

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

EXPLORING THE CAPABILITIES OF THE
AGILE BEAM 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

CHRISTOPHER DAVID CURTIS

Norman, Oklahoma

2009

EXPLORING THE CAPABILITIES OF THE
AGILE BEAM PHASED ARRAY WEATHER RADAR

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF ENGINEERING

BY

Dr. Tian-You Yu, Chair

Dr. Murad Özaydın

Dr. Robert Palmer

Dr. Yan Zhang

Dr. Dušan Zrnić

Acknowledgements

I am grateful to my adviser, Dr. Tian-You Yu, for sticking with me over the last several years. He was always willing to discuss my progress and provided helpful guidance on my dissertation even on short notice. I also appreciate Dr. Dušan Zrnić who helped bring me to the National Severe Storms Laboratory (NSSL) and encouraged me to work toward a PhD from the very beginning. He has always shown an interest in my research and has been very supportive of my educational endeavors. I am fortunate to have Dr. Robert Palmer on my committee who advised Dr. Yu and had Dr. Zrnić as a member of his PhD committee. He encouraged me to look at things in new ways and provided helpful suggestions when I was starting my research. I have known Dr. Murad Özaydın since I was an undergraduate at The University of Oklahoma. He helped me with my mathematics research then, taught some of my classes then and now, and has been gracious enough to serve as my outside member. One of my original committee members, Dr. Victor DeBrunner, is no longer a part of my committee but played an important role in teaching me the fundamentals of signal processing. I am also grateful to Dr. Yan Zhang for taking Dr. DeBrunner's place when called upon.

My colleagues at NSSL have been extremely supportive throughout this journey. My supervisor, Allen Zahrai, has been understanding as I took classes and started writing my dissertation. He always let me put my family first while I worked and studied. My friend, Sebastian Torres, was always there when I needed a pep-talk and provided both technical and practical assistance during the process. Ric Adams helped me collect several hours of radar data and encouraged me during the tough times. I also appreciate Igor Ivić for taking the time to run simulations from his dissertation for me even when he was busy trying to finish his own PhD.

Last but certainly not least, I am thankful for the support of my family. My youngest son, Ben, was not even born when I started this PhD, but playing with him when I was at my busiest helped me keep going. My oldest daughter, Emily, brightened my long nights working at the office by sending me text messages when she was going to bed. Hannah, my middle child, never complained about how much I had to be gone but always let me know that she wished I was at home. My parents, Steven and Janet, instilled in me the importance of education. I even had some practice helping my dad with his dissertation when he was working on his PhD while I was growing up. There are too many other family members to thank individually, but I appreciated the calls and the encouragement. Finally, I am so thankful that I married such a wonderful woman who picked up the slack when I was away but always made it so I wanted to be at home. Terri had to sacrifice the most, but I always knew she would be there for me when I finished.

Table of Contents

1. Introduction.....	1
2. Spectral Characterization of Ground Clutter.....	8
2.1 Introduction	8
2.2 Collection of Clutter Data.....	13
2.3 Data Processing and Preliminary Model Fitting.....	16
2.4 Quadratic Ground Clutter Model: Spectral Fitting.....	25
2.5 Fitting PAR Data with the Quadratic Model	31
2.6 Clutter Characterization Using the Quadratic Model	39
2.7 Effects of Beam Broadening	46
2.8 Modeling the Effects of Wind Speed, Foliage, and Terrain	49
2.9 Completing the Quadratic Clutter Model.....	59
2.10 Using the Quadratic Clutter Model to Study Clutter Width.....	61
2.11 Simulating Ground Clutter with the Quadratic Clutter Model	69
2.12 The Effect of Spectrum Length on Clutter Shape	75
2.13 Conclusion.....	79
3. Staggered PRT Beam Multiplexing.....	83
3.1 Introduction	83
3.2 Issues with BMX Pairs	89
3.3 A New BMX Strategy.....	92
3.4 Comparisons to Conventional Collection Strategies	94
3.5 Simulation Results.....	98
3.6 Weather Data Results	117
3.7 DC Removal for SBMX Data	126

3.8	Conclusion.....	130
4.	Strategies for Weather Detection Using a Phased Array Radar	133
4.1	Introduction	133
4.2	Detecting Weather Signatures	134
4.3	Detection Scanning Strategies	146
4.4	Additional Considerations	157
4.5	Conclusion.....	160
5.	Conclusions and Future Work	163
6.	References	169
	Appendix A. Estimating the Spectral Floor and Removing the Effect of System Noise	173
	Appendix B. Mean Correction for Log Spectra.....	175
	Appendix C. Mean Squared Error (MSE) Threshold.....	176
	Appendix D. Sample Spectra and MSE Histograms	178
	Appendix E. The d Parameter: Normalization and Window Correction	183

List of Tables

Table 2.1. A list of experiments for clutter characterization.....	15
Table 2.2. Estimated quadratic model parameters and f values for 14 ground clutter cases...	40
Table 2.3. Normalized d_n and estimated c and d for all 14 collected cases.	51
Table 2.4. Constant values for quadratic clutter model.	58
Table 2.5. Width values for the quadratic and linear models for all 14 cases.....	66
Table 3.1. Simulation results for SBMX, staggered PRT, and contiguous pulse scanning strategies using a true spectrum width of 4 m s^{-1}	104
Table 3.2. Simulation results for SBMX, staggered PRT, and contiguous pulse scanning strategies using a true spectrum width of 2 m s^{-1}	107
Table 4.1. Detection scanning strategies, scan times, and number of beams.....	156

List of Figures

Figure 1.1. NWR T system block diagram after modifications by personnel from the National Severe Storms Laboratory (NSSL).....	2
Figure 2.1. Comparison of ground clutter spectra from a stationary antenna. Rotation was simulated using the electronic scanning capability of the phased array antenna.....	9
Figure 2.2. Billingsley windblown clutter model with $r = 10$ and $\beta = 5.7$	12
Figure 2.3. A satellite view of the area surrounding the NWR T with 10° regions that extend from approximately 5 to 40 km showing the “Urban”, “Wooded,” and “Prairie” terrain types at azimuths of 0° , 90° , and 225° , respectively.....	14
Figure 2.4. Example of ground clutter spectrum from the “Full Foliage, Urban, Windy” case with a simple linear fit of the AC portion of the spectrum.....	18
Figure 2.5. Comparison of unaveraged periodogram of length 4096 and average of periodograms from seven overlapping sequences of length 1024.....	20
Figure 2.6. Comparison of unaveraged periodogram of length 4096 and average of periodograms from four interleaved sequences of length 1024.....	21
Figure 2.7. Example of ground clutter spectrum from the “Light Foliage, Urban, Breezy” case with a simple linear fit of the AC portion of the spectrum.	22
Figure 2.8. Two clutter spectra examples with both a linear and quadratic fit of the AC portion of the spectra.	24
Figure 2.9. Example spectral fitting applied to a spectrum from Figure 2.8 showing the quadratic fit, the spectral floor, and c and d parameters.....	26
Figure 2.10. Spectral fitting procedure using spectrum from Figure 2.9.	30
Figure 2.11. Histograms comparing quadratic and linear fits of the AC part of the spectra for the 1397 spectra where the fit converged and c was greater than 45 dB.....	32
Figure 2.12. Histogram of b parameter values for “Full Foliage, Urban, Breezy” case.	34
Figure 2.13. Histograms of a , b , c , and d_w parameter values for “Full Foliage, Urban, Breezy” case.	35
Figure 2.14. Nearly linear relationships between model parameters.	36
Figure 2.15. Histogram of $(c + d)$ showing the relationship between the spectral peak and spectral floor.....	38
Figure 2.16. Ground clutter spectrum from the “Full Foliage, Wooded, Windy” case with the adjusted spectral floor and 15 dB level.....	39
Figure 2.17. Sample ground clutter spectra for the “Light Foliage, Prairie” conditions. (Estimated model parameters: Light - $\hat{a} = 584.0$ dB/(m ² s ⁻²), $\hat{b} = 240.6$ dB/(m s ⁻¹), $\hat{c} = 61.4$ dB, $\hat{d} = 20.4$ dB, Breezy - $\hat{a} = 203.9$ dB/(m ² s ⁻²), $\hat{b} = 152.7$ dB/(m s ⁻¹), $\hat{c} = 71.0$ dB, $\hat{d} = 8.9$ dB, Windy - $\hat{a} = 14.8$ dB/(m ² s ⁻²), $\hat{b} = 47.7$ dB/(m s ⁻¹), $\hat{c} = 68.4$ dB, $\hat{d} = 6.8$ dB, Gale - $\hat{a} = 4.3$ dB/(m ² s ⁻²), $\hat{b} = 25.0$ dB/(m s ⁻¹), $\hat{c} = 65.9$ dB, $\hat{d} = 5.5$ dB.)...	43

Figure 2.18. Histograms of the MSE for both the quadratic and linear fits for the “Light Foliage, Prairie” conditions.....	44
Figure 2.19. Representative clutter shapes for all 14 cases showing both the AC and DC parts of the spectral shapes.....	45
Figure 2.20. Number of spectra used for modeling from all six cases plotted versus the beamwidth at the center of the shared 5 ° sector.....	48
Figure 2.21. Dependence of quadratic clutter model parameters on beamwidth.....	49
Figure 2.22. Relationship between d_n and wind speed, s	52
Figure 2.23. Comparison of quadratic and linear models for the “Full Foliage, Breezy, Urban” case.....	54
Figure 2.24. Relationship between f and wind speed, s	55
Figure 2.25. Relationship between the two slope parameters b and f	57
Figure 2.26. Example of linear extension of the quadratic clutter model.....	60
Figure 2.27. Comparison of the spectrum widths for both the representative spectra and the quadratic and linear clutter models for all four foliage-terrain pairs.....	64
Figure 2.28. Comparison of w for both collected data and quadratic and linear clutter models for all four foliage-terrain pairs.....	67
Figure 2.29. Comparison of w versus spectrum width for the Gaussian, quadratic, and linear clutter models. The quadratic and linear model parameters are based on the “Full Foliage, Urban, Prairie” conditions.....	69
Figure 2.30. Comparison of quadratic clutter models based on simulated parameters ($a = 37.2 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $b = 61.3 \text{ dB}/(\text{m s}^{-1})$, $c = 61.9 \text{ dB}$, and $d_n = 3.3 \text{ dB}$) and estimated parameters ($\hat{a} = 29.8 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 56.5 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 63.8 \text{ dB}$, and $\hat{d}_n = 3.3 \text{ dB}$) for the modified simulation.....	72
Figure 2.31. Comparison between two simulations which use discrete values from the clutter model to determine the power for the spectral coefficients. The red curve uses the first 64 samples from a simulation with 4096 values, and the blue curve uses all of the samples from a simulation with 64 values.....	73
Figure 2.32. Comparison between two simulations which use integration of the clutter model to determine the power in the Doppler bins. The red curve uses the first 64 samples from a simulation with 4096 values, and the blue curve uses all of the samples from a simulation with 64 values.....	74
Figure 2.33. Averaged spectra from simulations showing the effects of a rectangular window on time series data of differing lengths.....	76
Figure 2.34. Averaged spectra from simulations showing the effects of a Blackman window on time series data of differing lengths.....	77
Figure 3.1. A comparison between contiguous pulse data collection and BMX pairs.....	84
Figure 3.2. Correlation in time for a spectrum width of 1 m s^{-1} assuming a Gaussian spectrum.....	85

Figure 3.3. A simple azimuth-only beam multiplexing pattern.	87
Figure 3.4. A comparison of standard deviation of velocity for contiguous pulses and BMX pairs.	91
Figure 3.5. A simplified example showing first and second trip echoes. The solid rectangles depict two transmitted pulses, and the triangles are first trip echoes from each pulse. The red rectangle is a second trip echo from the first (red) pulse.	91
Figure 3.6. A scanning strategy that uses both staggered PRT's and BMX.....	93
Figure 3.7. Simulating contiguous pulse, staggered PRT, and SBMX time series data using uniform PRT samples.....	97
Figure 3.8. Biases for signal power, velocity, and spectrum width versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 4 m s^{-1} and a velocity of 15.25 m s^{-1}	100
Figure 3.9. Standard deviations of reflectivity, velocity, and spectrum width and successfully dealiased velocity percentages versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 4 m s^{-1} . Acquisition time for the contiguous pulse strategy is the sum of the long PRT acquisition time to needed to meet the reflectivity requirement and the short PRT acquisition time need to meet the velocity and spectrum width requirements.....	102
Figure 3.10. Standard deviations of reflectivity, velocity, and spectrum width and percentages of successfully dealiased velocities versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 2 m s^{-1} . Acquisition time for the contiguous pulse strategy is the sum of the long PRT acquisition time to needed to meet the reflectivity requirement and the short PRT acquisition time need to meet the velocity and spectrum width requirements.	106
Figure 3.11. Standard deviations of reflectivity, velocity, and spectrum width and percentages of successfully dealiased velocities versus true spectrum width for the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.	109
Figure 3.12. Estimated spectrum width plotted versus true spectrum width for the contiguous pulse and staggered strategies.....	111
Figure 3.13. Plots of reflectivity and spectrum width standard deviations for an array of true spectrum width and SNR values. The scanning strategies used in the simulations are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.....	113
Figure 3.14. Plots of velocity standard deviations and percentage of successfully dealiased velocities for an array of true spectrum width and SNR values. The scanning strategies used in the simulations are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.	115
Figure 3.15. Two 90° sectors collected on July 5, 2007 at 21:30 UTC (left) and July 13, 2007 at 14:27 UTC (right) close to the times of the experiments.....	117
Figure 3.16. Means and standard deviations from weather data collected on July 13, 2007 at 14:23 UTC. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT	

(30.72 ms), and SBMX (43.01 ms) collection strategies. The dashed vertical line delineates the range of the spectrum plotted in Figure 3.17.	120
Figure 3.17. Sample non-Gaussian spectrum from a range of 54.8 km for the July 13, 2007 experiment.	121
Figure 3.18. Means and standard deviations from weather data collected on July 5, 2007 at 21:34 UTC. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies. . The dashed vertical line delineates the range of the spectrum plotted in Figure 3.19.	123
Figure 3.19. Sample non-Gaussian spectrum from a range of 45.5 km for the July 5, 2007 experiment.	124
Figure 3.20. Percentage of successfully dealiased velocities computed from weather data collected on July 5, 2007 at 21:34 UTC along with percentage computed from simulations. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.....	125
Figure 3.21. Dependence of amount of power removed using DC removal on wind speed. Data simulated using the quadratic clutter model to match collected data from the “Full Foliage, Prairie” conditions. The dotted lines are one standard deviation from the mean of the simulated data.	128
Figure 3.22. Sample spectrum from the “Full Foliage, Prairie, Light” case displayed over the whole Nyquist interval.	129
Figure 4.1. Probability of detection (POD) and probability of false alarm (PFA) for independent and contiguous pulses using a spectrum width of 2 m s^{-1} , SNR of 2 dB, and PRT of 3 ms.	138
Figure 4.2. Probability of detection (POD) for independent and contiguous pulses using a spectrum width of 2 m s^{-1} , SNR of 5 dB, and PRT of 3 ms. The SNR threshold is computed so that the PFA has a constant value of 1×10^{-6}	140
Figure 4.3. SNR_{dB} for weather signatures having a POD of 0.5 and 0.9 including the effects of a continuity constraint.	142
Figure 4.4. The relationship between POD and spectrum width for both independent and contiguous samples at PRTs of 0.8 and 3 ms. For the top plot, $M = 2$ and SNR = 8.79 dB. For the bottom plot, $M = 4$ and SNR = 9.28 dB.....	144
Figure 4.5. Detection scan with square pattern for beam positions in angle space.....	148
Figure 4.6. Detection scan with square pattern for beam positions in sine space.....	150
Figure 4.7. Detection scan with square pattern for beam positions in sine space with constant elevation constraint.	151
Figure 4.8. Detection scan with hexagonal pattern for beam positions in sine space with constant elevation constraint.	153
Figure 4.9. Detection scan with hexagonal pattern for beam positions in sine space with constant elevation constraint and overlap.	155
Figure C.1. Histogram of MSE values computed using a time series simulation.	177

Figure D.1. Sample ground clutter spectra and MSE histograms for the “Full Foliage, Urban” conditions. (Estimated model parameters: Light - $\hat{a} = 121.0$ dB/(m² s⁻²), $\hat{b} = 106.2$ dB/(m s⁻¹), $\hat{c} = 58.9$ dB, $\hat{d} = 23.3$ dB, Breezy - $\hat{a} = 38.5$ dB/(m² s⁻²), $\hat{b} = 61.5$ dB/(m s⁻¹), $\hat{c} = 64.0$ dB, $\hat{d} = 16.8$ dB, Windy - $\hat{a} = 8.2$ dB/(m² s⁻²), $\hat{b} = 30.3$ dB/(m s⁻¹), $\hat{c} = 57.6$ dB, $\hat{d} = 15.2$ dB.) 179

Figure D.2. Sample ground clutter spectra and MSE histograms for the “Full Foliage, Prairie” conditions. (Estimated model parameters: Light - $\hat{a} = 238.4$ dB/(m² s⁻²), $\hat{b} = 145.8$ dB/(m s⁻¹), $\hat{c} = 53.6$ dB, $\hat{d} = 26.8$ dB, Breezy - $\hat{a} = 51.8$ dB/(m² s⁻²), $\hat{b} = 73.8$ dB/(m s⁻¹), $\hat{c} = 67.0$ dB, $\hat{d} = 11.8$ dB, Windy - $\hat{a} = 6.8$ dB/(m² s⁻²), $\hat{b} = 30.0$ dB/(m s⁻¹), $\hat{c} = 65.3$ dB, $\hat{d} = 7.0$ dB.) 180

Figure D.3. Sample ground clutter spectra for the “Light Foliage, Urban” conditions. (Estimated model parameters: Light - $\hat{a} = 785.7$ dB/(m² s⁻²), $\hat{b} = 239.0$ dB/(m s⁻¹), $\hat{c} = 50.6$ dB, $\hat{d} = 32.9$ dB, Breezy - $\hat{a} = 196.4$ dB/(m² s⁻²), $\hat{b} = 142.6$ dB/(m s⁻¹), $\hat{c} = 65.3$ dB, $\hat{d} = 14.5$ dB, Windy - $\hat{a} = 28.8$ dB/(m² s⁻²), $\hat{b} = 57.9$ dB/(m s⁻¹), $\hat{c} = 63.6$ dB, $\hat{d} = 12.4$ dB, Gale - $\hat{a} = 6.2$ / (m² s⁻²), $\hat{b} = 26.8$ dB/(m s⁻¹), $\hat{c} = 60.3$ dB, $\hat{d} = 13.9$ dB.) 181

Figure D.4. Histograms of the MSE for both the quadratic and linear fits for the “Light Foliage, Urban” conditions. 182

Figure E.1. Quadratic log-domain model with $a = 20$, $b = 50$, $c = 60$, and $d_{4096} = 20$ (or $d = 3.7$ db). 185

Figure E.2. Data and fit showing relationship between d_n used in the simulations and the estimated d_n 187

Abstract

Weather radar researchers have long been eager to exploit the capabilities of phased array antennas, but high cost and technical complexity have postponed their widespread use in radar meteorology. With the aging of the current network of operational Doppler weather radars, the possibility of replacing them with phased array radars has renewed interest in applying this technology to weather radar research. The main focus of this research is the “agile beam” or electronic scanning capability of phased array antennas. Three research areas that take advantage of this agile beam capability are addressed in this work: spectral characterization of ground clutter with phased array radar data, staggered PRT beam multiplexing (SBMX), and rapid weather detection. Most of the research on ground clutter filtering has been applied to rotating antennas, but the agile beam capability of the phased array allows the collection of data with a stationary antenna. Studying the characteristics of ground clutter spectra for a stationary antenna could lead to new techniques and improvements for clutter filtering with phased arrays. Ground clutter data were collected under varying wind conditions, foliage levels, and terrain types. The shapes of the ground clutter spectra are then characterized using a novel quadratic clutter model, and the dependence of the model parameters on different conditions is explored. The model is then applied to the examination of clutter width and the time series simulation of ground clutter. SBMX takes advantage of the ability of the phased array to scan the beam in a different direction on a pulse-to-pulse basis which can save time by collecting samples that are nearly independent. SBMX is compared to two conventional scanning strategies to assess its performance using both simulations and real data. It performs well at high signal-to-noise ratios and narrow spectrum widths, but the staggered PRT strategy performs comparably to SBMX, takes less time, and has proven strategies for clutter filtering. The last area of

research, rapid weather detection, looks at the use of beam multiplexing to improve the detection of weather signatures. A simple beam multiplexing strategy outperforms a contiguous pulse strategy because the probability of detection of weather signatures is constant for beam multiplexing while the probability of detection for contiguous pulses decreases at narrow spectrum widths. The effects of beam broadening on the scanning strategies are also examined.

1. Introduction

Weather radar researchers have been interested in capitalizing on the capabilities of phased arrays for years, but the costs have been too high for their widespread use in radar meteorology. The main advantage of phased arrays is the “agile beam” or electronic steering capability. This capability is described by Skolnik (2008) as follows:

The dramatic advantage of electronically steered phased arrays as compared to reflectors is provided by the time required to steer beams and the flexibility in steering. While prior radars took seconds to steer to a new location, phased arrays take microseconds. In addition, the new location can be anywhere in a hemisphere.

The agile beam phased array weather radar opens up several new areas of research that are not feasible when utilizing a conventional mechanically-rotating reflector antenna. The cost of phased array antennas are still high compared to conventional antenna technology, but projections of future cost reductions in phased array elements and the possibility of combining the functions of multiple radars into one phased array system have renewed interest in phased array weather radars (Zrnić et al. 2007).

As the National Weather Service’s current operational Weather Surveillance Radar - 1988 Doppler (WSR-88D; Crum and Alberty 1993) ages, the renewed interest in phased arrays has increased the likelihood of a future replacement incorporating phased array technology. The Next Generation Weather Radar (NEXRAD) network of WSR-88D radars was the subject of a 2002 report partially aimed at “identifying the most promising approaches for the development of its eventual replacement” (National Research Council 2002). One of the recommendations in this report articulates a clear need for phased array research: “Adaptive waveform selection, which may even be applied to present systems, and agile beam scanning strategies, which require an electronically scanned phased array system, should be explored to optimize performance in diverse weather.” Fortunately, the National

Weather Radar Testbed (NWRT) located in Norman, Oklahoma has provided an opportunity to begin studying some of these new areas of phased array research.

The NWRT includes a SPY-1A antenna which is a passive array with 4,352 elements and a beamwidth of 1.5° at broadside (Forsyth et al. 2005). A block diagram of the NWRT is shown in Figure 1.1 which was adapted from a drawing from the original NWRT user manual (Frasco and Katz 2003).

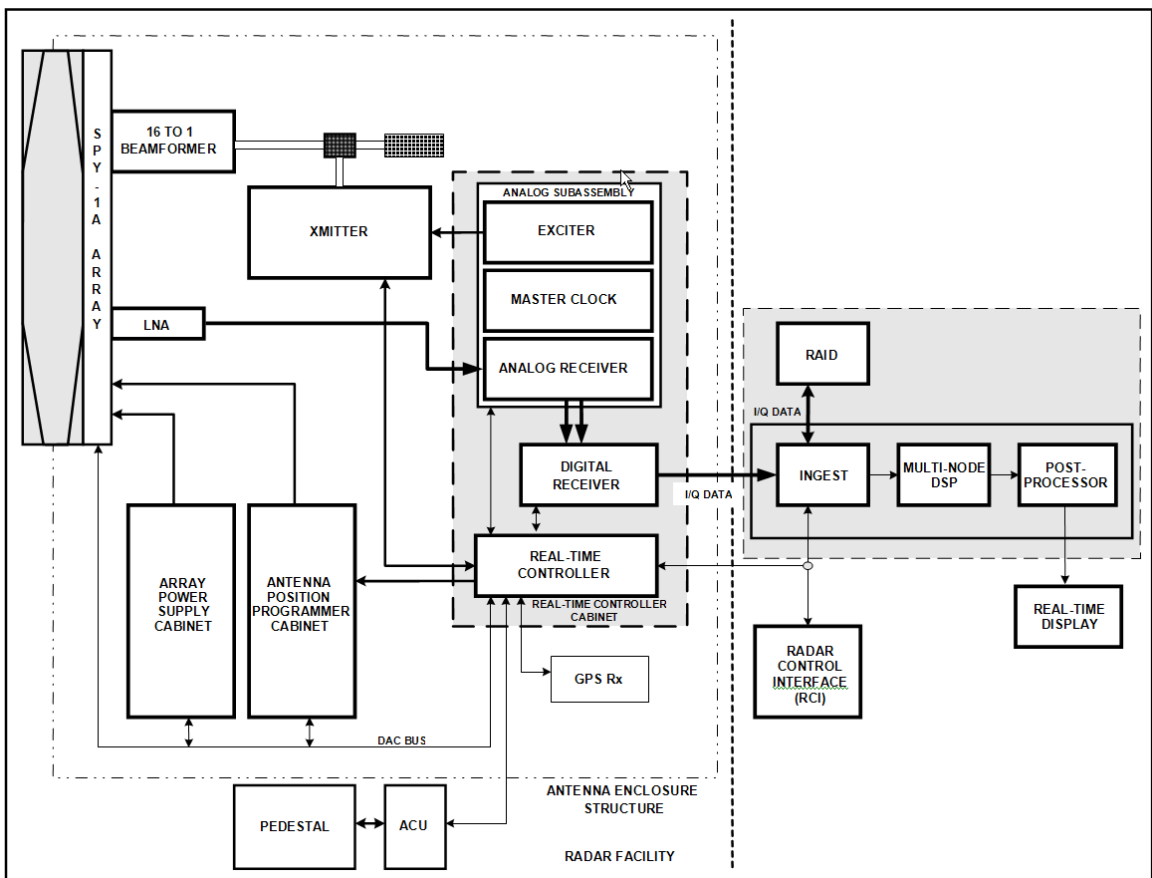


Figure 1.1. NWRT system block diagram after modifications by personnel from the National Severe Storms Laboratory (NSSL).

The presence of the NWRT has already encouraged phased array research in several areas. Zhang and Doviak (2007) have looked at splitting the array into two halves to implement spaced antenna interferometry. In addition, the NWRT has been applied to measuring

changes in refractivity (Cheong et al. 2008) and to examining the effects of rapid sampling of severe storms (Heinselmann et al. 2008). The unique capabilities of the phased array have also led to research on new ways to mitigate ground clutter (Le et al. 2009). The NWRT, also referred to as the PAR (phased array radar), was used to supply time series radar data for study in the following chapters.

The main advantage of phased arrays, electronic steering or scanning, has already been described, but there are other characteristics of phased array antennas that differ from conventional, mechanically-scanned antennas. When the phased array beam is electronically steered, the beamwidth increases as the beam is steered away from broadside. An approximate formula illustrating this effect is adapted from Skolnik (2008):

$$B(\theta_s) \approx \frac{B_0}{\cos \theta_s}, \quad (1.1)$$

where B_0 is the beamwidth at broadside and θ_s is the scanning angle away from broadside. For example, the 1.5° broadside beamwidth for the NWRT increases to about 2.1° when the beam is steered 45° from broadside. In addition to the beam broadening effect, the gain of the antenna decreases as the beam is steered away from broadside. An equation for this decrease in gain is shown below and is slightly modified from the equation given in Skolnik (2008):

$$G(\theta_s) \approx G_0 \cos \theta_s, \quad (1.2)$$

where G_0 is the gain at broadside and θ_s is the scanning angle away from broadside as in equation (1.1). Both of these effects need to be taken into account when designing a phased array radar if the performance at all scanning angles needs to meet certain requirements. These effects also need to be considered when conducting research with a phased array antenna.

The electronic steering capability of phased arrays can be utilized in a myriad of ways, but one of the simplest is scanning electronically while the antenna is stationary. In fact, the current design of the Multifunction Phased Array Radar (MPAR), a possible replacement for the NEXRAD radars, includes four stationary phased array antennas which will enable full volumetric scanning without any mechanical antenna rotation (Zrnić et al. 2007). In radars with mechanically-scanned antennas, the rotation causes a broadening of the spectrum. This broadening can have a significant effect on the spectrum width of stationary targets such as ground clutter. Because these broadening effects do not occur when using stationary phased arrays, spectral processing techniques such as spectrum-based ground clutter filters may need to be adapted to take advantage of the narrower ground clutter signatures. Characterizing the shape and width of clutter in the spectral domain could be helpful in improving the performance of ground clutter filters. The goal of Chapter 2 is to attempt to characterize the shape of ground clutter targets in the spectral domain for various wind conditions, terrain types, and both full and light foliage situations.

In addition to the possibility of a stationary antenna, electronic scanning also enables novel scanning strategies that are not feasible with a conventional antenna. The basic idea behind one particular class of strategies has been discussed since at least 1970. In a memorandum from Dennis and Smith (1970), they write that “we have investigated the possibility of improving the efficiency of the system by moving the antenna rapidly so that echoes could be obtained from other bearings while waiting for the scatterers in any given contributing region to reshuffle into an independent array.” One limitation of conventional scanning strategies that send a train of pulses in one particular direction is that the samples collected from these pulses are correlated because the scatterers have not moved or “reshuffled” enough from one sample to the next for the measurements to be independent.

When these samples are used to compute the power or lag-1 autocorrelation, the variance of the autocorrelation estimates is not reduced by the number of measurements because of the correlation between them. Dennis and Smith continue:

The basic idea of looking elsewhere while waiting for particle shuffling to occur in a given contributing region could be applied if we had an antenna capable of electronic scanning. Such antennas do exist; they were developed for applications such as missile tracking where very rapid scanning is required. They are generally known as phased-array antennas; most of them are designed to scan electronically in two dimensions...

This class of scanning strategies that decrease the variance of estimates by revisiting a particular location after time has passed is denoted by the term “beam multiplexing” or BMX. Some straightforward implementations of BMX have already been implemented on the NWRT (Yu et al. 2007), but a new BMX strategy is studied in Chapter 3 that addresses some of the limitations of the earlier strategies.

The most significant benefit of electronic scanning can be described with the term “adaptive scanning.” The recommendation from the National Research Council (2002) that was quoted earlier mentioned both “adaptive waveform selection” and “agile beam scanning strategies.” “Agile beam scanning strategies” include techniques such as BMX but could also include strategies for only scanning the regions where weather is occurring. For this study, “adaptive scanning” is used as an umbrella term that includes both “adaptive waveform selection,” “agile beam scanning strategies,” and other scanning techniques that reduce the overall scan time compared to a conventional antenna. The long-term goal of phased array weather radar development should be finding the most efficient way to surveil the weather using optimal scanning strategies that improve the warning lead times for hazardous weather and/or increase our understanding of meteorological phenomena. Adaptive scanning is a key element for reaching that goal.

In Chapter 4, some preliminary aspects of adaptive scanning are explored. If one goal of adaptive scanning is to only collect weather data where the weather is occurring, then some method needs to be employed to detect the weather. One possible approach for detecting weather signatures for more intensive surveillance is to scan the whole volume quickly using only a few pulses at each beam position. To efficiently accomplish this goal, another application of beam multiplexing is examined that attempts to maximize the probability of detecting weather signatures using the fewest possible pulses at a particular beam position. In addition, several scanning strategies are introduced for efficiently covering a significant portion of the total scanning volume. These strategies take into account the changing shape of the phased array beam as the beam is electronically scanned. Adaptive scanning is a rich area for future research and can also be extended to include other areas such as aircraft tracking in future multifunction phased array radars.

Even though researchers have been thinking about ways to utilize the capabilities of phased arrays for weather applications for years, much of the research is still in the early stages. The NWRT is a unique site dedicated to weather-related phased array research and provides a much-needed platform for studying this promising application of phased array technology. As the cost of phased arrays continues to decrease, the prospect of their use in weather applications will garner more interest. This is a pivotal time in the development of radar meteorology as the need to replace and improve the NEXRAD system coincides with a technology on the cusp of affordability. By combining the ideas from researchers who have been looking ahead to this moment with new ideas and hardware advances, the radar meteorology community can realize the full potential of agile beam phased array weather radar.

