

UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

Adaptive Processing for Fully Digital Polarimetric Phased Array Weather Radar

A DISSERTATION
SUBMITTED TO THE GRADUATE FACULTY
in partial fulfillment of the requirements for the
Degree of
DOCTOR OF PHILOSOPHY

By
YOON-SL KIM
Norman, Oklahoma
2026

Adaptive Processing for Fully Digital Polarimetric Phased Array Weather Radar

A DISSERTATION APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. David Schvartzman, Chair

Dr. Robert Palmer, Co-Chair

Dr. Tian-You Yu

Dr. Feng Nai

Dr. Christopher Curtis

Dr. Justin Metcalf

Dr. Ying Wang

Acknowledgments

I would like to express my sincere gratitude to my advisors, Dr. David Schwartzman and Dr. Robert Palmer, for their mentorship and guidance throughout my Ph.D. Their continuous support, insightful discussions, and regular research meetings have significantly broadened my research perspective and shaped the direction of this work.

I am especially grateful to Dr. David Schwartzman, who has advised and mentored me since my master's studies. His passion for research has consistently motivated me, and his encouragement has played an important role in helping me grow as a researcher. He has taught me a great deal about radar signal processing, and his expertise in this area, along with beam synthesis, has had a strong influence on my research. His guidance has been both technical and professional, and I greatly appreciate his continued support.

I would also like to thank Dr. Robert Palmer for his mentorship and guidance. As a pioneer in applying adaptive beamforming to profiler radar, his insights have been directly relevant to my Ph.D. research. His perspective and feedback have helped me better understand the broader context of this work and refine my approach.

I would like to thank my collaborators, Dr. Tian-Yu You, Dr. Feng Nai, and Dr. Christopher Curtis. They have been actively involved in my Ph.D. research, and our regular monthly meetings provided valuable discussions and diverse perspectives. Their support and feedback have been very helpful throughout this work.

Finally, I would like to acknowledge the ARRC engineering team for their effort and dedication in developing and deploying the Horus radar system. Their work made it pos-

sible to collect the high-quality data used in this dissertation, which was essential for the validation and analysis presented here.

Table of Contents

Acknowledgments	iv
Table of Contents	ix
List of Tables	x
List of Figures	xxiv
Abstract	xxv
1 Introduction	1
1.1 Background	1
1.2 Motivation	3
1.3 Contributions	7
1.4 Outline	9
2 Unique Capabilities of Phased Array Radar	11
2.1 Fundamentals of Phased Array Radar	12
2.2 Advantages of Phased Array Radar	14
2.3 Planar Array Geometry and Architecture	15
2.3.1 Subarrays	17
2.3.2 Sparse Array	18
2.4 Imaging with Digital Beamforming: Principles and Methods	22
2.5 Adaptive Beamforming	23

2.5.1	Capon Beamforming	24
2.5.2	Robust Capon Beamforming with Linear Constraints	27
2.5.3	Robust Capon Beamforming with Quadratic Constraints	29
2.5.4	Robust Capon Beamforming with Diagonal Loading	32
2.5.5	Minimum Mean-Square Error (MMSE) Beamforming	37
2.5.6	AI-Based Fast Capon Beamforming	40
2.5.7	Bayesian Adaptive Beamforming under DoA Uncertainty	42
2.6	Deterministic Null-Constrained Beamforming	44
3	Calibration for Adaptive Beamforming	49
3.1	Beam Pattern Calibration	51
3.2	Noise Calibration	57
3.3	Evaluation of Calibration for Adaptive Beamforming	58
3.3.1	Adaptive Beam Pattern Characteristics	60
3.3.2	Calibrated Polarimetric Variables	63
3.4	SNR Effects and Diagonal Loading	68
3.5	Cross-Correlation Coefficient After Adaptive Beamforming	70
3.6	H–V Channel Mismatch in Adaptive Beamforming	76
3.6.1	H–V Beamforming Weights Mismatch	77
3.6.2	H–V Beam Matching	80
3.7	Simulation-Based Evaluation with NEXRAD Data	86
4	Space-Time Adaptive Processing	92
4.1	Fundamentals of Space-Time Processing	95
4.2	Simulation-based Evaluation of Space-Time Processing	97
4.2.1	Spectrum Characteristics of Different Types of Targets	97
4.2.2	Performance Evaluation of Space-Time Processing	99
4.2.3	Clutter-to-Signal Ratio Impact on Space-Time Processing Performance	101

4.2.4	Impact of Sample Size on Space-Time Adaptive Processing Performance	103
4.2.5	Additional Case Studies	104
4.3	Improved Doppler Estimation via Space-Time Signal Processing	106
4.3.1	Simulation Model and Parameters	107
4.3.2	Performance Evaluation of Doppler Estimation	108
5	Case Study: Application to the Horus Phased Array Radar	117
5.1	Overview of the Horus Radar System	118
5.2	Horus Case Study 1: January 26, 2024	119
5.3	Horus Case Study 2: April 28, 2024	128
5.4	Horus Case Study 3: June 5, 2025	132
5.5	Horus Case Study 4: July 23, 2025	137
6	Practical Adaptive Processing for Operational Use	145
6.1	Conditional Application of Adaptive Beamforming	147
6.1.1	Application of Conditional Adaptive Beamforming on NEXRAD-Based Simulations	151
6.1.2	Application of Conditional Adaptive Beamforming on Horus Data .	154
6.2	Computationally Efficient Adaptive Processing with Reduced Covariance Matrix	157
6.2.1	Adaptive Digital Beamforming on Sparse Arrays	157
6.2.2	Adaptive Processing on Subarray	159
6.3	High Sidelobes Mitigation with Beam Synthesis	163
6.4	Adaptive Clutter Filter with Adaptive Processing	166
7	Conclusions and Future Work	177
7.1	Limitations and Open Issues	179
7.2	Suggestions for Future Research	179

References	182
Appendix List of Symbols	198
Appendix A Simulation of Weather-Like Signals for Phased Array Radar	204
A.1 Time-Series IQ Signal Simulation	204
A.2 Spoiled Beam Simulation	206
Appendix B Spectral Moment Estimation	208
Appendix C Properties of Covariance Matrices and Quadratic Forms	210
C.1 Hermitian Property of Covariance Matrices	210
C.2 Positive Definiteness	211
C.3 Inverse of a Positive Definite Matrix	211
C.4 Quadratic Form	211

List of Tables

4.1	The table illustrates the variation in dynamic range in STAP output based on the number of samples and the computational time required to compute STAP.	115
6.1	Performance time comparison between Fourier and Capon processing . . .	146
6.2	Performance time for Fourier + Capon processing with different thresholds	154
6.3	Computation time of STAP for different numbers of RX channels	162

List of Figures

1.1	Horus: Field deployment at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025.	4
1.2	Illustrations of the Fourier (non-adaptive) and Capon (adaptive) two-way beam patterns (dB). These patterns were implemented to Horus during the field deployment at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025. The black lines indicate the steering point of the receive beam, and the red dot marks the angular location of wind turbine interference at 2.5° azimuth.	5
2.1	Fundamental concepts of a PAR using a uniform linear array. Color shading represents the relative gain, where blue tones represent lower values and red tones represent higher values [43].	12
2.2	Four array configurations are evaluated: square, truncated, rectangular, and circular arrays. The upper panels provide the array layouts, the middle panels show the resulting 3D array factors, and the bottom panels provide the 1D azimuth cuts.	16
2.3	Non-overlapped subarrays consisting of 8 elements (1×8 columns), 32 elements (4×8 columns), and 64 elements (8×8 columns) are illustrated above with their corresponding array factors.	19

2.4	This figure shows the truncated PAR with different sparsity levels: 10%, 30%, 60%, and 90%. The first column shows the array configuration for each sparsity level, the second column shows the top view of the patterns, the third column shows the 3D view of the patterns, and the fourth column shows the 1D azimuth and elevation cuts. For the 1D pattern plots, the red dotted lines represent the beam patterns from the fully populated array. . . .	21
2.5	Simulated beam patterns obtained using the Fourier method (left) and the Capon method (right) are shown above. The blue arrows indicate the steering direction, and the black arrows indicate the direction of the interference.	26
2.6	The patterns from the Fourier (blue), Capon (red), linearly constrained Capon (yellow), and quadratically constrained Capon (purple) are simulated under the no-interference condition and plotted above.	30
2.7	Simulated beam patterns from the Fourier method (blue), Capon (orange), robust Capon with linear constraints (yellow, referred to as the LCC beam pattern), quadratic constraints with diagonal loading (purple, referred to as QCC beam pattern 1), quadratic constraints solved with an iterative method (green, referred to as QCC beam pattern 2), and quadratic constraints solved with CVX (light blue, referred to as QCC beam pattern 3) are presented above. Interference is present at 10° (indicated as a gray line in each plot). Beams are plotted for three cases: Case I, where the interference is located outside the mainlobe (beam steering at 18°), Case II, where the interference is located inside the mainlobe (beam steering at -2°), and Case III, where the beam is steered exactly at the interference (beam steering at -10°).	34

2.8	The measured 3-dB beam power (left plot) and the difference between the desired and actual peak directions (right plot) are presented above. Six beamformers are used: the Fourier method (blue), Capon (red), robust Capon with linear constraints (yellow, referred to as the LCC beam pattern), quadratic constraints with diagonal loading (purple, referred to as QCC beam pattern 1), quadratic constraints solved with an iterative method (green, referred to as QCC beam pattern 2), and quadratic constraints solved with CVX (light blue, referred to as QCC beam pattern 3) . . .	36
2.9	Simulation results from [62]. This figure shows examples of mean power estimates and formed beam patterns for two point targets. (a-1) True value of the point targets placed at azimuths of 0° and 0° , and elevations of 5° and 5° , respectively. (a-2) and (a-3) Mean power estimates for two targets in the results of Fourier and MMSE, respectively. (b-1) Formed beam pattern grazing at an azimuth of 0° and an elevation of 0° , respectively, along a line through two points. (b-2) and (b-3) Formed beam pattern grazing at an azimuth of 0° and an elevation of 0° in whole range.	39
2.10	Simulated results from [64], for the point target case are presented above. Left plot - DBF Power of Capon, Fourier and DLAF DBF and right plot - Antenna Pattern of Capon, Fourier and DLAF DBF. The Capon method is plotted in green, the Fourier method is plotted in blue, and DLAF is plotted in yellow.	41
2.11	Simulated results from [64], for the weather case are presented above. Left plot - Reflectivity field of the simulated Fourier DBF, middle plot - Reflectivity field of the simulated Capon DBF, and right plot - Reflectivity field of the simulated DLAF DBF.	42

2.12	Beam patterns from the Fourier beamformer with no constraint (black line), the beamformer with a zero-order constraint (yellow line), the beamformer with a first-order constraint (blue line), the beamformer with a second-order constraint (red) are plotted above. Their nulling constraints are at 10° , noted as a dotted line. Beams are pointed at 0°	48
3.1	The model polarimetric variables for Case 1 (blue), Case 2 (dotted red), and Case 3 (green) are presented. Case 1 is uniformly distributed targets, representing a center of the storm, Case 2 contains a stationary point target, and Case 3 has gradient changes, representing the edge of storms.	58
3.2	The configuration of the array is shown in the left plot, and the receive pattern power (dB) is shown in the right plot.	59
3.3	Beam patterns for the Fourier (black for H channel) and Capon (red for H channel and blue for V channel) methods across three different scenarios: the left figure corresponds to Case 1 (uniform distribution), the middle figure to Case 2 (point target), and the right figure to Case 3 (gradient change). The solid black line represents the beam steering angle, while the yellow line denotes the target power, serving as a reference.	61
3.4	The differences (in dB) between the H and V beams produced by the Capon method are analyzed for the three different cases and presented above. The x -axis represents the elevation angle, allowing for a direct comparison of beam mismatches as a function of target distribution. The y -axis represents the H and V mismatch in dB.	62
3.5	The calibrated polarimetric variables (rows) for the three cases (columns) using the Fourier (blue) and Capon (red) methods are presented. The true values are plotted with a black line in each plot.	64

3.6	Normalized probability density functions (PDFs) of bias in estimated polarimetric variables from 500 simulations of PAR time-series signals for three scenarios: Case 1 (uniform field), Case 2 (point target), and Case 3 (complex scene). Blue bars indicate bias from the Fourier method, and red bars represent the Capon method. Variables shown include P , v_r , Z_{DR} , ρ_{hv} , and ϕ_{DP}	67
3.7	Bias in the estimated polarimetric variables as a function of SNR. Results for Cases 1, 2, and 3 are shown in blue, orange, and green, respectively. The Fourier method is plotted with solid lines, and the Capon method with dotted lines.	69
3.8	Mean bias of ρ_{hv} from the Capon method as a function of Diagonal Loading (the x -axis) and SNR (the y -axis).	70
3.9	The correlation between the Capon beamforming weights in the H and V channels is plotted for different diagonal loading levels: -50 dB (blue), -40 dB (orange), -30 dB (yellow), -20 dB (purple), and -10 dB (green).	71
3.10	Differences between the Capon beam patterns in the H and V channels for Cases 1, 2, and 3 (see Figure 3.1).	76
3.11	The amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods, when the radar is steered at broadside ($\theta = 0^\circ$, $\phi = 0^\circ$), are shown above.	79
3.12	The amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods, when the radar is steered off broadside ($\theta = 20^\circ$, $\phi = 0^\circ$), are shown above.	80

3.13	Correlation between the Capon beamforming weights in the H and V channels as a function of SNR (shown on the y -axis) for Cases 1, 2, and 3. The x -axis represents the angular coordinate and corresponds to the x -axis in Figure 3.1.	81
3.14	The bias in Z_{DR} estimation as a function of diagonal loading (the y -axis) and ρ_{hv} (the x -axis) for the H–V matched case (left plot) and the H–V independent case (right plot) are presented.	83
3.15	The bias in Z_{DR} estimation as a function of SNR (the y -axis) and ρ_{hv} (the x -axis) for the H–V matched case (left plot) and the H–V independent case (right plot) are presented.	84
3.16	The bias in ρ_{hv} estimation as a function of diagonal loading (top row) and SNR (bottom row), with the intrinsic ρ_{hv} shown on the x -axis, for the H–V matched case (left column) and the H–V independent case (right column).	85
3.17	The geometry of the PAR is illustrated in the left plot. The middle plot shows the transmit beam configuration, the degree of overlap between adjacent transmit beams, and the number of receive beams formed within each transmit beam. The right plot illustrates how the rotating PAR scanning strategy was simulated.	87
3.18	The original NEXRAD Level-II data (left column), simulated Fourier DBF (middle column), and Capon beamforming (right column) are presented. The case is from the KTLX radar on May 6, 2013 at 20:00:50.96 UTC. The rows from top to bottom show reflectivity (Z_H), radial velocity (v_r), differential reflectivity (Z_{DR}), correlation coefficient (ρ_{hv}), and differential phase (ϕ_{DP}).	88
3.19	Azimuthal line plots at 26.875 km range comparing reflectivity (top-left), velocity (top-right), differential reflectivity (bottom-left), and correlation coefficient (bottom-right) for the truth (black), Fourier (blue), Capon with matched H–V beams (green), and Capon with independent beams (red).	90

4.1	Example STAP application to remove simulated WTC from actual digital PAR weather data (AIR); (left) weather only, (middle) WTC contamination on weather, and (right) recovered weather [70].	93
4.2	Doppler–time plot from the field deployment of Horus at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025.	94
4.3	Illustration of the digital PAR data cube and the space-time processing framework.	95
4.4	Modeled Doppler spectra for stationary ground clutter (left), weather-like targets (middle), and their combination (right). The x -axis represents radial velocity (m/s), and the y -axis represents power spectral density (PSD) in dB.	99
4.5	STP and STAP results for stationary ground clutter and weather-like targets. The x -axis represents radial velocity (m/s), and the y -axis represents azimuth ($^{\circ}$).	100
4.6	STP and STAP results for the combined scenario of stationary ground clutter and weather-like targets. The 2D plots show the joint spatial-temporal response. The 1D cuts correspond to velocity at -38° azimuth (top) and azimuth at 0 m/s (bottom), indicated by red lines in the 2D plots.	101
4.7	STP and STAP results for different clutter-to-signal ratios (CSR). The clutter is located at -8° azimuth with zero mean velocity, and the weather-like signals are distributed across azimuth with varying mean velocities.	102
4.8	STP (red) and STAP (blue) results for different numbers of samples. The computation time for each case is shown at the bottom.	103
4.9	STP and STAP results for four cases involving stationary and moving point targets embedded in weather-like echoes.	105
4.10	STP and STAP results for simulated wind turbines at different orientation angles (0° , 30° , 60° , and 90°).	106

4.11	The normalized spoiled transmit and receive beam patterns are plotted on a dB scale, with the x -axis and y -axis representing azimuth and elevation dimensions in degrees.	108
4.12	Normalized powers using STP (left column) and STAP (right column) with two configurations: fully digital (top row) and 1×8 subarray digital (bottom row) arrays. The outputs are normalized and plotted on a dB scale.	109
4.13	1D cuts of STP and STAP powers with two configurations shown in Figure 4.12. The top plot shows the Doppler cut at 0.6° azimuth, and the bottom plot shows the azimuth cut at -20 m/s velocity.	111
4.14	Estimated Doppler velocity for fully digital (top) and subarray (bottom) systems.	112
4.15	Quantitative RMSE analysis of Doppler velocity estimation.	114
5.1	The truncated array configuration consisting of 13 panels, with each panel forming an 8×8 subarray.	120
5.2	Reflectivity data captured by the Horus system on January 26, 2024 at 15:57:14 UTC. The right column shows full-scene views: the top plot is a front view at the range of maximum reflectivity, and the bottom plot is a PPI scan. The dataset includes weather, clutter, and aircraft. The left column shows zoomed-in views highlighting aircraft targets.	121
5.3	Aircraft tracks detected by Horus on January 26, 2024 at 16:32:10 UTC (source: ADS-B Exchange). The scanning sector is indicated by dotted black lines, and aircraft information is shown in the left panels.	122
5.4	Zoomed-in PPI and RHI plots of power (dB) from the dataset collected on January 26, 2024 at 16:32:10 UTC. The PPI is shown at 8.125° elevation, and the RHI is shown at 31.3° azimuth to capture the aircraft highlighted in Figure 5.3.	122

5.5	Normalized power after DBF with the Fourier method (left column) and Capon method (right column) are plotted in dB scale. The black box in each plot indicates the transmit beam used for STP and STAP in the later step.	124
5.6	The STP and STAP results are plotted on a dB scale, with Doppler velocity (m/s) on the x -axis and azimuth (degrees) on the y -axis at 8.125° elevation. The power levels are normalized to 0 dB.	124
5.7	The output power map from space-time processing is presented above. The x -axis represents Doppler velocity, and the y -axis represents azimuth. The figure is shown as a contour plot, with all elevation slices stacked together. The color scale denotes power in dB.	126
5.8	The output power map from space-time processing is presented above. The x -axis represents Doppler velocity, and the y -axis represents elevation. The figure is shown as a contour plot, with all azimuth slices stacked together. The color scale denotes power in dB.	127
5.9	Estimated velocity over the azimuth from STP (blue), STAP (red), STP with filtering (dotted yellow), and STAP with filtering (dotted purple) are shown above. The x -axis represents azimuth in degrees, and y -axis represents the estimated velocity in m/s.	128
5.10	Measured beam patterns. The left column shows the product of the transmit beam pattern and the subarray pattern. The middle column shows the array factor formed using the Fourier method (top row) and the Capon method (bottom row), and the right column shows the resulting beam pattern. The beams are steered to 0° azimuth and 15° elevation.	129
5.11	1D azimuth cut of the Capon beam pattern. The gray line indicates the steering angle (-20° azimuth, 14.25° elevation). The dotted lines indicate the transmit-beam boundaries, and the red line marks the location of stationary clutter.	130

5.12	Reflectivity fields produced by the Fourier and Capon DBF methods. The Horus data were collected at 02:24:24 UTC, while the KTLX data were collected at 02:24:54 UTC.	131
5.13	Velocity fields produced by the Fourier and Capon DBF methods.	132
5.14	The truncated array configuration, consisting of 26 subarrays with each subarray containing 4×8 elements, is illustrated in the left plot. The transmit beam used during data collection is shown in the right plot.	133
5.15	Measured two-way patterns using the Fourier (left) and Capon (right) methods. The black lines indicate the steering directions (-3.25° elevation, 0° azimuth).	134
5.16	RHI scan at 0.375° azimuth using Fourier (left) and Capon (right) DBF Bottom plots are showing the zoomed-in view from the top plots.	134
5.17	NEXRAD (left), Horus Fourier DBF (middle), and Horus Capon DBF (right). Rows from top to bottom show reflectivity (Z), radial velocity (v_r), differential reflectivity (Z_{DR}), correlation coefficient (ρ_{hv}), and differential phase (ϕ_{DP}).	135
5.18	Left: 2D normalized transmit beam (dB) used by Horus for data collection. Right: array configuration, where blue dots represent the 832 antenna elements, red dots indicate subarray centers used for DBF, and the red box illustrates the subarray layout.	138
5.19	Normalized two-way beam patterns (dB) at range 2.2 km: Fourier method (left) and Capon method (right). The black lines indicate the steering direction of the receive beam, and the red dot marks the angular location of wind turbine interference at azimuth 2.5°	138
5.20	DBF power (dB) from RHI scans at azimuth 2.5° . Top: Fourier method; Bottom: Capon method. The x -axis represents range (km).	139
5.21	The RHI scans of calibrated Fourier and Capon DBFs for the H and V channels at an azimuth of 0.5° are presented above.	140

5.22	Calibrated power (dB) at azimuth 2.5° and elevation 8.5° as a function of range. Fourier DBF results are shown in blue, and Capon results are shown in red.	141
5.23	The micro-Doppler plots of wind turbines, with one at a low elevation angle (1.5°) and one at a higher elevation angle (6.5°), are presented above. The azimuth angle is fixed at 0.5°	143
5.24	Space-time processing, both STP and STAP, results are presented above. The x -axis represents the Doppler dimension and y -axis represents the angular dimension.	144
6.1	The modeled power (dB) of two cases, a point target and distributed targets, with different gradient amount (plotted in different colors in each plot) are presented in the left plots. The bias measured in each polarimetric variables after the Fourier DBF (top row) and the Capon DBF (bottom row).	148
6.2	Bias measured in polarimetric variables after the Fourier DBF (first row) and the Capon DBF (second row) at different power gradient (the x -axis) and SNR (the y -axis) for a sharp power gradient case. The third row shows which DBF method produces less bias.	149
6.3	Bias measured in polarimetric variables after the Fourier DBF (first row) and the Capon DBF (second row) at different power gradient (the x -axis) and SNR (the y -axis) for a smooth power gradient case. The third row shows which DBF method produces less bias.	150
6.4	Bias map of a sharp gradient and smooth gradient case from Figure 6.2 and Figure 6.3 with gradient slope (dB/deg) on the x -axis.	151
6.5	Reflectivity fields from the Fourier DBF, conditional application of Fourier and Capon DBF, and Capon DBF.	152
6.6	Velocity fields from the Fourier DBF, conditional application of Fourier and Capon DBF, and Capon DBF.	153

6.7	RHI scans of the calibrated reflectivity obtained using the Fourier DBF (left) and the Capon DBF (right), from Horus data collected on June 5, 2025 (details in Section 5.4).	155
6.8	The storm traces from the Fourier DBF (black) and the Capon DBF (red) are plotted in the top plot and the magnitude of the differences between the two results (in dB) is plotted in the bottom plot.	156
6.9	The left plot shows the variance of power, with traces indicating regions where the variance exceeds 5 dB. The right plot shows the output from the conditional Capon DBF.	157
6.10	Estimated bias in polarimetric variable estimation after the Fourier DBF (top-row) and the Capon DBF (bottom-row), as a function of sparsity (the y -axis) and elevation (the x -axis) for a point target case are presented. . . .	158
6.11	Estimated bias in polarimetric variable estimation after the Fourier DBF (top-row) and the Capon DBF (bottom-row), as a function of sparsity (the y -axis) and elevation (the x -axis) for a distributed target case are presented. . . .	159
6.12	The computation time required to perform the Capon DBF for different sparsity levels is presented.	160
6.13	Fully digital (left column), non-overlapped subarrays consisting of 8 elements (middle column), and non-overlapped subarrays consisting of 16 elements (right column) are evaluated. The second-row plots show the power map after STP, and the third-row plots show the power map after STAP. . . .	161
6.14	Measured computation time to perform STAP as a function of number of receivers (RX).	162

6.15	The two-way patterns of the fully digital array (top row) and the 1×4 subarray system (bottom row) using two different transmit beam patterns: deterministic and optimized. Each plot contains the two-way patterns for multiple steering angles (plotted in different colors), from -10° to 10° with a 1° sampling rate. The dotted lines represent the mask that was used when the transmit beams were synthesized.	165
6.16	The flowchart of the proposed classification algorithm is illustrated above. .	168
6.17	Horus data from January 26, 2024: The power (dB) after the Fourier DBF for stationary ground clutter at 5.99 km is presented in the top plot. The second row shows plots of each transmit beam that is detected after Fourier DBF. The third row shows plots of the transmit beam after the Capon DBF.	170
6.18	Horus data from January 26, 2024: The power (dB) after the Fourier DBF for moving clutter at 24.41 km is presented in the top plot. The second row shows plots of each transmit beam that is detected after Fourier DBF. The third row shows plots of the transmit beam after the Capon DBF.	171
6.19	Space-time processing output for the second case at a range of 24.41 km and TX54.	173
6.20	1D plot of PSD at $\phi = 29^\circ$ and $\theta = 6^\circ$ is presented above. The blue line illustrates the spectrum before applying the filter, while the red dotted line shows the spectrum after the filter has been applied.	174
6.21	Estimated Doppler velocity before (top plot) and after (bottom plot) the filtering is illustrated above. The x -axis and y -axis represent the azimuth and elevation in degrees.	175

A.1	The framework is shown above, consisting of three main components: time-series IQ signal simulation, PAR simulation, and digital beamforming. The red boxes represent inputs to each stage and the blue boxes represent outputs from each stage. The outputs from the digital beamforming process, using both Fourier and Capon methods, are evaluated and compared to the initial model values.	205
-----	--	-----

Abstract

This dissertation investigates advanced adaptive signal processing techniques for fully digital polarimetric phased array weather radar (PAR) systems, with an emphasis on improving data quality, clutter mitigation, and operational feasibility. Conventional weather radars based on mechanically steered parabolic antennas are fundamentally limited in temporal resolution, restricting their ability to observe rapidly evolving atmospheric phenomena. Fully digital PAR systems overcome these limitations through electronic beam steering and digital beamforming (DBF), enabling rapid volumetric scanning and flexible beam synthesis.

This work focuses on adaptive beamforming and space-time adaptive processing (STAP) as key enabling technologies for next-generation weather radar. While adaptive methods such as Capon beamforming offer improved angular resolution, sidelobe suppression, and interference mitigation compared to conventional Fourier-based techniques, their application in polarimetric weather radar introduces significant challenges, including calibration complexity, computational cost, and sensitivity to sample support.

To address these challenges, a comprehensive framework for adaptive processing, specifically the Capon (also called minimum variance distortionless response) method, is developed, combining theoretical analysis, simulation studies, and validation using real data from the fully digital S-band Horus radar system. A novel calibration methodology is proposed to preserve dual-polarization consistency under dynamically varying adaptive beam patterns. The performance of adaptive beamforming is qualitatively evaluated against

conventional methods in terms of clutter suppression, sidelobe reduction, and polarimetric variable estimation accuracy. It is also evaluated quantitatively with the impact of signal-to-noise ratio (SNR), diagonal loading, and number of pulses. In addition, space-time processing techniques are investigated for mitigating complex, non-stationary clutter sources such as wind turbines, demonstrating improved separation of meteorological signals in both spatial and Doppler domains.

Practical considerations for operational deployment are also addressed. Computationally efficient implementations are developed through reduced-dimension covariance estimation, subarray-based processing, and conditional application of adaptive methods. The impact of array architecture, signal-to-noise ratio, and environmental conditions on adaptive processing performance is systematically analyzed.

Results demonstrate that adaptive processing significantly enhances radar data quality, particularly in challenging environments with strong interference and spatial gradients, while improving polarimetric performance for several variables, including reflectivity, differential reflectivity, radial velocity, and differential phase, when properly calibrated. It shows that adaptive processing mitigates contamination from complex clutter that includes moving point targets and/or wind turbines, which are challenging to mitigate. However, the limitations are that (1) adaptive processing performance depends on SNR, which requires at least 0 dB SNR from our study, (2) it requires a large number of samples, typically $2N$, twice the number of receivers, which can lead to longer dwell time, (3) there is a computational burden from adaptive processing; in one experiment, the Capon DBF took 28 times more time than the Fourier DBF to perform a full Plan Position Indicator (PPI) scan, and (4) there remains a challenge in ρ_{hv} estimation, arising from different optimization in each H and V channel. This work establishes a pathway toward the practical integration of adaptive beamforming and space-time processing in next-generation weather radar systems.

Chapter 1

Introduction

1.1 Background

Radar plays a significant role in advancing weather surveillance capabilities and improving the understanding of atmospheric science. Radar measurements quantify how electromagnetic (EM) waves are affected by propagation and reflection when interacting with hydrometeors such as raindrops, snowflakes, hailstones, or graupel. Through these measurements, bulk properties of hydrometeors, such as average concentration, size, shape, motion, orientation, and phase composition, can be studied. This information is essential for identifying hydrometeor types, estimating precipitation amounts [1], and understanding storm dynamics.

Polarimetric Doppler radar is the most common type of weather radar. It transmits horizontally (H) and vertically (V) polarized electric fields and uses a coherent receiver. The dual-polarization capability enables the measurement of spectral moments (reflectivity Z_H , Doppler velocity v_r , spectrum width σ_v) and polarimetric variables (differential reflectivity Z_{DR} , cross-correlation coefficient ρ_{hv} , differential phase ϕ_{DP}). These variables provide valuable information about hydrometeor microphysics, enhancing quantitative precipitation estimation, microphysical process studies, and target classification between meteorological and non-meteorological echoes such as birds and insects [2–7].

Z_H quantifies the power-weighted average radar return from scatterers within the resolution volume, while v_r measures their motion toward or away from the radar. Z_{DR} is one of the most informative polarimetric variables [8], as it expresses the ratio of the reflectivity factors at horizontal and vertical polarizations. Because Z_{DR} depends on the dielectric constant and shape of hydrometeors, it is useful for distinguishing between liquid and frozen particles. For example, a raindrop typically has a higher Z_{DR} value than a snowflake or graupel of the same size. ϕ_{DP} represents the difference between the phases of the horizontal and vertical components of the backscattered wave. ρ_{hv} quantifies the degree of phase coherence between the complex voltages measured in the horizontal and vertical polarization channels, reflecting how consistently the phase difference between H and V remains stable over time. For most weather scatterers, ρ_{hv} is close to 1 because the fluctuations occur almost in unison. This makes ρ_{hv} a useful variable for distinguishing weather echoes from non-meteorological scatterers such as ground clutter, insects, and birds.

The current Next-Generation Weather Radar (NEXRAD) program consists of 160 S-band, dish-based Weather Surveillance Radar-1988 Doppler (WSR-88D) systems distributed across the contiguous United States and operated by the NOAA National Weather Service (NWS) [9, 10]. This network provides real-time and archived datasets widely used to study precipitation, storm structure, and severe weather. Each WSR-88D radar employs mechanical rotation of a parabolic dish antenna to achieve volumetric coverage [11]. Although these systems achieve high angular resolution through narrow beamwidths (approximately 1°) and low antenna sidelobe levels, their mechanically steered design fundamentally limits temporal resolution. The fastest volume scans currently achieved by the WSR-88D are approximately 4 minutes, which is often too slow for capturing rapidly evolving phenomena such as tornado genesis, microbursts, and convective storm development [12–14].

High temporal resolution is critical for observing fast-evolving storm features, improv-

ing hazard prediction, and reducing errors in the retrieval of vertical air motion. Previous studies have shown that a temporal resolution of about 1 minute is needed for accurate analyses [15, 16]. NOAA has also specified a 1-minute update rate as a functional requirement for the next-generation weather radar. However, achieving such rapid updates is challenging for conventional dish-based systems due to the inertia of the rotating antenna and the need for a sufficient number of independent samples to obtain statistically robust estimates.

1.2 Motivation

Phased array radar (PAR) technology offers a promising pathway for next-generation atmospheric observations and has the potential to overcome the fundamental limitations of mechanically steered dish-based systems. A PAR is composed of many radiating elements whose amplitude and phase can be controlled independently, enabling electronic beam steering, beamforming, and agile scanning capabilities, which are not possible with conventional reflector antennas [17]. Electronic steering allows the radar to scan a region of interest within seconds, significantly improving temporal resolution. Digital PARs extend these capabilities by providing element-level digital amplitude/phase control. This architecture enables flexible beam synthesis, including the transmission of pencil beams, spoiled beams, and multiple simultaneous beams. Beam multiplexing further exploits beam agility by rapidly pointing the beam in widely separated directions while signals are de-correlated, increasing the number of independent samples without increasing dwell time [18]. Multiple simultaneous transmit beams, achieved through amplitude-based or phase-only synthesis, also provide rapid volumetric coverage while maintaining antenna sensitivity [19, 20].

Several PAR systems have been developed for atmospheric research. These include the National Weather Radar Testbed (NWRT), which demonstrated the feasibility of single-polarization PAR for weather observations [21]; the Advanced Technology Demonstrator



Figure 1.1: Horus: Field deployment at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025.

(ATD), the first operational-sized S-band dual-polarization PAR [22–24]; the Polarimetric Atmospheric Imaging Radar (PAIR), a C-band imaging radar [25]; and Japan’s dual-polarization X-band PAR [26]. A particularly advanced system is the Horus radar, a fully digital, S-band polarimetric PAR developed at the University of Oklahoma with funding from NOAA (see Figure 1.1) [27]. Its flexible architecture supports pencil beams, imaging modes, beam multiplexing, and subarray configurations. This flexibility makes Horus one of the most promising candidates for next-generation multi-mission weather and aircraft surveillance radar.

The imaging mode is a unique PAR capability in which the radar transmits a spoiled beam and forms multiple narrow receive beams simultaneously within the illuminated volume using digital beamforming (DBF). This approach enables large-sector atmospheric sampling with improved temporal resolution. In the conventional approach, DBF is implemented with predetermined beamforming weights using the Fourier method. This data-independent method produces beams with fixed sidelobe levels and fixed angular resolu-




Figure 1.2: Illustrations of the Fourier (non-adaptive) and Capon (adaptive) two-way beam patterns (dB). These patterns were implemented to Horus during the field deployment at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025. The black lines indicate the steering point of the receive beam, and the red dot marks the angular location of wind turbine interference at 2.5° azimuth.

tion, which has higher risk of sidelobe contamination. In contrast, adaptive methods such as the minimum variance distortionless response (MVDR), also known as the Capon method, compute data-dependent weights based on the scanning environment [28–30], thereby minimizing sidelobe contamination. The Capon method minimizes power from all directions except the desired steering direction and places nulls toward undesired high-power returns, reducing contamination and improving data quality. Figure 1.2 shows the measured two-way beam pattern by Horus during the field deployment at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025, illustrating Fourier (non-adaptive) and Capon (adaptive) beamforming methods. These beams are steered at -1° elevation and -4° azimuth, marked as black lines in each plot. The wind turbine was located at -2° elevation and 2.5° azimuth, marked as a red dot in each plot. The Fourier pattern has high power in the direction of the wind turbine, while the Capon pattern has a null (< -30 dB) in the direction of the wind turbine. These plots clearly show the advantage of clutter mitigation with PAR using adaptive beamforming, particularly for non-stationary clutter, which is difficult to remove using Doppler-only filtering with parabolic-dish-based radar systems. More analysis of

this dataset will be discussed in Chapter 5. In addition to clutter mitigation, adaptive beamforming has the potential to improve spatial data quality near storm edges or in regions with sharp power gradients.

Another unique capability of digital PAR systems is space-time processing (STP), which jointly processes spatial samples across the array and temporal samples across pulses. Non-adaptive STP uses Fourier-based weights, while space-time adaptive processing (STAP) employs data-dependent weights to minimize contamination in both spatial and temporal domains [31]. STAP has been widely used in airborne and spaceborne moving-target-indication radars [32–34]. STAP can provide distinct advantages for clutter mitigation, particularly for complex clutter such as wind turbines, which contaminate data widely in both temporal and spatial domains [35]. Their rotating blades generate large and time-varying radar cross sections and Doppler signatures that overlap with meteorological signals, making them difficult to filter using conventional Doppler-based methods. This contamination can lead to false storm signatures, erroneous rainfall estimates, and degraded data quality near wind farms [36, 37]. Furthermore, wind turbines have rapidly increased in number and size, creating a significant source of complex clutter for weather radars and becoming an increasingly serious issue for weather observations.

However, these benefits of PAR do not come for free. PARs require calibration of hundreds to thousands of beams, including mutual coupling calibration, co-polar calibration, noise calibration, and beam-pattern calibration [38–40]. In addition, advanced adaptive processing techniques such as adaptive beamforming and STAP introduce an additional layer of calibration challenges. When adaptive beamforming is applied, the beam pattern changes with every scan, producing spatially varying noise levels and complicating dual-polarization variable estimation. As a result, new calibration strategies are required before adaptive beamforming or STAP can be reliably applied to weather measurements.

