1 again shows the APC estimates for the $\kappa = 1.5$ when using beamspoiling, model uncertainty (to approximate the errors due to range straddling), and a unity constraint. Here, the dominant scatterer in the middle does not cohere to a single point and induces a roll-off due to straddling, masking the neighboring signatures. On the right side of the same figure, the same APC configuration was used, however, rather than approximating the straddling error with the generic model uncertainty parameter, the straddling compensation from (6.4) was used. Here all five sources are easily distinguished, and the roll-off from the dominant source is no longer present. Looking at the range-varying noise-power in Figure 6.2, we see that incorporation of straddling likewise offers a significant improvement in SNR as the structured covariance is more representative of the true measurements. For the $\kappa = 6$ estimates and noise powers in Figures 6.3 and 6.4 respectively, a similar trend is observed where the incorporation of straddle uncertainty improves the source separability and less overall noise gain.

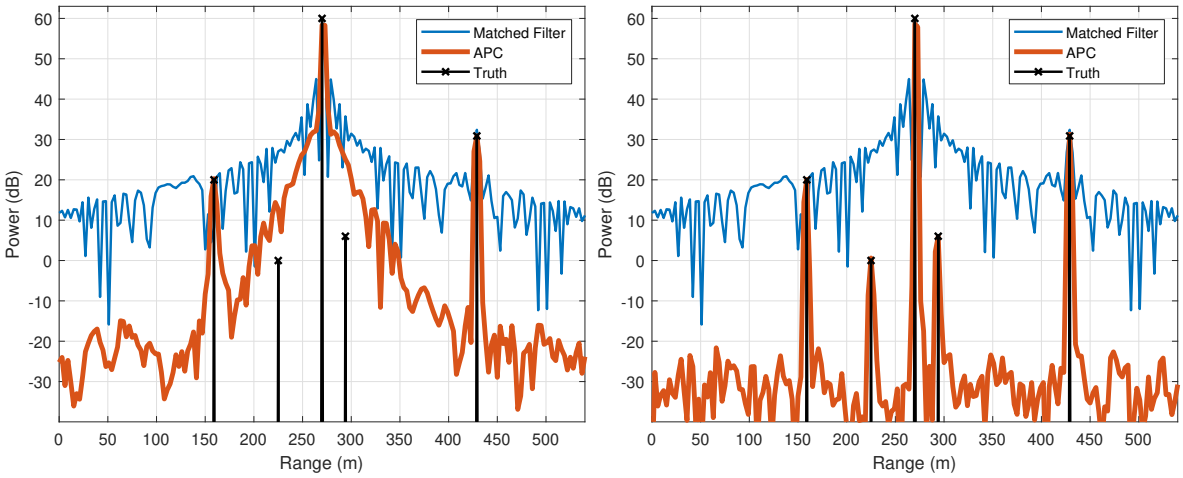

Figure 6.1: Simulated APC estimate for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right)

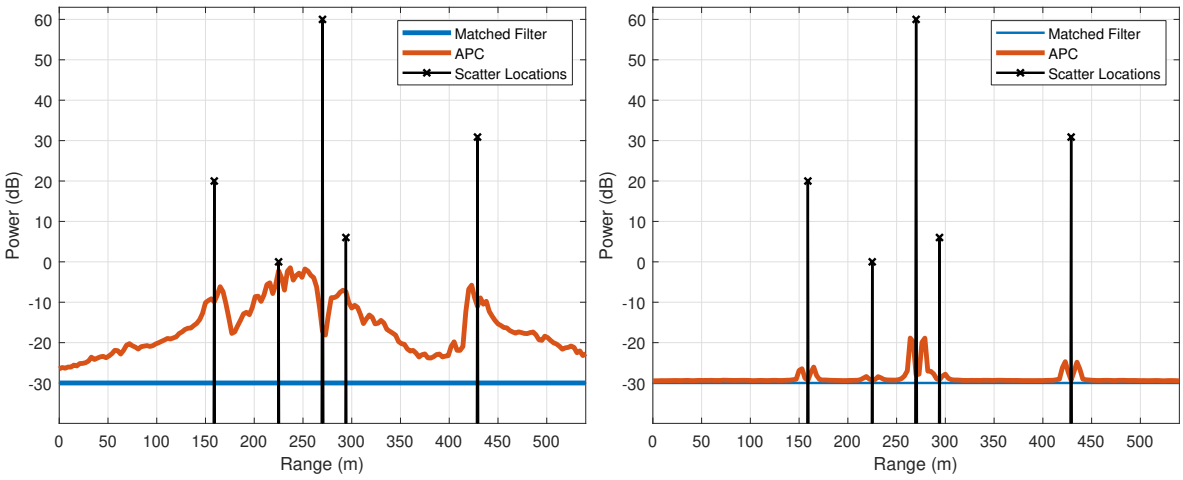

Figure 6.2: Simulated APC noise power for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right)

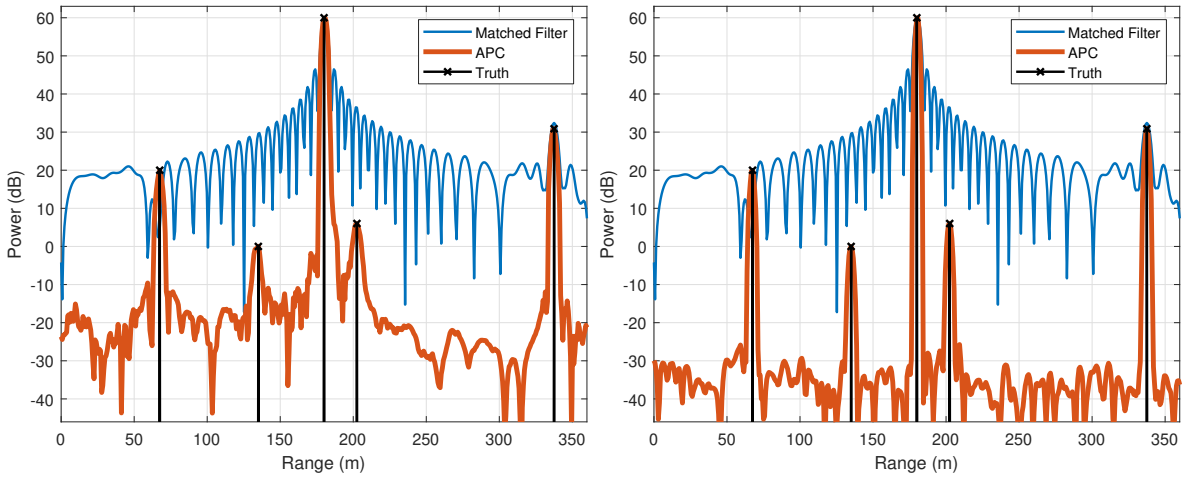

Figure 6.3: Simulated APC estimate for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right)

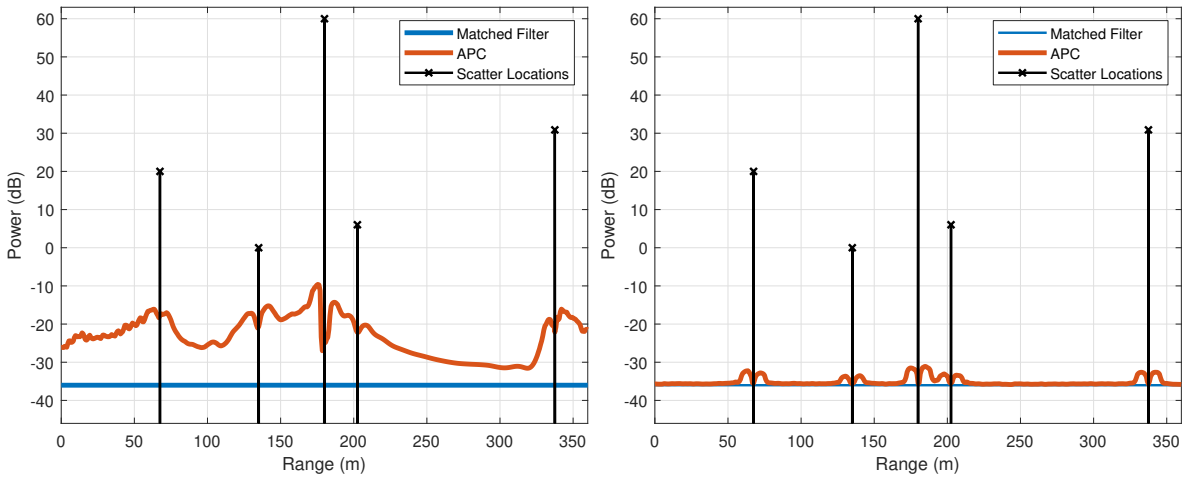

Figure 6.4: Simulated APC noise power for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs beamspoiling, straddle uncertainty, and a unity constraint (right)

6.1.1 Filter Beamspoiling

Rather than the covariance downdate/beamspoiling from Section 2.4.2, a "filter beamspoiling" approach akin to the template-based LS MMF framework can likewise be posed to help mitigate straddling degradation and MML. In [93] a localized "beamspoiling matrix" approach was examined using the structured covariance model of

$$\tilde{\mathbf{R}}_{\mathbf{s}} = \tilde{\mathbf{S}}\tilde{\mathbf{P}}\tilde{\mathbf{S}}^{H} \tag{6.10}$$

where

$$\tilde{\mathbf{S}} = \begin{bmatrix} \mathbf{S}_{\frac{D-N}{2}-Z} & \cdots & \mathbf{S}_{\frac{D-N}{2}} & \cdots & \mathbf{S}_{\frac{D-N}{2}+Z} \end{bmatrix} \tag{6.11}$$

are the center $2Z+1$ columns of $\mathbf{S}$ corresponding to the matched filter mainlobe (i.e., $Z = ceil\{\kappa\} - 1$) and $\tilde{\mathbf{P}}$ is a diagonal matrix with

$$diag\{\tilde{\mathbf{P}}\} = \begin{bmatrix} p_{\frac{D-N}{2}-Z} & \cdots & p_{\frac{D-N}{2}} & \cdots & p_{\frac{D-N}{2}+Z} \end{bmatrix} \tag{6.12}$$

are the respective $2Z+1$ source components where $\mathbf{P}_{i,i} = p_i$ which excludes the delay index ℓ for notational convenience. Note, this localized correlation matrix is similar to the downdate form in (2.65), however, this localized covariance additionally includes the center ($\frac{D-N}{2}$) elements. This beamspoiling matrix was then incorporated into the RMMSE filters as

$$\mathbf{w} = \frac{1}{\mathbf{v}^{H}\mathbf{v}}\mathbf{R}_{\mathbf{yy}}^{-1}\tilde{\mathbf{R}}_{\mathbf{s}}\mathbf{v} \tag{6.13}$$

where the beamspoiling matrix is effectively undoing the superresolution (i.e. decorrelating) effect imposed by the matrix inversion so that the estimated response does not collapse to nonphysical

point scatterers. Observe that when Z is set to zero (6.13) simplifies to (2.64) via

$$\begin{aligned}\mathbf{w} &= \frac{1}{\mathbf{v}^H\mathbf{v}}\mathbf{R}_{\mathbf{yy}}^{-1}\left(\mathbf{v}\rho\left(\ell\right)\mathbf{v}^H\right)\mathbf{v} \\ &= \rho\left(\ell\right)\mathbf{R}_{\mathbf{yy}}^{-1}\mathbf{v}\end{aligned}.$$

This approach can be generalized by relating (6.13) back to the original RMMSE objective function via

$$\begin{aligned}J &= \mathscr{E}\left\{\left\|\mathbf{g}^H\tilde{\mathbf{x}} - \mathbf{w}^H\mathbf{y}\right\|_2^2\right\} \\ &= \mathbf{w}^H\mathscr{E}\left\{\mathbf{y}\mathbf{y}^H\right\}\mathbf{w} - \mathbf{g}^H\mathscr{E}\left\{\tilde{\mathbf{x}}\mathbf{y}^H\right\}\mathbf{w} - \mathbf{w}^H\mathscr{E}\left\{\mathbf{y}\tilde{\mathbf{x}}^H\right\}\mathbf{g} + \mathbf{g}^H\mathscr{E}\left\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^H\right\}\mathbf{g}\end{aligned} \tag{6.14}$$

where $\tilde{\mathbf{x}}$ are the respective scatters corresponding to (6.12) and $\mathbf{g} = \frac{1}{\mathbf{v}^H\mathbf{v}}\tilde{\mathbf{S}}^H\mathbf{v}$ corresponds to the MF mainlobe and acts as the desired correlation between the scatters in $\tilde{\mathbf{x}}$. Effectively, (6.14) is a template-based RMMSE objective function where $\mathbf{g}$ is the desired filter response or scatterer correlation akin to $\mathbf{p}$ for the LS MMFs. Thus, $\mathbf{g}$ can replaced with any arbitrary desired response as opposed to simply the MF mainlobe as proposed in [93]. (6.14) then has the optimal filter

$$\mathbf{w} = \mathbf{R}_{\mathbf{yy}}^{-1}\tilde{\mathbf{S}}\tilde{\mathbf{P}}\mathbf{g} \tag{6.15}$$

where the equivalence with (6.13) can be easily shown by evaluating $\mathbf{g} = \frac{1}{\mathbf{v}^H\mathbf{v}}\tilde{\mathbf{S}}^H\mathbf{v}$.

While this template-based approach can be combined with covariance beamspoiling and a unity constraint, selecting an appropriate desired response and using the straddle compensated structured-covariance should provide some natural robustness towards the super-resolution and range straddling effects that typically hinder APC efficacy. Additionally, excluding these legacy robustness measures enables a significant reduction to computational complexity which will be examined further in Section 6.3. That being said, when a unity constraint is not incorporated, the template based-formulation can oversuppress small signatures, though, the degree of oversuppression can somewhat be controlled via an appropriate template selection.

The template-based APC formulation was then applied to the simulated APC measurements

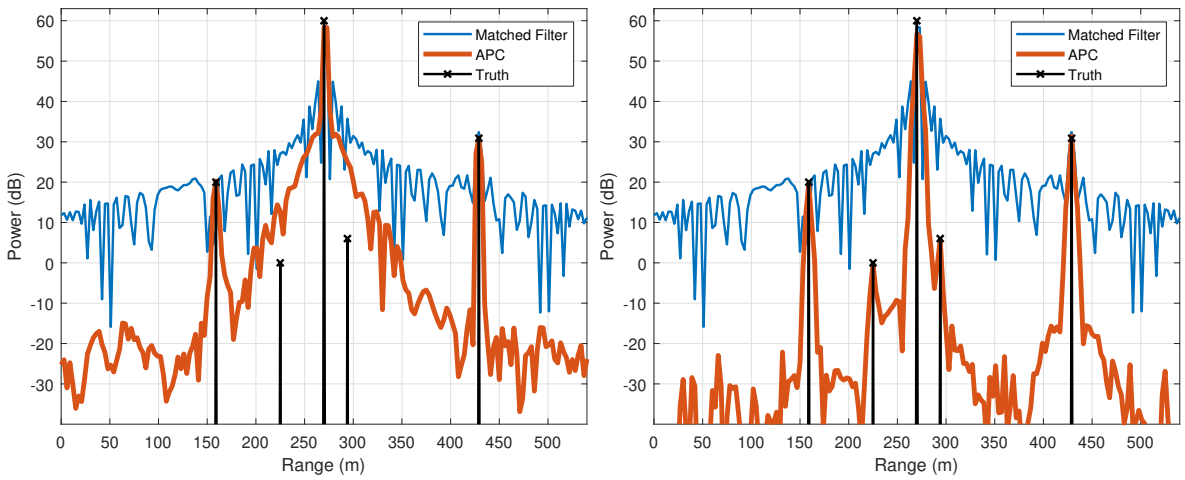

Figure 6.5: Simulated APC estimate for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right)

from Section 2.3. Here straddle uncertainty was incorporated in the same manner as Section 6.1 and the desired response corresponded to a Gaussian power-spectrum with the same 3dB bandwidth as the LFM waveform. Figures 6.5-6.6 illustrates the respective estimate and noise powers for the $\kappa = 1.5$ configuration while 6.7-6.8 correspond to the $\kappa = 6$ case. Again, the left side of each figure illustrates the previous "standard" of robust APC processing where beamspoiling, model uncertainty, and a unity constraint all all incorporated, while the right side of each figure corresponds to this new template-based approach with straddle uncertainty. For the $\kappa = 1.5$ estimate, the template-based approach sees improved source isolation as the dominant scatterer does not have the same extended roll-off, though, it is not able to clearly resolve the small scatterers. As there is no unity constraint, the estimate floor is unbounded and occasionally goes much lower than the noise floor. For the respective noise powers, the template-based approach yields lower average noise powers, but tends to have a notable noise gain centered about scatter locations. When subsequent Doppler processing is performed, this increase in expected noise power has the potential to induce a streaking effect across Doppler (since noise is not coherent) that may subsequently mask small scatterers. With the $\kappa = 6$ case, similar trends are observed, though, the increased oversampling has led to a tighter, smoother, and lower grouping of this noise gain.

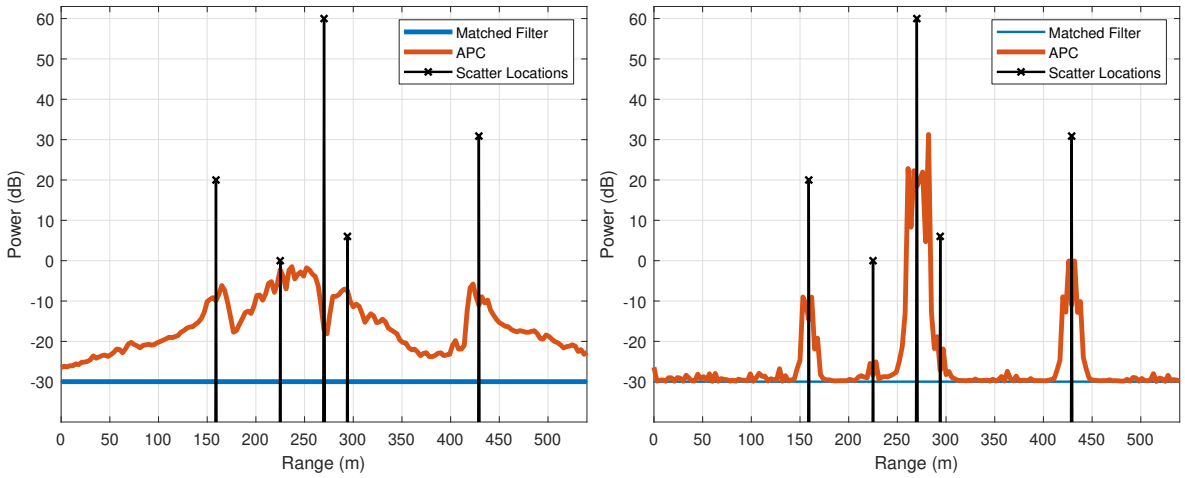

Figure 6.6: Simulated APC noise power for $\kappa = 1.5$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right)

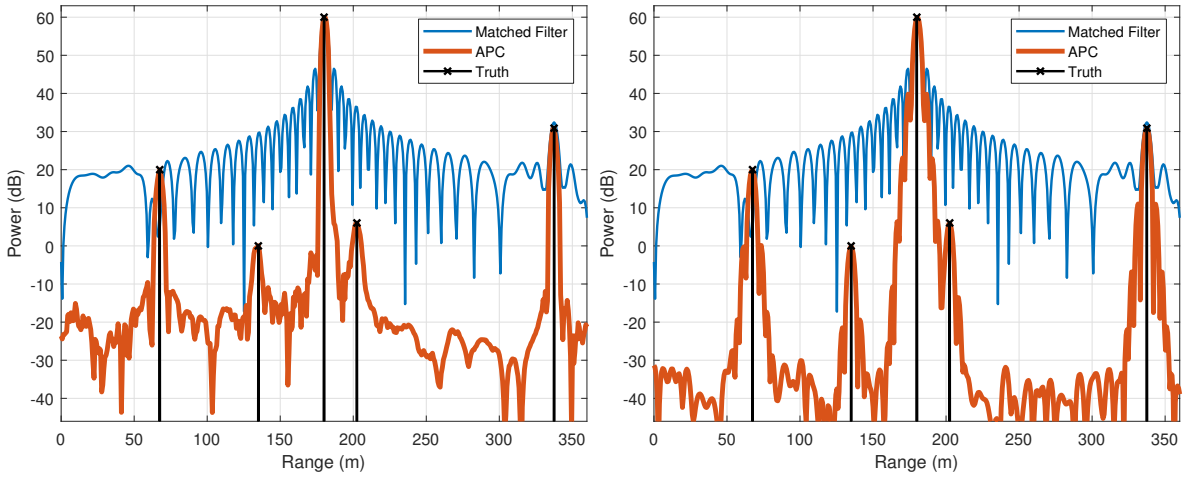

Figure 6.7: Simulated APC estimate for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right)

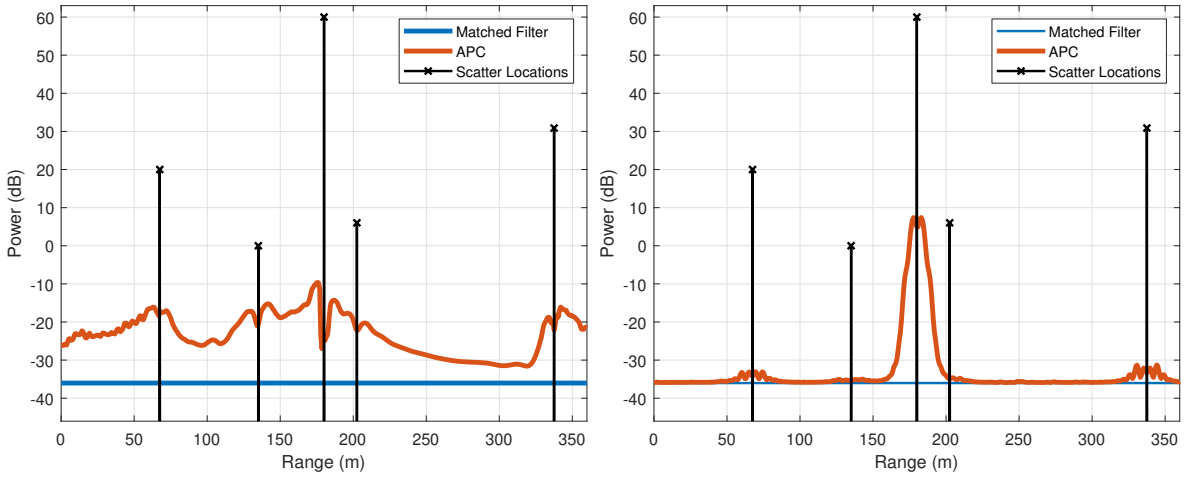

Figure 6.8: Simulated APC noise powers for $\kappa = 6$ recieve measurements using beamspoiling, model uncertainty, and a unity constraint (left) vs template-based objective and straddling uncertainty (right)

6.2 Full APC representation and relationship to snapshot-form

While the original APC formulation utilized a $D = N$ filter representation, the general snapshot notation (arbitrary D) used in this work can be used to extend the APC model to longer filters. This naturally increases the adaptive degrees of freedom, and can offer improved performance for less MML. The subsequent implementation of APC requires construction and inversion of a unique $D \times D$ structure covariance for every ℓ. With brute force, this corresponds to a complexity of $\sim\mathscr{O}(D^3)$[1] for both construction and inversion. Thus, in aggregate the complexity will be $\sim\mathscr{O}(D^3 L)$ every iteration.

This general philosophy of an extended filter lengths was examined in [140, p. 194-198] where a "block" processing methodology proposed to avoid forming and inverting a unique matrix for every delay index, thus, reducing computational cost and memory requirements. As the filter length is taken to the limit, the "full" $L \times 1$ linear receive model from (2.21)

$$\mathbf{u} = \mathbf{V}\mathbf{x} + \mathbf{n}$$

[1]The $\sim$ is included as incorporation of beamspoiling and unity constraints correspond to a small, but non-zero, increase in complexity. Likewise, construction has a cost of $\mathscr{O}(D^2(D+N-1))$ which after dropping the smaller terms yields $\mathscr{O}(D^3)$.

is used where the *i*th APC estimate is performed as

$$\hat{\mathbf{x}} = \mathbf{W}^H \mathbf{u} \tag{6.16}$$

with $\mathbf{W}$ corresponding to the $L \times L - N + 1$ bank of RMMSE filters computed via

$$\mathbf{W} = \mathbf{R}_{\mathbf{uu}}^{-1} \mathbf{V} \mathbf{P} \tag{6.17}$$

where $\mathbf{R}_{\mathbf{uu}}$ is the current estimate of the structured-covariance matrix. For the "base form" (i.e. prior to incorporating straddling or model uncertainties), the "full" structured covariance is then found as

$$\mathbf{R}_{\mathbf{uu}} = \mathbf{V}^H \mathbf{P} \mathbf{V} + \mathbf{R}_{\mathbf{n}}. \tag{6.18}$$

Uncertainties, beamspoiling, and a gain constraint can likewise be incorporated through a similar procedure as the snapshot-based model to yield a generic representation of

$$\mathbf{R}_{\mathbf{uu}} = \mathbf{R}_{\mathbf{V}} + \mathbf{R}_{\mathbf{z}} + \mathbf{R}_{\mathbf{n}} \tag{6.19}$$

where $\mathbf{R}_{\mathbf{V}}$, $\mathbf{R}_{\mathbf{z}}$, and $\mathbf{R}_{\mathbf{n}}$ are the structured-covariance matrices for the waveform/sources, model error uncertainty, and noise respectively.

Solving (6.17) necessitated explicitly forming $\mathbf{R}_{\mathbf{uu}}$ and multiple $L \times L$ linear systems, which through brute force corresponds to $\mathscr{O}(L^3)$ for each respective operation. However, as $\mathbf{V}$ is banded and $\mathbf{P}$ is diagonal, the "full" structured covariance matrix $\mathbf{R}_{\mathbf{uu}}$ will correspond to hermitian banded matrix with a matrix-bandwidth of N and can be constructed for $\mathscr{O}(N^2 L)$. The Cholesky decomposition of the full form will likewise have $\mathscr{O}(N^2 L)$ complexity and subsequent forward-backward substitution will cost $\mathscr{O}(NL)$.