The outline for this dissertation follows the structure laid out in the previous paragraphs. Chapter 2 examines the shape of ground clutter spectra when using a stationary antenna. Data were collected under different wind, foliage, and terrain conditions and processed to aid in the development of a flexible clutter model. BMX is the subject of Chapter 3, and a new beam multiplexing technique is introduced to address some of the limitations of earlier beam multiplexing strategies. The new scanning strategy is compared to two existing scanning strategies using WSR-88D requirements. In Chapter 4, BMX is applied to weather detection; weather detection strategies are important to fully take advantage of the adaptive scanning capabilities of phased array radars. In addition, strategies for efficiently covering the scanning volume are investigated. Lastly, Chapter 5 includes conclusions and suggested future work.

2. Spectral Characterization of Ground Clutter

2.1 Introduction

Agile-beam phased array weather radars provide several advantages over conventional reflector-based radars including the ability to electronically scan while the antenna is stationary. Radars with parabolic dish antennas need to rotate mechanically in both azimuth and elevation to collect volumetric weather data. This antenna rotation has the effect of broadening the spectra of ground clutter and therefore changing the spectral shape. As mentioned in Chapter 1, the current design of the Multifunction Phased Array Radar (MPAR) includes four stationary phased array antennas which will enable full volumetric scanning without any mechanical antenna rotation (Zrnić et al. 2007). The effects of antenna rotation on the shape of the clutter spectrum are exemplified in Figure 2.1. The rotation was simulated electronically using the PAR by transmitting 4096 pulses evenly spaced across a 1.5° sector corresponding to the beamwidth. The data for the stationary spectrum was collected by transmitting 4096 pulses at the center of the sector. Both spectra were processed as described in Section 2.3. The shape of the spectrum is greatly affected by rotation.

Early studies that investigated the effects of ground clutter on mechanically-rotated Doppler weather radars focused on the spectrum width (Zrnić and Hamidi 1981). In the Zrnić and Hamidi study, the mean of the spectrum width was shown to increase by approximately a factor of four when the antenna rotated at a rate of 10° s^{-1} compared to 1° s^{-1} . The distribution of spectrum widths was significantly different at the slower rotation rate of 1° s^{-1} with a majority of the measurements having spectrum widths smaller than 0.1 m s^{-1} . This early work demonstrates that the spectrum widths of ground clutter measured

using a stationary antenna (with an approximately 1° beamwidth) differ significantly from the spectrum widths from rotating antennas.

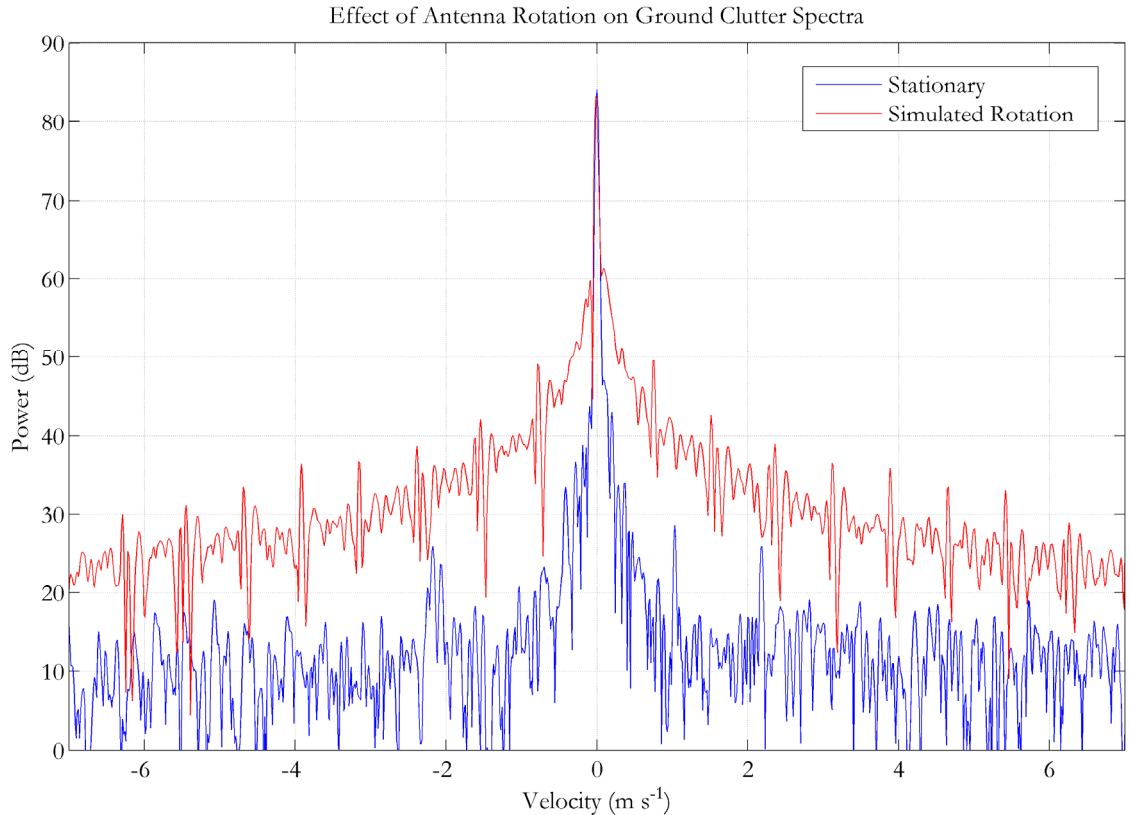


Figure 2.1. Comparison of ground clutter spectra from a stationary antenna. Rotation was simulated using the electronic scanning capability of the phased array antenna.

One of the goals for this work is to quantify the spectrum width of ground clutter observed using a stationary antenna. Note that the width can be measured using different methods; some of the spectrum width estimators implemented on operational radars assume a Gaussian spectral shape such as the covariance (or pulse pair) method (Doviak and Zrnic 1993). The shape of the clutter spectrum and the amount of departure from a Gaussian shape can affect the results of those spectrum width estimators. The shape can also affect the performance of spectral ground clutter filters because the amount of power removed depends on the amount of power in the Doppler bins around zero velocity. Some filters

such as the Gaussian Model Adaptive Processing filter (GMAP; Siggia and Passarelli 2004) assume a Gaussian shape for the clutter spectrum which may not be accurate for a stationary antenna. These considerations drive another goal of this work: finding the shape of the ground clutter spectrum for a stationary antenna and developing a width estimator related to the effectiveness of an ideal spectral ground clutter filter.

The shape of the clutter spectrum changes under different conditions such as wind speed, terrain type, and foliage level. At high wind speeds, the clutter shape is wider than at low wind speeds because the objects contributing to the clutter spectrum are moving more. The terrain and seasonal differences in foliage also affect the shape of the clutter spectrum because of differences in the amount of clutter scatterers that move, such as leaves. By examining clutter spectra collected under a wide variety of conditions, the range of possible clutter shapes can be studied and relationships between the conditions and the clutter shapes can be developed. To better understand these relationships, times series data were collected using the PAR at the NWRT under several different wind, terrain, and foliage conditions. The spectra computed from the data were used to model different clutter spectral shapes, and a representative clutter shape was computed for each set of conditions. Then, the effects of wind speed, terrain, and foliage level on the clutter shape were studied.

Moreover, the modeled clutter spectrum can be used in the simulation of time series data based on the standard simulation technique introduced by Zrnić (1975). This simulated data can then be utilized to investigate the performance of various clutter filtering techniques or to study novel clutter filtering strategies with applications to BMX or other scanning strategies. In this study, time series simulations are used to assess the effect of sequence length on the characterization of spectral clutter shapes.

For clutter characterization, finding an appropriate model for fitting the clutter spectra is essential. A promising model in *Low Angle Radar Land Clutter: Measurement and Empirical Models* by Billingsley (2002) was developed to characterize windblown clutter spectra. This book is one of the most comprehensive looks at ground clutter spectra using a stationary antenna. Measurements of ground clutter spectra at several frequencies and under varied wind conditions are presented. Billingsley proposes a two-part spectral model to fit the windblown clutter power spectral density (PSD):

$$P_{tot}(v) = \frac{r}{r+1} \delta(v) + \frac{1}{r+1} P_{ac}(v). \quad (2.1)$$

The first part is a DC (or zero-velocity) term that is added to an AC term; the ratio of the DC power to the AC power is given by r . Billingsley offers a simple explanation for the two-part model: “Consider a radar spatial resolution cell containing windblown trees. Such a cell contains both fixed scatterers (ground, rocks, tree trunks) and moving scatterers (leaves, branches). The returned signal correspondingly contains both a constant (or steady) and a varying component.” Billingsley also gives an analytic expression for the dependence of r on wind speed, $r = 489.8w^{-1.55}f_0^{-1.21}$, where w is the wind speed in mph, and f_0 is the radar frequency. The value of r decreases with wind speed as the power shifts from the DC term to the AC term. The AC (varying) term has a two-sided exponential shape:

$$P_{ac}(v) = \frac{\beta}{2} \exp(-\beta|v|). \quad (2.2)$$

When looking at the clutter PSD in a dB scale, the two-sided exponential term becomes linear with the β parameter determining the slope of both sides. In the log domain, the model spectrum described in equation (2.1) consists of an inverted “V” that represents the varying part of the spectrum with a DC component that is added at zero velocity. An

example is shown in Figure 2.2 with $r = 10$ and $\beta = 5.7$. Billingsley states that this value of beta is typical for windy conditions (wind speeds from 15-30 mph).

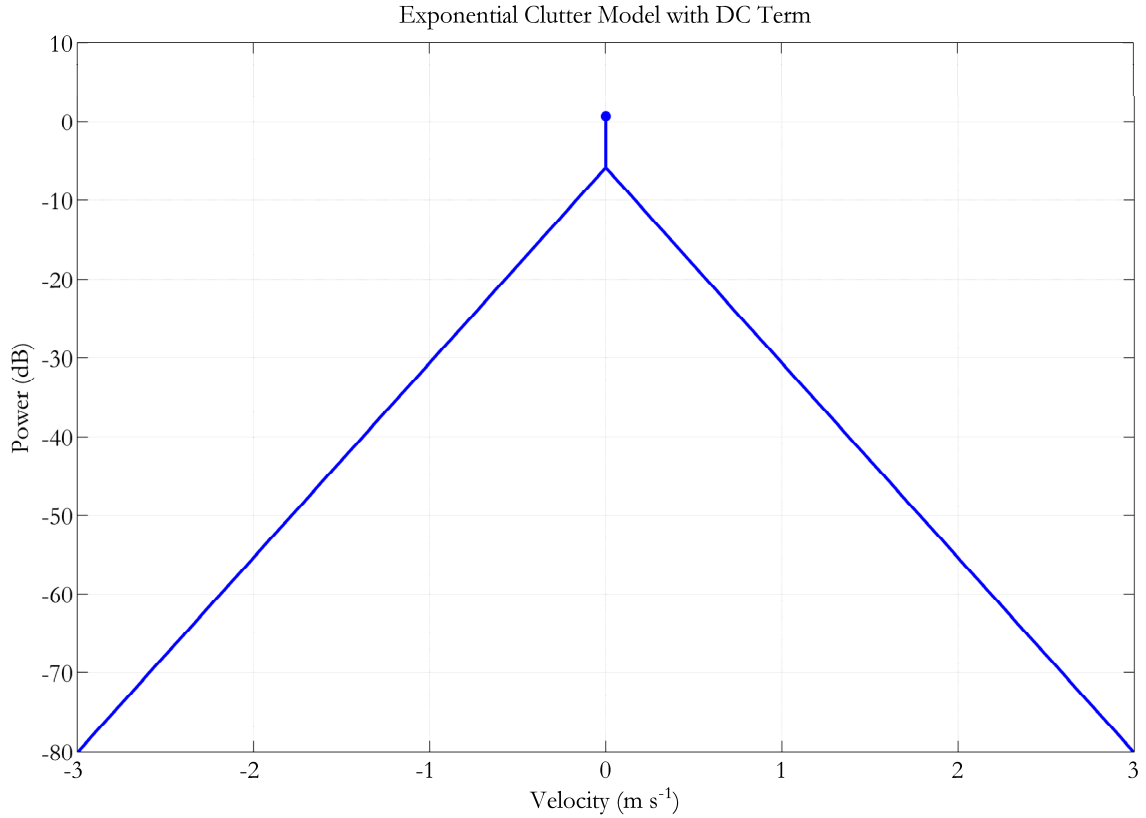


Figure 2.2. Billingsley windblown clutter model with $r = 10$ and $\beta = 5.7$.

The power is plotted using a dB scale to show the linear shape of the AC part of the model. Focusing on the shape of the clutter spectra in the log domain has one main advantage compared to modeling in the linear domain: fitting Billingsley's model or another similar model in the log domain results in a more straightforward linear or near-linear fit. For this study, a new ground clutter model was introduced which is similar to Billingsley's two-part exponential model but adds some additional complexity to encompass the spectral shapes in the collected PAR data.

The rest of this chapter can be divided into three main parts. Sections 2 through 7 describe the collection of time series data, introduce the new ground clutter model, explain the procedure used to fit the clutter spectra, report the results of the spectral fitting, and examine the effect of beam broadening. In Section 8, the effects of wind speed, terrain, and foliage level are studied. The next four sections, Sections 9 through 12, address some issues with the new model and with simulating time series data, examine the clutter width and the effects of the length of the spectrum on the clutter spectral shape. Section 13 summarizes the findings from the three main parts and suggests future work and associated data collections that could help answer some questions that emerged during the research.

2.2 Collection of Clutter Data

To study the effects of different factors on the shape of ground clutter spectra, varying wind conditions, terrain types, and foliage levels were considered when determining the collection parameters. Based on the previous work on ground clutter spectra by Billingsley (2002), clutter data were collected with the PAR under the same three sets of wind conditions to facilitate comparisons. The three categories are defined as follows: Light Air (1-7 mph), Breezy (7-15 mph), and Windy (15-30 mph). A particularly windy day also led to one dataset collected at the low end of “Gale Force” conditions (30-60 mph). Plots from the Oklahoma Mesonet were saved during the data collection to provide an account of wind speed and direction. Data were also collected from three different terrain types—the same ones considered in the study by Zrnić and Hamidi (1981). The PAR is located very close to the former site of the Norman Doppler radar which was the one used by Zrnić and Hamidi. The first terrain type, referred to as “Urban” in this work, is to the north of the PAR in the direction of Oklahoma City (0° azimuth). The second terrain type is to the east towards Lake

Thunderbird and is identified as “Wooded” (90° azimuth). The third terrain type is “Prairie” and is located to the southwest of the PAR (225° azimuth). Figure 2.3 shows a satellite view of the area with the NWRT marked with an “A” and 10° regions denoting the three terrain types overlaid on the map.




Figure 2.3. A satellite view of the area surrounding the NWRT with 10° regions that extend from approximately 5 to 40 km showing the “Urban”, “Wooded,” and “Prairie” terrain types at azimuths of 0°, 90°, and 225°, respectively.

In addition to several wind speeds and terrain types, data collection was performed at two different times of the year in order to study the effects of different foliage levels. The

“Light Foliage” data were collected in late winter before the trees had leafed, and the “Full Foliage” data were collected in summer or very early autumn before the leaves started to fall from the trees. Using the original three wind categories, three terrain types, and two foliage levels, 18 cases were collected. An additional three cases were collected during the “Gale Force” conditions for a total of 21 cases. The “Gale Force” conditions occurred during a “Light Foliage” period, and all three terrain types were collected. Table 2.1 shows all 21 cases along with collection dates.

Table 2.1. A list of experiments for clutter characterization.

Foliage	Terrain	Wind	Collection Dates
Full Foliage	Urban	Light	August 30, 2007
		Breezy	August 27, 2007
		Windy	June 3, 2008
	Wooded	Light	August 30, 2007
		Breezy	August 27, 2007
		Windy	September 30, 2007
	Prairie	Light	August 30, 2007
		Breezy	August 27, 2007
		Windy	September 17, 2007
Light Foliage	Urban	Light	February 8, 2008
		Breezy	March 7, 2008
		Windy	February 25, 2008 (2:54 – 3:43 p.m. CST)
		Gale	February 25, 2008 (4:18 – 5:07 p.m. CST)
	Wooded	Light	February 8, 2008
		Breezy	March 7, 2008
		Windy	February 25, 2008 (2:54 – 3:43 p.m. CST)
		Gale	February 25, 2008 (4:18 – 5:07 p.m. CST)
	Prairie	Light	February 8, 2008
		Breezy	March 7, 2008
		Windy	February 25, 2008 (2:54 – 3:43 p.m. CST)
		Gale	February 25, 2008 (4:18 – 5:07 p.m. CST)

For each case, three 90° sectors were collected consecutively. This allowed for redundancy if something was wrong with any of the data and provided more instances of

ground clutter to examine. The center 10° sector was outlined on the map because only beam positions from the center 10° of each 90° sector were used when characterizing the ground clutter spectra. Each 90° sector was collected at the lowest allowable elevation angle for the NWRT, 0.5° , and was made up of 90 beam positions separated by 1° in azimuth. The data were collected using the electronic scanning capability (i.e., step scan) of the phased array so the antenna was stationary. At each beam position, 4096 pulses were transmitted in the same azimuth direction using a PRT (Pulse Repetition Time) of 0.8 ms. Thus, each sector took a approximately five minutes to collect. In most situations, all three terrain types were collected over a roughly 60 minute time period under the same wind and foliage conditions. For the “Full Foliage, Windy” cases, each of the three cases was collected on a different day as noted in Table 2.1. In addition, nine 90° sectors were collected on September 22, 2008 at 0° , 355° , 350° , 345° , 340° , 335° , 330° , 325° , and 320° azimuth to study the effects of beam broadening. These data are examined in Section 2.7.

2.3 Data Processing and Preliminary Model Fitting

In this section, recorded time series data from a few cases are processed, and the associated spectra are displayed. Various methods for calculating the spectra are discussed, and two different models for the AC part of the spectrum are examined and compared to determine the one to be used for spectral fitting. The goal is to use the preliminary results to identify possible issues that could affect the fitting process and then apply that knowledge towards improving the recommended model.

The processing is similar to conventional PAR data processing, but the large number of collected pulses allows certain additional approaches to be used. After reading in the IQ data, the DC bias is removed first. Although the DC bias is relatively small, it can affect the

accuracy of the noise measurement. The DC bias is determined by averaging IQ values from further ranges where the power is less than a specific threshold (roughly 20 dB above the mean system noise). Next, a digital matched filter is applied. Note that the inherent range resolution is about 240 m based on the 1.57 μ s pulse length with the time series data oversampled at 60 m intervals. The matched filter values were computed by looking at strong point-clutter targets in the oversampled data. A digital matched filter of length four is utilized to improve the signal-to-noise ratio (SNR) and produce nearly uncorrelated range samples.

After the initial processing is completed, a spectrum is calculated at each range gate. Because of the large number of samples, a periodogram should be sufficient to accurately capture the shape of the ground clutter. The initial attempt used all 4096 samples to compute the spectrum, but a couple of additional methods for averaging periodograms are also examined. With spectral processing, windowing the data is necessary to ensure that the spectral shape is optimally preserved, but several factors need to be taken into account when choosing a window. Billingsley mentions a “system or processing noise floor” in a discussion about a “wideband noise-like background component” that some researchers have postulated as part of their clutter models (White 1994). To ensure that the shape of the clutter signal was not obscured before reaching either a “processing floor” or a “wideband noise-like background component,” an aggressive window was chosen. In particular, a 4096-point, DFT-even, Chebyshev window with 150 dB of sidelobe attenuation was employed. This window should have a negligible effect on the smoother portions of the spectrum but will cause broadening of any peaks such as the DC component at zero velocity. The maximum width of the main lobe of the window is estimated to extend from -0.08 to 0.08

m s^{-1} so the broadening of a strong DC component should not obscure a significant portion of the spectrum even though the Chebyshev window is aggressive.

A few clutter spectra from different cases are presented to examine the limitations of the Billingsley model (2.1). In general, the spectra are plotted with power in dB versus velocity in m s^{-1} . A simple linear fit was applied in Figure 2.4 to the AC portion of the clutter spectrum that extends above the spectral floor; the details of this procedure are addressed in Section 2.4.

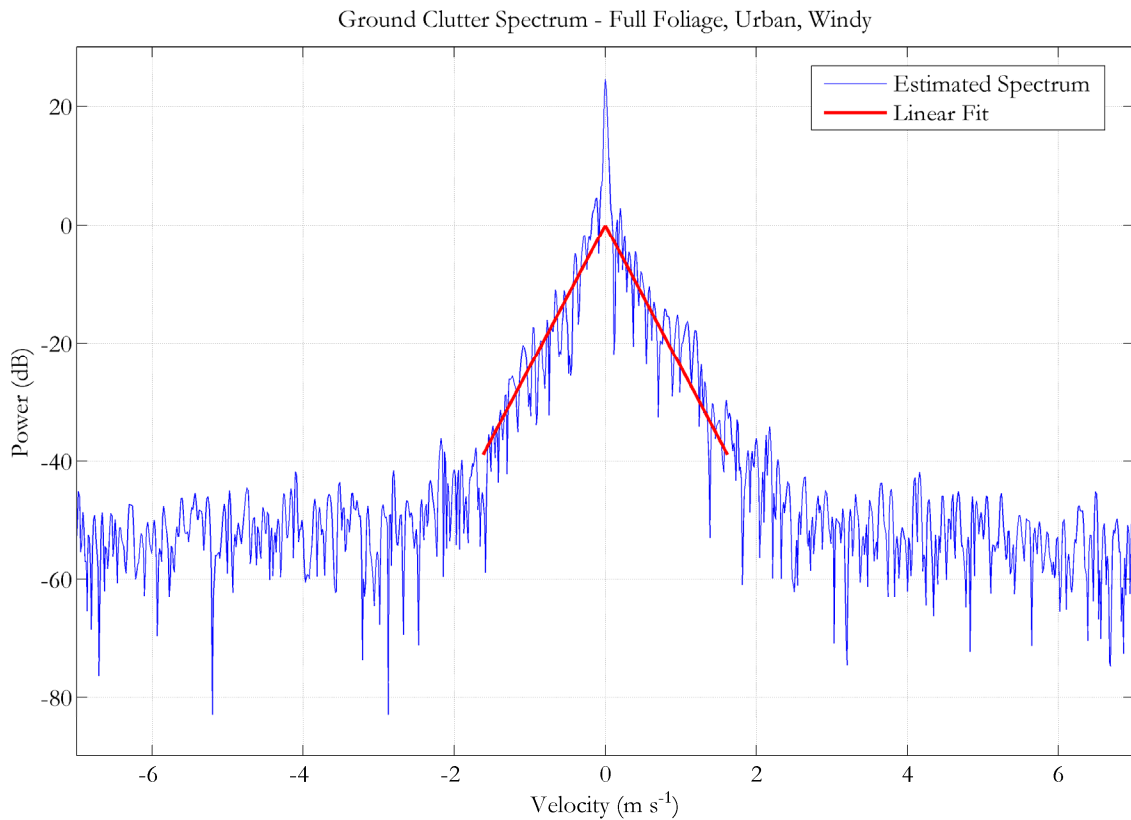


Figure 2.4. Example of ground clutter spectrum from the “Full Foliage, Urban, Windy” case with a simple linear fit of the AC portion of the spectrum.

Rather than scaling the power to one as in Billingsley’s model, the power at the peak of the linear fit is subtracted from the spectrum to shift the tip of the inverted “V” to 0 dB.

Because of the normalization (i.e. shift) of the spectrum, the DC term directly corresponds to the portion of the spectrum that extends above 0 dB. The spectrum in Figure 2.4 is a good representative of the “Full Foliage, Urban, Windy” case and is relatively free of contamination. Nevertheless, some of the spectra are contaminated with non-clutter signals. It is desirable to develop a fitting procedure that prevents corrupted spectra from producing unreliable estimates of the model parameters.

The Billingsley model fits the particular spectrum from Figure 2.4 very well. The AC or varying part of the spectrum is very close to linear as predicted by the model. A separate DC term can be seen extending about 25 dB above the peak of the AC portion of the spectrum. There also appears to be some sort of spectral floor about 80 dB below the zero velocity peak. This floor will be referred to as the spectral floor from now on and appears to be noise-like, but there are indications that it is not always a system noise floor. If all the spectra had the same shape as the one in Figure 2.4, the Billingsley model would be sufficient to model the ground clutter collected with the PAR. It is shown in Section 2.8 that clutter spectra collected under windy conditions (15-30 mph) tend to fit the two-part exponential model best.

The variance of the spectral coefficients of the periodogram in Figure 2.4 is a concern. The periodogram of length 4096 that was used to compute the spectrum has a high resolution, but the variance is the same as any unaveraged periodogram. Increasing the length of the periodogram does not reduce the variance of the estimates of the spectral coefficients; only the frequency resolution is increased (Doviak and Zrnić 1993). One way to decrease the variance is by averaging periodograms. A standard approach is to use Welch’s method and average periodograms from overlapping sequences of data. For example, the sequence of 4096 samples could be divided into seven overlapping sequences of length 1024.

The main drawback to this approach is that the frequency resolution is decreased. In this case, the main lobe of the aggressive Chebyshev window increases in width by a factor of four from the previously cited value of $\pm 0.08 \text{ m s}^{-1}$ to about $\pm 0.32 \text{ m s}^{-1}$. This obscures the transition between the peak of the AC part of the spectrum and the DC component as shown in Figure 2.5. Because of this drawback, this type of averaging was not employed.




Figure 2.5. Comparison of unaveraged periodogram of length 4096 and average of periodograms from seven overlapping sequences of length 1024.

Another type of averaging which was also employed by Billingsley is to divide the length-4096 sequence into four interleaved sequences of length 1024 where every fourth sample is utilized. This decreases the unambiguous velocity by a factor of four compared to the original sequence from about 29.28 m s^{-1} to about 7.32 m s^{-1} which is still sufficient for fitting the relatively narrow clutter spectra but can also lead to aliasing. An advantage of this

approach is that the resolution stays the same, and the main lobe width of the Chebyshev window does not increase. In Figure 2.6, the spectra for both the unaveraged and averaged periodograms are presented. The averaging does not reduce the variance significantly which is due to the strong correlation between the four interleaved sequences. The narrow spectrum width of ground clutter leads to a high correlation between adjacent samples of the original sequence. This explanation would also predict that the variance of the noise would be reduced more than the clutter since the noise should not be correlated from sample to sample, and this can also be seen in Figure 2.6.

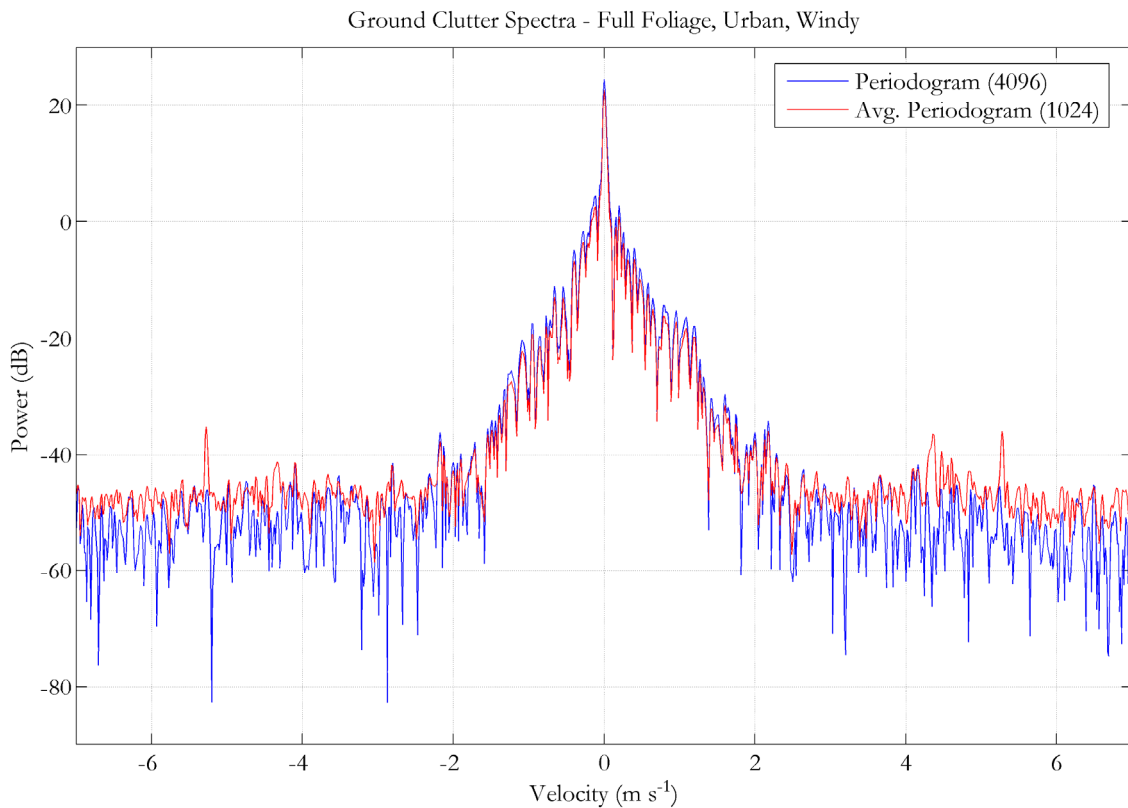


Figure 2.6. Comparison of unaveraged periodogram of length 4096 and average of periodograms from four interleaved sequences of length 1024.

The increased aliasing for the averaged periodogram does lead to a higher noise level and the appearance of a few additional signals. The minimal improvement in the variance and the

increased aliasing both point towards the use of the original length-4096 periodogram for estimating the spectra, and this method was employed for all of the processed spectra.

An additional spectrum is provided in Figure 2.7 that shows contamination of the spectrum with non-clutter signals that are likely from weather.




Figure 2.7. Example of ground clutter spectrum from the “Light Foliage, Urban, Breezy” case with a simple linear fit of the AC portion of the spectrum.

Contamination from non-clutter signals can interfere with fitting the spectrum, especially the spectral floor. In this case, the simple linear fit still performs adequately for fitting the AC part of the spectrum even though the spectrum is significantly narrower than the spectrum in Figure 2.4. A direct attempt to fit the spectral floor could be biased by the signals on the

left side of the spectrum; this is addressed in Section 2.4. The effects of lower wind speed and light foliage are studied in more detail in Section 2.8.

Although the AC part of the spectrum has been nearly linear when examining the first couple of example spectra, this does not necessarily hold true for all of the spectra. Two additional examples are shown in Figure 2.8 that were collected under two different sets of conditions. One is from the “Full Foliage, Urban, Breezy” case and the other from the “Light Foliage, Prairie, Breezy,” case. In both cases, the AC part of the spectrum departs from linear. Because of the curved shape, a more general quadratic model was developed to fit the AC portion of the spectra with the results also shown in Figure 2.8. The AC part of the spectrum seems to be linear or curved with a slightly concave-up shape which differs from the concave-down shape of the Gaussian function used to model weather spectra. The quadratic fit matches the slightly curved shape better than the linear fit, and the quadratic fit also allows a better fit of the peak of the AC portion of the spectrum. The better fit at the peak should also lead to more accurate estimation of the DC term because this term is calculated as the power at zero velocity extending above the peak of the AC part of the spectrum. In Figure 2.8, the spectra were shifted so that the peak of the quadratic fit is at 0 dB instead of the peak of the linear fit. The differences between these peaks are on the order of 10 dB or more which leads to a significant difference in the estimation of the two DC terms.

Based on the examples from this section and the examination of several other spectra, a new, more general clutter model is introduced in the next section that uses a quadratic model for the AC portion of the spectrum. Moreover, the spectral floor is included in the model because the spectral floor is not always the same as the system noise floor.