Furthermore, adaptive processing methods are computationally expensive and require

a large number of independent and identically distributed samples [41], which may limit their practical application in weather observations. Adaptive processing is also sensitive to noise and the scanning environment, causing performance to vary depending on radar configuration and operating conditions. Therefore, the performance of adaptive processing under different settings and radar parameters should be systematically evaluated to better understand its impact on weather data quality and measurement accuracy.

1.3 Contributions

This dissertation addresses the challenges of adaptive beamforming and STAP through theoretical development, simulation analysis, and validation using real fully digital radar data from Horus, providing new insights for the design and implementation of adaptive processing in next-generation weather radar networks. Several key challenges motivate this work. First, adaptive methods introduce significant calibration complexity because beam patterns change with each scan, making accurate polarimetric estimation difficult. Second, the computational burden remains a major limitation, particularly for large arrays and real-time applications. This includes the need to support different array configurations such as subarray and sparse systems, as well as selective operation depending on target characteristics. Third, adaptive processing requires a sufficient number of independent samples, which is difficult to achieve in practice because increasing dwell time is undesirable. Finally, the lack of real-world validation for fully digital weather PAR systems has limited the assessment of adaptive techniques under realistic clutter and environmental conditions. By addressing these issues, this dissertation advances the practical feasibility of adaptive beamforming and space-time processing for future operational weather radar systems. The research seeks to answer the following key questions:

- How can adaptive beamforming be effectively calibrated to preserve dual-

polarization data quality in a digital PAR system?

- What performance improvements can adaptive beamforming offer over conventional Fourier beamforming in terms of clutter suppression, sidelobe mitigation, and polarimetric integrity?
- How do different array architectures and processing configurations influence the behavior and performance of adaptive processing for meteorological targets?
- Can adaptive processing mitigate complex clutter sources, such as wind turbines, that challenge current weather radar operations?
- What are the operational implications and feasibility of applying adaptive processing in real-world, time-constrained radar environments?

The overarching goal of this dissertation is to investigate and demonstrate the benefits of adaptive processing for digital phased array weather radar, emphasizing both algorithmic development and system-level evaluation. This work establishes the following objectives:

- Develop and validate a novel adaptive beamforming calibration framework that improves dual-polarization consistency while preserving adaptive pattern flexibility.
- Quantitatively compare adaptive and deterministic beamforming techniques through controlled simulations and real-world experiments using the Horus radar.
- Assess the influence of array architectures on adaptive performance through simulation-based parametric studies.
- Demonstrate the application of adaptive beamforming and STAP for mitigating non-meteorological interference such as ground clutter, moving point targets, wind turbines, using Horus dual-polarization data.

1.4 Outline

The remainder of this dissertation is organized as follows. Chapter 2 provides the background needed for understanding phased array weather radar and adaptive beamforming. It begins with the fundamentals of PAR, followed by principles of DBF and a detailed discussion of adaptive beamforming, including Capon beamforming and constrained Capon. Simulation results are presented to illustrate their behavior and limitations. Additional beamforming methods including minimum mean squared error (MMSE), Bayesian approaches, and AI-based beamforming, are reviewed through a literature survey. The chapter then introduces deterministic null-constrained beamforming, which could be an alternative beamforming method if the clutter location is known.

Chapter 3 focuses on calibration challenges in adaptive beamforming. Section 3.1 introduces a beam-pattern calibration method derived from the weather radar equation, followed by Section 3.2, which presents a noise estimation technique essential for adaptive beamforming due to its dynamically changing noise levels. Simulation studies and quantitative analysis are provided throughout. Differences between horizontally and vertically polarized channels are examined, and the effects of variations in beamforming weights and patterns on polarimetric performance are studied. Furthermore, calibrated polarimetric variables from the simulation are evaluated to validate the proposed calibration method and to facilitate a comparison between adaptive and deterministic beamforming, thereby demonstrating both improvements and remaining challenges.

Chapter 4 presents the application of STAP. It starts with the method of STAP, then uses simulations to investigate how array architecture, required independent samples, diagonal loading, and transmit beam synthesis influence STAP performance. The simulated data contain both weather and a moving point target, demonstrating how STAP reduces point target contamination and improves meteorological velocity estimates. This chapter

also includes wind-turbine simulations and demonstrates how STAP mitigates wind-turbine clutter.

Chapter 5 validates the proposed methods using measurements from the Horus radar. Adaptive processing was implemented in real data, and the proposed calibration method was applied. The dataset includes various scenarios, such as weather with airplanes, weather with ground clutter, severe storms, and wind turbines. The use of dual-polarized PAR data enables evaluation of polarimetric variables, which is a unique capability.

Chapter 6 focuses on improving the feasibility of adaptive processing for real-world implementation. To address computational cost, several approaches were explored, including conditional adaptive processing (applied only when needed), and reducing the covariance matrix using subarrays and sparse arrays. Beam synthesis was also investigated for grating lobe mitigation by shaping the transmit beam to reduce sidelobes and grating lobes; this is a unique capability of digital PAR. Finally, a clutter filtering algorithm combined with adaptive processing was proposed.

Finally, Chapter 7 provides concluding remarks and summarizes the key contributions of this dissertation, including the implementation of adaptive processing for dual-polarized PAR data, addressing calibration issues of adaptive beam patterns, and analyzing the performance of adaptive processing both qualitatively and quantitatively across various environments and radar configurations. It also explores potential approaches for real-time implementation. Remaining limitations are discussed, including challenges in ρ_{hv} bias and constraints on real-time feasibility, such as the large number of samples required and the resulting longer dwell time. Lastly, directions for future research are outlined. While adaptive processing itself is a mature technique, the focus is placed on its practical use in next-generation weather radar systems.

Chapter 2

Unique Capabilities of Phased Array Radar

PAR technology was first conceived in the early 1900s, but major development was driven by the need for advanced air defense systems during World War II [42]. Over the past several decades, PAR technology has matured extensively in military air-surveillance and defense applications, making it increasingly feasible for civilian applications such as weather observation [43]. The first PAR adapted for meteorological observations appeared in 1997, when the S-band SPY-1 radar aboard a United States Navy (U.S. Navy) destroyer was used to observe weather phenomena [44, 45]. Following these initial demonstrations, the National Severe Storms Laboratory (NSSL) partnered with government agencies, private industry, and universities to develop the NWRT with a modified SPY-1 antenna to evaluate PAR technology for next-generation operational weather radar networks [17]. NWRT has a flat-panel phased array antenna, slotted-waveguide architecture, vertically polarized radiators, a rotating platform, and a modified NEXRAD transmitter. Around the same time, other single-polarization PAR systems were developed, including mobile PARs, enabling observations of severe weather with unprecedented temporal resolution [14, 46, 47]. These early demonstrations highlighted the potential advantages of PAR for meteorological observations, particularly the ability to rapidly scan storms and capture fast-evolving weather processes.

2.1 Fundamentals of Phased Array Radar

PAR fundamentally differs from conventional parabolic-reflector-based weather radars in the way it forms and steers beams. Instead of relying on a mechanically rotated reflector antenna, a PAR consists of many radiating elements, each capable of transmitting and receiving electromagnetic waves. In a uniform linear array with N antenna elements, the spacing between adjacent antenna elements is typically $d = \lambda/2$ to mitigate grating lobes, as illustrated in Figure 2.1.




Figure 2.1: Fundamental concepts of a PAR using a uniform linear array. Color shading represents the relative gain, where blue tones represent lower values and red tones represent higher values [43].

When the antenna elements are excited with sinusoidal signals, the array produces directions of constructive and destructive interference [43]. Close to the antenna elements, the beam forms in the reactive near field. As the wave moves farther away, it propagates through the radiative near field. At larger distances, the wavefronts become nearly planar and propagate in the far field. The angular distribution of radiated power in the far

field is called the antenna radiation pattern, which characterizes the electromagnetic radiation properties of the antenna. Reflector antenna radars also have radiation patterns with limited flexibility in their shapes. In contrast, digital PAR systems can synthesize beams electronically and steer them almost instantaneously by precisely controlling the phase and amplitude at each element [17].

The direction of maximum constructive interference is called the mainlobe, and it determines the beam pointing direction. Other directions with partial constructive interference form sidelobes. Sidelobes are undesirable because they can contaminate measurements observed through the mainlobe. All practical antennas, including those used in WSR-88D radars, have a mainlobe surrounded by sidelobes. To steer the mainlobe of a PAR electronically, the phase difference between adjacent antenna elements is changed. The phase of a signal is proportional to its propagation time, expressed as $\phi = 2\pi ft$, where f is the radar frequency and t is time. A phase difference between elements can therefore be interpreted as a time delay. To steer the beam to an angle θ relative to the broadside direction, the phase fronts from all elements must align to produce constructive interference. The broadside direction means the direction perpendicular to the array plane. This is achieved by introducing a constant phase difference between adjacent elements. The required phase difference is $\varphi = kd \sin \theta$, where $k = 2\pi/\lambda$ is the wavenumber. When the beam is steered away from the broadside direction, the width of the mainlobe increases. This increase in beamwidth leads to reduced angular resolution. The change occurs because the apparent antenna aperture becomes smaller when scanning off broadside. When the radar scans at broadside, the full antenna aperture is used. When scanning at an angle, only the projection of the aperture onto a plane perpendicular to the steering direction is used.

2.2 Advantages of Phased Array Radar

The flexibility of the electronically controlled elements leads to two primary advantages. First, PAR enables agile beam steering, with repositioning times on the order of microseconds [43]. This agility makes PAR especially promising for weather surveillance, where the rapid evolution of severe storms demands faster temporal sampling than what mechanically scanned radars can provide. Second, digital PAR supports imaging modes, where multiple receive beams can be formed simultaneously from a single wide transmit beam. Imaging is performed through DBF, which coherently combines element-level signals with complex weights chosen to point the beam in specific directions.

Imaging is one of the most powerful scanning modes enabled by PAR. Unlike a reflector radar, which transmits and receives a single narrow pencil beam, an imaging PAR transmits a broad illumination, sometimes referred to as a “spoiled” or wide transmit beam, and then forms many narrow receive beams simultaneously using DBF [23, 43, 46]. This approach dramatically increases the sampling rate across a sector, improving temporal resolution while maintaining spatial sampling. This capability allows imaging PARs to capture rapidly evolving features such as low-level divergence, rotation in reintensifying supercells, early microburst signals that cannot be resolved by a reflector radar such as WSR-88D, and the internal structures of hailstorms associated with rapid growth and downdraft initiation [12, 48, 49].

One example of a potential artifact from the use of electronically steered beams is beam steering bias, which occurs when the beam is steered away from broadside and is caused by co-polar and cross-polar antenna patterns. This bias can affect reflectivity and polarimetric variables, such as differential reflectivity, differential phase, and the cross-correlation coefficient. Another example occurs when radar imaging is used, which can increase two-way sidelobe levels and introduce large reflectivity biases when scanning regions with strong

gradients. Furthermore, when adaptive beamforming is performed, beam shapes change at each scanning angle depending on the scanning environment.

DBF can be performed using either deterministic or adaptive methods. Deterministic beamforming produces fixed patterns by the Fourier method, based purely on the array layout and the desired steering angle. Adaptive methods, such as the Capon method, Minimum Mean Square Error (MMSE), and Bayesian approaches, generate their weights from the received signals to minimize contamination from interference [30]. Therefore, adaptive patterns are not fixed and change dynamically depending on the environment. There are also other robust adaptive beamforming methods. In this dissertation, the primary focus is on Fourier and Capon DBF methods, with additional emphasis on STAP, while MMSE and Bayesian approaches are included mainly for contextual comparison. Fundamentals of DBF are discussed in more detail in later sections of this chapter, including the principles of imaging and conventional DBF, adaptive beamforming methods and their limitations, and deterministic null-constrained techniques.

2.3 Planar Array Geometry and Architecture

Array geometry, such as the overall aperture shape, diameter, and element spacing, is one of the fundamental factors that determine the resulting beam shape. The beam shape is dictated by the spatial arrangement of the elements, referred to as the array factor. Under the assumption of ideal elements, where each element radiates the same electromagnetic field pattern, the array factors for different array architectures are illustrated in Figure 2.2. The top row shows the element layouts for (1) a square 40×40 array, (2) a truncated cross-shaped array, (3) a rectangular 8×40 array, and (4) a circular array.

The middle row shows their corresponding three-dimensional (3D) array factors (in dB), where the influence of aperture shape on the overall beam structure is evident. The




Figure 2.2: Four array configurations are evaluated: square, truncated, rectangular, and circular arrays. The upper panels provide the array layouts, the middle panels show the resulting 3D array factors, and the bottom panels provide the 1D azimuth cuts.

square array exhibits a fully symmetric pattern in both angular dimensions. The truncated array shows a broadened mainlobe, as its effective aperture diameter is reduced. It also produces a different sidelobe structure compared to the square array because removing elements from the corners changes the aperture shape, resulting in an uneven and less symmetric weighting function. The rectangular array produces a narrow mainlobe along the long dimension (x -axis) and a wide mainlobe along the short dimension (y -axis). When more elements are arranged along a particular dimension, the radiated fields from those elements add coherently, effectively increasing the aperture in that direction. As a result, the mainlobe becomes narrower where more elements are summed, and wider where fewer elements contribute. The circular aperture yields a symmetric pattern with the lowest sidelobes and a slightly broadened mainlobe compared to the square array. This behavior follows from the fact that array patterns are obtained from the Fourier transform of the aperture weighting function [50]. The sharp discontinuities at the edges of the square aperture lead to increased

sidelobe levels, whereas the circular aperture effectively introduces a smoother transition at the boundary, acting like a natural taper in the angular domain and therefore reducing sidelobes. This tapering effect also slightly broadens the mainlobe, since the effective aperture diameter is reduced.

The bottom row provides one-dimensional (1D) cuts of the array factors, clearly showing how the mainlobe width and sidelobe behavior change with geometry. These comparisons highlight that different array architectures play a critical role in shaping the beam pattern, directly impacting mainlobe characteristics, sidelobe structure, and ultimately the performance of observations.

In the following subsections, subarray and sparse array will be discussed. Subarray and sparse array are one ways to reduce the covariance dimension, which becomes important later when operational feasibility is considered.

2.3.1 Subarrays

A subarray system groups multiple elements together, summed in analog, and the outputs are digitized for further processing and beamforming. Depending on the system requirements, subarrays can be a practical alternative to a fully digital PAR. Fully digital PAR systems often suffer from the massive volume of received data, which leads to challenges in data transfer, storage, and real-time processing. Using subarrays can dramatically reduce the processing burden. For example, performing adaptive digital beamforming with a subarray architecture reduces the size of the covariance matrix, which directly lowers the computational cost of matrix inversion—one of the major limitations that inhibits real-world adaptive processing. This reduction in covariance dimension becomes particularly important later when operational feasibility is considered in Chapter 6. However, subarray architectures can introduce undesired grating lobes because the effective spacing between

subarray phase centers becomes larger than $\lambda/2$. This violates the spatial Nyquist sampling condition, which requires the array to sample the wavefield at least twice per wavelength to avoid aliasing. When this condition is not met, spatial aliasing appears in the form of grating lobes in the array factor.

Figure 2.3 shows the subarray system formed from the truncated array. Subarray systems can be constructed in various ways. Here, the subarray configuration is based on the Horus system, which consists of 13 panels, each panel containing 40×40 elements. The top row shows a system with non-overlapped subarrays consisting of 8 elements (1×8 columns). Its effective element spacing is 4λ in the vertical direction and $\lambda/2$ in the horizontal direction, so grating lobes appear in the θ_y direction. The middle row shows a system with non-overlapped subarrays consisting of 32 elements (4×8 columns), with spacings of 4λ vertically and 2λ horizontally, which results in grating lobes appearing in both (θ_x, θ_y) directions. The bottom row shows a system with non-overlapped subarrays consisting of 64 elements (8×8 columns), with spacings of 4λ vertically and 4λ horizontally.

Even though grating lobes appear when beamforming is performed on subarray systems, these lobes can be suppressed in the two-way pattern as long as they lie outside the transmit beam's mainlobe (i.e., outside the radar field of view). Therefore, subarrays can still be an effective option for reducing computational time in adaptive beamforming and lowering the overall data volume, provided they are used in combination with properly designed transmit beams.

2.3.2 Sparse Array

Sparse arrays are another architecture that can reduce the data volume and the computational burden of adaptive beamforming while still maintaining a desired beamwidth from a large effective aperture. A sparse array uses fewer physical elements than a fully populated




Figure 2.3: Non-overlapped subarrays consisting of 8 elements (1×8 columns), 32 elements (4×8 columns), and 64 elements (8×8 columns) are illustrated above with their corresponding array factors.

array, but the elements are typically placed non-uniformly across the aperture. The idea is that by optimizing the element positions, the array can preserve mainlobe and sidelobe characteristics while dramatically reducing the number of elements that need to be digitized and processed. A well-designed sparse array can offer a good compromise between aperture size and computational load.

Designing sparse arrays typically involves optimizing the element positions under constraints such as the total number of elements, aperture dimensions, and minimum allowed element spacing. Unlike traditional thinning methods, which remove elements from a uniform grid, sparse array optimization provides more degrees of freedom and can therefore achieve lower peak sidelobe level (PSL) when the same aperture and element count are used. Several studies have demonstrated that sparse planar arrays optimized using genetic-algorithm-based methods can significantly reduce PSL while maintaining the aperture constraints [51, 52]. These methods avoid infeasible element placements (such as violating spacing constraints) and efficiently search for element configurations that balance aperture coverage with sidelobe performance.

Furthermore, sparsity can be useful for adaptive beamforming and STAP. Sparse-constrained methods enforce sparsity on the adaptive weight vector, improving convergence and reducing the need for large data sets. These approaches have shown potential to lower computational requirements and improve robustness in practical airborne phased-array radar environments [53].

Figure 2.4 shows the truncated PAR with different sparsity levels: 10%, 30%, 60%, and 90%, which indicate the percentage of elements removed. This sparsity is applied at the panel level to resemble the Horus architecture. Each panel drops the same percentage of elements, depending on the sparsity level, and the element removal is done randomly. The first column shows the array configuration for each sparsity level, the second column shows the top view of the patterns, the third column shows the 3D view of the patterns, and the




Figure 2.4: This figure shows the truncated PAR with different sparsity levels: 10%, 30%, 60%, and 90%. The first column shows the array configuration for each sparsity level, the second column shows the top view of the patterns, the third column shows the 3D view of the patterns, and the fourth column shows the 1D azimuth and elevation cuts. For the 1D pattern plots, the red dotted lines represent the beam patterns from the fully populated array.

fourth column shows the 1D azimuth and elevation cuts. For the 1D pattern plots, the red dotted lines represent the beam patterns from the fully populated array, allowing a direct comparison of how the beams change with sparsity.

The results indicate that the mainlobe does not become noticeably distorted even when 90% of the elements are removed. Although the mainlobe broadens slightly, the change is less than 1° . The sidelobe levels within $\pm 30^\circ$ remain relatively stable up to 60% sparsity, with no significant increase. However, when 90% of the elements are removed, the sidelobe level increases by roughly 10 dB. Overall, these results demonstrate the potential of sparse arrays. They suggest that sparsity can be a viable approach for reducing the size of the covariance matrix, improving convergence and reducing the overall data volume while still maintaining reasonable beam quality.

2.4 Imaging with Digital Beamforming: Principles and Methods

Imaging is a technique in which the radar transmits a wide beam and forms narrow receive beams in different directions simultaneously using DBF [43]. This approach allows the radar to scan a large sector of the atmosphere, thereby improving temporal resolution. DBF can be performed using either deterministic or adaptive methods. The conventional approach is deterministic, where the beamforming weights are pre-calculated based on the array element geometry and the steering angle. As a result, the produced beam is independent of the scanning environment and maintains a fixed beam shape at each steering angle, although the beam becomes broadened as the radar steers away from broadside.

In this conventional method, the beamforming weights are computed by applying the necessary phase shifts to steer the beam in the desired direction, $\mathbf{k}$, which is referred to as

the Fourier method. The desired direction is defined as

$$\mathbf{k} = \frac{2\pi}{\lambda} [\cos \theta \sin \phi \quad \cos \theta \cos \phi \quad \sin \theta],$$

where θ and ϕ denote elevation and azimuth angles, respectively. The Fourier beamforming weights are then defined as

$$\mathbf{w}_f = [e^{j\mathbf{k} \cdot \mathbf{D}_1} \quad e^{j\mathbf{k} \cdot \mathbf{D}_2} \quad \dots \quad e^{j\mathbf{k} \cdot \mathbf{D}_N}]^T, \quad (2.1)$$

where $\mathbf{D}_i$ is the position vector of the i^{th} element ($i = 1, 2, \dots, N$), and T denotes the transpose operation. These weights form the conventional steering vector.

Deterministic DBF techniques use a fixed beam pattern, which makes real-world implementation much easier than adaptive beamforming because they are computationally inexpensive and relatively easy to calibrate. However, they also have inherent limitations due to the fixed beam shape, which does not respond to changes in the environment. These fixed sidelobe levels often lead to contamination from interference and ground clutter. This problem becomes even more significant for weather radars, where targets are distributed within the wide transmit beam. In this case, leakage between adjacent receive beams becomes unavoidable and can significantly degrade the spatial quality of the received data. These limitations motivate the need for adaptive beamforming methods, which are introduced next to address environment-dependent interference and improve spatial data quality.

2.5 Adaptive Beamforming

Adaptive beamforming addresses the limitations of deterministic DBF by allowing the receive beam patterns to change depending on the received data. Instead of using a fixed steering vector, adaptive methods, such as the Capon method, estimate the covariance ma-

trix of the received signals and choose weights that minimize the output power while preserving gain in the desired steering direction. In practice, this enables the beamformer to place deep nulls toward strong interference, suppress clutter or grating-lobe responses, and sharpen the mainlobe. As a result, adaptive beamforming can substantially improve the signal-to-interference-plus-noise ratio (SINR) and the effective angular resolution.

There are many adaptive beamforming methods that share the same core idea—minimizing the output power—but differ in how the optimization is formulated and how robustness is enforced. In this chapter, we focus on several representative approaches: the Capon beamformer, robust Capon extensions with linear, quadratic, and diagonal-loading constraints, MMSE beamforming, Bayesian adaptive beamforming under direction-of-arrival (DoA) uncertainty, and AI-based fast Capon. Together, these methods illustrate both the potential of adaptive beamforming to improve spatial quality and the practical trade-offs in terms of computational cost, sensitivity to steering-vector mismatch, and covariance-estimation requirements.

2.5.1 Capon Beamforming

The Capon method is one of the most well-known adaptive beamforming methods [30]. It improves resolution, and its SINR performance is much better than that of conventional data-independent beamformers. It adaptively calculates weights by solving an optimization problem that minimizes the received power in all directions except the one of interest, thereby suppressing sidelobes and improving the quality of spatial data. In a theoretically ideal situation where the array, representing the spatial sampling window, is infinitely large, the beam would be a delta function. However, in real-life applications the array size is limited; therefore, a sinc function is expected with lowered sidelobe levels, or even nulls, in the direction of high-power return. The method starts with the optimization problem

stated below [28, 30]:

$$\min_{\mathbf{w}} \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{subject to} \quad \mathbf{w}^H \mathbf{a}(\theta) = 1, \quad (2.2)$$

where H denotes the Hermitian (conjugate transpose), $\mathbf{R}$ is the $N \times N$ autocovariance matrix of the received signals ($\mathbf{V}$, size of $N \times M$) from N receivers, defined as $\mathbf{R} = \mathbb{E}\{\mathbf{V}\mathbf{V}^H\}$, where $\mathbb{E}\{\cdot\}$ denotes the statistical expectation operator. $\mathbf{a}(\theta)$ is the steering vector for the desired angle θ , equivalent to that defined in (2.1).

This optimization problem is solved with the Lagrange multiplier (Λ) method [30]:

$$\mathcal{L}(\mathbf{w}, \Lambda) = \mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda(1 - \mathbf{w}^H \mathbf{a}(\theta)) \quad (2.3)$$

Taking the gradient of $\mathcal{L}$ with respect to $\mathbf{w}$ and setting it to zero, results in:

$$\nabla_{\mathbf{w}} \mathcal{L} = \frac{\partial}{\partial \mathbf{w}} (\mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda(1 - \mathbf{w}^H \mathbf{a}(\theta))) = 0 \quad (2.4)$$

$$2\mathbf{R}\mathbf{w} - \Lambda\mathbf{a}(\theta) = 0 \quad (2.5)$$

$$\mathbf{w} = \frac{\Lambda}{2} \mathbf{R}^{-1} \mathbf{a}(\theta) \quad (2.6)$$

Using the constraint, $\mathbf{w}^H \mathbf{a}(\theta) = 1$, we solve for the Lagrange multiplier:

$$\left(\frac{\Lambda}{2} \mathbf{R}^{-1} \mathbf{a}(\theta) \right)^H \mathbf{a}(\theta) = 1 \quad (2.7)$$

$$\frac{\Lambda}{2} \mathbf{a}^H(\theta) \mathbf{R}^{-1} \mathbf{a}(\theta) = 1 \quad (2.8)$$

$$(2.9)$$

Therefore, the Lagrange multiplier becomes:

$$\Lambda = \frac{2}{\mathbf{a}^H(\theta)\mathbf{R}^{-1}\mathbf{a}(\theta)}. \quad (2.10)$$

Replacing the Lagrange multiplier in (2.6) with (2.10), the resulting optimized weights by the Capon method are:

$$\mathbf{w}_c = \frac{\mathbf{R}^{-1}\mathbf{a}(\theta)}{\mathbf{a}^H(\theta)\mathbf{R}^{-1}\mathbf{a}(\theta)} \quad (2.11)$$




Figure 2.5: Simulated beam patterns obtained using the Fourier method (left) and the Capon method (right) are shown above. The blue arrows indicate the steering direction, and the black arrows indicate the direction of the interference.

Simulated beam patterns by the Fourier and Capon methods are illustrated in Figure 2.5. The figure shows the 3-dimensional view of the beams and the top view, with the blue arrows indicating the direction of the desired steering angle and the black arrows indicating the direction of the interference. The patterns are calculated by

$$b(\mathbf{k}) = \left| \sum_{n=1}^N \mathbf{w}_n e^{-j\mathbf{k} \cdot \mathbf{D}_n} \right|^2. \quad (2.12)$$

The Fourier beam depends only on the desired steering angle and the array architecture.

Therefore, even when interference is present within the mainlobe, the Fourier beam shape is not affected. In contrast, the Capon beam places a deep null in the direction of the interference. Since the interference is located inside the mainlobe, the mainlobe shape becomes distorted, unlike the Fourier beam, and the overall sidelobe structure is also altered. Note that this figure is intended only to illustrate the conceptual difference between the Capon and Fourier methods; more detailed beam pattern results are presented later.

2.5.2 Robust Capon Beamforming with Linear Constraints

The Capon beamformer provides higher angular resolution and substantially improved interference suppression compared with conventional data-independent (e.g., Fourier) beamformers, provided that the steering vector of the signal of interest is accurately known. However, its performance is well known to be highly sensitive to steering-vector mismatch. In particular, when a strong interferer lies within or close to the mainlobe region, the adaptive optimization tends to place a deep null in the direction of the interference. As a consequence, the mainlobe can become severely distorted (see Figure 2.5, for example). Furthermore, the beam pattern may no longer have its peak aligned with the desired steering direction, even though the distortionless constraint still enforces unity gain in the desired steering direction. Another issue resulting from this distorted mainlobe is that the 3-dB beamwidth can change. This issue is particularly critical for weather observations, as the targets are distributed hydrometeors, and weather radar processing relies on the averaged volume information within the beam, which can lead to biased or inconsistent meteorological measurements.

To preserve the shape of the mainlobe, meaning maintaining the peak at the desired steering angle and the 3-dB beamwidth, additional constraints can be added to the Capon optimization method. Common approaches include linearly constrained methods, quadrat-

ically constrained methods, and diagonal loading.

The robust Capon method with linear constraints starts with the optimization problem stated as follows [54]:

$$\min_{\mathbf{w}} \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{subject to} \quad \mathbf{w}^H \mathbf{C} = \mathbf{g}, \quad (2.13)$$

where $\mathbf{C} = [\theta_1, \theta_2, \dots, \theta_I]$ denotes the constraint matrix, which has full rank equal to I , and $\mathbf{g}$ is defined as the I -dimensional vector of desired gains. In the original Capon method, the constraint enforces a unity gain in the desired steering direction. The linearly constrained Capon method imposes constraints at specified angles with specified gains $\mathbf{g}$.

Using the method of Lagrange multipliers. (Λ):

$$\mathcal{L}(\mathbf{w}, \Lambda) = \mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda(1 - \mathbf{w}^H \mathbf{C}), \quad (2.14)$$

Taking the gradient of $\mathcal{L}$ with respect to $\mathbf{w}$ and setting it to zero

$$\nabla_{\mathbf{w}} \mathcal{L} = \frac{\partial}{\partial \mathbf{w}} (\mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda (1 - \mathbf{w}^H \mathbf{C})) = 0, \quad (2.15)$$

$$2\mathbf{R} \mathbf{w} - \Lambda \mathbf{C} = 0, \quad (2.16)$$

$$\mathbf{w} = \frac{\Lambda}{2} \mathbf{R}^{-1} \mathbf{C}. \quad (2.17)$$

Using the constraint, $\mathbf{w}^H \mathbf{C} = \mathbf{g}$:

$$\left(\frac{\Lambda}{2} \mathbf{R}^{-1} \mathbf{C} \right)^H \mathbf{C} = \mathbf{g}, \quad (2.18)$$

$$\frac{\Lambda}{2} \mathbf{C}^H \mathbf{R}^{-1} \mathbf{C} = \mathbf{g}. \quad (2.19)$$

$$(2.20)$$

Therefore, the Lagrange multiplier becomes:

$$\Lambda = \frac{2g}{C^H \mathbf{R}^{-1} C}. \quad (2.21)$$

The resulting optimized weights are:

$$\mathbf{w}_{\text{LCC}} = \frac{\mathbf{R}^{-1} C}{C^H \mathbf{R}^{-1} C} g, \quad (2.22)$$

2.5.3 Robust Capon Beamforming with Quadratic Constraints

Another robust Capon method that preserves the mainlobe shape uses quadratic constraints [55–57]. Unlike the linearly constrained method, the quadratically constrained method imposes bounded constraints. Its optimization problem is formulated as:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{subject to} \quad \mathbf{w}^H \mathbf{a}(\theta) = 1, \\ & \mathbf{w}^H \mathbf{C}_i \mathbf{w} \leq \gamma_i, \quad i = 1, \dots, I \end{aligned} \quad (2.23)$$

where γ denotes the bound constraint. This optimization problem combines both linear and quadratic constraints, which leads to an infinite number of feasible solutions. To obtain a practical (“suboptimal”) solution, the Karush-Kuhn-Tucker conditions [58] should be implemented.

To pose this method, first we define the Lagrangian with two multipliers (Λ, μ) :

$$\mathcal{L}(\mathbf{w}, \Lambda, \mu_i) = \mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda (1 - \mathbf{w}^H \mathbf{a}(\theta)) + \sum_{i=1}^I \mu_i (\mathbf{w}^H \mathbf{C}_i \mathbf{w} - \gamma_i). \quad (2.24)$$

Taking the gradient of $\mathcal{L}$ under KKT conditions leads to:

$$\nabla_{\mathbf{w}} \mathcal{L} = \frac{\partial}{\partial \mathbf{w}} \left(\mathbf{w}^H \mathbf{R} \mathbf{w} + \Lambda (1 - \mathbf{w}^H \mathbf{a}(\theta)) + \sum_{i=1}^M \mu_i (\mathbf{w}^H \mathbf{C}_i \mathbf{w} - \gamma_i) \right) = 0. \quad (2.25)$$

$$2\mathbf{R}\mathbf{w} - \lambda \mathbf{a}(\theta) + \frac{\partial}{\partial \mathbf{w}} \left(\sum_{i=1}^M \mu_i (\mathbf{w}^H \mathbf{C}_i \mathbf{w} - \gamma_i) \right) = 0. \quad (2.26)$$

Solving this requires an iterative process. Through the iterative process, it optimizes the μ that satisfies the quadratic constraint, namely, $|\mathbf{C}^H \mathbf{w}| - 1 < \gamma$. At every iteration, μ is updated: $\mu_i = \mu_{i-1} + c\gamma$.

The “suboptimized” weight is:

$$\mathbf{w}_{\text{QCC}} = \frac{\mathbf{R}^{-1} \mathbf{a}(\theta) + \mu \mathbf{C}}{\mathbf{a}^H(\theta) \mathbf{R}^{-1} \mathbf{a}(\theta) + \mu \mathbf{a}^H(\theta) \mathbf{R}^{-1} \mathbf{a}(\theta)}. \quad (2.27)$$




Figure 2.6: The patterns from the Fourier (blue), Capon (red), linearly constrained Capon (yellow), and quadratically constrained Capon (purple) are simulated under the no-interference condition and plotted above.

Figure 2.6 shows the patterns from the Fourier (blue), Capon (red), linearly constrained Capon (yellow), and quadratically constrained Capon (purple) under the no-interference

condition. In the left plot, the linear constraint was added at the 3-dB beamwidth ($\mathbf{C} = [a(\theta), 10^\circ, -10^\circ]$) with $\mathbf{g} = [1, 0.5, 0.5]$. However, because we attempted to place constraints exactly at the 3-dB beamwidth to preserve the beamwidth, the resulting beam pattern becomes unstable, producing a very large sidelobe immediately after the constraint points, as the energy is forced away from the constrained angles. This sidelobe becomes even higher than the mainlobe, and the resulting beamwidth is not preserved either. This suggests that linear constraints may not be the best choice for preserving the mainlobe shape, since linear constraints act more like fixing individual points (discontinuities). They may, however, be useful for other purposes, such as synthesizing beam patterns with multiple mainlobes or enforcing fixed nulls at specified directions.

In the right plot, we applied linear constraints at three different steering directions, $\mathbf{C} = [a(\theta), 42^\circ, -42^\circ]$ with $\mathbf{g} = [1, 1, 1]$. In this case, the result successfully forms three mainlobes in the three specified directions. This demonstrates the potential usefulness of linear constraints for certain types of beam synthesis [20]. The quadratic constraints for both plots are placed at the 3-dB beamwidth ($\mathbf{C} = [10^\circ, -10^\circ]$) with bound constraint of $\gamma = 1 \times 10^{-10}$ for both constraint points, in addition to the unity-gain constraint in the steering direction. The solution was optimized with 1000 iterations, and the initial Lagrange multiplier for the bound constraints (μ) was set to 0.1. In this spatially white, no-interference condition, the Fourier, Capon, and quadratically constrained Capon methods produce essentially the same beam patterns.

Another recent approach to constraining the mainlobe and avoiding beam pattern distortion introduces additional constraints to the Capon beamforming optimization problem [59]. These constraints are designed to preserve the shape and effective area of the beam-pattern mainlobe so that the resulting adaptive beams remain consistent with the calibration assumptions used in weather radar signal processing. Two types of constraints are imposed in this work. First, a mainlobe area constraint ensures that the integrated power

within the mainlobe region of the adaptive beam matches that of a reference beam formed using tapered Fourier beamforming. This guarantees that the adaptive beam produces the same calibration constant as the reference beam. Second, angular constraints are applied at selected angles within the mainlobe region to ensure that the adaptive beam closely follows the mainlobe shape of the reference beam and prevents shifts in the peak direction of the beam.

2.5.4 Robust Capon Beamforming with Diagonal Loading

In practical adaptive beamforming, the estimated covariance matrix, $\hat{\mathbf{R}}$, is often formed from a limited number of received data. When the number of independent samples is small, the estimated matrix becomes poorly conditioned, and the resulting eigenvalue spread is dominated by randomness in the noise subspace. As discussed in [60], this large noise-eigenvalue spread leads to undesirable subtraction of randomly shaped noise eigenbeams during adaptation, distorting both the mainlobe and the sidelobe structure of the beam pattern.

To mitigate this instability, the covariance matrix is modified using *diagonal loading*, expressed as

$$\mathbf{R}_L = \hat{\mathbf{R}} + L\mathbf{I}, \quad (2.28)$$

where L is a non-negative loading level and $\mathbf{I}$ is the identity matrix. Adding L increases all eigenvalues uniformly while leaving the eigenvectors unchanged. Strong interference eigenvalues remain essentially unaffected, whereas small noise eigenvalues are increased and compressed. This reduces the influence of noise eigenbeams in the adaptive solution and stabilizes the subsequent weight computation.

Simulation results in [60] show that diagonal loading significantly improves mainlobe fidelity and reduces highest sidelobe levels, particularly when the number of sample is

small. With loading, the adapted pattern approaches the ideal quiescent beam shape while still maintaining deep nulls in the interference directions. Moderate loading levels produce substantial robustness improvements without significantly degrading interference suppression.

Diagonal loading, however, introduces a fundamental tradeoff. By raising the noise-eigenvalue floor, the array becomes slightly less sensitive to weak interferences whose eigenvalues lie near the loading level. This desensitization is typically acceptable, since large interference eigenvalues are minimally affected, and the resulting robustness enhancements outweigh the loss in nulling capability against very weak sources. Overall, diagonal loading remains an effective and widely adopted technique for improving the stability and performance of adaptive beamformers under imperfect covariance estimation.