The Toeplitz properties of $\mathbf{V}$ can also be leveraged to further reduce the cost of constructing $\mathbf{R}_{\mathbf{uu}}$ (as will be discussed in Section 6.2.1), but the Cholesky decomposition of the full representation still has $\mathscr{O}(N^2 L)$ complexity. However, extracting a $D \times D$ submatrix of $\mathbf{R}_{\mathbf{uu}}$ that correspond to the

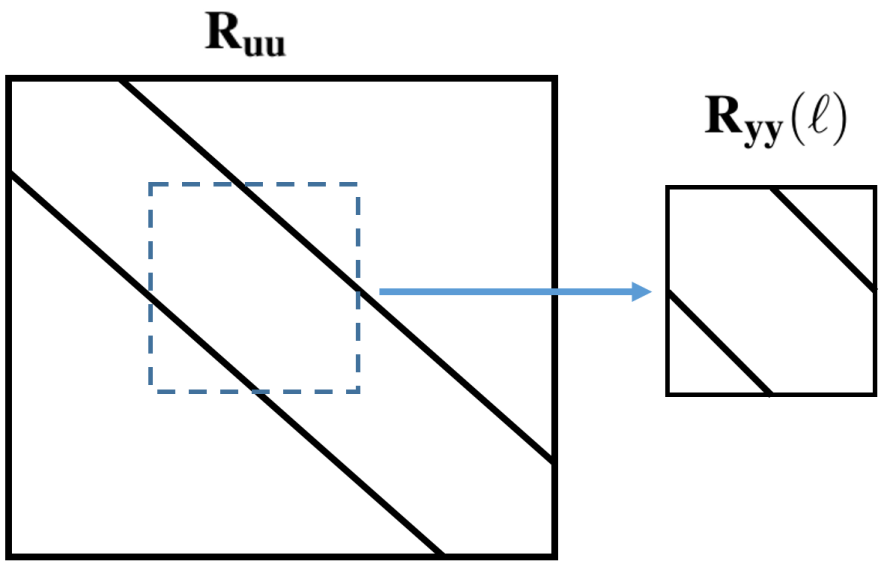

Figure 6.9: Illustration of Extracting the Snapshot Structured-Covariance from the "Full" Structured Covariance

respective indices of $\mathbf{y}(\ell)$ as

$$\mathbf{R_{yy}}(\ell) = \mathbf{J}(\ell)\mathbf{R_{uu}}\mathbf{J}^T(\ell) \tag{6.20}$$

where $\mathbf{J}(\ell)$ is a $D \times L$ extraction matrix with

$$\mathbf{J}(\ell)\mathbf{x} = \left[x_{\ell-1-\frac{D-N}{2}} \quad x_{\ell-\frac{D-N}{2}} \quad \cdots \quad x_{\ell+D-\frac{D-N}{2}-2} \right]^T,$$

will exactly yield the respective structured covariance matrix for the snapshot model, though, truncation or padding of $\mathbf{R_{yy}}(\ell)$ may be necessary to produce appropriate sized filters for the "edge" delay indices [87]. See Figure 6.9 for a visual illustration of this extraction. Thus, if construction of $\mathbf{R_{uu}}$ can be performed efficiently, extracting each $\mathbf{R_{yy}}(\ell)$ will not require any additional cost (previously $\mathcal{O}(D^2(D+N-1)) \approx \mathcal{O}(D^3)$ for every delay).

6.2.1 Constructing the Structured-Covariance Matrices via diagonal Convolution

From (6.5), we observe that the structured covariance follows a 2-Dimensional convolutional relationship with respect to $\mathbf{R}_{\gamma,\mathbf{z}}$ and the diagonal source covariance matrix

$$\mathbf{P}(\ell) = \begin{bmatrix} \rho\left(\ell - N + 1 - \frac{D-N}{2}\right) & & & & \\ & \ddots & & & \\ & & \rho(\ell) & & \\ & & & \ddots & \\ & & & & \rho\left(\ell + D - 1 + \frac{D-N}{2}\right) \end{bmatrix} \tag{6.21}$$

w.r.t. ℓ, which can be equivalently expressed as

$$\begin{aligned} \mathbf{R}_{\mathbf{sz}}(\ell) &= \mathbf{R}_{\mathbf{s}}(\ell) + \mathbf{R}_{\mathbf{z}}(\ell) \\ &= \sum_{i=-N+1}^{D} p_{i+N}(\ell)\, \mathbf{J}_i \mathbf{R}_{\gamma,\mathbf{z}} \mathbf{J}_i^T \\ &= \sum_{i=-N+1}^{D} \rho\left(\ell + i - \frac{D-N}{2}\right) \mathbf{J}_i \mathbf{R}_{\gamma,\mathbf{z}} \mathbf{J}_i^T \end{aligned} \tag{6.22}$$

where $p_i(\ell) = \mathbf{P}_{i,i}(\ell)$ corresponds to the ith element on the main diagonal of source correlation matrix. As the shifting operations are symetric with respect to the rows and columns, the diagonals of $\mathbf{R}_{\mathbf{sz}}(\ell)$ will correspond the convolution $\rho(\ell)$ and the diagonals of $\mathbf{R}_{\gamma,\mathbf{z}}$

Furthermore, as a change in ℓ will simply shift the respective $p_i(\ell)$,

$$diag\{\mathbf{R}_{\mathbf{sz}}(\ell)\}_d = \begin{cases} \mathbf{U}_d \mathbf{p}_{\mathbf{x}}(\ell) & 0 \le d \le N-1 \\ diag\{\mathbf{R}_{\mathbf{sz}}(\ell)\}_{-d}^{*} & -N+1 \le d \le -1 \\ \mathbf{0} & else \end{cases} \tag{6.23}$$

where $\mathbf{U}_d$ is the $D-d \times D+N-1$ convolution matrix

$$\mathbf{U}_d = \begin{bmatrix} u_{N-d,d} & \cdots & u_{1,d} & 0 & \cdots & 0 \\ 0 & u_{N-d,d} & & u_{1,d} & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & u_{N-d,d} & \cdots & u_{1,d} \end{bmatrix}, \tag{6.24}$$

$$u_{i.k} = \left[\mathbf{R}_{\gamma,\mathbf{z}}\right]_{i,i+k} \tag{6.25}$$

correspond to the diagonals of $\mathbf{R}_{\gamma,\mathbf{z}}$, and

$$\mathbf{p}_{\mathbf{x}}(\ell) = diag\{\mathbf{P}(\ell)\}_0$$

is a $D+N-1 \times 1$ vector corresponding to the main-diagonal of the source-correlation matrix.

The diagonals of the "full" representation can likewise be calculated via

$$diag\left\{\mathbf{R}_{\mathbf{Vz}}\right\}_d = \begin{cases} \mathbf{U}_d \mathbf{p}_{\mathbf{x}} & 0 \leq d \leq N-1 \\ diag\left\{\mathbf{R}_{\mathbf{Vz}}\right\}^*_{-d} & -N+1 \leq d \leq -1 \\ \mathbf{0} & else \end{cases} \tag{6.26}$$

where $\mathbf{U}_d$ is now the $L-d \times L-N+1$ convolution matrix

$$\mathbf{U}_d = \begin{bmatrix} u_{1,d} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ u_{N-d,d} & & \ddots & 0 \\ 0 & \ddots & & u_{1,d} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & u_{N-d,d} \end{bmatrix} \tag{6.27}$$

and

$$\mathbf{p_x} = \hat{\mathbf{x}} \odot \hat{\mathbf{x}}^*.$$

is the power estimate for the entire $L+N-1$ range interval. As $\mathbf{p}(\ell)$ will correspond to a shifted segment of the "full" $\mathbf{p}$, $\mathbf{R_{sz}}(\ell)$ can be constructed using portions of the "full" structured covariance diagonals via (6.20). Thus, the aggregate cost constructing all structured covariance be $\mathscr{O}((N-1)(L)log(L))$ where the "$N-1$" stems from the hermitian symmetry of $\mathbf{R}_{\gamma,\mathbf{z}}$ and noting that the convolution matrix for the $(N-1)$th diagonal itself is a diagonal-matrix and does not require frequency-domain convolution.

6.3 Reduced Complexity Full APC Implementations

While the "full" APC structured covariance can leverage efficient banded Cholesky decomposition

$$\mathbf{R_{uu}} = \mathbf{L_{uu}}\mathbf{L_{uu}}^H \tag{6.28}$$

for $\mathscr{O}(N^2L)$ complexity, solving for the bank of RMMSE filters

$$\mathbf{L_{uu}}\mathbf{L_{uu}}^H\mathbf{W} = \mathbf{VP} \tag{6.29}$$

via forward-backward substitution with $\mathbf{L_{uu}}$ will be $\sim\mathscr{O}(NL^2)$[2] and can be computationally prohibitive. Furthermore, incorporating beamspoiling and a unity constraint further increases complexity by $\sim\mathscr{O}(kNL^2)$ and $\sim\mathscr{O}(L^2)$ respectively. For the unconstrained case (e.g., no unity con-

[2]Here we note that this order of complexity is only approximately $\mathscr{O}(NL^2)$ as $\mathbf{V}$ has many zero elements that can be ignored during the forward/backward substitutions. These zeros may also be leveraged to modestly improve efficiency of the unity constraint and beamspoiling operations.

straint) with no beamspoiling, (6.16) can alternatively be expressed as

$$\begin{aligned}\hat{\mathbf{x}} &= \mathbf{W}^H\mathbf{u} \\ &= \mathbf{P}\mathbf{V}^H\mathbf{R}_{\mathbf{uu}}^{-1}\mathbf{y} \\ &= \mathbf{P}\mathbf{V}^H\tilde{\mathbf{y}}\end{aligned} \tag{6.30}$$

which only requires a single linear system solve

$$\mathbf{L}_{\mathbf{uu}}\mathbf{L}_{\mathbf{uu}}^H\tilde{\mathbf{y}} = \mathbf{y} \tag{6.31}$$

for $\mathscr{O}(NL)$ and the remaining matrix multiply with $\mathbf{P}\mathbf{V}^H$ can be computed efficiently using frequency domain convolution and elementwise-multiplication. However, beamspoiling and unity-constraints are generally necessary for a robust estimate and can not be easily incorporated in this formulation. Thus, to retain efficiency we must take an alternative approach.

6.3.1 Reduced Complexity Template-Based Approach

First, the template-based RMMSE implementation is assessed as it does not explicitly require a unity constraint or beamspoiling updates. For an arbitrary delay invariant $2Z+1\times 1$ desired response vector

$$\mathbf{g} = \begin{bmatrix} g_{-Z} & g_{-Z+1} & \cdots & g_{-1} & g_0 & g_1 & \cdots g_Z \end{bmatrix}^T \tag{6.32}$$

with

$$g_i = g_{-i}^*,$$

the corresponding RMMSE filters will have the form

$$\mathbf{W} = \mathbf{R}_{\mathbf{uu}}^{-1}\mathbf{V}\mathbf{P}\mathbf{G} \tag{6.33}$$

where

$$\mathbf{G} = \begin{bmatrix} g_0 & g_{-1} & \cdots & g_{-Z} & 0 & \cdots & 0 \\ g_1 & g_0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & & \ddots & 0 \\ g_Z & & \ddots & \ddots & \ddots & & g_{-Z} \\ 0 & \ddots & & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & \ddots & g_{-1} \\ 0 & \cdots & 0 & g_Z & \cdots & g_1 & g_0 \end{bmatrix} \tag{6.34}$$

is a $(L-N+1)\times(L-N+1)$ Hermitian-Toeplitz matrix. Applying (6.33) to the measurement vector is then expressed as

$$\begin{aligned} \hat{\mathbf{x}} &= \mathbf{W}^H \mathbf{y} \\ &= \mathbf{G}^H \mathbf{P} \mathbf{V}^H \mathbf{R}_{\mathbf{uu}}^{-1} \mathbf{y} \\ &= \mathbf{G}^H \mathbf{P} \mathbf{V}^H \tilde{\mathbf{y}} \end{aligned} \tag{6.35}$$

which again only requires a single linear system solve

$$\mathbf{R}_{\mathbf{uu}} \tilde{\mathbf{y}} = \mathbf{y}. \tag{6.36}$$

Leveraging the hermitian banded structure of $\mathbf{R}_{\mathbf{uu}}$, (6.36) can be solved with $\mathscr{O}(N^2L+NL)$ complexity. The subsequent multiplication with $\mathbf{G}^H\mathbf{P}\mathbf{S}^H$ can then be performed efficiently using frequency domain convolution and elementwise-multiplications for $\mathscr{O}(L\log(L)+(L-N+1)+(L+Z)\log(L+Z))$. Including the efficient formulation of the structured covariance, "full" template-based RMMSE will have a complexity of

$$\mathscr{O}((N-1)L\log(L)+N^2L+NL+L\log(L)+(L-N+1)+(L+Z)\log(L+Z))$$

per iteration which no longer scales with L^2 or N^3.

6.3.2 Outlier-Based Low-Rank Approximate Implementation

As the APC initialization is typically formed using the matched filter, the first few iterations are virtually guaranteed to incorporate sidelobe contributions into the source correlation estimate. Additionally, even after convergence, the scatterer estimates whose sidelobe contributions would be below the system noise floor will likewise be incorporated. Thus, APC inherently performs several redundant operations as suppressing "sidelobes" from sidelobes (e.g., suppressing things that aren't really there) or from low-powered scattering should not provide an improved estimate. Thus, treating $\mathbf{P}$ as a full rank matrix does not inherently improve the efficacy or convergence of APC.

Alternatively, a pseudo-APC iteration may be posed by excising all redundant elements from $\mathbf{P}$ and only retaining the signatures that are both high power and not likely to be sidelobes. The resulting APC structured covariance matrix can then be represented as diagonal plus low-rank, enabling efficient implementations via Woodbury's identity [141] or low-rank/sparse factorization updates [121]. Selection of what components to suppress requires an additional "detection" stage, resulting APC framework implementation is now akin to matching pursuit [142] or CLEAN [106, 143, 144], though, the use of a filter-based objective and statistical model uncertainty does differentiate these philosophies.

While there are many different approaches one could take to select the non-zero source directions in range or range-Doppler space, here we examine a inexpensive ordered statistics approach [145] in range as a proof of concept, leaving assessment and analysis of alternative or improved detection routines for future work. The proposed outlier selection routine requires two stages. First a length N moving median is applied to the current estimate source powers as

$$\mathbf{p}_{med} = moveMed\{|\hat{\mathbf{x}}|^2\} \tag{6.37}$$

where $\mathbf{p}_{med}$ provides an ordered statistical estimate of the noise plus sidelobe floor for each range bin. This moving vector then is used to determine which indices correspond to outliers above the

noise plus median floor via a simple threshold check of

$$\Omega = i \in \left\{ \gamma \mathbf{p}_{med,i} < \mathbf{p}_i; \quad \gamma \frac{\sigma_{\mathbf{n}}^2}{\mathbf{s}^H \mathbf{s}} < \mathbf{p}_i \right\} \tag{6.38}$$

where Ω is a length r_Ω range of outlier directions where $\gamma \sim 20$ used to incorporation of "false-alarms" into the structured covariance. This range of directions is then used to form the low-rank approximate source correlation matrix

$$\begin{aligned} \mathbf{R}_{\mathbf{V}} &= \mathbf{V}^H \mathbf{P} \mathbf{V} \\ &\approx \mathbf{R}_{\mathbf{V},\Omega} \\ &\approx \mathbf{V}_\Omega^H \mathbf{P}_\Omega \mathbf{V}_\Omega \end{aligned} \tag{6.39}$$

where $\mathbf{V}_\Omega$ and $\mathbf{P}_\Omega$ are $L \times r$ and $r_\Omega \times r_\Omega$ matrices corresponding to the outlier elements of $\mathbf{V}$ and $\mathbf{P}$ respectively. Note, $\mathbf{V}_\Omega$ is likewise sparse as each column of $\mathbf{V}$ will have only N non-zero elements. To incorporate straddling uncertainty, one must simply form (6.39) via convolution with a sparse $\mathbf{p}_\mathbf{x}$ using (6.26) or leverage the low-rank straddling approximations from Section 3.1.1 as

$$\sum_{a=1}^{A} \mathbf{V}_{a,\Omega}^H \mathbf{P}_\Omega \mathbf{V}_{a,\Omega}$$

where $\mathbf{V}_{a,\Omega}$ is from the A dominant eigenvectors and eigenvalues of (3.15).

The resulting approximate structured covariance

$$\mathbf{R}_{\mathbf{uu},\Omega} = \mathbf{R}_{\mathbf{V},\Omega} + \mathbf{R}_{\mathbf{z},\Omega} + \mathbf{R}_{\mathbf{n}} \tag{6.40}$$

can then be used to form the bank of APC filters efficiently, though, the order of complexity is dependent on straddle-compensation methodology, number of outliers, how close the outliers are spaced, and the solver used. For Cholesky decomposition the cost will generally be much less than $\mathcal{O}(N^2 L)$, but the exact order is a function of the outlier spacing and does not have a closed form solution. When using Woodbury's the order of complexity again does not follow a closed form

solution, but will largely be a function of $\sim\mathscr{O}(r_\Omega^3 L)$. The sparsity of $\mathbf{V}_\Omega$ and $\mathbf{V}$ offers potential for further complexity reduction but will not be assessed here for brevity. Both the Cholesky and Woodbury methodologies are likewise amenable to inexpensive beamspoiling inverse downdates, though it is no longer necessary to apply for every direction since $\mathbf{P}$ is sparse.

6.4 Robust and Efficient APC Experimental Evaluation

We now seek to apply these efficient NIMPC implementation to the open-air experimental PROFM CPIs used in Sections 3.4, 4.6, and 5.3. Here four different APC implementations will be applied to the open-air measurements to evaluate the efficacy of the proposed approaches. The first case will leverage the traditional beamspoiling, model uncertainty, and a unity constraint using the full model. Here the model uncertainty was estimated using by measuring the variance across a repeated CPI collected in loopback. The second test case leverages the same beamspoiling, model uncertainty, and unity constraint, but the interpolation based straddle compensation from Section 6.1 was also incorporated. For the third test case, the template-based APC approach from Section 6.3.1 was used, where the same desired response from the template-based LS MMFs in Section 3.4 was used. The final test case also leveraged beamspoiling, straddle uncertainty, model uncertainty, and a unity constraint, however, here the structured covariance was approximated using the outlier/detection methodology proposed in Section 6.3.2, where the threshold parameter was set to $\gamma = 10$. The corresponding parameters and dimensionality used for this experiment are provided in Table 6.1. Here, each of the APC approaches were applied to 3 iterations and (when it was applied) the beamspoiling was configured to match the nominal MF resolution. This table likewise includes the average number of outliers (across pulses) for each iteration when using the low-rank approximate implementations.

Next, Table 6.2 illustrates the order of complexity for constructing the full structured covariance using a direct (sparse matrix multiplication), convolution, or only outliers for the final APC iteration. Note, here the outlier construction was performed via matrix multiplication, but it can likewise be formed via convolution if it would provide a lower cost. Here the convolution-based

Measurement Parameters		
Parameter	Variable	Value
Sample Rate	f_s	50MHz
3dB Bandwidth	B	33.3MHz
Pulse Duration	T	4.5μseconds
Samples in Waveform	N	276
Number of Pulses	M	1000
Fast-time Samples processed	L	2500
Fast-Time Filter Samples	D	N
Number of APC Iterations	I	3
Avg Number of Outliers Iter 1	r_{out}	38
Avg Number of Outliers Iter 2	r_{out}	92
Avg Number of Outliers Iter 3	r_{out}	112

Table 6.1: Table of Open-Air Experimental Parameters for PROFM CPI of 1000 waveforms

Construction order of Complexity	
Brute Force	1.90E+8
Convolution	7.76E+6
Outlier (Final Iter)[a]	5.87E+6

Table 6.2: Complexity for constructing the Structured Covariance Matrix

[a]Can also be formed via convolution method, and in general does not explicitly need to be constructed.

approach sees a little more than an order of magnitude in cost.

The computational complexity for the Cholesky decomposition of the full and approximate outlier structured covariance matrices are then provided in Table 6.3. As a point of comparison, the aggregate computational complexity for the decompositions of each delay with the traditional snapshot APC formulation is likewise provided. Here, requiring only a single matrix decomposition with the full formulation sees an improvement by over two order of magnitude over the snapshot case as the banded structure enables an efficient decomposition (illustrated in Figure 6.10). Additionally, an approximate order of complexity based on the number of non-zeros for the outlier case (the final iteration has the most outliers) is formed and sees another improvement of two orders of magnitude. Note, this complexity roughly acts as an upper bound and parallelization techniques may further improve this cost. To visualize this complexity reduction, Figures 6.11-6.13 illustrate the first pulse's detected outliers (left) and corresponding structured covariance

Order of Complexity for Cholesky Decompositions	
Snapshot $D = N$	5.26E+10
Full	1.90E+08
Outlier (Final Iter)	2.83E+06

Table 6.3: Complexity for Cholesky Decompositions

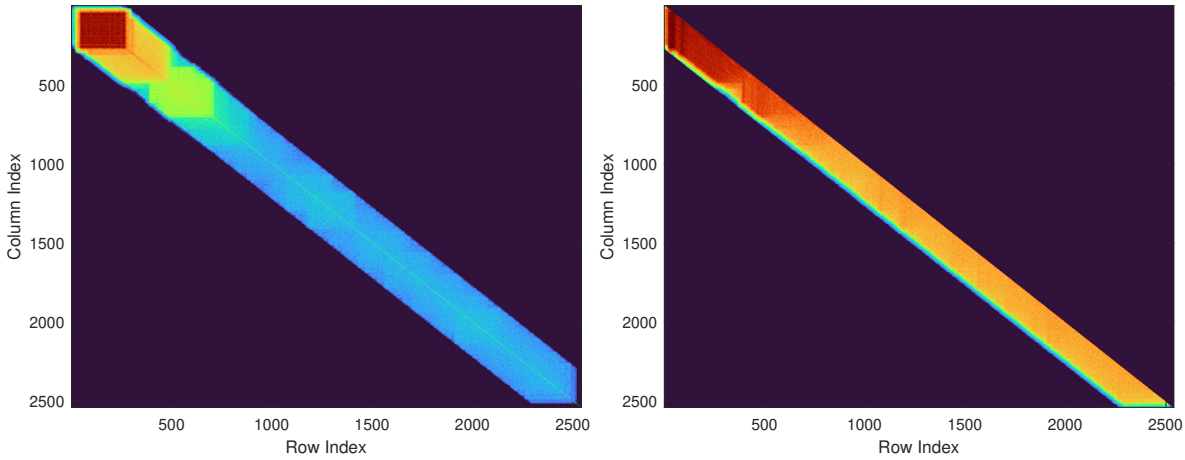

Figure 6.10: Full Structured Covariance (left) and it's Lower Triangular Cholesky Decomposition (right)

matrix for each iteration of APC.

The order of complexity for applying performing an APC iteration for each of the respective approaches is then listed in Table 6.4. In aggregate, use of the full model theoretically improves complexity by an order of magnitude over the snapshot approach, though alternative approximate implementations can potentially offer a reduction in complexity for either approach [85, 146–148]. When solving only a single linear system via the template-based approach in Section 6.3.1, the complexity is reduced by an order of magnitude. For the outlier approach, a Woodbury implementation sees a reduction in cost by another order of magnitude, but may run into stability concerns when using single precision (as opposed to double). That being said, for this experimental assessment, use of Woodbury's identity did not create any numerical precision issues. The Cholesky approaches see a slight increase in theoretical complexity but will generally be more stable. Again, here we do not seek to provide an exact "flop" count for each respective approach, but merely illustrate a sense of scale.

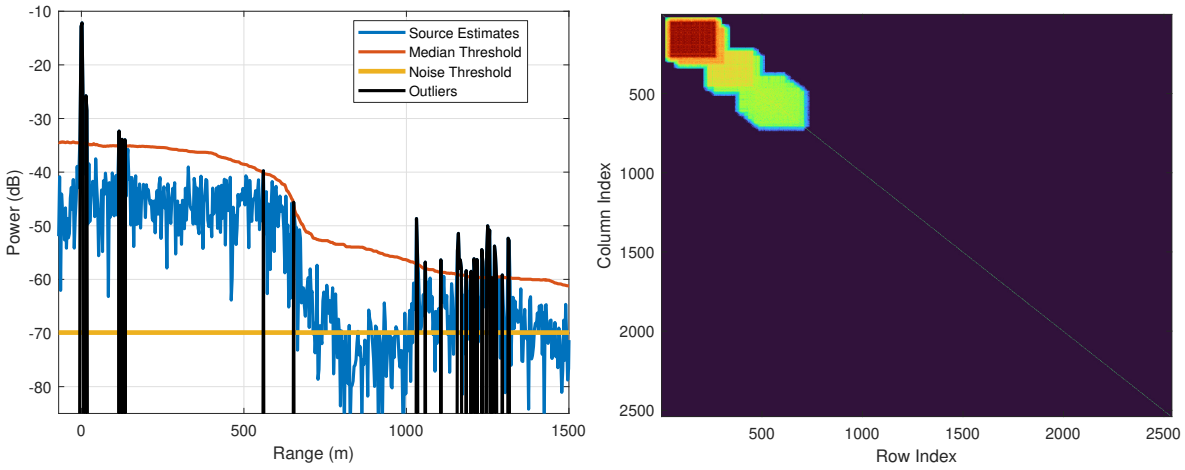

Figure 6.11: Detected outliers (left) and Structured Covariance (right) for first iteration

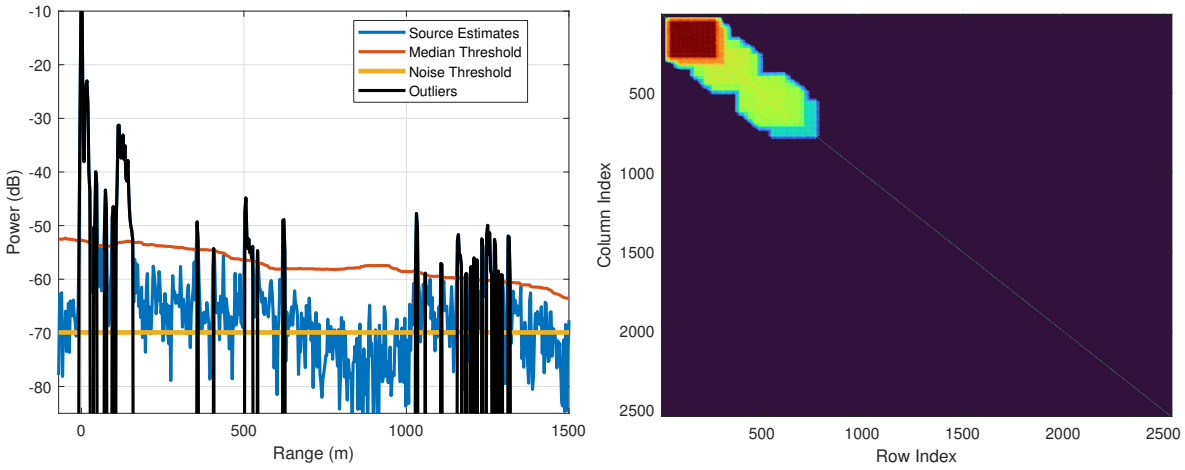

Figure 6.12: Detected outliers (left) and Structured Covariance (right) for second iteration

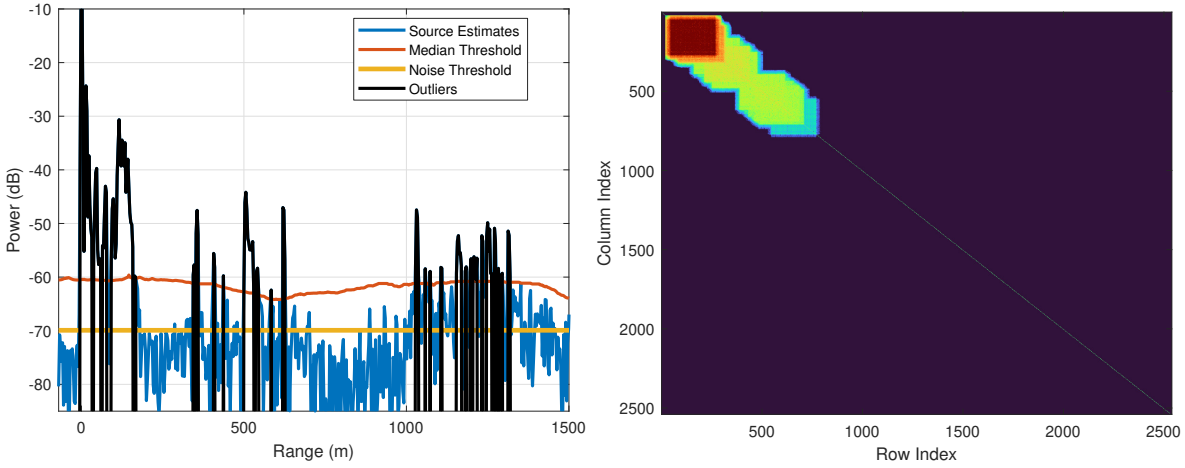

Figure 6.13: Detected outliers (left) and Structured Covariance (right) for third/final iteration

Order of Complexity	
Snapshot $D = N$ w Conv Construnction	5.28E+10
Full with Unity Constraint	1.92E+09
Full Template-Based Approach	1.99E+08
Outlier Approach (Woodbury's)	1.52E+07
Outlier Approach (Cholesky)	3.43E+07

Table 6.4: Average Complexity for single APC Iteration

Next, the RMS range estimates for each approach are illustrated in the left side of Figure 6.14 while the right side corresponds to the expected noise power of the output filters. Relative to the MF, each of the APC approaches sees a significant reduction in sidelobe level and does not yield the extended range-response visible with the LS MMF approaches. In terms of the noise power, the approaches that leverage a unity constraint typically see 0.5dB or less MML, though there is a spiking behavior corresponding to strong scatters. For the template based approach, the noise gain is generally higher ($\sim$3dB) but sees a fairly significant noise gain when the estimates are below the noise floor or the scattering power is high. A change in template may help mitigate this effect, but template "tuning" is currently a heuristic process. Focusing on the sub-plots corresponding to the direct path, as expected the template-based APC approach sees a fairly broad roll-off initially commensurate with the template-based LS MMF. The APC approach that does not have straddling compensation likewise sees a broadened response due to the likely sub-sample offset of the direct path. Both the straddle-compensated exact and approximate outlier implementations are robust to this broadening effect.