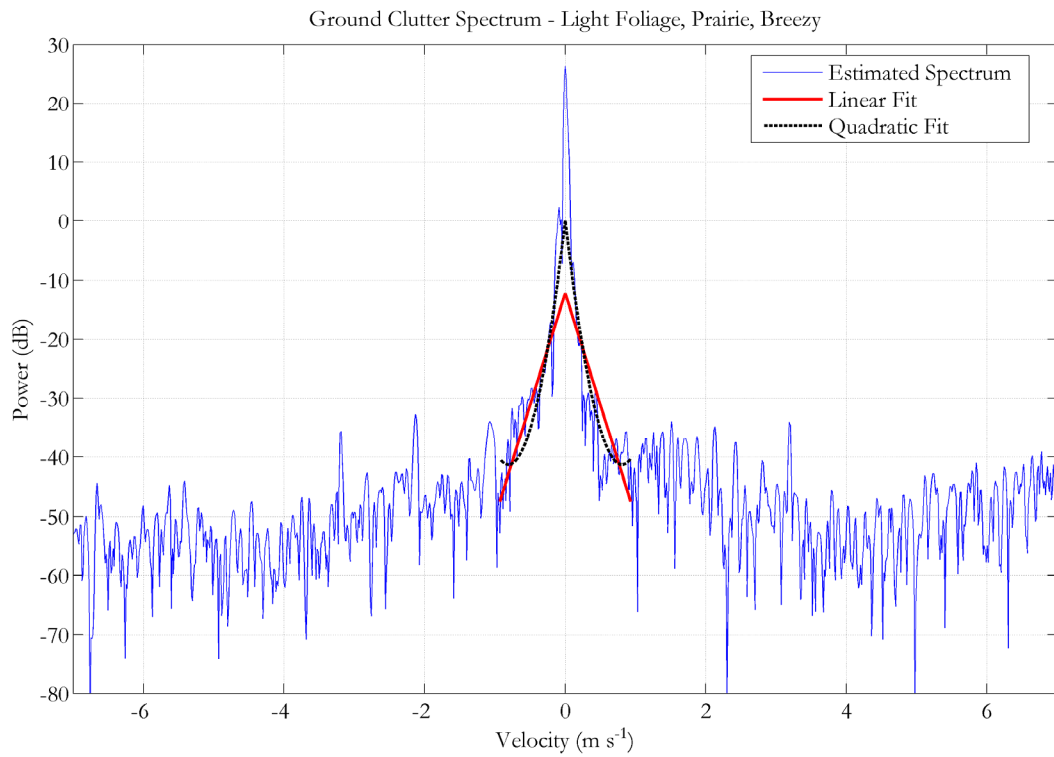



Figure 2.8. Two clutter spectra examples with both a linear and quadratic fit of the AC portion of the spectra.

The behavior of the quadratic fit as it approaches the spectral floor will need to be addressed at some point; the quadratic model starts to curve up in the “Light Foliage, Prairie, Breezy” example which does not match the expected decrease in the AC part of the spectrum as the velocity increases. These issues will be accounted for later in the chapter, but the first step in defining the new quadratic clutter model is to determine all of the model parameters.

2.4 Quadratic Ground Clutter Model: Spectral Fitting

In this section, a new quadratic ground clutter model is defined, and a procedure for fitting spectra is developed. The quadratic ground clutter model has three parts: the AC portion of the model, the DC term, and the spectral floor. This is similar to the Billingsley model except that the AC part of the spectrum is modeled with a quadratic function instead of a linear function and a spectral floor parameter is included. The modeling is also implemented in the log domain to simplify the quadratic fit. The log domain model for the AC and DC terms is given as follows:

$$P_{dB}(v) = d \cdot \delta(v) + P_{AC}(v). \quad (2.3)$$

The d parameter is measured in dB and determines the extent of the DC term above the peak of the AC part of the spectrum. The spectral floor is not included in (2.3) because it cannot be added directly in the log domain, but the c parameter, the distance in dB from the peak of the AC portion of the spectrum to the spectral floor, relates the spectral floor to the AC part of the spectrum. The quadratic portion of the model has two parameters which determine its shape:

$$P_{AC}(v) = -b|v| + \frac{a}{2}v^2 \quad (2.4)$$

where b is the magnitude of the slope in dB/(m s⁻¹) at $v = 0$, and a represents the rate of change of the slope. The quadratic model includes the two-part exponential model as a

special case when $a = 0 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, and it is assumed that the spectrum is normalized so that the peak of the AC portion is at 0 dB. The shift of the spectrum does not affect the spectral shape but does simplify the interpretation of the model parameters. The DC term extends d dB above the peak, and the spectral floor is located at $-c$ which is c dB below the peak of the AC part of the spectrum.

Figure 2.9 shows how these four parameters combine to model the shape of an example spectrum. One of the spectra from Figure 2.8, the “Full Foliage, Urban, Breezy” case, is used. The details of the fitting procedure are given later, but this example serves as an introduction. After the fitting procedure is completely defined, it can be applied to the recorded time series data.

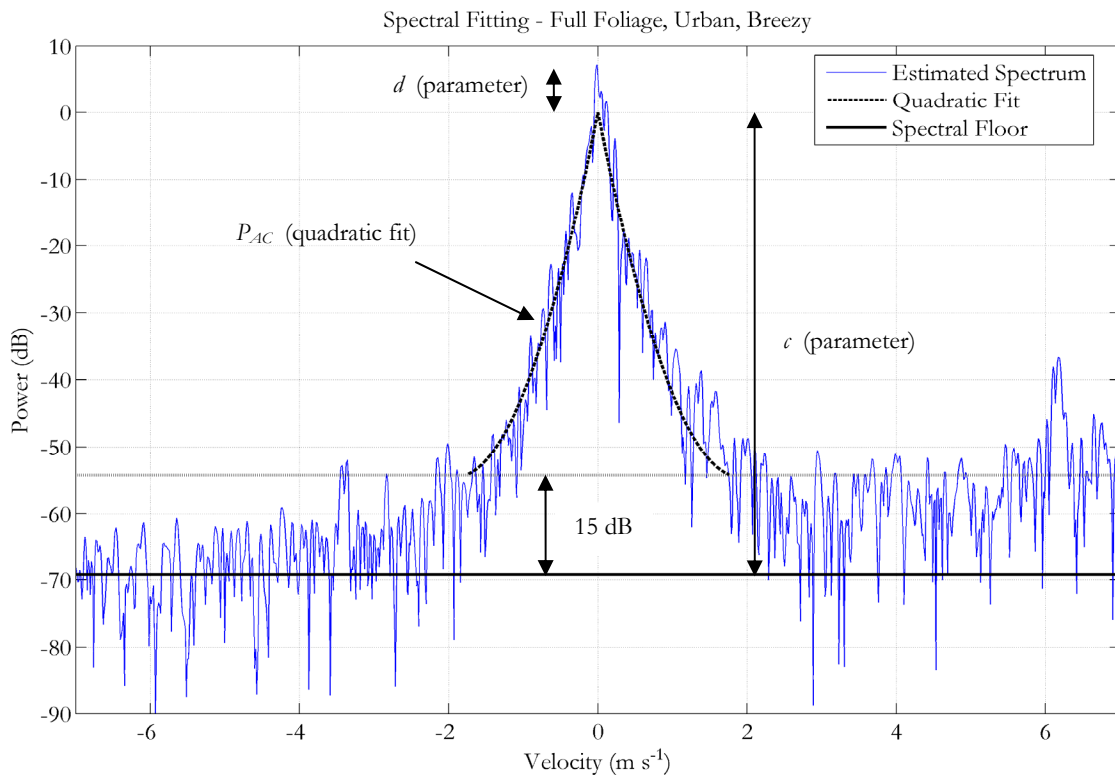


Figure 2.9. Example spectral fitting applied to a spectrum from Figure 2.8 showing the quadratic fit, the spectral floor, and c and d parameters.

The c and d parameters directly match their earlier definitions. The quadratic fit is also straightforward except that the fit only extends to 15 dB above the spectral floor. There are a few reasons for limiting the range of the quadratic fit. The first is that spectra can have non-clutter contamination near the spectral floor which can affect the fitting procedure. The second reason is that the curvature near the spectral floor is partially due to the addition of noise and thus does not represent the inherent shape of the clutter spectrum. The third is that the quadratic part of the model is limited to only one value, a , that determines the curvature over the range of the fit. By using this method for constraining the fit, the fit at the top of the AC part of the spectrum, the more important part, is improved. A way to fill in the missing part of the spectrum is addressed in Section 2.9 where the model is extended for use in clutter width calculations and simulations.

The next step is to describe the process for estimating the model parameters and starts with estimating the spectral floor. The goal is to estimate the spectral floor and shift the spectrum so the spectral floor is at 0 dB. This is a different spectral shift than the one used to normalize the peak of the AC portion of the spectrum to 0 dB. This shift is used to aid in the quadratic fit procedure and also to estimate the c parameter. If the spectral floor is at 0 dB, then the constant term returned from the quadratic fit is c . The details of the quadratic fit will be discussed after the spectral floor is estimated. The spectral floor includes system noise, but transmitter phase noise can also contribute to the spectral floor (Free and Patel 2005). The goal of this study is to characterize the shape of clutter spectra without including system noise. Only spectra with phase noise on the order of system noise or greater than system noise were included in the set of spectra to be processed. A technique similar to the “Dynamic noise power” technique used in the GMAP (Gaussian Model Adaptive Processing) clutter filter was utilized to measure the spectral floor (Siggia and Passarelli

2004). The details of the spectral floor and system noise estimation procedure can be found in Appendix A. The resulting shift is called the adjusted spectral floor, or F_{A} , which is in dB and is subtracted from the spectrum.

A final correction needs to be applied to the spectrum before the quadratic fit procedure. When fitting a spectrum in the dB domain, the logarithm function causes a shift in the spectrum. One way to see this is in terms of a spectral floor or noise floor. If the noise powers are averaged in the linear domain, the mean power matches the noise power. If the noise powers are converted to dB and averaged, the resulting power in dB is smaller than the true noise power (if converted to dB). The amount of shift can be quantified by looking at the distribution of the logarithm of exponential variates. This correction factor, ω , needs to be added to the spectrum before the quadratic fit. The correction factor and the distribution are discussed in more detail in Appendix B; the value of $\omega \approx 2.5068$ dB.

Figure 2.10 provides a flow chart for the fitting procedure. Step 1 is calculating the PSD or spectrum, and Step 2 is the application of the shifts to the spectrum for the adjusted spectral floor and the correction factor. Step 3 is estimating the a , b , and c parameters from the model, and Step 4 is estimating the d parameter. Because of the broadening of the DC component due to the Chebyshev window, four samples on either side of the zero-velocity sample are disregarded when performing the quadratic fit. An initial quadratic fit is performed using data starting with the fifth sample on either side of the zero-velocity sample and extending to a sample that is estimated to be near the 15 dB level. Because of the shifts, this should be near an absolute value of 15 dB with the adjusted spectral floor at 0 dB. The clutter spectrum is assumed to be symmetric, and both sides of the spectrum are fit simultaneously using a standard quadratic fitting algorithm. The quadratic fit is then repeated with either more or fewer samples until enough data are included so that the fit extends just

below the 15 dB level. The three values from the standard quadratic fit, $a_2v^2 + a_1v + a_0$, are then assigned to a , b , and c : the constant term, a_0 , is assigned to c which represents the distance between the peak of the AC portion and the adjusted spectral floor, the magnitude of the first degree term (the slope at zero), $|a_1|$, is assigned directly to b , and twice the quadratic term, $2a_2$, is assigned to a .

The final step is the calculation of the d parameter which is the value of the zero-velocity spectral coefficient minus c (assuming the spectral floor has been shifted to 0 dB). As the spectrum is being processed, four constraints are used to determine whether the parameters from a spectrum are included in the dataset. The first is that the zero-velocity spectral component is greater than or equal to 15 dB above the spectral floor before beginning the fitting process. This guarantees that there is significant power at zero velocity which indicates a possible ground clutter signal. The second constraint is imposed on the AC part of the spectrum. To ensure that enough of the AC part of the spectrum is used when fitting the spectrum, the c parameter must be greater than or equal to 45 dB. Since about 15 dB of data above the spectral floor are excluded from the fit, this condition ensures that at least the top 30 dB of the AC portion of the spectrum is utilized for fitting. The third constraint is that at least 10 spectral samples are used for the fit. Since the model is symmetrical, five samples from each side of the spectrum are required. The final check is based on a goodness of fit test to ensure the model reasonably matches the data; the mean squared error (MSE) is used to measure the error between the quadratic fit and the fitted data. The fourth constraint is that only spectra with MSE values below a certain threshold, 45 dB^2 , are kept for further processing. The threshold was determined using simulated data, and a full discussion can be found in Appendix C.

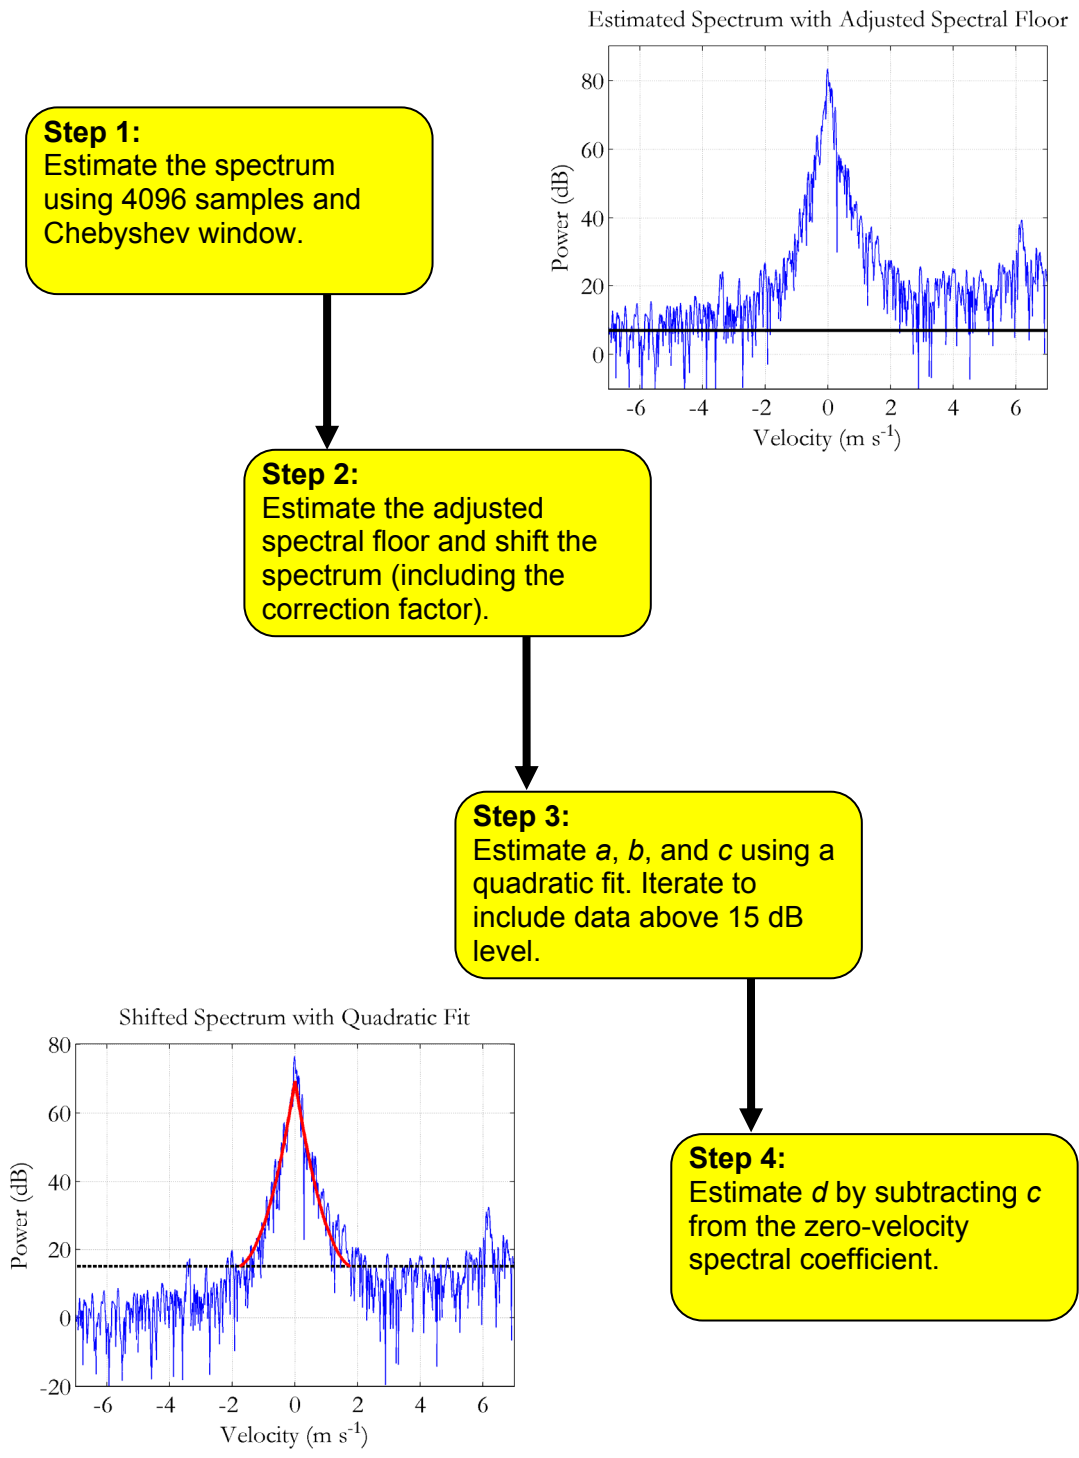


Figure 2.10. Spectral fitting procedure using spectrum from Figure 2.9.

2.5 Fitting PAR Data with the Quadratic Model

In this section, the recorded data from the “Full Foliage, Urban, Breezy” case are processed. There are two main goals: the first is to find a representative set of parameters for the case, and the second is to find possible relationships between the parameters to aid in modeling the dependence of each parameter on wind speed. The representative set of parameters determines a “representative clutter shape” for the case. One way to test the strength of the relationships between parameters is to investigate the parameter correlation matrix. If a high correlation between two parameters exists, the wind speed model can define one of the parameters in terms of the second. In other words, instead of finding two separate functions to determine the values of two model parameters based on wind speed, only one function is found and the relationship between the parameters can be used to determine the other parameter.

First, the required steps for processing the data from the “Full Foliage, Urban, Breezy” case are presented, and the representative clutter shape is determined. As discussed in Section 2.2 on data collection, only the center 10 azimuth positions of the three collected 90° sectors were processed. This ensures that possible beam broadening effects are avoided, that a sufficient number of spectra are included, and that the processed radials contain a significant amount of the desired terrain type. For the “Urban” terrain type, the 10 azimuth positions nearest 0° azimuth were processed. Because three sectors were collected and 10 radials are processed from each sector, 30 total radials were processed. These particular radials include clutter data from the Oklahoma City area and the metropolitan areas in between Oklahoma City and Norman as indicated in Figure 2.3.

After the processing was completed, 1959 spectra satisfied the weakest condition that the zero-velocity component was greater than 15 dB above the spectral floor and the condition

that the adjusted spectral floor was greater than the system noise floor. Of these, the fit failed to converge for 42 of the spectra leaving 1917. In addition, a significant number of the spectra, 520, had a ϵ parameter less than 45 dB. Especially at larger ranges, not enough of the clutter protrudes above the system noise floor to allow 30 dB of the AC part of the spectrum to be fit. Of the 1397 spectra where the fit converged and the ϵ parameter was greater than or equal to 45 dB, 1126 or 80.6% had an MSE below 45 dB². This shows that when enough of the shape extends above the spectral floor, the quadratic model successfully fits the data over 75% of the time (at least for this case). To compare the performance of the quadratic model to the linear model, the MSE values from both models are shown in Figure 2.11.

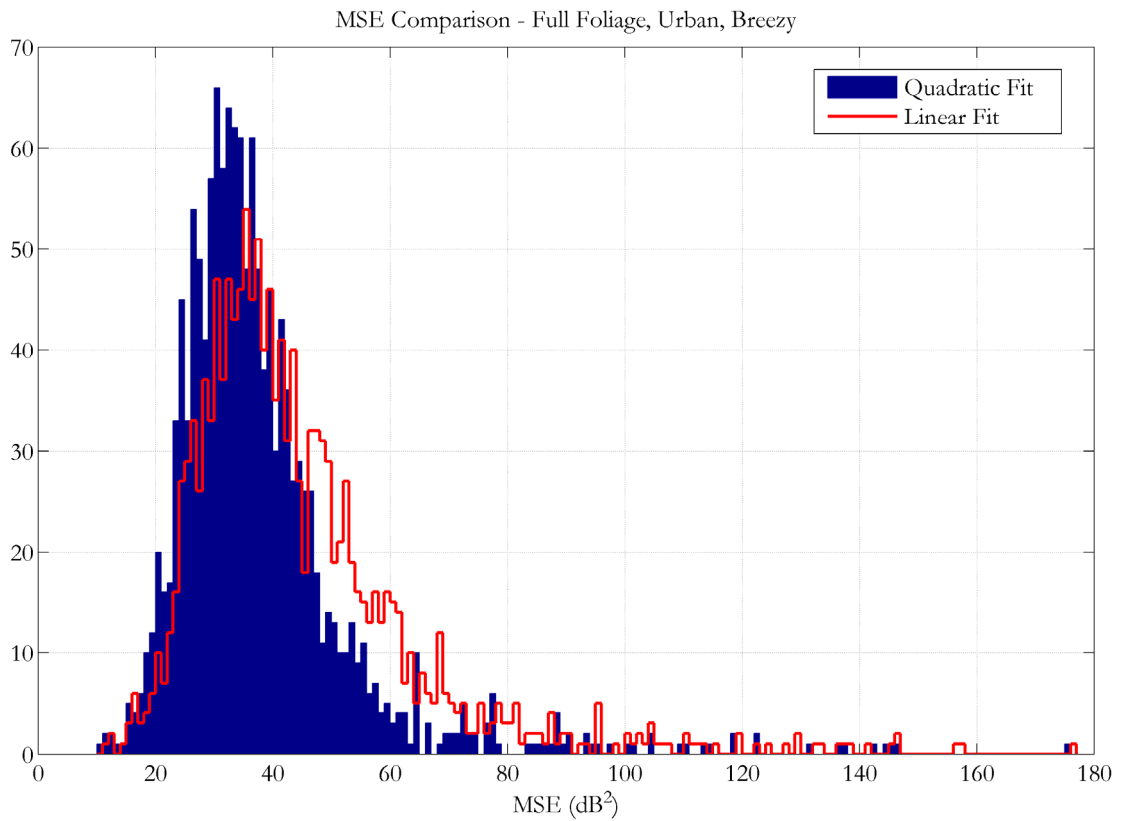


Figure 2.11. Histograms comparing quadratic and linear fits of the AC part of the spectra for the 1397 spectra where the fit converged and ϵ was greater than 45 dB.

The quadratic fit outperformed the linear fit as guaranteed by the additional degree of freedom, but 511 of the 1397 spectra (or 36.6%) did not meet the MSE threshold of 45 dB² using the linear fit compared to 271 (or 19.4%) for the quadratic fit. In addition, the mean MSE was 38.0 dB² for quadratic fit compared to 44.3 dB² for the linear fit. These results indicate that the quadratic model better characterizes the clutter spectra compared to the linear model.

After applying the fitting procedure to the spectra that met all of the conditions, 1126 sets of parameters were obtained. The next step is to examine the sets of parameters in order to find a representative and to find possible relationships between them. Approaches for determining and removing outliers are also needed because outliers can cause estimates of the representative parameters and the parameter correlation matrix to be skewed. For example, Figure 2.12 shows a histogram of the 1126 b values that were computed from the data. These values represent the magnitude of the slope of the AC portion of the spectra at zero velocity. The extreme positive values would skew the mean if included in the computation. These types of outliers occur because of failures in the fitting process and can affect all of the parameters. After several attempts at ad hoc outlier removal including removing outliers independently from each dimension, an established multivariate approach was employed. An implementation of the L1-median was utilized which identifies a median set that minimizes the sum of the Euclidean distances to all of the parameter sets (Hössjer and Croux 1995). This has a certain intuitive appeal because the L1-median is “close” to the other values and should act as a good representative of the set. Its use also avoids certain complications compared to other approaches that compute the mean of a set with outliers removed. The L1-median for this set of observations results in the following estimates: $\hat{a} = 37.2 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 61.3 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 61.9 \text{ dB}$, and $\hat{d} = 16.2 \text{ dB}$.

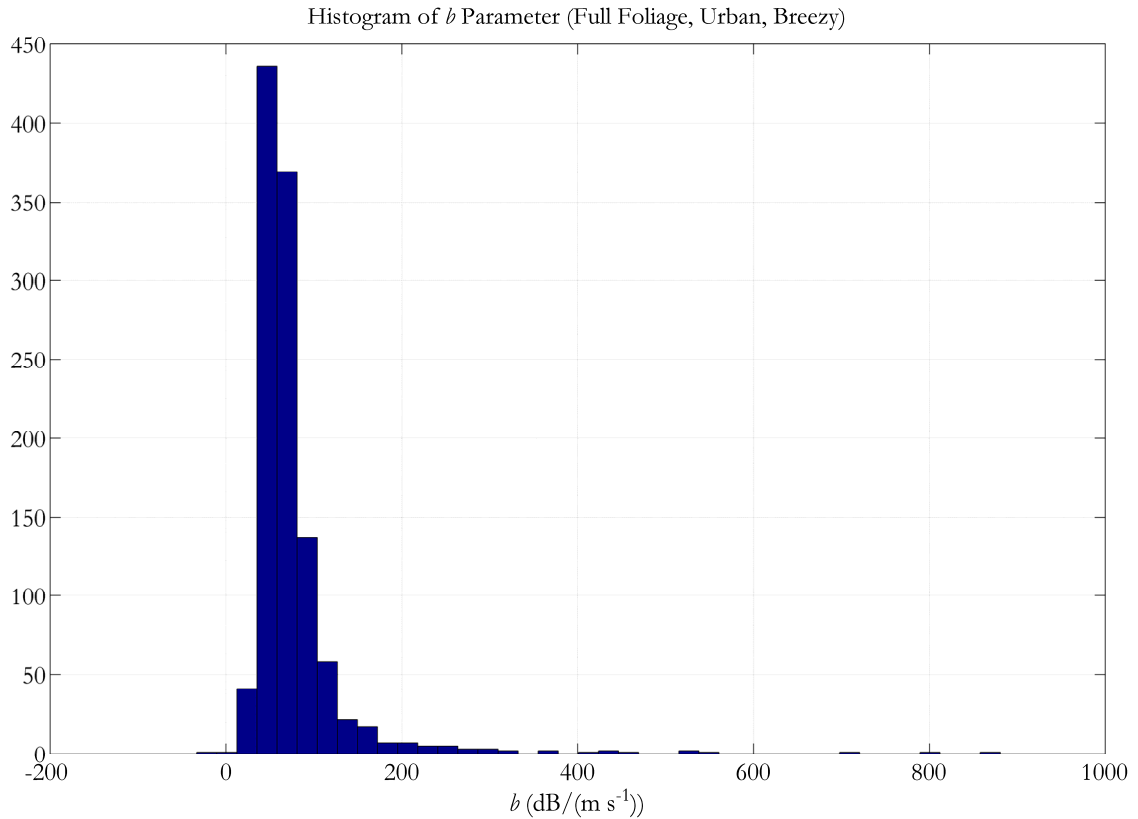


Figure 2.12. Histogram of b parameter values for “Full Foliage, Urban, Breezy” case.

The L1-median determines a reasonable representative clutter shape for this case, but another goal of this section is to seek possible relationships between the parameters to aid in modeling the dependence of the parameters on wind speed. To that end, a method for computing the correlation matrix for these parameters was sought, and another established multivariate approach was chosen. The minimum covariance determinant method (MCD) finds the b observations whose covariance matrix has the smallest determinant (Rousseeuw and Van Driessen 1999) where b corresponds to a certain percentage of the total number of observations (normally 75%). Each observation corresponds to a set of parameters calculated from a particular spectrum. After the b observations are determined, a chi-square test is applied to eliminate outliers. Figure 2.13 shows histograms for all four parameters

after the outliers have been removed. Of the 1126 sets of parameters, 227 outlier sets were removed leaving 899 sets included in the histograms. Although the distributions are not necessarily Gaussian or even symmetric, they seem to be reasonably well behaved.

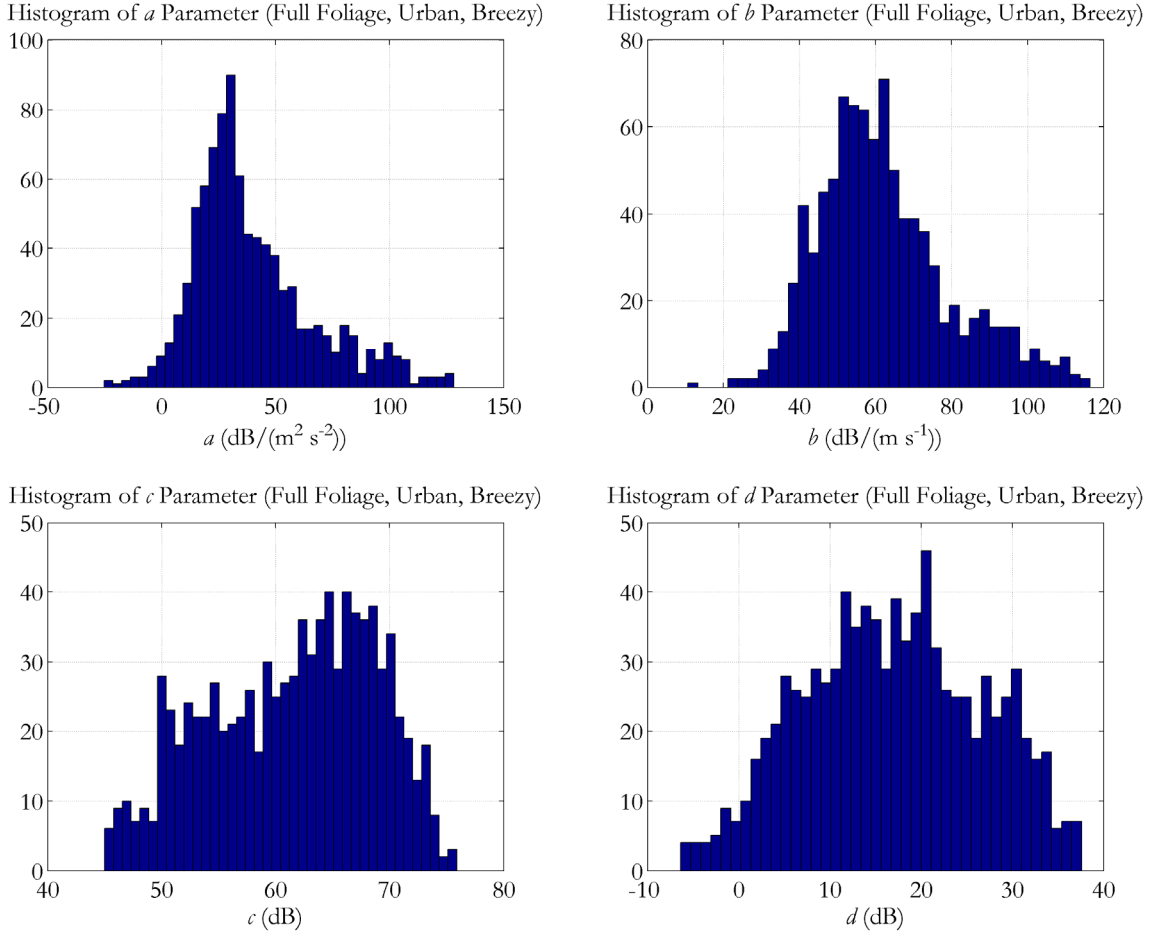


Figure 2.13. Histograms of a , b , c , and d parameter values for “Full Foliage, Urban, Breezy” case.