Simulations were conducted to generate beam patterns, shown in Figure 2.7, using several DBF methods, including: the Fourier method (blue), Capon (orange), robust Capon with linear constraints (yellow, referred to as the LCC beam pattern), quadratic constraints with diagonal loading (purple, referred to as QCC beam pattern 1), quadratic constraints solved with an iterative method (green, referred to as QCC beam pattern 2), and quadratic constraints solved with CVX (light blue, referred to as QCC beam pattern 3). In these simulations, interference is present at 10° (indicated as a gray line in each plot). CVX is a system for disciplined convex programming [61]. For the linear constraints, multiple constraint points are placed along the 3-dB beamwidth, with manually selected gains to preserve both the 3-dB beamwidth and the mainlobe shape. Beams are plotted for three different cases: Case I, where the interference is located outside the mainlobe (beam steering at 18°); Case II, where the interference is located inside the mainlobe (beam steering at -2°); and Case III, where the beam is steered exactly at the interference (beam steering at -10°). In all of these scenarios, the Fourier beam pattern remains unaffected by the interference. Because it is data-independent, it does not mitigate the contamination and simply

maintains its predetermined beam pattern defined by the array architecture and the desired steering angle.




Figure 2.7: Simulated beam patterns from the Fourier method (blue), Capon (orange), robust Capon with linear constraints (yellow, referred to as the LCC beam pattern), quadratic constraints with diagonal loading (purple, referred to as QCC beam pattern 1), quadratic constraints solved with an iterative method (green, referred to as QCC beam pattern 2), and quadratic constraints solved with CVX (light blue, referred to as QCC beam pattern 3) are presented above. Interference is present at 10° (indicated as a gray line in each plot). Beams are plotted for three cases: Case I, where the interference is located outside the mainlobe (beam steering at 18°), Case II, where the interference is located inside the mainlobe (beam steering at -2°), and Case III, where the beam is steered exactly at the interference (beam steering at -10°).

In Case I, where the interference is located outside the mainlobe, the Capon beamformer correctly places a null at the interference direction. As a result, the sidelobe immediately adjacent to the interference becomes higher than the corresponding sidelobe in the Fourier pattern. However, the mainlobe remains undistorted because the interference is sufficiently

far from it. The LCC beam pattern is able to preserve the shape of the mainlobe and the 3-dB beamwidth. However, since multiple linear constraint points are concentrated only around the 3-dB beamwidth, the energy outside this constrained region becomes excessively large, which is a known limitation of linearly constrained Capon methods. At this interference location, placing a null does not particularly improve data quality. QCC beam patterns 1 and 3 both place a null at the interference direction and have overall shapes similar to the Capon method. QCC beam pattern 2 also places a null, but its mainlobe broadens noticeably.

In Case II, where the interference lies inside the mainlobe, the Capon method again places a null, which leads to a distorted mainlobe. All three QCC beam patterns also place a null at the interference location, and all of them exhibit distorted mainlobes as well. Among these, QCC beam pattern 2 shows the most severe distortion. QCC beam pattern 3 produces a pattern that closely resembles that of the Capon method. QCC beam pattern 1 yields a narrower mainlobe compared to the others, and the sidelobe adjacent to the mainlobe increases in height. In Case III, the beam is steered exactly at the interference direction. In this situation, the desired behavior is to avoid placing a null at the steering direction and to maintain unity gain. All methods satisfy these requirements, producing a mainlobe at the steering angle without nulls. QCC beam patterns 1 and 3 produce beam patterns similar to both the Capon and Fourier methods. However, QCC beam pattern 2 shows differences in its sidelobe structure.

As shown in Figure 2.7, the adaptive beam patterns change depending on the location of the interference and the steering angle. Their beamwidths vary, and the peak of the mainlobe does not always align with the desired steering angle. To analyze their performance quantitatively, we measure the half-power beamwidth (3-dB) and compute the difference between the desired pointing angle and the actual peak direction of each beam. Figure 2.8 presents these results: the measured 3-dB beamwidth (left plot) and the difference between

the desired and actual peak directions (right plot).




Figure 2.8: The measured 3-dB beam power (left plot) and the difference between the desired and actual peak directions (right plot) are presented above. Six beamformers are used: the Fourier method (blue), Capon (red), robust Capon with linear constraints (yellow, referred to as the LCC beam pattern), quadratic constraints with diagonal loading (purple, referred to as QCC beam pattern 1), quadratic constraints solved with an iterative method (green, referred to as QCC beam pattern 2), and quadratic constraints solved with CVX (light blue, referred to as QCC beam pattern 3)

The 3-dB beamwidth of the Fourier beamformer (blue) changes only slightly as it steers away from broadside. This behavior is a fundamental characteristic of PARs: the beam broadens as it steers off broadside. The LCC beam pattern (yellow) maintains a constant 3-dB beamwidth because it is strictly constrained to preserve the mainlobe shape. The Capon (red) and QCC3 (dotted light blue) beam patterns exhibit the largest variations when steering near the interference location, specifically within approximately $\pm 5^\circ$ of the interference direction. In this region, the interference enters the mainlobe and directly distorts it. By placing nulls inside the mainlobe, the beamwidth narrows and the corresponding 3-dB beamwidth changes significantly.

QCC1 (purple), which uses diagonal loading, is more robust than the Capon method and produces less mainlobe distortion when steering near the interference. The behavior of

QCC2 (green) appears less predictable; when steering away from the interference, it tends to place nulls in the sidelobes, which in turn affects the shape of the mainlobe.

In the right plot of Figure 2.8, the Fourier beamformer (blue) maintains a value of zero across all pointing angles, meaning that the peak of the beam always aligns exactly with the steering angle. This is the ideal case, with no bias, and clearly indicates that the beam is pointing precisely where it is commanded. QCC1, which uses diagonal loading (purple), also has a value of zero across all pointing angles. This indicates that even though there is interference and the mainlobe becomes distorted (as shown in the left plot of Figure 2.8), the peak still remains at the steering direction. This demonstrates robustness and an improvement compared to the original Capon method. The LCC beam pattern (yellow) shows a small but constant bias of about 0.1° to 0.2° . This is due to sampling issues arising from the discrete constraint points; because the constraints are not continuous, the mainlobe becomes slightly shifted. The Capon method (red) and QCC3 (dotted light blue), which again overlap each other, show the largest bias near the interference location. This is expected, as the beam is most distorted when the interference lies inside the mainlobe. Lastly, QCC2 (green) shows the most fluctuating results, the peak of the mainlobe changes dramatically as the beam scans, indicating instability in the solution.

2.5.5 Minimum Mean-Square Error (MMSE) Beamforming

The MMSE digital beamforming method provides an adaptive approach for forming robust beams while suppressing interference and sidelobe contamination. Unlike the Capon method, which imposes an explicit unity-gain constraint in the steering direction, the MMSE method incorporates the desired steering direction and noise statistics directly into the optimization objective. As a result, the MMSE solution balances mainlobe preservation and sidelobe minimization through a statistically optimal weighting of the received signals.

Following [62], the MMSE weight vector for a given azimuth–elevation pair $(\theta_{m_x}, \phi_{m_y})$ is defined as

$$\mathbf{w}_{\text{MMSE}}(m_x, m_y) = \frac{\mathbf{R}^{-1} \mathbf{s}(\theta_{m_x}, \phi_{m_y})}{\mathbf{s}^H(\theta_{m_x}, \phi_{m_y}) \mathbf{R}^{-1} \mathbf{s}(\theta_{m_x}, \phi_{m_y})}, \quad (2.29)$$

where $\mathbf{s}(\theta, \phi)$ is the steering vector of the 2-D planar array and $\mathbf{R}$ is the covariance matrix of the received data. The normalization in the denominator ensures an appropriate gain level in the desired direction while maintaining robustness against noise and distributed clutter.

The covariance matrix $\mathbf{R}$ is modeled as the sum of the signal covariance and noise covariance,

$$\mathbf{R} = \mathbf{S} \mathbf{R}_s \mathbf{S}^H + \mathbf{R}_n, \quad (2.30)$$

where $\mathbf{R}_s$ represents the spatial covariance of distributed precipitation targets and $\mathbf{R}_n = \sigma_n^2 \mathbf{I}$ is the noise covariance, assumed to be spatially white. Because $\mathbf{R}_s$ depends on the unknown precipitation structure, an iterative approach is adopted: an initial estimate is obtained using Fourier beamforming, and subsequent iterations refine the MMSE weight vector using updated covariance estimates.

Compared to traditional methods such as Fourier beamforming, the MMSE method demonstrates superior sidelobe suppression and stability. Simulation results presented in [62] show that MMSE produces a stable mainlobe and adaptively places nulls toward interfering directions. For point-target scenarios, MMSE achieves accurate power estimation even when closely spaced strong and weak targets are present in Figure 2.9.

The MMSE beamformer provides superior sidelobe suppression and robust performance. However, these benefits come with several trade-offs. The MMSE method requires matrix inversion and iterative covariance estimation, making it significantly more computationally expensive, approximately 100 times slower than Fourier beamforming in the simulation from [62]. Its performance strongly depends on the accuracy of the co-




Figure 2.9: Simulation results from [62]. This figure shows examples of mean power estimates and formed beam patterns for two point targets. (a-1) True value of the point targets placed at azimuths of 0° and 0° , and elevations of 5° and 5° , respectively. (a-2) and (a-3) Mean power estimates for two targets in the results of Fourier and MMSE, respectively. (b-1) Formed beam pattern grazing at an azimuth of 0° and an elevation of 0° , respectively, along a line through two points. (b-2) and (b-3) Formed beam pattern grazing at an azimuth of 0° and an elevation of 0° in whole range.

variance matrix, and is therefore sensitive to sample availability, modeling assumptions, and array calibration errors. Unlike Capon, MMSE does not impose strict mainlobe constraints, which can lead to slight mainlobe distortion or underestimation of polarimetric variables, particularly at large off-boresight scan angles where projection effects widen the main beam. These limitations make MMSE computationally demanding and requiring careful design for real-time polarimetric PAR applications.

2.5.6 AI-Based Fast Capon Beamforming

Recent advancements in machine learning have enabled the development of fast, data-driven approximations to computationally expensive adaptive beamforming methods. The Capon beamformer provides superior angular resolution and sidelobe suppression but is difficult to implement in real time due to the need to invert the covariance matrix for every beam and at every range gate. To address this limitation, a Deep Learning Adaptive digital beamForming (DLAF) model has been proposed as a computationally efficient surrogate for the Capon method [63].

DLAF is a complex-valued neural network that learns the nonlinear mapping from received time-series IQ data to the corresponding Capon beamforming weights. Instead of explicitly computing (2.11), the DLAF network is trained using ground-truth Capon weights generated offline. Once trained, the network outputs the adaptive beamforming weights directly from the input signals, skipping covariance-matrix inversion entirely. The architecture consists of complex-valued dense layers, complex batch normalization, $\tanh$ activations, and dropout regularization, enabling the network to preserve and utilize phase information essential for accurate beamforming.

The model is trained with simulated point-target data and weather signals derived from NEXRAD observations, using the Horus radar parameters as the array configuration. As




Figure 2.10: Simulated results from [64], for the point target case are presented above. Left plot - DBF Power of Capon, Fourier and DLAF DBF and right plot - Antenna Pattern of Capon, Fourier and DLAF DBF. The Capon method is plotted in green, the Fourier method is plotted in blue, and DLAF is plotted in yellow.

shown in the point-target experiments (Figure 2.10), DLAF and the Capon produce beam patterns that place a null at the location of the point target while the Fourier beamformer has high sidelobe level at that location. In weather-signal experiments (Figure 2.11), DLAF captures the spatial structure of the weather more similar compared to the truth (NEXRAD data), while Fourier DBF power was contaminated due to the high sidelobes and the Capon DBF lost some structure of the storm due to the changing beam patterns and lack of calibration.

A key advantage of DLAF is computational efficiency. The computation time for DLAF grows approximately linearly with array size, while Capon scales exponentially due to matrix inversion. The measurements in [64] demonstrate that for a 1024 element array, DLAF requires approximately 1.53 s, compared to 9.24 s for Capon. This reduction enables practical real-time implementation of adaptive digital beamforming on large aperture polarimetric PARs.

Overall, AI-based fast Capon beamforming via DLAF has shown potential as an alternative solution to adaptive beamforming, offering Capon-like performance with a reduced




Figure 2.11: Simulated results from [64], for the weather case are presented above. Left plot - Reflectivity field of the simulated Fourier DBF, middle plot - Reflectivity field of the simulated Capon DBF, and right plot - Reflectivity field of the simulated DLAF DBF.

computational cost.

2.5.7 Bayesian Adaptive Beamforming under DoA Uncertainty

Steering-vector uncertainty arises whenever the presumed angle of arrival differs from the actual one and can significantly degrade performance in the presence of noise and interference. Such mismatch can be caused by DoA estimation errors, local scattering and source spreading, near-far spatial-signature differences, waveform distortion, imperfect array calibration, and deformations or modeling errors in the antenna array geometry.

The Bayesian adaptive beamforming approach incorporates uncertainty in the DoA by treating the DoA as a random parameter with a prescribed prior distribution. Using received data, the method updates this prior through Bayes' rule to obtain a posterior distribution, from which a composite steering vector is formed. This posterior-averaged steering vector enables the beamformer to adapt naturally to the data while maintaining robustness under conditions of directional uncertainty. A Bayesian framework behaves like a “responsive” processor under high-SNR conditions, while reverting to a more “robust” processor under low-SNR conditions [65].

In this approach, the DoA is modeled as a random parameter with prior probability den-

sity function (PDF) $q(\theta)$, representing the a priori knowledge or uncertainty of the signal direction. Given L number of received data, the beamformer updates this prior using Bayes' rule to obtain the posterior PDF $p(\theta | \mathbf{x}_L)$. The posterior distribution reflects both the prior uncertainty and the information contained in the received data. To avoid costly numerical integration, the DoA space is discretized into a finite set of P candidate directions,

$$\boldsymbol{\theta} = \{\theta_1, \dots, \theta_P\}. \quad (2.31)$$

Unlike the Capon method, which enforces a distortionless constraint at a single steering vector, the Bayesian formulation imposes a distortionless constraint on the *posterior-averaged* steering vector,

$$\bar{\mathbf{a}} = \sum_{i=1}^P a(\theta_i) p(\theta_i | \mathbf{x}_L) = A\mathbf{p}, \quad (2.32)$$

where $A = [a(\theta_1), \dots, a(\theta_P)]$ is the steering matrix and $\mathbf{p}$ is the posterior probability vector. The weight vector is then obtained by solving the Lagrange multiplier with this composite steering vector,

$$\mathbf{w}_{\text{Bayes}} = \frac{\hat{\mathbf{R}}^{-1} A \mathbf{p}}{\mathbf{p}^T A^H \hat{\mathbf{R}}^{-1} A \mathbf{p}}. \quad (2.33)$$

This expression resembles the Capon solution, but with the composite steering vector replacing the single-direction steering vector. Increased number of DoA points does not reduce the adaptive degrees of freedom, because the DoAs are averaged before being applied in the constraint.

The posterior PDF is computed by modeling the received data as a zero-mean complex Gaussian process. For each candidate DoA θ_i , the likelihood function takes the form

$$p(\mathbf{x}_L | \theta_i) \propto \frac{1}{|\mathbf{R}_x(\theta_i)|^L} \exp\left(-L a(\theta_i)^H \hat{\mathbf{R}}^{-1} a(\theta_i)\right), \quad (2.34)$$

where $\mathbf{R}_x(\theta_i) = \sigma_s^2 a(\theta_i) a(\theta_i)^H + \sigma_n^2 I$ is the model correlation matrix [65]. Applying Bayes' rule yields

$$p(\theta_i | \mathbf{x}_L) = \frac{q(\theta_i) \exp \left[-L\beta a(\theta_i)^H \hat{\mathbf{R}}^{-1} a(\theta_i) \right]}{\sum_{k=1}^P q(\theta_k) \exp \left[-L\beta a(\theta_k)^H \hat{\mathbf{R}}^{-1} a(\theta_k) \right]}, \quad (2.35)$$

where β is a monotonic function of the SNR. The parameter β can be viewed as a tuning parameter that adjusts the beamformer's responsiveness to changes between slow time samples of the received data.

When the SNR is high, the posterior PDF becomes sharply peaked near the true DoA, making $\bar{\mathbf{a}} \approx a(\theta_{\text{true}})$. The Bayesian beamformer then behaves similarly to a responsive processor such as the maximum-likelihood (ML) beamformer, achieving near-optimal gain. Under low SNR conditions, the posterior PDF spreads out and approximates the prior PDF, causing the composite steering vector to represent an average over the uncertainty region. This effectively widens the mainlobe and produces a more robust beam pattern.

Although MMSE, AI-based fast Capon, and Bayesian adaptive beamforming provide important alternative frameworks, they are not developed experimentally in this dissertation. Instead, the focus is placed on Capon beamforming and STAP, as these methods are mature, physically interpretable, and can be directly implemented and validated using Horus radar measurements.

2.6 Deterministic Null-Constrained Beamforming

Adaptive digital beamforming, regardless of the specific algorithm used, almost always involves trade-offs. Some of the key limitations include the computational burden associated with mathematically expensive operations (such as matrix inversion), which becomes critical for real-time implementation. Another limitation is the unpredictability of the resulting

beam shapes, since they change depending on the targets, interference, and surrounding environment. This introduces uncertainty in the actual steering direction and leads to variations in beamwidth. Such beamwidth variations are especially critical for weather observations, which rely on accurate measurements of distributed targets.

One alternative to adaptive beamforming is to design a deterministic beam pattern with nulls placed in the directions of a known interference sources. This approach only works when the interference directions are clearly known, and it may not yield an “optimized” pattern in the adaptive sense. However, it can still serve as a practical solution to improve the spatial quality of the data by mitigating known interference.

Nulling constraints can be defined as follows [50]:

$$\mathbf{w}^H \mathbf{a}(\mathbf{k}_i) = 0, \quad i = 1, 2, \dots, N - 1. \quad (2.36)$$

This condition specifies that the beam pattern has nulls at selected angular locations. A PAR with N elements has $N - 1$ degrees of freedom available for independent constraints. These are referred to as zero-order constraints (or zero-order nulls).

Taking the first derivative of the beam pattern with respect to the steering parameter (k) yields another type of constraint, called the first-order constraint:

$$\left. \frac{d}{dk} B(\mathbf{k}) \right|_{k=k_i} = \mathbf{w}^H \left[\left. \frac{d}{dk} \mathbf{a}(k) \right|_{k=k_i} \right] \triangleq \mathbf{w}^H \mathbf{d}_1(k_i), \quad i \in \Omega_1, \quad (2.37)$$

where Ω_1 is a subset of the I angular locations at which the derivative constraint is imposed (i.e., where the derivative of the beam pattern is forced to be zero).

This derivative process can be repeated. In general, we can impose an n^{th} -order deriva-

tive constraint:

$$\left. \frac{d^n}{dk^n} B(\mathbf{k}) \right|_{k=k_i} = \mathbf{w}^H \left[\frac{d^n}{dk^n} \mathbf{a}(k) \right] \Big|_{k=k_i} \triangleq \mathbf{w}^H \mathbf{d}_n(k_i), \quad i \in \Omega_n. \quad (2.38)$$

where n denotes the order of the derivative constraint, $\mathbf{d}_n(k_i)$ represents the n^{th} derivative of the steering vector evaluated at k_i , and Ω_n is the set of angular locations at which the n^{th} -order derivative constraint is imposed.

However, in practical implementations, constraints beyond the second derivative are rarely used.

The constraint matrices for zero-order ($\mathbf{C}_0$), first-order ($\mathbf{C}_1$), and second-order ($\mathbf{C}_2$) constraints are given by

$$\mathbf{C}_0 = \begin{bmatrix} \mathbf{a}(\mathbf{k}_1) & \mathbf{a}(\mathbf{k}_2) & \cdots & \mathbf{a}(\mathbf{k}_{I_0}) \end{bmatrix}, \quad (2.39)$$

$$\mathbf{C}_1 = \begin{bmatrix} \mathbf{d}_1(\mathbf{k}_1) & \mathbf{d}_1(\mathbf{k}_2) & \cdots & \mathbf{d}_1(\mathbf{k}_{I_1}) \end{bmatrix}, \quad (2.40)$$

$$\mathbf{C}_2 = \begin{bmatrix} \mathbf{d}_2(\mathbf{k}_1) & \mathbf{d}_2(\mathbf{k}_2) & \cdots & \mathbf{d}_2(\mathbf{k}_{I_2}) \end{bmatrix}, \quad (2.41)$$

where I_0 , I_1 , and I_2 denote the number of constraint locations for the zero-, first-, and second-order constraints, respectively.

The total constraint matrix is then

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}_0 & \mathbf{C}_1 & \mathbf{C}_2 \end{bmatrix}, \quad (2.42)$$

which has size $N \times I_c$, where I_c is the total number of constraints and $I_c < N$. Here, we assume that the columns of $\mathbf{C}$ are linearly independent. This assumption ensures that

$\mathbf{C}^H \mathbf{C}$ is non-singular, which is required for the subsequent optimization process.

We now optimize the squared weighting error subject to the null constraint (2.36). Using a Lagrange multiplier and taking the gradient with respect to $\mathbf{w}$, the optimized weights for the null constraint ($\mathbf{w}_o$) are [50]:

$$\mathbf{w}_o^H = \mathbf{w}_d^H \left(\mathbf{I}_N - \mathbf{C} (\mathbf{C}^H \mathbf{C})^{-1} \mathbf{C}^H \right), \quad (2.43)$$

where $\mathbf{w}_d$ is the steering weight vector (same as $\mathbf{a}(\theta)$). The optimized weight vector is the component of the steering weight vector in the subspace orthogonal to the constraint subspace. Thus, we can rewrite (2.43) using a projection matrix:

$$\mathbf{w}_o^H = \mathbf{w}_d^H \mathbf{P}_C^\perp, \quad (2.44)$$

where $\mathbf{P}_C^\perp = \mathbf{I}_N - \mathbf{P}_C$.

The projection matrices for zero-order, first-order, and second-order constraints are defined as:

$$\mathbf{P}_{C_0} = \mathbf{C} (\mathbf{C}^H \mathbf{C})^{-1} \mathbf{C}^H, \quad (2.45)$$

$$\mathbf{P}_{C_1} = \mathbf{C}_1 (\mathbf{C}_1^H \mathbf{C}_1)^{-1} \mathbf{C}_1^H, \quad (2.46)$$

$$\mathbf{P}_{C_2} = \mathbf{C}_2 (\mathbf{C}_2^H \mathbf{C}_2)^{-1} \mathbf{C}_2^H. \quad (2.47)$$

Therefore, the optimum weights with different constraint sets are:

$$\mathbf{w}_{o,0} = (\mathbf{I}_N - \mathbf{P}_{C_0}) \mathbf{w}_d, \quad (2.48)$$

$$\mathbf{w}_{o,1} = (\mathbf{I}_N - \mathbf{P}_{C_1}) \mathbf{w}_d, \quad (2.49)$$

$$\mathbf{w}_{o,2} = (\mathbf{I}_N - \mathbf{P}_{C_2}) \mathbf{w}_d. \quad (2.50)$$




Figure 2.12: Beam patterns from the Fourier beamformer with no constraint (black line), the beamformer with a zero-order constraint (yellow line), the beamformer with a first-order constraint (blue line), the beamformer with a second-order constraint (red) are plotted above. Their nulling constraints are at 10° , noted as a dotted line. Beams are pointed at 0° .

Figure 2.12 shows simulations from deterministic null-constrained beamforming. The Fourier beamformer (with no constraints) is plotted in black, the beamformer with a zero-order constraint in yellow, the beamformer with a first-order constraint in blue, and the beamformer with a second-order constraint in red. All constraints are applied at 10° angle (denoted by the dotted line), and the beams are formed at broadside (0°). The results show that the width of the null gradually increases from zero-order to first-order to second-order. The second-order null constraint lowers the sidelobe below -40 dB from about 7° to 13° . The trade-off for forming a wider null is a broadened mainlobe, although the sidelobe levels remain similar across all beamformers. This approach is included as a conceptual baseline and as a practical alternative when clutter locations are known a priori, providing contrast to adaptive methods discussed later. Depending on the scenario (e.g., type of interference), each constraint order has its own benefits, and choosing the appropriate constraint order can help optimize performance. For example, if the interference is coming from a broader angular region, the second-order null constraint may work better to mitigate interference contamination. Another example is when the radar operates for weather surveillance, where zero-order or first-order nulls may be preferable (especially to suppress returns from ground clutter targets) as the mainlobe distortions are minimal.

Chapter 3

Calibration for Adaptive Beamforming

Temporal resolution for reflector-based radars is limited by mechanical rotation. In contrast, PARs can significantly improve temporal resolution through imaging and other advanced beamforming methods. Imaging is a unique capability of PAR in which the radar transmits a wide beam and simultaneously forms multiple narrow receive beams. By doing so, the radar can scan a large sector without mechanical rotation, thereby achieving substantially improved temporal resolution. However, imaging increases sidelobe levels in the two-way pattern because the transmit beam is intentionally broadened. These elevated sidelobe levels increase the risk of contamination from interference and clutter. Furthermore, the fixed receive beam shapes from the deterministic DBF can lead to additional contamination when the power distribution changes rapidly, such as in the presence of strong clutter or near the edges of storms.

Adaptive beamforming can mitigate the fixed sidelobe challenge associated with imaging. As discussed in Chapter 2, adaptive beamforming adjusts its beam patterns according to the scanning environment. When strong power is present in undesired directions, the adaptive beamformer attempts to suppress antenna gain in those directions. As a result, adaptive beamforming can reduce sidelobe or mainlobe contamination and improve spatial data quality.

However, adaptive beamforming introduces several important trade-offs. It is compu-

tationally expensive and is often impractical to implement in real time for weather surveillance applications. Adaptive methods also require a large number of independent samples to estimate the covariance matrix reliably. In addition, because the receive beam patterns change adaptively, meticulous calibration is required to ensure reliable retrieval of meteorological and polarimetric variables.

Calibration of beam patterns is critical for weather radars because hydrometeors are distributed throughout the radar resolution volume, and the received signals represent an integrated response from many scatterers within the beam. If the beam shapes change, these variations must be properly calibrated to avoid biases caused by beam-to-beam differences. Furthermore, as the beam shape changes, the noise characteristics after beamforming also change, requiring noise to be re-estimated after beam formation. In addition, most weather radars operate in dual-polarization mode, where beams are formed independently for horizontal and vertical polarizations. Any mismatch between the H and V beam patterns directly affects the estimation of polarimetric variables, particularly differential reflectivity, cross-correlation coefficient, and differential phase.

This chapter addresses the calibration challenges introduced by adaptive beamforming. First, beam pattern calibration is discussed to account for changing receive beams. Next, noise re-estimation after beamforming is described. Differences between the H and V channels caused by adaptive beamforming are then examined, along with their impact on polarimetric variable estimation. Finally, both conventional and adaptive beamforming are evaluated using simulated data based on NEXRAD observations to validate the proposed calibration methods and to assess the resulting improvements, remaining challenges, and associated trade-offs.

3.1 Beam Pattern Calibration

Beam pattern calibration is primarily needed to recover reflectivity from the received power (e.g., Z and Z_{DR}). The received power P_r depends on several key radar parameters, including the transmitted power P_t , antenna gain g , and wavelength λ , which determine the amount of energy transmitted and received by the radar. The reflectivity η characterizes the ability of distributed targets to scatter electromagnetic energy back toward the radar (i.e., radar cross section per unit volume).

The received power also decreases with the square of the range r , and additional losses l account for attenuation due to atmospheric absorption, scattering, and system inefficiencies. The range-weighting function $W_s(r)$ describes how the transmitted pulse interacts with targets along the range dimension and how the receiver characteristics influence the sampled I/Q data. Furthermore, the two-way antenna beam pattern $f^4(\theta, \phi)$ describes the directional distribution of transmitted and received energy and must be integrated over elevation and azimuth angles. Together, these factors determine the received radar power.

To understand how beam pattern calibration is affected by adaptive beamforming, it is necessary to examine the relationship between the transmitted power, the antenna pattern, and the received signal. The radar transmits power P_t , and the resulting power density at range r is given by

$$S = \frac{P_t}{4\pi r^2} G(\theta, \phi), \quad (3.1)$$

where $G(\theta, \phi) = g_t f^2(\theta, \phi)$ is the antenna gain pattern, g_t is the transmit gain, and $f^2(\theta, \phi)$ is the normalized one-way antenna pattern.

When the transmitted wave interacts with a target having backscattering cross section σ_b , the scattered power density at the radar location becomes

$$S_b = \frac{P_t G(\theta, \phi)}{(4\pi)^2 r^4} \sigma_b. \quad (3.2)$$

The received power is obtained by multiplying this scattered power density by the effective antenna aperture

$$A_e(\theta, \phi) = \frac{g_r \lambda^2}{4\pi} f^2(\theta, \phi), \quad (3.3)$$

where g_r is the receive antenna gain. For a monostatic radar, $g_t = g_r = g$, and the received power becomes

$$P_r = \frac{P_t g^2 \lambda^2 \sigma_b f^4(\theta, \phi)}{(4\pi)^3 r^4 l^2}. \quad (3.4)$$

This expression represents the classical radar equation for a point target.

Weather radars observe hydrometeors such as raindrops, snowflakes, and hailstones, which are distributed randomly within a sampling volume. Unlike a single point target, the radar return from precipitation is the superposition of echoes from a large number of individual scatterers whose positions, sizes, and orientations vary randomly. Because of this randomness, weather radar measurements are typically described in terms of statistical averages rather than instantaneous values. Radar signal processing therefore relies on the expected (ensemble-averaged) received power, which represents the mean contribution of many scatterers within the radar resolution volume.

The expected power contribution from each scatterer to the mean echo power can be expressed as

$$\mathbb{E}[P_i] = \alpha^{1/2} |W_s|^2 \mathbb{E}[|A_i|^2], \quad (3.5)$$

where α is a proportionality constant and $\alpha^{1/2} \mathbb{E}[|A_i|^2]$ represents the expected echo

power from the i^{th} scatterer at the receiver. For a small volume element dV containing many hydrometeors, the expected differential received power becomes

$$\mathbb{E}[dP] = \alpha^{1/2} \sum_i |W_s|^2 \mathbb{E}[|A_i|^2] \quad (3.6)$$

$$\mathbb{E}[dP] = \alpha^{1/2} \sum_i |W_s|^2 \frac{P_t g^2 \lambda^2 \sigma_b}{(4\pi)^3 r^4 l^2} f^4(\theta, \phi) \quad (3.7)$$

$$\mathbb{E}[dP] = \alpha^{1/2} \int |W_s(r_0, r)|^2 C \frac{\sigma_b}{r^4 l^2(r)} f^4(\theta - \theta_0, \phi - \phi_0) dV \quad (3.8)$$

$$\mathbb{E}[dP] = I(r_0, r) dV \int_0^\infty \sigma_b(D) N(D, r) dD, \quad (3.9)$$

where

$$I(r_0, r) = |W_s(r_0, r)|^2 C \frac{1}{r^4 l^2(r)} f^4(\theta - \theta_0, \phi - \phi_0)$$

is defined as the illumination factor, and (θ_0, ϕ_0) denote the beam pointing direction.

Here, $\sigma_b(D)$ is the backscattering cross section of a hydrometeor with diameter D , and $N(D, r)$ is the particle size distribution (PSD). The term $\sigma_b(D)N(D, r)dD$ therefore represents the expected backscattering cross section per unit volume contributed by particles with diameters between D and $D + dD$. The integral

$$\int_0^\infty \sigma_b(D) N(D, r) dD$$

defines the volume reflectivity $\eta(r)$.

Rewriting the expected received power using the differential volume element in spherical coordinates, $dV = r^2 \sin \theta dr d\theta d\phi$, gives

$$\bar{P}(r_0) = \int_0^{r_2} \int_0^\pi \int_0^{2\pi} \eta(r) I(r_0, r) r^2 \sin \theta \, dr \, d\theta \, d\phi, \quad (3.10)$$

where r_0 denotes the center of the sampled range gate. The factor $r^2 \sin \theta$ arises from the spherical coordinate representation of the radar resolution volume.

Both η and l are functions of range r , but they can be assumed approximately constant over the region where $|W_s(r_0, r)|^2 f^4(\theta, \phi)$ has significant value. Under this assumption, they can be evaluated at $r = r_0$, yielding

$$\bar{P}(r_0) = \frac{P_t g^2 \lambda^2 \eta(r_0)}{(4\pi)^3 r_0^2 l^2(r_0)} \int_0^{r_2} |W_s(r_0, r)|^2 dr \int_0^\pi \int_0^{2\pi} f^4(\theta, \phi) \sin \theta \, d\theta \, d\phi. \quad (3.11)$$

In the hypothetical case where a perfect square pulse is transmitted and there are no receiver losses or gain variations, the system-weighting-function term becomes

$$\int_0^{r_2} |W_s(r)|^2 dr = \frac{c\tau}{2}, \quad (3.12)$$

where c is the speed of light and τ is the transmitted pulse duration. In real radar systems, receiver losses (l_r) and filter characteristics modify the effective weighting function.

For simplicity, the antenna beam pattern is often approximated by a Gaussian function, assuming that the beam shape and width remain fixed. Under this assumption, the two-way antenna pattern integral becomes

$$\int_0^\pi \int_0^{2\pi} f^4(\theta, \phi) \sin \theta \, d\theta \, d\phi = \frac{\pi \theta_b^2}{8 \ln 2}, \quad (3.13)$$

where θ_b denotes the 3-dB beamwidth.

However, for PARs, the 3-dB beamwidth changes as the beam steers away from broadside. Therefore, beam pattern calibration must account for this beam broadening effect.

When the conventional Fourier beamforming method is used, the beam pattern is deterministic and calibration can be precomputed prior to radar operation. In contrast, adaptive beamforming methods such as the Capon beamformer produce array patterns that vary depending on the received data. Consequently, the antenna pattern and beamwidth change dynamically, requiring beam pattern calibration to be performed in real time. In such cases, the conventional Gaussian beam assumption is no longer generally valid.

To account for the effect of the antenna beam pattern on the received signal, we incorporate the beam pattern into the received power formulation. Assuming uniform reflectivity within the mainlobe, the reflectivity can be expressed in terms of the equivalent reflectivity factor, Z , as

$$\eta = \frac{\pi^5}{\lambda^4} |K_w|^2 Z, \quad (3.14)$$

where K_w is the complex dielectric factor of water. Substituting (3.12) and (3.14) into (3.11), a calibrated relationship between reflectivity and other radar parameters can be obtained:

$$Z = P_r r_0^2 l^2 \frac{2(4\pi)^3 \lambda^2}{P_t g^2 c \tau \pi^5 |K_w|^2} \frac{1}{\int_0^\pi \int_0^{2\pi} f^4(\theta, \phi) \sin \theta \, d\theta d\phi}. \quad (3.15)$$

The beam pattern calibrated reflectivity Z can be written in dBZ as

$$\begin{aligned} Z(\text{dBZ}) = & 10 \log_{10} P_r + 20 \log_{10} r_0 + 20 \log_{10} l + 10 \log_{10} C \\ & - 10 \log_{10} \left(\int_0^\pi \int_0^{2\pi} f^4(\theta, \phi) \sin \theta \, d\theta d\phi \right), \end{aligned} \quad (3.16)$$

where $20 \log_{10} l$ accounts for all loss terms, and $10 \log_{10} C$ represents the remaining system constants, with

$$C = \frac{2(4\pi)^3 \lambda^2}{P_t g^2 c \tau \pi^5 |K_w|^2}. \quad (3.17)$$

The last term, $10 \log_{10} \left(\int_0^\pi \int_0^{2\pi} f^4(\theta, \phi) \sin \theta \, d\theta d\phi \right)$, accounts for beam-pattern calibration.

The two-way pattern, $f^4(\theta, \phi)$, includes both the transmit and receive patterns. The transmit pattern $f_t^2(\theta, \phi)$ is the product of the transmitting array factor (g_{at}) and the element pattern (g_e). The receive pattern is the product of the element pattern, the subarray pattern (g_s), if subarrays are used, and the receive array factor (g_{ra}):

$$f^4(\theta, \phi) = [g_{at}(\theta, \phi) \times g_e(\theta, \phi)] \times [g_e(\theta, \phi) \times g_s(\theta, \phi) \times g_{ra}(\theta, \phi)]. \quad (3.18)$$

In practice, the integration range for the beam pattern can be limited to within the 9 dB mainlobe of the two-way pattern, rather than over the full angular domain $[0, \pi]$ in elevation and $[0, 2\pi]$ in azimuth. This simplification is valid under the assumption of uniform reflectivity within the mainlobe and reasonably behaved receive patterns. It is worth noting that adaptive beam patterns may not behave reasonably for point targets, as they tend to suppress sidelobes aggressively. However, for weather targets, which are distributed targets, the adaptive beam pattern cannot suppress sidelobes to the same extent as it can for point targets. Furthermore, when no strong localized interference is present, the sidelobe structure of the Capon beam pattern tends to resemble that of the conventional Fourier beamformer. For weather observations, hydrometeors are distributed targets whose power is relatively uniform in the spatial domain. In such cases, the covariance matrix becomes approximately proportional to the identity matrix, $\mathbf{R} \approx \sigma^2 \mathbf{I}$, and the Capon weight vector $\mathbf{w} \propto \mathbf{R}^{-1} \mathbf{a}$ reduces to a form proportional to the steering vector. As a result, the beam pattern approaches that of the Fourier beamformer defined by the array geometry. Therefore, the assumption remains acceptable in this context.

3.2 Noise Calibration

In addition to beam pattern calibration, the noise power after adaptive beamforming is not constant, unlike in the conventional Fourier beamforming method. This occurs because the Capon beamformer continuously adjusts the beamforming weights based on the received data. As a result, the effective noise gain of the beamformer changes, causing the output noise level to fluctuate. These fluctuations must be properly calibrated to ensure reliable estimates of polarimetric variables and to maintain overall data integrity.