Next, the range processed measurements for the respective APC approaches were used to perform Doppler estimation. First, examining the "baseline" implementation with no straddling compensation in Figure 6.15 we see that APC offers RSM suppression on par with MiCRFt and NIMPC, but does not require the reduction in unambiguous Doppler or joint-dimension cancellation. However, as with the LS MMF and MiCRFt ISL cases (without compensation), we see a streaking behavior in Doppler corresponding to dominant sources. This effect is largely due to the same mainlobe modulation induced by the beamspoiling present with the LS MMS, however, the

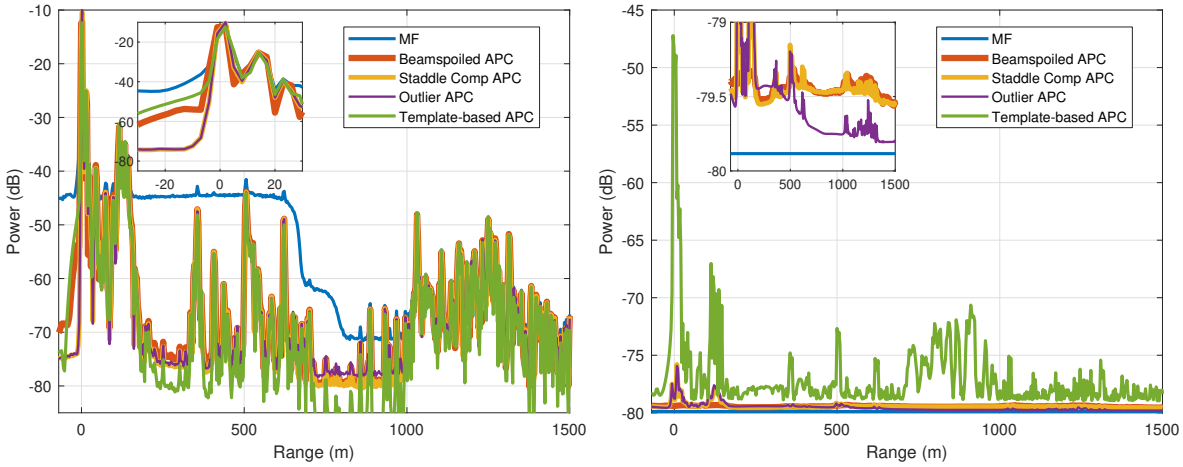

Figure 6.14: RMS range estimates (left) and noise powers (right) for the Full APC implementations

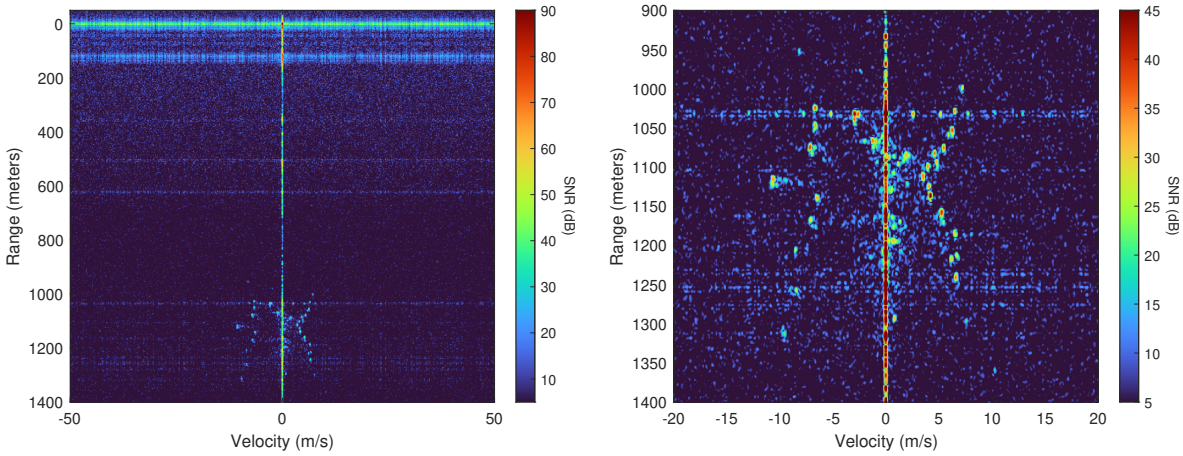

Figure 6.15: Baseline Range-Doppler Estimate for PROFM CPI using APC

range-varying noise power will produce a similar streaking behavior.

Next, the straddle compensated APC range-Doppler estimates are shown in Figure 6.16. Here, the compensation has resulted in a significant reduction in mainlobe modulation and the streaking in Doppler has been significantly reduced. Here, the residual streaks still likely correspond to mainlobe modulation as the range-varying noise power is roughly commensurate with the MF as seen in Figure 6.14.

For the template-based APC shown in Figure 6.17, the clutter and vehicle estimates are commensurate with the approaches that leverage a unity constraint. However, the lack of a constraint induces significant background floor fluctuation and likewise creates a streaking behavior in Doppler

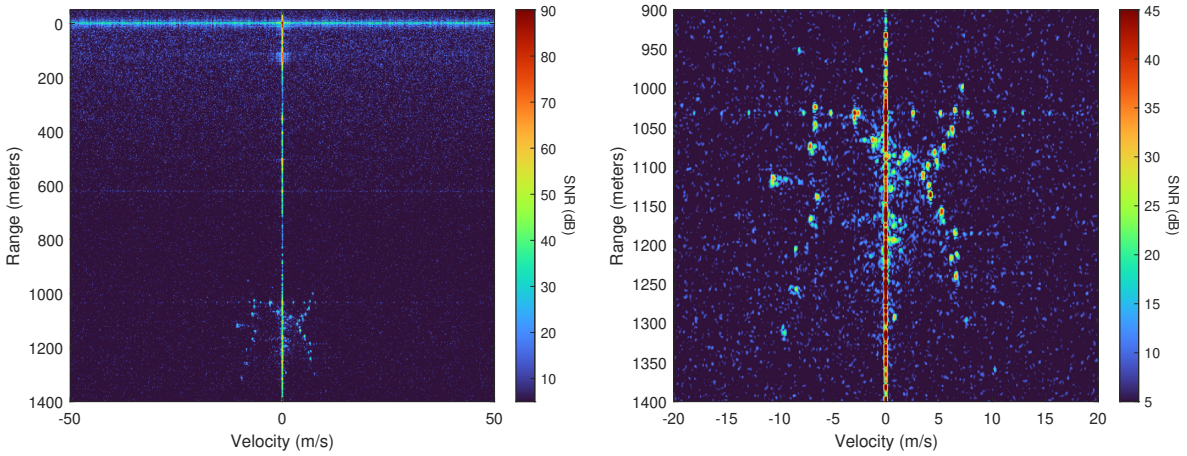

Figure 6.16: Range-Doppler Estimate for PROFM CPI using APC with Straddling Compensation

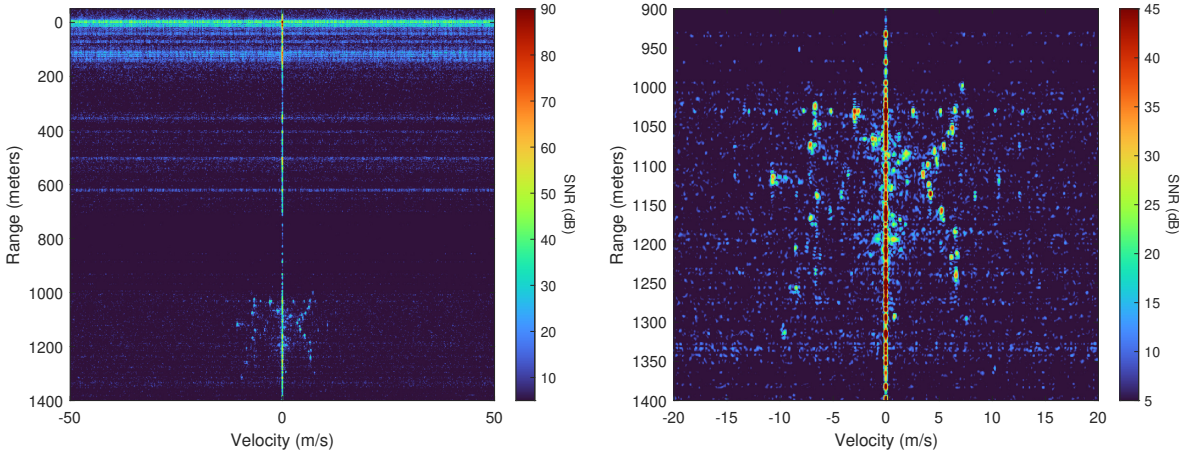

Figure 6.17: Range-Doppler Estimate for PROFM CPI using APC with template-based objective

corresponding to the dominant scatters. Unlike the beamspoiled case, this streaking in Doppler is not due to mainlobe modulation but rather relatively the large noise-gain seen in Figure 6.14. Thus, while the template-based approach offers improved efficiency, it may not be well suited for applications that require a subsequent processing stage (e.g. Doppler estimation after range processing).

Lastly, estimates from the outlier based approximate implementation are shown in Figure 6.18. Here, the final iteration of APC only incorporates 112 outliers (effectively suppressing only $\sim$5% of the possible directions) and yields significant improvement to the RSM background for both the intersection and direct path. However, the approximate covariance was not able to achieve quite the

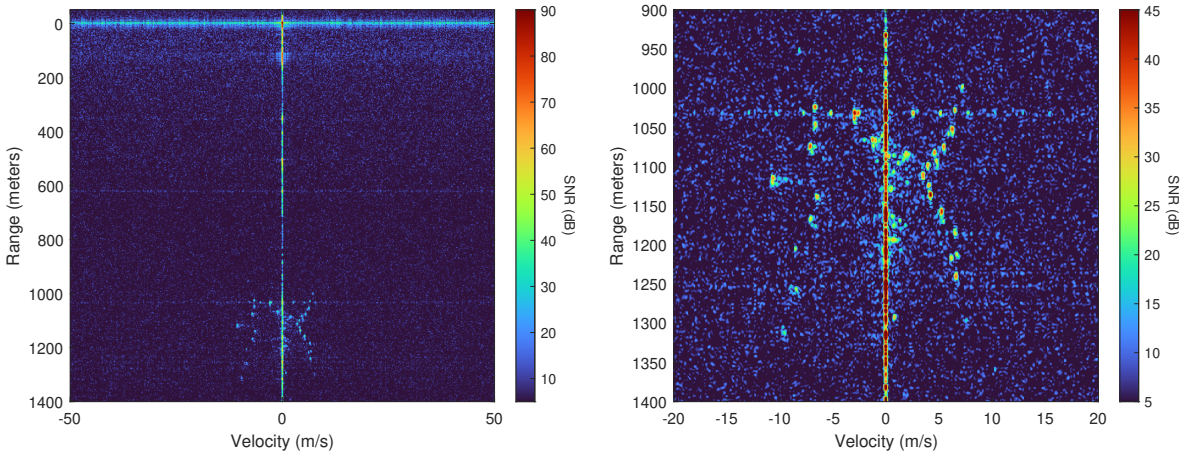

Figure 6.18: Range-Doppler Estimate for PROFM CPI using Approximate Outlier APC with Straddling Compensation

performance of the exact formulation and the intersection RSM levels are 2dB higher. That being said, it achieved this estimate with a two orders of magnitude reduction in complexity and a very minimalistic outlier detection approach. Additionally, since the approach incorporates straddling compensation via a low-rank approximation to $\mathbf{R}_\gamma$, the Doppler streaking is fairly minimal and on par with the exact implementation.

Chapter 7

Generalized SINR Loss Constrained LS MMFs

As mentioned in Section 2.4, the LS MMFing approaches effectively approximate the receive snapshot model as delay invariant and design the filter MMF response to control range-sidelobes. Previous work [72, 74, 75, 78] leveraged this delay invariant form to pose LS MMF QCQPs that directly control MML as opposed to heuristic tuning of the regularization term σ^2. However, these QCQPs were restricted to the sidelobe minimization MMF formulation (i.e., ISL or peak sidelobe level minimization) and white Gaussian noise. While system thermal noise is predominantly white, hardware effects such as spurs from analog-to-digital-converters will introduce colored noise. Additionally, as the current RF environment is congested with commercial communication systems, modern radars often must contend with additive in band interference. Direct solvers to the QCQP also tend to be computationally intensive. While previous work attempted to address this complexity, the proposed techniques required an expensive eigen decomposition or a approximation to the filter can degrade filter efficacy.

Here, generalizations to the QCQP are proposed yield a tractable optimization problem for an SINR constraint that addresses colored noise and in-band interference. Furthermore, these constraints are extended to the template-based objective function. The algorithmic framework to solve these QCQPs via the Lagrange Dual problem is then outlined. Extensions to solver methodologies from [74, 75] for the SINR formulation are then examined and a new computationally efficient solver is introduced.

7.1 Generalized SINR Loss Constraint

Consider the receive signal model comprised of D samples corresponding to a particular point in range,

$$\mathbf{y} = \gamma\mathbf{v} + \mathbf{c} + \mathbf{i} + \mathbf{n}, \tag{7.1}$$

where $\gamma\mathbf{v}$ is the desired scattering response to be estimated, and γ is the complex scattering coefficient. The vectors $\mathbf{c}$, $\mathbf{i}$, and $\mathbf{n}$ represent clutter (signal-dependent scattering from other ranges), external interference, and noise, respectively. Note, here the delay index ℓ has been excluded as the LS MMF formulation inherently assumes the statistics are delay independent.

After application of the filter $\mathbf{w}$ to the signal model $\mathbf{y}$ (i.e., $\mathbf{w}^H\mathbf{y}$), the output SINR can be defined as

$$\overline{SINR}\{\mathbf{w}\} = \frac{|\gamma|^2|\mathbf{w}^H\mathbf{v}|^2}{\mathscr{E}\{|\mathbf{w}^H\mathbf{c} + \mathbf{w}^H\mathbf{i} + \mathbf{w}^H\mathbf{n}|^2\}}. \tag{7.2}$$

Assuming the clutter, interference, and noise are independent and zero-mean, the SINR can be rewritten as

$$\overline{SINR}\{\mathbf{w}\} = \frac{|\gamma|^2|\mathbf{w}^H\mathbf{v}|^2}{\mathbf{w}^H(\mathbf{R}_\mathrm{c} + \mathbf{R}_\mathrm{i} + \mathbf{R}_\mathrm{n})\mathbf{w}}, \tag{7.3}$$

where $\mathbf{R}_\mathrm{c} = \mathscr{E}\{\mathbf{c}\mathbf{c}^H\}$, $\mathbf{R}_\mathrm{i} = \mathscr{E}\{\mathbf{i}\mathbf{i}^H\}$, and $\mathbf{R}_\mathrm{n} = \mathscr{E}\{\mathbf{n}\mathbf{n}^H\}$, are the $D \times D$ clutter, interference, and noise covariances, respectively. In theory, if all covariance structures are known, the optimal SINR filter can be determined by maximizing this quantity with respect to $\mathbf{w}$. If the covariance matrices are assumed constant with range, then the max SINR filter for each range bin are identical and can be efficiently implemented using a convolution.

Define the interference-plus-noise covariance as

$$\mathbf{R} = \mathbf{R}_\mathrm{i} + \mathbf{R}_\mathrm{n}. \tag{7.4}$$

Here, it is assumed that $\mathbf{i}$ and $\mathbf{n}$ can be treated as stationary random processes resulting in a covariance with Hermitian Toeplitz structure for $\mathbf{R}$ and therefore does not alter with range. Within the

context of radar-communications cohabitation, this stationarity assumption is made as communication systems typically operate on longer timescales than radar. Furthermore, it is also assumed that the noise covariance $\mathbf{R}_{\mathrm{n}}$ is full-rank while the interference covariance $\mathbf{R}_{\mathrm{i}}$ has rank $r \ll L$ which can be incorporated efficiently via low-rank updates (see Appendix A.1). This specific definition is used to differentiate spectrally narrowband interference (low-rank) from "broadband" interference (not easily represented as low-rank) which would be subsumed in $\mathbf{R}_{\mathrm{n}}$. Here, the covariance $\mathbf{R}$ is assumed known or estimated via the collection of "quiet" (TX not active) measurements and can be utilized for the design of $\mathbf{w}$.

The clutter energy component in the denominator of (7.3)

$$\mathbf{w}^H \mathbf{R}_{\mathrm{c}} \mathbf{w} = \mathbf{w}^H \left(\mathbf{S}\mathbf{P}\mathbf{S}^H - \gamma^2 \mathbf{v}\mathbf{v}^H \right) \mathbf{w}$$

represents range-sidelobe energy from scatterers at offset ranges coupling into our range of interest via the filter $\mathbf{w}$. While clutter range-sidelobe energy can be treated as an interference source in this SINR formulation, in general, the relative weighting of the "interference" sidelobes is a function of range based on the delay profile of the illuminated scene (e.g. more akin to the MMSE form of (2.55)). Therefore, the resulting clutter covariance would likewise vary with range and would not be Toeplitz in structure, breaking the range invariance of max SINR filter design. Instead, we can address the clutter power by directly minimizing its contribution via the MMF design methods discussed prior (i.e., reducing range sidelobes). However, it is still desired that the range sidelobes are addressed within the context of the interference and noise environments.

"Noise gain" or MML are terms used to describe the loss SNR of mismatched filtering due to deviation from the matched filter. Previous works have focused on limiting white noise MML by placing constraints on loss in SNR relative to the MF (optimal for WGN) [74, 75]. Here, this framework is generalized to both interference and colored noise by constraining the SINR loss given a known interference-plus-noise matrix $\mathbf{R}$. The filter SINR loss is derived solely on the interference and noise contributions, while the clutter (i.e., range sidelobe contribution) is

minimized via the objective function akin to [71].

The reformulated SINR at the output of the filter $\mathbf{w}$ (considering only interference and noise) is

$$SINR\{\mathbf{w}\} = \frac{|\gamma|^2 \left|\mathbf{w}^H \mathbf{v}\right|^2}{\mathbf{w}^H \mathbf{R} \mathbf{w}}. \tag{7.5}$$

The filter that maximizes this quantity is well-known [84] and can be computed as

$$\mathbf{w}_{\max} \propto \mathbf{R}^{-1} \mathbf{v}, \tag{7.6}$$

yielding a maximum SINR of

$$SINR_{\max} = |\gamma|^2 \mathbf{v}^H \mathbf{R}^{-1} \mathbf{v}. \tag{7.7}$$

Therefore, the SINR loss relative to $SINR_{\max}$ is defined as[1]

$$L_{\text{SINR}}\{\mathbf{w}\} = \frac{SINR\{\mathbf{w}\}}{SINR_{\max}} = \frac{|\mathbf{w}^H \mathbf{v}|^2}{(\mathbf{v}^H \mathbf{R}^{-1} \mathbf{v})(\mathbf{w}^H \mathbf{R} \mathbf{w})} \tag{7.8}$$

which is bounded as $0 \leq L_{\text{SINR}}\{\mathbf{w}\} \leq 1$. Note, for a white-noise-only case (i.e., $\mathbf{R}$ is a scaled identity matrix), the max SINR filter is simply the matched filter and the SINR loss becomes the SNR loss formulation from [74, 75].

In general, the max SINR filter from (7.6) may not have desirable range sidelobes depending on the nature of the waveform modulation and interference structure. However, by allowing the MMF to deviate from (7.6) (thus necessitating a loss in SINR), we can improve the sidelobe performance of $\mathbf{w}$. Furthermore, by bounding the SINR loss by a user-specified value ρ, we can control how much $\mathbf{w}$ deviates from the max SINR form. Therefore, to meet this constraint the QC MMF must satisfy

[1] Note that this definition of loss is technically a gain since it is less than unity; however, this definition of loss is commonplace in the literature.

$$\frac{|\mathbf{w}^H\mathbf{v}|^2}{(\mathbf{v}^H\mathbf{R}^{-1}\mathbf{v})(\mathbf{w}^H\mathbf{R}\mathbf{w})} \geq \rho \tag{7.9}$$

where $0 \leq \rho \leq 1$ is the degree of acceptable loss relative to the maximum SINR.

Since the solution of $\mathbf{w} = \mathbf{R}^{-1}\mathbf{v}$ yields a value of unity (corresponding to 0 dB of SINR loss), the QC will always have a feasible solution for any $0 \leq \rho \leq 1$ and will guarantee a tractable QCQP. While perhaps less practical, for the sake of generality, loss constraints relative to maximum SNR, such as

$$\frac{SINR\{\mathbf{w}\}}{SNR_{\max}} = \frac{|\mathbf{w}^H\mathbf{v}|^2}{\left(\mathbf{v}^H\mathbf{R}_{\mathrm{n}}^{-1}\mathbf{v}\right)(\mathbf{w}^H\mathbf{R}\mathbf{w})},$$

or any such variation can likewise be formulated. However, such constraints are not guaranteed to be feasible for all $\rho > 0$ and, as such, the primary focus of this work utilizes the $L_{\mathrm{SINR}}\{\mathbf{w}\}$ formulation from (7.8).

After trivial algebraic manipulation, (7.9) can be expressed in standard QC form as

$$\mathbf{w}^H\left(\mathbf{B} - \mathbf{v}\mathbf{v}^H\right)\mathbf{w} \leq 0, \tag{7.10}$$

for

$$\mathbf{B} = \rho\left(\mathbf{v}^H\mathbf{R}^{-1}\mathbf{v}\right)\mathbf{R}. \tag{7.11}$$

With the template-based MMF formulation, this quadratic constraint is necessarily non-convex as the matrix $\mathbf{B} - \mathbf{v}\mathbf{v}^H$ will always have a single non-positive eigenvalue. Therefore, adding this SINR loss constraint to the template-based MMF problem results in a non-convex problem statement. Despite this non-convexity, certain QCQPs with indefinite quadratic constraints are tractable [80] and the globally optimal MMF can still be found.

For the direct ISL minimization problem, all feasible solutions must satisfy the equality con-

straint $\mathbf{w}^H\mathbf{v} = 1$, thus (7.10) can be simplified as

$$\mathbf{w}^H\mathbf{B}\mathbf{w} - 1 \leq 0. \tag{7.12}$$

Since $\mathbf{B}$ is PD, (7.12) is a quadratic constraint. Therefore, the addition of the SINR loss constraint does not alter the convexity of the direct ISL minimization problem statement.

7.2 Generalized Form of the QCQP

As both the ISL and template constrained objectives may be expressed as a standard QCQP, here the ISL objective is expressed as

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H\tilde{\mathbf{T}}\mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H\mathbf{C}\mathbf{w} + a \leq 0 \\ & \mathbf{w}^H\mathbf{v} = 1 \end{aligned} \tag{7.13}$$

and template objective as

$$\begin{aligned} \min_{\mathbf{w}^*} \quad & \mathbf{w}^H\mathbf{T}\mathbf{w} - 2\,\mathrm{Re}\left\{\mathbf{w}^H\mathbf{b}\right\} \\ \text{s.t.} \quad & \mathbf{w}^H\mathbf{C}\mathbf{w} \leq 0 \end{aligned} \tag{7.14}$$

such that both Lagrangians may be represented in a notationally consistent manner as

$$\mathscr{L}\{\mathbf{z};\lambda\} = \mathbf{z}^H\left(\mathbf{D} + \lambda\mathbf{Q}\right)\mathbf{z} - 2\,\mathrm{Re}\{\mathbf{z}^H\mathbf{b}\} + \lambda a \tag{7.15}$$

where

$$\mathbf{D} = \begin{bmatrix} \tilde{\mathbf{T}} & \mathbf{v} \\ \mathbf{v}^H & 0 \end{bmatrix}, \quad \mathbf{z} = \begin{bmatrix} \mathbf{w} \\ \mu \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix},$$

$$a = -1, \quad \mathbf{Q} = \begin{bmatrix} \mathbf{C} & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix}, \quad \mathbf{C} = \mathbf{B}$$

for ISL and

$$\mathbf{D} = \mathbf{T}, \quad \mathbf{z} = \mathbf{w}, \quad \mathbf{b} = \mathbf{S}^H \mathbf{p},$$

$$a = 0, \quad \mathbf{Q} = \mathbf{C}, \quad \mathbf{C} = \mathbf{B} - \mathbf{v}\mathbf{v}^H$$

for template[2].

Again, note that for the template formulation $\mathbf{C}$ has a single negative eigenvalue by definition. While this constraint makes the problem non-convex, it has been shown the globally optimal solution can be determined for non-PSD QCQP with a single QC as long as the Hessian of the Lagrangian remains PD [108, p. 653]

$$\mathbf{T} + \lambda \mathbf{C} \succeq 0, \tag{7.16}$$

and Slater's condition holds (i.e., there exist a $\mathbf{w}$ such that the constraints are strict), which can be easily proved when

$$\mathbf{w} = \frac{\mathbf{R}^{-1}\mathbf{v}}{\mathbf{v}^H \mathbf{R}^{-1}\mathbf{v}}.$$

Therefore, strong duality will hold for both (7.13) and (7.14) such that the optimal Lagrange multipliers $\lambda_\star$ (and $\mu_\star$) can be determined via the Lagrange dual problem.