The correlation matrix computed using the MCD method is presented below:

$$\mathbf{R} = \begin{bmatrix} 1 & \rho_{a,b} & \rho_{a,c} & \rho_{a,d} \\ \rho_{a,b} & 1 & \rho_{b,c} & \rho_{b,d} \\ \rho_{a,c} & \rho_{b,c} & 1 & \rho_{c,d} \\ \rho_{a,d} & \rho_{b,d} & \rho_{c,d} & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0.94 & -0.18 & 0.24 \\ 0.94 & 1 & -0.00 & 0.09 \\ -0.18 & -0.00 & 1 & -0.92 \\ 0.24 & 0.09 & -0.92 & 1 \end{bmatrix}. \quad (2.5)$$

The elements of the matrix are correlations between pairs of parameters with $\rho_{b,c}$ denoting the correlation between b and c , for example. The diagonal elements are ones because the correlation of any variable with itself is 1. There are strong correlations between a and b and between c and d . To see these nearly linear relationships, the scatter plots of these two pairs of parameters are presented in Figure 2.14. The reasons for these two relationships differ significantly and are examined separately.

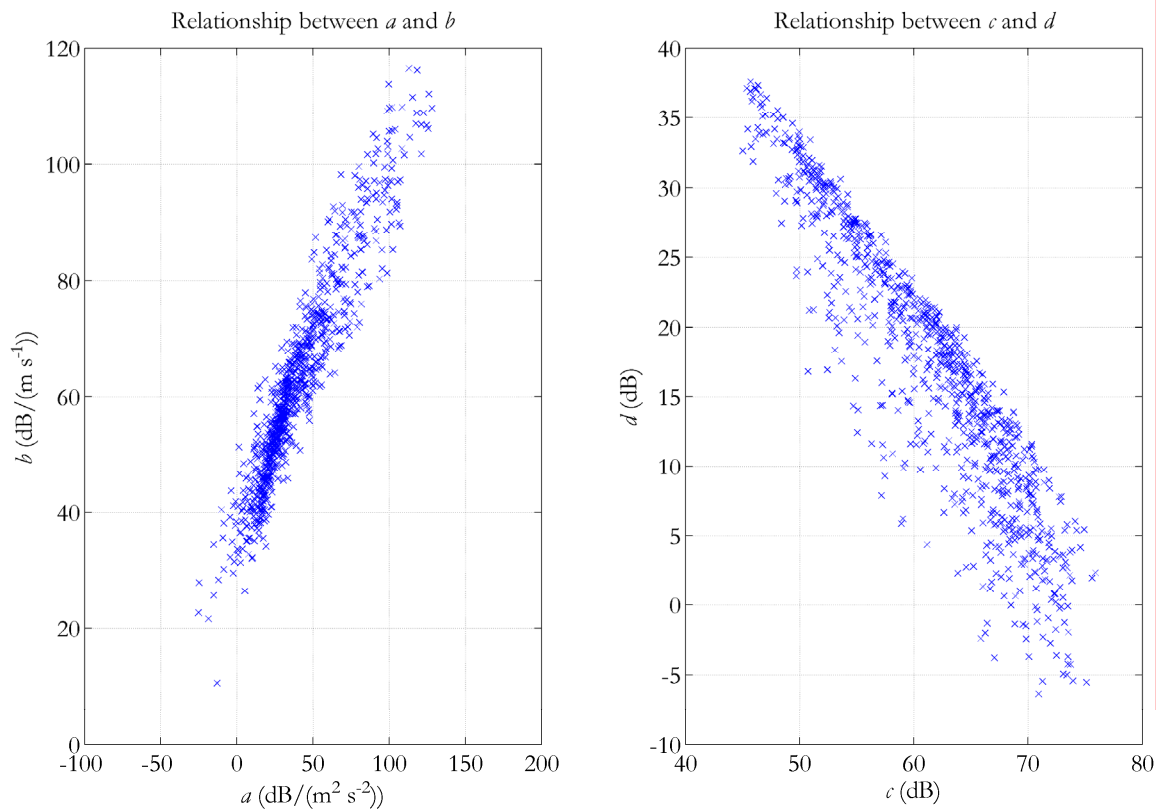


Figure 2.14. Nearly linear relationships between model parameters.

The strong correlation between a and b can be attributed to the similar widths of the AC portions of the spectra. If a line is fit to these values, $b = ma + f$, the b -intercept, f , is 37.5 dB/(m s⁻¹) and the slope, m , is 0.62 m s⁻¹. It can be shown that any pair of (a, b) values that falls on this line pass through the same point given by $(2m, -2mf)$, and f is the slope of the

line from the origin (or peak of the AC part of the spectrum) through this point. The f value can be viewed as the slope of the linear clutter model ($a = 0 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$) that passes through this point. For this particular case, the point is $(1.24 \text{ m s}^{-1}, -46.4 \text{ dB})$. The distance between the top of the AC part of the spectrum and the spectral floor, or c , was estimated to be 61.9 dB. If the AC portion is fit to 15 dB above the spectral floor, this would correspond to -46.9 dB after the spectrum is shifted. The constraint on the model to fit down to this level leads to values of a and b that fall near this line since the widths for the case are similar. The variability on either side of the line is an indicator of the variability of the width of the AC portion of the spectrum 15 dB above the spectral floor. Even though the model restricts the values of a and b to a certain extent in order to ensure that the model extends below the 15 dB level, the curvature of the AC part of the spectrum still seems to be fit reasonably well.

The second strong correlation, the nearly linear relationship between c and d , is related to the spectral floor. This can more easily be observed by examining the histogram of $(c + d)$ which represents the distance from the spectral peak to the spectral floor in dB. Figure 2.15 is a histogram of $(c + d)$ values using the parameters obtained from the MCD method. The effects of system noise have been removed which leaves a system spectral floor that is directly related to the strength of the spectral peak. This spectral floor can most likely be attributed to phase noise and is an inherent part of the radar system. Free and Patel (2005) quantified the phase noise in the WSR-88D system; after the SNR exceeds about 55 dB, the phase noise floor increases one dB for each dB of increase in signal power. This is consistent with the behavior of ground clutter data collected with the PAR. In addition, they describe how the Open Radar Data Acquisition (ORDA) signal processor adjusts the noise floor for censoring using an algorithm based on the amount of clutter power removed through

filtering. For strong clutter signals, the effects of phase noise are significant which is a key justification for including the spectral floor in the model.

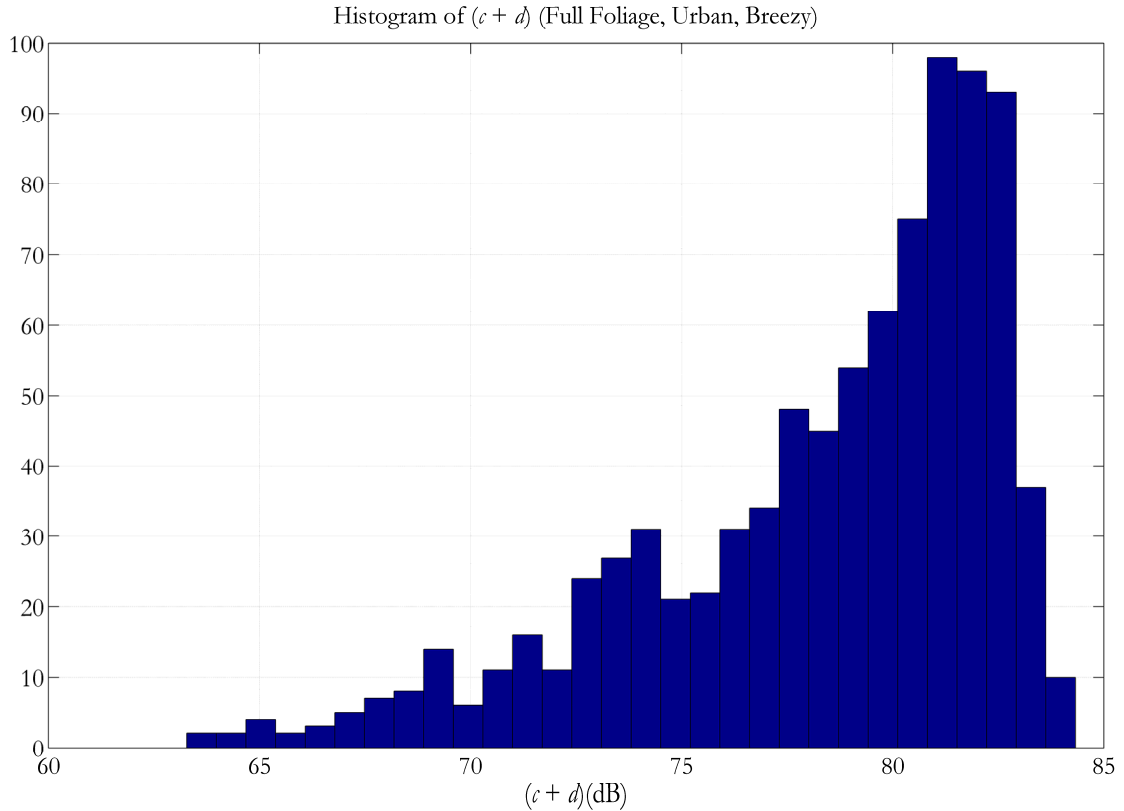


Figure 2.15. Histogram of $(c + d)$ showing the relationship between the spectral peak and spectral floor.

Based on the parameters corresponding to the representative clutter shape, the estimated value of $(\hat{c} + \hat{d})$ is 78.1 dB which is very close to the mean of $(c + d)$ with outliers removed using the MCD technique, 78.4 dB. For these spectra computed from time series data with 4096 samples, the spectral floor is roughly 78 dB below the spectral peak. The estimated d values contains two separate effects that will need to be corrected for later. The first is the windowing effect; the broadening of the spectral peak leads to an underestimation of the “true” d parameter. The second effect is that the value of d depends on the size of the

spectrum. A normalized version of d is developed in Appendix E to remove the dependence on spectral length.

2.6 Clutter Characterization Using the Quadratic Model

As stated earlier, 21 different ground clutter cases were collected with different wind speeds, foliage levels, and terrain types. The objective of this section is to examine the results from the processed spectra using all the cases and investigate the impact of wind, terrain, and foliage conditions on the model parameters. All of the cases were processed successfully except for the case of wooded terrain where the power levels were significantly lower than the other two terrain types. An example of one of the wooded terrain spectra is shown in Figure 2.16.

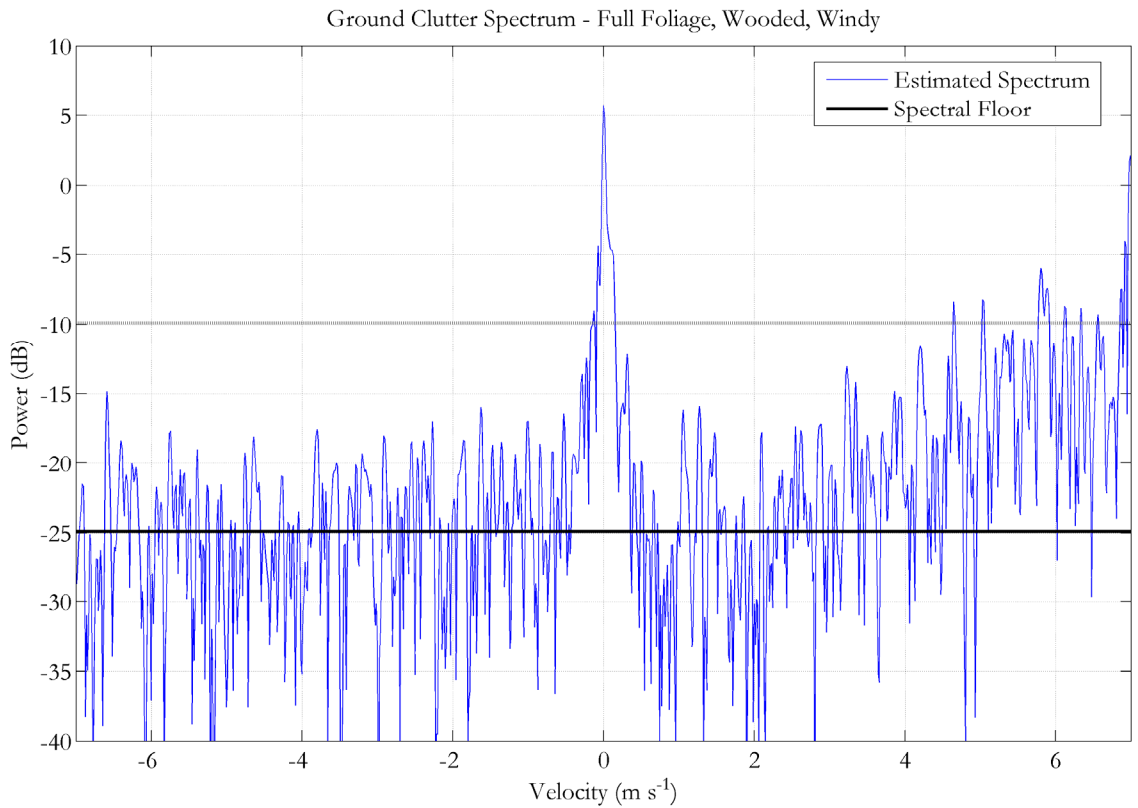


Figure 2.16. Ground clutter spectrum from the “Full Foliage, Wooded, Windy” case with the adjusted spectral floor and 15 dB level.

Only a handful or in some cases no spectra satisfied the constraints imposed when fitting the clutter model. After several attempts at relaxing the constraints in order to include more spectra, the spectra from the wooded area were excluded from processing. Only a small portion of the clutter samples appear above the spectral floor. Even if some of the restraints were relaxed, it would still be difficult to apply the spectral fitting procedure outlined in Section 2.4. Fortunately, all 14 cases from the other two terrain types could be successfully processed.

The parameters corresponding to each of the representative clutter shapes for the 14 processed cases are presented in Table 2.2.

Table 2.2. Estimated quadratic model parameters and f values for 14 ground clutter cases.

Foliage	Terrain	Wind	mph	\hat{a}	\hat{b}	\hat{c}	\hat{d}	\hat{f}
Full Foliage	Urban	Light	4.0	119.8	104.5	56.8	20.3	64.7
		Breezy	9.0	37.2	61.3	61.9	16.2	37.5
		Windy	16.9	9.4	29.4	57.6	15.4	19.2
	Prairie	Light	4.0	245.8	146.7	54.7	25.9	86.9
		Breezy	6.5	50.9	75.1	65.6	12.2	44.4
		Windy	13.0	7.3	30.2	65.7	7.7	19.5
Light Foliage	Urban	Light	2.7	775.6	256.1	54.3	26.0	133.8
		Breezy	9.3	194.9	145.1	62.4	15.9	91.9
		Windy	21.3	27.0	58.3	60.5	12.2	43.1
		Gale	32.5	6.7	27.0	59.5	14.9	19.4
	Prairie	Light	4.0	584.6	227.1	56.3	25.9	124.0
		Breezy	11.0	202.8	154.3	68.6	11.1	87.4
		Windy	19.0	13.2	46.1	68.8	7.6	33.9
		Gale	30.7	3.4	25.7	66.0	5.9	21.3

The foliage levels, terrain types, and wind speed categories are included in the first three columns. The next column shows the wind speed in miles per hour (mph) which was computed from an average of five-minute Oklahoma Mesonet data collected by the Norman Mesonet station (found at <http://www.mesonet.org/>). This station is located within a few

hundred feet of the NWRT. The wind speed for the “Full Foliage, Prairie, Windy” case is slightly under the stated windy range of 15-30 mph but is still far enough from the wind speed of the “Breezy” case to provide useful information. The next four columns contain the estimated model parameters as described in the previous section including the f parameter. This parameter is derived from a linear fit of the a and b parameters and gives a measure of the width of the spectrum. It is similar to the β parameter from Billingsley’s model.

Several trends are clear from the model parameters. The parameters associated with the quadratic fit, a , b , and f , all decrease with increasing wind speed. The decreases in both b and f roughly agree with the behavior of the slope parameter, β , in Billingsley’s model as wind speed increases, but the relationship in Billingsley was derived for wooded terrain (2002). The decreases in the slope parameters reflect a widening of the spectrum as the wind speed increases. The corresponding decrease in the curvature parameter, a , and its relationship to the decreases in b and f are examined in Section 2.8. The d values also seem to decrease as wind speed increases except for a slight increase in the “Light Foliage, Urban, Gale” case. This decrease in d is also expected because power from the DC component is spread into the AC part of the spectrum as the wind increases. The c parameter seems to be negatively correlated with the d parameter as predicted from the spectra studied in the previous section. This relationship is also studied more closely in Section 2.8 where the relationships between the parameters and their dependence on the wind speed are used to extend the quadratic clutter model to capture most of the effects of terrain type and foliage level. This model can then be used to in simulations to faithfully characterize ground clutter data signatures.

The urban and prairie terrain types have very similar clutter shapes overall except the urban terrain has slightly larger values of d especially at higher wind speeds. A possible

explanation is that the urban terrain has more hard targets such as buildings that contribute to the DC or non-varying part of the spectrum. The attribute that makes the most difference is the foliage level. Figure 2.8 has two examples of spectra under breezy conditions, but the “Full Foliage” spectrum is significantly wider even though the wind speed in the “Light Foliage, Prairie, Breezy” case is 11 mph compared to 9 mph for the “Full Foliage, Urban, Breezy” case. Width comparisons can also be made using the f values from Table 2.2. The f values for the “Light Foliage, Gale” cases were similar to the “Full Foliage, Windy” values. The change from light foliage to full foliage seems to increase the width by roughly the same amount as increasing one wind speed category while keeping the foliage level the same.

Spectra from all four of the foliage/terrain-type sets of conditions, “Full Foliage, Urban,” “Full Foliage, Prairie,” “Light Foliage, Urban,” and “Light Foliage, Prairie,” have been shown previously except for the “Full Foliage, Prairie” set. To see the effects of the conditions on real ground clutter spectra, sample spectra and histograms of MSE values for the “Full Foliage, Prairie” set of conditions are presented in Figures 2.17 and 2.18; figures for the other cases are included in Appendix D. The sample spectra in Figure 2.17 were found by minimizing the Euclidean distance between the representative set of parameters from Table 2.2 for each case and the model parameters for each individual spectrum. These spectra exhibit some of the contamination that led to limiting the fitting of the AC part of the model to 15 dB above the spectral floor. The spectra for the four wind speed ranges are normalized so that the peak of the AC portion of the spectrum is at 0 dB, and all of the axes are identically scaled to facilitate comparisons among the spectra. The spectral floor and quadratic fit are also included in the spectra with an additional line denoting the level 15 dB above the spectral floor. The clutter width increases with wind speed as expected, and the DC component and corresponding d parameter decrease with wind speed. The decrease in d

roughly corresponds to an increase in ϵ as predicted by the analysis of the high correlation between them.




Figure 2.17. Sample ground clutter spectra for the “Light Foliage, Prairie” conditions.

(Estimated model parameters: Light - $\hat{a} = 584.0$ dB/(m² s⁻²), $\hat{b} = 240.6$ dB/(m s⁻¹), $\hat{c} = 61.4$ dB, $\hat{d} = 20.4$ dB, Breezy - $\hat{a} = 203.9$ dB/(m² s⁻²), $\hat{b} = 152.7$ dB/(m s⁻¹), $\hat{c} = 71.0$ dB, $\hat{d} = 8.9$ dB, Windy - $\hat{a} = 14.8$ dB/(m² s⁻²), $\hat{b} = 47.7$ dB/(m s⁻¹), $\hat{c} = 68.4$ dB, $\hat{d} = 6.8$ dB, Gale - $\hat{a} = 4.3$ dB/(m² s⁻²), $\hat{b} = 25.0$ dB/(m s⁻¹), $\hat{c} = 65.9$ dB, $\hat{d} = 5.5$ dB.)

The sample spectra illustrate the general effects of wind speed on clutter spectral shapes, but histograms of the MSE for both quadratic and linear fits give some idea of the value of the quadratic fit. Figure 2.18 includes histograms of the MSE for the same cases that were included in Figure 2.17. The numbers in parentheses in the legend correspond to the mean

values of the MSE for both the quadratic and linear fits of the AC part of the spectrum. The value of N in the titles of the subfigures represents the number of spectra that were fit. Recall that only spectra with an MSE less than 45 dB^2 were used to determine the representative set of parameters.

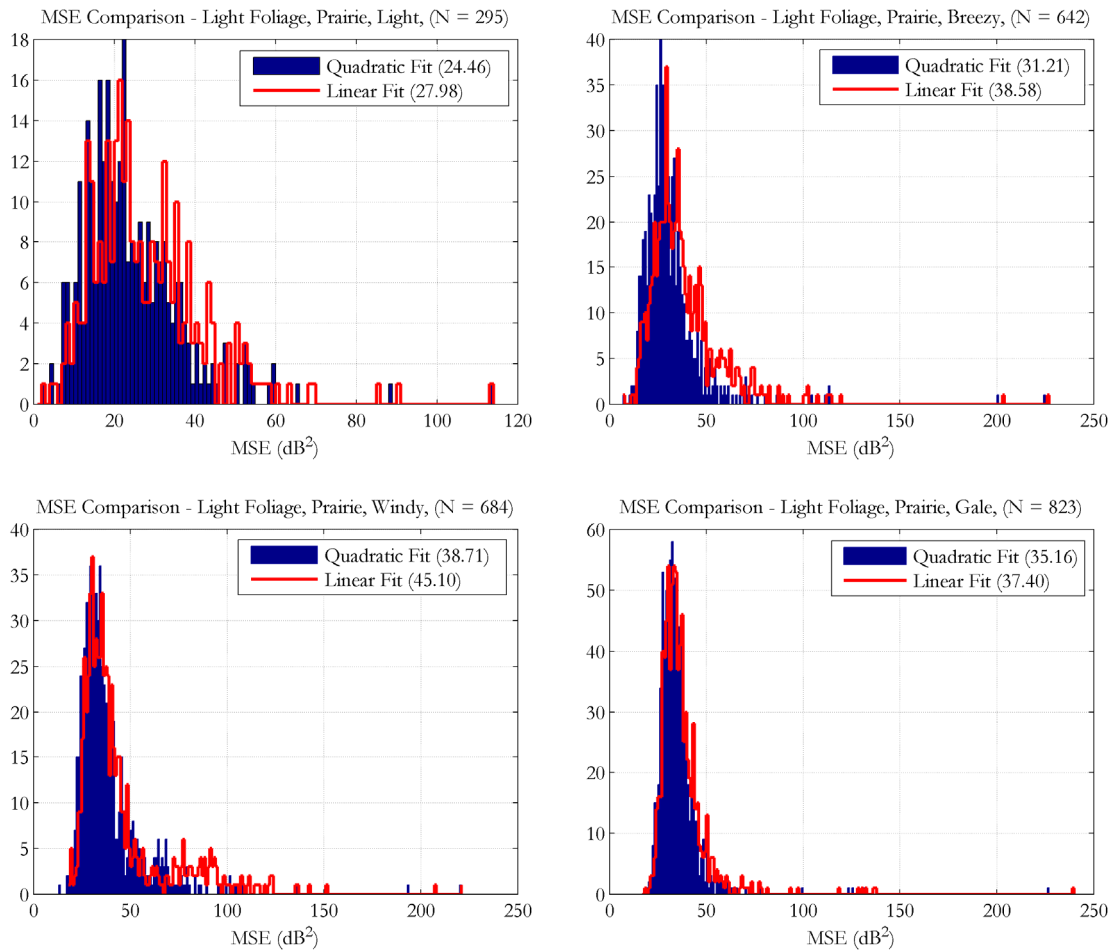


Figure 2.18. Histograms of the MSE for both the quadratic and linear fits for the “Light Foliage, Prairie” conditions.

The difference in the mean MSE values is smaller for the “Light” and “Gale” wind conditions and larger for the “Breezy” and “Windy” conditions. For the “Gale” conditions, the spectra tend to be more linear as can be seen in the sample spectrum from Figure 2.17. The “Light” case also looks relatively linear, but the “Breezy” and “Windy” cases also seem

to be nearly linear based on the sample spectra. After taking all of the cases into account, the “Breezy” conditions lead to the largest discrepancies between the mean MSE for the quadratic and linear models, and the AC portion of the spectra has the most curvature.

Although the parameters in Table 2.2 determine the representative clutter shapes for all 14 cases, it can be difficult to compare the clutter shapes from different conditions. In Figure 2.19, the shapes corresponding to a particular foliage level and terrain type are shown in their own subfigure window.

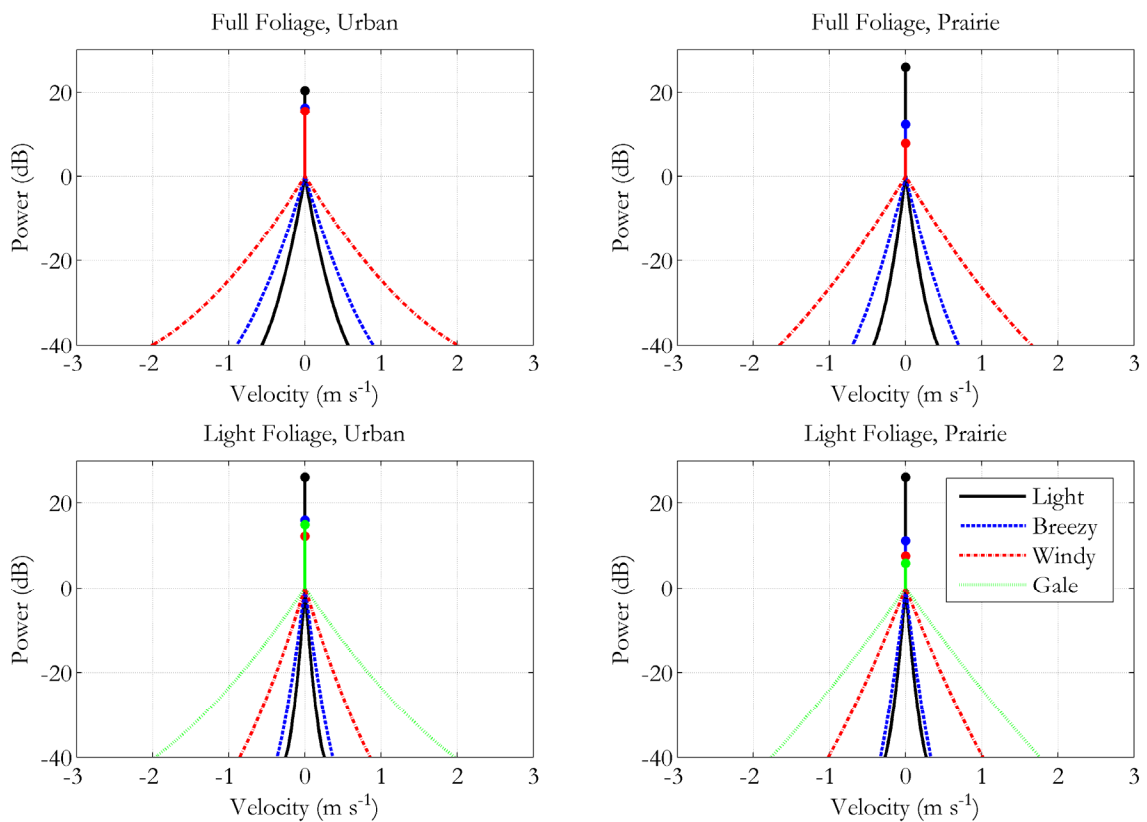


Figure 2.19. Representative clutter shapes for all 14 cases showing both the AC and DC parts of the spectral shapes.

The spectral floors are not included so that the widths of the AC parts of the spectra are clearly depicted. Many of the trends in the parameters that were described earlier are now

more apparent. The “Light Foliage, Gale” shapes are very similar in width to the “Full Foliage, Windy” ones. The d values at higher wind speeds are smaller for the “Prairie” terrain than the “Urban” terrain but are larger for the “Light” wind conditions. The discrepancy in “Light” conditions may come from the difficulty of fitting such narrow spectra. Now that a representative clutter shape has been computed for all of the cases, the effects of wind speed, foliage level, and terrain can be modeled. Before this modeling is undertaken in Section 2.8, the effects of beam broadening are addressed.

2.7 Effects of Beam Broadening

The objective of the section is to examine the effects of beam broadening on the clutter shape. In Chapter 1, beam broadening was introduced as one of the unique attributes of phased array antennas. The possible effects of beam broadening on ground clutter need to be considered. As stated in Section 2.2, nine 90° sectors centered at azimuths of 0° , 355° , 350° , 345° , 340° , 335° , 330° , 325° , and 320° were collected. Because the centers of these sectors are successively shifted by 5° , a 5° sector can be chosen where all nine of the 90° sectors overlap, but the beamwidths vary for each sector. For example, a 5° sector centered at an azimuth of 2° would be near the center of the 90° sector centered at 0° but very close to the edge of the 90° sector centered at 320° . By examining the data from the different 90° sectors, the effects of beam broadening can be studied. The data were collected when the wind was relatively stable over the collection period with the average wind speeds varying from 8 to 13 mph with a mean value close to 10 mph. An additional 5° sector was chosen with a center at 317° which also contains data from all 9 of the 90° sectors. To increase the number of cases, two 5° sectors with 8 overlapping 90° sectors and two more with 7 overlapping 90° sectors were considered. These four 5° sectors have centers at 7° , 312° , 12° ,

and 307° , respectively. This gives a total of six cases to study with beamwidths at the center of each 5° sector varying from approximately 1.5° to slightly larger than 2.05° . The beamwidth was estimated using the equation from Chapter 1. Because of minor pointing discrepancies in the NWRT pedestal, the actual pointing directions differ by about 0.2° from the commanded pointing directions. For each case, the same fitting procedure described in Section 2.4 was applied to the overlapping portion of each of the 90° sectors and a representative set of parameters was computed using the L1-median.

In addition to the broadening of the beam, the gain of the antenna also decreases as the beam is steered away from broadside. This should lead to a decrease in the power of clutter targets near the edges of the 90° sectors, but the beam broadening also leads to more clutter targets falling in the beam. To study these competing effects, the number of clutter spectra that passed all of the thresholds for inclusion in the model was plotted against the beamwidth for all six cases in Figure 2.20. Overall, in five of the six cases the total number of included spectra did decrease at the largest beamwidth compared to the smallest beamwidth, but the effects were not especially significant. In only half of the cases was the number of spectra smallest for the largest beamwidth. The effects of reduced gain and beam broadening seemed to nearly cancel out with the decreased gain having a somewhat larger effect than the beam broadening. The number of spectra included in the cases centered at 316.8° , 311.8° and 306.8° was lower than the number of spectra included in the cases from 1.8° , 6.8° and 11.8° because the amount of urban ground clutter is greatest to the north of the NWRT and decreases as the azimuth departs from this. This leads to only about 60 to 90 spectra being included in the case centered at 306.8° which is the case farthest from the high clutter areas north of the radar.

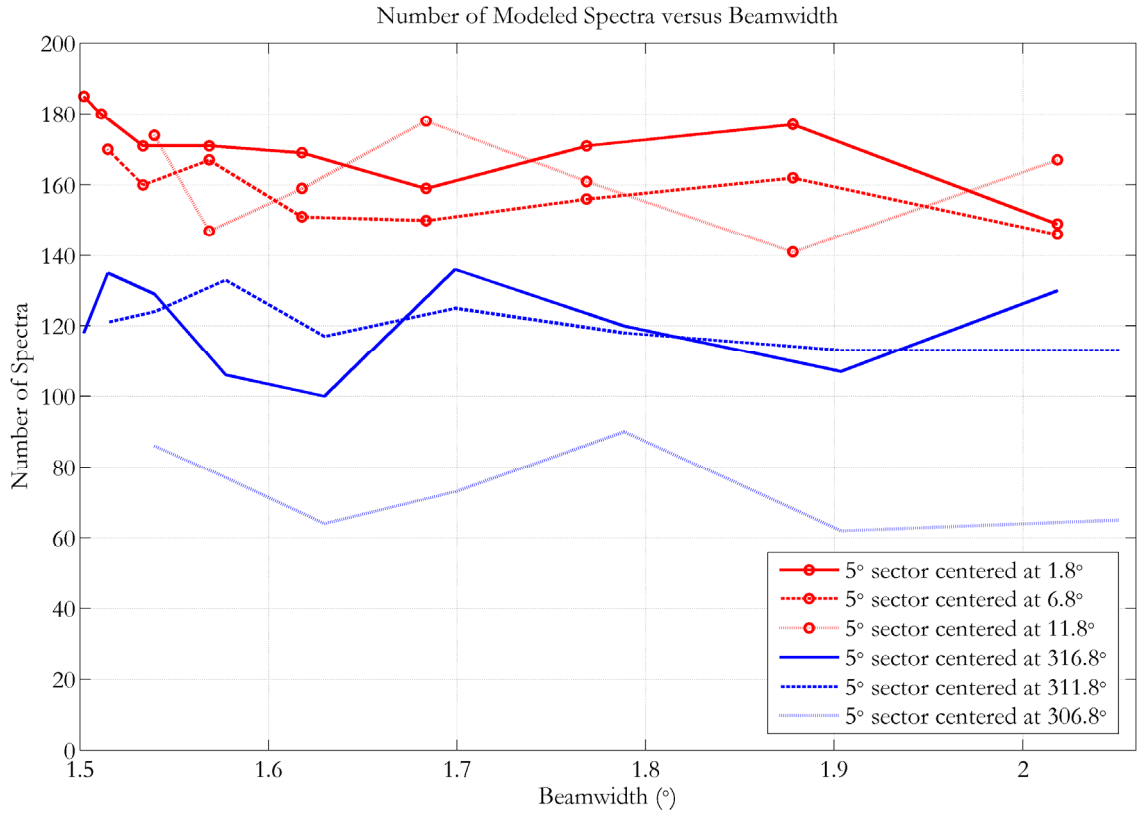


Figure 2.20. Number of spectra used for modeling from all six cases plotted versus the beamwidth at the center of the shared 5 ° sector.