The estimated power after Capon beamforming is given by

$$\hat{P}_c = \mathbf{w}_c^H \hat{\mathbf{R}} \mathbf{w}_c = \mathbf{w}_c^H \hat{\mathbf{R}}_{\text{sig}} \mathbf{w}_c + \mathbf{w}_c^H \hat{\mathbf{R}}_{\text{noise}} \mathbf{w}_c, \quad (3.19)$$

where $\hat{\mathbf{R}}_{\text{sig}}$ and $\hat{\mathbf{R}}_{\text{noise}}$ represent the estimated signal and noise covariance matrices.

The noise covariance matrix is typically modeled as

$$\hat{\mathbf{R}}_{\text{noise}} = \sigma_n^2 \mathbf{I}, \quad (3.20)$$

where σ_n^2 is the receiver noise power and $\mathbf{I}$ is the identity matrix.

In practice, the noise power is estimated independently for each receiver channel, for example through spectral analysis of noise-only samples. Because the adaptive weights $\mathbf{w}_c$ vary with the received data, the contribution of noise to the total power estimate also changes. Therefore, the effective noise power must be computed using the adaptive weights at each range gate where adaptive beamforming is applied. Without this adaptive noise calibration, the resulting variable estimates become scene-dependent and unreliable.

3.3 Evaluation of Calibration for Adaptive Beamforming

To assess the performance of the Fourier and Capon beamforming techniques, two types of simulations are performed. First, three simple distinct cases are simulated to examine the behavior of the beam patterns and the resulting polarimetric variables under controlled conditions. Next, simulations using NEXRAD data are conducted to evaluate the potential of Capon beamforming to improve polarimetric weather measurements in a more realistic scenario.

For the controlled simulations, three distinct cases are modeled: (Case 1) uniformly distributed reflectivity, representing a stratiform region in a real-world scenario; (Case 2) a point target embedded in Case 1, representing clutter (e.g., buildings and airplanes) within weather targets; and (Case 3) non-uniformly distributed reflectivity with a gradient variation, representing the storm edge. These cases are illustrated in Figure 3.1. Details of the simulation setup are provided in Appendix A.




Figure 3.1: The model polarimetric variables for Case 1 (blue), Case 2 (dotted red), and Case 3 (green) are presented. Case 1 is uniformly distributed targets, representing a center of the storm, Case 2 contains a stationary point target, and Case 3 has gradient changes, representing the edge of storms.

Figure 3.1 shows the modeled polarimetric variables (Z , v_r , σ_v , Z_{DR} , ρ_{hv} , ϕ_{DP}) for the three simulation cases. Different colors represent each case: blue for Case 1, red for Case 2, and green for Case 3. The targets are continuously distributed in a two-dimensional angular domain (azimuth and elevation). DBF is performed in one dimension along elevation, from 0° to 20° , resembling a range–height indicator (RHI) scan.

The time-series signals were simulated with 64 samples and 832 receivers to match the Horus radar architecture. Figure 3.2 shows the array configuration and the corresponding receive beam pattern.




Figure 3.2: The configuration of the array is shown in the left plot, and the receive pattern power (dB) is shown in the right plot.

Note that a diagonal loading level of -10 dB was applied to stabilize the covariance matrix during inversion in the Capon beamforming method. The impact of diagonal loading will be discussed in Section 3.4. After applying diagonal loading, the resulting SNR is approximately 30 dB.

Using this simulation scenario, the analysis will begin by examining the beam patterns formed by each method in the H and V channels. It will then investigate the calibrated polarimetric variables obtained from each DBF method, with particular focus on how the Capon method impacts the estimation of polarimetric variables through its adaptively varying beam patterns and the associated beam pattern calibration. The estimation of polari-

metric variables is described in Appendix B.

Next, the Fourier and Capon DBF methods are applied to simulations using NEXRAD data to evaluate the potential of Capon DBF to improve polarimetric weather measurements. The PAR time-series signals are simulated in the same manner as in the previous simulation cases using polarimetric observations from NEXRAD. For these simulations, NEXRAD data from 05/20/2013 20:03:56.96 UTC from KTLX are used, which contain tornadic supercell storms. In this case, the focus is on examining the overall quality of the polarimetric measurements.

Finally, the analysis examines the impact of H–V beam matching on the performance of adaptive beamforming.

3.3.1 Adaptive Beam Pattern Characteristics

The beam patterns for the Fourier (black) and Capon methods (red for the H-channel and blue for the V-channel) across three different cases are presented in Figure 3.3. The left figure corresponds to Case 1 (uniform distribution), the middle figure to Case 2 (point target), and the right figure to Case 3 (gradient change). In each figure, the target power is plotted in yellow as a reference. The Fourier method produces identical beam patterns in both H and V channels (although in real-world applications, they may differ due to element patterns, which is beyond the scope of this study). Therefore, only the H-channel beam pattern is plotted in black. In contrast, the Capon method produces distinct beam patterns in the H and V channels when $Z_{DR} \neq 0$. The solid gray line represents the beam steering angle, and DBF is performed from 0° to 20° .

In Case 1, the beams are steered at 9° . No noticeable nulls are present, since the power of the targets is uniformly distributed across the entire space. The mainlobe of the Capon beam pattern does not align with that of the Fourier method. This is because the Capon




Figure 3.3: Beam patterns for the Fourier (black for H channel) and Capon (red for H channel and blue for V channel) methods across three different scenarios: the left figure corresponds to Case 1 (uniform distribution), the middle figure to Case 2 (point target), and the right figure to Case 3 (gradient change). The solid black line represents the beam steering angle, while the yellow line denotes the target power, serving as a reference.

method minimizes power from all directions except the desired steering direction. The Capon method produces a beam pattern similar to that of the Fourier method when the signal is either spectrally white (i.e., containing noise across all frequencies) or spatially white (i.e., fully and uniformly distributed across the entire space). Although the targets in Case 1 are uniformly distributed, this condition of spatial whiteness is not satisfied because the targets are simulated at discrete 1° intervals. As a result, the Capon beam patterns deviate from those of the Fourier method.

In Case 2, the beams are steered to 6° , while the point target is located at 8° . Since the point target falls within the mainlobe, it introduces significant contamination in the Fourier method due to its beamwidth (approximately 2.8°), highlighting the limitation of spatial resolution.

In contrast, the Capon method effectively suppresses the power of the point target by placing a null. Additionally, the Capon method places a deeper null in the V channel, as the point target has a negative Z_{DR} value. Differences between the H- and V-channel beam patterns occur when the target characteristics differ between the two channels, i.e., when $Z_{DR} \neq 0$. A larger magnitude of Z_{DR} leads to greater differences between the H- and

V-channel beams produced by the Capon method, as illustrated in Figure 3.4.

The H–V beam differences are shown in Figure 3.4, where Case 1 is plotted in blue, Case 2 in red, and Case 3 in yellow. The difference is computed by subtracting the V-channel beam from the H-channel beam. At the location of the point target, there is approximately a -1.7 dB difference. Furthermore, differences are observed in the surrounding region, as the Capon method adaptively modifies the beam shape to minimize power from the undesired direction.

This demonstrates the Capon method’s ability to adaptively reshape the beam pattern, including the mainlobe, to optimize power in the desired direction.




Figure 3.4: The differences (in dB) between the H and V beams produced by the Capon method are analyzed for the three different cases and presented above. The x -axis represents the elevation angle, allowing for a direct comparison of beam mismatches as a function of target distribution. The y -axis represents the H and V mismatch in dB.

In Case 3, the power of the targets increases gradually with elevation angle, and the beams are steered to 1° . The Fourier method produces a fixed beam shape. In this scenario, it introduces significant contamination from the sidelobes on the right-hand side, where

higher target powers are located. As a result, accurate power estimation becomes more difficult, since the measured power represents an integral of both the beam pattern and the target power distribution in the scanning environment.

In contrast, the Capon method suppresses the sidelobes on the right-hand side, effectively mitigating this contamination. Figure 3.4 shows the H–V beam differences for Case 3 (yellow line). At steering angles between 0° and 6° , the differences are larger in magnitude, about -0.5 dB. This occurs because, at these steering angles, the Capon method suppresses power from higher elevation angles (right-hand side), where the targets have larger Z_{DR} values.

3.3.2 Calibrated Polarimetric Variables

Figures 3.3 and 3.4 show that the beam patterns generated by the Capon method vary with the target distribution. Therefore, beam pattern calibration is required for each steering direction, range bin, and both H and V channels to obtain meaningful polarimetric measurements. In addition, noise calibration is necessary. These beam pattern and noise calibration procedures are implemented as described in Sections 3.1 and 3.2.

Qualitative Results

Figure 3.5 presents the calibrated polarimetric variables from the Fourier (blue) and Capon (red) methods, along with the true values (black). The results shown correspond to a single realization of the simulation. Since the time-series signal simulation is stochastic, the outputs exhibit variability; these fluctuations would be reduced by averaging over multiple realizations. However, a single realization is presented here to illustrate realistic estimation behavior in the presence of noise. The Capon method may exhibit larger fluctuations due to its higher sensitivity to noise.




Figure 3.5: The calibrated polarimetric variables (rows) for the three cases (columns) using the Fourier (blue) and Capon (red) methods are presented. The true values are plotted with a black line in each plot.

In Case 1 (uniform field), both the Fourier and Capon methods estimate P , Z_{DR} , v_r , and ϕ_{DP} accurately, closely matching the true values. This confirms that the calibration procedures, including beam pattern calibration and noise estimation, have been properly applied. However, for ρ_{hv} , a small bias of approximately 0.008 is observed in the Capon estimate, while the Fourier estimate remains nearly unbiased. This difference arises from the beamforming weights: the Fourier method uses identical weights for H and V channels, preserving correlation, whereas the Capon method produces different complex value weights for each channel, introducing bias in ρ_{hv} even in well-behaved conditions.

In Case 2 (point target), the Fourier method exhibits noticeable bias near the target location for all variables (P , Z_{DR} , v_r , ϕ_{DP} , and ρ_{hv}). These biases result from sidelobe contamination and limited spatial resolution (approximately 2.8° in this simulation) and cannot be fully corrected through calibration. In contrast, the Capon method demonstrates improved spatial resolution, effectively reducing bias in P and Z_{DR} , and estimating v_r with minimal error. Although both methods show approximately 5° bias in ϕ_{DP} at the target location, the Capon method still performs better overall. For ρ_{hv} , the Fourier method is more accurate in uniform regions (0° – 2° and 12° – 20°), while the Capon method outperforms it near the point target (2° – 12°), demonstrating its adaptability to localized features.

Case 3 further highlights the limitations of the Fourier method in non-uniform scenes. From 0° to 10° , the Fourier method introduces consistent bias across all variables due to its fixed sidelobe structure (see Figure 3.3). In contrast, the Capon method adapts its beam pattern to the scanning environment, reducing bias across the full angular range. Although the Capon estimate of ρ_{hv} still exhibits some bias due to H–V beam weight decorrelation, the magnitude remains within acceptable limits for weather radar applications.

Overall, the results indicate that while the Fourier method benefits from consistent H–V beam matching—particularly for ρ_{hv} —its fixed beam pattern limits performance in non-uniform target environments. The Capon method, although more sensitive to noise and

calibration, provides improved spatial resolution and reduced bias in most variables when properly calibrated.

Quantitative Results

For the quantitative assessment, 500 realizations of the PAR time-series signals are generated. The bias for each polarimetric variable is computed by subtracting the estimated values from the true values (indicated by the black lines in Figure 3.5), and the results are presented as histograms in Figure 3.6. In all plots, the blue bars represent the bias distribution of the Fourier method, and the red bars represent that of the Capon method.

In Case 1, both methods show bias distributions centered near zero for all variables. The Capon method exhibits a slight mean bias of approximately 0.008 in ρ_{hv} , while the Fourier method remains nearly unbiased. The standard deviations of the bias distributions are similar for both methods, indicating comparable estimation stability in a uniformly distributed scene.

Case 2, which includes a point target, highlights the performance differences more clearly. The Capon method continues to produce narrow, near-zero-mean bias distributions for all variables except ρ_{hv} , which retains a small positive bias. In contrast, the Fourier method exhibits significantly larger and more variable biases, particularly in P and Z_{DR} . The bias peaks in P appear at 0, -3 , -6 , and -9 dB, corresponding to the sidelobe structure of the fixed Fourier beam pattern. This structural bias contaminates the estimation near the target. Across all variables, the Fourier method shows wider variance, whereas the Capon method maintains a tight, approximately Gaussian distribution centered near zero, demonstrating robustness to clutter and target interference.

Case 3 further confirms these observations in a more complex scene. The Fourier method introduces consistent negative bias in P , v_r , and ϕ_{DP} , particularly in the 0° – 10° region, reflecting the impact of fixed sidelobes (see Figure 3.3). In contrast, the Capon




Figure 3.6: Normalized probability density functions (PDFs) of bias in estimated polarimetric variables from 500 simulations of PAR time-series signals for three scenarios: Case 1 (uniform field), Case 2 (point target), and Case 3 (complex scene). Blue bars indicate bias from the Fourier method, and red bars represent the Capon method. Variables shown include P , v_r , Z_{DR} , ρ_{hv} , and ϕ_{DP} .

method adapts its beam pattern to the environment, reducing overall bias in most variables. Although ρ_{hv} from the Capon method still exhibits a slight bias (approximately 0.03), it remains within an acceptable range for practical applications. The tighter distributions produced by the Capon method across all variables reinforce the advantages of adaptive beamforming for high-resolution polarimetric weather radar applications.

3.4 SNR Effects and Diagonal Loading

The Capon method is sensitive to SNR [66], whereas the Fourier method is not. This is because Fourier beamforming uses fixed weights that are independent of the received signal, while Capon weights are data-dependent and therefore influenced by the SNR. As the SNR decreases, the Capon beam pattern approaches that of the Fourier method. However, even if the beam patterns appear similar, the beamforming weights are not identical to the Fourier weights unless the covariance matrix corresponds to spatially white noise. In that case, the covariance matrix reduces to a scaled identity matrix.

At lower SNRs, the H- and V-channel beamforming weights can become less correlated, which degrades the estimation of polarimetric variables such as ρ_{hv} . In addition, low SNR conditions can cause the covariance matrix to become poorly conditioned or nearly singular, making inversion unstable. This highlights the importance of techniques such as diagonal loading to ensure numerical stability in adaptive beamforming [67].

In this section, the impact of SNR on DBF performance is assessed. Since SNR directly affects the covariance estimation, it plays a critical role in the performance of the Capon method. Furthermore, the effect of diagonal loading—an essential component for stabilizing adaptive processing—is analyzed in terms of its impact on polarimetric variable estimation.

Figure 3.7 presents the bias in the estimated polarimetric variables as a function of

SNR. Results for Cases 1, 2, and 3 are shown in blue, orange, and green, respectively. The Fourier method is plotted with solid lines, while the Capon method is shown with dotted lines. The results indicate that ρ_{hv} is the most sensitive variable to SNR. To achieve acceptable estimation quality using Capon DBF, the SNR must exceed approximately 15 dB. In contrast, the Fourier method maintains reasonable estimation performance even at lower SNR levels (e.g., 5 dB), demonstrating its robustness to noise.




Figure 3.7: Bias in the estimated polarimetric variables as a function of SNR. Results for Cases 1, 2, and 3 are shown in blue, orange, and green, respectively. The Fourier method is plotted with solid lines, and the Capon method with dotted lines.

Figure 3.8 shows the mean bias of ρ_{hv} as a function of diagonal loading (x -axis) and SNR (y -axis). In practical scenarios, SNR is determined by system and environmental conditions. Therefore, diagonal loading must be applied to ensure numerical stability when SNR is low. However, selecting an appropriate loading level is not straightforward; it remains an active area of research and depends on system characteristics, particularly SNR. This figure provides practical guidance for selecting diagonal loading levels under different SNR conditions to minimize bias in ρ_{hv} , assuming that SNR can be estimated.

The reduced bias in ρ_{hv} observed with larger diagonal loading arises because the beamforming weights in the H and V channels become more correlated as the loading level increases. Figure 3.9 illustrates the correlation between the H- and V-channel beamforming weight amplitudes for different diagonal loading levels: -50 dB (blue), -40 dB (orange),

−30 dB (yellow), −20 dB (purple), and −10 dB (green). Larger diagonal loading levels (e.g., −10 dB) result in significantly higher correlation between the H- and V-channel weights compared to lower loading levels. This increased correlation improves polarimetric consistency and reduces bias in the estimated ρ_{hv} .




Figure 3.8: Mean bias of ρ_{hv} from the Capon method as a function of Diagonal Loading (the x -axis) and SNR (the y -axis).

3.5 Cross-Correlation Coefficient After Adaptive Beamforming

Despite proper beam and noise calibration, the correlation between the H- and V-channel beamforming weights directly impacts the estimation of the cross-correlation coefficient ρ_{hv} . In Fourier beamforming, this weight correlation remains close to unity, provided that the element patterns are accurately calibrated and matched between polarizations. As a result, the bias introduced in ρ_{hv} estimation is minimal. In contrast, for Capon beamforming, the weights are adaptively determined from the data covariance matrices. Therefore, the correlation between the H- and V-channel weights is no longer guaranteed to be unity and can vary between 0 and 1 depending on the scanning environment, clutter conditions, and target characteristics. This variability introduces a systematic bias in the estimation of ρ_{hv} .




Figure 3.9: The correlation between the Capon beamforming weights in the H and V channels is plotted for different diagonal loading levels: -50 dB (blue), -40 dB (orange), -30 dB (yellow), -20 dB (purple), and -10 dB (green).

This issue represents a fundamental limitation of unconstrained adaptive beamforming for polarimetric weather radar and is revisited in Chapter 7.

To understand more about this concept, consider the following formulation. First, Let the array-domain IQ data be

$$\mathbf{X}_h, \mathbf{X}_v \in \mathbb{C}^{N \times M}, \quad (3.21)$$

where N is the number of antenna elements and M is the number of pulses.

Define the beamforming weights

$$\mathbf{w}_h, \mathbf{w}_v \in \mathbb{C}^{N \times 1}. \quad (3.22)$$

The beamformed signals are

$$y_h = \mathbf{w}_h^H \mathbf{X}_h, \quad y_v = \mathbf{w}_v^H \mathbf{X}_v, \quad (3.23)$$

where $y_h, y_v \in \mathbb{C}^{1 \times M}$.

Define the covariance matrices as

$$\mathbf{R}_{hh} = \mathbb{E}\{\mathbf{X}_h \mathbf{X}_h^H\}, \quad (3.24)$$

$$\mathbf{R}_{vv} = \mathbb{E}\{\mathbf{X}_v \mathbf{X}_v^H\}, \quad (3.25)$$

$$\mathbf{R}_{hv} = \mathbb{E}\{\mathbf{X}_h \mathbf{X}_v^H\}, \quad (3.26)$$

$$\mathbf{R}_{vh} = \mathbb{E}\{\mathbf{X}_v \mathbf{X}_h^H\}. \quad (3.27)$$

The cross-correlation coefficient is defined as

$$\rho_{hv}^{\text{bf}} = \frac{|\mathbb{E}\{y_v y_h^H\}|}{\sqrt{\mathbb{E}\{y_h y_h^H\} \mathbb{E}\{y_v y_v^H\}}}. \quad (3.28)$$

Substituting the beamformed signals,

$$\rho_{hv}^{\text{bf}} = \frac{|\mathbb{E}\{(\mathbf{w}_v^H \mathbf{X}_v)(\mathbf{w}_h^H \mathbf{X}_h)^H\}|}{\sqrt{\mathbb{E}\{(\mathbf{w}_h^H \mathbf{X}_h)(\mathbf{w}_h^H \mathbf{X}_h)^H\} \mathbb{E}\{(\mathbf{w}_v^H \mathbf{X}_v)(\mathbf{w}_v^H \mathbf{X}_v)^H\}}}. \quad (3.29)$$

Using $(\mathbf{AB})^H = \mathbf{B}^H \mathbf{A}^H$,

$$(\mathbf{w}_h^H \mathbf{X}_h)^H = \mathbf{X}_h^H \mathbf{w}_h, \quad (3.30)$$

we obtain

$$\rho_{hv}^{\text{bf}} = \frac{|\mathbb{E}\{\mathbf{w}_v^H \mathbf{X}_v \mathbf{X}_h^H \mathbf{w}_h\}|}{\sqrt{\mathbb{E}\{\mathbf{w}_h^H \mathbf{X}_h \mathbf{X}_h^H \mathbf{w}_h\} \mathbb{E}\{\mathbf{w}_v^H \mathbf{X}_v \mathbf{X}_v^H \mathbf{w}_v\}}}. \quad (3.31)$$

Conditioning on the beamforming weights $\mathbf{w}_h$ and $\mathbf{w}_v$, which are assumed known for

the following derivation,

$$\rho_{\text{hv}}^{\text{bf}} = \frac{|\mathbf{w}_{\text{v}}^H \mathbf{R}_{\text{vh}} \mathbf{w}_{\text{h}}|}{\sqrt{(\mathbf{w}_{\text{h}}^H \mathbf{R}_{\text{hh}} \mathbf{w}_{\text{h}})(\mathbf{w}_{\text{v}}^H \mathbf{R}_{\text{vv}} \mathbf{w}_{\text{v}})}}. \quad (3.32)$$

Recalling the Capon beamforming weights for the H and V channels are

$$\mathbf{w}_{\text{h,Capon}} = \frac{\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}, \quad (3.33)$$

$$\mathbf{w}_{\text{v,Capon}} = \frac{\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}, \quad (3.34)$$

where $\mathbf{a}_{\text{h}}$ and $\mathbf{a}_{\text{v}}$ denote the steering vectors for the H and V channels.

Substituting these weights into the beamformed cross-correlation coefficient expression (3.32),

$$\rho_{\text{hv}}^{\text{bf, Capon}} = \frac{\left| \left(\frac{\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \right)^H \mathbf{R}_{\text{vh}} \left(\frac{\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}} \right) \right|}{\sqrt{\left[\left(\frac{\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}} \right)^H \mathbf{R}_{\text{hh}} \left(\frac{\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}} \right) \right] \left[\left(\frac{\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \right)^H \mathbf{R}_{\text{vv}} \left(\frac{\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \right) \right]}}. \quad (3.35)$$

Since the normalization terms in the Capon weights are scalars, define

$$\alpha = \frac{1}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}, \quad (3.36)$$

$$\beta = \frac{1}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}, \quad (3.37)$$

so that

$$\mathbf{w}_{\text{h,Capon}} = \alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}, \quad (3.38)$$

$$\mathbf{w}_{\text{v,Capon}} = \beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}. \quad (3.39)$$

Then,

$$\rho_{\text{hv}}^{\text{bf, Capon}} = \frac{|(\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{R}_{\text{vh}} (\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})|}{\sqrt{[(\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H \mathbf{R}_{\text{hh}} (\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})] [(\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{R}_{\text{vv}} (\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})]}}. \quad (3.40)$$

Using the Hermitian transpose property,

$$(\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H = \beta^* \mathbf{a}_{\text{v}}^H (\mathbf{R}_{\text{vv}}^{-1})^H, \quad (3.41)$$

$$(\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H = \alpha^* \mathbf{a}_{\text{h}}^H (\mathbf{R}_{\text{hh}}^{-1})^H. \quad (3.42)$$

Because $\mathbf{R}_{\text{hh}}$ and $\mathbf{R}_{\text{vv}}$ are Hermitian covariance matrices (see Appendix C), their inverses are also Hermitian:

$$(\mathbf{R}_{\text{hh}}^{-1})^H = \mathbf{R}_{\text{hh}}^{-1}, \quad (\mathbf{R}_{\text{vv}}^{-1})^H = \mathbf{R}_{\text{vv}}^{-1}. \quad (3.43)$$

Therefore,

$$\rho_{\text{hv}}^{\text{bf, Capon}} = \frac{|\alpha \beta| |\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{R}_{\text{vh}} \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}|}{\sqrt{|\alpha|^2 (\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{R}_{\text{hh}} \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}) |\beta|^2 (\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{R}_{\text{vv}} \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})}}. \quad (3.44)$$

Using $\mathbf{R}^{-1} \mathbf{R} = \mathbf{I}$,

$$\rho_{\text{hv}}^{\text{bf, Capon}} = \frac{|\alpha \beta| |\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{R}_{\text{vh}} \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}|}{\sqrt{|\alpha|^2 (\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}) |\beta|^2 (\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})}}. \quad (3.45)$$

Since α and β are real positive scalars for positive definite covariance matrices, the scalar factors cancel, yielding

$$\rho_{\text{hv}}^{\text{bf, Capon}} = \frac{|\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{R}_{\text{vh}} \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}|}{\sqrt{(\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}) (\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})}}. \quad (3.46)$$

For conventional beamforming with the Fourier method using steering vectors,

$$\rho_{\text{hv}}^{\text{bf, Fourier}} = \frac{|\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vh}} \mathbf{a}_{\text{h}}|}{\sqrt{(\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}} \mathbf{a}_{\text{h}})(\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}} \mathbf{a}_{\text{v}})}} = \frac{|C_{\text{hv}}^{\text{Fourier}}|}{\sqrt{S_{\text{h}}^{\text{Fourier}} S_{\text{v}}^{\text{Fourier}}}}. \quad (3.47)$$

These expressions highlight a fundamental structural difference between Fourier and Capon beamforming in the estimation of ρ_{hv} . In the Fourier case, both the numerator and denominator are formed by direct projections of the covariance matrices onto the steering vectors. As a result, the same deterministic weighting is applied to both channels, preserving the intrinsic H–V correlation structure. Under proper calibration, this leads to an unbiased estimate of ρ_{hv} .

In contrast, the Capon formulation introduces inverse covariance weighting, where $\mathbf{R}_{\text{hh}}^{-1}$ and $\mathbf{R}_{\text{vv}}^{-1}$ independently shape the H- and V-channel responses. Consequently, the numerator involves a cascaded operation,

$$\mathbf{R}_{\text{vv}}^{-1} \mathbf{R}_{\text{vh}} \mathbf{R}_{\text{hh}}^{-1},$$

which alters the cross-correlation structure rather than simply projecting it. Because these weights are data-dependent and generally differ between polarizations, the relative alignment between H and V channels is no longer preserved.

This mismatch propagates into the normalization terms as well, where the denominator is no longer proportional to the beamformed power but instead depends on quadratic forms involving the inverse covariance matrices. As a result, the Capon estimate of ρ_{hv} is inherently biased, with the magnitude of the bias governed by the degree of disparity between $\mathbf{R}_{\text{hh}}$ and $\mathbf{R}_{\text{vv}}$, as well as the surrounding scattering and clutter environment. This sensitivity to scene-dependent statistics makes the bias difficult to characterize or correct consistently, representing a key limitation of Capon beamforming in polarimetric radar applications.

3.6 H–V Channel Mismatch in Adaptive Beamforming

Weather radars transmit and receive signals independently in the H and V polarization channels. In theory, when DBF is performed using a deterministic method, the beam patterns in the H and V channels should be identical. However, in practice, differences can arise due to variations in element patterns caused by hardware imperfections. These discrepancies can be mitigated through element pattern calibration, which involves measuring and correcting the individual responses of each element. Another common approach is H–V beam matching, where beamforming weights derived for one channel are applied to both channels to maintain consistency. This matching is important for polarimetric variable estimation, as both channels must sample the same target volume.

In contrast, in adaptive beamforming, the beamforming weights are computed independently for each polarization channel based on the received data. As a result, the beams formed in the H and V channels are not guaranteed to be identical, and their differences depend on both the target characteristics and the surrounding environment. Figure 3.10 illustrates these differences in H and V beams for Cases 1, 2, and 3 from the simulated models in Figure 3.1. For uniformly distributed targets (Case 1), the Capon method produces nearly identical beam patterns in both channels. In the point-target case (Case 2) with a negative Z_{DR} , the H and V beams differ noticeably near the target, as the adaptive beamformer suppresses power differently in each channel. In Case 3, where both reflectivity and Z_{DR} vary gradually, the largest differences occur in transition regions where these




Figure 3.10: Differences between the Capon beam patterns in the H and V channels for Cases 1, 2, and 3 (see Figure 3.1).

quantities change rapidly.

Beamforming weights are complex-valued. In the Fourier method, the weights typically have uniform amplitude across all elements, with phases varying linearly according to the steering direction. In contrast, the Capon method derives weights through an optimization process, resulting in more complex amplitude and phase structures. These weights do not follow a simple analytical form and are often irregular, making them difficult to predict or model. Consequently, applying H–V beam matching in the Capon method can introduce bias between the channels, making it an unsuitable approach.

This section examines the differences in H and V beam patterns and beamforming weights, and analyzes how these differences affect the estimation of polarimetric variables.

3.6.1 H–V Beamforming Weights Mismatch

In addition to beam pattern differences, adaptive beamforming weights are complex-valued, and the beam pattern represents the coherent sum of these complex weights. Even when the beam shapes of the H and V channels appear similar, the complex-valued beamforming weights can differ in both phase and amplitude. This highlights that beam matching in adaptive polarimetric processing extends beyond matching beam shapes, and requires consistency in the phase-sensitive behavior of the beamforming weights themselves. These differences directly impact polarimetric measurements that rely on the joint behavior of the H and V signals. In particular, the cross-correlation coefficient (ρ_{hv}), which quantifies the correlation relationship between the H and V channels, is sensitive to mismatches in the complex-valued beamforming weights. Any inconsistency between the adaptive beams in the two channels can therefore introduce bias and increased variability in polarimetric variables. This makes the calibration of H–V beam consistency a critical challenge in adaptive beamforming for weather radar applications.

Figure 3.11 shows the amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods when the radar is steered at broadside ($\theta = 0^\circ$, $\phi = 0^\circ$). The beams are formed in a uniformly distributed target environment so that the Capon beamformer produces a beam shape similar to that of the Fourier method. When the radar is steered at broadside, the phase of the Fourier beamforming weights is zero for all elements, and the amplitudes are identical across the array. In contrast, although the Capon beamformer produces a beam shape that closely matches the Fourier beam, the corresponding beamforming weights differ substantially and exhibit more irregular behavior in both amplitude and phase.

Figure 3.12 shows the amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods when the radar is steered off broadside ($\theta = 20^\circ$, $\phi = 0^\circ$). As before, the beams are formed in a uniformly distributed target environment. When the radar steers away from broadside, the Fourier beamformer electronically tilts the beam by applying different phase shifts depending on the element location. This results in a clear linear relationship between the phase of the beamforming weights (the desired steering angle) and the element locations, as discussed in Section 2.4 and described in (2.1). In contrast, the Capon beamforming weights exhibit more dynamic behavior across the elements, with variations in both amplitude and phase from element to element. This is because the Capon method derives its weights through a higher-order optimization process, resulting in significantly more complex amplitude and phase structures.

Figure 3.13 shows the correlation between the Capon beamforming weights in the H and V channels as a function of SNR (y -axis) for Cases 1, 2, and 3. The x -axis represents the angular coordinate and corresponds to the x -axis in Figure 3.1. Overall, the results indicate high H–V weight correlation at higher SNR levels (greater than 20 dB). However, in Case 2 (point-target scenario), regions near the point target—where non-zero Z_{DR} values




Figure 3.11: The amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods, when the radar is steered at broadside ($\theta = 0^\circ$, $\phi = 0^\circ$), are shown above.

are present—exhibit reduced H–V correlation even at high SNR (> 20 dB). In general, increasing SNR improves the correlation between the Capon beamforming weights. This trend is consistent with the ρ_{hv} results shown in Figure 3.9, where improved estimation performance is also observed at higher SNR levels.

3.6.2 H–V Beam Matching

H–V beam matching is a widely used technique in weather radar to ensure accurate estimation of polarimetric variables. However, as discussed earlier, strict H–V beam matching is not always desirable for the Capon method, since the adaptive beamforming weights are derived independently for each polarization channel based on the received data. In adaptive beamforming, beam matching extends beyond aligning beam shapes; the beamforming weights themselves are complex and phase-sensitive. Because the Capon method optimizes these weights for the signal in each channel, applying weights derived from one




Figure 3.12: The amplitude (left column) and phase (right column) of the beamforming weights produced by the Fourier (top row) and Capon (bottom row) methods, when the radar is steered off broadside ($\theta = 20^\circ, \phi = 0^\circ$), are shown above.




Figure 3.13: Correlation between the Capon beamforming weights in the H and V channels as a function of SNR (shown on the y -axis) for Cases 1, 2, and 3. The x -axis represents the angular coordinate and corresponds to the x -axis in Figure 3.1.

channel (e.g., H) to another channel (e.g., V) can introduce errors and distortions. In this subsection, we examine how H–V beam matching affects the estimation of Z_{DR} and ρ_{hV} .

Unlike ρ_{hV} , the bias in Z_{DR} can be significantly reduced after proper beam pattern calibration, even when adaptive beamforming is applied. This is because Z_{DR} is fundamentally a differential power measurement between the H and V channels, and does not depend on the cross-correlation structure between channels.

The beamformed Z_{DR} is defined as

$$Z_{\text{DR}}^{\text{bf}} = \frac{S_{\text{h}}}{S_{\text{v}}} = \frac{\mathbb{E}\{|\mathbf{w}_{\text{h}}^H \mathbf{V}_{\text{h}}|^2\}}{\mathbb{E}\{|\mathbf{w}_{\text{v}}^H \mathbf{V}_{\text{v}}|^2\}} \quad (3.48)$$

$$= \frac{\mathbb{E}\{(\mathbf{w}_{\text{h}}^H \mathbf{V}_{\text{h}})(\mathbf{w}_{\text{h}}^H \mathbf{V}_{\text{h}})^H\}}{\mathbb{E}\{(\mathbf{w}_{\text{v}}^H \mathbf{V}_{\text{v}})(\mathbf{w}_{\text{v}}^H \mathbf{V}_{\text{v}})^H\}} \quad (3.49)$$

$$= \frac{\mathbb{E}\{\mathbf{w}_{\text{h}}^H \mathbf{V}_{\text{h}} \mathbf{V}_{\text{h}}^H \mathbf{w}_{\text{h}}\}}{\mathbb{E}\{\mathbf{w}_{\text{v}}^H \mathbf{V}_{\text{v}} \mathbf{V}_{\text{v}}^H \mathbf{w}_{\text{v}}\}} \quad (3.50)$$

$$= \frac{\mathbf{w}_{\text{h}}^H \mathbf{R}_{\text{hh}} \mathbf{w}_{\text{h}}}{\mathbf{w}_{\text{v}}^H \mathbf{R}_{\text{vv}} \mathbf{w}_{\text{v}}}. \quad (3.51)$$

For Capon beamforming, substituting (3.38) and (3.39) into the beamformed Z_{DR} expression gives

$$Z_{\text{DR}}^{\text{bf, Capon}} = \frac{(\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H \mathbf{R}_{\text{hh}} (\alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})}{(\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{R}_{\text{vv}} (\beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})} \quad (3.52)$$

$$= \frac{\alpha^* (\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H \mathbf{R}_{\text{hh}} \alpha \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\beta^* (\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{R}_{\text{vv}} \beta \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \quad (3.53)$$

$$= \frac{|\alpha|^2 (\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H \mathbf{R}_{\text{hh}} \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{|\beta|^2 (\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{R}_{\text{vv}} \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \quad (3.54)$$

$$= \frac{|\alpha|^2 (\mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}})^H \mathbf{a}_{\text{h}}}{|\beta|^2 (\mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}})^H \mathbf{a}_{\text{v}}} \quad (3.55)$$

$$= \frac{|\alpha|^2 \mathbf{a}_{\text{h}}^H (\mathbf{R}_{\text{hh}}^{-1})^H \mathbf{a}_{\text{h}}}{|\beta|^2 \mathbf{a}_{\text{v}}^H (\mathbf{R}_{\text{vv}}^{-1})^H \mathbf{a}_{\text{v}}}. \quad (3.56)$$

Since covariance matrices are Hermitian (see Appendix C), their inverses are also Hermitian. Therefore,

$$Z_{\text{DR}}^{\text{bf, Capon}} = \frac{|\alpha|^2 \mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{|\beta|^2 \mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}. \quad (3.57)$$

Substituting the definitions of α and β ,

$$Z_{\text{DR}}^{\text{bf, Capon}} = \frac{\left| \frac{1}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}} \right|^2 \mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}{\left| \frac{1}{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} \right|^2 \mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}} = \frac{\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}}{\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}}. \quad (3.58)$$

This expression shows that $Z_{\text{DR}}^{\text{Capon}}$ is sensitive to the beamforming weights formed independently in each channel. Specifically, the numerator and denominator are governed by the interactions $\mathbf{a}_{\text{v}}^H \mathbf{R}_{\text{vv}}^{-1} \mathbf{a}_{\text{v}}$ and $\mathbf{a}_{\text{h}}^H \mathbf{R}_{\text{hh}}^{-1} \mathbf{a}_{\text{h}}$. As a result, the ratio depends directly on the channel-specific inverse covariance weighting rather than the beamformed power itself. If H-V beam matching is imposed by replacing $\mathbf{a}_{\text{v}}$ with $\mathbf{a}_{\text{h}}$, this structure is no longer preserved. Such substitution introduces a mismatch between the steering vector and the corresponding covariance matrix, i.e., $\mathbf{a}_{\text{h}}$ interacting with $\mathbf{R}_{\text{vv}}^{-1}$, which is not physically

consistent with the V-channel statistics. The resulting Z_{DR} estimate becomes distorted.