7.3 The Lagrange Dual Problem

The dual function $G\{\lambda\}$ is found via an infimum of the Lagrangian with respect to $\mathbf{z}$ as

$$G\{\lambda\} = \inf_{\mathbf{z} \in \mathscr{D}} \mathscr{L}\{\mathbf{z}, \lambda\}, \tag{7.17}$$

where $\mathscr{D}$ is the domain of $\mathbf{z}$ where the objective function exists and the quadratic constraint is satisfied. Taking the derivative of $\mathscr{L}\{\mathbf{z}, \lambda\}$ w.r.t. $\mathbf{z}^*$, setting equal to zero, and solving for $\mathbf{z}$ reveals

[2]Note, while here the derivations focus on base LS MMF notation from Section 2.2, the straddle compensation approaches from in Section 3.1 can likewise be introduced without loss of generality.

the form of the optimal mismatch filter as

$$\mathbf{z}_{opt}\{\lambda\} = \mathbf{G}^{-1}\{\lambda\}\mathbf{b} \tag{7.18}$$

where

$$\mathbf{G}\{\lambda\} = \mathbf{D} + \lambda\mathbf{Q}. \tag{7.19}$$

Thus, maximization of the dual function

$$G\{\lambda\} = \begin{cases} -\mathbf{b}^H\mathbf{G}^{-1}\{\lambda\}\mathbf{b} + a\lambda & \lambda \in (\lambda^-, \lambda^+) \\ -\infty & otherwise \end{cases} \tag{7.20}$$

will yield $\mathbf{z}_\star$ and is strictly concave in the valid range of $\lambda \in (\lambda^-, \lambda^+)$.

7.3.1 Bounds for the Dual

By definition, the Lagrange multiplier inherently has a lower bound, $\lambda^- = 0 \leq \lambda$ and the upper bound (denoted λ^+) is determined according to the PSD requirement ($\tilde{\mathbf{T}} + \lambda\mathbf{C} \succeq 0$ or $\mathbf{T} + \lambda\mathbf{C} \succeq 0$). Therefore, the valid regions of λ for the ISL objective is the domain of all real numbers (since $\mathbf{B}$ is PD and $\tilde{\mathbf{T}}$ is PSD), while the valid regions for the template correspond to when

$$\mathbf{T} + \lambda\mathbf{C} = \mathbf{T} + \lambda\left(\mathbf{B} - \mathbf{v}\mathbf{v}^H\right) \succeq 0\,. \tag{7.21}$$

As $\left(\mathbf{B} - \mathbf{v}\mathbf{v}^H\right)$ has only a single negative eigenvalue and $\mathbf{T}$ is PSD, λ^+ corresponds to the only positive λ such that

$$\det\{\mathbf{T} + \lambda\left(\mathbf{B} - \mathbf{v}\mathbf{v}^H\right)\} = 0 \tag{7.22}$$

where $\det\{\cdot\}$ takes the determinant of the matrix. (7.22) may be expressed as the generalized eigenvalue problem [149]

$$\mathbf{T}\phi_i = \mu_i\left(\mathbf{B} - \mathbf{v}\mathbf{v}^H\right)\phi_i \tag{7.23}$$

where ϕ_i and μ_i are the ith generalized eigenvector and eigenvalue respectively. Since $\left(\mathbf{B}-\mathbf{v}\mathbf{v}^H\right)$ is nonsingular for any $0<\rho<1$, (7.23) can alternatively be expressed as the standard eigenvalue problem

$$\left(\mathbf{B}-\mathbf{v}\mathbf{v}^H\right)^{-1}\mathbf{T}\phi_i=\mu_i\phi_i \tag{7.24}$$

and λ^+ will correspond to the only negative eigenvalue of $\left(\mathbf{B}-\mathbf{v}\mathbf{v}^H\right)^{-1}\mathbf{T}$. However, in practice solving for λ^+ directly is computationally expensive and is not strictly necessary for maximizing the dual since any $\hat{\lambda}^+$ that satisfies $\lambda_\star\leq\hat{\lambda}^+\leq\lambda^+$ may be treated as an effective upper bound due to the concave behavior in the valid regions of λ.

7.3.2 Derivatives of the Dual Function

The first derivative of (7.20) in the feasible region can be found via chain rule as

$$\begin{aligned}G^{'}\{\lambda\}&=\frac{\delta}{\delta\lambda}\left(-\mathbf{b}^H\mathbf{G}^{-1}\mathbf{b}+a\lambda\right)\\&=-\mathbf{b}^H\frac{\delta}{\delta\lambda}\left(\mathbf{G}^{-1}\right)\mathbf{b}+a\end{aligned} \tag{7.25}$$

where the functional notation of λ for $\mathbf{G}$ has been dropped for brevity. Next, the derivative of an inverse matrix [141] is used to find

$$\frac{\delta}{\delta\lambda}\left(\mathbf{G}^{-1}\right)=-\mathbf{G}^{-1}\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)\mathbf{G}^{-1} \tag{7.26}$$

where only scalar dependencies on λ are present, thus,

$$\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)=\mathbf{Q}. \tag{7.27}$$

The second derivative is likewise computed so that concave nature of the dual may be utilized to achieve superlinear convergence near the maxima via Newton's method. Through applying chain

rule

$$\begin{aligned} G^{''}\{\lambda\} &= \frac{\delta^2}{\delta\lambda^2}\left(-\mathbf{b}^H\mathbf{G}^{-1}\mathbf{b}+a\lambda\right) \\ &= -\mathbf{b}^H\frac{\delta^2}{\delta\lambda^2}\left(\mathbf{G}^{-1}\right)\mathbf{b} \end{aligned} \tag{7.28}$$

which can be expanded as

$$\begin{aligned} \frac{\delta^2}{\delta\lambda^2}\left(\mathbf{G}^{-1}\right) = 2\mathbf{G}^{-1}\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)\mathbf{G}^{-1}\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)\mathbf{G}^{-1} \\ -\mathbf{G}^{-1}\frac{\delta^2}{\delta\lambda^2}\left(\mathbf{G}\right)\mathbf{G}^{-1} \end{aligned} \tag{7.29}$$

where

$$\frac{\delta^2}{\delta\lambda^2}\left(\mathbf{G}\right) = \mathbf{0}. \tag{7.30}$$

Thus, (7.29) reduces to

$$\frac{\delta^2}{\delta\lambda^2}\left(\mathbf{G}^{-1}\right) = 2\mathbf{G}^{-1}\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)\mathbf{G}^{-1}\frac{\delta}{\delta\lambda}\left(\mathbf{G}\right)\mathbf{G}^{-1} \tag{7.31}$$

which does not require additional matrix inversions beyond what was found with the first derivative.

Evaluating the dual, its derivatives, and MMF yields

$$\begin{aligned} \mathbf{w}\{\lambda\} &= \frac{\left(\tilde{\mathbf{T}}+\lambda\mathbf{B}\right)^{-1}\mathbf{v}}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{B}\right)^{-1}\mathbf{v}} \\ G\{\lambda\} &= \frac{1}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{B}\right)^{-1}\mathbf{v}} - \lambda \\ G^{'}\{\lambda\} &= \mathbf{w}^H\{\lambda\}\mathbf{B}\mathbf{w}\{\lambda\} - 1 \\ G^{''}\{\lambda\} &= -2\mathbf{w}^H\{\lambda\}\mathbf{B}\left(\tilde{\mathbf{T}}+\lambda\mathbf{B}\right)^{-1}\mathbf{B}\mathbf{w}\{\lambda\} \\ &\quad -2\frac{\mathbf{w}^H\{\lambda\}\mathbf{B}\mathbf{w}\{\lambda\}}{\left(\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{B}\right)^{-1}\mathbf{v}\right)^2} \end{aligned} \tag{7.32}$$

for the ISL formulation and

$$\begin{aligned}
\mathbf{w}\{\lambda\} &= (\mathbf{T}+\lambda\mathbf{C})^{-1}\mathbf{v} \\
G\{\lambda\} &= -\mathbf{v}^H\mathbf{w}\{\lambda\} \\
G^{'}\{\lambda\} &= \mathbf{w}^H\{\lambda\}\mathbf{C}\mathbf{w}\{\lambda\} \\
G^{''}\{\lambda\} &= -2\mathbf{w}^H\{\lambda\}\mathbf{C}\,(\mathbf{T}+\lambda\mathbf{C})^{-1}\,\mathbf{C}\mathbf{w}\{\lambda\}
\end{aligned} \tag{7.33}$$

for the template formulation (where $\mathbf{C} = \mathbf{B} - \mathbf{v}\mathbf{v}^H$). Explicit algebraic proofs for the ISL case are given in Appendix A.3.

7.3.3 Maximizing the Dual

While Newton's method can converge to the optimal super-linearly when $\hat{\lambda}$ is close to $\lambda_\star$ [76, p. 44], near the boundaries $\mathbf{T}+\lambda\mathbf{C}$ may ill conditioned which may lead to steps that exceed the feasible region. To ensure convergence, the proposed maximization procedure utilizes an effective upper and lower bound (λ_a, λ_b) for the range of feasible λ as

$$\lambda^- \leq \lambda_a \leq \lambda_\star \leq \lambda_b \leq \lambda^+.$$

A Newton's step is taken from to the boundary point (either λ_a or λ_b) that corresponds to the larger Dual value. If the step yields a new point $\lambda \in (\lambda_a, \lambda_b)$, then the boundary of the range is updated based off the sign of $G^{'}\{\lambda\}$. When Newton's exceeds the range, a regula falsi update [150] (akin to secant method) is used to update the range. If both Newton's and regula falsi do not sufficiently shrink the range of the bounds, bisection is used. The algorithm is stopped if either the initial $\lambda_a \approx 0$ satisfies the QC, or when $|G^{'}\{\lambda\}|^2$ is within a desired tolerance.

As the dual is a scalar value, the search for the next query λ is computationally inexpensive and computational cost per iteration is driven by a single evaluation of the dual and it's derivatives. A basic outline of the maximization procedure is provided in Algorithm 1 where the $EvaluateDual\{\cdot\}$ function corresponds to evaluating (7.32) or (7.33) depending on the ob-

jective function. Thus, to ensure fast and stable convergence one must set an initial effective range that spans the optimal (i.e., $\lambda_\star \in (\lambda_a, \lambda_b)$) and is contained within the valid region (i.e., $(\lambda_a, \lambda_b) \in (\lambda^-, \lambda^+)$).

As the lower bound for λ is zero, the initial λ_a may be set either to zero or a small value

$$\lambda_a = \frac{tol}{\rho\left(\mathbf{v}^H \mathbf{R}^{-1} \mathbf{v}\right)} \approx 0 \tag{7.34}$$

to ensure $\mathbf{T} + \lambda \mathbf{C}$ is not singular. The search for the upper boundary is somewhat more challenging as a λ_b too large may result in slow convergence or may exceed λ^+. For the ISL case, $\lambda^+ = \infty$ so λ_b can simply be estimated iteratively via Newton's step from λ_a or a multiplicative increase (e.g., $\lambda_b = 10\lambda_a$) where the sign of the derivative will inform if the new estimate updates λ_a or λ_b. As mentioned in Section 7.3.1, the upper bound for the template objective can be explicitly computed via a generalized eigen decomposition, however, this can be computationally intensive. Alternatively, λ_b can be iteratively estimated using the procedure outlined in Section 7.4.2 which will asymptotically approach λ^+ and will rapidly yield a viable λ_b (empirically one or two iterations are typically sufficient).

7.4 Computational Cost and Efficient Implementations

For both objective functions, the computational complexity is primarily driven by the inversion and multiplication of a Hermitian Toeplitz matrix which we expressed as

$$\mathbf{T}\{\alpha\} = \mathbf{T} + \alpha \mathbf{R}_\mathbf{n} \tag{7.35}$$

where α is a strictly positive variable that is a function of λ. The remaining components (e.g., beamspoiling updates for the ISL case, subtraction of $\lambda \mathbf{v}\mathbf{v}^H$ for the template case, and addition of low-rank interference) can later be incorporated for a relatively low computational cost via low-rank matrix updates. Woodbury's is the natural choice for these updates (Appendix A.1) but,

Algorithm 1 Dual Maximization

$\lambda_a \leftarrow \hat{\lambda}^-$

$\left[\hat{\mathbf{w}}_a, G_a, G_a^{'}, G_a^{''}\right] \leftarrow EvaluateDual\{\lambda_a\}$

if $SINR_{Loss}\{\hat{\mathbf{w}}_a\} \geq \rho$ **then**

  $\mathbf{w} \leftarrow \hat{\mathbf{w}}_a$

  Return $\mathbf{w}$

$\lambda_b \leftarrow EstimateUpperBound\{\lambda_a\}$

$\left[\hat{\mathbf{w}}_b, G_b, G_b^{'}, G_b^{''}\right] \leftarrow EvaluateDual\{\lambda_b\}$

if $\left|G_b^{'}\right|^2 < tol$ **then**

  $\mathbf{w} \leftarrow \hat{\mathbf{w}}_b$

  Return $\mathbf{w}$

while $G_b^{'} > 0$ **do**

  $\left[\hat{\mathbf{w}}_a, G_a, G_a^{'}, G_a^{''}, \lambda_a\right] \leftarrow \left[\hat{\mathbf{w}}_b, G_b, G_b^{'}, G_b^{''}, \lambda_b\right]$

  $\lambda_b \leftarrow EstimateUpperBound\{\lambda_a\}$

  $\left[\hat{\mathbf{w}}_b, G_b, G_b^{'}, G_b^{''}\right] \leftarrow EvaluateDual\{\lambda_b\}$

if $G_a > G_b$ **then**

  $\left[\hat{\mathbf{w}}, G, G^{'}, G^{''}, \lambda\right] \leftarrow \left[\hat{\mathbf{w}}_a, G_a, G_a^{'}, G_a^{''}, \lambda_a\right]$

else

  $\left[\hat{\mathbf{w}}, G, G^{'}, G^{''}, \lambda\right] \leftarrow \left[\hat{\mathbf{w}}_b, G_b, G_b^{'}, G_b^{''}, \lambda_b\right]$

while $\left|G^{'}\right|^2 < tol$ **do**

  $\lambda \leftarrow \lambda - \frac{G^{'}}{G^{''}}$

  if $\lambda \notin (\lambda_a, \lambda_b)$ **then**

    $\lambda \leftarrow \frac{G_b^{'}\lambda_a - G_a^{'}\lambda_b}{G_b^{'} - G_a^{'}}$

  if $\frac{\lambda - \lambda_a}{\lambda_b - \lambda_a} < tol$ or $\frac{\lambda_b - \lambda}{\lambda_b - \lambda_a} < tol$ **then**

    $\lambda \leftarrow \frac{\lambda_a + \lambda_b}{2}$

  $\left[\hat{\mathbf{w}}, G, G^{'}, G^{''}\right] \leftarrow EvaluateDual\{\lambda\}$

  if $G^{'} > 0$ **then**

    $\left[\hat{\mathbf{w}}_a, G_a, G_a^{'}, G_a^{''}, \lambda_a\right] \leftarrow \left[\hat{\mathbf{w}}, G, G^{'}, G^{''}, \lambda\right]$

  else

    $\left[\hat{\mathbf{w}}_b, G_b, G_b^{'}, G_b^{''}, \lambda_b\right] \leftarrow \left[\hat{\mathbf{w}}, G, G^{'}, G^{''}, \lambda\right]$

Return $\mathbf{w}$

depending on the numeric precision and matrix conditioning, low-rank inverse Cholesky updates [121] may be necessary to ensure numeric stability. While there are several operations involving $\mathbf{T}^{-1}\{\alpha\}$, the Hermitian Toeplitz structure enables efficient representation and multiplication by computing just the first column of the inverse (see Sections 2.7.3 and 2.7.4). Thus, the computational complexity per iteration will be dominated by a single linear system solve as the inverse Toeplitz matrix-vector multiplications can be computed with $\mathscr{O}(Dlog(2D))$ cost and the low-rank inverse updates $\sim\mathscr{O}\left(r^3\right)$, will not affect the order of complexity in $EvaluateDual(\cdot)$.

For K $EvaluateDual(\cdot)$ instances in Algorithm 1, the order of complexity will follow $\mathscr{O}\left(KD^3\right)$ for a brute force linear system solve as a unique system is formed for every λ. However, as (7.35) is inherently Hermitian Toeplitz, Levinson-Durbin recursion can be used to reduce the complexity to $\mathscr{O}\left(KD^2\right)$ [109, p. 81-85]. There are more computationally efficient solvers [111, 112] that may push $\mathscr{O}(KDlog(D))$ for larger matrices, but numeric instability or a high constant complexity may hinder their use for this application.

Additionally, it is worth noting that $\mathbf{T}$ is banded with a matrix-bandwidth of N. Thus, when $\mathbf{R}_n$ is also banded, then the computational complexity can be reduced to $\mathscr{O}\left(K\left(z^2L+Dlog(2D)\right)\right)$ via a banded Cholesky decomposition where z is the max matrix-bandwidth between $\mathbf{T}$ and $\mathbf{R}_n$. Further complexity reduction may be possible by utilizing a banded Levinson implementation [120] or banded Circulant update [151], though again, these solvers may face challenges with numeric instability or high constant cost.

Iterative solvers such as preconditioned conjugate gradient (PCG) may also be used to solve the linear system to within some fixed tolerance. For Toeplitz systems, Circulant preconditioners [122] can be incorporated rather inexpensively and are often very effective when the spectral content of signal and noise is not sparse (see Section 3.2.1). For an average of J_C PCG iterations per system, the computational complexity of PCG with a Circulant preconditioner is $\mathscr{O}(KJ_CDlog(2D))$ and $J_C\sim(\eta log(D)+1)$ [152] where η is the condition number after preconditioning and can often greatly outpace a direct solver with negligible error (see Figure 7.4 and associated discussion in Section 7.6.1).

[74] proposed performing an eigen decomposition of $\mathbf{T}$ and pre-multiply out the eigenvectors to diagonalize subsequent calculations resulting in a computational complexity of $\mathscr{O}\left(D^3+DK\right)$. While the approach in [74] was proposed for white noise MML (i.e., a diagonal $\mathbf{B}$ and no interference or colored noise), a generalized implementation can be applied for an SINR constraint as derived in Section 7.4.1. However, the $\mathscr{O}\left(D^3\right)$ cost of eigen decomposition may still be prohibitive for longer filters.

[75, 153] suggested that $\mathbf{T}$ could be approximated as a Circulant matrix so that the system can be diagonalized by a fast Fourier transform (see Section 2.7.2). Likewise, this philosophy may be extended to form Circulant approximations for $\mathbf{T}$ and $\mathbf{B}$ to diagonalize SINR constrained form. This in turn greatly reduces computational complexity to $\mathscr{O}\left(Dlog(D)+DK\right)$, but yields an infinite impulse response (IIR) approximation to the true finite impulse response (FIR) MMF that may not be efficacious in all situations. Truncation of the IIR MMF may alleviate the spreading of model error, but said truncated will induce additional model mismatch [81]. Regardless, certain scenarios are amenable to the IIR Circulant approximate MMF and may even improve performance as the number of filter taps is equal to the number of delay indices.

As the eigenvalues of a Toeplitz matrix and it's Circulant approximation are asymptotically equivalent as $D \rightarrow \infty$, the approximate QCQP may be first solved at low cost to provide an initial estimate for $\lambda_\star$. If this initialization is then close to the optimal for the Toeplitz QCQP, the required number of iterations may be greatly reduced.

Here, the current effective bounds of λ and the structure of the Toeplitz Matrix inverse are leveraged to develop informed preconditioners to solve the current (7.35). For initial system corresponding to $\lambda_a \approx 0$, the inverse Toeplitz matrix is either computed via PCG with a Circulant preconditioner or through a direct solver depending on the spectral content of $\mathbf{s}$. The search for λ_b is then performed using PCG using the inverse Toeplitz matrix of the λ_a system as the preconditioner. Once the range (λ_a, λ_b) has been established, novel Upper-Lower Bound (ULB) preconditioners constructed via an optimal a pareto weighting of the current upper and lower Toeplitz inverses are utilized to update all future Toeplitz systems. Section 7.4.3 outlines this ULB preconditioner in

more detail and provides performance bounds for specific cases.

This approach requires solving initial Toeplitz systems corresponding to λ_a yielding a constant complexity of $\mathscr{O}(T(x))$ where $T(x)$ denotes the cost of the initial solve (can be direct solution $\sim\mathscr{O}\left(D^2\right)$ or PCG $\sim\mathscr{O}(2J_C D\log(2D))$). The search for λ_b then has a cost of $\mathscr{O}(8FJ_1 D\log(2D))$ where F corresponds to the number of calls to $EstimateUpperBound\{\cdot\}$ that were required and J_1 was the average PCG iterations for the inverse Toeplitz preconditioners. Subsequent instances of (7.35) are solved using the ULB preconditioners yielding a cost of

$$\mathscr{O}(16K_G J_2 D\log(2D))$$

where K_G is the number of iterations required for the Dual search to converge and J_2 was the average number of PCG iterations for the ULB preconditioners. In total, the order of complexity is for a pure PCG implementation is $\sim\mathscr{O}((2J_C+8FJ_1+16KJ_2)D\log(2D))$. Further assessment and comparisons regarding convergence and complexity of the ULB preconditioner are examined in Section 7.6.1 and Section 7.4.3.

7.4.1 Extending previous efficient solvers to the General SINR QCQP

In [74, 75] the QC is configured to constrain MML for white noise, meaning $\mathbf{B}$ is a uniform weighted diagonal and $\lambda\mathbf{B}$ effectively acts as diagonal loading on the eigenvalues of $\mathbf{T}$. Thus, [74] suggests performing an initial eigen decomposition

$$\mathbf{T}=\mathbf{U}\Lambda\mathbf{U}^H \tag{7.36}$$

and rotating all parameters by the eigenvectors $\mathbf{U}$, e.g,

$$\tilde{\mathbf{w}}=\mathbf{U}^H\mathbf{w},$$

the rotated Dual becomes a function of the sum of diagonal matrices

$$\mathbf{U}^H(\mathbf{T}+\lambda\mathbf{B})\mathbf{U} = \Lambda + \lambda\mathbf{B}. \tag{7.37}$$

This QCQP can then be solved efficiently using element wise vector-vector operations to solve for $\tilde{\mathbf{w}}_\star$, and a final rotation of

$$\mathbf{w}_\star = \mathbf{U}\tilde{\mathbf{w}}_\star$$

will yield an equivalent optimal solution.

This procedure does help improve computational complexity of the Dual maximization, but the eigen decomposition has a computational complexity of $\mathscr{O}(D^3)$ which can still be a computational hindrance for longer filters. [75] leveraged this same eigen decomposition philosophy, but instead replaced $\mathbf{T}$ with it's Circulant approximation such that the eigen vectors corresponded to the DFT matrix (see Section 2.7.2). The eigen decomposition and parameter rotation can then be performed with $\mathscr{O}(Dlog(D))$ complexity, however, the resulting filter is effectively an IIR approximation to the true FIR MMF.

For the generalized SINR QC ($\mathbf{B}$ not necessarily diagonal), but eigen decomposition based diagonalization is still possible. When both $\mathbf{T}$ and $\mathbf{B}$ are replaced with their Circulant approximations, both sets of eigenvectors will correspond to the DFT, thus, the diagnonalized form is simply

$$\mathbf{U}^H(\mathbf{T}+\lambda\mathbf{B})\mathbf{U} = \Lambda + \lambda\Lambda_{\mathbf{B}} \tag{7.38}$$

which still has complexity of $\mathscr{O}(DlogD)$. For the Toeplitz case, the Cholesky decomposition of

$$\mathbf{B} = \mathbf{L}_{\mathbf{B}}\mathbf{L}_{\mathbf{B}}^H \tag{7.39}$$

can be used to express the hessian as

$$\mathbf{T}+\lambda\mathbf{B} = \mathbf{L}_{\mathbf{B}}\left(\mathbf{L}_{\mathbf{B}}^{-1}\mathbf{T}\mathbf{L}_{\mathbf{B}}^{-H}+\lambda\mathbf{I}\right)\mathbf{L}_{\mathbf{B}}^H \tag{7.40}$$

such that the eigen decomposition of

$$\mathbf{L}_{\mathbf{B}}^{-1}\mathbf{T}\mathbf{L}_{\mathbf{B}}^{-H} = \tilde{\mathbf{U}}\tilde{\Lambda}\tilde{\mathbf{U}}^{H} \tag{7.41}$$

and variable rotation with respect to

$$\mathbf{L}_{\mathbf{B}}\tilde{\mathbf{L}}^{H}$$

results in a diagonalized QCQP Dual. Doing so does require additional decompositions and matrix-matrix multiplies, all with the same expensive $\mathscr{O}(D^3)$ order of complexity.

7.4.2 Rapid Estimation of the Template Upper Bound

Computing the upper bound for the template form requires determining when

$$\mathbf{T} + \lambda\left(\mathbf{B} - \mathbf{v}\mathbf{v}^{H}\right) \succeq 0\,. \tag{7.42}$$

For any $\rho > 0$ there exists a feasible λ_a such that

$$\mathbf{T} + \lambda_a\left(\mathbf{B} - \mathbf{v}\mathbf{v}^{H}\right) = \mathbf{A}\{\lambda_a\} - \lambda_a\mathbf{v}\mathbf{v}^{H} \succeq 0$$

and

$$G^{'}\{\lambda_a\} > 0.$$

Given a small (or zero) initial guess for λ_a (i.e., (7.34)) where

$$\det\{\mathbf{A}\{\lambda_a\}\} > 0 \tag{7.43}$$

and

$$\det\left\{\mathbf{A}\{\lambda_a\} - \lambda_a\mathbf{v}\mathbf{v}^{H}\right\} > 0, \tag{7.44}$$

the downdate weighting β that satisfies

$$\begin{aligned} det\{\mathbf{A}\{\lambda_a\} - \beta\mathbf{v}\mathbf{v}^H\} &= 0 \\ &= \left(1 - \beta\mathbf{v}^H\mathbf{A}^{-1}\{\lambda_a\}\mathbf{v}\right) det\{\mathbf{A}\{\lambda_a\}\} \end{aligned} \tag{7.45}$$

can be computed via a rank-1 determinant update formula. This weighting is agnostic to $det\{\mathbf{A}\{\lambda_a\}\}$ and can be computed as

$$\beta = \frac{1}{\mathbf{v}^H\mathbf{A}^{-1}\{\lambda_a\}\mathbf{v}} = \lambda_b. \tag{7.46}$$

Since $\mathbf{B}$ is PD,

$$\mathbf{A}\{\lambda_b\} > \mathbf{A}\{\lambda_a\}$$

and

$$\mathbf{A}^{-1}\{\lambda_b\} < \mathbf{A}^{-1}\{\lambda_a\}$$

such that

$$\left(1 - \lambda_b\mathbf{v}^H\mathbf{A}^{-1}\{\lambda_b\}\mathbf{v}\right) > \left(1 - \lambda_b\mathbf{v}^H\mathbf{A}^{-1}\{\lambda_a\}\mathbf{v}\right) = 0$$

meaning the estimate upper bound will asymptotically approach λ^+ but never reach it. Fortunately, any λ_b that satisfies

$$G^{'}\{\lambda_b\} < 0$$

may be treated as an effective upper bound since the problem is concave. Thus, this asymptotic update procedure is used to iteratively estimate an effective upper bound as the $EstimateUpperBound\{\cdot\}$ function in Algorithm 1. An illustrative example is shown in Figure 7.1 where the first upper bound estimate satisfied $G^{'}\{\lambda_b\} < 0$.