The next figure, Figure 2.21, illustrates the beamwidth dependence of the a , b , c , and d parameters computed using the quadratic clutter model. The c parameter is the most stable as the beamwidth varies with the d parameter a close second. The a and b parameters appear to fluctuate more, but the fluctuations look random without a strong dependence on the beamwidth. Overall, the results lead to a couple of conclusions: the procedure used to choose the representative parameters is stable, and the effects of beam broadening on the shapes of the clutter spectra are not significant.

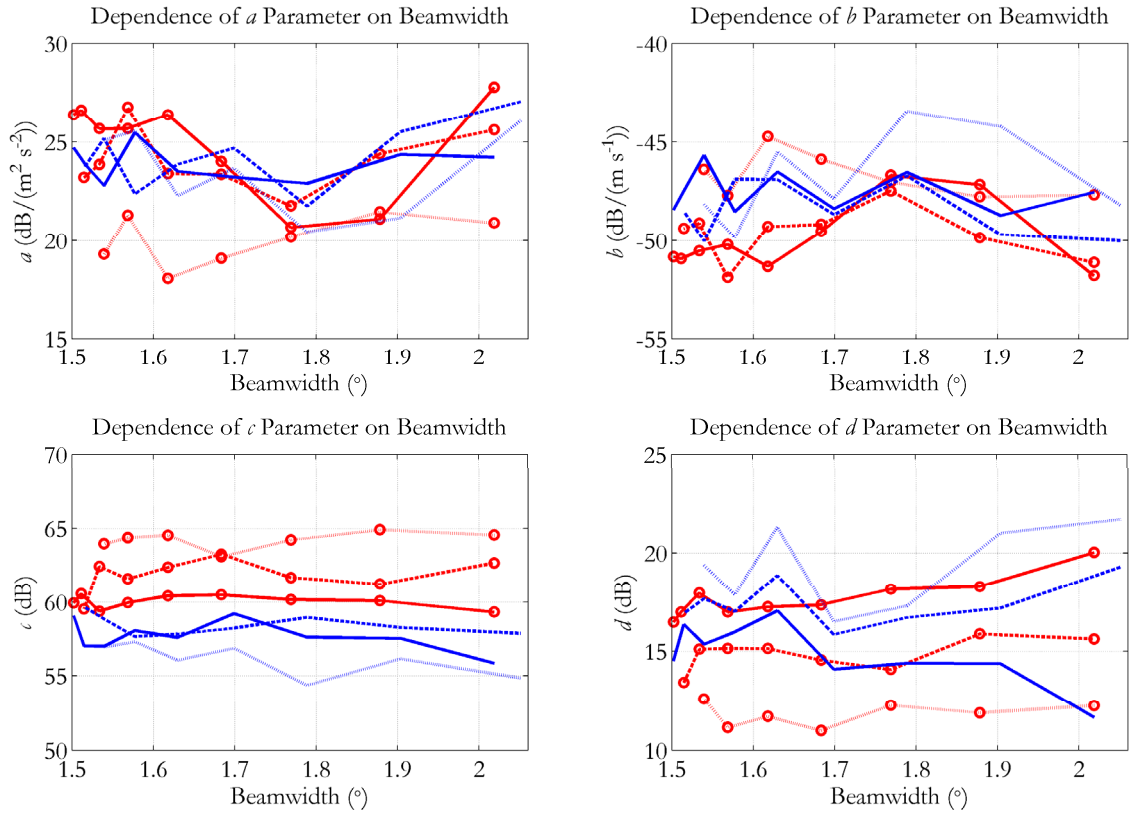


Figure 2.21. Dependence of quadratic clutter model parameters on beamwidth.

2.8 Modeling the Effects of Wind Speed, Foliage, and Terrain

In Section 2.6, several general trends were discussed based on the representative clutter shapes of the cases. For example, all of the parameters tend to decrease as the wind speed increases. These trends can be modeled and utilized to extend the quadratic clutter model to capture the change in the shape of the clutter spectrum with changing wind speed. By taking advantage of some of the relationships between the parameters, a relatively simple model can be developed. The similarities between clutter spectra because of foliage level and terrain type can also be exploited. In this section, the quadratic clutter model is extended based on the wind speed trends, relationships between the parameters, and the effects of foliage and terrain type.

The a , b , c , and f parameters from Section 2.6 and a normalized version of the estimated d parameter are used to construct the extended quadratic clutter model. As was already mentioned, the d parameter measured from the windowed spectrum needs to be corrected because of broadening and the fact that d varies with the length of the spectrum. A relationship between a normalized version of the parameter, d_n , and the estimated d is developed in Appendix E along with an equation to find the spectrum-length dependent values of d_N based on the normalized d_n where N is the length of the spectrum. The two equations used to find d_n are given below:

$$d_w = 10 \log_{10} \left(P_N \left(10^{\frac{d}{10}} - 1 \right) + 1 \right) \quad d_n = \left(0.4 d_w^{1.3} + 2 d_w \right)^{\frac{1}{1.3}}, \quad (2.6)$$

where d is the value estimated from the spectrum, d_w is the normalized version of d with window effects, and d_n is corrected for window effects. P_N is the power in the zero-velocity bin from the AC part of the spectrum and is discussed further in Appendix E. The normalized value is smaller than the value estimated from the length-4096 spectrum, and both values are shown in Table 2.3. The quadratic clutter model will be based on d_n so that the model is not tied to one particular spectrum length.

The first relationship to be examined is the one between the c and d_n parameters. As noted in Section 2.5, the value of $(c + d)$ seems to be relatively constant at least within a particular case. This should also carry over to the normalized d_n . The estimate of $(\hat{c} + \hat{d})$ over all 14 cases varies from 60.5 dB to 69.9 dB with a mean of 65.3 dB. Although there is some variation in this value over the 14 cases, using the mean value of 65.3 dB is a straightforward way to capture the strong correlation between c and d_n . Because d_n is normalized, the distance between the peak of the spectrum and the spectral floor depends on the number of elements

in the spectrum. For a particular spectrum size, the distance is given by $(c + d_N)$. The equation for d_N in terms of d_n that was derived in Appendix E is given below:

$$d_N = 10 \log_{10} \left(\frac{1}{P_N} \left(10^{\frac{d_n}{10}} - 1 \right) + 1 \right) \quad (2.7)$$

For example, the distance between the peak and the floor when $N = 4096$ is about 80.2 dB if the mean value of 65.3 dB is used for $(c + d_n)$. This is close to the value of $(\hat{c} + \hat{d}) = 78.1$ dB for the “Full Foliage, Urban, Breezy” case that was mentioned in Section 2.5. This computation of the spectral floor due to phase noise needs to be made independently for each individual radar since it depends on the transmitter and other parts of the system.

Table 2.3. Normalized d_n and estimated c and d for all 14 collected cases.

Foliage	Terrain	Wind	mph	\hat{c}	\hat{d}	\hat{d}_n
Full Foliage	Urban	Light	4.0	56.8	20.3	5.8
		Breezy	9.0	61.9	16.2	3.3
		Windy	16.9	57.6	15.4	3.0
	Prairie	Light	4.0	54.7	25.9	10.4
		Breezy	6.5	65.6	12.2	1.7
		Windy	13.0	65.7	7.7	0.7
Light Foliage	Urban	Light	2.7	54.3	26.0	10.2
		Breezy	9.3	62.4	15.9	3.0
		Windy	21.3	60.5	12.2	1.7
		Gale	32.5	59.5	14.9	2.7
	Prairie	Light	4.0	56.3	25.9	10.2
		Breezy	11.0	68.6	11.1	1.3
		Windy	19.0	68.8	7.6	0.7
		Gale	30.7	66.0	5.9	0.5

Because of the strong correlation between c and d_n , finding a model for either parameter will determine both. It seems more natural to look at the relationship between d_n and wind speed because the spectral floor varies with the spectral peak or DC term. Figure 2.22 shows the relationship along with a shifted exponential fit for the urban and prairie terrain types.

The shifted exponential is defined as a constant added to a multiple of an exponential function.

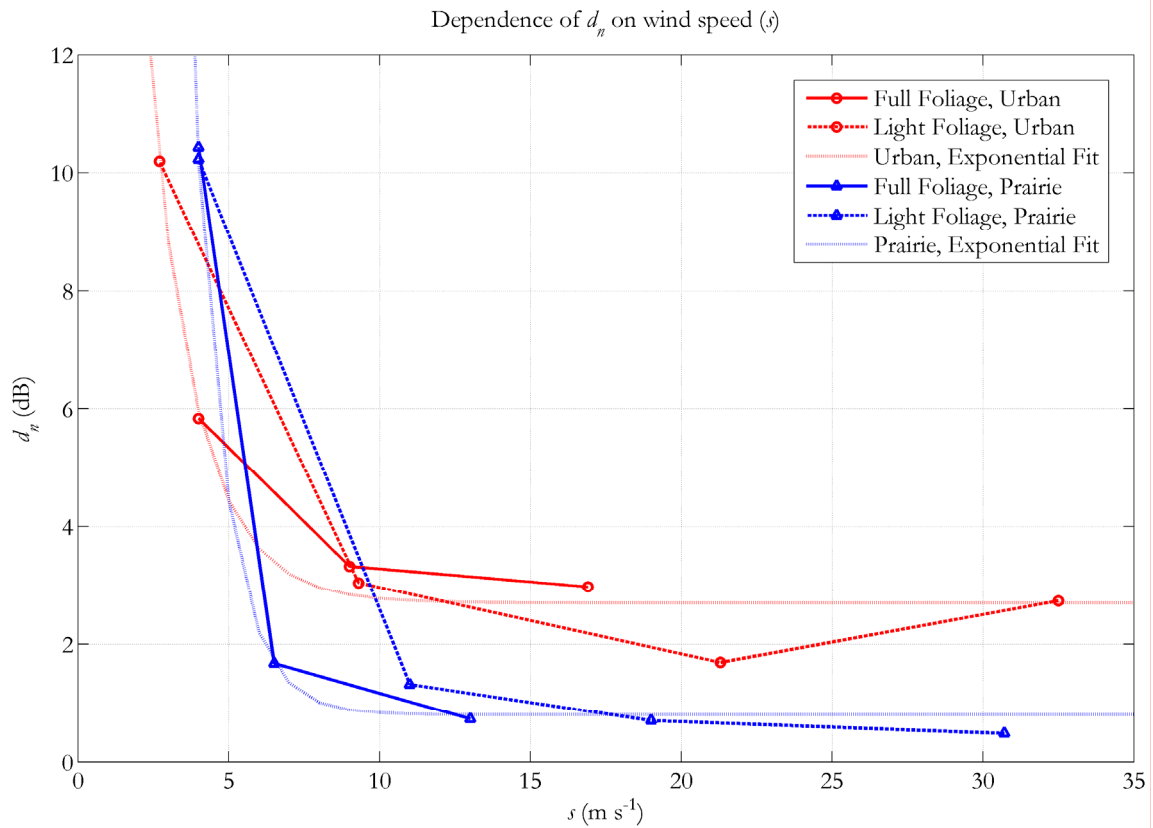


Figure 2.22. Relationship between d_n and wind speed, s .

The shifted exponential fit captures the overall trend in d_n , but there are some discrepancies. The largest errors are on the order of 1 dB which is not extremely significant. The most interesting finding is that both urban cases are similar and both prairie cases are similar with a lower value of d_n for the prairie case at high wind speeds. The reason that the d_n values at low wind speeds are larger in the prairie cases is not apparent. As discussed earlier, the general argument for larger values in the urban case is that urban clutter has more hard targets such as buildings that drive up the power in the DC component. This argument matches the results at high wind speeds, but not at low wind speeds. This is another situation

where more data collected at very low wind speeds could help sort out what is occurring. In general, the detailed behavior at low wind speeds is not as significant as the behavior at high wind speeds because so much of the power is limited to a small region around DC. The ground clutter spectra at low wind speeds are very similar to DC spikes or delta functions.

The final expressions for d_n in terms of s for both urban and prairie terrain are given as follows:

$$d_n(s) = \begin{cases} 40.4 \exp(-0.63s) + 2.7 & s = 1 \text{ to } 35 \text{ mph, Urban} \\ 417.7 \exp(-0.95s) + 0.8 & s = 1 \text{ to } 35 \text{ mph, Prairie} \end{cases} \quad (2.8)$$

Non-linear least squares minimization was used to find the coefficients of the shifted exponential functions. Given a d_n value, an estimate of c for the model can be found by subtracting the mean value of 65.3 dB that was calculated earlier.

The second relationship to be studied is the one between a and b . For the DC part of the model, d_n was modeled because it is directly related to the spectral floor. For the AC part of the spectrum, the most important part of the model also needs to be captured. After trying several different approaches for modeling a and b , one that is based on the width of the AC portion of the spectrum was chosen. When processing data from the “Full Foliage, Urban, Breezy” case, a strong relationship between a and b values was discovered that determined a particular point that most of the curves closely approached. This point, $(2m, -2mf)$ where m is the slope and f is the intercept of the linear least squares fit of a and b , captures the basic width of the AC part of the spectrum. The point could be used in the model directly, but another approach is more useful for producing both a quadratic model for ground clutter and a simpler linear model which is closely related to Billingsley’s exponential model from equation (2.2). The intercept of the linear least squares fit, f , is also the magnitude of the slope of the line that passes through the peak of the AC part of the spectrum (located at the

origin) and the point, $(2m, -2mf)$. The f parameter will be modeled based on the data from the 14 ground clutter cases and will then be used to compute a and b for the quadratic model or used directly for the simpler linear model. To more clearly understand what f represents, Figure 2.23 shows the quadratic fit of the AC part of the representative spectrum for the “Full Foliage, Urban, Breezy” case along with the point, $(2m, -2mf)$, and the line passing through it and the origin. The points, denoted by black “x” markers, at $(\pm 2m, -2mf)$ clearly fall on the linear model but do not fall exactly on the curve representing the quadratic model. This occurs because f and m represent a linear least squares fit for all of the spectra in the case rather than just for the representative quadratic model.

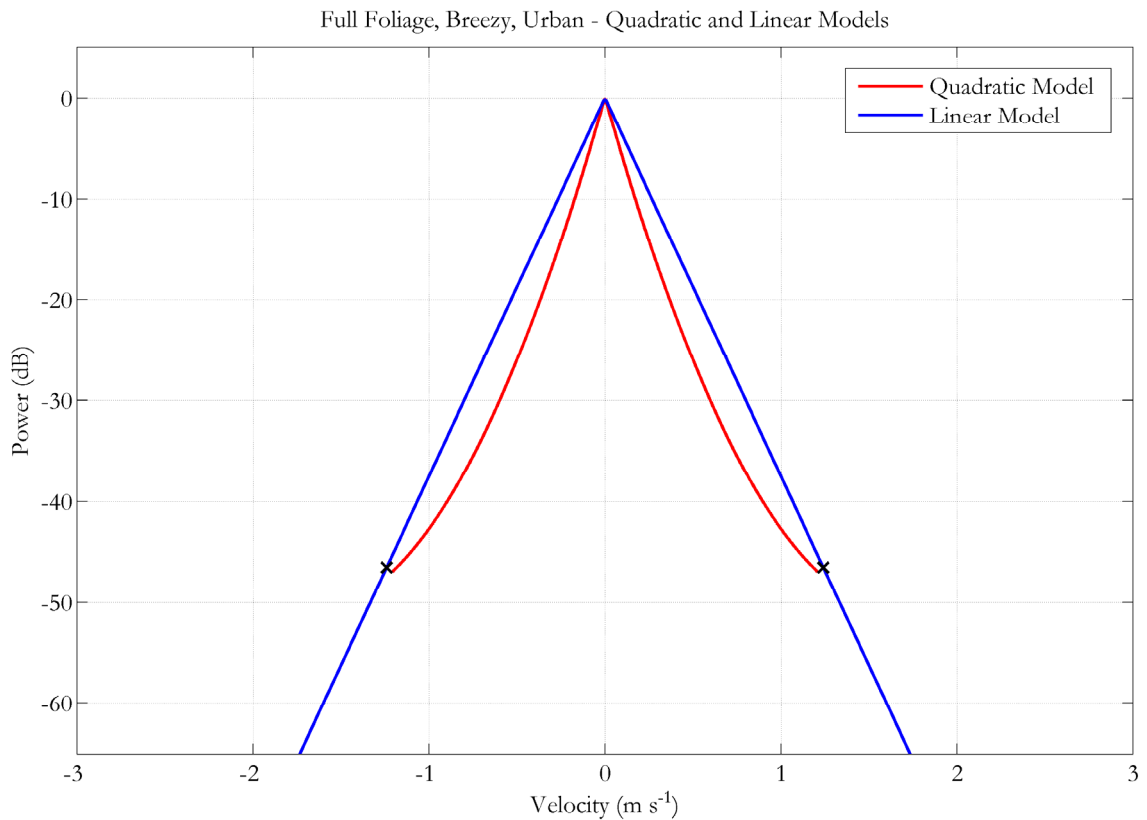


Figure 2.23. Comparison of quadratic and linear models for the “Full Foliage, Breezy, Urban” case.

The next step is to find a relationship between the wind speed, s , and the estimate of the slope, f , for the linear model. After examining the relationship between f and wind speed, f seems to have a linear dependence on the logarithm of the wind speed. Figure 2.24 illustrates this relationship. The “Light Foliage, Urban” cases line up very well with a linear fit of all of the “Light Foliage” data. The “Full Foliage” cases have smaller f values in general which corresponds to wider spectra given the same wind speed. This is consistent with more moving parts (leaves) in the “Full Foliage” case which leads to more power in the AC part of the spectrum. The range of wind speeds fit in the “Full Foliage” case is from 1 to 17 mph in order to match the range of wind speeds in the data. The “Light Foliage” linear fit extends from 1 to 33 mph because of the additional cases collected over 30 mph.

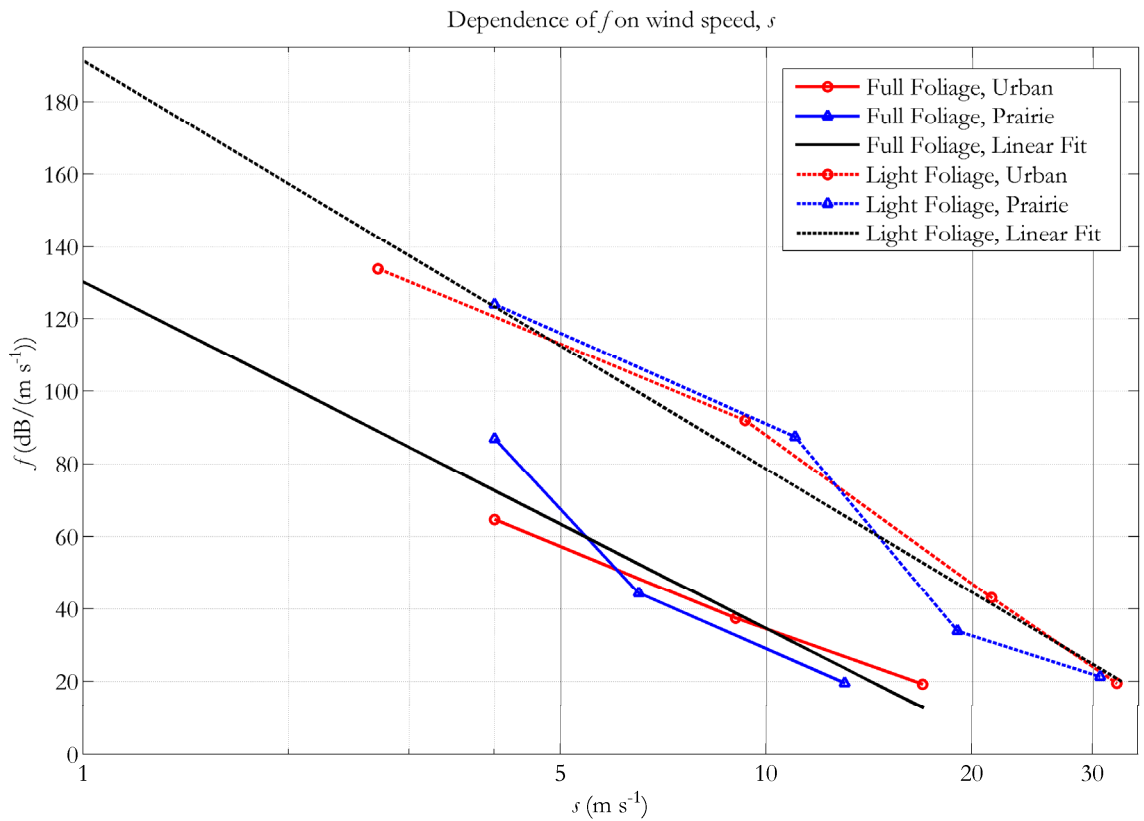


Figure 2.24. Relationship between f and wind speed, s .

The final linear fit values which can be used for the model are given below:

$$f(s) = \begin{cases} 130.3 - 41.6 \cdot \log(s) & s = 1 \text{ to } 17 \text{ mph, Full Foliage} \\ 191.3 - 49.0 \cdot \log(s) & s = 1 \text{ to } 33 \text{ mph, Light Foliage} \end{cases} \quad (2.9)$$

The difference between the intercept values, 130.3 and 191.3, clearly illustrates a strong difference between the light and full foliage conditions. The slope values, -41.6 and -49, are similar, but there does seem to be a slight change in slope between the light and full foliage conditions.

Now that a formula for f in terms of the wind speed has been determined, the next step is to find a and b . In the case of the b parameter, the behavior is very similar to that of f . In fact, an independent linear fit of b in terms of the log of the wind speed could be used for the model, but this also implies that b should be nearly linearly related to f . For this model, b was defined in terms of f with the relationship shown in Figure 2.25. Although the overall shape seems very linear, small variations in the fit can have large consequences when b and f are similar in size. In most cases, b is larger or steeper than f which implies a positive value for a . This particular linear model also predicts that in some cases b can be smaller than f which corresponds to a negative value of a . In the data from the “Full Foliage, Breezy, Urban” case, there are a handful of negative a values for individual spectra, but there are significantly more when the representative a value is smaller such as in the “Full Foliage, Windy” cases and the “Light Foliage, Gale” cases. Additional data collection would be very helpful in studying the behavior of b and f in these cases, and collecting data with “Full Foliage” at wind speeds between 20 and 30 mph would be more feasible than with “Light Foliage” at wind speeds greater than 35 mph. The formula corresponding to the linear least squares fit of b in terms of f is defined as follows: $b = 1.92f - 14.2$. When $a = 0$, then $b = f$

which occurs at the following value: $b = f = 15.4 \text{ dB}/(\text{m s}^{-1})$. So for $f > 15.4 \text{ dB}/(\text{m s}^{-1})$, a is positive and for $f < 15.4 \text{ dB}/(\text{m s}^{-1})$, a is negative.




Figure 2.25. Relationship between the two slope parameters b and f .

Given a particular wind speed, b , c , d_n , and f are defined. To complete the clutter model, a formula for a needs to be found. As mentioned earlier, nearly all of the curves for a particular case pass near the point, $(2m, -2mf)$. This point is also on the linear model which has slope f . To find a , a point on this line needs to be chosen so that the curve passes through the point. The exact location of $(2m, -2mf)$ is not known, but $-2mf$ should be close to 15dB above the spectral floor. The point (v_0, P_0) can be defined as the point on the linear model that is 15 dB above the spectral floor or $P_0 = 15 - c$. Thus (v_0, P_0) is on the line from the origin which has slope f , and $v_0 = -P_0/f$. Once (v_0, P_0) is known and it is assumed that

the quadratic part of the model passes through this point, then a formula for a can be found in terms of b, f and P_0 . This completes the extended version of the quadratic clutter model which is defined as follows:

$$\begin{aligned} d_n(s) &= d_0 \exp(d_1 s) + d_2 & b(s) &= b_1 f + b_0 \\ c(d) &= c_0 - d_n & P_0(c) &= q_0 - c \\ f(s) &= f_0 + f_1 \log(s) & a(b, f, P_0) &= \frac{2f(f-b)}{P_0} \end{aligned} \quad (2.10)$$

The formula for a shows that it is a quadratic function of f (or b). All of the lowercase subscripted variables ($b_0, b_1, c_0, d_0, d_1, d_2, f_0, f_1$, and q_0) are constants. Using equation (2.10), all of the parameters for the model can now be found given the wind speed, s .

Table 2.4 contains values for the constants that were found by fitting the ground clutter data. Some of the constants such as b_0, b_1, c_0 , and q_0 stay the same for all of the cases. The parameters that determine d_n (d_0, d_1 , and d_2) depend on the terrain type while the parameters that determine f (f_0 and f_1) depend on the foliage level.

Table 2.4. Constant values for quadratic clutter model.

Foliage	Terrain	b_0	b_1	c_0	d_0	d_1	d_2	f_0	f_1	q_0	s (mph)
Full Foliage	Urban	14.2	1.92	65.3	40.4	-0.63	2.7	130.3	-41.6	15	1-17
	Prairie	14.2	1.92	65.3	417.7	-0.95	0.8	130.3	-41.6	15	1-17
Light Foliage	Urban	14.2	1.92	65.3	40.4	-0.63	2.7	191.3	-49.0	15	1-33
	Prairie	14.2	1.92	65.3	417.7	-0.95	0.8	191.3	-49.0	15	1-33

The quadratic clutter model utilizes the dependence of f and d_n on wind speed, and all of the other parameters depend either directly or indirectly on those two. Alternative values for f_0 and f_1 could be found in other foliage situations; values between those found in the table might match conditions between full and light foliage. The c_0 value is radar-dependent and captures a significant amount of the dependence of c on d_n . If a new value of c_0 is used, the q_0

value should be adjusted to that P_0 stays constant. This ensures that the part of the spectrum modeled with the quadratic term stays the same. For example, if the new c_0 was increased by 5 dB, q_0 would also have to be increased by 5 dB. The d_0 , d_1 , and d_2 values result in a reasonable fit of d_n , but additional data could lead to values for more terrain types or even a better relationship between d_n and wind speed. The simpler linear model is completely determined by f and d_n (along with the constant parameter, c_0).

2.9 Completing the Quadratic Clutter Model

The quadratic clutter model was used to characterize collected data in Section 2.6, and the model was extended to reflect the effects of wind speed, foliage, and terrain in Section 2.8. The next step is to complete the model so that it can be used to study clutter width and simulate time series data. As shown in Figure 2.9, the AC portion of the spectrum is only fit to a level 15 dB above the spectral floor. This leads to a hole in the spectrum that needs to be filled, but the behavior of the spectrum is unknown in this region. In this study, the quadratic portion of the model will be extended linearly to fill the unknown part of the spectrum. This results in the same decay at large velocities as Billingsley's model but is more accurate for the part of the spectrum that was characterized.

The final model is similar to equation (2.3) except that the DC term is quantified using the normalized version of the d parameter, d_n .

$$P_{dB}(v) = d_n \cdot \delta(v) + P_{AC}(v) \quad (2.11)$$

A clutter spectrum based on the quadratic clutter model is already defined from the peak of the AC portion down to $(\pm v_0, P_0)$. The simplest way for the spectrum to be continuous and smooth at (v_0, P_0) is to extend the model in a linear fashion with a slope equal to the quadratic model slope at (v_0, P_0) . The AC part of the model is defined as follows:

$$P_{AC}(v) = \begin{cases} -b|v| + \frac{a}{2}v^2 & \text{when } |v| < v_0 \\ S(|v| - v_0) + P_0 & \text{when } |v| \geq v_0 \end{cases} \quad (2.12)$$

where S is the slope at (v_0, P_0) . The value of the slope is calculated as $S = -b + av_0$. The quadratic model is mainly concerned with the spectral shape so it is still assumed that the spectra are normalized so the peak of the AC portion is at zero.

Figure 2.1 shows an example of the linear extension of the quadratic model including the spectral floor and the level 15 dB above the spectral floor.

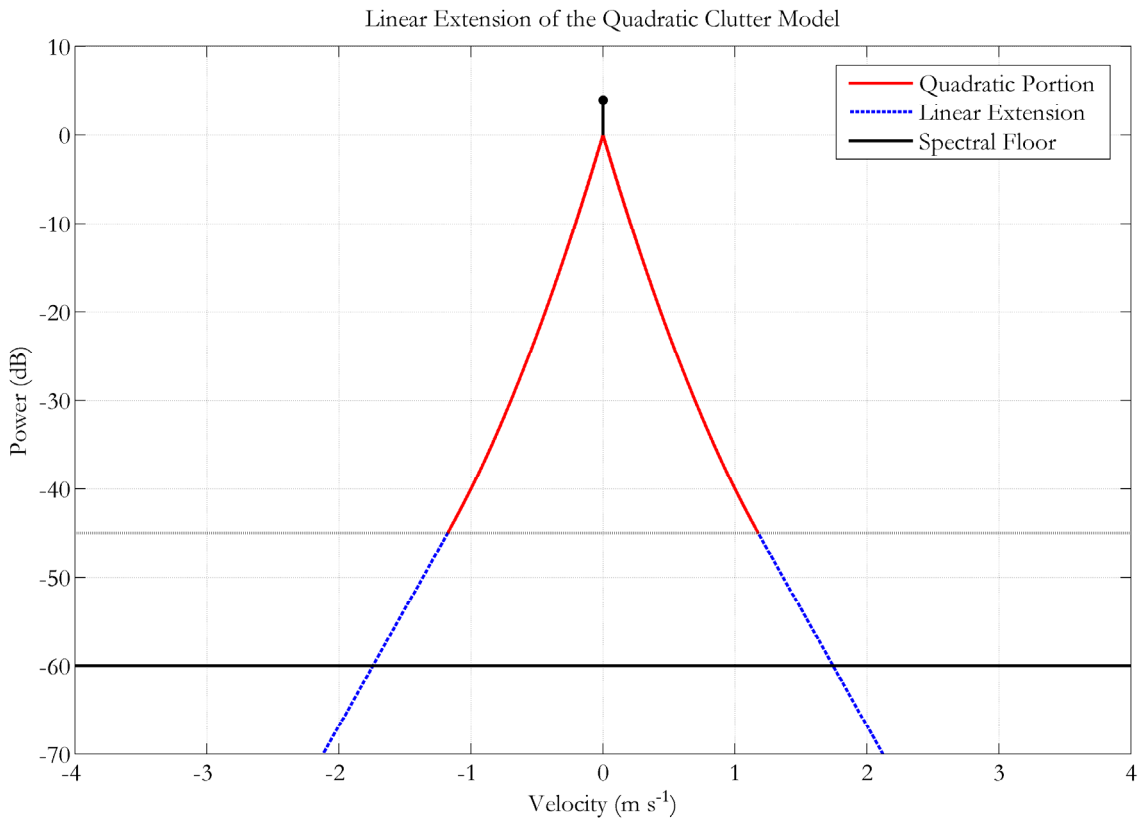


Figure 2.26. Example of linear extension of the quadratic clutter model.