Now, we examine how H–V beam matching impacts the estimation of Z_{DR} using simulated data from the simple model (see Figure 3.1). Two Capon beamforming cases are considered: one where the H and V beams are formed independently, and one where the H and V beams are matched. Figure 3.14 shows the bias in Z_{DR} estimation as a function of diagonal loading (y -axis) and ρ_{hv} values (x -axis) for the H–V matched case (left) and the H–V independent case (right). The bias is computed using 5000 simulated samples, and the plotted values represent the estimated mean bias.




Figure 3.14: The bias in Z_{DR} estimation as a function of diagonal loading (the y -axis) and ρ_{hv} (the x -axis) for the H–V matched case (left plot) and the H–V independent case (right plot) are presented.

The results show that, after proper beam pattern calibration, the H–V independent approach introduces negligible bias in Z_{DR} , with values close to 0 dB. In contrast, when H and V beams are matched, the resulting bias varies significantly with both the diagonal loading level and the true ρ_{hv} value. The bias generally decreases with increasing diagonal loading and/or higher ρ_{hv} , indicating improved performance when the H and V signals are more strongly correlated.

Then, we examine the performance of Z_{DR} estimation under H–V matched and non-matched conditions for different SNR levels and ρ_{hv} values.




Figure 3.15: The bias in Z_{DR} estimation as a function of SNR (the y -axis) and ρ_{hv} (the x -axis) for the H–V matched case (left plot) and the H–V independent case (right plot) are presented.

Figure 3.15 shows the bias in Z_{DR} estimation as a function of SNR and ρ_{hv} for the H–V matched case (left plot) and the H–V independent case (right plot). Consistent with the results in Figure 3.14, the H–V independent approach introduces little bias in Z_{DR} , with values remaining close to 0 dB. A similar trend is observed with respect to SNR, where higher SNR levels result in reduced bias.

The performance of H–V beam matching depends on three key factors: the diagonal loading level, the SNR, and the intrinsic ρ_{hv} . At low SNR (approximately 5 dB), H–V beam matching can provide some improvement in Z_{DR} bias, but the benefit is relatively modest. Introducing diagonal loading in combination with H–V beam matching leads to more noticeable bias reduction. However, even with increased diagonal loading, the region in which the bias remains below 0.03 dB does not expand significantly. This indicates that while diagonal loading-assisted H–V beam matching can mitigate bias to some extent, it does not fully overcome the fundamental limitations imposed by adaptive beam differences between the H and V channels.

Next, the impact of H–V beam matching on the estimation of ρ_{hv} is examined. The same experiment is repeated, evaluating performance over different diagonal loading levels,

SNR conditions, and true ρ_{HV} values.




Figure 3.16: The bias in ρ_{HV} estimation as a function of diagonal loading (top row) and SNR (bottom row), with the intrinsic ρ_{HV} shown on the x -axis, for the H–V matched case (left column) and the H–V independent case (right column).

Figure 3.16 shows the bias in ρ_{HV} estimation as a function of diagonal loading (top row) and SNR (bottom row), with the intrinsic ρ_{HV} shown on the x -axis, for the H–V matched case (left column) and the H–V independent case (right column). Unlike Z_{DR} , the estimation of ρ_{HV} improves when the H and V beams are matched at low SNR and low diagonal loading, as expected from the discussion in Section 3.5. However, for diagonal loading levels above -30 dB and SNR greater than approximately 10 dB, the matched and independent H–V cases produce similar results. Overall, trends similar to those observed for

Z_{DR} are evident: higher SNR and increased diagonal loading both reduce the bias in ρ_{hv} estimation.

For the Capon method, it is generally preferable to use H–V independent beams, since the beamformer optimizes the weights based on the received target power in each polarization channel. However, computing and calibrating adaptive beams independently for both channels can be computationally expensive. As an alternative, H–V beam matching can still be applied by appropriately tuning the diagonal loading levels based on the expected ρ_{hv} values and the SNR, providing a practical compromise between computational efficiency and polarimetric accuracy.

3.7 Simulation-Based Evaluation with NEXRAD Data

PAR time-series signals were simulated using NEXRAD data and the Horus PAR geometry. Details of the simulation are provided in Appendix A. This configuration yields an approximate receive beamwidth of 2.8° , whereas the NEXRAD beamwidth is approximately 1° . This difference in spatial resolution should be considered when interpreting the results.

The original NEXRAD data were sampled at 0.5° azimuthal resolution; therefore, DBF (both Fourier and Capon) was performed at 1° resolution. The range resolution of both NEXRAD and the simulation is 250 m. DBF was performed in the azimuth direction. For the elevation dimension, low-power signals with a uniform distribution (i.e., noise-like) were simulated to prevent the Capon pattern from becoming overly impulsive in the vertical direction.

A Gaussian-shaped transmit beam with a beamwidth of approximately 15° was used to simulate spoiled-beam signals. Within each transmit beam, 10 receive beams were formed with a 0.5° sampling interval. A rotating PAR was simulated, with the radar rotating every

10° to cover 360° . This PAR geometry and scanning concept are illustrated in Figure 3.17.




Figure 3.17: The geometry of the PAR is illustrated in the left plot. The middle plot shows the transmit beam configuration, the degree of overlap between adjacent transmit beams, and the number of receive beams formed within each transmit beam. The right plot illustrates how the rotating PAR scanning strategy was simulated.

Figure 3.18 shows a side-by-side comparison of radar-variable products generated using the Fourier and Capon DBF methods, with the original NEXRAD moment data shown in the left column for reference. This case was collected by the KTLX radar on May 6, 2013 at 20:00:50.96 UTC.

In the reflectivity panel (top row), both DBF methods preserve the general storm structure seen in the original product, but the Capon method shows a slightly cleaner and sharper reflectivity pattern compared to the Fourier DBF. The Fourier result appears smeared, with higher reflectivity values estimated near the edges of storms compared to the original. This overestimation and smoothing are due to the fixed beam pattern, which can blur fine-scale structures.

For the velocity field (second row), both methods generally capture the large-scale flow, but the Capon DBF exhibits smoother transitions, especially in regions with complex velocity signatures and velocity aliasing, such as around the 270° azimuth where the hook echo is present. Compared to the Fourier DBF, the Capon method reduces noisy patches and better defines shear and convergence zones, likely due to its adaptive nature that suppresses




Figure 3.18: The original NEXRAD Level-II data (left column), simulated Fourier DBF (middle column), and Capon beamforming (right column) are presented. The case is from the KTLX radar on May 6, 2013 at 20:00:50.96 UTC. The rows from top to bottom show reflectivity (Z_H), radial velocity (v_r), differential reflectivity (Z_{DR}), correlation coefficient (ρ_{hv}), and differential phase (ϕ_{DP}).

sidelobes and noise more effectively.

Looking at Z_{DR} (third row), the Capon result is noticeably less noisy than the Fourier DBF, particularly in lower SNR regions. The Fourier DBF shows clear signs of sidelobe contamination and striping artifacts, while Capon maintains a more continuous and physically realistic pattern.

In the ρ_{hv} field (fourth row), the Fourier method preserves high correlation values across the main storm regions, except near the storm edges where signal power is weak. In contrast, the Capon method shows degraded correlation due to uncorrelated beamforming weights between the H and V channels, as discussed in earlier sections. This highlights a limitation of the Capon method.

Lastly, the ϕ_{DP} field (bottom row) from the Capon method appears smoother and more consistent throughout the entire Plan Position Indicator (PPI) scan compared to the result from the Fourier method. This improvement is attributed to the adaptive beamforming of the Capon method, whereas the fixed beam pattern of the Fourier method leads to increased contamination.

Figure 3.19 shows the azimuthal profile of the estimated polarimetric variables from the PPI scan at a fixed range of 26.875 km, which intersects the tornado hook echo region. The four panels represent: Z_H (top-left), v_r (top-right), Z_{DR} (bottom-left), and ρ_{hv} (bottom-right). The black line shows the truth field from the simulation, the blue line is the result from Fourier beamforming, the red line represents Capon beamforming with independently computed H and V channel weights, and the green line shows the result from matched H–V beamforming using Capon method. Results from H–V beam matching are shown only for Z_{DR} and ρ_{hv} , as only these two variables are affected.

In the reflectivity field, Capon beamforming provides sharper transitions and more accurate localization of features, particularly in the region of strong reflectivity between 270° and 320° azimuth, which corresponds to the hook echo structure. Compared to the Capon




Figure 3.19: Azimuthal line plots at 26.875 km range comparing reflectivity (top-left), velocity (top-right), differential reflectivity (bottom-left), and correlation coefficient (bottom-right) for the truth (black), Fourier (blue), Capon with matched H–V beams (green), and Capon with independent beams (red).

result, the Fourier method shows smearing and less defined gradients due to its fixed beam shape and higher sidelobe contamination. The ability of the Capon method to suppress sidelobe leakage results in finer spatial resolution and clearer structural definition. However, sharp azimuthal transitions appear every 10° in the Capon result; this artifact arises from the simulation setup where transmit beams are spaced at 10° intervals. At the edges of these transmit beams, the Capon beamforming performance degrades due to discontinuities in the data and lack of overlapping information, which leads to unstable beamforming behavior. This represents a limitation of the simulation. It also explains the increased bias produced by the Capon method around 260° , where large power differences occur between adjacent transmit beams.

In the radial velocity panel, Capon beamforming demonstrates notable advantages in two key areas. First, at the storm edge near 75° azimuth, Capon beamforming resolves a sharp velocity gradient more clearly than the Fourier method, which exhibits approximately 5° of beam smearing. Second, around the hook echo region, where velocity aliasing occurs due to intense rotational shear, Capon beamforming effectively captures the rapid phase wrapping transitions, whereas the Fourier output fails to resolve these transitions. These

improvements highlight the potential of Capon beamforming for enhanced storm-edge detection and improved characterization of vortex features.

For differential reflectivity, both the Fourier and Capon beamforming produce generally consistent outputs, although the Capon beamforming captures some finer details with higher variability. This suggests that the Capon beamforming may provide enhanced data quality in certain areas. However, in regions where the data become sparse or noisy, the Capon estimates can become more erratic due to sensitivity to the underlying signal structure.

Finally, the cross-correlation coefficient reveals a significant limitation of Capon beamforming due to the reduced correlation between the H and V beamforming weights compared to Fourier beamforming. While the truth and Fourier estimates maintain high values (> 0.95), the Capon beamforming estimate struggles to preserve coherence, especially in regions of reflectivity gradients or data discontinuity. These results indicate a fundamental tension between unconstrained adaptive beamforming and the preservation of polarimetric coherence, rather than simply a calibration issue. Therefore, calibration and further methodological adjustments are required to make Capon beamforming suitable for reliable ρ_{hv} estimation.

In the next chapter, space-time processing will be introduced and discussed. Space-time processing is a joint technique in both spatial and temporal dimensions, representing another capability of PAR. This technique will be studied and evaluated for weather radar applications, with particular focus on its benefits under clutter conditions.

Chapter 4

Space-Time Adaptive Processing

Meteorological radar observations are frequently contaminated by clutter, which includes stationary targets on the ground such as buildings and communication towers, as well as moving targets such as wind turbines and airplanes. Conventional filtering algorithms used in reflector-based weather radars operate primarily in the Doppler spectrum and are effective at removing stationary non-meteorological scatterers. However, these algorithms can induce significant biases in radar-variable estimates when they attempt to identify and remove moving clutter targets. One widely used clutter-filtering method for weather radars is Gaussian Model Adaptive Processing (GMAP), which relies mainly on Doppler processing techniques [68]. Since stationary ground clutter has predictable signal characteristics in the Doppler domain (i.e., near 0 m/s velocity), GMAP applies a Gaussian clutter model to remove spectral components associated with stationary clutter. However, GMAP does not effectively mitigate moving clutter because its signal characteristics are not predictable in the Doppler domain. This limitation arises from its one-dimensional (Doppler-only) perspective, whereas moving clutter often requires a joint spatial–temporal representation, which motivates the use of STAP.

Fully digital PAR provides maximum flexibility for investigating signal-processing methods that may overcome the challenge of mitigating moving clutter targets. One such method is STAP[69]. STAP adaptively adjusts the two-dimensional space-time fil-

ter response through the joint use of DBF and Doppler spectral processing. It is a well-established technique for clutter mitigation in airborne and spaceborne moving target indication (MTI) radars [32–34]. More recently, STAP has also shown potential for mitigating wind turbine clutter (WTC) [70], as illustrated in Figure 4.1.




Figure 4.1: Example STAP application to remove simulated WTC from actual digital PAR weather data (AIR); (left) weather only, (middle) WTC contamination on weather, and (right) recovered weather [70].

Wind turbine clutter is particularly challenging for conventional reflector-based weather radars because it contaminates data across both spatial and temporal domains. Figure 4.2 shows a Doppler-time plot from a field deployment of Horus at the Redbed Wind Farm in Tuttle, Oklahoma, on July 23, 2025. The figure demonstrates that the spectral contamination is very broad, extending across much of the Doppler spectrum, where the maximum unambiguous velocity is approximately $\pm 48 \text{ m s}^{-1}$. This dataset will be further studied and analyzed in Chapter 5.

In this chapter, the application of space-time processing to weather radar observations is evaluated, especially its ability to reduce contamination from moving clutter, thereby enhancing weather observations. Space-time processing will be performed using deterministic (STP) and adaptive (STAP) space-time weights on simulated weather data in clutter environments. The trade-offs between non-adaptive space-time processing and adaptive processing will be studied. In addition, different PAR architectures will be assessed, including fully digital and subarray systems, to understand their impacts on the quality of




Figure 4.2: Doppler–time plot from the field deployment of Horus at the Redbed Wind Farm in Tuttle, Oklahoma on July 23, 2025.

weather estimates after space-time processing in cluttered environments. Finally, Doppler velocity estimates from meteorological returns are evaluated using STP and STAP. The proposed simulation enables precise control of radar and weather/clutter signal settings such as pulse number, antenna configurations, clutter velocity, desired weather characteristics, SNR, clutter-to-signal ratio (CSR), among others. This level of control allows for an in-depth investigation of the impact of space-time processing on weather observations in cluttered environments across various PAR configurations.

4.1 Fundamentals of Space-Time Processing

space-time processing involves a two-dimensional filtering operation that jointly exploits spatial (array) and temporal (Doppler) dimensions through DBF and Doppler processing. Figure 4.3 illustrates the overall framework.




Figure 4.3: Illustration of the digital PAR data cube and the space-time processing framework.

The process begins with a fully digital (or highly digitized) PAR data cube containing time-series IQ signals from multiple receivers, with dimensions $N \times M \times L$, where N is the number of receivers, M is the number of temporal samples, and L is the number of range bins. For a fixed range bin l , the data are reshaped into a space-time snapshot vector

$\mathbf{y} \in \mathbb{C}^{MN \times 1}$:

$$\mathbf{y} = \begin{bmatrix} V_1(1) \\ \vdots \\ V_1(M) \\ V_2(1) \\ \vdots \\ V_N(M) \end{bmatrix}, \quad (4.1)$$

where $V_n(m)$ denotes the complex IQ sample at element n and time index m .

The sample covariance matrix is estimated as

$$\hat{\mathbf{S}} = \mathbf{y}\mathbf{y}^H, \quad (4.2)$$

which is an $MN \times MN$ matrix. For a single snapshot, this estimate is rank-1. In practice, averaging over multiple snapshots is required to obtain a full-rank covariance matrix.

The temporal (Doppler) steering vector is defined as

$$\mathbf{a}_t(f_D) = [1, e^{-j2\pi f_D T}, \dots, e^{-j2\pi(M-1)f_D T}]^T, \quad (4.3)$$

and the spatial steering vector is

$$\mathbf{a}_s(\theta) = [1, e^{-jkd \sin \theta}, \dots, e^{-j(N-1)kd \sin \theta}]^T, \quad (4.4)$$

where $f_D = 2v_r/\lambda$ is the Doppler frequency, $k = 2\pi/\lambda$ is the wavenumber, d is the element spacing, and θ is the steering angle.

The space-time steering vector is formed using the Kronecker product:

$$\mathbf{t} = \mathbf{a}_t(f_D) \otimes \mathbf{a}_s(\theta), \quad (4.5)$$

with $\mathbf{t} \in \mathbb{C}^{MN \times 1}$.

For the deterministic method (STP), the output power is computed as

$$P_{\text{STP}} = \mathbf{t}^H \mathbf{S} \mathbf{t}. \quad (4.6)$$

For the adaptive method (STAP), adaptive weights are obtained using the Capon formulation:

$$\mathbf{h} = \frac{\mathbf{S}^{-1} \mathbf{t}}{\mathbf{t}^H \mathbf{S}^{-1} \mathbf{t}}, \quad (4.7)$$

where $\mathbf{h} \in \mathbb{C}^{MN \times 1}$. The corresponding output power is

$$P_{\text{STAP}} = \mathbf{h}^H \mathbf{S} \mathbf{h}. \quad (4.8)$$

4.2 Simulation-based Evaluation of Space-Time Processing

space-time processing is evaluated using simulated data. Different types of targets, including stationary ground clutter, moving point targets, weather-like distributed targets, and wind turbines, are modeled to assess the impact and benefits of space-time processing. The overall simulation procedure is described in Appendix A. To utilize this framework, the Doppler spectrum of each target type must be modeled appropriately to reflect its physical characteristics.

4.2.1 Spectrum Characteristics of Different Types of Targets

Ground clutter consists of both direct-current (DC) and alternating-current (AC) components. Due to the high temporal correlation and slow temporal variation of clutter signals, a commonly used model for the ground clutter Doppler spectrum, $P_{GC}(v)$, is a triangle-

shaped spectrum given by [71]

$$P_{GC}(v) = \frac{r}{r+1} \delta(v) + \frac{1}{r+1} P_{ac}(v), \quad (4.9)$$

where the first term represents the DC component at zero velocity, corresponding to stationary scatterers such as terrain and buildings. The second term represents the AC component, associated with moving scatterers such as vegetation.

The parameter r denotes the ratio of DC to AC power and is defined as

$$r = 489.8 w^{-1.55} f_o^{-1.21}, \quad (4.10)$$

where w is the wind speed and f_o is the radar operating frequency. The AC spectrum is modeled as

$$P_{ac}(v) = \frac{\beta}{2} \exp(-\beta|v|), \quad (4.11)$$

where $\beta = 5.7$ corresponds to moderate wind conditions (approximately 15–30 mph).

The Doppler spectrum of a moving point target is modeled using a Gaussian distribution [72], with specified mean radial velocity, spectrum width, and signal power at a given elevation (θ) and azimuth (ϕ).

Weather signals are characterized by nonzero mean velocities (except when motion is perpendicular to the radar beam) and broader spectrum widths compared to clutter. They also exhibit spatial continuity over extended regions. The radial velocity, v_r , of a laminar wind field is modeled as [73]

$$v_r = w \sin \theta + u \sin \phi \cos \theta + v \cos \phi \cos \theta, \quad (4.12)$$

where u , v , and w are the Cartesian components of the wind velocity vector.

A Gaussian Doppler spectrum is then generated with a prescribed spectrum width. Weather spectrum widths vary depending on storm type, with typical values around 2–4 m/s [74]. In this study, a representative value of 3 m/s is used.

A broader spectrum width improves separability between weather and point-target clutter in the Doppler domain. However, even when the weather spectrum is relatively narrow, space-time processing exploits both Doppler and spatial information, enabling improved discrimination between different target types.

4.2.2 Performance Evaluation of Space-Time Processing

Three cases are simulated: stationary ground clutter, weather-like targets, and a combination of both. The data are generated with 1024 temporal samples and an SNR of 30 dB. Figure 4.4 shows the modeled Doppler spectra: stationary ground clutter (left), weather-like targets (middle), and their combination (right). The x -axis represents radial velocity (m/s), and the y -axis represents PSD in dB. The ground clutter is simulated at -38° azimuth with zero mean radial velocity, while the weather-like targets are distributed from -45° to 45° in azimuth with varying mean radial velocities.




Figure 4.4: Modeled Doppler spectra for stationary ground clutter (left), weather-like targets (middle), and their combination (right). The x -axis represents radial velocity (m/s), and the y -axis represents power spectral density (PSD) in dB.

space-time processing is performed using a uniform linear array with both non-adaptive (STP) and adaptive (STAP) space-time weights. The resulting outputs for the ground clutter and weather-like target cases are shown in Figure 4.5. The x -axis represents radial velocity (m/s), and the y -axis represents azimuth ($^{\circ}$), providing a joint view of the temporal and spatial domains.

The STP results exhibit noticeable sidelobe contamination in both spatial and temporal dimensions, whereas STAP significantly suppresses this contamination. In the stationary ground clutter case, STAP reduces sidelobes and enables more accurate localization of clutter in both angle and Doppler. However, STP preserves higher output SNR due to the absence of adaptive suppression.




Figure 4.5: STP and STAP results for stationary ground clutter and weather-like targets. The x -axis represents radial velocity (m/s), and the y -axis represents azimuth ($^{\circ}$).

Figure 4.6 presents the results for the combined scenario, where stationary ground clutter and weather-like targets coexist. The 2D plots show the joint spatial-temporal response, while the 1D cuts correspond to velocity at -38° azimuth (top) and azimuth at 0 m/s (bottom). These cut locations are indicated by red lines in the 2D plots. In the 1D plots, STP and STAP results are shown in red and blue, respectively.

This case highlights the impact of sidelobe leakage in STP. The 1D cuts show that

sidelobes from STP extend into the weather-like target region, degrading the quality of the measurements. In contrast, STAP suppresses this leakage and effectively separates clutter from weather-like targets in both spatial and Doppler dimensions.




Figure 4.6: STP and STAP results for the combined scenario of stationary ground clutter and weather-like targets. The 2D plots show the joint spatial-temporal response. The 1D cuts correspond to velocity at -38° azimuth (top) and azimuth at 0 m/s (bottom), indicated by red lines in the 2D plots.

4.2.3 Clutter-to-Signal Ratio Impact on Space-Time Processing Performance

The performance of space-time processing is evaluated for varying CSR. The azimuth domain is simulated from -10° to 10° . The clutter is located at -8° azimuth with zero mean velocity, while the weather-like signals are distributed across azimuth with mean velocity varying as a function of azimuth.

Figure 4.7 shows the STP and STAP results for CSR values of 0 dB, 10 dB, 20 dB, 30 dB, 40 dB, and 50 dB. At 0 dB CSR, contamination from clutter is minimal. However,

once the clutter power exceeds the signal power ($\text{CSR} > 0 \text{ dB}$), sidelobe leakage from STP begins to contaminate the weather signals. As CSR increases, this contamination becomes more severe. At CSR values of 30 dB and above, it becomes difficult to recover weather signals using STP. Although the clutter is localized at -8° , its sidelobes extend across the entire azimuth domain, leading to widespread contamination.

In contrast, STAP effectively suppresses clutter contamination even at high CSR values, up to 50 dB. This demonstrates the capability of STAP to mitigate sidelobe leakage from strong clutter sources. Furthermore, when clutter and weather signals overlap in both spatial and Doppler domains, STAP exploits their joint space-time characteristics to improve separation. This enables effective clutter suppression even in scenarios where conventional spatial processing (e.g., DBF) or Doppler filtering alone would fail.




Figure 4.7: STP and STAP results for different clutter-to-signal ratios (CSR). The clutter is located at -8° azimuth with zero mean velocity, and the weather-like signals are distributed across azimuth with varying mean velocities.

4.2.4 Impact of Sample Size on Space-Time Adaptive Processing Performance

STAP requires a sufficient number of independent and identically distributed (i.i.d.) samples [41] to obtain a well-conditioned covariance matrix for stable inversion. However, using a large number of samples is not desirable for weather radar operations. Increasing the number of samples increases the dwell time, and the computational burden of STAP also increases due to covariance estimation and matrix inversion. This issue will be discussed further in Chapter 6. Here, we evaluate STAP performance for different sample sizes to examine how the number of samples affects both performance and computation time.

Figure 4.8 shows results from STP (red lines) and STAP (blue lines) for different numbers of samples, $M = 256, 512, 1024, 2048$. The top plots show 1D azimuth cuts, and the bottom plots show 1D velocity cuts. The computation time for STAP, measured on a MacBook Pro equipped with an Apple M1 processor, is listed for each case.




Figure 4.8: STP (red) and STAP (blue) results for different numbers of samples. The computation time for each case is shown at the bottom.

The results show that STP performance is largely unaffected by the number of samples, since its weights are determined by a fixed filter response and are not data-dependent. In contrast, STAP performance is strongly dependent on the number of samples. As the number of samples decreases, the covariance matrix becomes less accurately estimated and increasingly ill-conditioned, which degrades the adaptive weight computation. As a result, the effective SNR after STAP decreases, making it more difficult to detect weaker targets such as weather-like signals. This highlights that STAP performance is fundamentally constrained by sample support, making it highly sensitive to practical data availability.

Furthermore, the computation time of STAP increases with the number of samples due to the increased cost of covariance estimation and matrix inversion, while the computation time of STP remains nearly constant. Notably, even the smallest sample size considered in this study (256) already exceeds typical weather radar operations, where the number of samples is often 64 or fewer, underscoring a gap between theoretical performance and operational feasibility. In addition to increased computational cost, using more samples also increases dwell time, which is undesirable for practical weather radar applications. These requirements on the number of samples indicate that full STAP, in its present form, is difficult to deploy routinely within typical weather radar dwell-time constraints.

4.2.5 Additional Case Studies

Additional examples of STP and STAP are presented in Figure 4.9 (moving point targets) and Figure 4.10 (wind turbines). Figure 4.9 shows cases where point targets (both stationary and moving) overlap with weather-like targets. In this simulation, the SNR is set to 30 dB and the CSR is set to 20 dB.

Wind turbines with four different orientation angles (0° , 30° , 60° , and 90°) are simulated [35], and the corresponding STP and STAP results are shown in Figure 4.10. The




Figure 4.9: STP and STAP results for four cases involving stationary and moving point targets embedded in weather-like echoes.

orientation angle represents the relative angle between the radar line-of-sight and the turbine blade rotation plane.

When the radar points directly at the wind turbine and the blades are parallel to the radar beam (orientation angle 0°), the turbine appears as a nearly stationary distributed target, since the blade motion has minimal radial component. In contrast, when the orientation angle is 90° , the blade motion is primarily along the radial direction, producing significant Doppler velocities.

The results show that wind turbines can introduce broad contamination in the Doppler domain, particularly at 90° orientation. In STP, this contamination spreads widely due to sidelobe leakage. STAP reduces this sidelobe contamination; however, strong turbine returns can still dominate the temporal (Doppler) domain, especially when the turbine motion aligns with the radar line-of-sight.

Across all simulated cases (Figures 4.9–4.10), STAP consistently provides the greatest




Figure 4.10: STP and STAP results for simulated wind turbines at different orientation angles (0° , 30° , 60° , and 90°).

benefit when clutter exhibits both spatial and Doppler diversity, allowing joint filtering to suppress contamination more effectively than STP. The improvement is most evident in scenarios with strong sidelobe leakage or moving clutter, where STP is fundamentally limited by its fixed spatial response. However, the results also highlight key limitations of STAP, including its sensitivity to sample support, increased computational cost, and reduced effectiveness when clutter dominates the Doppler domain or closely overlaps with weather signals. These findings indicate that while STAP offers clear performance gains in complex environments, its practical application must balance these benefits against operational constraints.

4.3 Improved Doppler Estimation via Space-Time Signal Processing

Space-time processing, which jointly processes spatial and temporal domains, can be beneficial for Doppler velocity estimation. Doppler velocity estimation is important for weather radars because it provides insight into storm motion characteristics, which are essential for meteorological analysis. For instance, Doppler velocity estimates can be used to identify

wind rotation patterns within storms, indicating the presence of mesocyclones. In this chapter, clutter mitigation is the primary objective, with improved Doppler velocity estimation emerging as a direct consequence of more effective clutter suppression.

Doppler velocity can be estimated as the first moment of the Doppler spectrum using the following equation [73],

$$\hat{v} = \int_{-v_a}^{v_a} v S_n(v) dv, \quad (4.13)$$

where $S_n(v)$ is the normalized power spectral density resulting from either STP or STAP. In this section, Doppler velocity estimation is evaluated using simulated data. The simulation setup is described in detail in the following subsection.

4.3.1 Simulation Model and Parameters

The simulated data comprise 1024 samples (M), which is notably higher than what is typically used in practical weather radar operations due to the long dwell time required. However, a well-known challenge in STAP is the requirement for a large number of independent and identically distributed samples to enable robust covariance matrix inversion. To address this, diagonal loading, Doppler spectrum averaging, and range averaging are applied to improve covariance matrix conditioning and ensure numerical stability.

The transmit beam pattern used in the simulations is illustrated in Figure 4.11. It consists of a spoiled beam with a 6.5° beamwidth, designed for radar imaging. The simulated PAR operates at 3 GHz (S-band), and its configuration is a 1×40 uniform linear array with element spacing of $\lambda/2$, consistent with the horizontal dimension of the Horus array. This setup models scanning exclusively in the azimuth direction. Therefore, a 1D Horus radar configuration is adopted to reduce the computational burden associated with processing large volumes of data.

To investigate the impact of different digital PAR backend architectures on STP and




Figure 4.11: The normalized spoiled transmit and receive beam patterns are plotted on a dB scale, with the x -axis and y -axis representing azimuth and elevation dimensions in degrees.

STAP performance, five configurations are implemented: a fully digital system and non-overlapped subarray systems of 1×20 , 1×10 , 1×8 , and 1×5 (corresponding to 20, 10, 8, and 5 elements per subarray).

The weather and point target simulation setup is defined as follows. The weather spectrum is modeled using a Gaussian distribution with a spectrum width of 3 m/s and a 3D wind vector defined by $w = 20$ m/s, $u = 10$ m/s, and $v = 10$ m/s. Weather scatterers are simulated to be continuously distributed across the azimuth domain with uniform power.

A total of three point targets are simulated. The targets located at -1.82° with a velocity of -18 m/s and at 0.61° with a velocity of -19 m/s are positioned within the main lobe of the transmit beam, while a third target at 7.88° with a velocity of -22 m/s is located outside the main lobe. The spectra of the point targets are modeled using narrow Gaussian distributions with a spectrum width of 0.5 m/s. The weather signal is simulated with an SNR of 10 dB, and the point targets are simulated with a CSR of 5 dB.

4.3.2 Performance Evaluation of Doppler Estimation

The PSD results from STP and STAP using a fully digital PAR and a 1×8 subarray configuration are shown in Figure 4.12. The power levels are normalized to 0 dB and displayed

over a range from -30 dB to 0 dB. The top row shows results from the fully digital array, and the bottom row shows results from the subarray system. The left column presents STP outputs, while the right column presents STAP outputs.




Figure 4.12: Normalized powers using STP (left column) and STAP (right column) with two configurations: fully digital (top row) and 1×8 subarray digital (bottom row) arrays. The outputs are normalized and plotted on a dB scale.

For both architectures, STP exhibits significant sidelobe contamination in both Doppler and azimuth dimensions from point targets. This contamination leaks into the weather signals and degrades spectral moment estimation. In contrast, STAP mitigates this contamination by adaptively suppressing high-power signals from undesired directions. As a result, STAP provides improved target separation and characterization, whereas STP makes it more difficult to distinguish targets, particularly when multiple targets are closely spaced and have similar velocities.

The bottom-row figures highlight a key limitation of subarray systems, spatial aliasing

caused by grating lobes. This occurs when the effective spacing between receivers exceeds $\lambda/2$, leading to false target detections and significant errors in meteorological analysis. Since weather echoes are continuously distributed over a wide azimuth range, spatial aliasing can distort the observed structure and significantly degrade data accuracy. In addition, it becomes difficult to detect point-target clutter located outside the main lobe when using subarray systems with STP. This limitation reduces the effective alias-free field of view compared to the fully digital system.

However, subarray systems reduce processing time by forming beams at the subarray level, thereby reducing the amount of data to be processed. In this case, the subarray system reduces the data volume by approximately a factor of five compared to the fully digital system.

In Figure 4.13, 1D cuts of STP and STAP power are shown for the two configurations presented in Figure 4.12. The top plot shows the Doppler cut at 0.6° azimuth, and the bottom plot shows the azimuth cut at -20 m/s Doppler velocity.

The top plot shows that STAP, for both the fully digital configuration (red line) and the subarray configuration (purple dotted line), effectively suppresses most of the sidelobe contamination to approximately -22 dB around -20 m/s. In contrast, noticeable sidelobe contamination of approximately -13 dB is present in the STP results (blue line for fully digital and dotted yellow line for subarray). This contamination extends into the weather signal region, approximately from 0 m/s to 20 m/s, which can degrade the accuracy of weather analysis.

In the bottom plot, which shows the 1D azimuth cut at -20 m/s Doppler velocity, aliasing effects are visible in both STP and STAP outputs from the subarray system. Another important observation is that STAP is able to separate two point targets that are close in location and have similar velocities, indicating improved angular resolution compared to STP.




Figure 4.13: 1D cuts of STP and STAP powers with two configurations shown in Figure 4.12. The top plot shows the Doppler cut at 0.6° azimuth, and the bottom plot shows the azimuth cut at -20 m/s velocity.

To assess the feasibility of STP and STAP for weather surveillance, Doppler velocity estimation is performed using (4.13). Doppler velocities are estimated from both the raw outputs of STP and STAP and from their outputs after filtering point targets. The point targets are filtered by manually identifying their locations in each STP and STAP output and interpolating the gaps with noise.

The results are shown in Figure 4.14. The top plot shows the results from the fully digital PAR, and the bottom plot shows the results from the subarray system. The estimated velocities from STP and STAP without filtering are plotted as dotted green and yellow




Figure 4.14: Estimated Doppler velocity for fully digital (top) and subarray (bottom) systems.

lines, while the estimates with filtering are shown as solid blue and red lines. The true wind velocity is plotted as a black line for reference.

With a transmit beam of approximately 6.5° , both fully digital and subarray systems provide reasonably accurate velocity estimates after point-target filtering. However, STAP

consistently outperforms STP. This improvement arises because sidelobe contamination is more difficult to mitigate when using non-adaptive filter weights in STP. At the edges of the point targets (approximately -5° to -3° and 3° to 5°), STP shows significant deviations from the true values, whereas STAP retrieves velocities closer to the true values.

The subarray system exhibits larger errors due to contamination from grating lobes. It is also evident that clutter significantly affects Doppler velocity estimation in both STP and STAP results prior to filtering.

A quantitative evaluation of velocity estimation using STP and STAP is conducted by varying key parameters, including the number of samples, number of receivers, CSR, and SNR. For each test, the remaining parameters are held constant: a fully digital configuration ($N = 40$), 512 samples, 5 dB CSR, and 10 dB SNR.

The results are presented in Figure 4.15, which shows the root mean square error (RMSE) between the true and estimated velocities across the azimuth domain for STP (blue), STAP (red), STP with filtering (yellow), and STAP with filtering (purple). The top-left plot illustrates the effect of the number of samples on Doppler velocity estimation. Overall, the RMSE shows minimal variation across different sample sizes.

Top-left plot: increasing the number of independent samples enhances STP and STAP performance by establishing a more robust covariance matrix, resulting in stable covariance estimation and improved Doppler resolution. However, this enhancement was not directly reflected in Doppler velocity estimation, especially for unfiltered data. A marginal improvement was observed for filtered data, attributed to the finer resolution in filtering point targets. More significantly, the number of samples directly impacted STAP performance itself. The dynamic range of the STAP PSD expanded with increasing sample sizes, with an associated increase in computation time. Details regarding the dynamic range of STAP and the measured computation times are outlined in Table 4.1. The computations were performed using the Apple M1 chip processor. This chip features an ARM-based ar-




Figure 4.15: Quantitative RMSE analysis of Doppler velocity estimation.

chitecture with an 8-core CPU, comprising 4 high-performance cores and 4 high-efficiency cores. Furthermore, these metrics were measured using single threading. Although this table solely illustrates the correlation between computation time and the number of samples (M), it is important to note that the computation time of STAP is influenced by additional factors, including spatial (N) and range (L) samples. These factors contribute to an expo-

nential growth pattern rather than the linear growth portrayed in the table [75].

Table 4.1: The table illustrates the variation in dynamic range in STAP output based on the number of samples and the computational time required to compute STAP.

Number of samples (M)	Dynamic range (dB)	Computation time (s)
64	12.3	12.68
256	24.2	55.22
512	34.5	110.74
1024	52.8	222.57
2048	61.6	444.79

Bottom-left plot: five different PAR configurations were tested. Fully digital (total number of receivers $N = 40$), 1×2 subarray ($N = 20$), 1×4 subarray ($N = 10$), 1×5 subarray ($N = 8$), and 1×8 subarray ($N = 5$) are implemented. Aliasing was observed with the 1×4 subarray, 1×5 subarray, and 1×8 subarray systems as velocity was estimated from -20° to 20° in azimuth. The aliasing issues lead to significant challenges in achieving accurate Doppler velocity estimation. Consequently, the lowest RMSE was observed for the fully digital system and the highly digitized subarray system configured as 1×2 .

Top-right plot: the impact of CSR on Doppler velocity estimation is studied. Lower CSR levels exhibited nearly zero RMSE in all scenarios. This is attributed to the minimal power of point targets compared to weather signals, resulting in minimal contamination from point targets. Conversely, higher CSR levels presented challenges in accurately estimating Doppler velocity, leading to elevated RMSE even with filtering algorithms applied. Notably, at a CSR of 30 dB, both STP and STAP, with and without filtering, produced similar results. This indicates limited improvement in Doppler velocity estimation through filtering at high CSR levels, as contamination from point targets dominates the weather data.