7.4.3 Lower-Upper Preconditioner

Rather than strictly using a Circulant preconditioner for to solve each Toeplitz system, information from the previous iterations can be leveraged to develop more effective preconditioners. Once an

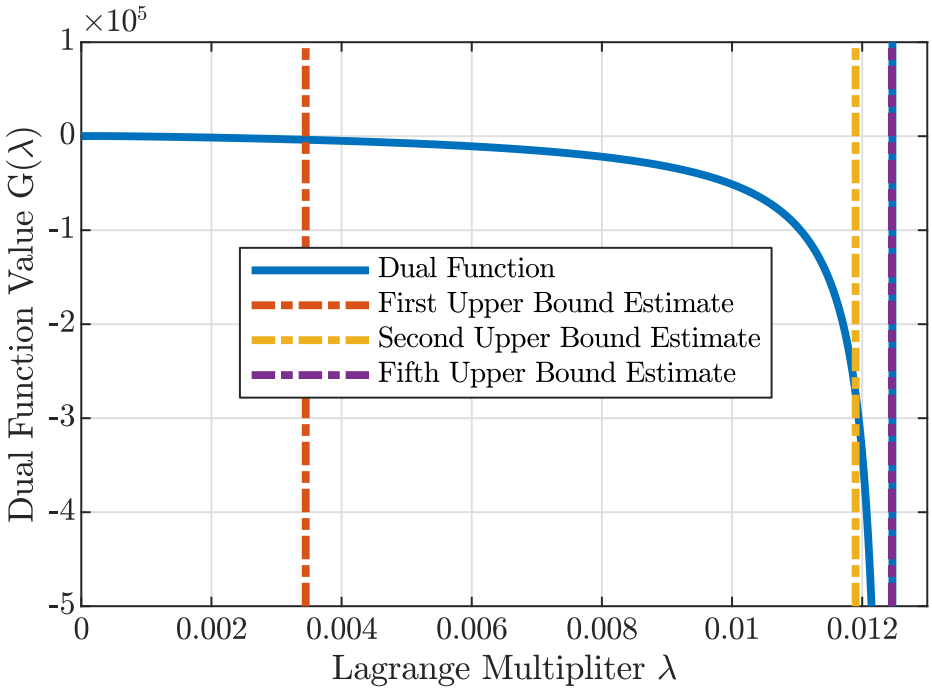

Figure 7.1: Iterative estimate of the template-based dual upper bound

upper and lower bound (λ_a and λ_b respectively) for the Dual has been established, PCG attempts to solve for the inverse of

$$\mathbf{Z}\{\lambda\} = \mathbf{T} + \lambda\mathbf{B} \tag{7.47}$$

where $\lambda_a \leq \lambda \leq \lambda_b$ defines the bounds and both $\mathbf{Z}^{-1}\{\lambda_a\}$ and $\mathbf{Z}^{-1}\{\lambda_b\}$ were previously computed to evaluate the dual for previous iterations. As $\mathbf{B}$ is PD, the eigenvalues of $\mathbf{Z}\{\lambda\}$ have a similar bounded relationship as λ, thus a preconditioner $\mathbf{M}^{-1}\{\lambda\}$ that involves linear combinations of $\mathbf{Z}^{-1}\{\lambda_a\}$ and $\mathbf{Z}^{-1}\{\lambda_b\}$ can be an effective preconditioner as long as it clusters the eigenvalues of the preconditioned system, has low complexity for matrix-vector multiplication, and remains hermitian PD.

A preconditioner of the form

$$\mathbf{M}^{-1}\{\lambda, \gamma\} = \gamma\mathbf{Z}^{-1}\{\lambda_a\} + (1-\gamma)\mathbf{Z}^{-1}\{\lambda_b\} \tag{7.48}$$

can be applied with $\mathcal{O}(16Dlog(2D))$ complexity and (in a similar vein to the analysis in Section

7.4.1) can be shown to have eigenvalues of

$$f\{\mu,\gamma\} = \gamma\frac{\mu+\lambda}{\mu+\lambda_a} + (1-\gamma)\frac{\mu+\lambda}{\mu+\lambda_b} \tag{7.49}$$

where μ corresponds to the eigenvalues of (7.41). A similar eigen functional relationship can be drawn for alternative linear combinations of $\mathbf{Z}^{-1}\{\lambda_a\}$ and $\mathbf{Z}^{-1}\{\lambda_b\}$, but for brevity and simplicity here the focus is on a simple pareto combination. Additionally, upper and lower bounds for μ can be determined as

$$0 \leq \mu \leq \hat{\mu}_u$$

where the lower bound stems from the PSD nature of $\mathbf{T}$ and $\hat{\mu}_u$ is calculated via the max eigen values of the Circulant approximate system of $\mathbf{T}$ divided by the white noise variance since true max eigenvalue $\mu_u \leq \hat{\mu}_u$ and are asymptotically equivalent as $D \to \infty$.

Taking the derivative of (7.49) with respect to μ can be used to determine the positive extrema (at most two) by rearranging into a single fraction and using the quartic equation of the numerator. These extrema, $\mu = 0$, and $\mu = \hat{\mu}_u$ can then be used to determine the condition number for various γ. The best possible pareto weight can then quickly be determined by minimizing the extrema distance through an inexpensive quarry search. For the optimal γ, clearly

$$f\{\mu,\gamma\} \leq \frac{\mu+\lambda}{\mu+\lambda_a}$$

and

$$f\{\mu,\gamma\} \leq \frac{\mu+\lambda}{\mu+\lambda_b}$$

will hold as a value of 0 or 1 will yield the respective preconditioned functions as illustrated the examples in Figures 7.2 and 7.3 where the largest eigenvalue is greater than and less than λ_b respectively and the normalized range of (λ_a, λ_b) corresponds to $\lambda_b = 10^5\lambda_a$ which is a larger (and worse performing) range than would realistically show up in Algorithm 1.

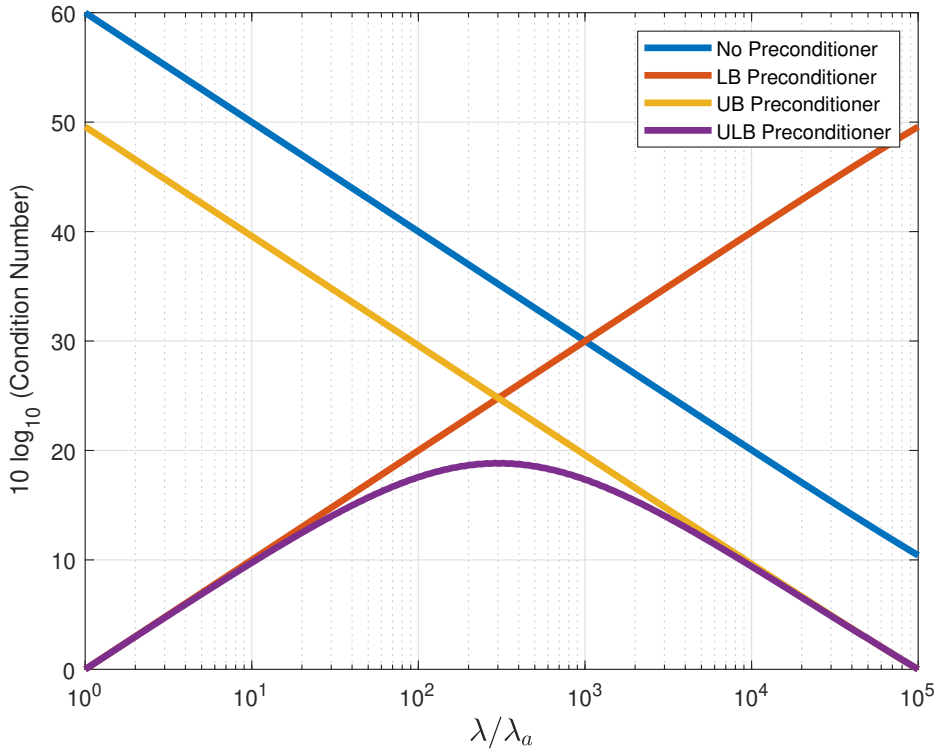

Figure 7.2: Condition Number Bounds for the ULB Preconditioner when $\mu_u > \lambda_b$

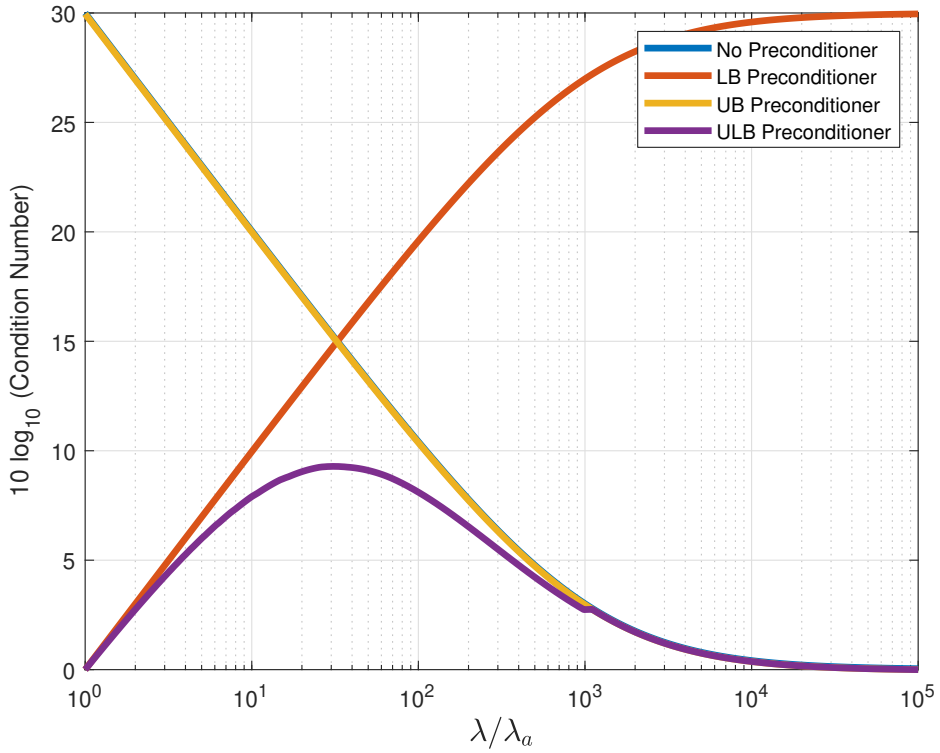

Figure 7.3: Condition Number Bounds for the ULB Preconditioner when $\mu_u < \lambda_b$

7.5 A Note on Multi-Waveform Generalizations and Quadratically Constrained MiCRFt

For a waveform agile CPI, some waveforms may have a better natural fit meaning the SINR loss may vary waveform-to-waveform. Rather than optimizing each MMF to satisfy the constraint individually, (7.13) and (7.14) may be modified to instead constrain the aggregate SINR loss for M waveforms as

$$\begin{aligned} \min_{\mathbf{w}_1^*,\cdots\mathbf{w}_M^*} \quad & \sum_{m=1}^{M} \mathbf{w}_m^H \tilde{\mathbf{T}}_m \mathbf{w}_m \\ \text{s.t.} \quad & \sum_{m=1}^{M} \mathbf{w}_m^H \mathbf{C}_m \mathbf{w}_m - M \leq 0 \\ & \mathbf{w}_m^H \mathbf{v}_m = 1 \quad m = 1, \cdots, M \end{aligned} \tag{7.50}$$

for the ISL objective or

$$\begin{aligned} \min_{\mathbf{w}_1^*,\cdots\mathbf{w}_M^*} \quad & \sum_{m=1}^{M} \mathbf{w}_m^H \mathbf{T}_m \mathbf{w}_m - 2\,\mathrm{Re}\left\{\mathbf{w}_m^H \mathbf{b}_m\right\} \\ \text{s.t.} \quad & \sum_{m=1}^{M} \mathbf{w}_m^H \mathbf{C}_m \mathbf{w}_m \leq 0 \end{aligned} \tag{7.51}$$

for the template objective. The dual for (7.50) and (7.51) then have an analogous form to the single waveform case where optimal set of filters can be solved by maximizing with respect to a single Lagrange multiplier.

A SINR QC constraint can likewise be incorporated into MiCRFt by simply extending the interference plus noise covariance to a block formulation via

$$\bar{\mathbf{R}} = \bar{\mathbf{R}}_{\bar{\mathbf{i}}} + \bar{\mathbf{R}}_{\bar{\mathbf{n}}} \tag{7.52}$$

with

$$\bar{\mathbf{R}}_{\bar{\mathbf{i}}} = \mathscr{E}\left\{\bar{\mathbf{i}}\bar{\mathbf{i}}^H\right\}$$

and

$$\bar{\mathbf{R}}_{\bar{\mathbf{n}}} = \mathscr{E}\left\{\bar{\mathbf{n}}\bar{\mathbf{n}}^H\right\}$$

where $\bar{\bar{\mathbf{i}}} = \left[\mathbf{i}_1^T\ \mathbf{i}_2^T\ \cdots\ \mathbf{i}_Q^T\right]^T$ and $\bar{\bar{\mathbf{n}}} = \left[\mathbf{n}_1^T\ \mathbf{n}_2^T\ \cdots\ \mathbf{n}_Q^T\right]^T$ are the concatenated interference and noise vectors for the Q subsets respectively. The remaining QCQP can then be set up using this interference-plus-noise matrix to form a block QC and the QCQP can then be solved using Algorithm 1 by replacing each respective vector/matrix with its block counterpart. Depending on how the SINR matrix is posed, MiCRFt may attempt to suppress interference via a complementary cancellation across subsets, e.g. the block interference-plus-noise covariance is non-zero on the off-diagonal-blocks. As the complementary condition requires presumming across slow-time, MiCRFt inherently operates on a longer time-scale than the LS MMFs and the LTI assumption of the interference may not be as accurate. Thus, it may be beneficial to enforce that the interference-plus-noise matrix is strictly a block-diagonal such that the interference is addressed on a per-filter basis (e.g., on a shorter-time scale) and the complementary condition only addresses sidelobes.

7.6 QC LS MMF Closed-Loop Assessment

To evaluate the efficiency of the proposed QCQP algorithm and efficacy of the resulting filters, two coherent processing intervals (CPIs) of FM waveforms were generated on a Tektronix arbitrary waveform generator (AWG). Both CPIs consisted of 500 pulses with a center frequency of 3.45 GHz, 10 micro-second pulse duration, a 3dB bandwidth of 50 MHz (time-bandwidth product of 500), and a 10 kHz PRF (10% duty cycle). The first CPI followed a standard repeated LFM chirp waveform framework, while the second used a waveform agile set of non-repeating optimized random FM (RFM) waveforms with a Gaussian power spectral density [73]. The AWG was connected to a bandpass filter and class A amplifier to model TX amplification, followed by an attenuator, bandpass filter, low-noise amplifier (LNA), and then a Rhode & Schwarz spectrum analyzer with a sample rate of 65 MHz (κ=1.3) so that closed-loop measurements of the waveforms on real hardware could be collected. Note, relative to the first experimental test, these waveforms correspond to an increased *TB* such that the respective MMFs have more DoF. This particular dataset was selected as the LS MMFs from the first experiment induced higher intersection RSM due to

Dual Evaluations				
tol	10^{-3}	10^{-6}	10^{-9}	10^{-12}
# calls	33	45	48	52

Table 7.1: Calls to $EvaluateDual\{\cdot\}$ - ISL MMF D=4N, ρ=-2dB, k=3

the extended filter length. This dataset on the otherhand, can not see the intersection without the use of LS MMFs.

Using these loopback measurements, a multi-variable MMF performance surface can be created to assess the contribution of the multitude of free variables (template choice, degree of beamspoiling, filter length, mismatch loss, etc.). However, as general trends for these MMF parameters have been examined previously [55, 154] and since the QCQP form guarantees the resulting filter is optimal with respect to the objective function, the remainder of this section is restricted to a select few examples to illustrate particular aspects of the contributions from this work. A more thorough tabularized assessment of MMF performance with respect to the free variables is examined using open-air experimental measurements in Section 7.7.

7.6.1 Algorithm Performance

First, to demonstrate the efficiency of the proposed algorithm a Toeplitz-model ISL QCQP with a diagonal interference-plus-noise covariance (i.e. a white noise MML only form) was posed for the LFM waveform with filter size $D = 4N$ four times the duration of the transmit waveform, a SINR loss constraint of $\rho = -2$dB, and mainlobe beamspoiling of $k = 3$. The dual maximization algorithm was then run for this QCQP with different stopping/PCG tolerances. The resulting MMF responses are illustrated in Fig. 7.4 where even the largest tolerance ($tol = 10^{-3}$) induces only small errors and smaller tolerances ($tol \leq 10^{-6}$) all yield effectively identical results. The total number of $EvaluateDual\{\cdot\}$ calls in Algorithm 1 for the respective tolerance is shown in Table 7.1 where $\sim$45 evaluations were required to effectively convergence to the fidelity of the model.

Next to test the rate of convergence for the proposed preconditioners, the same QCQP was solved and total number of PCG iterations was recorded for conjugate gradient (CG) without a

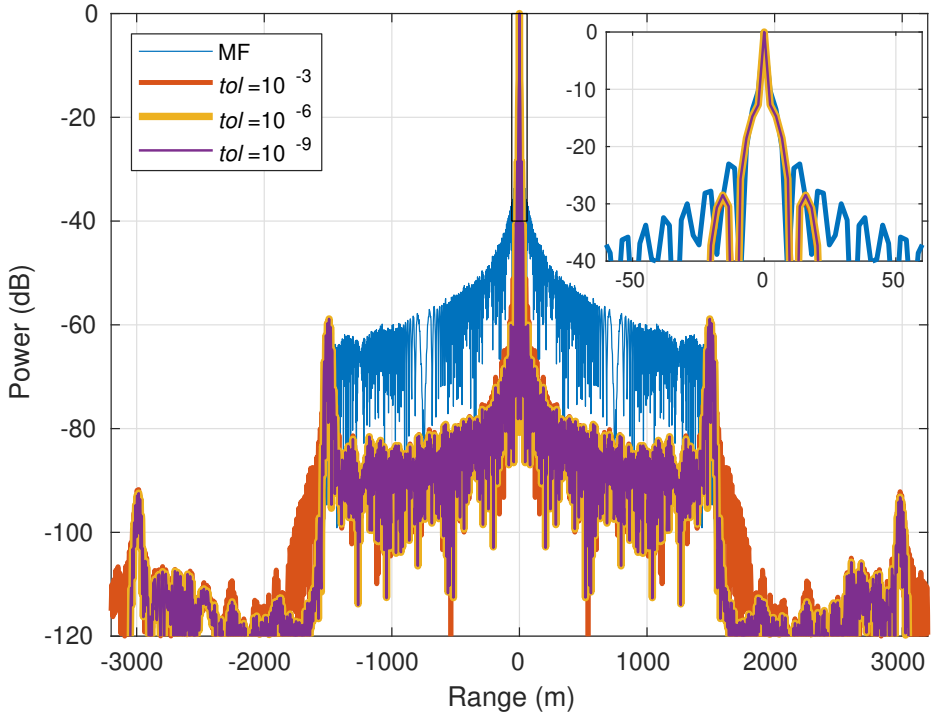

Figure 7.4: Resulting MMFs for different tolerances

preconditioner, the optimal Circulant preconditioner for each system (PCG Circ), and the proposed approach. These results are shown in Table 7.2 where the proposed approach consistently required $\sim 4\times$ fewer iterations than either of the alternatives. However, as the proposed approach requires more and often larger FFTs, the order of complexity to solve the dual for these PCG approaches was likewise calculated and compared to other proposed Toeplitz solvers as shown in Table 7.3. Here the "brute force"" heading corresponds to a solver that does not utilize the Toeplitz structure (such as [155]), "Eigen"" corresponds to the eigen decomposition approach in [74], while "Levinson" utilizes Levinson-Durbin recursion [109, p. 81-85] as the Toeplitz solver. Here the order of complexity for solvers that utilize the Toeplitz show a significant reduction in cost, though, it should be noted that the eigen decomposition approach of [74] front-loads the computational burden which can offer significant performance improvement when multiple QCQPs are posed for same Toeplitz components, (e.g., solving multiple filters for the same waveform with different ρ, k, etc.).

While this order of complexity metric demonstrates that the proposed approach has the low-

Number of PCG iterations				
tol	10^{-3}	10^{-6}	10^{-9}	10^{-12}
CG	3366	18954	40906	66283
PCG Circ	668	1417	2359	3412
Proposed	141	239	380	487

Table 7.2: Total number of PCG iterations - ISL MMF D=4N, ρ=-2dB, k=3

Order of Complexity				
tol	10^{-3}	10^{-6}	10^{-9}	10^{-12}
Brute Force [155]	6.70E+11	9.14E+11	9.74E+11	1.06E+12
Eigen [74]	2.03E+10	2.03E+10	2.03E+10	2.03E+10
Levinson [109]	2.48E+08	3.38E+08	3.60E+08	3.91E+08
CG	1.14E+8	6.42E+08	1.39E+09	2.24E+09
PCG Circ	4.52E+07	9.60E+07	1.60E+08	2.31E+08
Proposed	3.85E+07	6.59E+07	1.06E+08	1.35E+08

Table 7.3: Order of Complexity - ISL MMF L=4N, ρ=-2dB, k=3

est theoretical cost, typical MMF lengths are amenable to parallelization which can significantly improve runtime of the PCG solvers relative to the direct approaches when utilizing parallel processing on a GPU. In particular, a Levinson approach is inherently recursive and necessitate D^2 sequential loops which can create a bottleneck, while the PCG approaches are largely limited by the number of sequential fast Toeplitz matrix-vector multiplies (from preconditioning and/or multiplication with the CG hessian). To provide a semblance of potential for parallelization Table 7.4 provides the total number of sequential Toeplitz multiplies required for the PCG solvers compared with the number of non-parallelizable Levinson loops, where parallelization has the potential to improve runtime by orders of magnitude.

Sequential Toeplitz multiplies/Levinson recursions				
tol	10^{-3}	10^{-6}	10^{-9}	10^{-12}
Levinson [109]	90024	122760	130944	141856
CG	3366	18954	40906	66283
PCG Circ	1336	2834	4718	6824
Proposed	471	816	1313	1669

Table 7.4: Total number of non-parallelizable Toeplitz multiplies/Levinson recursions - ISL MMF D=4N, ρ=-2dB, k=3

7.6.2 Closed-Loop MMF Results

Using the loopback captured waveforms, ISL QCQPs were posed for each waveform to examine the range and frequency characteristics of the filter responses with and without in band interference for a fixed SINR loss of $\rho = -2$dB, filter length of $D = 4N$ and variable degree of beamspoiling: k=3 - MMF 1 (orange); k=5 - MMF 2 (yellow); k=13 - MMF 3 (purple). Figure 7.5 illustrates the RMS range response to give a sense of per-pulse sidelobe levels while Fig. 7.5 illustrates the mean frequency response across the CPI. The top row of subplots in each figure correspond MF (blue) and MMFs for the LFM waveform with an SNR loss constraint (a.) and with an SINR loss constraint (b.) relative to the optimal filter (green) and interference (black) generated via a 50 ms OFDM communication waveform with 4 MHz of bandwidth and is 20dB above the noise floor, i.e. the interference-to-noise ratio (INR) is 20dB. The bottom row of subplots (c. and d.) correspond to the same QC applied to each of the PROFM waveforms individually[3].

For the SNR constrained filters, the rage-sidelobe levels trend lower as the degree of beamspoiling increases while the spectral shape becomes "smoother" until the beamspoiling subsumes the first sidelobe where a rippling effect starts to appear present and a shoulder-lobe behavior appears in the range domain. The relatively small degree of oversampling on receive (κ=1.3) and fairly flat spectral content likewise induces a "ringing" behavior at multiple of N samples, which is a well known artifact [55]. The relatively flat frequency characteristics of the LFM waveform provide a convenient surface in which the MMF can easily induce a smooth frequency response (and thus low sidelobes), while the random and spikey nature of the PROFM waveforms typically yield higher RMS sidelobes. However, due to the non-repeating nature of the agile CPI, these range sidelobes will actually achieve a noise-like coherent averaging effect proportional to the number of pulses ($500 \sim 27dB$) and the zero-Doppler slice will have a range sidelobes corresponding to mean frequency response. That being said, as the range-sidelobes change pulse-to-pulse performing Doppler processing will result in a "smearing" of ambiguity due to RSM.

[3]For a waveform agile CPI, it is likewise possible and likely desirable to apply the QC in terms of the aggregate SINR loss across the whole CPI, rather than for each waveform individually. Doing so simply involves the summation of several components and only necessitates trivial modification of the dual and its derivatives via chain rule.

When the interference is incorporated into the QCQP, the MMFs attempt to place a notch in the region of the interference which in turn induces a sinc-like sidelobe roll-off in range. A similar phenomena has been observed when spectral notching is incorporated on transmit and is due to the inverse Fourier relationship between sinc and rectangular functions. In terms of range-sidelobe levels, this spectral notch significantly degrades sidelobe suppression for the LFM waveform until a large degree of beamspoiling is incorporated. Interestingly, outside of the sinc-like roll-off, the RMS PROFM range-response does not degrade as significantly as the LFM even though the average frequency responses behave similarly.

7.7 QC LS MMF Open-Air Experimental Results

Next the attenuator from the closed-loop setup was removed and the respective transmitter and receiver chains were connected to parabolic antennas. This test setup was taken to the roof of Nichols hall at the University of Kansas and placed in a pseudo monostatic configuration where TX and RX happens simultaneously (see Fig. 7.7). The antennas were pointed at an intersection with moving vehicles $\sim$1000 meters away and the two CPIs from the previous section were sequentially transmitted to collect open-air experimental measurements of a scene with moving vehicles and stationary clutter. Note, these the selected pulse duration and test setup resulted in direct path sidelobes that extend $\sim$1500 meters which subsume the intersection.

For the first open-air experiment, a white noise MML constraint QC was used to generate Toeplitz MMFs for the LFM and PROFM CPIs using the ISL MMF 2 parameters from Section 7.6.2 (ρ=-2dB; $D = 4N$; k=5) as well as template objective MMFs with ρ=-2dB, $D = 4N$, and $\mathbf{p}$ defined as a Gaussian times a Tukey window in the frequency domain corresponding to a desired 3dB oversampling of $\kappa_{\mathbf{w}} = 1.3$ where the Tukey component was incorporated to coarsely model the anti-aliasing filter[4]. These MMFs were then applied to the respective measurements and Doppler

[4]While a template design designed to be gaussian may have lower sidelobes, accounting for the sharp rolloff induced by the anti-aliasing filter typically enables the waveform-filter response to better match the template for less mismatch loss. Additionally, distortion an model errors due to range straddling is most significant on the edge of the band, thus avoiding placing energy there naturally provide some robustness towards range straddling.