As discussed when the model was first introduced, the spectral floor is not included in equation (2.11) because it cannot be added directly in the log domain. Instead, the spectral

floor is added as white noise in the linear domain. The noise level is set ϵ_0 dB below the peak of the DC term. The spectral floor is radar-dependent and is not part of the underlying clutter spectrum. In an ideal system without phase noise or quantization errors, there would be no spectral floor, and the power would continue to decrease at large velocities.

2.10 Using the Quadratic Clutter Model to Study Clutter Width

In this section, two different measures of the clutter width are examined. The first is the spectrum width, σ_v , which is defined as the standard deviation of the spectrum. This is a standard measure of width and will be used to observe the dependence of the clutter width on wind speed. The second width measure, w , is introduced to study the clutter width in terms of an ideal ground clutter filter. A simple spectral ground clutter filter can be implemented by notching the spectrum around zero velocity; w will be defined in terms of the clutter power removed by an ideal version of this notch filter. The quadratic and linear models will be compared to the Gaussian clutter model which is commonly used to model ground clutter spectra.

The spectrum width, σ_v , is defined as the square root of the variance of the spectrum. Because the clutter spectrum is assumed to be symmetric about zero velocity, the mean velocity is zero and the variance can be written as follows (Raghavan 2003):

$$\sigma_v^2 = \frac{\int v^2 S(v) dv}{\int S(v) dv} \quad (2.13)$$

where v is the velocity and $S(v)$ is the power spectrum. To calculate the spectrum width, the limits of integration need to be determined. The velocity could be limited to the Nyquist interval, but the underlying clutter shape is independent of the PRT used when collecting data. The velocity could also be limited to the portion of the spectrum above the spectral

floor, but the spectral floor is radar dependent. In this case, the velocity was integrated from $-\infty$ to ∞ with no spectral floor. This results in a conservative estimate of the spectrum width that may be smaller in practice depending on the treatment of the spectral floor.

Before computing the spectrum width using the quadratic clutter model, the power spectrum needs to be converted from the log domain to the linear domain for integration. The resulting power spectrum is derived from equation (2.11) and is presented below:

$$S(\nu) = \begin{cases} \delta(\nu) \left(10^{\frac{d}{10}} - 1 \right) + 10^{\frac{-b|\nu| + \frac{a}{2}\nu^2}{10}} & \text{when } |\nu| \leq \nu_0 . \\ 10^{\delta(|\nu| - \nu_0) + P_0} & \text{when } |\nu| > \nu_0 \end{cases} \quad (2.14)$$

For $|\nu| \leq \nu_0$, the first term includes all of the power from the DC component that extends above the AC part of the spectrum, and the second term includes the power from the AC part of the spectrum. For $|\nu| > \nu_0$, the term includes the power from the AC part of the spectrum for the linear extension of the model beyond ν_0 . Clearly, the integral of the DC part is $10^{\frac{d}{10}} - 1$ because of the delta function. The integral of the AC part for $|\nu| \leq \nu_0$ can be found using equation (E.2) from Appendix E while the integral for $|\nu| > \nu_0$ is a straightforward integral of an exponential. For the following computations, numerical integration was used for the integral of the spectrum and for the spectrum width.

Because the linear model is less complicated than the quadratic model, there are closed form solutions for most of the integrals. The power spectrum is similar to equation (2.14) except that b is replaced with f and the transition from the quadratic part of the model to the linear extension is not needed:

$$S(\nu) = \delta(\nu) \left(10^{\frac{d}{10}} - 1 \right) + 10^{\frac{-f|\nu|}{10}} . \quad (2.15)$$

After substituting this spectrum into equation (2.13) and integrating from $-\infty$ to ∞ , the spectrum width for the linear model can be written in terms of f and d_r :

$$\sigma_v = \sqrt{\frac{4\left(\frac{10}{f \log(10)}\right)^3}{\left(10^{\frac{d_r}{10}} - 1\right) + 2\left(\frac{10}{f \log(10)}\right)}}. \quad (2.16)$$

If the DC component is zero, the spectrum width is inversely related to f . As f increases and the AC portion of the spectrum gets steeper, the width decreases as expected.

With an expression for the spectrum width for the linear model and an expression for the spectrum for the quadratic model, both spectrum widths for all 14 cases were calculated. These values are collected in Table 2.5 along with the values for ν . The dependence of spectrum width on wind speed can also be examined using the full quadratic and linear models. Figure 2.27 shows the spectrum widths calculated for each of the 14 representative clutter spectra along with the spectrum widths for the linear and quadratic models. For most wind speeds, the spectrum width is less than 0.3 m s^{-1} , and all of the spectrum widths from collected data are smaller than 0.3 m s^{-1} . Currently, the GMAP clutter filter used on the WSR-88D uses a spectrum width of 0.4 m s^{-1} for determining the number of spectral coefficients to remove (Ice et al. 2007). With a stationary antenna, a smaller value could possibly be used which would lessen the effects of clutter filtering on spectral moment estimation of weather signatures.

The curves for the “Full, Foliage, Prairie” and “Light Foliage, Urban” conditions match best. The linear model is closer to three out of four of the quadratic representative spectrum widths for the “Light Foliage, Prairie” case, but both models fall below their respective representative spectrum widths. The largest discrepancy occurs for the “Full Foliage, Urban, Windy” case. The spectrum width for the quadratic model overestimates the width from the

representative spectrum by approximately 0.25 m s^{-1} . The f value for the model is $12.6 \text{ dB}/(\text{m s}^{-1})$ while the f value for the representative spectrum is $19.2 \text{ dB}/(\text{m s}^{-1})$. This difference in f values causes most of the discrepancy, but the difference also causes the a parameter to be negative for the quadratic clutter model leading to an even larger error. The quadratic clutter model is very sensitive as the wind speed increases because the values of b and f are similar in magnitude, and negative a values are possible. Additional clutter data at high wind speeds would be helpful in better characterizing the behavior of the slope parameters in this problem area.

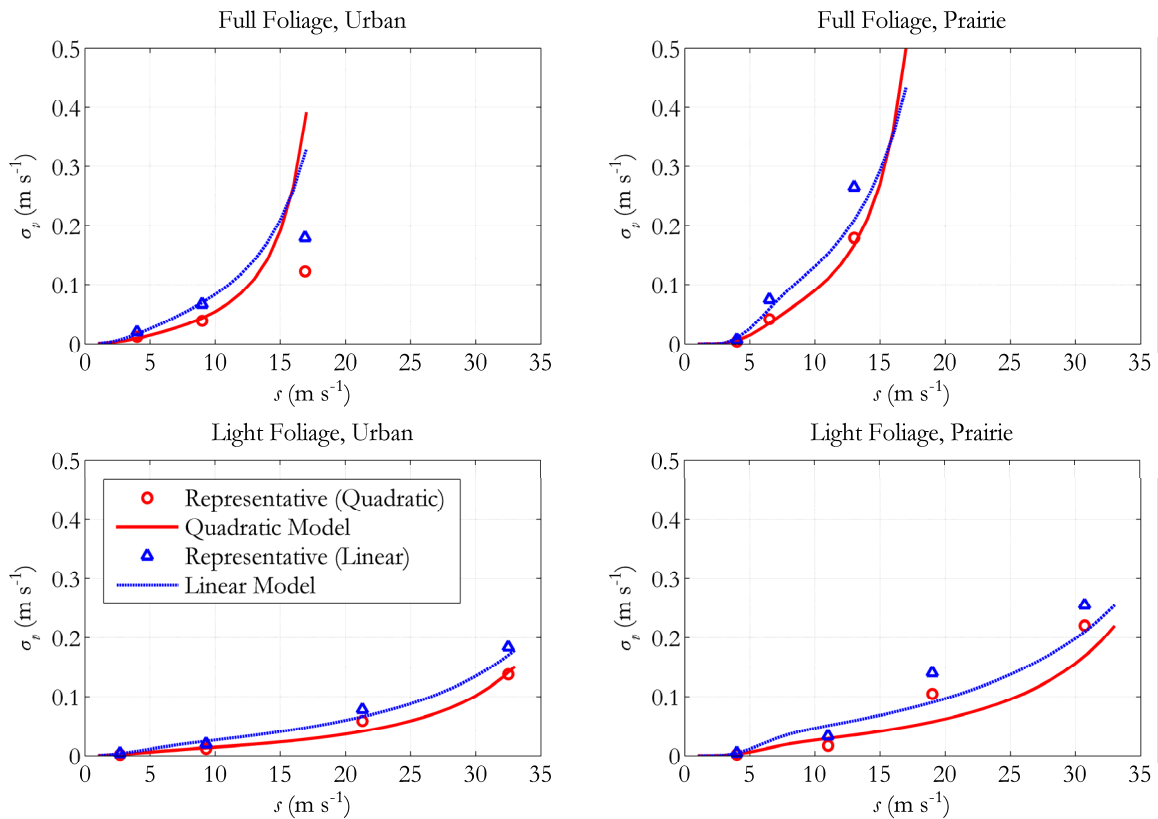


Figure 2.27. Comparison of the spectrum widths for both the representative spectra and the quadratic and linear clutter models for all four foliage-terrain pairs.

The next measure of the clutter width is based on the amount of power removed by an ideal clutter filter. This width can be defined as the value, w , such that the ratio of the total

clutter power to the residual power left after removing all power between $-w$ and w is 50 dB. The value of 50 dB was chosen since this is a standard value used to determine the effectiveness of clutter filters. A natural interpretation of this width, w , is that it can be used to determine the number of bins on either side of the zero-velocity bin that need to be removed when implementing a spectral clutter filter. Windowing can also have a large effect on ground clutter spectra, but w was computed for the ideal case without windowing and assuming an ideal clutter filter. Windowing effects are examined in Section 2.12.

The computation of w depends only on integrals of the power spectrum, $S(v)$:

$$10 \log_{10} \left(\frac{P_{tot}}{P_{tot} - P_{rem}} \right) = 10 \log_{10} \left(\frac{\int_{-\infty}^{\infty} S(v) dv}{\int_{-\infty}^{\infty} S(v) dv - \int_{-w}^w S(v) dv} \right) = th_{dB} \quad (2.17)$$

where th_{dB} is the threshold for the width which is set to 50 dB, P_{tot} is the total power of the clutter model from $-\infty$ to ∞ , and P_{rem} is the clutter power removed between $-w$ and w . For the quadratic clutter model, the integrals were computed using numerical integration. Because of this, a zero-finding method such as the bisection method needs to be employed to find the particular w that satisfies equation (2.17). For the linear model, a closed form solution for w can be found:

$$w = \frac{10}{f} \log_{10} \left(\frac{10^{\frac{th_{dB}}{10}}}{1 + \frac{f \log(10)}{20} \left(10^{\frac{d_n}{10}} - 1 \right)} \right), \quad (2.18)$$

where th_{dB} is again set to 50 dB, and d_n and f are the parameters that define the linear clutter model.

Using the formulas for σ_p and w , all of the width values for the representative spectra for both the quadratic and linear clutter models were calculated and are presented in Table 2.5. One of the most notable observations from Table 2.5 is that the values of w are much larger than the values of σ_p . The relationship between the two widths will be examined later in this section. When comparing the quadratic and linear models, the w values are very similar with the largest discrepancy occurring for the “Full Foliage, Urban, Windy” case. For σ_p , the values computed from the linear model are always greater than the ones computed from the quadratic clutter model. The spectrum width, σ_p , seems to be more sensitive to the difference between the linear and quadratic clutter spectral shapes.

Table 2.5. Width values for the quadratic and linear models for all 14 cases.

Foliage	Terrain	Wind	mph	w		σ_p	
				Quadratic	Linear	Quadratic	Linear
Full Foliage	Urban	Light	4.0	0.56	0.56	0.012	0.020
		Breezy	9.0	1.16	1.13	0.040	0.067
		Windy	16.9	3.13	2.34	0.123	0.180
	Prairie	Light	4.0	0.28	0.34	0.004	0.007
		Breezy	6.5	0.94	1.01	0.042	0.075
		Windy	13.0	2.40	2.49	0.179	0.266
Light Foliage	Urban	Light	2.7	0.15	0.21	0.002	0.004
		Breezy	9.3	0.38	0.43	0.012	0.020
		Windy	21.3	1.06	1.04	0.058	0.078
		Gale	32.5	2.63	2.33	0.137	0.184
	Prairie	Light	4.0	0.16	0.23	0.002	0.004
		Breezy	11.0	0.40	0.50	0.017	0.033
		Windy	19.0	1.32	1.41	0.105	0.140
		Gale	30.7	2.31	2.30	0.220	0.255

Because σ_p and w seem to measure different properties of the spectra, the relationship between w and wind speed may differ from the relationship between σ_p and wind speed.

Figure 2.28 shows the calculated w values for the representative spectra and both clutter models. The models match the representative spectra much more closely when considering w rather than σ_r . Interestingly, the largest discrepancy again occurs for the “Full Foliage, Urban” case. This time the linear model significantly overestimates the width compared to the representative spectrum, and the quadratic clutter model has a smaller error than the linear model. The largest discrepancies still occur at higher wind speeds which again suggests that additional data could improve the model in this area.

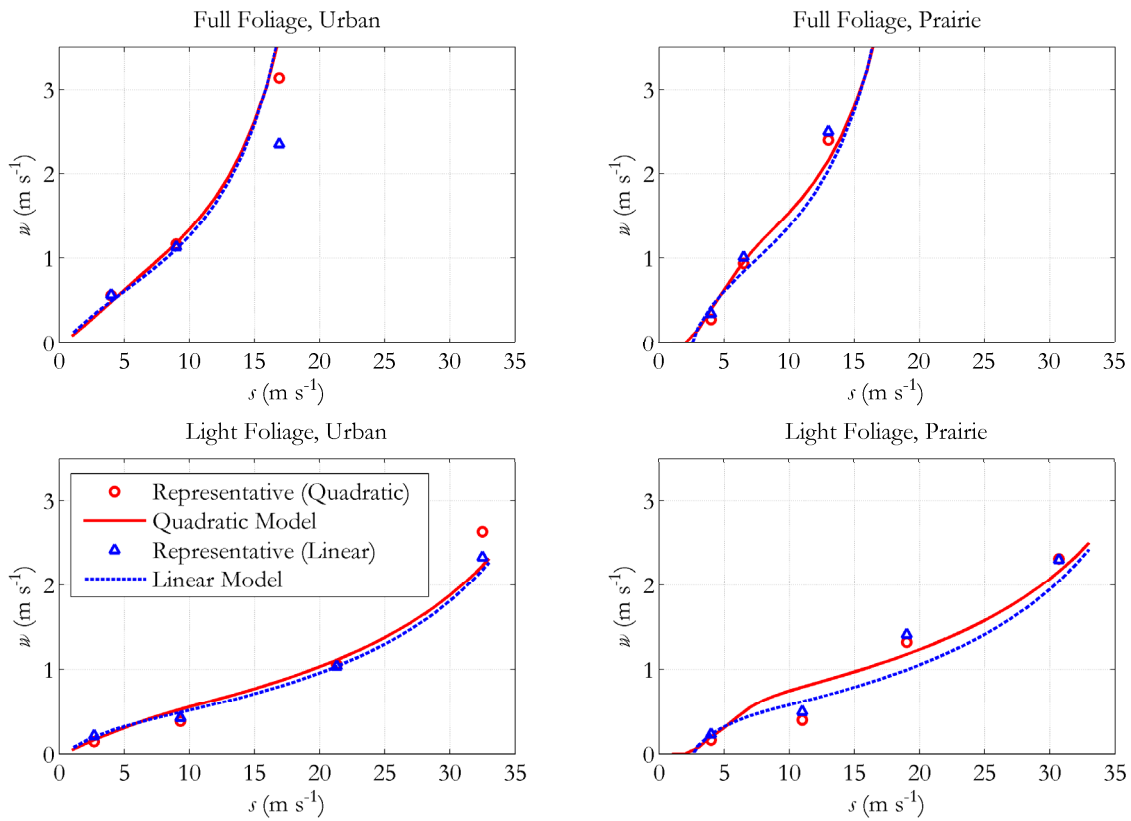


Figure 2.28. Comparison of w for both collected data and quadratic and linear clutter models for all four foliage-terrain pairs.

As mentioned previously, the Gaussian model is often used to model ground clutter spectra. The mean velocity is set to zero and an appropriate spectrum width is chosen. For

stationary clutter, the spectrum width could be chosen using the relationship shown in Figure 2.27 so that the spectrum width of the Gaussian model would be the same as the spectrum width of the quadratic or linear clutter model. Although the spectrum widths may be equal, the shapes of the spectra are significantly different. Another measurement of the width such as w might be affected by the difference in the shape. The relationship between σ_v and w for all three clutter models will be utilized to illustrate that different width measurements are sensitive to different properties of the spectra.

A method for determining σ_v and w for the quadratic and linear models has already been presented. For the Gaussian model, the spectrum width is an intrinsic parameter so only a formula for w needs to be derived:

$$w = \sigma_v \sqrt{2} \operatorname{erfc}^{-1} \left(10^{-\frac{th_{dB}}{10}} \right). \quad (2.19)$$

In this formula, erfc^{-1} is the inverse of the complementary error function, and th_{dB} is again set to 50 dB. It is clear from equation (2.19) that w is linearly related to σ_v and for $th_{dB} = 50$ dB, w is about 4.4 times larger. Figure 2.29 shows the relationship between σ_v and w for the Gaussian, quadratic, and linear clutter models. The models with the largest w values for a given spectrum width are the quadratic and linear models using the parameters derived for the “Full Foliage, Prairie” conditions. The quadratic model has a more complicated shape compared to the linear model, but the linear model does a reasonably good job of capturing the basic relationship. In general, the w values are larger for the linear and quadratic models than the values for the Gaussian model given the same measured spectrum width. This could lead to underestimates of the number of Doppler bins that need to be filtered using a spectral clutter filter if the Gaussian model is used to make the determination.




Figure 2.29. Comparison of w versus spectrum width for the Gaussian, quadratic, and linear clutter models. The quadratic and linear model parameters are based on the “Full Foliage, Urban, Prairie” conditions.

2.11 Simulating Ground Clutter with the Quadratic Clutter Model

Given a particular spectrum, simulating time series data using a method similar to the one found in Zrnić (1975) is straightforward. In this section, a couple of additional issues will be addressed: the first is appropriately simulating the spectral floor, and the second is accurately simulating time series data with smaller numbers of samples. The simulation developed in this section will be employed in Section 2.12 to examine the effect of spectrum length on the clutter shape.

Equation (2.13) defines the shape of the clutter spectrum in the log domain. The spectral floor component cannot be added directly to the log spectrum so another method will need

to be developed. A simple method for producing a spectral floor is to add noise ($c + d_n$) dB below the fixed spectral peak. The main drawback to that approach is that the fixed spectral floor does not change with the realization-to-realization variations of the spectral peak. This could lead to a lower correlation between c and d than seen in the real data. To test this hypothesis, data were simulated using parameters that match those from the “Full Foliage, Urban, Breezy” case studied in Section 2.4. This allows a direct comparison between the correlation matrix, \mathbf{R} , from the real data and the one calculated from the simulated data with a fixed spectrum, \mathbf{R}_F . For this test, 1000 realizations of 4096 samples were simulated and processed identically to the real data. The simulation parameters used were the same as the estimates from the “Full Foliage, Urban, Breezy” case ($a = 37.2$ dB/(m² s⁻²), $b = 61.3$ dB/(m s⁻¹), $c = 61.9$ dB, and $d_n = 3.3$ dB). The parameters estimated from the simulation were close to the input parameters ($\hat{a} = 32.1$ dB/(m² s⁻²), $\hat{b} = -58.4$ dB/(m s⁻¹), $\hat{c} = 61.4$ dB, and $\hat{d}_n = 3.4$ dB), but a slight bias shows up in the estimates of the a and b parameters. The two correlation matrices are shown below:

$$\mathbf{R} = \begin{bmatrix} 1 & 0.94 & -0.18 & 0.24 \\ 0.94 & 1 & -0.00 & 0.09 \\ -0.18 & -0.00 & 1 & -0.92 \\ 0.24 & 0.09 & -0.92 & 1 \end{bmatrix}, \quad \mathbf{R}_F = \begin{bmatrix} 1 & 0.97 & 0.81 & -0.40 \\ 0.97 & 1 & 0.91 & -0.44 \\ 0.81 & 0.91 & 1 & -0.47 \\ -0.40 & -0.44 & -0.47 & 1 \end{bmatrix}. \quad (2.20)$$

The correlation between a and b is very similar between the real and simulated data. The correlation between c and d is much smaller for the simulated data as hypothesized. There is also a strong negative correlation between b and c in the simulated data that does not occur in the real data.

Because of these differences, a change to the simulation was implemented that adjusts the simulated spectral floor based on the variations of the DC peaks in the spectra. Instead

of adding noise that is $(c + d_N)$ dB below the fixed spectral peak, noise is added $(c + d_N)$ dB below the spectral peak for each realization. After each spectral realization is simulated, the maximum spectral component is identified. Noise is then added to the spectrum so that the spectral floor is a fixed value below the spectral peak. This version of the simulation was run with the same parameters as before and with the spectral floor set to $(c + d_n) = 65.3$ dB. The expression $(c + d_n)$ is the same as c_0 from equation (2.10) which represents the distance from the spectral peak to the spectral floor in terms of d_n . The parameters estimated from the modified simulation ($\hat{a} = 29.8$ dB/(m² s⁻²), $\hat{b} = 56.5$ dB/(m s⁻¹), $\hat{c} = 63.8$ dB, and $\hat{d}_n = 3.3$ dB) differed from the input parameters more than the fixed simulation, but the correlation matrices are much more similar:

$$\mathbf{R} = \begin{bmatrix} 1 & 0.94 & -0.18 & 0.24 \\ 0.94 & 1 & -0.00 & 0.09 \\ -0.18 & -0.00 & 1 & -0.92 \\ 0.24 & 0.09 & -0.92 & 1 \end{bmatrix}, \mathbf{R}_s = \begin{bmatrix} 1 & 0.96 & -0.10 & 0.06 \\ 0.96 & 1 & -0.03 & -0.01 \\ -0.10 & -0.03 & 1 & -0.95 \\ 0.06 & -0.01 & -0.95 & 1 \end{bmatrix}. \quad (2.21)$$

For the modified simulation, the correlation matrix computed from the simulated data, \mathbf{R}_s , exhibited the strong negative correlation between c and d that was missing in the original simulation. The correlation between b and c is also much closer based on the data from the modified simulation. To see the extent of the differences between the simulated parameters and the estimated parameters, spectra from the quadratic clutter model corresponding to both are shown in Figure 2.30. Although the parameters differ, the overall shapes of the clutter models are nearly identical. In the rest of the chapter, the modified simulation is used to simulate ground clutter because of the significantly better correspondence to the correlation matrix computed from real data.

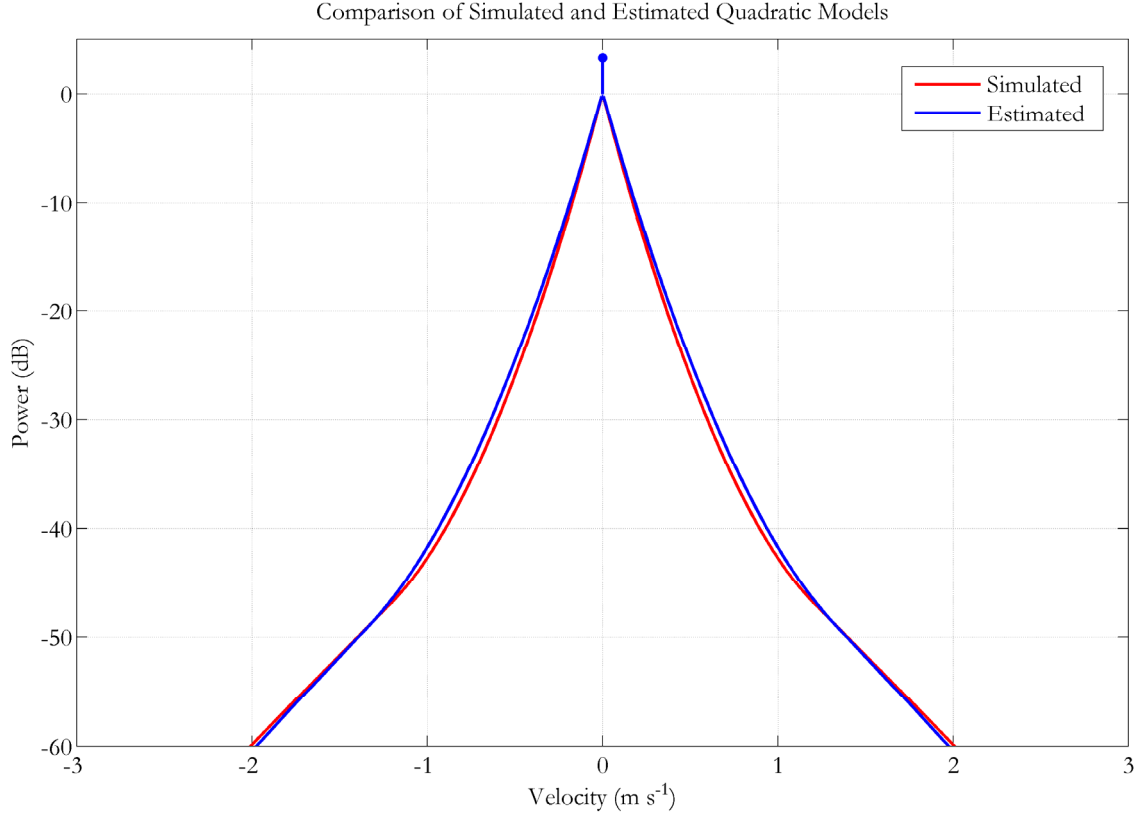


Figure 2.30. Comparison of quadratic clutter models based on simulated parameters ($a = 37.2 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $b = 61.3 \text{ dB}/(\text{m s}^{-1})$, $c = 61.9 \text{ dB}$, and $d_n = 3.3 \text{ dB}$) and estimated parameters ($\hat{a} = 29.8 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 56.5 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 63.8 \text{ dB}$, and $\hat{d}_n = 3.3 \text{ dB}$) for the modified simulation.

In addition to finding an appropriate way to simulate the spectral floor, there is also an issue when implementing the simulation for spectra with a small number of coefficients. When the number of spectral coefficients is large, approximating the power for each coefficient using discrete values from the quadratic clutter model seems to be sufficient. When the number of spectral coefficients is smaller, it is more accurate to integrate the model over the extent of each associated Doppler bin to calculate the power. Figure 2.31 shows two spectra of length 64. The first is the average of the spectra using the first 64 samples of a simulation with 4096 simulated samples and 1000 realizations. The second is a simulation of 64 samples, also with 1000 realizations.




Figure 2.31. Comparison between two simulations which use discrete values from the clutter model to determine the power for the spectral coefficients. The red curve uses the first 64 samples from a simulation with 4096 values, and the blue curve uses all of the samples from a simulation with 64 values.

To properly simulate windowing effects, both simulations start with a spectrum consisting of $3N$ samples. The first set of time series data is based on a spectrum with $3 \cdot 4096 = 12,288$ samples. The second set is based on a spectrum with $3 \cdot 64 = 192$ samples. Both spectra are processed with a rectangular window, and the parameters from the “Full Foliage, Urban, Breezy” case were utilized. No additional noise was added so the spectral floor is from the model and not a system noise floor. For this simulation, the power for the Doppler coefficients is determined using discrete values calculated from the clutter model. The spectra match relatively well, but the peak is about 3.5 dB larger when only 64 samples are simulated, and the skirts are lower by about 2 db.

To address the differences between the spectra in Figure 2.31, the simulation was modified to take into account the discrepancy caused by discrete sampling. Instead of using discrete values from the clutter model, the power for each spectral coefficient is computed by integrating the model over the extent of the associated Doppler bin using numerical integration. This results in the appropriate amount of power for each spectral coefficient. After the model spectrum is computed, the total power can also be scaled to a particular value. This allows the simulation of time series data with a specified mean power. Figure 2.32 shows a comparison similar to the previous figure except that the new simulation is used with the power for each spectrum set to 10 dB.




Figure 2.32. Comparison between two simulations which use integration of the clutter model to determine the power in the Doppler bins. The red curve uses the first 64 samples from a simulation with 4096 values, and the blue curve uses all of the samples from a simulation with 64 values.

Integrating to find the power for each spectral coefficient results in a much better match between the spectra. The only discrepancy occurs at the far edges of the skirts. The spectra in this figure are shifted compared to Figure 2.31 because the power is set to 10 dB instead of normalizing the spectra so that the peaks of the AC portions of the spectra are at the origin.

2.12 The Effect of Spectrum Length on Clutter Shape

Most of the spectra in this chapter are periodograms using 4096 time series samples. It is assumed in this section that the length of the spectrum is the same as the length of the time series. The earlier spectra were processed using an aggressive Chebyshev window to ensure that the spectral floor was undisturbed. These circumstances are uncommon when simulating time series data from a weather radar. Operational scanning strategies seldom use more than tens of pulses for a particular dwell, and less aggressive windows are normally utilized. The objective of this section is to examine the effect of the spectrum length on the shape of the clutter spectrum when using different data windows.

Four different spectrum lengths were used for the comparison: 64, 256, 1024, and 4096. The quadratic clutter model was utilized to simulate a spectrum from the “Full Foliage, Urban, Breezy” case ($a = 37.2 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $b = 61.3 \text{ dB}/(\text{m s}^{-1})$, $c = 61.9 \text{ dB}$, and $d_n = 3.3 \text{ dB}$). For comparison, the linear and Gaussian clutter models were also employed. The linear clutter model utilized the value of f for the “Full Foliage, Urban, Breezy” case from Table 2.2, or $f = 37.5 \text{ dB}/(\text{m s}^{-1})$ and the same value of d_n as the quadratic model. The Gaussian clutter model used the spectrum width calculated in Section 2.10 for the quadratic clutter model, $\sigma_p = 0.04 \text{ m s}^{-1}$. For all three models, the power was set to 10 dB and 1000 realizations with 4096 samples were simulated. A spectral floor was introduced to the

Gaussian clutter model using an SNR of 45.5 dB. The first 4096, 1024, 256, or 64 samples were processed with a rectangular window. The periodograms were calculated and averaged, and the results are shown in Figure 2.33.




Figure 2.33. Averaged spectra from simulations showing the effects of a rectangular window on time series data of differing lengths.

For all of the lengths greater than 64, the shape of the spectrum from the quadratic model is most similar to the spectrum from the linear model. When the length is 64, the Gaussian spectrum more closely matches the quadratic spectrum. The spectrum width of the linear model spectrum was computed to be 0.067 m s^{-1} in Section 2.10 which is larger than the 0.04 m s^{-1} computed for the quadratic model. The spectrum width of the Gaussian model was chosen to match the quadratic clutter model, and it appears that the spectrum width dominates at shorter spectrum lengths. Nearly any symmetric spectrum with the same

spectrum width would closely match both the quadratic and Gaussian spectra if the spectrum had only 64 samples. Apparently, the shape of the spectrum is not significant for shorter spectra, but the shape becomes more and more important as the number of samples increases.

The next window that was applied to the data was a Blackman window. This window is used when a more aggressive window is needed to control spectral leakage. For example, a Blackman window is used by the WSR-88D when implementing the GMAP clutter filter (Ice et al. 2007). Figure 2.34 shows the same spectra as Figure 2.33 except that the data were processed with a Blackman window instead of a rectangular window.