Bottom-right plot: the impact of SNR on the filtering process is evaluated over 20 iterations (different seeds) with random noise power. Unlike the other factors, SNR with random

noise is tested across 20 iterations. For SNR values from -15 to 30 dB, performance varied significantly. At low SNR, all cases exhibited poor performance, with the mean RMSE close to 20 m/s and a high standard deviation of about 4 m/s. This degradation can be attributed to the noisy background, which hinders accurate estimation. At higher SNR levels, STAP outperformed STP without filtering. Furthermore, both STAP and STP achieved close to 0 m/s RMSE with the filtering algorithm. Additionally, an extremely small (close to zero) standard deviation was observed at high SNR values, indicating minimal impact of noise on the estimation process.

The results demonstrate that STAP effectively detects point targets and mitigates side-lobe contamination, enabling improved target localization, easier filtering, and more accurate Doppler velocity estimation. The analysis further shows that the number of samples and receivers has limited impact on estimation accuracy compared to CSR and SNR, which directly influence performance.

A key limitation of STAP is its requirement for a large number of samples (approximately 512 in this study), resulting in impractically long dwell times for operational systems. This can be partially mitigated through techniques such as diagonal loading and covariance averaging, although at the cost of reduced resolution.

These findings highlight the impact of STP and STAP on Doppler velocity estimation and weather analysis, as well as the importance of filtering and the trade-offs associated with different PAR configurations. In practice, the high computational burden and sample requirements of STAP remain significant challenges. Approaches such as sparse STAP [76] and reduced-dimension STAP [77, 78] have been proposed to address these issues and more will be discussed in Chapter 6. Overall, while STAP provides substantial performance gains in clutter suppression and Doppler estimation, its sample requirements and computational cost become the dominant constraints when transitioning from simulation to real radar operation.

Chapter 5

Case Study: Application to the Horus Phased Array Radar

In this chapter, adaptive beamforming and STAP will be applied to Horus measurements to validate the proposed calibration method and evaluate their performance in improving weather data quality. While simulations provide a controlled and ideal environment where all parameters are known, real measurements include uncertainties and system imperfections. Therefore, it is important to apply the proposed method to actual measurements to ensure its robustness and reliability for real-world applications.

First, the Horus radar will be introduced in terms of its specifications, operating frequency, architecture, and capabilities in the following section. Then, four cases are presented. These cases are designed to highlight different aspects of performance, including moving clutter (aircraft), ground clutter and polarimetric calibration, storm-edge behavior with increased degrees of freedom, and wind-turbine contamination with STAP.

The first case, from January 26, 2024, is one of the initial measurements from Horus. It was collected in single polarization with a non-overlapped subarray architecture consisting of 13 subarrays with 8×8 elements. The dataset contains aircraft flying over storms, providing a suitable case to investigate adaptive beamforming and STAP for moving clutter, particularly in scenarios where it overlaps with weather echoes.

The second case, from April 28, 2024, was measured in dual polarization with the same non-overlapped subarray architecture (13 subarrays with 8×8 elements). It includes strong

ground clutter and a high-reflectivity storm, including a supercell. This case demonstrates the advantages of adaptive beamforming, enables evaluation of the proposed calibration method, and allows analysis of the estimated polarimetric variables.

The third case, from June 5, 2025, was also collected in dual polarization but with smaller subarrays using a non-overlapped architecture of 26 subarrays with 4×8 elements. This configuration provides increased degrees of freedom for adaptive beamforming. The dataset focuses on RHI scans and includes ground clutter, while also showing how adaptive beamforming can still be effective near storm edges even in the absence of strong clutter.

The final case, from July 23, 2025, is a dual-polarization measurement of wind turbines, collected using 52 non-overlapped subarrays with 4×4 elements. Wind turbines represent one of the most challenging types of clutter due to their wide spatial contamination. Although their locations are known, contamination is unavoidable when weather passes through them. In this case, adaptive beamforming and STAP are applied to evaluate how effectively these techniques can mitigate contamination and improve weather data quality.

5.1 Overview of the Horus Radar System

The Horus radar is a fully digital, polarimetric phased array weather radar developed by the Advanced Radar Research Center at the University of Oklahoma for weather observation [27]. It operates in the S-band (2.7-3.1 GHz), which is commonly used for weather radar because it provides low attenuation and reliable scattering characteristics for precipitation. The system is built on a mobile, truck-based platform and uses a planar phased array composed of 1600 dual-polarized antenna elements arranged in 25 panels, where each panel contains an 8×8 grid of elements. Each element supports both horizontal and vertical polarizations, allowing the radar to measure key polarimetric variables such

as differential reflectivity, correlation coefficient, and differential phase, which are used for precipitation estimation and hydrometeor classification. For the case studies presented in this dissertation, 13-panel configurations with different subarray groupings such as 13 (Figure 5.1), 26 (Figure 5.14), 52 (Figure 5.18) subarray layouts are used, representing different active configurations of the same Horus platform tailored to specific experimental and computational constraints.

A defining feature of Horus is its fully digital architecture, where each antenna element has an independent transmit and receive chain and is digitized individually. Beamforming is performed digitally using FPGA-based processing rather than analog phase shifters. This design maximizes the flexibility in beam agility, including rapid electronic steering, formation of multiple beams, and adaptive beamforming techniques. The system supports high temporal resolution by enabling much faster scanning compared to traditional mechanically steered radars, making it possible to observe rapidly evolving weather phenomena. It also provides dense sampling in elevation, reducing gaps in vertical coverage that are common in conventional radar systems. The radar is designed as a software-defined system, meaning that waveform generation, beamforming weights, and scanning strategies are controlled through software. This allows reconfiguration for different observation modes or future requirements without hardware modifications.

Overall, the Horus radar integrates phased array technology, dual-polarization capability, and digital signal processing to support high-resolution and flexible weather observations.

5.2 Horus Case Study 1: January 26, 2024

A unique dataset was collected with Horus using the 13-panel configuration shown in Figure 5.1. The 13 panels correspond to subarrays consisting of 8×8 elements. Horus was




Figure 5.1: The truncated array configuration consisting of 13 panels, with each panel forming an 8×8 subarray.

operated at 3.070 GHz (S-band) using a spoiled transmit beam (shown in Figure 4.11), collecting 64 samples per dwell. This represents a strongly sample-limited regime for STAP, especially compared to the larger sample support typically required for stable covariance estimation discussed in Chapter 4. The radar conducted a 3D volume scan using DBF, generating narrow receive beams from -45° to 45° in azimuth and from 0.6250° to 19.2750° in elevation, with a sampling interval of 1.25° . A total of 64 spoiled transmit beams (16 in azimuth and 4 in elevation) were transmitted sequentially, first from bottom to top along the y -axis and then from left to right along the x -axis to cover the scanning sector. The range coverage spans approximately 5 km to 90 km with a range resolution of 38.4 m.

The data collected on January 26, 2024 at 15:57:14 UTC are shown in Figure 5.2. The right column displays reflectivity (Z in dBZ), where the top panel presents a front view at the range of maximum reflectivity and the bottom panel shows the PPI scan. This

dataset represents a complex environment consisting of weather echoes, ground clutter, and aircraft. Several aircraft are detected, and their tracks are indicated in the figure for reference. The left column provides zoomed-in views to highlight the aircraft signatures within the dataset.




Figure 5.2: Reflectivity data captured by the Horus system on January 26, 2024 at 15:57:14 UTC. The right column shows full-scene views: the top plot is a front view at the range of maximum reflectivity, and the bottom plot is a PPI scan. The dataset includes weather, clutter, and aircraft. The left column shows zoomed-in views highlighting aircraft targets.

For the STP and STAP assessments, a specific scan containing aircraft embedded within weather was selected. Figure 5.3 illustrates the scanning sector, including two aircraft and their flight paths, along with additional metadata (source: ADS-B Exchange). The location of Horus is marked by a black dot, and the scanning sector is delineated by dotted black lines. This dataset was collected later on the same day at 16:32:10 UTC and is visualized in Figure 5.4.

In Figure 5.4, the left panel shows the PPI of signal power (dB) at an elevation angle of 8.125° , while the right panel shows the RHI of power (dB) at an azimuth of 31.3° . An




Figure 5.3: Aircraft tracks detected by Horus on January 26, 2024 at 16:32:10 UTC (source: ADS-B Exchange). The scanning sector is indicated by dotted black lines, and aircraft information is shown in the left panels.




Figure 5.4: Zoomed-in PPI and RHI plots of power (dB) from the dataset collected on January 26, 2024 at 16:32:10 UTC. The PPI is shown at 8.125° elevation, and the RHI is shown at 31.3° azimuth to capture the aircraft highlighted in Figure 5.3.

aircraft (highlighted in yellow in Figure 5.3) is detected at approximately 24.41 km range, 31.3° azimuth, and 8.125° elevation, with power levels around 46 dB. Due to the Horus receive beamwidth of approximately 3° (see Figure 4.11), the high-power return from the aircraft spreads spatially from approximately 29° to 36° in azimuth and from 5° to 13° in elevation.

Firstly, DBF using both the non-adaptive (Fourier) and adaptive (Capon) approaches was applied to the data at a range of 24.4135 km, as shown in Figure 5.5. DBF was conducted with oversampling of 1.1° in both azimuth and elevation to achieve finer angular sampling and a better understanding of the scene. It is important to note that, due to the finer sampling rate employed in the DBF process, the results may exhibit slight variations compared to Figure 5.4, which used sampling intervals of 1.4° in azimuth and 1.2° in elevation.

All power outputs are normalized to a maximum of 0 dB for ease of comparison. High-power signals resembling point targets are observed from approximately 30° to 36° in azimuth and from 7° to 13° in elevation. Although these signals appear to represent multiple closely spaced point targets, the data in Figure 5.4 indicate a large point target extending into multiple transmit beams. The location of these high-power signals corresponds well with the aircraft location indicated in Figure 5.4.

Additionally, weaker power is detected at 42° azimuth and at 18° and 3° elevation. This weaker power is detected through sidelobes, and its shape matches the receive beam pattern (refer to Figure 4.11) with a significant sidelobe span of about 10° . The Capon output does not show these sidelobe signatures, as it effectively suppresses sidelobe contamination through adaptive changes in the beam pattern.

One of the transmit beams, spanning from 22.9031° to 27.1031° in azimuth and from 5.625° to 9.3750° in elevation, was selected for further analysis using STP and STAP. This transmit beam, delineated by a black box in Figure 5.5, was chosen because it contains




Figure 5.5: Normalized power after DBF with the Fourier method (left column) and Capon method (right column) are plotted in dB scale. The black box in each plot indicates the transmit beam used for STP and STAP in the later step.

weather data inside the main lobe and a point target (the aircraft) outside the main lobe. This selection aligns with the study's focus on enhancing weather data analysis by mitigating clutter contamination.




Figure 5.6: The STP and STAP results are plotted on a dB scale, with Doppler velocity (m/s) on the x -axis and azimuth (degrees) on the y -axis at 8.125° elevation. The power levels are normalized to 0 dB.

STP and STAP were implemented with spatial filtering applied in both elevation and azimuth, using sampling rates of 1.4° and 1.2° , respectively, to match the Horus system. Additionally, 10 range samples ($L = 10$) were averaged to estimate the space-time covariance matrix. This step was necessary to improve the performance of STAP, since it involves matrix inversion and requires a sufficient number of independent samples. However, averaging over 10 range samples reduces the range resolution from 38.4 m to 384 m, which

exceeds the desirable range resolution of 250 m for weather observations. This limitation could be mitigated by collecting more independent samples from either the spatial (N) or temporal (M) domain.

The outputs of STP and STAP are three-dimensional, including azimuth, elevation, and Doppler. For ease of interpretation, a 2D output consisting of azimuth and Doppler at 8.125° elevation is shown in Figure 5.6. In the STP and STAP results, two distinct point targets are identified at approximately 31.3° and 16.3° in azimuth. The aircraft is localized at 31.3° azimuth (see Figure 5.4 and Figure 5.5), while a false target is detected at 16.3° due to grating lobes.

Additionally, power detected within the main-lobe region (22.9031° to 27.1031° in azimuth) exhibits a broader spectrum width along the Doppler axis. These targets are widely distributed across azimuth and show aliasing effects outside the main lobe. Based on these characteristics, they are attributed to weather. The STAP output demonstrates improved point-target localization and enhanced data quality by mitigating sidelobe interference, enabling more accurate target characterization.

Figure 5.7 shows the output power map in azimuth and Doppler, where all elevation slices are stacked into a contour plot. The color scale represents power in dB. Figure 5.8 presents a similar visualization, but in the elevation-Doppler domain, where all azimuth slices are stacked. These representations provide a compact way to visualize the three-dimensional space-time information.

From these figures, a point target is clearly observed at approximately 11° elevation and 25° azimuth with a radial velocity of about 5 m/s. The weather echoes appear to be distributed at lower elevations ($< 5^\circ$), with radial velocities spanning a wide range from approximately -10 m/s to 3 m/s.

A critical limitation of STAP is also observed, namely a constrained dynamic range of power due to an insufficient number of independent samples; in this dataset, only 64 sam-




Figure 5.7: The output power map from space-time processing is presented above. The x -axis represents Doppler velocity, and the y -axis represents azimuth. The figure is shown as a contour plot, with all elevation slices stacked together. The color scale denotes power in dB.

ples were available. It is important to note that Horus is still under software development, and this dataset is not ideal for STAP.

The velocity characteristics of point targets and weather are illustrated in the STP and STAP power plots shown in Figure 5.6. The point target exhibits a velocity of approximately -15 m/s, which is markedly different from the mean weather velocity of around 5 m/s with a spectrum width of 2 m/s. However, the azimuth location of the point target overlaps with the weather region. Therefore, these velocity disparities introduce challenges for accurate wind velocity estimation of the weather component.




Figure 5.8: The output power map from space-time processing is presented above. The x -axis represents Doppler velocity, and the y -axis represents elevation. The figure is shown as a contour plot, with all azimuth slices stacked together. The color scale denotes power in dB.

The Doppler estimation improvement method, proposed in Section 4.3, is applied. The estimated velocities from STP, STAP, STP with filtering, and STAP with filtering are presented in Figure 5.9. Notably, the velocity estimates from STP and STAP without filtering exhibit significant discontinuities and abrupt transitions from 5 m/s to -15 m/s near the point-target locations. After manually filtering the point targets (used here as a proof-of-concept demonstration rather than an operational workflow) and re-estimating the velocity, the wind field is recovered, and the estimates become significantly smoother, resulting in a clearer representation of the weather structure.




Figure 5.9: Estimated velocity over the azimuth from STP (blue), STAP (red), STP with filtering (dotted yellow), and STAP with filtering (dotted purple) are shown above. The x -axis represents azimuth in degrees, and y -axis represents the estimated velocity in m/s.

Furthermore, STP exhibits a wider range of discontinuity regions (blue line), spanning approximately 14° to 21° and 28° to 36° , whereas the erroneous velocity regions from STAP (red line) are narrower, spanning about 14° to 19° and 28° to 34° . Despite the suboptimal STAP performance due to the limited number of samples, the velocity estimation using STAP after filtering shows clear improvement over STP.

5.3 Horus Case Study 2: April 28, 2024

Horus was deployed at the Lloyd Noble Center parking lot at the University of Oklahoma in Norman, Oklahoma ($35^\circ 11' 10.7''\text{N}$, $97^\circ 26' 33.7''\text{W}$). The data were collected using a truncated array configuration consisting of 13 panels, with each panel forming an 8×8 subarray, as illustrated in Figure 5.1. The 2D transmit beam was used, illustrated in Figure 4.11, and its beamwidth was approximately 10° in azimuth and elevation.

First, the measured beam patterns are presented in Figure 5.10. The left column shows the product of the transmit beam pattern and the subarray pattern. The middle column shows the array factor formed using the Fourier method (top row) and the Capon method (bottom row), and the right column shows the resulting beam pattern. These beams are steered to 0° azimuth and 15° elevation, plotted as black lines in each plot. The array factor plots clearly show grating lobes, since the centers of the subarrays are spaced farther apart than $\lambda/2$. The location of ground clutter is plotted in a red line in the array plot and beam pattern plot. Differences between the Fourier and Capon array factors are evident, as the Capon beamformer attempts to place a null at the clutter location. Therefore, this leads to differences in the resulting beam pattern shapes.




Figure 5.10: Measured beam patterns. The left column shows the product of the transmit beam pattern and the subarray pattern. The middle column shows the array factor formed using the Fourier method (top row) and the Capon method (bottom row), and the right column shows the resulting beam pattern. The beams are steered to 0° azimuth and 15° elevation.

To further examine the behavior of the Capon beam, a 1D azimuth cut of the beam pattern is shown in Figure 5.11. The gray line indicates the steering angle; in this case,

the beam is steered to -20° azimuth and 14.25° elevation. The dotted lines indicate the transmit beam boundaries, and the red line marks the location of stationary clutter. As shown in the figure, the Capon method places a deep null in the direction of the clutter, demonstrating its ability to adaptively suppress strong localized interference.




Figure 5.11: 1D azimuth cut of the Capon beam pattern. The gray line indicates the steering angle (-20° azimuth, 14.25° elevation). The dotted lines indicate the transmit-beam boundaries, and the red line marks the location of stationary clutter.

After applying the proposed calibration method, the reflectivity fields produced by the Fourier and Capon DBF methods are presented in Figure 5.12, with corresponding KTLX data shown for comparison. The Horus data were collected at 02:24:24 UTC, while the KTLX data were collected at 02:24:54 UTC. Due to differences in observation geometry, scan strategy, and timing between Horus and KTLX, the comparison should be interpreted qualitatively rather than as a point-by-point validation.

Several notable differences are observed between the two beamforming approaches. First, the overall storm structure retrieved using the Fourier DBF exhibits signs of contamination, with storm features appearing spatially compressed. This loss of spatial extent

is indicative of sidelobe contamination around the edge of the storm. In contrast, the reflectivity estimated using Capon beamformer shows substantially reduced contamination, particularly near the edges of the storm, resulting in a more coherent storm structure.

Another notable difference is the presence of clutter. A strong localized return, likely associated with a nearby building, is clearly visible in the Fourier-based reflectivity as a contamination artifact. In the Capon-based reflectivity, this clutter return is effectively suppressed and does not appear in the final product, as the adaptive beamformer places a null in the direction of the clutter and minimizes power from undesired directions. Note that the clutter appears to be completely removed. This occurs because the Capon beamformer does not sample in the direction of the clutter. Therefore, the clutter is fully suppressed.




Figure 5.12: Reflectivity fields produced by the Fourier and Capon DBF methods. The Horus data were collected at 02:24:24 UTC, while the KTLX data were collected at 02:24:54 UTC.

The velocity fields produced by the Fourier and Capon methods are presented in Figure 5.13. Velocity is a radar variable that does not require beam pattern calibration, however, it is still useful for illustrating the benefits of adaptive beamforming. Differences are




Figure 5.13: Velocity fields produced by the Fourier and Capon DBF methods.

observed near the boundaries of the storm, but the most notable distinction is the presence of clutter contamination in the Fourier-based velocity field. Specifically, near-zero velocity returns associated with stationary clutter are evident in the Fourier result, whereas this contamination is effectively suppressed in the Capon-based velocity field. As a result, the Capon method retrieves the velocity field with substantially reduced clutter contamination.

5.4 Horus Case Study 3: June 5, 2025

For another storm observation, Horus was deployed at the same location at the Lloyd Noble Center parking lot at the University of Oklahoma in Norman, Oklahoma ($35^{\circ}11'10.7''\text{N}$, $97^{\circ}26'33.7''\text{W}$). The data were collected using a truncated array configuration consisting of 26 subarrays, with each panel forming a 4×8 subarray, as illustrated in Figure 5.14. A vertically wide transmit beam (right plot in Figure 5.14) was used to achieve finer resolution in elevation. Two-dimensional imaging scanning was performed, with 20 receive beams formed in elevation and 2 receive beams formed in azimuth within each transmit beam. Horus was oriented toward the southwest, scanning elevations from 0.167° to 19.83° with 0.33° sampling (oversampling) and azimuths from -44.625° to 44.625° with 0.75° sampling (oversampling).




Figure 5.14: The truncated array configuration, consisting of 26 subarrays with each subarray containing 4×8 elements, is illustrated in the left plot. The transmit beam used during data collection is shown in the right plot.

First, the measured two-way beam patterns produced by the Fourier (left) and Capon DBF (right) methods are presented in Figure 5.15. The black lines indicate the steering directions; in this case, the beams are steered to -3.25° elevation and 0° azimuth. The Fourier two-way beam pattern is deterministic, exhibiting nearly identical shapes throughout the scan, except for mainlobe broadening as the beam steers away from broadside. In contrast, the Capon two-way beam pattern differs substantially, as it adapts to the scanning environment. Both the mainlobe and sidelobe structures vary with the scene while maintaining unity gain in the steering direction.

The beamformed power obtained using the Fourier (left plot) and Capon (right plot) methods is presented in Figure 5.16. The figure shows an RHI scan at 0.375° azimuth. There are three transmits in this scan. The overall storm structure measured by the Fourier and Capon methods is noticeably different, especially near the edges. Furthermore, at closer ranges, approximately from 1 km to 9 km, relatively strong contamination is observed, likely from stationary ground targets. This contamination starts at lower elevation angles but leaks up to approximately 20° elevation. In contrast, the Capon DBF shows




Figure 5.15: Measured two-way patterns using the Fourier (left) and Capon (right) methods. The black lines indicate the steering directions (-3.25° elevation, 0° azimuth).




Figure 5.16: RHI scan at 0.375° azimuth using Fourier (left) and Capon (right) DBF. Bottom plots are showing the zoomed-in view from the top plots.

clutter at lower elevation angles; however, it suppresses contamination at higher elevation angles as well as along the range. Although the Capon method is applied in the angular domain (along the azimuth and/or elevation dimension, depending on how DBF is performed), it can also improve data quality along the range dimension.

After DBF, the proposed beam-pattern calibration from Chapter 3 is applied to achieve polarimetric variables. The resulting polarimetric variables are presented in Figure 5.17.

KTLX 22:44:22 UTC
 EL 0.5° , Beam height 1 km
 Beam width : 1°

Horus 22:46:00 UTC
 EL 1.5° , Beam height: 1 km
 Beam width : 3°




Figure 5.17: NEXRAD (left), Horus Fourier DBF (middle), and Horus Capon DBF (right). Rows from top to bottom show reflectivity (Z), radial velocity (v_r), differential reflectivity (Z_{DR}), correlation coefficient (ρ_{hv}), and differential phase (ϕ_{DP}).

The left column shows the NEXRAD Level-II data, the middle column shows the Horus data processed with Fourier DBF, and the right column shows the Horus data processed with Capon DBF. The KTLX measurements were collected on June 5, 2025 at 22:44:22 UTC, while the Horus measurements were collected at 22:46:00 UTC, resulting in an approximately 2-minute time difference. In addition, the KTLX data correspond to a 0.5° elevation PPI scan, whereas the Horus data are from a 1.5° elevation scan. This elevation selection was chosen to approximately match the beam height at the storm location (1 km). It should also be noted that KTLX has a beamwidth of approximately 1° , whereas Horus has a broader beamwidth of approximately 3° .

The reflectivity fields show that the Fourier-based result exhibits noticeable loss near the edges of the storm. This loss is caused by sidelobe contamination as the beam scans across the storm boundary, where energy leaks from higher-power regions into the mainlobe. In contrast, the Capon-based reflectivity suppresses this contamination, allowing the storm structure to be retrieved more accurately. This improvement at storm edges is consistent with the gradient-based simulation results discussed in Chapters 3 and 6, where adaptive beamforming demonstrated clear advantages in resolving sharp spatial transitions.

A similar behavior is observed in the velocity field. Although direct comparison with KTLX is limited by differences in viewing geometry, the Fourier-based velocity field shows contamination near the storm edges, whereas the Capon result exhibits reduced contamination.

Differential reflectivity (Z_{DR}) is more challenging to estimate accurately due to its sensitivity to noise and calibration. The Capon-based Z_{DR} field appears noisier; however, its values are more consistent with the KTLX observations, and the overall spatial structure is comparable between the two methods.

The cross-correlation coefficient (ρ_{hv}) is the polarimetric variable most affected by Capon DBF, as discussed in Section 3.5. While good agreement is observed in the storm

core, the Capon-based ρ_{hv} decorrelates more rapidly near the storm edges.

Finally, differential phase (ϕ_{DP}), which does not require beam-pattern calibration, shows the most pronounced improvement with Capon DBF. The Capon-based ϕ_{DP} field is notably cleaner and less contaminated than the Fourier-based result, highlighting the advantage of adaptive beamforming for this variable.

5.5 Horus Case Study 4: July 23, 2025

The Horus radar was deployed at the Redbed Wind Farm located in Tuttle, Oklahoma ($35^{\circ}15'35''N$, $97^{\circ}49'07''W$) on July 23, 2025 at 15:03:49 UTC. The observations were conducted under clear-air conditions, with no other detectable targets within the scanning volume. The transmit beam and the PAR architecture used for this measurement are illustrated in Figure 5.18. This measurement employs a 2D transmit beam with a beamwidth of approximately 10° in both elevation and azimuth. The data were collected using 52 non-overlapped subarrays of 4×4 elements. These 52 subarrays provide more degrees of freedom compared to the previous subarray setup measurements.

The dataset collected from this deployment was processed using both the conventional Fourier beamforming method and the adaptive Capon method. The initial analysis focuses on comparing the receive beam patterns obtained from each approach. While the Fourier method produces a fixed beam pattern defined by the array geometry, the Capon method adaptively adjusts its beam pattern in response to the signal environment, particularly in the presence of wind turbine targets. Following this comparison, the adaptive Capon beams were calibrated to retrieve accurate target reflectivity values using the proposed method. The subsequent analysis qualitatively examines how the Capon approach mitigates contamination from wind turbine interference in both the Doppler and spatial domains.

Figure 5.19 shows the normalized two-way beam patterns (dB) for the Fourier method




Figure 5.18: Left: 2D normalized transmit beam (dB) used by Horus for data collection. Right: array configuration, where blue dots represent the 832 antenna elements, red dots indicate subarray centers used for DBF, and the red box illustrates the subarray layout.




Figure 5.19: Normalized two-way beam patterns (dB) at range 2.2 km: Fourier method (left) and Capon method (right). The black lines indicate the steering direction of the receive beam, and the red dot marks the angular location of wind turbine interference at azimuth 2.5° .

(left) and the Capon method (right) at a range of 2.2 km. Tapering has been applied to the Fourier method to reduce sidelobe levels. The black lines indicate the steering direction, and the red dot marks the location of wind turbine interference at azimuth 2.5° . The Fourier method produces a predetermined beam pattern based on the array configuration,

with fixed mainlobe and sidelobe characteristics. In contrast, the Capon method adaptively modifies the beam pattern by suppressing sidelobes in the direction of strong wind turbine interference and redistributing power to other directions. Since no other significant targets are present, the Capon method aggressively redistributes power, resulting in noticeable changes in the mainlobe shape.




Figure 5.20: DBF power (dB) from RHI scans at azimuth 2.5° . Top: Fourier method; Bottom: Capon method. The x -axis represents range (km).

The DBF power by the Fourier and Capon methods is shown in Figure 5.20. This figure presents an RHI scan at azimuth 2.5° , where the x -axis represents range in km. The top panel corresponds to the Fourier method, and the bottom panel corresponds to the Capon method. In the Fourier result, higher power is observed at higher elevation angles, which is attributed to sidelobe contamination. In contrast, the Capon method effectively suppresses this contamination. Furthermore, applying Capon beamforming at each range gate improves data quality along the range dimension. From the Fourier RHI scan, the wind turbine clutter appears to be located at 1.2 km, 1.7 km, 2.2 km, and 5 km. However, the Capon RHI scan indicates that the features at 1.2 km and 1.7 km are likely due to sidelobe

contamination, as their power is suppressed in the Capon scan.

After DBF, the proposed beam pattern calibration (see Chapter 3) has been applied. The calibrated reflectivity obtained using Fourier DBF (top row) and Capon DBF (bottom row) for the H channel (left column) and V channel (right column) is shown in Figure 5.21. This figure presents an RHI scan at an azimuth of 0.5° . It is noticeable that sidelobe contamination is minimized in the Capon DBF along both the elevation and range dimensions, whereas the Fourier DBF does not achieve the same level of suppression.

Figure 5.22 shows the calibrated power at azimuth 2.5° and elevation 8.5° as a function of range. Several strong returns are observed in both methods, corresponding to wind turbines at ranges of approximately 1.2 km, 1.51 km, 2.2 km, 3.3 km, 3.6 km, and 4.8 km. At ranges of 2.9 km, 4.1 km, and 5.3 km, elevated power levels observed in the Fourier method are likely due to sidelobe contamination, as these features are effectively suppressed in the Capon result, approaching the noise floor. Additionally, near turbine locations, the Fourier method shows a broader spread of high power, indicating contamination when the beam is steered around strong targets. In contrast, the Capon method suppresses these artifacts and provides more localized responses. In regions without targets, both methods exhibit consistent noise levels, indicating proper calibration.

As discussed previously, wind turbine clutter causes wide contamination in the Doppler




Figure 5.21: The RHI scans of calibrated Fourier and Capon DBFs for the H and V channels at an azimuth of 0.5° are presented above.




Figure 5.22: Calibrated power (dB) at azimuth 2.5° and elevation 8.5° as a function of range. Fourier DBF results are shown in blue, and Capon results are shown in red.

spectrum. Figure 5.23 shows the micro-Doppler features of the wind turbine from Fourier DBF and Capon DBF in time-Doppler plots. GMAP, a well-known filtering algorithm introduced in Chapter 3, was applied to both DBF methods. Both H and V channels are presented. These micro-Doppler plots are taken at the wind turbine location, with one at a low elevation angle (1.5°) and one at a higher elevation angle (6.5°); the azimuth angle is fixed at 0.5° . At the higher elevation (6.5°), the wind turbine is not present, so contamination is not expected.

At 1.5° elevation, strong contamination is observed in the Fourier DBF micro-Doppler map. A noticeable strong power (approximately 75 dB) appears at 0 m/s Doppler, representing the stationary body of the wind turbine. In contrast, the Capon DBF suppresses this contamination, and the 0 m/s component is reduced compared to the Fourier DBF. At the higher elevation angle (6.5°), the Fourier DBF still shows a wide range of contamination along the Doppler dimension, although it is weaker than at 1.5° . This indicates sidelobe contamination. In contrast, the Capon method effectively suppresses this contamination, with only a small residual power near 0 m/s Doppler. Furthermore, the plots show that GMAP removes the stationary clutter at 0 m/s but does not suppress other components, as discussed earlier, highlighting its limitation.

Lastly, STAP and STP have been applied to the data at a range of 3.7198 km and are

presented in Figure 5.24. The wind turbine clutter is estimated to be located at an azimuth of 0.5° . In Figure 5.24, the first two plots correspond to STP (top) and STAP (second) at an elevation of 1.5° and the last two plots at an azimuth of 0.5° , showing the Doppler and azimuth dimensions. The STP result is widely contaminated in the Doppler spectrum due to the rotation of the wind turbine blades and is also contaminated in the angular domain due to sidelobes and limited angular resolution (approximately 3°). In contrast, the STAP results show that these contamination are significantly reduced in both the Doppler and angular domains. In particular, the angular resolution is improved due to the adaptive beamforming weights used in STAP. As a result, it is clearer to distinguish the wind turbine, including the radial velocity of its blades and its location. Furthermore, if weather echoes overlap with the wind turbine, STP would cause significant contamination in both domains, whereas STAP shows strong potential to minimize interference and preserve the quality of the weather data. This case provides one of the clearest real-world demonstrations in this dissertation of the need for STAP, particularly in environments where clutter exhibits both spatial and Doppler variability that cannot be effectively addressed by Doppler-only methods such as GMAP.

In the following chapter, practical approaches for implementing adaptive processing (including adaptive beamforming and STAP) for operational use will be discussed. As discussed earlier, adaptive processing presents several challenges for application in polarimetric weather radars. One major challenge is the high computational cost, which limits real-time implementation. This issue will be examined in detail, along with several possible alternatives.




Figure 5.23: The micro-Doppler plots of wind turbines, with one at a low elevation angle (1.5°) and one at a higher elevation angle (6.5°), are presented above. The azimuth angle is fixed at 0.5°




Figure 5.24: Space-time processing, both STP and STAP, results are presented above. The x -axis represents the Doppler dimension and y -axis represents the angular dimension.

Chapter 6

Practical Adaptive Processing for Operational Use

Adaptive processing, including both adaptive beamforming and STAP, has several key limitations. The first is the calibration issue, which was discussed in Chapter 3. The second limitation is computation load. This computational burden primarily arises from the inversion of the covariance matrix. The computational complexity of matrix inversion is $\mathcal{O}(n^3)$, where n denotes the matrix dimension. As a result, the computation time increases cubically with matrix size, which severely limits practical real-time operation. For adaptive beamforming, the covariance matrix has a size of $N \times N$, where N is the number of array elements. For STAP, the covariance matrix size increases to $NM \times NM$, where M is the number of pulses. As the radar samples more densely in either the spatial or temporal domain to improve data quality, the computation time increases to a level that limits practical use.

Table 6.1 shows the measured computation time for DBF using the Fourier and Capon methods for the NEXRAD-based simulated case introduced in Chapter 3.7. The measurements were obtained using a MacBook Pro with an Apple M3 processor. The purpose of this comparison is to evaluate the relative computational performance of each method, rather than their absolute execution times.

The second column shows the measured time required to perform DBF for a single transmit beam, while the third column shows the time required to process an entire PPI

Table 6.1: Performance time comparison between Fourier and Capon processing

	Performance Time (1 TX)	Performance Time (Entire TX, 36×150)	Dwell Time (64 pulse, 1 ms, 36 TX)	Total Performance Time
Fourier	0.002 s	10.8 s	2.3 s	12.3 s
Capon	0.0617 s	333.18 s	2.3 s	335.48 s

scan. The PPI scan consists of 36 transmit beams along azimuth (covering 360°) and 150 range gates. For a full PPI scan, the Capon method is approximately 31 times slower than the Fourier DBF. The fourth column shows the dwell time, which is approximately 2.3 s and corresponds to 64 pulses with a pulse repetition time of 1 ms for 36 transmit beams. The total processing time is shown in the fifth column. Overall, the Fourier DBF requires approximately 12.3 s to process a full PPI scan, while the Capon DBF requires approximately 335.48 s, making the Capon method about 27.3 times slower.

In this chapter, the computational burden of adaptive beamforming and alternative approaches to mitigate it are examined. Several approaches can be considered to address this issue. First, adaptive beamforming can be applied only when needed. As discussed in Chapter 2, adaptive beamforming is most beneficial when the target power is not uniformly distributed within the beam.

In Section 6.1, we quantify a threshold for the conditional application of adaptive beamforming and examine how this strategy can reduce the computational burden of adaptive processing while still improving DBF performance compared to deterministic DBF. Second, the size of the covariance matrix can be reduced to accelerate matrix inversion. This can be achieved through array architectures such as sparse arrays or subarrays. These approaches reduce the dimensionality of the adaptive problem and therefore significantly lower computational cost. This topic is explored in Section 6.2. Lastly, the application of adaptive processing to generalize clutter filtering is explored in Section 6.4. This clutter filtering algorithm is designed to mitigate various types of clutter, including stationary ground clutter, moving point clutter, and complex clutter that exhibits wide distribution in

both spatial and temporal dimensions.

6.1 Conditional Application of Adaptive Beamforming

To apply adaptive beamforming only when it is needed, it is first necessary to define the conditions under which adaptive processing provides a clear benefit. Two types of targets are simulated: a sharp reflectivity gradient, representing a point target in real-world scenarios, and a smooth reflectivity gradient, representing the edge of a storm. For both cases, different gradient magnitudes are considered.

Figure 6.1 illustrates these two cases. The left plots show the modeled power (dB) for the two target types with different gradient magnitudes, plotted in different colors, ranging from 5 dB to 35 dB. Digital beamforming is then performed using both the Fourier and Capon methods. Beam-pattern calibration, as discussed in Chapter 3, is applied prior to estimating the radar variables.

The top row of plots shows the bias measured in each polarimetric variable (power, velocity, differential reflectivity, differential phase, and cross-correlation) as a function of gradient magnitude (y -axis) and elevation angle (x -axis) after applying the Fourier DBF. The bottom row shows the corresponding biases after applying the Capon DBF.

For the point target case, the results show that the Fourier DBF produces significant bias over an increasingly wide elevation range as the power gradient increases. This trend is clearly visible along the y -axis, where higher gradient magnitudes lead to broader angular regions of bias across all polarimetric variables. In contrast, for the Capon DBF, the angular extent of the bias does not increase with the power-gradient magnitude. The main exception is the cross-correlation variable, ρ_{hv} . As discussed in Section 3.5, bias in cross-correlation remains a challenging issue for adaptive beamforming.

For the distributed target case, a similar trend is observed. The angular region of bias




Figure 6.1: The modeled power (dB) of two cases, a point target and distributed targets, with different gradient amount (plotted in different colors in each plot) are presented in the left plots. The bias measured in each polarimetric variables after the Fourier DBF (top row) and the Capon DBF (bottom row).

after Fourier DBF becomes wider as the reflectivity gradient increases. The bias is more pronounced in the lower power regions, which are contaminated by sidelobes originating from adjacent high-power regions. For the Capon DBF, most polarimetric variables remain largely unaffected by the gradient magnitude, except for cross-correlation. As the gradient increases, the bias in ρ_{hv} becomes larger in the low-power region. This occurs because the adaptive beams become more aggressive in suppressing the high-power region, which alters the beam shape. This adaptive behavior was discussed in Chapter 2.