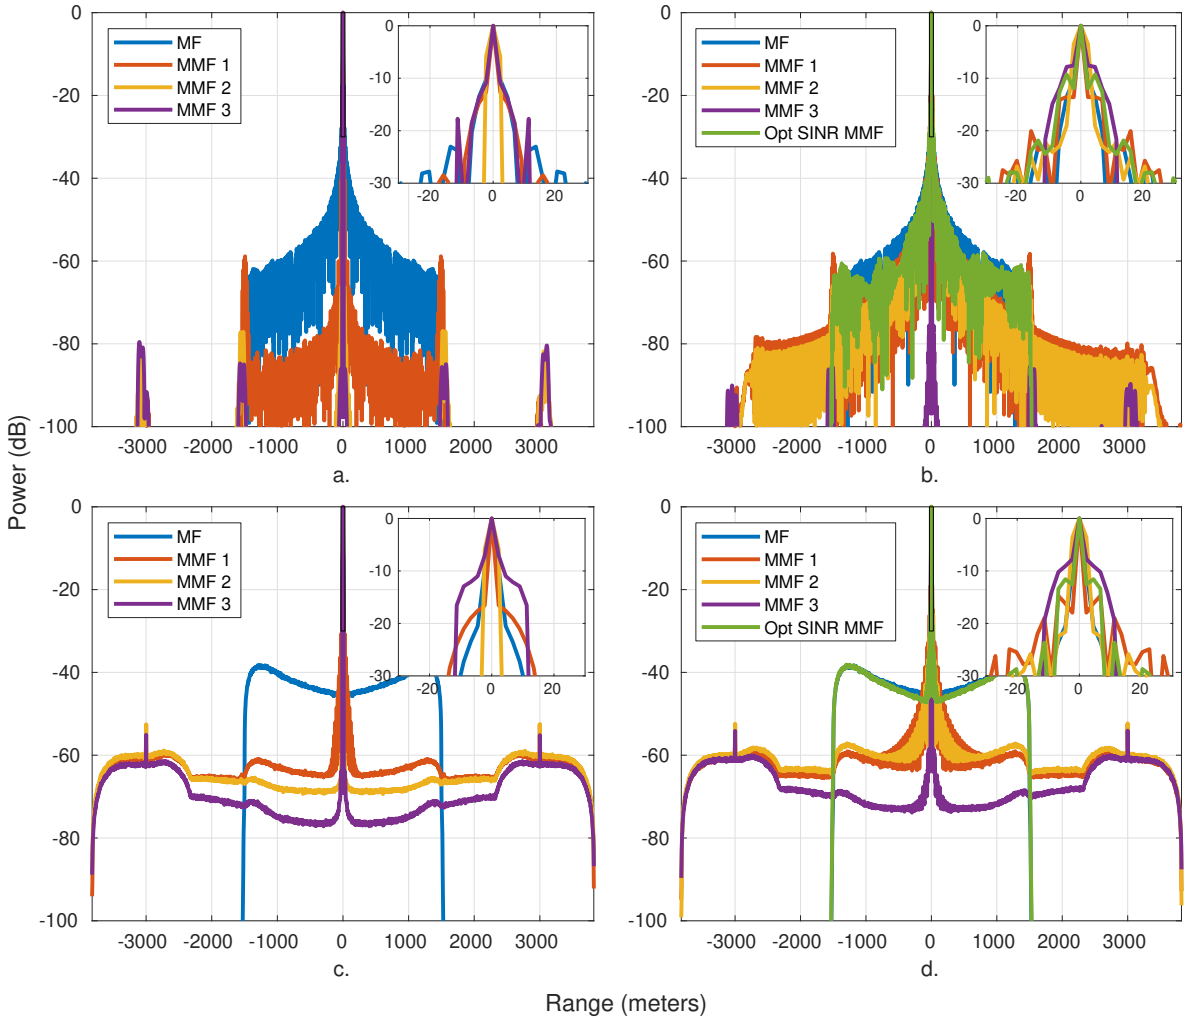

Figure 7.5: RMS Filter Response for MF and MMFs with ρ=-2dB; $D = 4N$; k=3 (MMF 1), k=5 (MMF 2), or k=13 (MMF 3)
a. - LFM with SNR constraint; b. - LFM with SINR constraint; c. - PROFM with SNR constraint; d. - PROFM with SINR constraint;

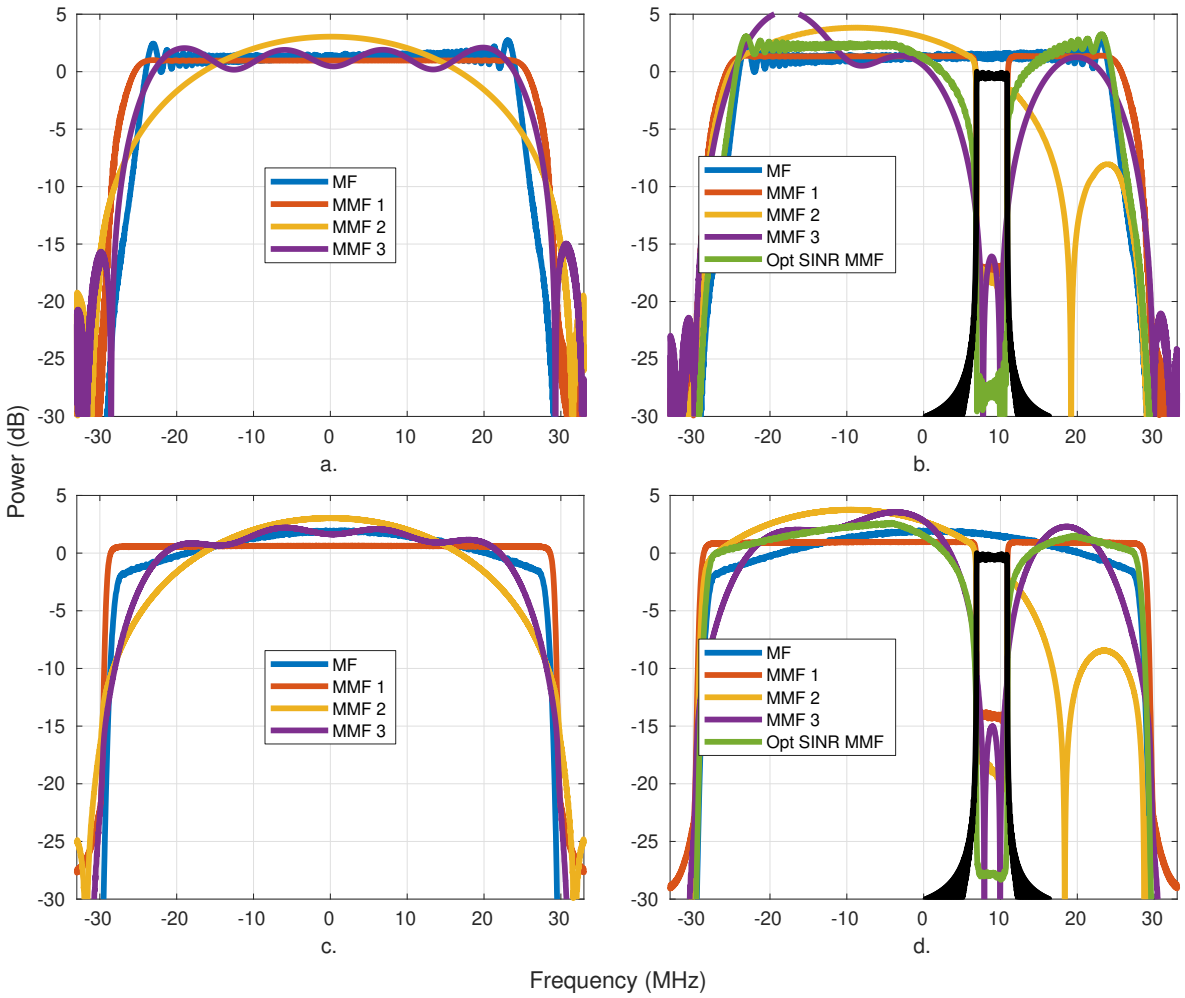

Figure 7.6: Mean Frequency Response for MF and MMFs with ρ=-2dB; $D = 4N$; k=3 (MMF 1), k=5 (MMF 2), or k=13 (MMF 3)
a. - LFM with SNR constraint; b. - LFM with SINR constraint; c. - PROFM with SNR constraint; d. - PROFM with SINR constraint;

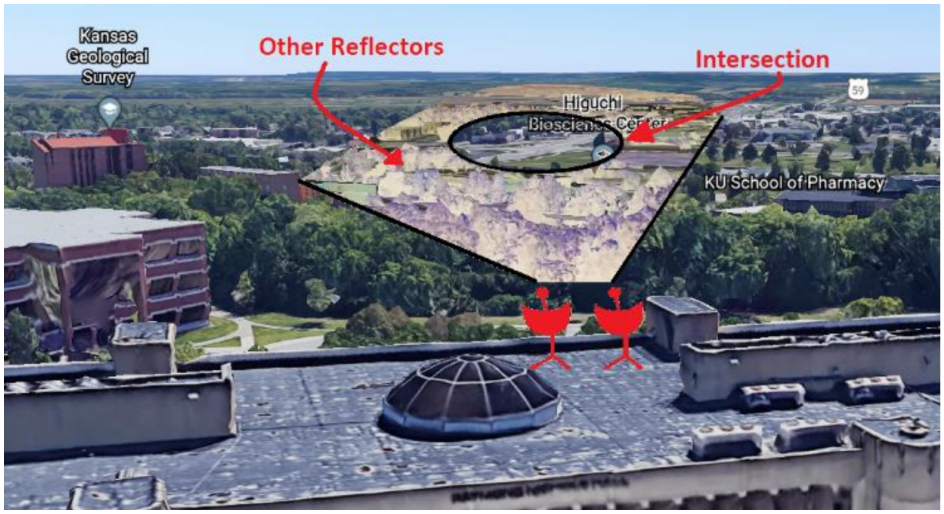

Figure 7.7: Open-Air Experimental Test Setup

processing with a Taylor window was performed. Figure 7.8 illustrates the range-Doppler maps of the intersection where the top row of subplots (a.-c.) correspond to the LFM MF, ISL MMF, and template MMF respectively while the bottom row (d.-f.) correspond to the same QCQP setup but for the PROFM CPI. As each filters was normalized to yield a pulse-compression peak of unity, the color axes (and all subsequent color axes) are normalized such that 0 dB corresponds to the MF noise floor. For the LFM CPIs, the range-sidelobes extending from the direct path (seen at zero velocity and $\sim$1050 meters) and targets are clearly reduced using the MMF approaches, though the background floor has come up slightly due to the MML. For the PROFM MF, the direct path RSM appears as a significantly raise background floor which masks all the movers and makes detection unlikely. After applying the MMFs to the PROFM measurements, the background floor is push down $\sim$24 dB which correlates well with what was expected with the loopback assessment from Fig. 7.5. While the direct-path RSM floor is still higher than the noise floor, many of the movers are now detectable which was not possible with the MF.

Next, the synthetic OFDM signal from Section 7.6.2 was added to the open-air measurements with an INR of 20 dB and the corresponding Toeplitz QC MMFs were generated using the same parameters as before. Additionally, the optimal SINR filters with the same filter length as the MF ($D = N$) were generated. Figure 7.9 illustrates the output after range-Doppler processing when

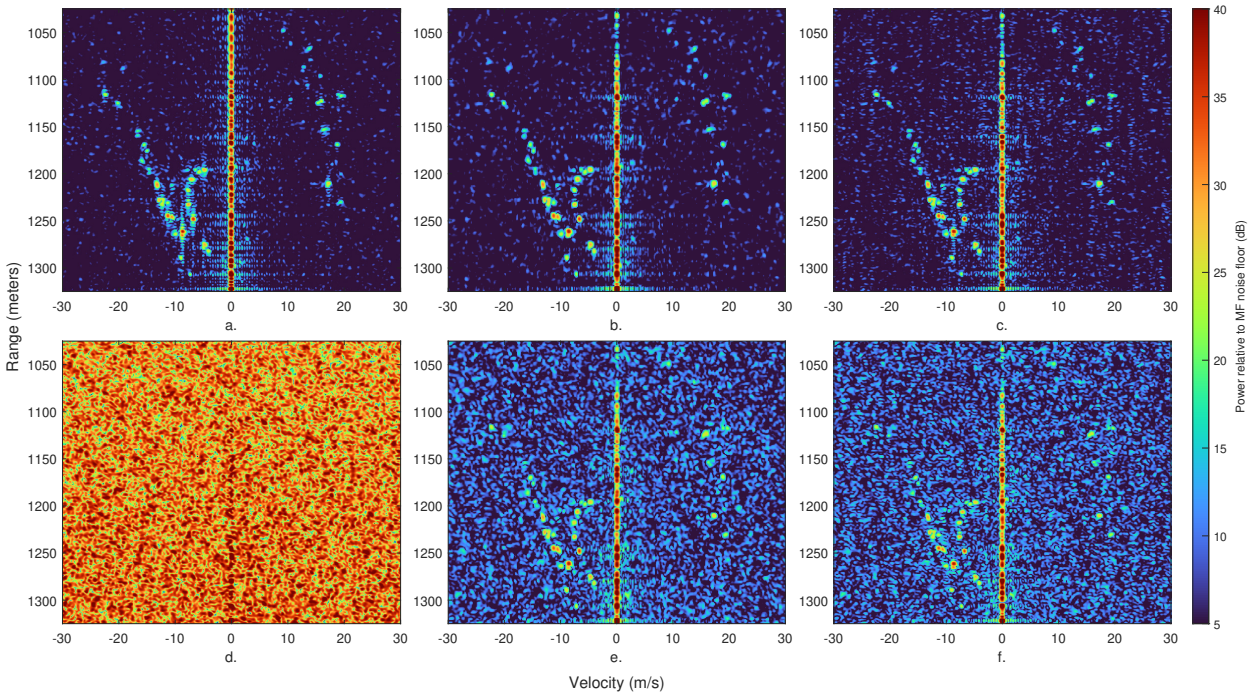

Figure 7.8: Toeplitz Model D=4N, ρ=-2dB, k=5 (ISL) or $\kappa_{\mathbf{w}} = 1.3$ (template) with no interference

using the MF (a., c.) and optimal SINR filter (b., d.) for the LFM (top row a.-b.) and the PROFM (bottom row c.d.) CPIs respectively. For the LFM MF (a.), the additive interference source leads to a raised background floor which masks most of the movers. After applying the optimal SINR filter (b.) the movers are again visible, though they have an extended range response due to the spectral notching behavior seen in Fig. 7.6 and noise floor has raised slightly. The application of the optimal SINR filters to the PROFM measurements (d.) does not provide much benefit relative to the MF (c.) as the direct-path RSM is unaffected by the optimal filter. Figure 7.10 provides a histograms of the intersection background floor for the MF when no interference is present (left) as well as a comparison between the MFs and optimal SINR filters when interference is present (right). Here it is clear that the optimal SINR filters are able to largely retain the performance of the MF with no interference present for this OFDM interferer.

Figure 7.11 then provides a visualization the range-Doppler maps of the intersection after the QC MMFs are applied where the top row of subplots (a.-c.) correspond to the LFM optimal SINR filter, ISL MMF, and template MMF respectively while the bottom row (d.-f.) correspond to

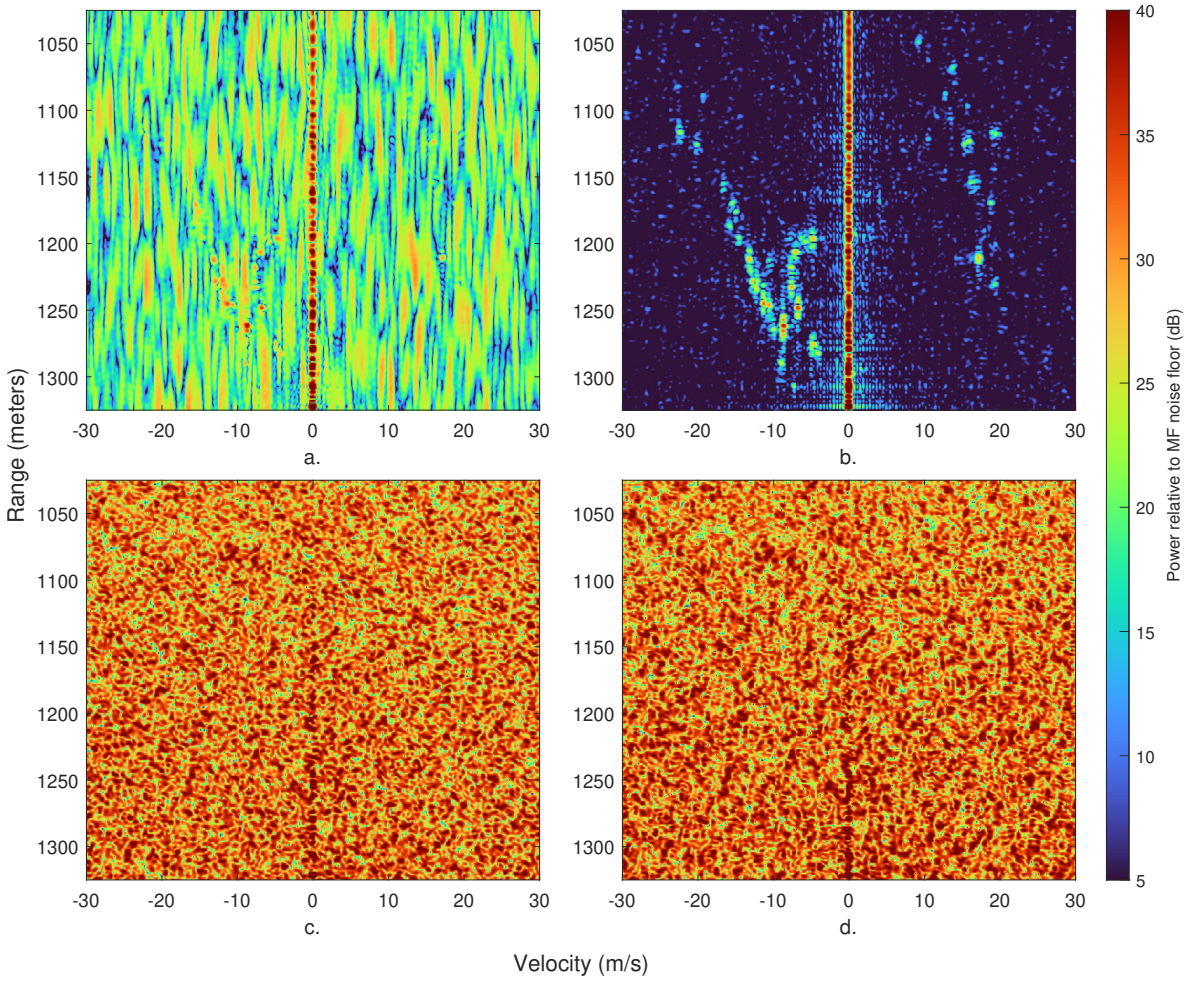

Figure 7.9: Matched Filter vs Optimal SINR filters for OFDM interference with 20dB INR

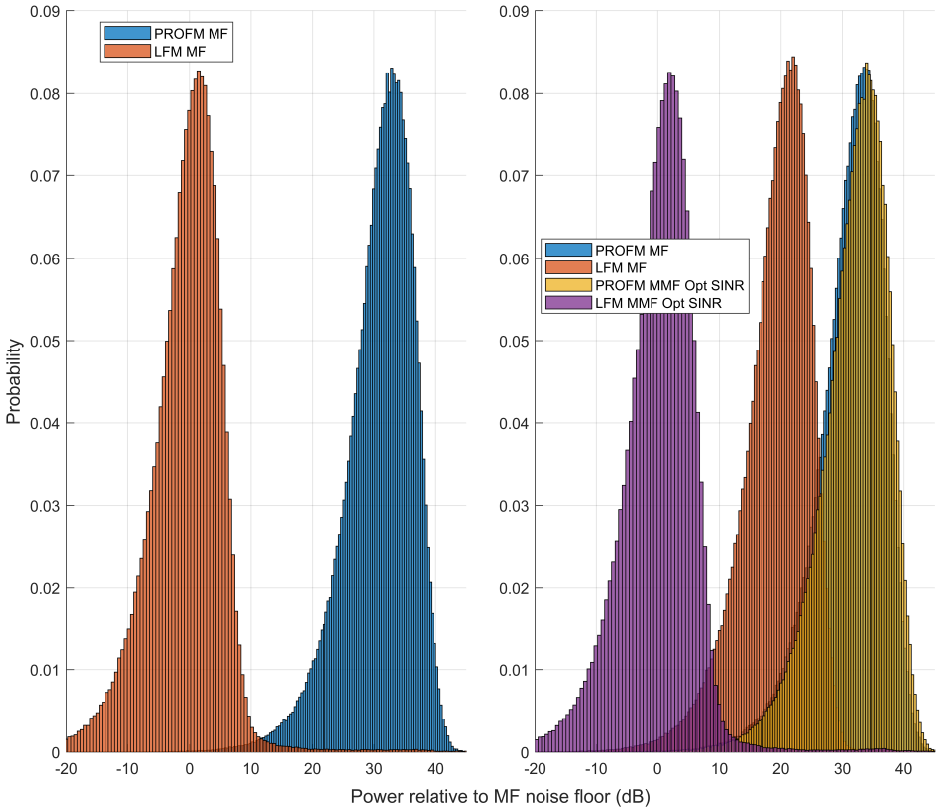

Figure 7.10: Histogram of intersection for PROFM and LFM Waveform sets with (right) and without (left) interference

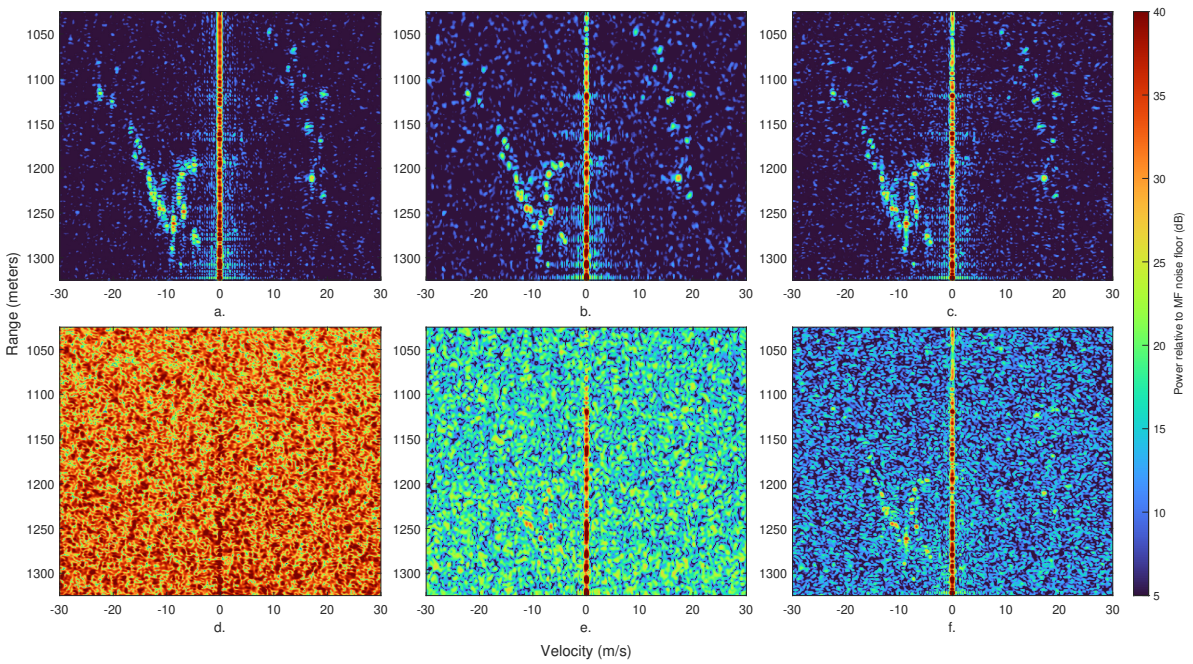

Figure 7.11: Toeplitz Model D=4N, ρ=-2dB, k=5 (ISL) or $\kappa_{\mathbf{w}} = 1.3$ (template), OFDM interference with 20dB INR

same respective the PROFM MMFs. Here the LFM MMFs alleviate much of the extended range response, though some near-in sidelobes remain. The PROFM MMFs again provide a significant reduction in RSM relative to the optimal SINR filter and matched filter where the background floor has been reduced by $\sim$15 dB and $\sim$20 dB for the ISL and template objectives respectively. While the improved performance of the template objective may seem unintuitive at first since the spectral notch is diametrically opposed to the Gaussian template, it can be explained through a philosophy similar to [156]. Where the ISL objective seeks to find the MMF that has the lowest sidelobes for every pulse, the template objective tries to match all filter responses to the same template. When incorporating an interference covariance, this desired response effectively changes to account for the frequency characteristics of the interference, but the per-pulse desired response largely remains the same. Thus, the template MMFs have a response with range-sidelobes that are higher (due to the spectral notch) but are correlated on a per-pulse basis. This correlation then results in a reduction of RSM as the pulse-to-pulse range-sidelobes coherently combine when performing Doppler processing.

Toeplitz MMFs - median of intersection (dB)																		
	MF	Opt SINR			Template Objective							ISL Objective						
D	N	N	$2N$	$4N$	N	$2N$	$4N$	$4N$	$4N$	$4N$	$4N$	N	$2N$	$4N$	$4N$	$4N$	$4N$	$4N$
ρ (dB)	0	0	0	0	-2	-2	-3	-1	-2	-2	-2	-2	-2	-3	-1	-2	-2	-2
Beamspoiling	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=2	κ_w=3	k=5	k=5	k=5	k=5	k=5	k=3	k=7
PROFM	31.66	31.66	31.66	31.66	30.16	22.43	7.03	7.03	7.03	6.69	7.41	30.16	22.39	6.76	6.70	6.73	11.72	7.34
LFM	0.12	0.12	0.12	0.12	1.41	1.61	2.27	0.81	1.60	1.00	1.75	0.83	0.77	0.78	0.76	0.77	1.55	1.94
PROFM w intf	32.05	32.79	31.52	31.42	30.85	23.24	9.22	12.74	9.61	10.19	15.59	33.68	26.58	13.99	19.91	15.92	13.30	10.57
LFM w intf	20.06	0.70	0.56	0.56	1.36	1.70	2.45	0.84	1.57	1.37	1.82	2.10	2.21	2.97	1.16	2.06	1.82	2.37
PROFM w tr	31.66	31.66	31.66	31.66	30.16	22.43	14.70	14.70	14.70	14.65	15.50	30.16	22.39	14.73	14.69	14.71	16.50	14.67
LFM w tr	0.12	0.12	0.12	0.12	1.41	1.61	2.27	0.82	1.61	1.00	1.75	0.83	0.77	0.78	0.76	0.77	1.55	1.93
PROFM w intf + tr	32.05	32.79	31.52	31.47	30.85	23.24	15.70	17.08	15.87	16.13	18.82	33.68	26.58	20.07	21.87	20.19	17.50	18.69
LFM w intf + tr	20.06	0.70	0.56	0.56	1.36	1.70	2.44	0.84	1.57	1.37	1.82	2.10	2.21	2.97	1.16	2.06	1.82	2.37

Table 7.5: Toeplitz MMFs - median of intersection (dB)

To give a sense of MMF performance vs parameter selection, Table 7.5 was generated using the median of intersection power to provide an ordered statistical estimate of the background floor. In the table, the various filter parameters are tested with additive interference (denoted with the intf label) as well as with a simulated front-end transmit-receive (TR) switch in which the first 1% of samples are zeroed out ($\sim$10% of the direct path) each PRI (denoted with the tr label). In practice, a TR switch would likely zero out the entire direct path to avoid transmitting and receiving simultaneously. However, this experimental test setup is largely SNR limited (output power $\sim$1Watt) and the close proximity of the intersection causes a more representative implementation to zero out a significant portion of the mover energy. Thus, while not a perfect representation of reality, this TR implementation to at least provides a semblance MMF degradation when partial pulse-eclipsing is present.