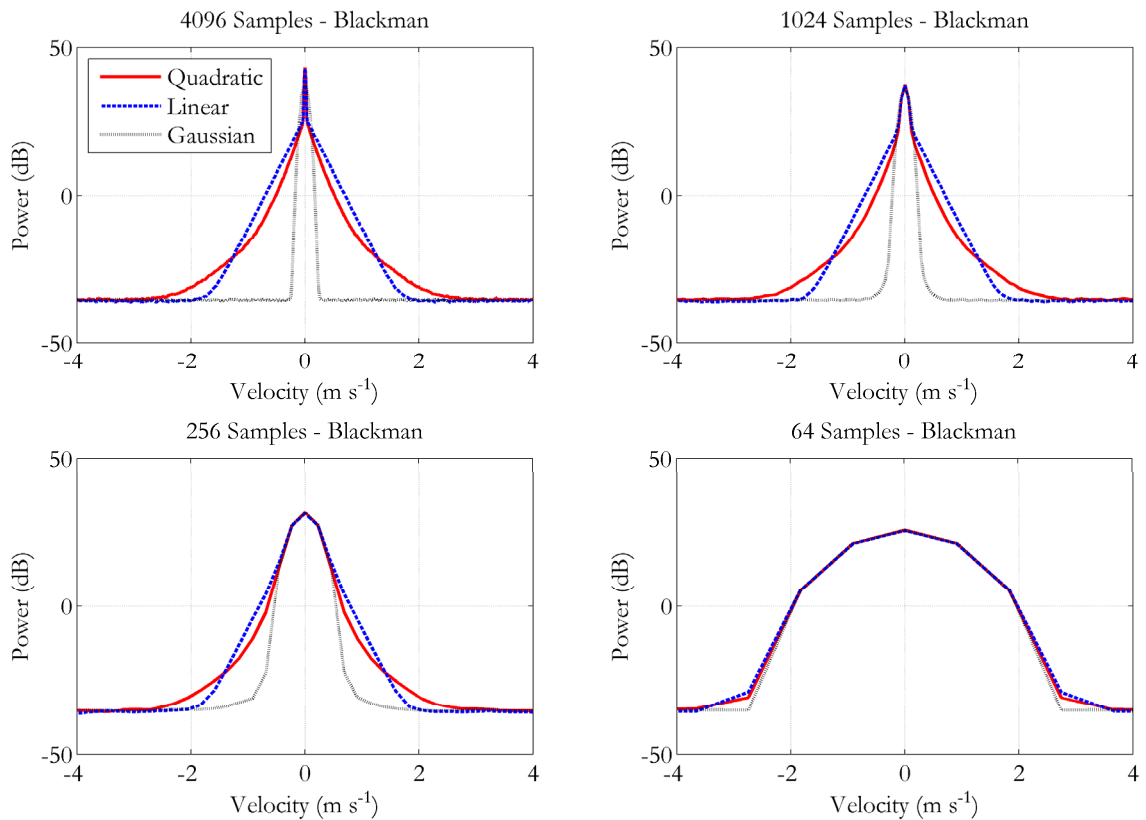


Figure 2.34. Averaged spectra from simulations showing the effects of a Blackman window on time series data of differing lengths.

The spectral shape is very recognizable when the length is 4096. The Blackman window is not quite as aggressive as the Chebyshev window used earlier in the chapter, but the spectral floor and the point where the linear and quadratic models cross are easily identified. The Gaussian spectrum is very narrow when the length is 4096 or 1024. Even at 256, the portion of the spectrum near the spectral floor departs from the quadratic and linear spectra. The 64 sample case is similar to the 64 sample case with the rectangular window. The Gaussian and quadratic spectra match very closely except for minor differences near the spectral floor. With the Blackman window, the linear spectrum is also very similar to other two spectra when the length is 64. The similarities between the Gaussian and quadratic spectra at shorter spectral lengths suggest that a Gaussian clutter model can be used to simulate ground clutter with acceptable accuracy as long as the spectrum is not too long.

One other issue that has not been addressed is the simulation of short spectra when using the Gaussian model. A standard simulation was used that samples the Gaussian spectral shape at discrete points. When this type of simulation was used to simulate Gaussian spectra of length 64, the spectrum width did not match the requested spectrum width. Apparently, for very narrow spectra and small spectral lengths, the amount of power in each Doppler bin needs to be determined by integrating to find the power instead of sampling the spectrum at discrete points. To avoid developing a Gaussian simulation similar to the quadratic model simulation from Section 2.11, 4096 samples were used in all of the simulations with subsets being used for the shorter spectra. This approach can be used as a workaround when narrow spectra need to be simulated.

2.13 Conclusion

After collecting several sets of ground clutter data under different terrain and wind conditions, the spectra were fit in the dB domain using a newly developed quadratic clutter model. The model better captured the visible curvature in the AC part of the spectrum compared to Billingsley's two-part exponential model (2002). A DC component similar to the one used by Billingsley was also included. An additional feature, a spectral floor, was incorporated into the model to capture radar-dependent phase noise. Fourteen cases were processed, and a representative set of model parameters was produced for each case. The parameters for those cases were then compared against the wind speed, and the quadratic clutter model was extended to capture wind speed effects. This model also captures differences between full and light foliage and urban and prairie terrain.

A section on beam broadening showed that increased beamwidths do not have a strong effect on the parameters estimated for the clutter models. The clutter is similar enough over the broader beam that the parameters do not change significantly. The effects of lower gain off broadside seem to be balanced out by the increased beamwidth so that the number of clutter targets that were modeled stay relatively constant. The beam broadening data also allowed an opportunity to look at the stability of the parameter estimation, and the c and d parameters were more stable over the collected cases than the a and b parameters. The stability of d leads to the stability of c because of the nearly constant phase noise. The a and b parameters are computed from the AC part of the spectrum which may vary more compared to the DC term. Overall, the model fitting seems to reliably capture the basic shape of the clutter spectra.

Section 2.10 examined the effect of wind speed on the spectrum width for both the quadratic model and the simpler linear model. A width measure, w , was developed that

measured how much power was removed by an ideal clutter filter. The effect of wind speed on w was also studied for both models. This measure is especially sensitive to the shape of the clutter and was used to compare the quadratic, linear, and commonly used Gaussian clutter models. The relationship between the spectrum width and w was examined which shows that the quadratic and linear clutter models exhibit larger w values for a given spectrum width compared to the Gaussian clutter model. This is partly due to the different shape of the AC portion of the spectrum and partly due to the DC component. In some cases, a larger part of the clutter spectrum may need to be filtered than would be predicted from the Gaussian model.

The section on simulating time series data using the quadratic model introduced the use of integration to better calculate the amount of power in the Doppler bins. This becomes more necessary when fewer samples are simulated. Two ways for adding in the spectral floor were compared to evaluate the effects on the correlation between the c and d parameters. The spectral floor seems to change based on the peak of the spectrum, and this feature was captured by adjusting the spectral floor for each simulation realization based on the magnitude of the DC component. The estimated correlation matrix was used to measure the effectiveness of the modified simulation which showed a better correspondence to the correlation matrix estimated from real data. Adding the spectral floor separately for each realization could also be implemented with the Gaussian clutter model, but the results were not studied.

The effects of data windowing when simulating spectra of various lengths showed that the Gaussian model may be sufficient for modeling ground clutter in some cases. The addition of a spectral floor to the Gaussian model enabled it to nearly match the shape of the quadratic model when the spectral length was 64. The additional accuracy from using the

quadratic or linear clutter models is obscured when simulating fewer samples, but the ability to simulate time series data that correspond to a particular terrain type, foliage level, and wind speed is still useful.

Even though progress was made on characterizing the spectra of ground clutter, additional data collections could lead to a better understanding of what happens when wind speeds increase especially in “Full Foliage” conditions. The behavior of the b and f parameters when a is close to and possibly less than zero is still an open question. Finding a way to collect data from wooded terrain with either the PAR or another radar could extend the model even further. Tighter controls on the terrain at particular range bins may be necessary to fine tune the model because mixed terrain data is the norm. Although this may have had some effect on modeling the parameters, the overall relationships seemed to follow the expected behavior for ground clutter spectra. The DC part of the model decreased as the width of the AC part increased. The DC component was also larger for urban clutter compared to prairie clutter which is consistent with the higher prevalence of hard targets in urban clutter. The linear approximation of the f parameter with the log of the wind speed matched the data quite well but differs from Billingsley’s model of the β parameter which varies inversely with the log of the wind speed. Part of this could come from the differences in the linear fit for the two models. The f parameter is more closely related to an intrinsic clutter width at a particular wind speed rather than being a direct fit of the data.

In general, the basic trends in the shapes of ground clutter spectra were verified, but certain additional details and relationships were discovered. Incorporating a spectral floor in the model captures an important radar-dependent attribute of ground clutter spectra that can have a practical effect on processing data that contain ground clutter. Ground clutter filtering may be improved by tailoring the filter based on the spectral shapes predicted from

the quadratic model. In the long run, better characterization of ground clutter can only help in finding new ways to mitigate the effects of clutter when estimating weather parameters from real data, and taking advantage of narrower spectra from a stationary antenna could provide additional opportunities for improving clutter mitigation performance.

3. Staggered PRT Beam Multiplexing

3.1 Introduction

Beam multiplexing (BMX) is a scanning strategy that takes advantage of the electronic steering capabilities of phased array radars to collect independent weather samples. When contiguous pulses are transmitted with a traditional weather radar using a parabolic dish, the samples from the weather return may be highly correlated especially at narrower spectrum widths. In contrast, a phased array radar can transmit a smaller number of pulses or even a single pulse at a particular beam position and then return to that same beam position at a later time. If there is sufficient time between these short data collection periods, the data are nearly independent and the errors after averaging will be reduced significantly. If the time between data collection periods is used to collect data at other beam positions, the radar can be utilized continuously resulting in significant time savings.

An example of a BMX scanning strategy is shown in Figure 3.1. In the case of contiguous pulses, a sequence of pulses is transmitted at a single beam position, and the associated data are collected. In this particular BMX example, two pulses are transmitted at four different beam positions. This cycle is repeated so that several pairs of pulses are transmitted at each of four beam positions. By adjusting the number of beam positions included in the strategy, the time between pairs at a particular beam position can be selected so that the data are nearly independent. Also note that the radar is being used continuously so that the time between pulses at a particular beam position is being used to collect data at other beam positions. This particular strategy of sending pairs of pulses at each beam position will be referred to as BMX pairs.

Contiguous Pulses:



Beam Multiplexing Pairs:




Figure 3.1. A comparison between contiguous pulse data collection and BMX pairs.

The basic idea behind BMX is straightforward, but there are additional issues that need to be addressed before an actual scanning strategy is implemented. The first is how much time is needed between data collection periods so that the samples are nearly independent. If the power spectrum density of weather signals is assumed to have a Gaussian shape, the samples of the return are correlated in time with the following correlation coefficient (Doviak and Zrnić 1993):

$$\rho(mT_s) = \exp\left(-8(\pi\sigma_p mT_s / \lambda)^2\right). \quad (3.1)$$

The correlation coefficient depends on the spectrum width, σ_p , the lag or time between samples, mT_s , and the wavelength, λ . The time between samples is determined by T_s , the pulse repetition time or PRT, and an integer, m , denoting a particular number of samples.

The spectrum width is an intrinsic property of the weather that can be estimated but not modified, and the wavelength is fixed when transmitting only one frequency. This leaves the time between samples as the only easily modifiable parameter. Because the samples are more correlated at narrower spectrum widths, the correlation at a spectrum width of 1 m s^{-1} is a good representative case for highly correlated samples. Figure 3.2 shows the correlation with respect to the time between samples for the NWRT wavelength of 9.37 cm :




Figure 3.2. Correlation in time for a spectrum width of 1 m s^{-1} assuming a Gaussian spectrum.

At 20 ms between samples, the correlation coefficient is approximately 0.027, and the data will be nearly independent. This gives a reasonable lower bound for the time between samples when collecting beam multiplexed data. Because the correlation depends on the

wavelength, a radar that uses a different frequency would have a different lower bound for the time between samples.

A related concern is the maximum allowable total dwell time; this is the time from when the first sample is collected at a particular beam position until the last sample is collected at that beam position. Beam multiplexing scanning strategies have a longer total dwell time at each beam position than traditional contiguous-pair scanning strategies even if the acquisition time is shorter. The acquisition time is defined as the amount of time spent collecting data at a particular beam position. For example, the acquisition time for the simple strategy illustrated in Figure 3.1 would be $N_C * T_C = 6$ ms where $N_C = 3$ is the number of collections at each beam position, and T_C is the collection time assuming a PRT of 1 ms. The total dwell time is based on the collection time at each beam position, $T_C = 2$ ms, and the number of beam positions, $N_B = 4$: $(N_C - 1) * N_B * T_C + T_C = 18$ ms. This distinction between acquisition time and total dwell time is important because the acquisition time is used to compare the efficiency of scanning strategies by looking at the tradeoff between errors of estimates and acquisition time.

The total dwell time determines the amount of weather advection that occurs in a resolution volume. Over time, advection leads to weather features moving from one resolution volume to another. If a maximum advection speed of 60 mph is assumed, it is possible to set reasonable limits on the dwell time for BMX scanning strategies. Both the WSR-88D and the NWRTP use the same transmitter with the shortest pulse corresponding to a pulse depth of 250 m. If the advection in the range cell is limited to one quarter of this length (a common rule-of-thumb), then the longest possible dwell time is the time it takes 60 mph weather to move 62.5 m. This turns out to be about 2.33 s. Assuming a beamwidth of 1° , the cross-range resolution is 250 m or greater at ranges of 14 km or more. This means

that the 2.3 s value will be sufficient in all cases except for 60 mph weather movement nearly perpendicular to the beam at less than 14 km. Because this covers nearly all of the possible conditions, 2.3 s seems to be a reasonable limit for the total dwell time at a particular beam position. Of course, a different radar with a different pulse width could have a shorter or longer dwell time limit.

The last concern to be considered is the return from second trip echoes if the BMX beam positions are too close together. For example, referring to the BMX strategy from Figure 3.1, the return from second trip echoes from the second pulse in the first pair could show up after the transmission of the first pulse of the first pair at the second beam position. In order to mitigate this contamination, it seems reasonable to try to avoid the main lobe and first sidelobe of the antenna pattern. For an antenna pattern with a 3 dB two-way beamwidth of less than 2.5° , keeping 6° or 7° between beam positions should be enough to avoid nearly all of these second-trip effects. Unusually strong third or fourth trip echoes could still cause problems.

There are certainly many different ways to meet the previously described conditions, but a straightforward one is shown below. This pattern assumes a 1° beamwidth with azimuth spacing of 1° .




Figure 3.3. A simple azimuth-only beam multiplexing pattern.

The above pattern starts with beam position one at 0° followed by the second beam position at 7° . The third beam position follows at 1° which gives a minimum angular distance of 6° for the pattern. This pattern could be repeated as many times as necessary to collect the data of interest as long as the total dwell time for each beam position was less than 2.3 s. The amount of time between collections at a particular beam position is determined by the number of beam positions in the data collection strategy. The pattern could then be shifted and repeated to cover the entire area of interest.

One example of a collection strategy that employs this pattern is transmitting two pulses during each visit to a beam position and repeating the pattern 16 times for a total of 32 pulses at each beam position. Using a PRT of 1 ms, the total time for collecting one pattern is 28 ms. This meets the lower bound of 20 ms before returning to the same beam position. The total time for 16 patterns is $16 * 28 \text{ ms} = 448 \text{ ms}$ which also meets the 2.3 ms condition. This total time is slightly larger than the total dwell time at each beam position which would be $(16-1) * 14 * T_C + T_C = 422 \text{ ms}$, but the total collection time is easier to compute and is an upper bound for the total dwell time. This strategy is one example of using a simple pattern to meet the previously stated conditions. An additional condition for measuring radial velocity (or spectrum width) is to collect at least two pulses during each visit to a beam position. The BMX-pairs strategy also satisfies this additional condition.

An approach very similar to the one above was implemented on the NWRT to show that beam multiplexing is feasible and that the predicted theoretical performance is realizable (Orescanin et al. 2005). This conference paper and another follow-up paper (Yu et al. 2007) showed the feasibility of BMX using a phased array radar and confirmed the theoretical predictions. Now that this initial work has been completed, it is important to look at the long-term viability of BMX. In the next section, some of the issues with BMX pairs will be

considered that keep it from being practical as the sole collection strategy for meeting the WSR-88D operational requirements. After that, an additional alternative will be presented and tested. Although BMX has shown promise, there are several hurdles that still need to be overcome before it can be used in place of other proven strategies.

After examining some of the issues with BMX pairs and introducing an alternative strategy, a simulation is used to compare the performance of the new strategy to a couple of conventional scanning strategies. The performance is assessed using a requirements-based approach starting with the WSR-88D requirements. Next, the simulation results are verified using weather data collected at a single beam position. A simple clutter filtering strategy is tested using the new BMX strategy, and concluding remarks are presented in the final section.

3.2 Issues with BMX Pairs

The limitations of BMX pairs fall into two main categories: general BMX limitations and limitations specific to BMX pairs. Two of the drawbacks to BMX in general include difficulty in clutter filtering and difficulty in spectral processing. Some of the difficulties are caused by the fact that BMX samples are not uniformly spaced. Most clutter filters rely on uniformly spaced samples, and spectral processing approaches using standard Fourier transforms also assume uniformly spaced samples. There are other clutter filtering techniques that do not rely on uniformly spaced samples such as regression filters (Torres and Zrnić 1999) and DC removal, but these types of filters also have limitations for BMX. The regression filters fit polynomials to the underlying clutter signals, but the long times between samples interfere with the fit. DC removal filters some of the clutter power, but the effectiveness is limited even when applied to uniform samples. A test of DC removal is

presented in Section 3.7. These drawbacks are significant because clutter filtering is important for accurate estimation of spectral moments especially at low elevation angles. Because of these difficulties, alternative methods for mitigating clutter should be considered for BMX such as the use of auxiliary receive channels (Le et al. 2009). In addition to the clutter filtering limitations, the difficulty in implementing spectral processing when using BMX is also significant because spectral processing has a lot of promise for improving the estimation of spectral moments and for providing additional information based on the shape of the spectrum (Fabry and Keeler 2003).

Another limitation of BMX is that most of the gains come at high signal-to-noise ratios (SNRs). Figure 3.4 shows simulation results of velocity errors for contiguous sampling with the number of pulses, $M = 54$, and BMX pairs with $M = 32$ (or 16 pairs). A more detailed description of the simulation method will be discussed later. The BMX collection takes less than 60% of the time for contiguous pairs, but the improvement only occurs above about 15 dB SNR. At very low SNRs, the errors are roughly twice as large for BMX pairs. Also note that the spectrum width for the simulations was set to 2 m s^{-1} . At larger spectrum widths, BMX may only outperform contiguous pulses if more pairs are collected or maybe not at all when using comparable collection times.

The main limitation specific to BMX pairs is that it is difficult to recover velocity and spectrum width from the second trip. Since only one pair of pulses is transmitted at each beam position, there are no second trip echoes in the data collected after the first pulse is sent. Figure 3.5 depicts a simplified example corresponding to a single pair of pulses. Second trip echoes do show up after the second pulse is sent (i.e. the open red rectangle), but the velocity and spectrum width for the second trip echoes cannot be computed.




Figure 3.4. A comparison of standard deviation of velocity for contiguous pulses and BMX pairs.



Figure 3.5. A simplified example showing first and second trip echoes. The solid rectangles depict two transmitted pulses, and the triangles are first trip echoes from each pulse. The red rectangle is a second trip echo from the first (red) pulse.

If second trip echoes are present, the lag-1 autocorrelation of first trip echoes will not be biased since the second trip echoes are not correlated with the first trip echoes. This means that the second trip echoes will show up as noise when computing the lag-1 autocorrelation.

This inability to compute velocity or spectrum width for second trip echoes is a significant limitation. The range unfolding that is used on the WSR-88D for split cuts and batch cuts is useless for data collected with BMX pairs. In the next section, a new BMX collection strategy will be introduced that addresses this particular problem. The general limitations of BMX will still apply to this new strategy, but these may be addressed in the future if the strategy is otherwise viable.

3.3 A New BMX Strategy

One way to recover second trip echoes would be to collect three pulses at each beam position instead of two. Single BMX pulses could be collected separately with a longer PRT to compute reflectivity. This strategy would be similar to batch mode except that the data would be collected using BMX instead of contiguous pulses. The problem with this strategy is that second trip echoes could only be computed from the second pair of pulses. This would mean that the same number of three-pulse collections would be needed as pairs. The time for this would be 1.5 times as long as collecting pairs which would result in a three-pulse strategy similar to the previous example taking 90% of the time compared to contiguous pulses based on the original strategy taking 60% of the time. This 10% time savings is not very significant considering that clutter filtering and spectral processing are much more difficult and that the benefits only occur at high SNR values. However, the previous example showed error results only for velocity. At narrow spectrum widths, significant gains can be made when measuring reflectivity, and BMX could also be useful for

high SNR situations. These situations where BMX excels should be studied further to find the best possible utilization of BMX strategies.

Instead of considering additional strategies that rely on second trip echoes, it makes sense to explore strategies that depend only on first trip echoes. Staggered PRT is one such strategy (Torres et al. 2004). With staggered PRT, longer PRTs are utilized for computing velocity and spectrum width instead of the short PRTs used for batch-type strategies, and velocities are dealiased rather than range unfolded. It may be possible to collect a few pulses using a staggered strategy but also use beam multiplexing so that the data are nearly independent.

Here is one possible scanning strategy that implements short, BMX groups of staggered PRT collections:



Figure 3.6. A scanning strategy that uses both staggered PRTs and BMX.

The vertical lines show when the pulses are transmitted, and the longer lines correspond to longer PRTs. This helps visualize the long PRTs compared to the shorter ones without trying to measure small changes in distances between the pulses. In this strategy, one pulse is transmitted using a long PRT followed by two shorter ones. BMX is used so that the data from the different collections are nearly independent.

A velocity can be computed from the first pair of pulses corresponding to the long PRT, and another velocity can be computed from the second pair of pulses corresponding to the shorter PRT. These two velocities can then be dealiased which results in a larger unambiguous velocity than the unambiguous velocities corresponding to each of the

individual PRTs. For example, the unambiguous velocity after velocity dealiasing when the shorter PRT is $2/3$ the length of the longer PRT is the same unambiguous velocity that would result from a contiguous pulse collection using a PRT that is $1/3$ the length of the longer PRT.

One advantage of staggered PRT beam multiplexing (SBMX) over standard staggered PRT is that the data from the first pulse does not have any overlaid echoes. This allows for the computation of uncontaminated reflectivity from the first pulse alone. Otherwise, the processing is nearly the same. The data collected from the third pulse could be contaminated with second trip echoes from the second pulse, but the second trip echoes are not needed for range unfolding when using a staggered strategy. The second trip echoes will act as noise when computing the velocity from the second pair, but the estimate will not be biased. In this case, the lack of second trip echoes in both pulses is an advantage. Unlike in BMX pairs where the lack of second trip echoes prohibits range unfolding; SBMX (and conventional staggered PRT) rely on first trip echoes because velocity dealiasing is utilized instead of range unfolding. By starting the SBMX collection with an appropriately long PRT, the returns from both the first and second pulses should have minimal occurrences of overlaid echoes.

3.4 Comparisons to Conventional Collection Strategies

When looking at a new data collection strategy, it is important to select appropriate strategies for comparison as well as a reasonable method to make the comparisons. Until recently, contiguous pulse strategies based on range unfolding were the only techniques used by the NEXRAD network to mitigate range and velocity ambiguities. These range unfolding techniques depend on a set of longer PRT pulses used to compute range unambiguous reflectivity along with a set of shorter PRT pulses used to compute velocity and spectrum

width (Torres and Zrnić 2004). The signal power from the longer PRT is used to range unfold the velocities and widths computed from the shorter PRT. The main drawback of this range unfolding technique is that range overlaid signals with similar powers cannot be separated. This can result in overlaid echoes that are displayed as purple on the WSR-88D display (commonly referred to as purple haze). More recently, SZ phase coding has been implemented at lower elevation angles to minimize the occurrence of overlaid echoes (Torres 2005). Because of the lack of phase coding capabilities on the NWRT, these strategies cannot be tested using real data. In terms of estimation errors, phase coding strategies perform very similarly to contiguous pulse strategies in most cases so including a range-unfolded contiguous pulse strategy should be sufficient. Another scanning strategy that was developed to mitigate range and velocity ambiguities is conventional staggered PRT processing (Torres 2006). Conventional staggered PRT collections will be implemented on the WSR-88D in the next few years and have also been implemented on the NWRT. As mentioned earlier, staggered PRT techniques use an alternating pattern of two different PRTs along with velocity dealiasing to avoid range unfolding. A comparison between range-unfolded contiguous pulses, conventional staggered PRT, and SBMX should provide a reasonable test of the SBMX strategy.

In Yu et al. (2007), a metric termed improvement factor was introduced to compare the acquisition times between BMX and conventional strategies. Theoretical formulas and experimental results showed the benefits of BMX at narrow spectrum widths for reflectivity estimates and at high signal-to-noise ratios for velocity estimates. These conditions highlight the strengths of BMX and illustrate the situations where BMX can produce the most benefits. Another way to examine radar performance is from a requirements-based perspective. For example, the WSR-88D was designed to meet certain requirements to

ensure an acceptable level of data quality. The NEXRAD requirements were established for a particular set of conditions. For example, the requirement for reflectivity is a standard deviation of less than 1 dBZ at an SNR of 10 dB and a true spectrum width of 4 m s^{-1} for 1 km range bins. For velocity and spectrum width, the requirements are standard deviations less than 1 m s^{-1} for 250 m range bins with a true spectrum width of 4 m s^{-1} at 8 dB and 10 dB SNR, respectively. These requirements can be found in the NEXRAD Technical Requirements (NTR 1991). The following scanning strategy comparisons will use these requirements as a starting point to examine performance. By utilizing this type of requirements-based comparison, certain shortcomings of this type of approach will be revealed especially when using strategies that differ significantly from conventional contiguous pulse collections.

Both the criteria for comparison and the strategies to be compared are determined. Now, the methodology to make the comparison needs to be established. Although there are theoretical formulas for the variance of estimates for some of these techniques, simulations of weather-like time series data avoid the assumptions of some theoretical formulas and ensure easily comparable results. The bias and standard deviations of the estimates will be used for the comparisons. In the next section, time series simulations will be used that assume a Gaussian shaped spectra. Even though this assumption does not always hold for actual weather signals, these types of simulations have been used for several years to test radar processing techniques and have been useful in predicting algorithm performance (Zrnić 1975). Because two of the techniques utilize staggered PRT data, the standard practice of simulating data using a shorter PRT that divides both of the staggered PRTs and then discards some of the samples will be used. For this study, the staggered PRTs have a ratio of $2/3$ which allows the use of a simulation PRT with $1/3$ the length of the longer staggered

PRT. When using conventional staggered processing techniques, the unambiguous range is determined by the longer PRT. To get a comparable unambiguous range with the range-unfolded contiguous pulse case, the longer (or surveillance PRT) will be the same as the longer of the staggered PRTs. The shorter PRT (or Doppler PRT) will be set to 1/3 of the surveillance PRT which will give velocity estimates with the same maximum unambiguous velocity as the dealiased unambiguous velocity for the staggered techniques. Because all of the PRTs are multiples of the same shortest PRT, one long set of uniform samples can be used to produce a simulation realization for all three scanning strategies.

Figure 3.7 demonstrates how the uniformly-spaced samples can be used to simulate data corresponding to each of the scanning strategies.




Figure 3.7. Simulating contiguous pulse, staggered PRT, and SBMX time series data using uniform PRT samples.

The samples corresponding to the surveillance PRT can be simulated by skipping two samples while the Doppler PRT samples already match the underlying simulation PRT. For the staggered PRT collection, a single sample is skipped to simulate the PRT that is 2/3 the length of the longer PRT while the longer PRT is simulated by skipping two samples. This same strategy is used for SBMX except that several samples are skipped in between each three-sample collection to simulate the time between data collections at a particular beam position. A single set of simulated samples of appropriate length can be used to produce a simulation realization for all three scanning strategies.

3.5 Simulation Results

Now that the simulation method has been determined, the appropriate simulation parameters need to be established. Even though the requirements for reflectivity, velocity, and spectrum width were developed for the WSR-88D which operates between 2700 and 3000 MHz, the operating frequency of the NWRT is close enough (3200 MHz) that the requirements are applicable. It also aids in comparisons with real data to use the NWRT transmit frequency in the simulations. The PRTs available on the NWRT vary from 800 to 3200 μs . If the longest PRT is three times the shortest PRT as discussed earlier and the shortest PRT used is the simulation PRT, the simulation PRT can vary from 800 to about 1066 μs . An additional restriction is that the PRT must be a multiple of 4 μs . For this study, a simulation PRT of 1024 μs was chosen. This results in an unambiguous velocity of $v_a = c/4T_s f = 22.87 \text{ m s}^{-1}$ where c is the speed of light, T_s is the PRT, and f is the transmit frequency. For the contiguous pulse strategy, the surveillance PRT is 3072 μs and the Doppler PRT is the same as the simulation PRT, 1024 μs . The staggered PRT and SBMX strategies both use a shorter PRT, T_1 , of 2048 μs and a longer PRT, T_2 , of 3072 μs . As stated

earlier, when the velocity is dealiased using the algorithm from Torres et al. (2004), the maximum unambiguous velocity for the staggered strategies matches the unambiguous velocity of the contiguous pulse strategy, 22.87 m s^{-1} .

The spectrum width used in the simulation is based on the requirement of a 4 m s^{-1} true spectrum width. Enough uniform samples were generated for each realization so that 25 three-pulse SBMX collections could be simulated with a sufficient number of samples between the collections to ensure the collections are nearly independent (531 samples). This results in a time of 21.5 ms between the first sample in one SBMX collection and the first sample of the next collection. Based on the earlier correlation discussion, this should be sufficient to ensure that the estimates are nearly independent and that the full benefits of beam multiplexing are realized. A total of 100,000 realizations of 531 samples each were simulated so that each realization could be used to simulate time series samples for all three data collection strategies. The time series data were simulated without noise added so that both the 8 dB and 10 dB SNR realizations could be produced by adding noise independently. The 10 dB SNR case was used to test the performance of reflectivity and spectrum width, and the 8 dB SNR case was used for velocity. Conventional pulse pair processing was utilized to compute the velocities and spectrum widths. The two velocities from the staggered strategies were both dealiased and averaged which is a small departure from the algorithm described in Torres et al. (2004). In all three cases, the spectrum width was computed using the ratio of signal power to lag-1 autocorrelation calculated from the shorter PRT.

The results from Figure 3.8 show the bias of each of the spectral moment estimates (S , v , and w) versus acquisition time in ms using a true spectrum width of 4 m s^{-1} in the simulation. S is the signal power, v the velocity, and w the spectrum width. As defined earlier, the

acquisition time is the time actually spent collecting data at a particular beam position. This is in contrast with the total dwell time which is the total time from the beginning of the data collection at a particular beam position until the end. For example, 25 SBMX three-pulse collections would have an acquisition time of $N_C * T_C = 179.2$ ms where $N_C = 25$ and $T_C = 3.072$ ms + $2 * 2.048$ ms = 7.168 ms is the time in ms for each collection. The total dwell time for this case is $(N_C - 1) * N_B * T_C + T_C = 523.3$ ms where $N_B = 3$ and $N_B * T_C$ is the time between the first sample in one SBMX collection and the first sample of the next collection. The number of beam positions, N_B , is set to 3 to ensure that the SBMX collections are nearly independent. There are $(N_C - 1) = 24$ time intervals of this length for the first 24 collections with one added on the end to compute the total dwell time.




Figure 3.8. Biases for signal power, velocity, and spectrum width versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 4 m s^{-1} and a velocity of 15.25 m s^{-1} .

The signal power is computed from the long PRT data in all three cases since the unambiguous range is largest. When using the staggered techniques, powers from the shorter PRT can be used to lower the errors of reflectivity estimates up to the maximum unambiguous range of the shorter PRT, but the worst-case errors occur in the data corresponding to the last third of the longer PRT which is why the estimates from only the longer PRT are used. The biases for signal power and velocity are small. There is a larger bias for the spectrum width especially for shorter acquisition times. This bias is described by Zrnić (1979) as “a bias inversely proportional to the number of sample pairs.” For the SBMX strategy, the largest spectrum width bias occurs when only two three-pulse collections are employed. For staggered PRT, the largest bias occurs when three pairs are collected. In the time it takes for the SBMX and staggered PRT strategies to acquire data for only two or three pairs, respectively, the contiguous pulse strategy can acquire 14 samples and 13 pairs to compute the lag-1 autocorrelation. Overall, the bias results look reasonable except for the expected larger spectrum width biases when using only a few sample pairs. This reveals a possible problem when shorter acquisition times are utilized with staggered PRT or SBMX.