To further quantify these effects, the bias is averaged over elevation and evaluated as a function of power gradient (the x -axis) and SNR (the y -axis). The results are shown in Figure 6.2 for the point target case and Figure 6.3 for the distributed target case. Note that the color bars for differential reflectivity (0 dB to 0.05 dB), differential phase (0° to 5°), and cross-correlation (0 to 0.3) differ from those in Figure 6.1 to better highlight the results.

In Figures 6.2 and 6.3, the first row shows the bias after Fourier DBF, the second row shows the bias after Capon DBF, and the third row indicates which DBF method produces lower bias for each polarimetric variable. The first-row plots indicate that the bias from Fourier DBF is largely insensitive to SNR and is primarily driven by the power gradient, except for cross-correlation. When the gradient is large, meaning that the target power changes rapidly across angle, the Fourier DBF produces stronger sidelobe contamination and therefore larger bias. In contrast, cross-correlation is sensitive to SNR: when the SNR is low, ρ_{HV} becomes biased because the noise contribution increases and reduces the correlation between the H and V channel signals.




Figure 6.2: Bias measured in polarimetric variables after the Fourier DBF (first row) and the Capon DBF (second row) at different power gradient (the x -axis) and SNR (the y -axis) for a sharp power gradient case. The third row shows which DBF method produces less bias.

The second-row plots show that the Capon DBF is affected by both SNR and power gradient. Similar trends are observed for both the point-target case (Figure 6.2) and the distributed-target case (Figure 6.3). The cross-correlation estimation improves with increasing SNR. In particular, once the SNR exceeds approximately 15 dB, the bias in ρ_{HV}




Figure 6.3: Bias measured in polarimetric variables after the Fourier DBF (first row) and the Capon DBF (second row) at different power gradient (the x -axis) and SNR (the y -axis) for a smooth power gradient case. The third row shows which DBF method produces less bias.

after Capon DBF drops below 0.05. Velocity and differential phase are the two variables that are least affected by beam-pattern calibration, since they are phase-based variables. For these variables, the Capon DBF outperforms the Fourier DBF in most cases.

The third-row plots indicate which DBF method produces less bias for each polarimetric variable. Figure 6.2 shows that, for the point target case, the Capon DBF outperforms the Fourier DBF for all variables across all SNR and gradient conditions, except for cross-correlation. For cross-correlation estimation in the point target case, the Capon DBF outperforms the Fourier DBF when the SNR is higher than 20 dB and when the target power gradient exceeds 5 dB (Figure 6.2). In contrast, for the distributed target case (Figure 6.3), the Fourier DBF consistently outperforms the Capon DBF for cross-correlation estimation.

For velocity and differential phase, which are phase-driven variables, the Capon DBF outperforms the Fourier DBF in most scenarios. Power estimation is more nuanced. For the point target case (Figure 6.2), the Capon DBF begins to outperform the Fourier DBF

when the target power gradient exceeds approximately 10 dB. For the distributed target case (Figure 6.3), the Capon DBF begins to outperform when the power gradient exceeds approximately 20 dB. Based on these results, practical thresholds can be defined to determine when to apply the Capon DBF: From the Fourier DBF power, approximately 10 dB power gradient for point target scenarios and 20 dB for distributed target scenarios. To further refine this decision, the gradient slope (dB/deg) can be used as a classification metric, as shown in Figure 6.4. A threshold between 2 dB/deg and 3 dB/deg provides a practical criterion for determining when adaptive beamforming should be applied.




Figure 6.4: Bias map of a sharp gradient and smooth gradient case from Figure 6.2 and Figure 6.3 with gradient slope (dB/deg) on the x -axis.

6.1.1 Application of Conditional Adaptive Beamforming on NEXRAD-Based Simulations

This conditional application of adaptive beamforming, using gradient-slope thresholds of 2 dB/deg, 2.5 dB/deg, and 3 dB/deg, is applied to the NEXRAD case shown in Figure 3.18. Figure 6.5 presents the resulting DBF outputs from the Fourier DBF (left column), the combined Fourier-Capon DBF (middle column), and the Capon DBF (right column). In each plot, a mask indicates the regions where the power-gradient slope exceeds the specified threshold. The first row corresponds to a threshold of 2 dB/deg, the second row corre-

sponds to 2.5 dB/deg, and the third row corresponds to 3 dB/deg.




Figure 6.5: Reflectivity fields from the Fourier DBF, conditional application of Fourier and Capon DBF, and Capon DBF.

Given that the beam patterns are properly calibrated (Chapter 2), Figure 6.5 shows that the Fourier and Capon DBF can be used together without introducing discontinuities between range gates. The mask highlights the regions where the Capon DBF is most beneficial, which are primarily located near the edges of the storm. Even though reflectivity is higher within the core of the storm, the Capon DBF does not provide significant additional benefit in those regions, since the hydrometeors are widely distributed and the adaptive beamformer is limited in its ability to suppress energy in specific directions. Instead, the Capon DBF provides its greatest benefit near the edges of the storm. Figure 6.6 shows the

velocity fields for the three DBF cases, using the same layout as Figure 6.5. Near the center of the hook-echo region, the velocity fields produced by the Fourier and Capon DBF differ significantly.




Figure 6.6: Velocity fields from the Fourier DBF, conditional application of Fourier and Capon DBF, and Capon DBF.

Finally, we evaluate the computation-time savings achieved by the conditional application of adaptive beamforming. Table 6.2 presents the measured computation time when the Fourier DBF is used as the primary method and the Capon DBF is applied only when the gradient slope exceeds a specified threshold. Three threshold values are evaluated: 2 dB/deg, 2.5 dB/deg, and 3 dB/deg.

For these three cases, the total computation times for performing DBF over the entire

Table 6.2: Performance time for Fourier + Capon processing with different thresholds

Fourier + Capon	Performance Time (Entire TX, 36×150)	Dwell Time (64 pulse, 1 ms, 36 TX)	Total Performance Time
Threshold 2 dB/deg	38.52 s	2.3 s	40.82 s
Threshold 2.5 dB/deg	23.14 s	2.3 s	25.44 s
Threshold 3 dB/deg	15 s	2.3 s	17.3 s

PPI scan are 38.52 s, 23.14 s, and 15 s, respectively. Compared to the computation time of the Fourier DBF reported in Table 6.1, these cases are approximately 3.3, 2.14, and 1.4 times slower than the Fourier DBF alone. These results demonstrate the potential for implementing Capon DBF in real-time applications through conditional processing.

6.1.2 Application of Conditional Adaptive Beamforming on Horus Data

The conditional adaptive beamforming is applied to Horus data collected on June 5, 2025 (details in Section 5.4). The top-row plots in Figure 6.7 show RHI scans of the calibrated reflectivity obtained using the Fourier DBF (left) and the Capon DBF (right). Three transmit beams are used along the elevation dimension. Discontinuities between the transmit beams are observed in the Capon DBF results, which is a natural behavior of the adaptive beamformer. This effect can be mitigated by transmitting a wider spoil beam or by overlapping the transmit beams, however, this is beyond the scope of this study. The bottom-row plots show the storm traces, where the storm is detected based on reflectivity values exceeding 30 dBZ.

Figure 6.8 shows the storm traces from the Fourier DBF (black) and the Capon DBF (red) in the top plot. The bottom plot presents the magnitude of the differences between the two results (in dB). Noticeable differences appear mainly near the edges of the storm, where the Fourier DBF detects a wider storm area compared to the Capon DBF.




Figure 6.7: RHI scans of the calibrated reflectivity obtained using the Fourier DBF (left) and the Capon DBF (right), from Horus data collected on June 5, 2025 (details in Section 5.4).

Another noticeable region is at closer ranges, between approximately 5 km and 10 km, where ground clutter is present. In this region, the strongest differences and the widest areas of disagreement are observed. This behavior is attributed to strong ground clutter (e.g., buildings), which causes significant contamination through the fixed sidelobes of the Fourier DBF. In contrast, the Capon DBF can suppress this contamination by adaptively changing the beam shape. These regions of difference are primarily caused by sidelobe contamination from the Fourier DBF.

The variance of power from Fourier DBF (Figure 6.7) along the elevation dimension is calculated and shown in the left plot of Figure 6.9. Five additional elevation gates are padded at each end of the profile. A sliding window of five elevation gates is then used to compute the variance, after which the window is shifted by one gate and the process is repeated. A threshold of 5 dB is applied, and regions exceeding this threshold are indicated by black outlines. The Capon DBF is then applied to these regions, and the resulting output is shown in the right plot.

These results show that combining the two DBF methods is feasible when proper cali-




Figure 6.8: The storm traces from the Fourier DBF (black) and the Capon DBF (red) are plotted in the top plot and the magnitude of the differences between the two results (in dB) is plotted in the bottom plot.

bration is applied. The effectiveness of this hybrid strategy arises from selectively applying Capon DBF in regions where Fourier sidelobe contamination is most severe, such as storm edges and near strong reflectivity gradients, while retaining Fourier DBF elsewhere. The approach demonstrates the potential to reduce sidelobe contamination without performing the Capon DBF everywhere, making the overall scan more practical. However, there is still room for improvement in determining where the Capon DBF should be applied. For example, the window used to calculate the variance could be reduced to allow more sensitive region selection.




Figure 6.9: The left plot shows the variance of power, with traces indicating regions where the variance exceeds 5 dB. The right plot shows the output from the conditional Capon DBF.

6.2 Computationally Efficient Adaptive Processing with Reduced Covariance Matrix

As discussed at the beginning of this chapter, the computational burden can be mitigated by reducing the size of the covariance matrix. The matrix dimension can be directly reduced by performing DBF or STAP on fewer elements, which corresponds to reducing the spatial sampling in the adaptive filter. Possible approaches include sparse arrays and subarrays. Sparse array and subarray architectures were introduced in Section 2.3. In this section, the impact of sparse arrays and subarrays on adaptive beamforming and STAP is evaluated.

6.2.1 Adaptive Digital Beamforming on Sparse Arrays

The impact of sparse arrays is evaluated by estimating the bias in polarimetric variable estimation after DBF. The point target and distributed target cases (from Figure 3.5) are used, and the results are presented in Figure 6.10 (point target case) and Figure 6.11 (distributed target case). The bias in polarimetric variable estimation is shown after the Fourier DBF (top-row plots) and the Capon DBF (bottom-row plots), as a function of sparsity (the y -axis) and elevation (the x -axis). The sparse array was simulated based on the Horus array structure, where elements were randomly deactivated in each panel according to the specified sparsity percentage. The sparse array architecture was discussed in Chapter 2, see

Figure 2.4. In this study, sparsity levels ranging from 0% (fully populated array) to 95% were evaluated.

Both the Fourier and Capon methods are affected by the percentage of deactivated elements (i.e., sparsity level). However, the Capon method is more sensitive to sparsity, exhibiting significant bias when the sparsity exceeds 80%. It is especially noticeable that the Capon DBF introduces significant bias in velocity and differential phase estimation. These are the variables whose estimation is typically improved by adaptive beamforming, since they do not require beam pattern calibration (as discussed in Chapter 3).




Figure 6.10: Estimated bias in polarimetric variable estimation after the Fourier DBF (top-row) and the Capon DBF (bottom-row), as a function of sparsity (the y -axis) and elevation (the x -axis) for a point target case are presented.

Reducing the spatial sampling also limits the degrees of freedom of the adaptive beamforming in placing nulls. Therefore, the degradation in adaptive beamforming performance becomes more evident in the distributed target case (Figure 6.11) than in the point target case (Figure 6.10). This highlights that sparse configurations fundamentally constrain the adaptability of the beamformer, particularly in complex scenes where interference is spatially distributed.

The computation time required to perform the Capon DBF is measured for different sparsity levels. Figure 6.12 shows the computation time on the x -axis and the sparsity




Figure 6.11: Estimated bias in polarimetric variable estimation after the Fourier DBF (top-row) and the Capon DBF (bottom-row), as a function of sparsity (the y -axis) and elevation (the x -axis) for a distributed target case are presented.

percentage on the y -axis for the point target case. The computation time was measured using an Apple M3 chip. Ten iterations were performed to estimate the computation time; the statistical accuracy could be improved by increasing the number of iterations.

The results show that the computation time decreases from approximately 0.14 s to close to 0.01 s when the sparsity reaches 90%. For a more practical case that preserves DBF performance quality, a sparsity level of 40% reduces the computation time to about 0.065 s. This is nearly half of the time required when the array is fully populated. Such reduction can be critical for practical operation, where the radar transmits numerous beams and performs DBF over a wide area. However, this reduction is not achieved without cost: as sparsity increases, the degradation in adaptive beamforming performance becomes more pronounced, indicating that reduced covariance dimension represents a trade-off between computational efficiency and beamforming accuracy rather than a simple acceleration strategy.

6.2.2 Adaptive Processing on Subarray

Another way to reduce spatial sampling is to perform adaptive processing on a subarray system. The subarray system was introduced in Section 2.3.1. When DBF or space-time




Figure 6.12: The computation time required to perform the Capon DBF for different sparsity levels is presented.

processing is performed on a subarray system, grating lobes will appear depending on the subarray size. In contrast, sparse arrays use non-uniform element spacing, which prevents coherent grating lobe formation but redistributes energy into elevated sidelobes. According to the Nyquist sampling theorem, grating lobes (aliasing) occur when the centers of adjacent subarrays are spaced farther than $\lambda/2$.

Figure 6.13 shows the impact of subarray systems on space-time processing. Three different array architectures are evaluated: fully digital (left column), non-overlapped subarrays consisting of 8 elements (middle column), and non-overlapped subarrays consisting of 16 elements (right column). The first-row plots show the layout of each array architecture. The second-row plots show the power map after STP (non-adaptive space-time processing), and the third-row plots show the power map after STAP (adaptive space-time processing). These results show that the subarray system impacts both STP and STAP in a similar way, by limiting the observable region due to grating lobes. As the subarray size increases, the spacing between subarray centers becomes larger, causing grating lobes to appear more prominently and reducing the effective viewing region. Specifically, for the subarray consisting of 1×8 elements, grating lobes appear approximately every 15° . For




Figure 6.13: Fully digital (left column), non-overlapped subarrays consisting of 8 elements (middle column), and non-overlapped subarrays consisting of 16 elements (right column) are evaluated. The second-row plots show the power map after STP, and the third-row plots show the power map after STAP.

the subarray consisting of 1×16 elements, grating lobes appear approximately every 7° . Overall, moving from a fully digital array to subarray processing reduces computational cost and sample requirements, but at the expense of fewer degrees of freedom and reduced nulling flexibility in adaptive processing.

Figure 6.14 and Table 6.3 show the computation time for STAP as a function of the number of receivers (RX), measured using an Apple M3 chip. The plot shows an almost exponential growth in computation time as the number of RX increases. This growth is much faster compared to the adaptive beamforming computation time, which increases more linearly. This is because space-time processing uses joint filter responses, resulting




Figure 6.14: Measured computation time to perform STAP as a function of number of receivers (RX).

Table 6.3: Computation time of STAP for different numbers of RX channels

Number of RX	Covariance Matrix Size	Computation Time (s)
64	1024×1024	209.26
96	1536×1536	432.14
104	1664×1664	506.73
128	2048×2048	771.67
256	4096×4096	3017.84
1024	16384×16384	49342.75

in a covariance matrix that is much larger than that of DBF.

Table 6.3 shows how the size of the covariance matrix increases. For example, when Horus performs STAP at the element level, the covariance matrix size is 16384×16384 , and the measured computation time is 13 hours, 42 minutes, and 22 seconds. In contrast, when Horus performs STAP at the subarray level, specifically using a 1×8 system, the computation time is reduced to 3 minutes and 29 seconds, which is significantly more practical.

In the following section, we discuss how to mitigate grating lobes and high sidelobes by synthesizing the transmit beam. While adaptive processing can also suppress these effects, transmit beam synthesis provides a more computationally efficient approach. By shaping

the transmit pattern, the grating lobes and sidelobes in the two-way pattern can be reduced, thereby minimizing contamination.

6.3 High Sidelobes Mitigation with Beam Synthesis

A reflector antenna can only produce a fixed shape beam once the aperture size is chosen. In contrast, a highly digitized PAR can reconfigure its beams, which opens the door to synthesizing different beam patterns for different scanning modes. Digital PAR can synthesize beams directly from the array of elements by controlling the complex weights at each element. This capability can be particularly beneficial for imaging, where a wide transmit beam is used to illuminate a large sector and multiple narrow receive beams are formed within that sector.

When a wide “spoiled” transmit beam is used, the two-way pattern has higher sidelobes and a broader mainlobe compared to a narrow pencil beam from a dish-based radar [23]. As a result, the two-way pattern (transmit $\times$ receive) will suffer some loss in sensitivity and sidelobe performance if the transmit beam is simply obtained by uniformly broadening the aperture. Transmit beam synthesis aims to partially recover this loss by shaping the wide beam so that the mainlobe is as flat and efficient as possible while controlling the sidelobe structure and, when needed, placing additional nulls (for example, at grating-lobe directions).

Recent work has shown that this can be done using phase-only or hybrid (magnitude/phase) optimization methods that operate directly on the array elements and explicitly include the measured embedded element patterns in the cost function. For the fully digital Horus radar, pattern-synthesis frameworks have been developed that search for element weights which satisfy a prescribed pattern envelope while limiting gain loss and integrated “contamination” outside the desired region [79]. These methods can synthesize spoiled

beams and beams with multiple simultaneous transmit peaks in different directions. Simulations indicate that multi-peak transmit patterns can improve two-way sidelobe isolation and produce two-way patterns that look much closer to conventional pencil beams while improving the temporal resolution by scanning multiple sector [20, 80].

Because fully digital PARs for weather applications often operate in dual-polarization modes, beam synthesis must also consider the matching between horizontally and vertically polarized beams. Phase-only optimization methods have been proposed to synthesize independent H/V transmit beams that are intentionally broadened for imaging while keeping the H/V copolar patterns well matched in gain and shape [81]. These methods parameterize the phase distribution across the array using a reduced set of control points (e.g., Bézier surfaces) and use global optimization (typically genetic algorithms) to minimize metrics such as integrated contamination and efficiency loss [81, 82]. Similar ideas have been applied to the PAIR system, where phase-only synthesis is used to generate C-band imaging beams while accounting for practical element patterns and polarimetric constraints [82].

A key practical challenge is that optimizing thousands of element phases directly is computationally expensive. To reduce the search space, the phase across the aperture is parameterized with smooth basis functions (Bézier surfaces or Non-Uniform Rational B-Splines, NURBS), so that only a relatively small number of control points are optimized instead of one variable per element [83]. These parameterization schemes help enforce smooth phase transitions across the array, which is beneficial for sidelobe control, and significantly reduce the dimensionality of the optimization problem, making beam synthesis feasible for large PAR apertures used in weather applications [81–83].

From the perspective of imaging scanning, transmit beam synthesis gives us extra freedom to improve the two-way pattern. For example, when subarray architectures or sparse apertures introduce grating lobes in the receive pattern, a synthesized transmit beam can be shaped so that its sidelobes (or local nulls) reduce the two-way pattern in those grating-

lobe directions. Likewise, for wide imaging beams that inevitably raise sidelobe levels compared to a dish-based pencil beam, a well-designed synthesized pattern can still push some of the excess energy away from critical angles, reducing the worst sidelobes. The two-way pattern will not match the performance of a narrow pencil beam, but transmit beam synthesis allows PAR systems to move closer to that ideal while preserving the high temporal resolution benefits of imaging scanning.




Figure 6.15: The two-way patterns of the fully digital array (top row) and the 1×4 subarray system (bottom row) using two different transmit beam patterns: deterministic and optimized. Each plot contains the two-way patterns for multiple steering angles (plotted in different colors), from -10° to 10° with a 1° sampling rate. The dotted lines represent the mask that was used when the transmit beams were synthesized.

Figure 6.15 shows simulated two-way patterns of the fully digital Horus array (top row),

and the 1×4 subarray system (bottom row) using two different transmit beam patterns: deterministic and optimized. Each plot contains the two-way patterns for multiple steering angles (plotted in different colors), from -10° to 10° with a 1° sampling rate. The dotted lines represent the mask that was used when the transmit beams were synthesized.

The deterministic transmit beam represents a typical imaging-transmit design, where the mainlobe is wide and the power decreases gradually away from boresight. The optimized transmit beam, on the other hand, is designed to maintain the same mainlobe width but aggressively suppress the pattern right outside the mainlobe so that the high sidelobes and grating lobes can be reduced.

As shown in the figure, the fully digital array has about -30 dB sidelobes with the deterministic transmit beam (top-left plot), while the optimized transmit beam reduces them to about -45 dB (top-right plot). The same mechanism applies to the subarray system: with the deterministic transmit beam (bottom-left), the sidelobes are around -20 dB, and with the optimized transmit beam (bottom-right), they are reduced to about -35 dB. Although the optimized transmit beam successfully reduces sidelobe levels, some sensitivity loss is observed when the receive beams steer close to the edge of the transmit beam, for example in the region from ± 10 to ± 5 degrees.

6.4 Adaptive Clutter Filter with Adaptive Processing

Numerous studies have focused on differentiating and eliminating the stationary clutter from meteorological returns particularly through spectral processing, given that most existing polarimetric weather radars primarily utilize dish-based antenna systems. For instance, techniques like clutter phase alignment (CPA), which measures the phase variability of returned time-series signals, have been developed to identify stationary clutter returns [84]. Weather typically exhibits a CPA close to 0, whereas ground clutter displays a CPA much

larger than 0. This difference arises because weather induces phase shifts due to Doppler effects, while ground clutter generates minimal phase shifts in return signals, as it does not involve Doppler effects. Another classification method employing a Bayesian classifier in the spectral domain using single-polarized time-series signals has been developed in [85]. An algorithm that uses a fuzzy logic approach and Bayesian classifier with dual-polarization measurements has been developed to classify ground clutter, as discussed in [86]. The widely utilized ground clutter filtering method, Gaussian Model Adaptive Processing (GMAP), operates by filtering ground clutter in the spectral domain, as outlined in [68]. Another approach for filtering ground clutter involves using an adaptive filter based on the autocorrelation spectral density, known as Clutter Environment Analysis using Adaptive Processing (CLEAN-AP), [87]. However, despite their effectiveness in mitigating stationary clutter, the methods mentioned above often struggle to address moving clutter adequately, primarily because of the Doppler effects it generates.

In this section, an adaptive clutter mitigation algorithm that leverages information from DBF and space-time processing is explored to extend clutter classification beyond purely Doppler-based methods within the scope of the cases studied in this dissertation and filter out clutter. The algorithm exploits the joint spatial and spectral characteristics of both clutter and weather phenomena to achieve effective classification. This differentiation allows for more accurate distinction between clutter and weather signals. Adaptive beamforming and space-time processing are selectively applied when clutter is detected, thus optimizing data quality while minimizing computational load by performing adaptive processing only when necessary.

The flowchart of the proposed clutter mitigation algorithm is illustrated in Fig. 6.16. The main steps of the classification process are as follows: First, perform DBF using the Fourier method. Next, identify the high-power areas and compute the detected area size. Next, if the detected area is larger than the threshold, it is classified as weather, if smaller




Figure 6.16: The flowchart of the proposed classification algorithm is illustrated above.

than the threshold, it is classified as a point target. If no point target is detected ($PT = 0$), the process ends. However, if a point target is detected ($PT = 1$), adaptive beamforming using the Capon method is performed on the specific transmit beam containing the high power area. This is followed by another stage with a smaller threshold. If the area is larger than the threshold, the particular transmit beam is classified as containing a point target ($PT = 1$). The rationale for undergoing another stage with adaptive beamforming is to address potential false target detection in the Fourier DBF caused by grating lobes and/or sidelobes. Adaptive beamforming can mitigate high-power signals originating from these undesired directions. Subsequently, STAP is performed to obtain Doppler information at the selected transmit beam with $PT = 1$. The algorithm identifies the maximum power along the Doppler axis, assuming a positive Clutter-to-Signal Ratio (CSR), and measures its 3 dB width. If this width is narrower than a predetermined threshold, the algorithm considers this peak as indicative of a point target. Subsequently, 65% of the power in this

peak is filtered out along the spatial axis. Finally, Doppler velocity is estimated to evaluate the classification of clutter and the effectiveness of the filtering. The following subsections will provide detailed explanations of how to apply each step.

In this algorithm, a three-stage process is utilized, with each stage occurring after specific processing steps: Fourier DBF, adaptive beamforming, and spectral estimation. Each stage requires a distinct threshold size for effective operation.

The initial threshold (Th_1) following Fourier DBF is calculated based on the beamwidth of the transmit and receive beams, given by:

$$Th_1 = 4(\theta_{tx_3}\phi_{tx_3}) \left(\frac{\theta_{rx_3}\phi_{rx_3}}{\theta_{tx_3}\phi_{tx_3}} + 0.3 \right) \quad (deg^2), \quad (6.1)$$

where $\theta_{tx_3}\phi_{tx_3}$ represents the transmit beam width in elevation and azimuth, and $\theta_{rx_3}\phi_{rx_3}$ denotes the receive beam width in elevation and azimuth. The factor of 4 accommodates situations where the target is near the transmit boundary, potentially causing power leakage into neighboring transmit beams. The constant term of 0.3 is added to adjust for additional power detected through sidelobes.

The second threshold (Th_2) following adaptive beamforming is also based on beamwidth but without the constant term, as adaptive beamforming suppresses high side-lobe power and minimizes power leakage into neighboring transmit beams:

$$Th_2 = (\theta_{tx_3}\phi_{tx_3}) \left(\frac{\theta_{rx_3}\phi_{rx_3}}{\theta_{tx_3}\phi_{tx_3}} \right) \quad (deg^2). \quad (6.2)$$

The final threshold (Th_3) after the adaptive spectral estimation is designed to distinguish point targets from weather in the spectrum domain based on the measured 3 dB width. Since weather typically exhibits a wider 3 dB width ($> 2 \text{ m/s}$), this threshold is set to $Th_3 = 1.3 \text{ (m/s)}$. This value is empirically selected based on Horus measure-




Figure 6.17: Horus data from January 26, 2024: The power (dB) after the Fourier DBF for stationary ground clutter at 5.99 km is presented in the top plot. The second row shows plots of each transmit beam that is detected after Fourier DBF. The third row shows plots of the transmit beam after the Capon DBF.

ments and simulation tuning to provide a conservative separation between narrow-band point-target signatures and broader weather spectra.

This algorithm is applied to Horus data collected on January 26, 2024 (details in Section 5.2). Figure 6.17 shows data at a range of 5.99 km, where stationary ground clutter is present. The top-row plot shows the power map, azimuth in the x -axis and elevation in the y -axis, after the Fourier DBF. The second-row plots show the selected transmit beams, some correspond to the stationary ground clutter, while others are associated with sidelobe and grating-lobe contamination. Grating lobes appears due to the subarray system. This particular system has an element spacing of $d = 4\lambda$, where λ represents the wavelength. Since this spacing satisfies $d > \lambda/2$, it violates the Nyquist sampling criterion in the spatial domain and results in spatial aliasing in the form of grating lobes. The location of the




Figure 6.18: Horus data from January 26, 2024: The power (dB) after the Fourier DBF for moving clutter at 24.41 km is presented in the top plot. The second row shows plots of each transmit beam that is detected after Fourier DBF. The third row shows plots of the transmit beam after the Capon DBF.

grating lobes (θ_g) can be calculated as follows:,

$$\theta_g = \sin^{-1} \left(\sin(\theta_0) \pm \frac{n_g \lambda}{d} \right), \quad (6.3)$$

where θ_0 represents the beam pointing angle and n_g is the grating lobe order. In this specific data, the occurrence of grating lobes is expected at angles of $\pm n_g 14.48^\circ$. Adaptive beamforming is then applied to the selected transmit beams, as shown in the third-row plots. Through adaptive beamforming, the sidelobe and grating-lobe contamination are mitigated. As a result, the accurate location of the ground clutter can be identified, which corresponds to TX17 and TX49.

The second case is from a range of 24.41 km, where an aircraft was observed flying

over weather echoes. The power from the Fourier DBF is shown in Fig. 6.18 (top-row plot), using the same grid setup as in the previous case. In this scenario, a smaller area of power is observed with a higher maximum power level of approximately 15 dB. Since the point target is located near the boundary of a transmit beam, power leaks into neighboring transmit beams through sidelobes.

In addition, a few areas with slightly lower power levels (around 8 dB) are detected. These lower-power regions are false detections caused by the subarray system. As shown in the top-row plot of Fig. 6.18, two grating lobes are detected diagonally northeast of the point target location.

The second-row plots show the transmit beams selected by the algorithm. Adaptive beamforming is applied only to these selected beams. As a result, the sidelobe and grating-lobe contamination are mitigated, as seen in TX58, TX59, and TX64. Consequently, the accurate location of the clutter is identified at 7° elevation and 21.5° elevation in TX54.

Next, space-time processing has been applied to the selected transmit beam after adaptive beamforming, which is TX54. Space-time processing produces a 3D data structure, including elevation, azimuth, and Doppler dimensions. Figure 6.19 shows the power output maps from space-time processing at different cuts. The top plot shows the 2D map of elevation (the x -axis) and Doppler velocity (the y -axis) at the azimuth angle of highest power ($\phi = 29^\circ$). The second plot shows the 2D map of azimuth (the x -axis) and Doppler velocity (the y -axis) at the elevation angle of highest power ($\theta = 6^\circ$). The bottom plot shows the 2D map of azimuth (the x -axis) and elevation (the y -axis) at the Doppler velocity of highest power (-15 m/s).

In these plots, three target regions are observed at approximately -15 m/s, 8 m/s, and 30 m/s Doppler velocity. The first one, at -15 m/s, corresponds to the airplane, as it exhibits high power and a relatively narrow spectral width. Strong sidelobe contamination is also observed from this high-power target. The second region, at 8 m/s, has lower power and




Figure 6.19: Space-time processing output for the second case at a range of 24.41 km and TX54.




Figure 6.20: 1D plot of PSD at $\phi = 29^\circ$ and $\theta = 6^\circ$ is presented above. The blue line illustrates the spectrum before applying the filter, while the red dotted line shows the spectrum after the filter has been applied.

wider Doppler distributions, which are more consistent with weather-like targets. The third region, at approximately 31 m/s, also exhibits lower power but a relatively narrow spectral width. This characteristic suggests that it could be another airplane located farther away and detected through sidelobes. Note that there were several airplanes flying on this day.

These plots demonstrate the benefit of space-time processing. If only DBF were applied, these three targets would be averaged along the Doppler dimension (see the bottom plot in Figure 6.19). As a result, the aircraft return would dominate and contaminate the weather data. In contrast, space-time processing provides a joint view of spatial and spectral dimensions, allowing us to distinguish that multiple different targets are present in specific regions.

Figure 6.20 shows a 1D plot of the PSD at $\phi = 29^\circ$ and $\theta = 6^\circ$. The point target is identified by its narrow spectrum width ($< 1.3 \text{ m/s}$), and higher power. After the point targets are detected, they are filtered out and interpolated. Fig. 6.20 displays PSD before (blue line) and after (red dotted line) filtering. In this case, two targets are detected, exhibiting Doppler velocities of -15 m/s and 30 m/s. In contrast, the weather, characterized




Figure 6.21: Estimated Doppler velocity before (top plot) and after (bottom plot) the filtering is illustrated above. The x -axis and y -axis represent the azimuth and elevation in degrees.

by lower power and a wider spectrum width, shows Doppler velocities ranging from 6 m/s to 11 m/s.

The algorithm concludes by estimating the Doppler velocity across the entire scene, thereby evaluating its effectiveness in improving the quality of weather data. Top plot in Fig. 6.21 illustrates the Doppler velocity estimation without employing this algorithm, whereas the bottom plot in Fig.6.21 depicts the Doppler velocity estimation post-application of the algorithm. Without the algorithm, biased observations are noted due to the clutter. Conversely, the algorithm successfully retrieves Doppler velocity information for weather by mitigating clutter, demonstrating continuous connectivity along the spatial axes.

Overall, the proposed classification algorithm offers a comprehensive solution for improving weather radar data quality by accurately identifying and filtering clutter in both spatial and spectral domains. Its adaptive approach ensures efficient processing, making it a valuable tool for real-time weather observation and analysis. The successful implemen-

tation and validation with the Horus radar system pave the way for broader application in polarimetric radar systems, enhancing clutter mitigation and weather signal classification capabilities. These results demonstrate that adaptive processing is most effective when applied selectively, indicating that practical deployment requires careful control over where, when, and at what dimensionality adaptive methods are used.

This chapter explores several approaches to mitigate the computational burden of adaptive beamforming, including conditional adaptive beamforming, where it is applied only when needed based on the spatial gradient of the scanning environment, and performing adaptive processing with reduced covariance matrices through subarrays and sparse arrays. Through controlled simulations, the results and bias in estimated polarimetric variables show that adaptive processing is beneficial when the target power is not uniformly distributed within the beam. Furthermore, the results show that adaptive processing can be applied selectively to certain regions without performing it over the entire scan, provided that proper calibration is applied. However, these results also demonstrate that reducing dimensionality and computational cost is not without trade-offs, and that adaptive processing must be carefully balanced against performance degradation. Overall, this chapter highlights that adaptive beamforming is most practical when applied selectively and in a sample-aware, computationally efficient manner.

Chapter 7

Conclusions and Future Work

This dissertation investigates the application of adaptive processing techniques, including adaptive beamforming and STAP, to polarimetric PAR for weather observations. The work is motivated by the emergence of PAR as a next-generation weather radar technology, exemplified by the Horus system, and by the need to improve spatial and temporal data quality while maintaining accurate polarimetric estimates.

A key challenge addressed in this work is the degradation of data quality associated with imaging mode operation. While digital PAR enables rapid volumetric updates, the resulting increase in sidelobe levels in the two-way beam pattern introduces significant contamination, particularly in clutter environments. This dissertation demonstrates that adaptive processing methods can mitigate these effects, improving spatial quality and reducing sidelobe interference compared to conventional Fourier-based approaches. However, these improvements are achieved under controlled conditions and do not directly imply full operational readiness without additional constraints on computation and data availability.

The results show that adaptive beamforming methods, including Capon-based and constrained formulations, provide substantial improvements in beam pattern control. However, their effectiveness is strongly dependent on accurate calibration. A calibration framework was developed and evaluated, demonstrating that improper calibration can introduce significant bias in polarimetric variables. The proposed calibration approach improves the relia-

bility of dual-polarized measurements under varying signal conditions, including changes in SNR, CSR, and target distribution. At the same time, it is important to distinguish between calibration-related errors, which can be reduced through improved system characterization, and fundamental limitations of adaptive processing, such as H-V weight mismatch and phase inconsistency, which cannot be fully eliminated without constraining the adaptation.

Space-time processing further enhances performance by exploiting both spatial and temporal degrees of freedom. The analysis shows that STAP provides improved clutter suppression and target discrimination in scenarios involving moving targets, overlapping weather signals, and complex clutter such as wind turbines. These improvements were validated using both simulated data and real measurements from the Horus radar, confirming the practical applicability of adaptive processing in operational environments. However, as shown in Chapter 4, these gains depend strongly on the availability of sufficient samples, which are often limited in real radar dwell times.

A critical limitation of adaptive processing is its computational cost and reliance on sufficient i.i.d. samples. This dissertation addresses these constraints by exploring strategies such as conditional adaptive processing, covariance matrix reduction through subarrays and sparse arrays, and transmit beam synthesis for sidelobe and grating lobe control. These approaches improve the feasibility of real-time implementation while preserving performance gains, but they also introduce trade-offs between computational efficiency, degrees of freedom, and estimation accuracy. In particular, the results show that adaptive processing does not need to be applied uniformly across the entire scan to achieve meaningful improvements. Instead, selectively applying adaptive beamforming in regions with strong gradients or high contamination can retain most of the performance benefits while significantly reducing computational cost. This selective strategy highlights that adaptive processing is most effective when treated as a controlled, scene-dependent operation rather

than a uniformly applied technique.

7.1 Limitations and Open Issues

The cross-correlation coefficient (ρ_{hv}) is particularly sensitive to adaptive beamforming. In adaptive methods, the beamforming weights are complex-valued and vary across snapshots, introducing phase distortions between the H and V channels. These phase inconsistencies can reduce the magnitude of the cross-correlation coefficient, leading to bias in polarimetric estimates. This effect is inherent to adaptive beamforming and cannot be completely eliminated without constraining the adaptation, which may in turn reduce interference suppression capability. This represents a fundamental limitation of unconstrained adaptive polarimetric processing, rather than a calibration issue alone.

In addition, adaptive processing relies on accurate estimation of the covariance matrix, which assumes the availability of sufficient i.i.d. samples. In practical weather scenarios, especially in rapidly evolving or heterogeneous environments, this assumption may not be strictly satisfied, potentially degrading performance. This limitation is consistent with the findings in Chapter 4, where even moderately sized sample sets significantly impact STAP performance.

7.2 Suggestions for Future Research

Future work should focus on developing operationally viable adaptive processing strategies for PAR systems. Both adaptive beamforming and STAP require inversion of the sample covariance matrix, which in turn depends on the availability of sufficient independent samples to achieve a full-rank and well-conditioned estimate. In practice, this requirement is often difficult to satisfy under typical weather radar dwell constraints.

For the Horus system in its current configuration, the array consists of 13 panels with

832 elements, which can be exploited as explored in Chapter 6. While reducing the array dimension lowers the sample requirement for covariance estimation, it also reduces the available degrees of freedom, limiting the ability to place nulls and suppress interference. Specifically, the number of controllable nulls decreases from $N - 1$, directly impacting adaptive performance.