As a final experiment, the efficacy of the Circulant approximate MMFs in the SINR QC context was examined. First, the experiment from Fig. 7.11 was recreated using the same SINR loss constraint ($\rho = -2$ dB) and beamspoiling/template parameters, however, the filters were formed using the IIR approximate Circulant MMFs where the filter length was equal to the number of samples in a PRI ($D = L_{PRI}$). This approximation was performed for both the QC MMFs and the optimal SINR filters. Figure 7.12 illustrates these results with the same ordering as Fig. 7.11. Here the Circulant MMFs largely retain or slightly outperform the Toeplitz MMFs as $L_{PRI} > 4N$ and there are no significant model errors from pulse eclipsing or multiple-time-around-clutter for this PRF an transmit power. Interestingly, the Circulant approximate optimal SINR filter resulted in a surprising reduction in RSM by $\sim$2 dB, which may be an artifact of the long filter length.

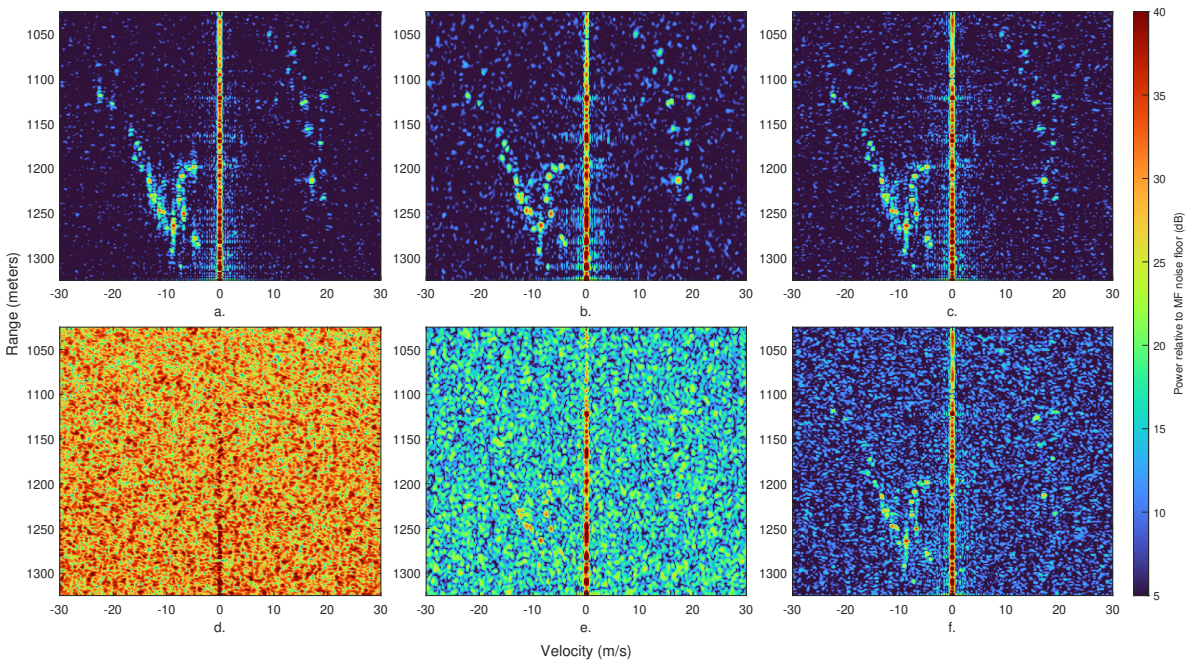

Figure 7.12: Approximate Circulant Model D=L_{PRI}, ρ=-2dB, k=3 (ISL) or **p**-gaussian shape, OFDM interference with 20dB INR

Circulant Approximate MMFs - median of intersection (dB)																		
	MF	Opt SINR			Template Objective							ISL Objective						
D	N	$0.5L_{PRI}$	$1.5L_{PRI}$	L_{PRI}	$0.5L_{PRI}$	$1.5L_{PRI}$	L_{PRI}	L_{PRI}	L_{PRI}	L_{PRI}	L_{PRI}	$0.5L_{PRI}$	$1.5L_{PRI}$	L_{PRI}	L_{PRI}	L_{PRI}	L_{PRI}	L_{PRI}
ρ (dB)	0	0	0	0	-2	-2	-3	-1	-2	-2	-2	-2	-2	-3	-1	-2	-2	-2
Beamspoiling	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=1.3	κ_w=2	κ_w=3	k=5	k=5	k=5	k=5	k=5	k=3	k=7
PROFM	31.66	31.66	31.66	31.66	3.65	3.58	3.55	3.54	3.55	3.48	4.44	3.81	3.77	3.81	3.65	3.74	10.27	4.30
LFM	0.12	0.12	0.12	0.12	1.64	1.64	2.05	0.83	1.64	0.92	1.75	0.77	0.77	0.77	0.76	0.77	1.56	1.47
PROFM w intf	32.03	43.07	30.84	31.31	30.32	6.26	6.43	8.38	6.43	6.56	10.17	30.37	14.71	11.28	18.18	13.89	11.69	6.70
LFM w intf	20.06	20.29	0.46	0.47	1.59	1.91	2.84	1.09	1.96	1.94	2.23	1.65	1.98	3.15	1.35	2.30	1.93	2.40
PROFM w tr	31.66	31.66	31.66	31.66	14.89	14.36	14.26	14.22	14.24	14.22	15.04	14.78	14.48	14.48	14.26	14.38	16.03	14.84
LFM w tr	0.12	0.12	0.12	0.12	1.64	1.64	2.05	0.83	1.64	0.92	1.75	0.77	0.77	0.77	0.76	0.77	1.56	1.47
PROFM w intf + tr	32.03	43.20	30.91	31.38	30.45	15.11	15.09	15.64	15.15	15.26	16.75	30.51	19.59	19.50	20.84	19.52	16.88	18.17
LFM w intf + tr	20.06	20.43	0.46	0.47	1.59	1.91	2.85	1.09	1.96	1.95	2.22	1.66	1.98	3.15	1.35	2.30	1.93	2.40

Table 7.6: Circulant Approximate MMFs - median of intersection (dB)

A tabulated performance evaluation of these Circulant approximations is provided in Table 7.6, where alternate filter lengths correspond to truncations or zero-padding of the PRI.

Chapter 8

Conclusions

In this work, various extensions and efficient implementations for structure-based estimation have been proposed to support waveform-agile MTI pulsed-radar operation. Algorithm efficacy has been tested in simulation, hardware-in-the-loop, and open-air experimental measurements and shown to generally enhance practical performance and robustness. Theoretical computational complexities have been provided, where the efficient implementations offer significant reduction in cost. Additionally, proposed iterative frameworks offer reductions in storage requirements and have the potential for hardware acceleration as the reformulated implementations only require vector-vector and FFT/IFFT operations. Below is an itemized summary of contributions from this work for each corresponding chapter.

Chapter 3 Contributions: Efficient Straddle-Robust Least Squares Mismatched Filtering

◇ Proposed and developed range straddling robustness measures for the LS MMF ISL QP

· Incorporation of straddling uncertainty yields ISL MMFs that avoid high frequency components

· On-grid samples are commensurate with continuum of possible sub-sample shifts

· Filter response continuum yields significantly lower sidelobes near zero-delay than the uncompensated ISL MMFs and does not suffer from the sinc-like off-grid ringing effect

· Resulting filters have a more restricted solution space, slightly raising on-grid peak-

sidelobes level ($\sim$1-2dB) but significantly reducing mainlobe modulation effect when using a waveform agile CPI

◇ Assessed viability and complexity of efficient MMF solvers

· Showed both the uncompensated and straddle compensated LS MMFs have Toeplitz structure and can be constructed and represented efficiently through summation across the diagonals of a signal-correlation matrix

· Inverse Toeplitz identity offers rapid beamspoiling updates and performs matrix-vector multiplications via elementwise multiplication and FFTs/IFFTs

· Demonstrated that the straddle compensated signal-correlation matrix can be well approximated as low-rank, enabling efficient beamspoiling updates

· Levinson recursion based implementations offer improved complexity over direct solvers

· Banded Cholesky approaches offer further complexity for filters that are much longer than the respective waveform

· Assessed preconditioners for conjugate gradient implementations which offer further improvements to computational complexity and storage requirements

· Additionally, as each PCG iteration is performed with vector-vector multiplications and FFTs, it is amenable to further parallelization and hardware acceleration that is not possible with the Levinson or Cholesky approaches

◇ Experimentally validated all approaches on hardware-in-the-loop and open-air measurements

· All LS MMFs show improved direct path RSM, however, the longer filter length extends the response in range and results in a higher intersection RSM floor

· Straddle compensated ISL and template-based LS MMF show significant improvement in mainlobe modulation and improved direct path roll-off

Chapter 4 Contributions: Mismatched Complementary on Receive Filtering

◇ Proposed the Mismatched Complementary on Receive Filtering framework for arbitrary FM waveforms

· Leverages the variable sidelobes structure of unique waveforms within a subset to design filters that achieve a complementary sidelobe cancellation after presuming is performed

· Accounts for spectral oversampling by incorporates beamspoiling or template-based objective

· Trades off Doppler sensitivity and unambiguous velocity to significantly increase design DoF

◇ Generalized straddling uncertainty for MiCRFt framework

· Filter response continuum yields significantly lower sidelobes near zero-delay than the uncompensated ISL MMFs and does not suffer from the sinc-like off-grid ringing effect

· Resulting filters have a more restricted solution space, yielding a broadened mainlobe response but significantly reducing mainlobe modulation effect when using a waveform agile CPI

◇ Assessed viability and complexity of efficient MiCRFt solvers

· Showed both the uncompensated and straddle compensated MiCRFt forms have a Toeplitz-block structure and can be constructed and represented efficiently through summation across the diagonals of a signal-correlation and signal-cross-correlation matrix

· Subsequent permutation of MiCRFt matrix yields Block-Toeplitz matrix structure

· Inverse Block-Toeplitz identity offers efficient beamspoiling updates and performs matrix-vector multiplications via multiplication with the smaller block-components and block FFTs/IFFTs

· Multi-channel Levinson recursion based implementations offer improved complexity over direct solvers

· Banded Cholesky approaches offer further complexity for filters that are much longer than the respective waveform

· Reposed MiCRFt filter form offers potential to further reduce cost when the number of waveforms within a subset is large ($Q > 4$)

· Assessed viability of conjugate gradient implementations which offer improvements to storage requirements but may need further developments to improve convergence

· That being said, the PCG iterations are still inexpensive to apply and leverage FFTs, thus, for some cases it may still outperform the Levinson or Cholesky approaches

◇ Experimentally validated all approaches on hardware-in-the-loop and open-air measurements

· MiCRFt approaches showed suppression of RSM to the model uncertainty floor for $\sim$1.5dB of mismatch loss, but halves the unambiguous velocity

Chapter 5 Contributions: Alternate Joint Range-Doppler Clutter Cancellation Representations and Efficient Implementations

◇ Computationally Efficient NIMPC Implementations

· Proposed alternative NIMPC filter solution to reduce computational complexity and memory requirements when the number of clutter slices is less than the number of pulses

· Alternative representation supports both multichannel levinson and banded Cholesky solvers

· Proposed reduced complexity method to solve for the NIMPC unity gain weights

· Proposed method to solve for NIMPC uniform noise weights and evaluated when this approach provides a complexity reduction

· Examined the efficacy of PCG for NIMPC and found that it primarily serves as a method to reduce memory requirements, and not necessarily improve computational cost for

the preconditioners evaluated

◇ Projection NIMPC Implementation

· Reposed NIMPC in a null-space projection framework

· Computational cost of Proj-NIMPC scales with the number of clutter directions rather than number of waveforms, and does not require computing and storing multiple filters

· The required linear system is Toeplitz-Block-Toeplitz and can be solved efficiently with multichannel levinson or banded Cholesky

· Proj-NIMPC is extremely amenable to PCG and inexpensive preconditioners can exceptional RSM suppression for only a few (5-10) iterations

◇ Evaluated the reformulated NIMPC and Proj-NIMPC using open-air experimental measurements

· All the NIMPC methods were able to suppress clutter RSM to the model uncertainty floor

· Reformulated NIMPC provided nearly ten orders of magnitude reduction in complexity relative to a brute force solve, and three orders of magnitude improvement over an efficient solver

· The Proj-NIMPC framework pushed the computational cost even lower with an additional two orders of magnitude improvement over the Reformulated NIMPC approaches

· PCG for Proj-NIMPC saw overall the lowest order of complexity and likely will see even more practical speedups over the multichannel levinson approach as it performs most operations elementwise or utilizing FFTs/IFFTs

Chapter 6 Contributions: Robust Adaptive Pulse Compression

◇ Proposed and developed measurement-based range straddling robustness measures for APC

· Incorporation of straddling uncertainty yields estimates that avoid sinc-like broadening effect seen when APC is applied on near critically sampled ($\kappa \approx 1$) architectures

· Additionally, provides robustness towards mainlobe modulation effect seen when APC is applied to waveform agile radar

◇ Generalized an APC filter beamspoiling technique to yield an arbitrary template-based RMMSE framework

· Like the template-based LS MMFs, this framework helps mitigate super-resolution induced MML and provides natural robustness to straddling

· Additionally, this template-based objective avoids additional covariance beamspoiling and unity constraint operations, which leads to more computationally efficient implementation

· However, the lack of a traditional unity constraint can still produce significant noise floor fluctuation and can produce an effect similar to Doppler streaking from mainlobe modulation

◇ Proposed Low-Rank approximate APC implementation

· Approach attempts to avoid incorporating redundant or uninformative information into the source covariance estimate

· The basic concept couples generic CFAR detection of scatterers, with the iterative refinement of APC

· Implementation of this work leverages an simple inexpensive moving median threshold detection routine, but likely can be generalized and improved with more sophisticated approaches

◇ Further evaluated the "block" or "full" APC implementations and developed efficient implementations for certain configurations

· Leveraged the "full" structured covariance to effeciently form every snapshot-based structured covariance simultaniously

· Demonstrated that construction of the structured covariance could be performed via diagonals convolution of a signal correlation matrix and the source covariance matrix, which significantly reduces computational burden

· Using the template-based RMMSE framework, a "full" APC iteration can be performed by solving a single banded hermitian linear system, and subsequent matrix operations can be performed efficiently by leveraging Toeplitz-vector frequency domain multiplication

· Alternatively, by leveraging the low-rank approximate implementation enables efficient filter solutions using low-rank factorizations or Woodbury's inverse identity

◇ Proposed APC techniques were then evaluated on open-air experimental measurements from a waveform agile CPI

· The straddling robust implementation of APC saw RSM suppression down to the model uncertainty floor, without the required dimensionality tradeoffs of MiCRFt and NIMPC

· Additionally, incorporating straddling robustness significantly reduced mainlobe modulation effects present with previous APC configurations

· The template-based APC implementation provided source estimates commensurate with the other approaches, however, the lack of a unity constraint induced more noise gain and Doppler streaking from high powered scattering

· The low-rank approximate implementation offered significant improvement in RSM levels over the MF when only incorporating 112 outliers which corresponded to a two orders of magnitude decrease to computational complexity

· However, this low-rank implementation could not fully suppress RSM to the same level as the "full-rank" approaches and the RSM background floor is 2dB higher

Chapter 7 Contributions: Generalized SINR Loss Constrained LS MMFs

◇ Proposed a generalized SINR Constrained QP for LS MMF formulation

· Approach builds off of previous work which examined a SNR constraint ISL LS MMF

· Here generalizations for a template based objective and an SINR constraint are introduced

· Generic framework provides the means to solve for the globally optimal constrained LS MMF via the Lagrange Dual problem

· Explicit bounds for the Dual are then provided via a generalized eigen decomposition, but in practice this expensive decomposition is not necessary to solve the optimal filter

· The Dual, its derivative, and second Derivative are then used to develop a maximization algorithm that yields the optimal MMF

· Extensions of the QCQP and dual maximization to account for aggregate SINR loss across multiple waveform and a QCQP implementations of MiCRFt are then discussed, but further analysis and development is left as future work

◇ Efficient Dual Maximization Contributions

· The efficient implementations from previous works (for the SNR constrained ISL objective) are generalized for this SINR constrained framework

· An rapid upper bound estimation technique for the template-based objective is developed to avoid performing an expensive generalized eigen decomposition

· Efficient direct solver routines that leverage levinson recursion and banded Cholesky are discussed

· An efficient and parallelizable PCG implementation is then developed where the solver leverages a novel upper-lower bound preconditioner that provides rapid and accelerating convergence

◇ Experimental evaluation of SINR constrained LS MMFs

· Template-based and ISL LS MMFs were both designed using the proposed maximization algorithm and efficient solvers using hardware in the loop measurements and were shown to offer an improvement in complexity two orders of magnitude lower than eigen decomposition based approach

· The proposed preconditioner likewise outperformed both Circulant preconditioners and efficient Levinson based implementation

· This preconditioner may offer further practical speedups as it converged in significantly fewer iterations than the circulant preconditioners and is likewise only requires Toeplitz matrix-vector multiplications that can be computed rapidly using FFTs/IFFTs

· Open-air experimental validation confirms that trading off minor SINR loss can yield a significant reduction in range sidelobes and reveal scatterers that were previously masked by RSM

Appendix A

Appendix

A.1 Toeplitz Plus Low Rank Linear Systems

A rank-r update to the inverse of matrix $\mathbf{A}$ (or linear system solve $\mathbf{Ad} = \mathbf{x}$) may be incorporated via the Sherman–Morrison–Woodbury (or Woodbury's for shorthand) formula

$$\begin{aligned} &\left(\mathbf{A}+\mathbf{U}\Lambda\mathbf{U}^H\right)^{-1}\mathbf{x} = \\ &\mathbf{A}^{-1}\mathbf{x}-\mathbf{A}^{-1}\mathbf{U}\left(\Lambda^{-1}+\mathbf{U}^H\mathbf{A}^{-1}\mathbf{U}\right)^{-1}\mathbf{U}^H\mathbf{A}^{-1}\mathbf{x} \end{aligned} \tag{A.1}$$

for an "update" (i.e. incorporating low rank interference $\mathbf{R}_{\mathrm{i}} = \mathbf{U}\Lambda\mathbf{U}^H$) or

$$\begin{aligned} &\left(\mathbf{A}-\mathbf{U}\mathbf{U}^H\right)^{-1}\mathbf{x} = \\ &\mathbf{A}^{-1}\mathbf{x}+\mathbf{A}^{-1}\mathbf{U}\left(\mathbf{I}-\mathbf{U}^H\mathbf{A}^{-1}\mathbf{U}\right)^{-1}\mathbf{U}^H\mathbf{A}^{-1}\mathbf{x} \end{aligned} \tag{A.2}$$

for a "downdate" (e.g., beamspoiling $\mathbf{U}\mathbf{U}^H = \tilde{\mathbf{S}}^H\tilde{\mathbf{S}}$), though they are functionally equivalent. For PD hermitian Toeplitz matrices, matrix-vector multiplication with $\mathbf{A}^{-1}$ can be applied with $\mathscr{O}(8L\log(2L))$ and the $r \times r$ "capacitance" matrix $\left(\ \left(\Lambda^{-1}+\mathbf{U}^H\mathbf{A}^{-1}\mathbf{U}\right)$ or $\left(\mathbf{I}-\mathbf{U}^H\mathbf{A}^{-1}\mathbf{U}\right)\ \right)$ is PD. Thus for the Toeplitz case, the order of complexity for both forms is $\mathscr{O}\left(16L\log(2L)+r^2L\right)$.

In theory, the inverse update will be accurate as long as the capacitance matrix is PD, which is always the case for (7.17) for any non-zero λ in the feasible region. In practice however, numerical precision and ill conditioning often cause instability when using (A.1) and (A.2) [157]. A heuristic rule of thumb that helps ensure the update is stable, is to compute the MMF with $\sim 2\times$ the numerical precision of the waveform $\mathbf{s}$, but in general this may not always be practical.

Alternatively, an iterative update procedure via PCG or preconditioned generalized minimum residual (PGMRES), where $\mathbf{A}^{-1}$ acts is treated as the preconditioner, may offer some robustness towards numeric instability. Both approaches have (theoretically) guaranteed convergence in at most $r+1$ steps as the preconditioned system has r non unity positive eigenvalues and at most $r+1$ clusters of eigenvalues, though iterations may be necessary in practice for poorly conditioned systems. The choice of PGMRES vs PCG largely stems from how well the capacitance matrix is conditioned, where PGMRES will likely offer the more robustness. Neither the order of complexity for PGMRES $\mathscr{O}\left(\frac{5}{2}(r+1)^2(r+2)^2L\log(2L)+r(r+1)L\right)$ nor PCG $\mathscr{O}\left(10(r+1)L\log(2L)+r(r+1)L\right)$ will significantly increase computational burden when $r<<L$, but they will add additional latency.

A.2 Inverse of the KKT matrix

As the KKT matrix has a block form, the block matrix inverse identity

$$\mathbf{D}^{-1}=\begin{bmatrix}\mathbf{T} & \mathbf{v}\\ \mathbf{v}^H & 0\end{bmatrix}^{-1}=\begin{bmatrix}\mathbf{A} & \mathbf{g}\\ \mathbf{g}^H & b\end{bmatrix} \tag{A.3}$$

where

$$\mathbf{A}=\mathbf{T}^{-1}-\mathbf{T}^{-1}\mathbf{v}\left(\mathbf{v}^H\mathbf{T}^{-1}\mathbf{v}\right)^{-1}\mathbf{v}^H\mathbf{T}^{-1}, \tag{A.4}$$

$$\mathbf{g}=\mathbf{T}^{-1}\mathbf{v}\left(\mathbf{v}^H\mathbf{T}^{-1}\mathbf{v}\right)^{-1}, \tag{A.5}$$

and

$$b=-\left(\mathbf{v}^H\mathbf{T}^{-1}\mathbf{v}\right)^{-1} \tag{A.6}$$

can be used to find the stationary points of the Lagrangian or simplify the QCQP Dual function [76, p. 456].

A.3 Algebraic Simplifications for the ISL Dual

For the ISL dual function, $\mathbf{G}^{-1}\mathbf{b}$ only requires computation of the second block column of (A.3) as

$$\begin{aligned}\mathbf{G}^{-1}\mathbf{b} &= (\mathbf{D}+\lambda\mathbf{Q})^{-1}\,\mathbf{b} \\ &= \begin{bmatrix} \tilde{\mathbf{T}}+\lambda\mathbf{C} & \mathbf{v} \\ \mathbf{v}^H & 0 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} \frac{\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}} \\ \frac{1}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}} \end{bmatrix}\end{aligned} \tag{A.7}$$

where the vector with $\mathbf{b}^H$ further simplifies the dual as

$$\mathbf{b}^H\mathbf{G}^{-1}\mathbf{b} = \frac{1}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}}. \tag{A.8}$$

Likewise, multiplication by $\mathbf{Q}$ allows for the simple expansions of

$$\mathbf{Q}\mathbf{G}^{-1}\mathbf{b} = \begin{bmatrix} \frac{\mathbf{C}\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}}{\mathbf{v}^H\left(\tilde{\mathbf{T}}+\lambda\mathbf{C}\right)^{-1}\mathbf{v}} \\ 0 \end{bmatrix} \tag{A.9}$$

which can be used to calculate the first derivative and reveals that the second derivative only requires an additional multiplication with the first block element of (A.3).

References

[1] S. D. Blunt, J. K. Jakabosky, C. A. Mohr, P. M. McCormick, J. W. Owen, B. Ravenscroft, C. Sahin, G. D. Zook, C. C. Jones, J. G. Metcalf, and T. Higgins, "Principles and applications of random fm radar waveform design," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 10, pp. 20–28, 2020.

[2] M. Richards, J. Scheer, J. Scheer, and W. Holm, *Principles of Modern Radar: Basic Principles, Volume 1*, ser. Electromagnetics and Radar. Institution of Engineering and Technology, 2010.

[3] M. Skolnik, *Radar Handbook, Third Edition*, ser. Electronics electrical engineering. McGraw-Hill Education, 2008.

[4] ——, *Introduction to Radar Systems*, ser. International student edition. McGraw-Hill, 1962.

[5] F. T. Ulaby, D. G. Long, and U. of Michigan. Press., *Microwave radar and radiometric remote sensing*. Ann Arbor: The University of Michigan Press, 2014.

[6] C. Lewis, H. Owen, D. Abi, J. Hecker and J. Sulzen, "Multi-waveform radar for ice sheet measurements and classroom demonstration," *2007 IEEE International Geoscience and Remote Sensing Symposium*, Barcelona, Spain, 2007.

[7] W. Melvin and J. Scheer, *Principles of Modern Radar: Radar Applications, Volume 3*, ser. Electromagnetics and Radar. Institution of Engineering and Technology, 2013.

[8] H. Griffiths, L. Cohen, S. Watts, E. Mokole, C. Baker, M. Wicks, and S. Blunt, "Radar spectrum engineering and management: Technical and regulatory issues," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 85–102, 2015.

[9] A. F. Martone, K. I. Ranney, K. Sherbondy, K. A. Gallagher, and S. D. Blunt, "Spectrum allocation for noncooperative radar coexistence," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 1, pp. 90–105, 2018.

[10] M. Labib, V. Marojevic, A. F. Martone, J. H. Reed, and A. I. Zaghloui, "Coexistence between communications and radar systems: A survey," *URSI Radio Science Bulletin*, vol. 2017, no. 362, pp. 74–82, 2017.

[11] S. D. Blunt and E. L. Mokole, "Overview of radar waveform diversity," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 11, pp. 2–42, 2016.

[12] M. Wicks, E. Mokole, S. D. Blunt, V. Amuso, and R. Schneible, eds., *Principles of Waveform Diversity & Design*, Raleigh, NC: SciTech Publishing, 2010.

[13] S. Pillai, K. Y. Li, I. Selesnick, and B. Himed, *Waveform Diversity: Theory & Applications*, New York, NY, USA: McGraw-Hill, 2011.

[14] F. Gini, A. De Maio, and L. K. Patton, *Waveform Design and Diversity for Advanced Radar Systems*, London, UK: IET Press, 2012.

[15] F. Liu, C. Masouros, A. P. Petropulu, H. Griffiths, and L. Hanzo, "Joint radar and communication design: Applications, state-of-the-art, and the road ahead," *IEEE Transactions on Communications*, vol. 68, no. 6, pp. 3834–3862, 2020.

[16] B. Ravenscroft, P. M. McCormick, S. Blunt, E. S. Perrins, C. Sahin and J. G. Metcalf, "Experimental Assessment of Tandem-Hopped Radar and Communications (THoRaCs)," *2019 International Radar Conference (RADAR)*, Toulon, France, 2019.

[17] S. D. Blunt, M. R. Cook and J. Stiles, "Embedding information into radar emissions via waveform implementation," *2010 International Waveform Diversity and Design Conference*, Niagara Falls, ON, Canada, 2010.

[18] C. Sahin, J. Jakabosky, P. M. McCormick, J. G. Metcalf and S. D. Blunt, "A novel approach for embedding communication symbols into physical radar waveforms," *2017 IEEE Radar Conference (RadarConf)*, Seattle, WA, USA, 2017.

[19] P. M. McCormick, S. D. Blunt and J. G. Metcalf, "Simultaneous radar and communications emissions from a common aperture, Part I: Theory," *2017 IEEE Radar Conference (RadarConf)*, Seattle, WA, USA, 2017.

[20] D. Ciuonzo, A. De Maio, G. Foglia, and M. Piezzo. Intrapulse radarembedded communications via multiobjective optimization. *IEEE Transactions on Aerospace & Electronic Systems*, 51, 4, 2960–2974 (Oct. 2015).

[21] A. Hassanien, M. G. Amin, Y. D. Zhang, and F. Ahmad, “Phase-modulation based dual-function radar-communications,” *IET Radar, Sonar & Navigation*, vol. 10, no. 8, pp. 1411–1421, 2016.

[22] A. Hassanien, M. G. Amin, Y. D. Zhang and F. Ahmad, "Dual-Function Radar-Communications: Information Embedding Using Sidelobe Control and Waveform Diversity," in *IEEE Transactions on Signal Processing*, vol. 64, no. 8, pp. 2168-2181, April15, 2016.

[23] S. Blunt and E. Perrins, *Radar and Communication Spectrum Sharing*. Institution of Engineering and Technology, 2018.

[24] P. S. Tan, J. M. Stiles and S. D. Blunt, "Optimizing sparse allocation for radar spectrum sharing," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[25] J. Owen, C. Mohr, B. Ravenscroft, S. Blunt, B. Kirk and A. Martone, "Real-Time Experimental Demonstration and Evaluation of Open-Air Sense-and-Notch Radar," *2022 IEEE Radar Conference (RadarConf22)*, New York City, NY, USA, 2022.

[26] B. H. Ferrell and W. C. Woody. Notched-spectrum swept-frequency generator and method therefore, U.S. Patent 5 852 418, Dec. 1998.

[27] K. Gerlach, M. R. Frey, M. J. Steiner, and A. Shackelford. Spectral nulling on transmit via nonlinear FM radar waveforms. *IEEE Transactions on Aerospace & Electronic Systems*, 47, 2, 1507–1515 (Apr. 2011).

[28] I. W. Selesnick and S. U. Pillai. Chirp-like transmit waveforms with multiple frequency-notches. *Proceedings of the IEEE Radar Conference*, Kansas City, MO, May 2011.

[29] L. K. Patton, C. A. Bryant, and B. Himed. Radar-centric design of waveforms with disjoint spectral support. *Proceedings of the IEEE Radar Conference*, Atlanta, GA, May 2012.

[30] C. Nunn and L. R. Moyer. Spectrally-compliant waveforms for wideband radar. *IEEE Aerospace & Electronic Systems Magazine*, 27, 8, 11–15 (Aug. 2012).

[31] A. Aubry, A. De Maio, Y. Huang, M. Piezzo, and A. Farina. A new radar waveform design algorithm with improved feasibility for spectral coexistence. *IEEE Transactions on Aerospace & Electronic Systems*, 51, 2, 1029–1038 (Apr. 2015).

[32] J. Jakabosky, B. Ravenscroft, S. D. Blunt, and A. Martone. Gapped spectrum shaping for tandem-hopped radar/communications & cognitive sensing. In *Proceedings of the IEEE Radar Conference*, Philadelphia, PA, May 2016.

[33] A.F. Martone, J.A. Kovarskiy, C.E. Thornton, B.H. Kirk, J.W. Owen. B. Ravenscroft, A. Egbert, A. Goad, R.M. Buehrer, R.M. Narayanan, S.D. Blunt, C. Baylis, K.D. Sherbondy, "Closing the Loop on Cognitive Radar for Spectrum Sharing," *IEEE Aerospace & Electronic Systems Magazine*, vol. 36, no. 9, pp. 44-55, Sept. 2021.

[34] A. F. Martone et al., "Practical Aspects of Cognitive Radar," *2020 IEEE Radar Conference (RadarConf20)*, Florence, Italy, 2020.

[35] A. Martone and M. Amin, “A view on radar and communication systems coexistence and dual functionality in the era of spectrum sensing,” *Digital Signal Processing*, vol. 119, p. 103135, 2021.

[36] G. Zook, P. M. McCormick, S. D. Blunt, C. Allen and J. Jakabosky, "Dual-polarized FM noise radar," *International Conference on Radar Systems (Radar 2017)*, Belfast, 2017.

[37] C. A. Mohr, P. M. McCormick, C. A. Topliff, S. D. Blunt, and J. M. Baden, “Gradient-based optimization of pcfm radar waveforms,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 2, pp. 935–956, 2021.

[38] P.M. McCormick, S.D. Blunt, "Nonlinear Conjugate Gradient Optimization of Polyphase-Coded FM Radar Waveforms," *IEEE Radar Conference*, Seattle, WA, 8-12 May 2017.

[39] G. Coxson and J. Russo. Efficient exhaustive search for optimal-peaksidelobe binary codes. *IEEE Transactions on Aerospace & Electronic Systems*, 41, 1, 302–308 (Jan. 2005).

[40] C. J. Nunn and G. E. Coxson. Best-known autocorrelation peak sidelobe levels for binary codes of length 71 to 105. *IEEE Transactions on Aerospace & Electronic Systems*, 44, 1, 392–395 (Jan. 2008).

[41] M. N. Cohen, M. R. Fox, and J. M. Baden. Minimum peak sidelobe pulse compression codes. In *Proceedings of the IEEE International Radar Conference*, Arlington, VA, May 1990.

[42] M. Labitt. Obtaining low sidelobes using non-linear FM pulse compression. *MIT Lincoln Lab Project Report*, ATC-223 (Nov. 1994).

[43] D. Tarchi, K. Lukin, J. Fortuny-Guasch, A. Mogyla, P. Vyplavin, and A. Sieber. SAR imaging with noise radar. *IEEE Transactions on Aerospace & Electronic Systems*, 46, 3, 1214–1225 (July 2010).

[44] S. R. J. Axelsson. Random noise radar/sodar with ultrawideband waveforms. *IEEE Transactions on Geoscience & Remote Sensing*, 45, 5, 1099–1114 (May 2007).

[45] N. Levanon and E. Mozeson, *Radar Signals*. Wiley – IEEE Press, 2004.

[46] J. Jakabosky, S. D. Blunt and B. Himed, "Optimization of "over-coded" radar waveforms," *2014 IEEE Radar Conference*, Cincinnati, OH, USA, 2014.

[47] B. Ravenscroft, J. Owen, J. Jakabosky, S. D. Blunt, A. F. Martone, and K. D. Sherbondy, "Experimental demonstration and analysis of cognitive spectrum sensing and notching for radar," *IET Radar, Sonar & Navigation*, 2018.

[48] D. B. Herr, P. S. Raju and J. M. Stiles, "Adaptive-on-transmit using information theoretic measures," *International Conference on Radar Systems (RADAR 2022)*, Hybrid Conference, Edinburgh, UK, 2022.

[49] D.B. Herr, P.S. Raju, J.M. Stiles, "Information theoretic waveform design with applications to adaptive-on-transmit radar," *IET Radar, Sonar & Navigation*, pp. 1–13, 2023.

[50] S. D. Blunt, M. Cook, J. Jakabosky, J. De Graaf, and E. Perrins, "Polyphase-coded FM (PCFM) radar waveforms, part I: implementation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 3, pp. 2218–2229, 2014.

[51] N. Bleistein and R. Handelsman, *Asymptotic Expansions of Integrals*. Holt, Rinehart and Winston, 1975.

[52] S. D. Blunt, J. Jakabosky, M. Cook, J. Stiles, S. Seguin, and E. L. Mokole, "Polyphase-coded FM (PCFM) radar waveforms, part II: optimization," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 3, pp. 2230–2241, 2014.

[53] A. d'Aspremont, and S. Boyd, "Relaxations and Randomized Methods for Nonconvex QCQPs," Stanford University, 2003.

[54] M. H. Ackroyd and F. Ghani, “Optimum mismatched filters for sidelobe suppression,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-9, no. 2, pp. 214–218, 1973.

[55] D. Henke, P. McCormick, S. D. Blunt and T. Higgins, "Practical aspects of optimal mismatch filtering and adaptive pulse compression for FM waveforms," *2015 IEEE Radar Conference (RadarCon)*, Arlington, VA, USA, 2015.

[56] M. Coutino, F. Uysal and L. Anitori, "Waveform-Aware Optimal Window Function Design for Mismatched Filtering," *2022 IEEE Radar Conference (RadarConf22)*, New York City, NY, USA, 2022.

[57] S. D. Blunt and K. Gerlach, "A novel pulse compression scheme based on minimum mean-square error reiteration [radar signal processing]," *2003 Proceedings of the International Conference on Radar (IEEE Cat. No.03EX695)*, Adelaide, SA, Australia, 2003.

[58] S. D. Blunt and K. Gerlach, "Adaptive pulse compression," *Proceedings of the 2004 IEEE Radar Conference (IEEE Cat. No.04CH37509)*, Philadelphia, PA, USA, 2004.

[59] T. Higgins, K. Gerlach, A. K. Shackelford and S. D. Blunt, "Aspects of Non-Identical Multiple Pulse Compression," *2011 IEEE RadarCon (RADAR)*, Kansas City, MO, USA, 2011.

[60] D. P. Scholnik, "Range-ambiguous clutter suppression with Pulse-diverse waveforms," *2011 IEEE RadarCon (RADAR)*, Kansas City, MO, USA, 2011.

[61] T. Higgins, S. D. Blunt and A. K. Shackelford, "Time-Range Adaptive Processing for pulse agile radar," *2010 International Waveform Diversity and Design Conference*, Niagara Falls, ON, Canada, 2010.

[62] M. Roberts, *Fundamentals of Signals and Systems*. McGraw-Hill Higher Education, 2008.

[63] M. Albulet, *RF Power Amplifiers*, ser. Electromagnetic Waves. Institution of Engineering and Technology, 2001.

[64] D. Pozar, *Microwave and RF Design of Wireless Systems*. Wiley, 2000.

[65] J. Echard, “Estimation of radar detection and false alarm probability,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 27, no. 2, pp. 255–260, 1991.

[66] K. Shanmugan and A. Breipohl, *Random Signals: Detection, Estimation and Data Analysis*. Wiley, 1988.

[67] W. Melvin and J. Scheer, *Principles of Modern Radar: Advanced Techniques, Volume 2*, ser. EBSCO ebook academic collection. Institution of Engineering and Technology, 2012.

[68] L. Harnett et al., "Optimal and adaptive mismatch filtering for stretch processing," *2018 IEEE Radar Conference (RadarConf18)*, Oklahoma City, OK, USA, 2018.

[69] P. M. McCormick, C. Sahin, S. D. Blunt and J. G. Metcalf, "FMCW Implementation of Phase-Attached Radar-Communications (PARC)," *2019 IEEE Radar Conference (RadarConf)*, Boston, MA, USA, 2019.

[70] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2003.

[71] A. C. O’Connor, J. M. Kantor and J. Jakabosky, "Space-time adaptive mismatch processing," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[72] O. Rabaste and L. Savy, "Mismatched filter optimization for radar applications using quadratically constrained quadratic programs," in *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 4, pp. 3107-3122, Oct. 2015.

[73] J. Jakabosky, S. D. Blunt and B. Himed, "Spectral-shape optimized FM noise radar for pulse agility," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[74] M. Chatzitheodoridi, A. Taylor, O. Rabaste and H. Oriot, "Fast Computation of the Optimal Integrated Sidelobe Level Filter for Large-Scaled Problems," *2022 IEEE Radar Conference (RadarConf22)*, New York City, NY, USA, 2022.

[75] L. Jin, J. Wang, Y. Zhong, and D. Wang, “Optimal mismatched filter design by combining convex optimization with circular algorithm,” *IEEE Access*, vol. 10, pp. 56 763–56 772, 2022.

[76] J. Nocedal and S. Wright, *Numerical Optimization*, ser. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.

[77] J. M. Kempf, "MISMATCHED FILTER EFFECTS ON SYNTHETIC APERTURE RADAR IMAGE QUALITY METRICS," *Masters Thesis*, Air Force Institute Of Technology, 2020.

[78] T. Aittomäki and V. Koivunen, "Mismatched filter design for radar waveforms by semidefinite relaxation," *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, South Brisbane, QLD, Australia, 2015.

[79] Z. Wang, S. Shi, Z. He, and J. Hu, “Radio frequency interference suppression filters design for hf radar based on socp,” *The Journal of Engineering*, vol. 2019, no. 21, pp. 7885–7889, 2019.

[80] J. Park and S. Boyd, “General heuristics for nonconvex quadratically constrained quadratic programming,” 2017.

[81] M. B. Heintzelman, J. W. Owen, S. D. Blunt, B. Maio and E. D. Steinbach, "Practical Considerations for Optimal Mismatched Filtering of Nonrepeating Waveforms," *2023 IEEE Radar Conference (RadarConf23)*, San Antonio, TX, USA, 2023.

[82] S. Haykin, *Adaptive Filter Theory: International Edition*. Pearson Education, 2014.

[83] H. Van Trees, K. Bell, and Z. Tian, *Detection Estimation and Modulation Theory, Part I: Detection, Estimation, and Filtering Theory*, ser. Detection Estimation and Modulation Theory. Wiley, 2013, no. pt. 1.

[84] J. Ward, "Space-time adaptive processing for airborne radar," MIT Lincoln Laboratory Technical Report 1015, 1994.

[85] S. D. Blunt and T. Higgins, "Achieving Real-Time Efficiency for Adaptive Radar Pulse Compression," *2007 IEEE Radar Conference*, Waltham, MA, USA, 2007.

[86] T. Higgins, S. D. Blunt and K. Gerlach, "Gain-constrained adaptive pulse compression via an MVDR framework," *2009 IEEE Radar Conference*, Pasadena, CA, USA, 2009.

[87] S. D. Blunt, K. Gerlach and E. L. Mokole, "Pulse compression Eclipsing-Repair," *2008 IEEE Radar Conference*, Rome, Italy, 2008.

[88] S.D. Blunt, K. Gerlach, "Adaptive Pulse Compression via MMSE Estimation," *IEEE Trans. Aerospace & Electronic Systems*, vol. 42, no. 2, pp. 572-584, April 2006.

[89] T. D. Cuprak and K. E. Wage, "Efficient doppler-compensated reiterative minimum mean-squared-error processing," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 53, no. 2, pp. 562–574, 2017.

[90] S. D. Blunt, T. Chan, and K. Gerlach, "Robust DOA Estimation: The Reiterative Superresolution (RISR) Algorithm," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 1, pp. 332–346, 2011.

[91] C. C. Jones et al., "Development & Experimental Assessment of Robust Direction Finding and Self-Calibration," *2022 IEEE Radar Conference (RadarConf22)*, New York City, NY, USA, 2022.

[92] S. D. Blunt, K. J. Smith and K. Gerlach, "Doppler-compensated adaptive pulse compression," *2006 IEEE Conference on Radar*, Verona, NY, USA, 2006.

[93] D. G. Felton, C. C. Jones, D. B. Herr, L. A. Harnett, S. D. Blunt and C. T. Allen, "Experimental Demonstration of Single Pulse Imaging (SPI)," *2023 IEEE Radar Conference (RadarConf23)*, San Antonio, TX, USA, 2023.

[94] P. M. McCormick and S. D. Blunt, "Shared-spectrum multistatic radar: Experimental demonstration using FM waveforms," *2018 IEEE Radar Conference (RadarConf18)*, Oklahoma City, OK, USA, 2018.

[95] S.D. Blunt, K. Gerlach, "Multistatic Adaptive Pulse Compression," *IEEE Trans. Aerospace & Electronic Systems*, vol. 42, no. 3, pp. 1043-1057, July 2006.

[96] C.C. Jones, L.A. Harnett, C.A. Mohr, S.D. Blunt, C.T. Allen, "Structure-Based Adaptive Radar Processing for Joint Clutter Cancellation and Moving Target Estimation," *IEEE International Radar Conference*, Washington, DC, 27-30 Apr. 2020.

[97] L.A. Harnett, B. Ravenscroft, S.D. Blunt, C.T. Allen, "Experimental Evaluation of Adaptive Doppler Estimation for PRI-Staggered Radar," *IEEE Radar Conference*, New York City, NY, 21-25 Mar. 2022.

[98] D. Montgomery, *Applied Statistics and Probability for Engineers, 6th Edition*. John Wiley and Sons, Incorporated, 2013.

[99] S.D. Blunt, T. Chan, K. Gerlach, "A New Framework for Direction-of-Arrival Estimation," *IEEE Sensor Array and Multichannel Processing Workshop*, Darmstadt, Germany, 21-23 July 2008.

[100] P. Stoica, Z. Wang, and J. Li, "Robust capon beamforming," *IEEE Signal Processing Letters*, vol. 10, no. 6, pp. 172–175, 2003.

[101] E. Hornberger, S. D. Blunt and T. Higgins, "Partially constrained adaptive beamforming for super-resolution at Low SNR," *2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Cancun, Mexico, 2015.

[102] T. Cuprak, "Applications of the Reiterative MMSE Algorithm to Underwater Acoustics using Covariance Matrix Tapers for Robustness," *Masters Thesis*, George Mason University, 2013.

[103] J.E. Quirk, R.J. Chang, J.W. Owen, S.D. Blunt, P.M. McCormick, "A Simple yet Effective Metric for Assessing Doppler Tolerance," *IEEE Trans. Radar Systems*, vol. 1, pp. 12-20, 2023.

[104] S. D. Blunt, A. K. Shackelford and K. Gerlach, "Single pulse imaging," *2006 International Waveform Diversity & Design Conference*, Lihue, HI, USA, 2006.

[105] S. D. Blunt, A. K. Shackelford, K. Gerlach, and K. J. Smith, "Doppler compensation & single pulse imaging using adaptive pulse compression," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 2, pp. 647–659, 2009.

[106] K. Kulpa, "The CLEAN type algorithms for radar signal processing," *2008 Microwaves, Radar and Remote Sensing Symposium*, Kiev, Ukraine, 2008.

[107] C. Sahin, J. G. Metcalf and B. Himed, "Reduced complexity maximum SINR receiver processing for transmit-encoded radar-embedded communications," *2018 IEEE Radar Conference (RadarConf18)*, Oklahoma City, OK, USA, 2018.

[108] S. Boyd and L. Vandenberghe, *Convex Optimization*, ser. Berichte über verteilte messysteme. Cambridge University Press, 2004, no. pt. 1.

[109] S. Marple, *Digital Spectral Analysis: Second Edition*, ser. Dover Books on Electrical Engineering. Dover Publications, 2019.

[110] A. Townsend, "Matrix-vector multiplication using the FFT," Teaching Notes, Massachusetts Institute of Technology, 2015.

[111] J. Xia, Y. Xi, and M. Gu, "A superfast structured solver for toeplitz linear systems via randomized sampling," *SIAM Journal on Matrix Analysis and Applications*, vol. 33, no. 3, pp. 837–858, 2012.

[112] G. S. Ammar and W. B. Gragg, "Superfast solution of real positive definite toeplitz systems," *SIAM Journal on Matrix Analysis and Applications*, vol. 9, no. 1, pp. 61–76, 1988.

[113] Y. Saad and M. H. Schultz, “Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 3, pp. 856–869, 1986.

[114] O. Rabaste and J. Bosse, “Robust mismatched filter for off-grid target,” *IEEE Signal Processing Letters*, vol. 26, no. 8, pp. 1147–1151, 2019.

[115] B. White, M. Heintzelman and S. D. Blunt, "Alternative “Bases” for Gradient-Based Optimization of Parameterized FM Radar Waveforms," *2023 IEEE Radar Conference (RadarConf23)*, San Antonio, TX, USA, 2023.

[116] E. R. Biehl, C. A. Mohr, B. Ravenscroft and S. D. Blunt, "Assessment of Constant Envelope OFDM as a Class of Random FM Radar Waveforms," *2020 IEEE Radar Conference (RadarConf20)*, Florence, Italy, 2020,.

[117] P. M. McCormick and S. D. Blunt, "Gradient-based coded-FM waveform design using Legendre polynomials," *International Conference on Radar Systems (Radar 2017)*, Belfast, 2017.

[118] A. W. Doerry. Generating nonlinear FM chirp waveforms for radar. *Sandia Report*, SAND2006-5856 (Sept. 2006).

[119] C. A. Mohr and S. D. Blunt, "Analytical Spectrum Representation for Physical Waveform Optimization Requiring Extreme Fidelity," *2019 IEEE Radar Conference (RadarConf)*, Boston, MA, USA, 2019.

[120] D. A. Bini and B. Meini, “Effective methods for solving banded toeplitz systems,” *SIAM Journal on Matrix Analysis and Applications*, vol. 20, no. 3, pp. 700–719, 1999.

[121] E. Baranoski, “Triangular factorization of inverse data covariance matrices,” in *1991 International Conference on Acoustics, Speech, and Signal Processing*, 1991, pp. 2245–2247 vol.3.

[122] T. F. Chan, “An optimal circulant preconditioner for toeplitz systems,” *SIAM Journal on Scientific and Statistical Computing*, vol. 9, no. 4, pp. 766–771, 1988.

[123] N. K. Govindaraju, B. Lloyd, Y. Dotsenko, B. Smith, and J. Manferdelli, “High performance discrete fourier transforms on graphics processors,” in *SC ’08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, 2008, pp. 1–12.

[124] B. Ravenscroft, J. W. Owen, S. D. Blunt, A. F. Martone and K. D. Sherbondy, "Optimal Mismatched Filtering to Address Clutter Spread from Intra-CPI Variation of Spectral Notches," *2019 IEEE Radar Conference (RadarConf)*, Boston, MA, USA, 2019.

[125] B. H. Kirk, A. F. Martone, K. A. Gallagher, R. M. Narayanan and K. D. Sherbondy, "Mitigation of Clutter Modulation in Cognitive Radar for Spectrum Sharing Applications," in *IEEE Transactions on Radar Systems*, vol. 1, pp. 193-204, 2023.

[126] B. Ravenscroft, J.W. Owen, B.H. Kirk, S.D. Blunt, A.F. Martone, K.D. Sherbondy, R.M. Narayanan, “Experimental Assessment of Joint Range-Doppler Processing to Address Clutter Modulation from Dynamic Radar Spectrum Sharing,” *IEEE International Radar Conference*, Washington, DC, 27-30 Apr. 2020.

[127] M. Golay, “Complementary series,” *IRE Transactions on Information Theory*, vol. 7, no. 2, pp. 82–87, 1961.

[128] C. A. Mohr, P. M. McCormick and S. D. Blunt, "Optimized complementary waveform subsets within an FM noise radar CPI," *2018 IEEE Radar Conference (RadarConf18)*, Oklahoma City, OK, USA, 2018.

[129] J. Bi and H. Rohling, “Complementary binary code design based on mismatched filter,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1793–1797, 2012.

[130] C. C. Jones and S. D. Blunt, "Mismatched Complementary-on-Receive Filtering of Diverse FM Waveform Subsets," *2019 International Radar Conference (RADAR)*, Toulon, France, 2019.

[131] C. C. Jones, C. A. Mohr, P. M. McCormick, and S. D. Blunt, "Complementary frequency modulated radar waveforms and optimised receive processing," *IET Radar, Sonar & Navigation*, vol. 15, no. 7, pp. 708–723, 2021.

[132] T. F. Chan and J. A. Olkin, "Circulant preconditioners for toeplitz-block matrices," p. 89–101, Mar. 1994.

[133] A. Jakobsson, S. Marple, and P. Stoica, "Computationally efficient two-dimensional capon spectrum analysis," *IEEE Transactions on Signal Processing*, vol. 48, no. 9, pp. 2651–2661, 2000.

[134] C. Jones et al., "Computationally Efficient Joint-Domain Clutter Cancellation for Waveform-Agile Radar," *2021 IEEE Radar Conference (RadarConf21)*, Atlanta, GA, USA, 2021.

[135] F. Colone, D. W. O'Hagan, P. Lombardo, and C. J. Baker, "A multistage processing algorithm for disturbance removal and target detection in passive bistatic radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 2, pp. 698–722, 2009.

[136] F. Colone, C. Palmarini, T. Martelli, and E. Tilli, "Sliding extensive cancellation algorithm for disturbance removal in passive radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 3, pp. 1309–1326, 2016.

[137] A. Farina and H. Kuschel, "Guest editorial special issue on passive radar (part i)," *IEEE Aerospace and Electronic Systems Magazine*, vol. 27, no. 10, pp. 5–5, 2012.

[138] C. Ning, J. Tian, K. Li, and S. Wu, "Modified adaptive pulse compression algorithm for

targets with range-straddling," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 5, pp. 3057–3070, 2021.

[139] C. Ning, J. Tian, S. Zheng, B. Zhang and W. Cui, "Robust Adaptive Pulse Compression Algorithm for Targets with Straddling," *2023 IEEE Radar Conference (RadarConf23)*, San Antonio, TX, USA, 2023.

[140] L. Harnett, "Mismatched Processing for Radar Interference Cancellation," *Doctoral dissertation*, University of Kansas, 2022.

[141] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," Oct. 2008.

[142] R. Baraniuk and P. Steeghs, "Compressive Radar Imaging," *2007 IEEE Radar Conference*, Waltham, MA, USA, 2007.

[143] R. Bose, A. Freedman, and B. Steinberg. Sequence CLEAN: A modified deconvolution technique for microwave imaging of contiguous targets. *IEEE Transactions on Aerospace & Electronic Systems*, 38, 1, 89–97 (Jan. 2002).

[144] H. Deng. Effective CLEAN algorithms for performance-enhanced detection of binary coding radar signals. *IEEE Transactions on Signal Processing*, 52, 1, 72–78 (Jan. 2004).

[145] S. Blake, "OS-CFAR theory for multiple targets and nonuniform clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, no. 6, pp. 785–790, 1988.

[146] P. M. McCormick, S. D. Blunt and T. Higgins, "A gradient descent implementation of adaptive pulse compression," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[147] S.D. Blunt, T. Higgins, "Dimensionality Reduction Techniques for Efficient Adaptive Radar Pulse Compression," *IEEE Trans. Aerospace & Electronic Systems*, vol. 46, no. 1, pp. 349-362, Jan. 2010.

[148] Z. Li, Y. Zhang, S. Wang, L. Li, and M. Mclinden, “Fast adaptive pulse compression based on matched filter outputs,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 1, pp. 548–564, 2015.

[149] B. Ghojogh, F. Karray, and M. Crowley, “Eigenvalue and generalized eigenvalue problems: Tutorial,” 2019.

[150] J. A. Ford, “Improved algorithms of illinois-type for the numerical solution of nonlinear equations.”

[151] A. Jain, “Fast inversion of banded toeplitz matrices by circular decompositions,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 26, no. 2, pp. 121–126, 1978.

[152] R. H. Chan and M.-C. Yeung, “Circulant preconditioners for complex toeplitz matrices,” *SIAM Journal on Numerical Analysis*, vol. 30, no. 4, pp. 1193–1207, 1993.

[153] A. C. O’Connor, "Frequency-Domain Synthesis of Advanced Pulse-Compression Filters," *2019 IEEE Radar Conference (RadarConf)*, Boston, MA, USA, 2019.

[154] W. A. Blake, "Multi-dimensional mismatch filter design optimization for radar waveforms," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[155] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.1,” http://cvxr.com/cvx, Mar. 2014.

[156] A. C. O’Connor, J. M. Kantor and J. Jakabosky, "Joint equalization filters that mitigate waveform-diversity modulation of clutter," *2016 IEEE Radar Conference (RadarConf)*, Philadelphia, PA, USA, 2016.

[157] E. L. Yip, “A note on the stability of solving a rank-p modification of a linear system by the sherman–morrison–woodbury formula,” *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 2, pp. 507–513, 1986.