Given that typical dwell times for weather radar provide on the order of 32 samples, there is a fundamental trade-off between array dimension (degrees of freedom), sample support (covariance estimation quality), and computational complexity. This trade-off directly links the sample limitations identified in Chapter 4 with the computational constraints analyzed in Chapter 6, emphasizing that both factors must be considered jointly for practical system design.

Future work should therefore focus on developing adaptive configuration strategies that dynamically select array size, processing method (e.g., conventional beamforming, adaptive beamforming, or STAP), and covariance estimation techniques based on the available sample support and scene characteristics. This dynamic, sample-aware, and dimension-aware processing framework represents one of the most important directions for enabling operational adaptive processing in PAR systems. The flexibility of digital PAR systems makes such adaptive operational schemes feasible, but systematic frameworks for optimizing these trade-offs remain to be developed.

Additional future research directions include robust covariance estimation under limited or non-i.i.d. samples, reduced-rank and subspace-based adaptive processing methods, real-time implementation strategies for large-scale arrays, integration of adaptive processing with multi-radar or networked systems.

Overall, this dissertation demonstrates that adaptive beamforming and STAP clearly improve spatial resolution and clutter mitigation in polarimetric weather radar; however, routine operational use will require selective, sample-aware, and computationally efficient

implementations that explicitly balance performance with practical constraints.

References

- [1] A. V. Ryzhkov and D. S. Zrnić, *Radar Polarimetry for Weather Observations*. Springer, 2019, vol. 486.
- [2] M. R. Kumjian, “Principles and Applications of Dual-Polarization Weather Radar. Part I: Description of the Polarimetric Radar Variables.,” *Journal of Operational Meteorology*, vol. 1, 2013.
- [3] D. S. Zrnic and A. V. Ryzhkov, “Polarimetry for Weather Surveillance Radars,” *Bulletin of the American Meteorological Society*, vol. 80, no. 3, pp. 389–406, 1999.
- [4] S. Lim, V. Chandrasekar, V. N. Bringi, and M. Maki, “Precipitation Classification and Quantification Using X-Band Dual-Polarization Weather Radar: Application in Japan,” *Journal of Atmospheric and Oceanic Technology*, vol. 30, no. 9, pp. 2108–2120, 2013. DOI: 10.1175/JTECH-D-12-00123.1.
- [5] Y. Wang and V. Chandrasekar, “A Prototype Quantitative Precipitation Estimation Algorithm for Operational S-Band Polarimetric Radar,” *Journal of Hydrometeorology*, vol. 20, no. 5, pp. 961–978, 2019. DOI: 10.1175/JHM-D-18-0071.1.
- [6] A. Ryzhkov, D. Zrnic, and P. Zhang, “Polarimetric Radar Quantitative Precipitation Estimation,” *Remote Sensing*, vol. 14, no. 7, p. 1695, 2022. DOI: 10.3390/rs14071695.

- [7] E. Saltikoff et al., “An Overview of Using Weather Radar for Climatological Studies: Successes, Challenges, and Potential,” *Bulletin of the American Meteorological Society*, vol. 100, no. 9, pp. 1739–1752, 2019.
- [8] D. S. Zrnić, A. Ryzhkov, J. Straka, Y. Liu, and J. Vivekanandan, “Testing a procedure for automatic classification of hydrometeor types,” *Journal of Atmospheric and Oceanic Technology*, vol. 18, no. 6, pp. 892–913, 2001.
- [9] T. D. Crum, R. E. Saffle, and J. W. Wilson, “An Update on the NEXRAD Program and Future WSR-88D Support to Operations,” *Weather and Forecasting*, vol. 13, no. 2, pp. 253–262, 1998.
- [10] W. H. Heiss, D. L. McGrew, and D. Sirmans, “NEXRAD: next generation weather radar (WSR-88D),” *Microwave Journal*, vol. 33, no. 1, pp. 79–89, 1990.
- [11] D. S. Zrnić, R. J. Doviak, V. N. Bringi, A. V. Ryzhkov, J. C. Hubbert, and V. Chandrasekar, “Update on the NEXRAD Program: Polarimetric Upgrade,” *Bulletin of the American Meteorological Society*, vol. 87, no. 9, pp. 1213–1214, 2006. DOI: 10.1175/BAMS-87-9-1213.
- [12] P. L. Heinselman, D. L. Priegnitz, K. L. Manross, T. M. Smith, and R. W. Adams, “Rapid Sampling of Severe Storms by the National Weather Radar Testbed Phased Array Radar,” *Weather and Forecasting*, vol. 23, no. 5, pp. 808–824, 2008.
- [13] D. S. Zrnić, C. D. Curtis, R. J. Doviak, and T.-Y. Yu, “Phased Array Radar for Weather Observations,” *Proceedings of the IEEE*, vol. 103, no. 10, pp. 722–746, 2015. DOI: 10.1109/JPROC.2015.2444734.

- [14] H. B. Bluestein, M. M. French, I. PopStefanija, R. T. Bluth, and J. B. Knorr, “A Mobile, Phased-Array Doppler Radar for the Study of Severe Convective Storms: The MWR-05XP,” *Bulletin of the American Meteorological Society*, vol. 91, no. 5, pp. 579–600, 2010.
- [15] R. Carbone, M. Carpenter, and C. Burghart, “Doppler Radar Sampling Limitations in Convective Storms,” *Journal of Atmospheric and Oceanic Technology*, vol. 2, no. 3, pp. 357–361, 1985.
- [16] P. L. Heinselman, D. S. LaDue, and H. Lazrus, “Exploring Impacts of Rapid-scan Radar Data on NWS warning decisions,” *Weather and forecasting*, vol. 27, no. 4, pp. 1031–1044, 2012.
- [17] D. Zrnica et al., “Agile-Beam Phased Array Radar for Weather Observations,” *Bulletin of the American Meteorological Society*, vol. 88, no. 11, pp. 1753–1766, 2007.
- [18] T.-Y. Yu, D. S. Zrnica, D. Richard, and R. Doviak, “Beam Multiplexing Using the Phased-Array Weather Radar,” *Journal of Atmospheric and Oceanic Technology*, vol. 24, no. 4, pp. 616–626, 2007. DOI: 10.1175/JTECH1989.1.
- [19] E. Yoshikawa, T. Ushio, and H. Kikuchi, “A Study of Comb Beam Transmission on Phased Array Weather Radars,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 8, pp. 6346–6356, 2021. DOI: 10.1109/TGRS.2021.3067423.
- [20] D. Schwartzman, R. D. Palmer, M. Herndon, and M. Yeary, “Enhanced Weather Surveillance Capabilities with Multiple Simultaneous Transmit Beams,” *IEEE Transactions on Radar Systems*, 2025.

- [21] P. Kollias et al., “Science Applications of Phased Array Radars,” *Bulletin of the American Meteorological Society*, vol. 103, no. 10, E2370–E2390, 2022.
- [22] J. E. Stailey and K. D. Hondl, “Multifunction Phased Array Radar for Aircraft and Weather Surveillance,” *Proceedings of the IEEE*, vol. 104, no. 3, pp. 649–659, 2016.
- [23] D. Schwartzman, S. M. Torres, and T.-Y. Yu, “Distributed Beams: Concept of Operations for Polarimetric Rotating Phased Array Radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 11, pp. 9173–9191, 2021. DOI: 10.1109/TGRS.2020.3047090.
- [24] D. Schwartzman, S. M. Torres, and T.-Y. Yu, “Motion-Compensated Steering: Enhanced Azimuthal Resolution for Polarimetric Rotating Phased Array Radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 12, pp. 10 073–10 093, 2021. DOI: 10.1109/TGRS.2021.3055033.
- [25] J. L. Salazar et al., “An Ultra-Fast Scan C-band Polarimetric Atmospheric Imaging Radar (PAIR),” in *2019 IEEE International Symposium on Phased Array System & Technology (PAST)*, IEEE, 2019, pp. 1–5.
- [26] T. Ushio, Y. Wada, and S. Yoshida, “Recent Advances in Phased Array Weather Radar,” *IEICE Transactions on Electronics*, vol. E107.C, no. 10, pp. 274–278, Oct. 2024. DOI: 10.1587/transele.2024MMI0001.
- [27] R. D. Palmer et al., “Horus—A Fully Digital Polarimetric Phased Array Radar for Next-Generation Weather Observations,” *IEEE Transactions on Radar Systems*, vol. 1, pp. 96–117, 2023. DOI: 10.1109/TRS.2023.3280033.

- [28] R. D. Palmer, S. Gopalam, T.-Y. Yu, and S. Fukao, “Coherent Radar Imaging Using Capon’s Method,” *Radio Science*, vol. 33, no. 6, pp. 1585–1598, 1998. DOI: 10.1029/98RS02217.
- [29] T.-Y. Yu, R. D. Palmer, and D. L. Hysell, “A Simulation Study of Coherent Radar Imaging,” *Radio Science*, vol. 35, no. 5, pp. 1129–1141, 2000. DOI: 10.1029/1999RS002274.
- [30] J. Capon, “High-Resolution Frequency-Wavenumber Spectrum Analysis,” *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, 1969.
- [31] W. L. Melvin, “A Space Time Adaptive Processing (STAP) Overview,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 19–35, 2004.
- [32] R. Klemm, “Introduction to Space-Time Adaptive Processing,” 1998.
- [33] M. C. Wicks, M. Rangaswamy, R. Adve, and T. B. Hale, “Space-Time Adaptive Processing: A Knowledge-based Perspective for Airborne Radar,” *IEEE Signal Processing Magazine*, vol. 23, no. 1, pp. 51–65, 2006.
- [34] H. Wang, “An Overview of Space-Time Adaptive Processing for Airborne Radars,” in *Proceedings of International Radar Conference*, IEEE, 1996, pp. 789–794.
- [35] D. Schwartzman, “Mitigation of Wind Turbine Clutter With Digital Phased Array Radar,” *IEEE Access*, vol. 11, pp. 12 911–12 924, 2023. DOI: 10.1109/ACCESS.2023.3242910.

- [36] R. J. Vogt, J. R. Reed, and T. Crum, “Impacts of Wind Farms on WSR-88D Operations and Policy Considerations,” in *Proceedings of the 23rd Conference on Interactive Information Processing Systems (IIPS)*, San Antonio, Texas, USA, Jan. 2007.
- [37] B. Isom et al., “Detailed Observations of Wind Turbine Clutter With Scanning Weather Radars,” *Journal of Atmospheric and Oceanic Technology*, vol. 26, pp. 894–910, 2009. DOI: 10.1175/2008JTECHA1136.1.
- [38] M. Weber et al., “Towards the Next Generation Operational Meteorological Radar,” *Bulletin of the American Meteorological Society*, vol. 102, no. 7, E1357–E1383, Jul. 2021. DOI: 10.1175/BAMS-D-20-0067.1.
- [39] I. Ivić and D. Schwartzman, “A First Look at the ATD Data Corrections,” in *Proceedings of the 39th International Conference on Radar Meteorology*, 2019, pp. 2–6.
- [40] C. Fulton et al., “Mutual Coupling-Based Calibration for the Horus Digital Phased Array Radar,” in *Proceedings of the IEEE International Symposium on Phased Array Systems and Technology (PAST)*, Oct. 2022, pp. 1–6.
- [41] J. R. Guerci, *Space-time Adaptive Processing for Radar*. Artech House, 2014.
- [42] R. L. Haupt and Y. Rahmat-Samii, “Antenna Array Developments: A Perspective on the Past, Present and Future,” *IEEE Antennas and Propagation Magazine*, vol. 57, no. 1, pp. 86–96, 2015.

- [43] R. Palmer et al., “A Primer on Phased Array Radar Technology for the Atmospheric Sciences,” *Bulletin of the American Meteorological Society*, vol. 103, no. 10, E2391–E2416, 2022.
- [44] T. Maese et al., “Dual-use Shipborne Phased Array Radar Technology and Tactical Environmental Sensing,” in *Proceedings of the 2001 IEEE Radar Conference (Cat. No. 01CH37200)*, IEEE, 2001, pp. 7–12.
- [45] S. D. Robinson, “Utility of Tactical Environmental Processor (TEP) as a Doppler at-sea Weather Radar,” Ph.D. dissertation, Monterey, California. Naval Postgraduate School, 2002.
- [46] B. Isom et al., “The Atmospheric Imaging Radar: Simultaneous Volumetric Observations Using a Phased Array Weather Radar,” *Journal of Atmospheric and Oceanic Technology*, vol. 30, no. 4, pp. 655–675, 2013. DOI: 10.1175/JTECH-D-12-00063.1.
- [47] E. Yoshikawa et al., “MMSE Beam Forming on Fast-Scanning Phased Array Weather Radar,” *IEEE Transactions on Geoscience and remote sensing*, vol. 51, no. 5, pp. 3077–3088, 2012.
- [48] J. M. Kurdzo, D. J. Bodine, B. L. Cheong, and R. D. Palmer, “High-temporal resolution polarimetric x-band doppler radar observations of the 20 may 2013 moore, oklahoma, tornado,” *Monthly Weather Review*, vol. 143, no. 7, pp. 2711–2735, 2015.
- [49] H. Kikuchi et al., “Initial observations for precipitation cores with x-band dual polarized phased array weather radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 5, pp. 3657–3666, 2020.

- [50] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation theory*. John Wiley & Sons, 2002.
- [51] K. Chen, Z. He, and C. Han, "Design of 2-dimension sparse arrays using an improved genetic algorithm," in *Fourth IEEE Workshop on Sensor Array and Multichannel Processing, 2006.*, 2006, pp. 209–213. DOI: 10.1109/SAM.2006.1706123.
- [52] K. Chen, X. Yun, Z. He, and C. Han, "Synthesis of Sparse Planar Arrays Using Modified Real Genetic Algorithm," *IEEE Transactions on Antennas and Propagation*, vol. 55, no. 4, pp. 1067–1078, 2007.
- [53] X. Yang, Y. Sun, T. Zeng, T. Long, and T. K. Sarkar, "Fast step method based on past with sparse constraint for airborne phased array radar," *IEEE Transactions on Signal Processing*, vol. 64, no. 17, pp. 4550–4561, 2016. DOI: 10.1109/TSP.2016.2569471.
- [54] O. Frost, "An Algorithm for Linearly Constrained Adaptive Array Processing," *Proceedings of the IEEE*, vol. 60, no. 8, pp. 926–935, 1972. DOI: 10.1109/PROC.1972.8817.
- [55] C.-Y. Chen and P. P. Vaidyanathan, "Quadratically Constrained Beamforming Robust Against Direction-of-Arrival Mismatch," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 4139–4150, 2007. DOI: 10.1109/TSP.2007.894402.
- [56] K. Bell and H. Van Trees, "Adaptive and non-adaptive beampattern control using quadratic beampattern constraints," in *Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020)*, vol. 1, 1999, 486–490 vol.1. DOI: 10.1109/ACSSC.1999.832377.

- [57] Z. Tian, K. Bell, and H. Van Trees, “A Recursive Least Squares Implementation for LCMP Beamforming under Quadratic Constraint,” *IEEE Transactions on Signal Processing*, vol. 49, no. 6, pp. 1138–1145, 2001. DOI: 10.1109/78.923296.
- [58] G. Gordon and R. Tibshirani, “Karush-Kuhn-Tucker conditions,” *Optimization*, vol. 10, no. 725/36, p. 725, 2012.
- [59] F. Nai, J. M. Kurdzo, S. M. Torres, and C. D. Curtis, “A constrained capon adaptive beamforming technique for phased-array weather radars,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 63, pp. 1–13, 2025.
- [60] B. Carlson, “Covariance matrix estimation errors and diagonal loading in adaptive arrays,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 24, no. 4, pp. 397–401, 1988. DOI: 10.1109/7.7181.
- [61] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming, version 2.1*, <https://cvxr.com/cvx>, Mar. 2014.
- [62] H. Kikuchi et al., “Performance of Minimum Mean-Square Error Beam Forming for Polarimetric Phased Array Weather Radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 5, pp. 2757–2770, 2017. DOI: 10.1109/TGRS.2017.2653816.
- [63] Y.-S. Kim, D. Schwartzman, R. D. Palmer, and T.-Y. Yu, “Fast Adaptive Beamforming Using Deep Learning for Digital Phased Array Radars,” in *2022 IEEE International Symposium on Phased Array Systems & Technology (PAST)*, 2022, pp. 01–07. DOI: 10.1109/PAST49659.2022.9975038.

- [64] Y.-S. Kim, D. Schwartzman, T.-Y. Yu, and R. D. Palmer, “Fast adaptive beamforming for weather observations with convolutional neural networks,” *Remote Sensing*, vol. 15, no. 17, p. 4129, 2023.
- [65] K. Bell, Y. Ephraim, and H. van Trees, “Robust Adaptive Beamforming under Uncertainty in Source Direction-of-Arrival,” in *Proceedings of 8th Workshop on Statistical Signal and Array Processing*, 1996, pp. 546–549. DOI: 10.1109/SSAP.1996.534935.
- [66] Q. Yang, H. Zhang, and W. Li, “Performance Analysis of MVDR Beamforming Under Low SNR Conditions,” *International Journal of Antennas and Propagation*, vol. 2022, p. 7 984 443, 2022. DOI: 10.1155/2022/7984443.
- [67] J. Liu, Y. Tan, and M. Zhao, “Robust Adaptive Beamforming Using Diagonal Loading for Mismatched Signal Steering Vector,” *IET Radar, Sonar & Navigation*, vol. 11, no. 6, pp. 889–896, 2017. DOI: 10.1049/iet-rsn.2017.0245.
- [68] A. Siggia and R. Passarelli, “Gaussian Model Adaptive Processing (GMAP) for Improved Ground Clutter Cancellation and Moment Calculation,” in *Proc. ERAD*, vol. 2, 2004, pp. 421–424.
- [69] J. Ward, “Space-Time Adaptive Processing for Airborne Radar,” 1998.
- [70] D. Schwartzman, “Mitigation of Wind Turbine Clutter With Digital Phased Array Radar,” *IEEE Access*, vol. 11, pp. 12 911–12 924, 2023. DOI: 10.1109/ACCESS.2023.3242910.
- [71] J. B. Billingsley, *Low-angle Radar Land Clutter: Measurements and Empirical Models*. IET, 2002.

- [72] D. S. Zrnić, “Simulation of Weatherlike Doppler Spectra and Signals,” *Journal of Applied Meteorology*, vol. 14, no. 4, pp. 619–620, 1975. DOI: 10.1175/1520-0450(1975)014<0619:SOWDSA>2.0.CO;2.
- [73] R. J. Doviak and D. S. Zrnić, *Doppler Radar and Weather Observations*. Courier Corporation, 2006.
- [74] M. Fang, R. J. Doviak, and V. Melnikov, “Spectrum Width Measured by WSR-88D: Error Sources and Statistics of Various Weather Phenomena,” *Journal of Atmospheric and Oceanic Technology*, vol. 21, no. 6, pp. 888–904, 2004.
- [75] Y. Kim, D. Schwartzman, R. D. Palmer, T.-Y. Yu, F. Nai, and C. Curtis, “Exploiting All Digital Phased Array Radars for Clutter Mitigation with Space-Time Adaptive Processing,” in *40th Conference on Radar Meteorology*, AMS, 2023.
- [76] Z. Yang, X. Li, H. Wang, and W. Jiang, “On Clutter Sparsity Analysis in Space-Time Adaptive Processing Airborne Radar,” *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 5, pp. 1214–1218, 2013.
- [77] L. Xie, Z. He, J. Tong, and W. Zhang, “A Recursive Angle-Doppler Channel Selection Method for Reduced-Dimension Space-Time Adaptive Processing,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 5, pp. 3985–4000, 2020.
- [78] P. Huang, Z. Zou, X.-G. Xia, X. Liu, and G. Liao, “A Novel Dimension-Reduced Space-Time Adaptive Processing Algorithm for Spaceborne Multichannel Surveillance Radar Systems Based on Spatial-Temporal 2-D Sliding Window,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–21, 2022.

- [79] D. Schwartzman, R. D. Palmer, M. Herndon, and M. Yearly, “Pattern Synthesis and Digital Beamforming Capabilities of the Fully Digital Horus Radar,” in *2024 IEEE Radar Conference (RadarConf24)*, 2024, pp. 1–6. DOI: 10.1109/RadarConf2458775.2024.10548659.
- [80] D. Schwartzman et al., “Advanced Weather Surveillance Capabilities of the Fully Digital Horus Phased Array Radar,” *Bulletin of the American Meteorological Society*, BAMS-D-24-0296.1, 2025. DOI: 10.1175/BAMS-D-24-0296.1.
- [81] D. Schwartzman, R. D. Palmer, T.-Y. Yu, and F. Nai, “Phase-Only Antenna Pattern Synthesis for Polarimetric Phased Array Radar,” in *2024 IEEE International Symposium on Phased Array Systems and Technology (ARRAY)*, 2024, pp. 1–6. DOI: 10.1109/ARRAY58370.2024.10880311.
- [82] D. Schwartzman, J. D. D. Díaz, J. L. Salazar-Cerreño, T.-Y. Yu, R. D. Palmer, and M. S. McCord, “A Hybrid Antenna Pattern Synthesis Method for the Polarimetric Atmospheric Imaging Radar (PAIR),” in *2022 IEEE Radar Conference (RadarConf22)*, 2022, pp. 01–06. DOI: 10.1109/RadarConf2248738.2022.9764359.
- [83] R. J. Reinke, D. Schwartzman, F. Nai, T.-Y. Yu, J. Salazar-Cerreño, and R. D. Palmer, “Evaluation of a spline-based parameterization scheme for phase-only antenna pattern synthesis,” in *2022 IEEE International Symposium on Phased Array Systems & Technology (PAST)*, 2022, pp. 01–08. DOI: 10.1109/PAST49659.2022.9975004.
- [84] J. Hubbert, M. Dixon, S. Ellis, and G. Meymaris, “Weather Radar Ground Clutter. Part I: Identification, Modeling, and Simulation,” *Journal of Atmospheric and Oceanic Technology*, vol. 26, no. 7, pp. 1165–1180, 2009.

- [85] Y. Li, G. Zhang, R. J. Doviak, L. Lei, and Q. Cao, "A New Approach to Detect Ground Clutter Mixed with Weather Signals," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 4, pp. 2373–2387, 2012.
- [86] M. A. Rico-Ramirez and I. D. Cluckie, "Classification of Ground Clutter and Anomalous Propagation Using Dual-Polarization Weather Radar," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 7, pp. 1892–1904, 2008.
- [87] S. M. Torres and D. A. Warde, "Ground Clutter Mitigation for Weather Radars Using the Autocorrelation Spectral Density," *Journal of Atmospheric and Oceanic Technology*, vol. 31, no. 10, pp. 2049–2066, 2014.
- [88] G. Galati and G. Pavan, "Computer simulation of weather radar signals," *Simulation Practice and Theory*, vol. 3, pp. 17–44, 1995. DOI: 10.1016/0928-4869(95)00009-I.
- [89] R. J. Doviak and D. S. Zrnić, *Doppler Radar and Weather Observations*, 2nd. Academic Press, 1993.

List of Acronyms

AC Alternating Current

ATD Advanced Technology Demonstrator

CLEAN-AP Clutter Environment Analysis using Adaptive Processing

CPA Clutter Phase Alignment

CSR Clutter-to-Signal Ratio

CVX Disciplined Convex Programming

DBF Digital Beamforming

DC Direct Current

DL Diagonal Loading

DLAF Deep Learning Adaptive digital beamForming

DoA Direction of Arrival

dBZ Decibel Reflectivity

EM Electromagnetic

Fourier Fourier Beamforming

GMAP Gaussian Model Adaptive Processing

H Horizontal Polarization

HV Horizontal–Vertical

IQ In-phase and Quadrature

i.i.d. Independent and Identically Distributed

LCC Linearly Constrained Capon

ML Machine Learning

MMSE Minimum Mean-Square Error

MTI Moving Target Indication

MVDR Minimum Variance Distortionless Response

NEXRAD Next-Generation Weather Radar

NOAA National Oceanic and Atmospheric Administration

NSSL National Severe Storms Laboratory

NURBS Non-Uniform Rational B-Splines

NWS National Weather Service

NWRT National Weather Radar Testbed

PAIR Polarimetric Atmospheric Imaging Radar

PAR Phased Array Radar

PDF Probability Density Function

PPI Plan Position Indicator

PSD Power Spectral Density

PSL Peak Sidelobe Level

QCC Quadratically Constrained Capon

RCS Radar Cross Section

RHI Range–Height Indicator

RMSE Root Mean Square Error

RX Receive Channel

SINR Signal-to-Interference-plus-Noise Ratio

SNR Signal-to-Noise Ratio

STAP Space-Time Adaptive Processing

STP Space-Time Processing

TX Transmit Beam

ULA Uniform Linear Array

UTC Coordinated Universal Time

V Vertical Polarization

WTC Wind Turbine Clutter

WSR-88D Weather Surveillance Radar–1988 Doppler

List of Symbols

A_e Effective antenna aperture

$b(\mathbf{k})$ Beam pattern

C Radar constant

c Speed of light

C Constraint matrix

C_i Quadratic constraint matrix

C_0 Zero-order constraint matrix

C_1 First-order constraint matrix

C_2 Second-order constraint matrix

d Element spacing

D Particle diameter

D_i Position vector of the i^{th} array element

Δr Range resolution

$\Delta \theta$ Elevation resolution

$\Delta\phi$ Azimuth resolution

f Radar frequency

f_D Doppler frequency

g Antenna gain

g_t Transmit gain

g_r Receive gain

H Hermitian (conjugate transpose)

$I(r_0, r)$ Illumination factor

$\mathbf{I}$ Identity matrix

j Imaginary unit

k Wavenumber

$\mathbf{k}$ Steering wavevector

K_w Dielectric factor of water

l Loss factor

L Number of range bins

M Number of temporal samples (pulses)

μ Multiplier associated with quadratic constraints

N Number of receivers (antenna elements)

$N(D, r)$ Particle size distribution

NM Space-time dimension

n_g Grating lobe order

Ω_n Set of angular locations for n^{th} -order derivative constraints

$\propto$ Proportional to

P Signal power

$P_{ac}(v)$ Alternating-current component of clutter spectrum

$P_{GC}(v)$ Ground-clutter Doppler spectrum

P_{STAP} Output power from STAP

P_{STP} Output power from STP

P_r Received power

P_t Transmitted power

$q(\theta)$ Prior probability density function of DoA

r Range

r_0 Range gate center

$\mathbf{R}$ Covariance matrix

$\hat{\mathbf{R}}$ Estimated covariance matrix

$\mathbf{R}_L$ Diagonally loaded covariance matrix

$\mathbf{R}_n$ Noise covariance matrix

$\mathbf{R}_s$ Signal covariance matrix

$\mathbf{R}_{hh}$ H-channel covariance matrix

$\mathbf{R}_{hv}$ Cross-covariance matrix

$\mathbf{R}_{vv}$ V-channel covariance matrix

S Power density

S_b Backscattered power density

$S_n(v)$ Normalized Doppler power spectral density

$\mathbf{S}$ Space-time covariance matrix

$\hat{\mathbf{S}}$ Estimated sample covariance matrix

t Time

T Pulse repetition time

v_r Radial velocity

$V_n(m)$ Complex IQ sample at element n and time index m

$\mathbf{V}$ Received signal vector

w Wind-speed component (e.g., vertical wind)

$\mathbf{w}_c$ Capon beamforming weight vector

$\mathbf{w}_d$ Desired / steering weight vector

$\mathbf{w}_f$ Fourier beamforming weight vector

$\mathbf{w}_{LCC}$ Linearly constrained Capon weight vector

$\mathbf{w}_{MMSE}$ MMSE beamforming weight vector

$\mathbf{w}_o$ Optimized weight vector under null constraints

$\mathbf{w}_{o,0}$ Optimized weight vector with zero-order constraints

$\mathbf{w}_{o,1}$ Optimized weight vector with first-order constraints

$\mathbf{w}_{o,2}$ Optimized weight vector with second-order constraints

$\mathbf{w}_{\text{QCC}}$ Quadratically constrained Capon weight vector

$\mathbf{w}_{\text{Bayes}}$ Bayesian beamforming weight vector

$\mathbf{y}$ Space-time snapshot vector

Z Reflectivity factor

Z_{DR} Differential reflectivity

α Proportionality constant

γ_i Bound for the i^{th} quadratic constraint

$\delta(v)$ Dirac delta function

η Volume reflectivity

λ Wavelength

ϕ Azimuth angle

ϕ_{DP} Differential phase

ρ_{hv} Cross-correlation coefficient

σ_b Backscattering cross section

σ_n^2 Noise variance

σ_s^2 Signal variance

σ_v Spectrum width

θ Elevation angle

θ_b 3-dB beamwidth

θ_g Grating lobe angle

θ_{rx} Receive beamwidth

θ_{tx} Transmit beamwidth

τ Pulse duration

$\otimes$ Kronecker product

Appendix A

Simulation of Weather-Like Signals for Phased Array Radar

The overall simulation framework is illustrated in Fig. A.1. It consists of three main components: (1) generation of time-series IQ signals based on polarimetric variables, (2) expansion of these IQ signals into PAR signals by applying appropriate phase shifts to the signal from each receiver, and (3) DBF using both the Fourier and Capon methods. This appendix provides the complete implementation details of the simulation framework, while key configurations and results used in Chapters 3 and 4 are referenced in the main text.

A.1 Time-Series IQ Signal Simulation

Time-series IQ signals (V), with samples indexed by $i = 1, 2, \dots, M$, are simulated based on signal power (S_h), mean radial velocity (v_r), and spectral width (σ_v). These parameters correspond to the idealized weather and clutter scenarios introduced in Chapter 3 and Chapter 4, while the full signal-generation procedure is detailed here for completeness. The process begins by constructing a Gaussian-shaped power spectral density (S_m) for each Doppler velocity bin v_m as follows [72]:

$$S_m = \frac{S_h}{\sqrt{2\pi}\sigma_v} \exp \left[-\frac{(v_m - v_r)^2}{2\sigma_v^2} \right]. \quad (\text{A.1})$$




Figure A.1: The framework is shown above, consisting of three main components: time-series IQ signal simulation, PAR simulation, and digital beamforming. The red boxes represent inputs to each stage and the blue boxes represent outputs from each stage. The outputs from the digital beamforming process, using both Fourier and Capon methods, are evaluated and compared to the initial model values.

To simulate random signal fluctuations, a total power spectral value P_m is generated using a random variable X_m with a uniform distribution in (0, 1):

$$P_m = -(S_m) \ln(X_m), \quad (\text{A.2})$$

where X_m introduces amplitude variability. Noise is not included at this stage to isolate the signal-generation process; it is introduced later in the spoiled beam simulation, consistent with the modeling assumptions used in Chapter 3. The phase term θ_m is drawn from a uniform distribution over $[-\pi, \pi]$. The time-domain IQ signal is obtained via the inverse discrete Fourier transform of the frequency domain data:

$$V_h(i) = \frac{1}{M} \sum_{m=1}^M \sqrt{P_m} \exp(j\theta_m) \exp(-j\frac{2\pi}{M}mi). \quad (\text{A.3})$$

An independent sequence time series, I_y , is generated by repeating the process in Equation (A.3). The V channel signal is then simulated using the differential reflectivity (Z_{DR}),

the correlation coefficient (ρ_{hv}), and the differential phase shift (ϕ_{DP}) as follows [88]:

$$V_v(i) = \left(\rho_{\text{hv}} V_h + \sqrt{(1 - \rho_{\text{hv}}^2) I_y} \right) \frac{e^{j\phi_{\text{DP}}}}{\sqrt{Z_{\text{DR}}}} \quad (\text{A.4})$$

A.2 Spoiled Beam Simulation

The received time series signals at multiple array elements, indexed by $i = 1, 2, 3, \dots, N$, are synthesized by applying the desired phase shifts determined by the range between the transmitting element and the target range gate (r_{tx}) and the range from the target range gate to each receiving element (r_{rx}). This stage corresponds to the PAR signal formation step used in the simulation studies presented in Chapters 3 and 4.

Next, the spoiled beam signals $V_{\text{h,v}}^{\text{spoiled}}$ are simulated by (1) extracting the corresponding received signals within the desired angular sector, (2) multiplying with the transmit beam pattern (f_t^2), then (3) coherently summing them to form the composite received signal. This approach enables the simulation of spatially distributed echo signals and follows the methodology described in [64].

$$V_{\text{h}}^{\text{spoiled}}(i) = \sum_{q=1}^{nt} f_t(q) V_{\text{h},q}(i) \exp \left[-j \frac{2\pi}{\lambda} (r_{tx} + r_{rx}) \right] + n_T(i), \quad (\text{A.5})$$

where $f_t(q)$ represents the amplitude of the transmit beam pattern and $n_T(i)$ is additive noise introduced at each receiver. The summation spans nt spatial gates, 360 spatial gates (-90° to 90°) are used for the first simple simulation and 25 spatial gates ($\pm 5^\circ$ TX center angle) are used for the second NEXRAD simulation. This allows replicating the spatial integration performed in actual phased array radar systems. This process is applied itera-

tively to construct spoiled beams that cover the desired field of view and is repeated for the V channel. For the preliminary results, a Gaussian shaped transmit beam with a desirable beamwidth is used, along with the Horus radar array geometry.

Appendix B

Spectral Moment Estimation

Spectral moments and polarimetric variables can be estimated from the expected autocorrelation functions (R_h , R_v) and the cross-correlation function (C_{hv}) of the received H and V polarized signals, as follows [89]:

$$\hat{S}_{h,v} = \hat{R}_{h,v}(0) - \hat{N}_{h,v}, \quad (\text{B.1})$$

$$\hat{v}_r = \frac{v_a}{\pi} \arg \left[\hat{R}_h(T_s) \right], \quad (\text{B.2})$$

$$\hat{\sigma}_v = \frac{v_a \sqrt{2}}{\pi} \left| \ln \left[\frac{\hat{S}_h}{|\hat{R}_h(T_s)|} \right] \right|^{1/2}, \quad (\text{B.3})$$

$$\hat{Z}_{\text{DR}} = \frac{\hat{S}_h}{\hat{S}_v}, \quad (\text{B.4})$$

$$\rho_{hv} = \frac{|\hat{C}_{hv}(0)|}{\sqrt{\hat{S}_h \hat{S}_v}}, \quad (\text{B.5})$$

$$\phi_{\text{DP}} = \arg[\hat{C}_{hv}(0)], \quad (\text{B.6})$$

In these equations, $\hat{R}_{h,v}(0)$ are the estimated autocorrelations of the H and V channel

signals at a temporal lag of zero, and $\hat{N}_{h,v}$ denotes the estimated noise power for each channel. The variable $\hat{R}_h(T_s)$ is the autocorrelation of the H channel at a temporal lag of one, while $\hat{S}_{h,v}$ refers to the estimated signal power after subtracting noise. $\hat{C}_{hv}(0)$ is the zero-lag cross-correlation between the H and V channels. The term v_a denotes the Nyquist (aliasing) velocity, and T_s is the pulse repetition time (PRT). These equations provide the basis for deriving the spectral moments and polarimetric variables from IQ time-series data in polarimetric weather radar systems.

Appendix C

Properties of Covariance Matrices and Quadratic Forms

C.1 Hermitian Property of Covariance Matrices

Let the covariance matrix be defined as

$$\mathbf{R} = \mathbb{E}\{\mathbf{X}\mathbf{X}^H\}. \quad (\text{C.1})$$

Taking the Hermitian transpose,

$$\mathbf{R}^H = (\mathbb{E}\{\mathbf{X}\mathbf{X}^H\})^H = \mathbb{E}\{(\mathbf{X}\mathbf{X}^H)^H\}. \quad (\text{C.2})$$

Using $(\mathbf{AB})^H = \mathbf{B}^H \mathbf{A}^H$,

$$(\mathbf{X}\mathbf{X}^H)^H = \mathbf{X}\mathbf{X}^H, \quad (\text{C.3})$$

thus,

$$\mathbf{R}^H = \mathbf{R}. \quad (\text{C.4})$$

Therefore, covariance matrices are Hermitian.

C.2 Positive Definiteness

For any nonzero vector $\mathbf{x}$,

$$\mathbf{x}^H \mathbf{R} \mathbf{x} = \mathbb{E}\{|\mathbf{x}^H \mathbf{X}|^2\} \geq 0. \quad (\text{C.5})$$

If $\mathbf{X}$ is not degenerate, then

$$\mathbf{x}^H \mathbf{R} \mathbf{x} > 0, \quad (\text{C.6})$$

which implies that $\mathbf{R}$ is positive definite.

C.3 Inverse of a Positive Definite Matrix

If $\mathbf{R}$ is Hermitian positive definite, then its inverse $\mathbf{R}^{-1}$ also exists and is Hermitian positive definite. In particular,

$$(\mathbf{R}^{-1})^H = (\mathbf{R}^H)^{-1} = \mathbf{R}^{-1}. \quad (\text{C.7})$$

C.4 Quadratic Form

For any nonzero vector $\mathbf{a}$,

$$\mathbf{a}^H \mathbf{R}^{-1} \mathbf{a} > 0. \quad (\text{C.8})$$

Furthermore, since $\mathbf{R}^{-1}$ is Hermitian,

$$(\mathbf{a}^H \mathbf{R}^{-1} \mathbf{a})^* = \mathbf{a}^H \mathbf{R}^{-1} \mathbf{a}, \quad (\text{C.9})$$

which shows that the quadratic form is real-valued.