Figure 3.9 shows the standard deviations of reflectivity, velocity, and spectrum width versus acquisition time in ms for the same simulation. The standard deviation of the reflectivity is estimated from the standard deviation of signal power using a formula similar to the one found in Doviak and Zrnić (1993):

$$\sigma_z \approx 10 \log \left[1 + \frac{SD(\hat{S})}{\bar{S} \sqrt{r}} \right], \quad (3.2)$$

where σ_z is the standard deviation of reflectivity, SD is the standard deviation function, \bar{S} is the signal power, and r is the number of samples averaged in range. The values of signal

power in range are assumed to be independent so that if r values are averaged, the standard deviation will be reduced by the square root of r . This is a reasonable assumption for range bins with a length similar to the length of the pulse. In order for the WSR-88D to meet the reflectivity requirement, the signal power from four range bins is averaged to produce reflectivity bins with 1 km resolution. If not otherwise stated, a range averaging value of $r = 4$ was used in all of the simulations that produce reflectivity error estimates.

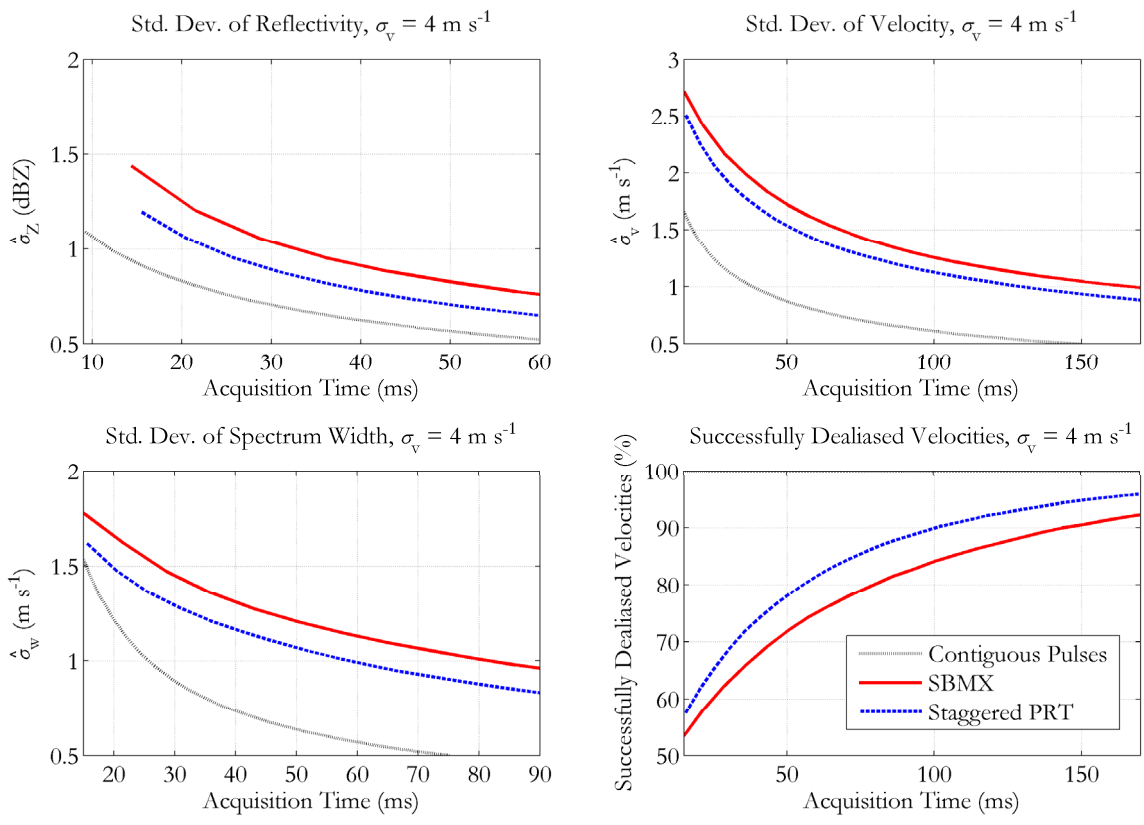


Figure 3.9. Standard deviations of reflectivity, velocity, and spectrum width and successfully dealiased velocity percentages versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 4 m s^{-1} . Acquisition time for the contiguous pulse strategy is the sum of the long PRT acquisition time to needed to meet the reflectivity requirement and the short PRT acquisition time need to meet the velocity and spectrum width requirements.

The final plot of successfully dealiased velocities was included because of the dealiasing failures that can occur when using staggered PRT dealiasing algorithms. Catastrophic errors

occur when the velocity is dealiased improperly because of errors in the velocity estimates. If the catastrophic errors were included in the standard deviation or bias computations, the results would be skewed. In order to study the effects of dealiasing, the simulated velocity was chosen to be twice the maximum unambiguous velocity for the longer PRT or 15.25 m s^{-1} . The unambiguous velocity range was divided up into two regions: Region 1, -22.87 m s^{-1} to 7.63 m s^{-1} , and Region 2, 7.63 m s^{-1} to 22.87 m s^{-1} . The velocities falling in Region 2 were considered to be dealiased correctly while the velocities in Region 1 were considered to be dealiased incorrectly and fell in the category of catastrophic errors. The plot of successfully dealiased velocities is the percentage of simulated velocities that fall in Region 2. This allows the properly aliased velocities to be compared appropriately while still taking into account the effects of catastrophic dealiasing errors.

The results in Figure 3.9 show that the contiguous pulse strategy performs much better at a true spectrum width of 4 m s^{-1} than either of the staggered strategies. It is important to note that the contiguous pulse strategy has two parts: the longer PRT used to estimate reflectivity and the shorter PRT used to estimate velocity and spectrum width. In order for the contiguous pulse strategy to meet the requirements, the acquisition time to meet the reflectivity requirement of 1 dBZ using the long PRT is added to the maximum of the two short PRT acquisition times needed to meet the 1 m s^{-1} requirements for both velocity and spectrum width. In this case, the acquisition time to satisfy the requirement for reflectivity with the long PRT is 12.29 ms and for velocity and spectrum width with the shorter PRT is 38.91 ms. For the staggered strategies, the same set of samples are used to compute all three spectral moments so the acquisition time is the maximum acquisition time needed to meet all the requirements. For SBMX, the acquisition time is 172.03 ms and for staggered PRT, 133.12 ms. All of these results are captured in Table 3.1. The first two columns show the

collection strategy and the acquisition time. The next column is the percentage of successfully dealiased velocities. For the contiguous pulse strategy, there are no catastrophic errors. Both staggered strategies do reasonably well avoiding catastrophic errors; both percentages are over 90%.

Table 3.1. Simulation results for SBMX, staggered PRT, and contiguous pulse scanning strategies using a true spectrum width of 4 m s^{-1} .

Collection Strategy	Acq. Time	Succ. Deal.	% of Time	$\hat{\sigma}_z$	$\hat{\sigma}_v$	$\hat{\sigma}_w$
SBMX	172.03 ms	92.5%	336%	0.46 dBZ	0.98 m s^{-1}	0.71 m s^{-1}
Staggered PRT	133.12 ms	93.5%	260%	0.44 dBZ	0.99 m s^{-1}	0.69 m s^{-1}
Contiguous Pulse	51.20 ms	100%	100%	0.99 dBZ	0.997 m s^{-1}	0.75 m s^{-1}

Some of these catastrophic errors could be fixed through comparisons with the surrounding data, but further study is needed to determine what level of catastrophic errors could be corrected. The next column shows the percentage of time each strategy takes compared to the acquisition time for the contiguous pulse strategy. The last three columns show the standard deviations of the spectral moment estimates for each of the strategies. The staggered techniques need significantly longer to meet the velocity requirement of 1 m s^{-1} at 8 dB SNR. For SBMX, this is consistent with the results for BMX pairs where the best performance is at high SNRs and narrow spectrum widths. Apparently, the staggered PRT strategy is similar but takes less time than the SBMX strategy. The fact that staggered PRT has more sample pairs for the same acquisition time seems to compensate for the fact that the SBMX samples are less correlated.

In the first test of SBMX, it comes up short compared to both the staggered PRT and contiguous pulse strategies. It should be noted that the staggered strategies will have fewer overlaid echoes than the contiguous pulse strategies because they excel at eliminating

range/velocity ambiguities. Although the staggered strategies have more catastrophic errors, that may be balanced out by the fact that more valid velocity measurements will be recovered. It would be interesting to see how the staggered strategies perform at a narrower spectrum width since that should be better for a BMX strategy, and staggered PRT seems to have similar performance. There are other independent justifications for using a narrower true spectrum width for the comparisons. The requirements in the NTR were set for a particular true spectrum width and SNR that may not be representative of most weather echoes. One argument for using a narrower true spectrum width for the requirements comes from research done after the NTR requirements were defined. In Fang et al. (2004), it was shown that the median spectrum width of most weather echoes is closer to 2 m s^{-1} than 4 m s^{-1} . This means that setting the true spectrum width to 2 m s^{-1} may be more representative for a variety of weather events. This also suggests a possible issue with a requirements-based comparison that examines strategies under a single set of conditions. Some strategies may perform better depending on the particular weather phenomena of interest, and it may be difficult to find a small number of spectrum widths and SNR values that capture these varying conditions.

To study the performance of the strategies at 2 m s^{-1} , another simulation was run with the narrower true spectrum width. Figure 3.10 shows the standard deviations of reflectivity, velocity, and spectrum width for all three strategies with the true spectrum width set to 2 m s^{-1} . Not surprisingly, SBMX and staggered PRT perform significantly better than they did with a true spectrum width of 4 m s^{-1} . They also perform more comparably to the contiguous pulse strategy, although the contiguous pulse strategy continues to perform better for reflectivity and velocity. Just as in the 4 m s^{-1} case, the acquisition time for both parts of the contiguous pulse strategy need to be added together to find the total acquisition

time. This turns out to be 65.54 ms which is longer than the corresponding acquisition time of 51.2 ms for the 4 m s^{-1} case. Most of this additional time comes from reflectivity requirement which increases the time for the longer PRT collection by roughly a factor of two. As expected, the SBMX strategy now requires less time, 43.01 ms, to meet the new requirements, and the staggered PRT strategy requires only 30.72 ms. The staggered PRT strategy still outperforms the SBMX strategy even though the narrower spectrum width improves the acquisition times for both. The number of catastrophic errors also decreases significantly compared to the 4 m s^{-1} spectrum width case.

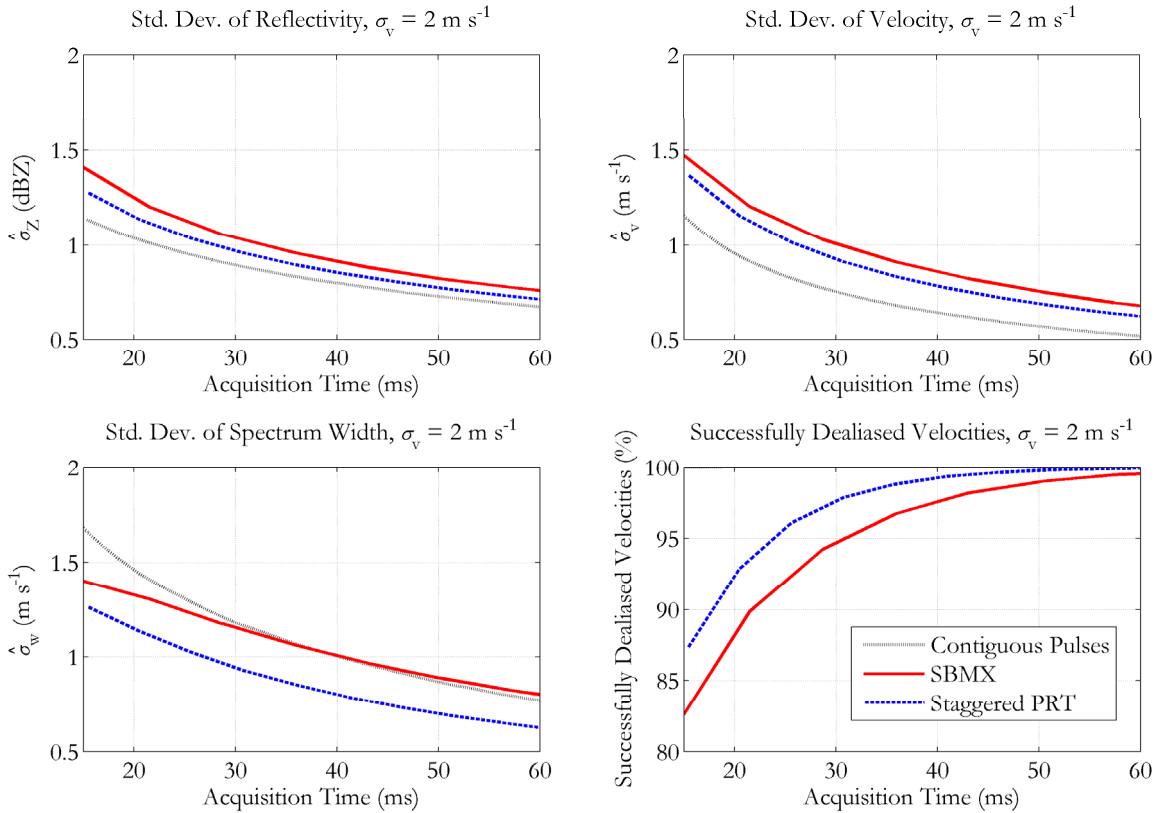


Figure 3.10. Standard deviations of reflectivity, velocity, and spectrum width and percentages of successfully dealiased velocities versus acquisition time for the contiguous pulse, staggered PRT, and SBMX collection strategies with a true spectrum width of 2 m s^{-1} . Acquisition time for the contiguous pulse strategy is the sum of the long PRT acquisition time to needed to meet the reflectivity requirement and the short PRT acquisition time need to meet the velocity and spectrum width requirements.

Table 3.2 captures the results from this latest simulation. It is surprising how a change of 2 m s^{-1} to the true spectrum width completely changes the relative performance of the scanning strategies. One difference compared to the 4 m s^{-1} case is that the spectrum width requirement drives the length of the acquisition times more than the velocity requirement.

Table 3.2. Simulation results for SBMX, staggered PRT, and contiguous pulse scanning strategies using a true spectrum width of 2 m s^{-1} .

Collection Strategy	Acq. Time	Succ. Deal.	% of Time	$\hat{\sigma}_z$	$\hat{\sigma}_v$	$\hat{\sigma}_w$
SBMX	43.01 ms	98.1%	65.6%	0.88 dBZ	0.82 m s^{-1}	0.97 m s^{-1}
Staggered PRT	30.72 ms	97.9%	46.9%	0.96 dBZ	0.91 m s^{-1}	0.93 m s^{-1}
Contiguous Pulse	65.54 ms	100%	100%	0.96 dBZ	0.63 m s^{-1}	0.99 m s^{-1}

The staggered PRT strategy also has an interesting property; its short acquisition time seems to rely on the fact that all three of the requirements are barely met. Staggered PRT seems to be well balanced for this particular set of parameters because acquisition time is not spent reducing errors significantly below the requirement level for any one of the spectral moments. In the case of the contiguous pulse strategy, extra acquisition time is needed to meet the spectrum width requirement, but the velocity requirement is surpassed by a relatively large margin. The catastrophic errors are also much less prevalent at the narrower spectrum width of 2 m s^{-1} .

Based on the results from the previous simulation and the requirement change of the true spectrum width from 4 m s^{-1} to 2 m s^{-1} , the staggered PRT collection strategy meets the modified requirements in less than half the time of the contiguous pulse strategy. The SBMX strategy needs a longer acquisition time than staggered PRT to meet the modified spectrum width requirement, but the acquisition time is still significantly shorter than the contiguous pulse strategy. The number of catastrophic errors for both staggered strategies is small

enough that there should be minimal impact under these particular conditions. The question then becomes, how will these staggered strategies perform over a wider range of conditions? Inherent in the definition of the NEXRAD requirements is the idea that performance at a true spectrum width of 4 m s^{-1} and an SNR of 8 or 10 dB also determines the performance at other spectrum widths and SNRs. The performance of contiguous pulse strategies over varying conditions is well-understood, and this overall performance can reasonably be measured by looking at one set of conditions. But looking at the results from the simulations, it is clear that the staggered PRT strategies do not follow the same performance assumptions. If that were true, the relative acquisition times among the strategies would not change under differing conditions. The best way to see this is to run a new simulation that fixes the acquisition times for the scanning strategies and varies either true spectrum width or SNR. In keeping with the first two simulations that looked at two values for the true spectrum width, a simulation that varies the spectrum width over a wider range of values while keeping the SNR values as defined in the requirements should help to better understand the differing performance characteristics of the scanning strategies.

Figure 3.11 shows the results of a simulation that varies the true spectrum width from 1 m s^{-1} to 8 m s^{-1} while fixing the acquisition times for the three scanning strategies. The acquisition times were chosen based on the values in Table 3.2 which correspond to strategies meeting the modified WSR-88D requirements utilizing a true spectrum width of 2 m s^{-1} . The staggered strategies perform very similarly to each other but differ significantly from the contiguous pulse strategy. The reflectivity errors are relatively well behaved over the whole range of true spectrum widths with SBMX performing better at narrow spectrum widths (less than 2 m s^{-1}) and the contiguous pulse strategy performing better at larger spectrum widths (greater than 4 m s^{-1}). The contiguous pulse strategy also performs better

when estimating the spectrum width at true spectrum widths greater than about 2 m s^{-1} . The greatest differences are seen in the velocity estimates. The errors of the contiguous pulse strategy are lower at all spectrum widths and are nearly twice as small over the range from 4 m s^{-1} to 8 m s^{-1} . Of greater concern is the fact that the staggered strategies have catastrophic velocity errors occurring over half the time at true spectrum widths larger than about 5 m s^{-1} . These levels of catastrophic errors would be very difficult to correct if widespread.

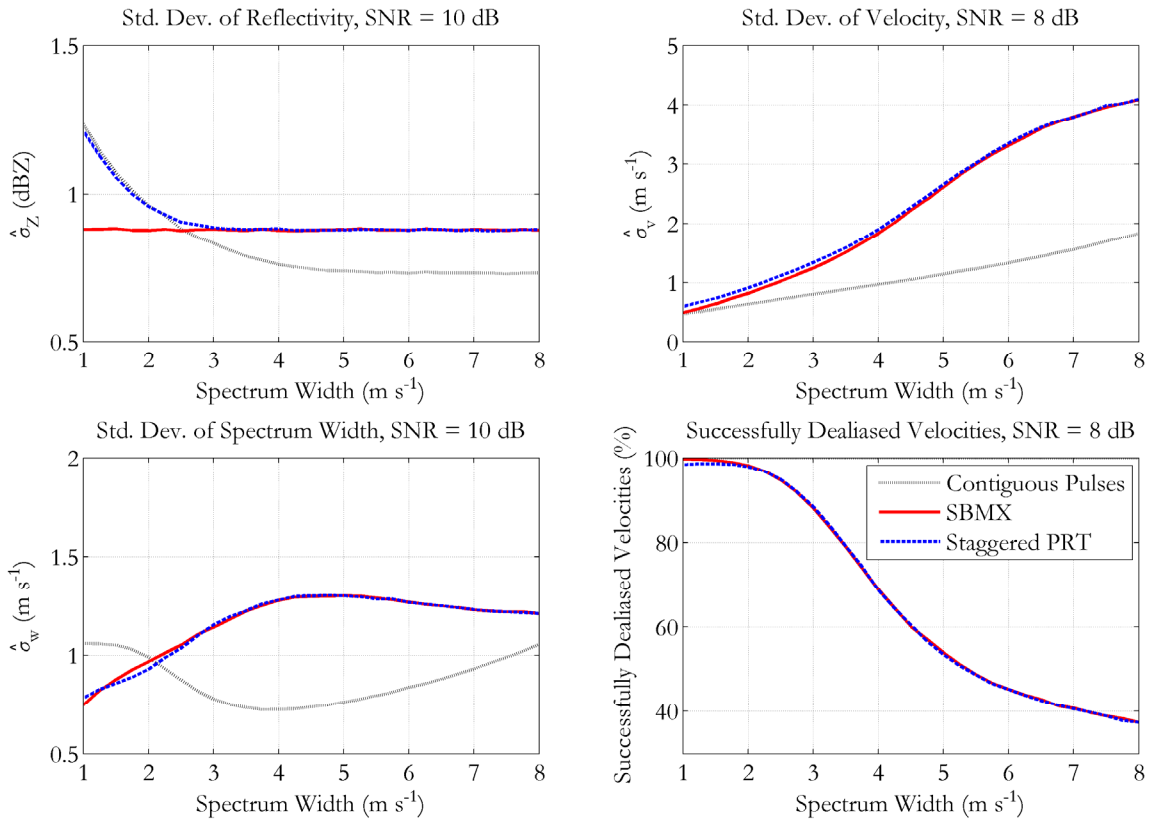


Figure 3.11. Standard deviations of reflectivity, velocity, and spectrum width and percentages of successfully dealiasd velocities versus true spectrum width for the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.

The differing shapes of the error curves also give insight into the differing performance of the scanning strategies. For example, the behavior of the velocity standard deviation helps to

explain the long acquisition times needed for the staggered strategies when using the original requirements based on a true spectrum width of 4 m s^{-1} .

The initial impression after looking at the performance of the three scanning strategies at different spectrum widths is that the contiguous pulse strategy outperforms the staggered strategies over a wide range of true spectrum widths. Although this is accurate, this initial impression does not tell the full story. The distribution of spectrum widths in actual weather data is also significant when comparing the strategies. As mentioned earlier, the median spectrum width for a wide range of weather conditions is close to 2 m s^{-1} . If the range of comparison was limited to true spectrum widths less than 3 m s^{-1} , this would correspond to significantly more than half of the data collected from all weather events. Over this more limited range, the staggered strategies perform similarly to the contiguous pulse strategy with some minor differences. SBMX performs much better when estimating reflectivity with true spectrum widths less than 2 m s^{-1} while staggered PRT is very similar to the contiguous pulse strategy when comparing reflectivity errors over this range. When estimating spectrum width, the staggered strategies excel at true spectrum widths less than 2 m s^{-1} while the contiguous pulse strategy does better between 2 m s^{-1} and 3 m s^{-1} . For velocity, the staggered strategies have larger errors over the range but not by a large amount. Finally, the percentage of successfully dealiased velocities is between 90% and 100% over this range for the staggered strategies which should be manageable for continuity-based error correction techniques. All of this leads to the conclusion that staggered strategies can compete with a conventional contiguous pulse strategy at smaller true spectrum widths while also cutting the acquisition time in half for the staggered PRT strategy. The staggered strategies also reduce the amount of overlaid echoes considerably compared to the more conventional range-unfolding approach used with the contiguous pulse strategy.

There is an additional issue with the staggered strategies when estimating spectrum width. Because the shorter of the two PRTs used in the staggered strategies is twice as long as the shorter PRT in the conventional strategy, the spectrum width “saturates” as the true spectrum width increases. This effect can be seen clearly in Figure 3.12. The spectrum width computed from contiguous pulses utilizing a PRT of 1024 μs has a small negative bias in some cases but successfully reflects the true spectrum width. The two staggered strategies perform similarly to each other with a relatively small negative bias at small spectrum widths and a significant negative bias at large spectrum widths.




Figure 3.12. Estimated spectrum width plotted versus true spectrum width for the contiguous pulse and staggered strategies.

For comparison, an additional case using contiguous pulses with a PRT of 2048 μs was included. This PRT matches the shorter PRT used to compute the spectrum width for the

staggered strategies. Although the saturation effect is evident for 40 contiguous samples, the negative bias is greater for the staggered strategies. This is due to the same bias mentioned earlier that is inversely proportional to the number of sample pairs used to compute the spectrum width. In that case the negative bias occurred at narrow spectrum widths, but the effects are even greater at large spectrum widths. Spectrum width estimates can be severely compromised when the true spectrum width is large, the PRT is relatively long, and the spectrum width is calculated from a small number of sample pairs.

In Figure 3.11, the effects of spectrum width on the performance of the three scanning strategies are illustrated, but the effects of varying SNR have been neglected. Because of the requirements-based perspective that has been used up to this point, the SNR has been fixed at 10 dB when looking at reflectivity and spectrum width errors and at 8 dB for velocity errors. To get a better overall picture of the estimation errors under different conditions, simulations at varying spectrum widths and SNR values were employed. The spectrum width was varied in increments of 1 m s^{-1} from 1 to 8 m s^{-1} , and the SNR was varied from 2 to 20 dB in 2 dB increments. The data are plotted in two dimensions with different colors signifying different error levels. In the first figure, Figure 3.13, the standard deviations for reflectivity and spectrum width are plotted. The second figure, Figure 3.14, shows the standard deviations for velocity along with the percentage of successfully dealiased velocities. The simulation parameters are the same as those used for Figure 3.11 except for the varying SNR values. The standard deviations for the row corresponding to 10 dB SNR also match the values from Figure 3.11.

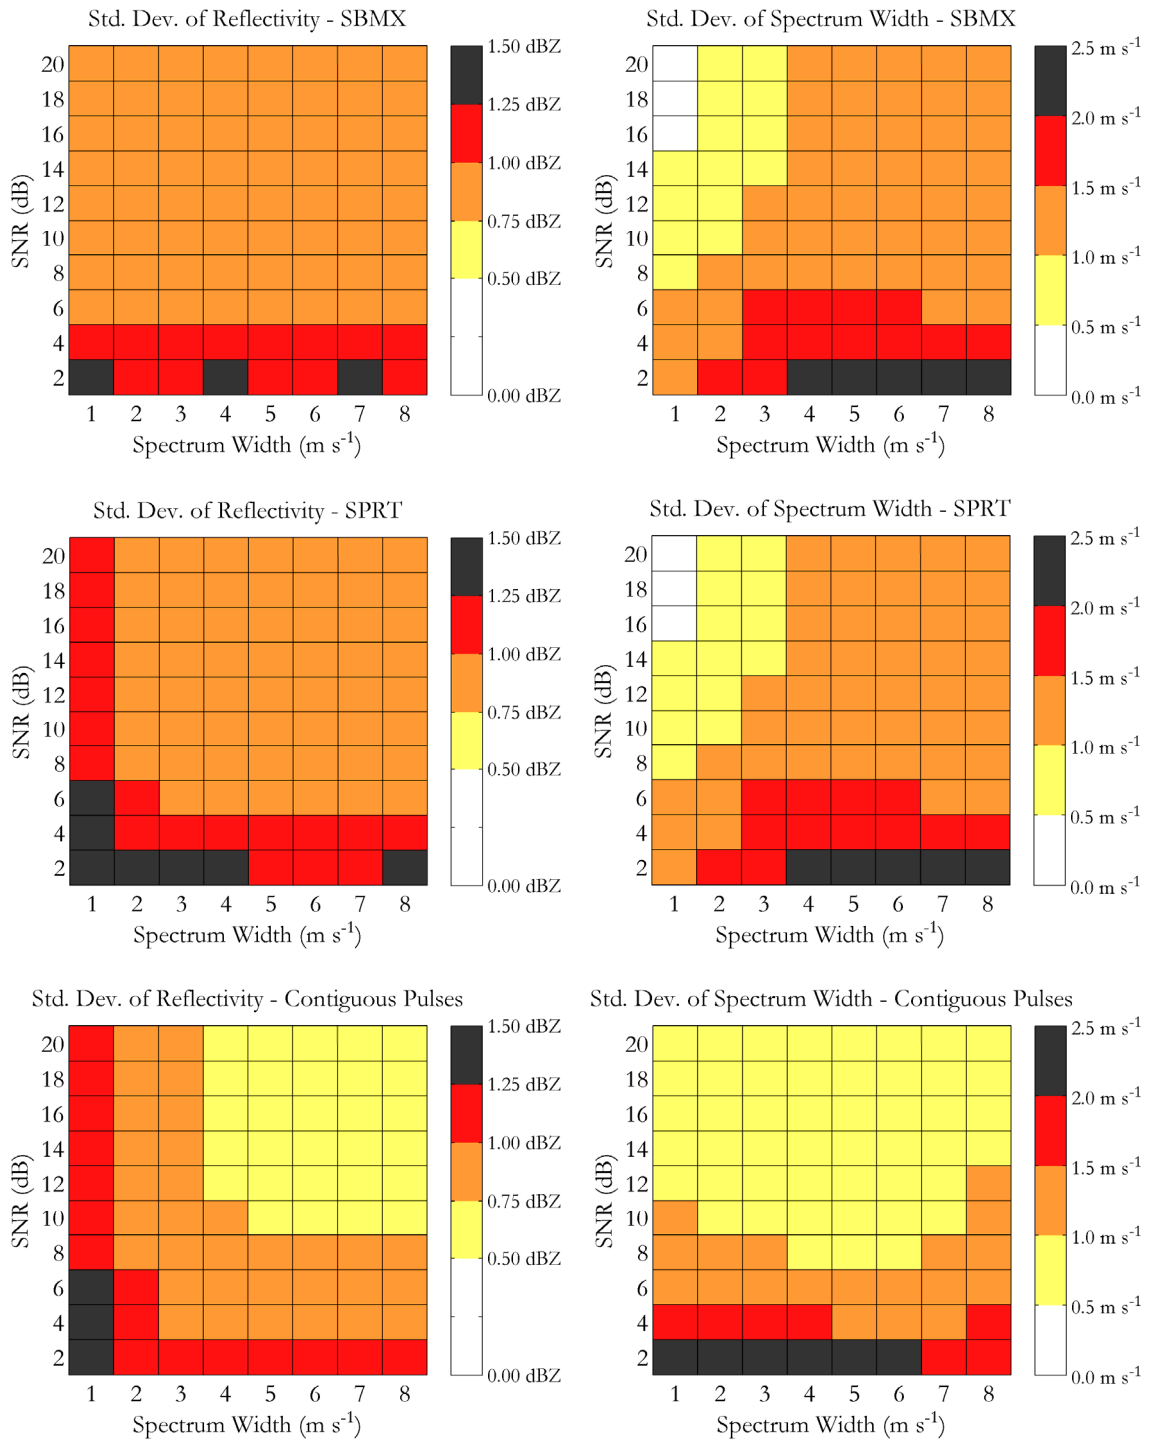


Figure 3.13. Plots of reflectivity and spectrum width standard deviations for an array of true spectrum width and SNR values. The scanning strategies used in the simulations are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.

For the reflectivity standard deviations, the same pattern that is seen at 10 dB SNR is also seen at other SNR values. The SBMX strategy performs the best at spectrum width values of 1 m s^{-1} and 2 m s^{-1} , and the contiguous pulse strategy performs better at larger spectrum widths. The contiguous pulse strategy has a greater number of samples which leads to lower standard deviations because samples that are close together in time are nearly independent because of the large spectrum widths. The two staggered strategies also seem to perform similarly for spectrum widths greater than or equal to 3 m s^{-1} . At 2 dB SNR, the staggered strategies have standard deviations close to 1.25 dBZ which leads to some of them falling just above 1.25 dBZ and some just below because of variability in the simulation results.

For the spectrum width standard deviations, the values at true spectrum widths less than 5 m s^{-1} are more relevant because of the saturation that occurs when using the staggered strategies with larger spectrum widths. At a true spectrum width of 1 m s^{-1} , the staggered strategies perform better than the contiguous pulse strategy especially at high SNRs, but they also have a small negative bias in this range because of the small number of sample pairs. At 2 m s^{-1} , all of the strategies perform similarly except that the staggered strategies perform somewhat better at low SNRs. The standard deviations for the contiguous pulse strategy are consistently better for large true spectrum widths, and the biases are also much smaller.

The velocity results are shown in Figure 3.14. The standard deviations are plotted in the left column, and the percentage of successfully dealiased velocities are shown on the right. The same method for determining catastrophic dealiasing errors that was described earlier was used for all of these simulations. The results that were shown for the 10 dB SNR case seem to hold for SNRs greater than 10 dB. At smaller SNRs, the standard deviations increase as expected, but the relative performance of the strategies stays the same.

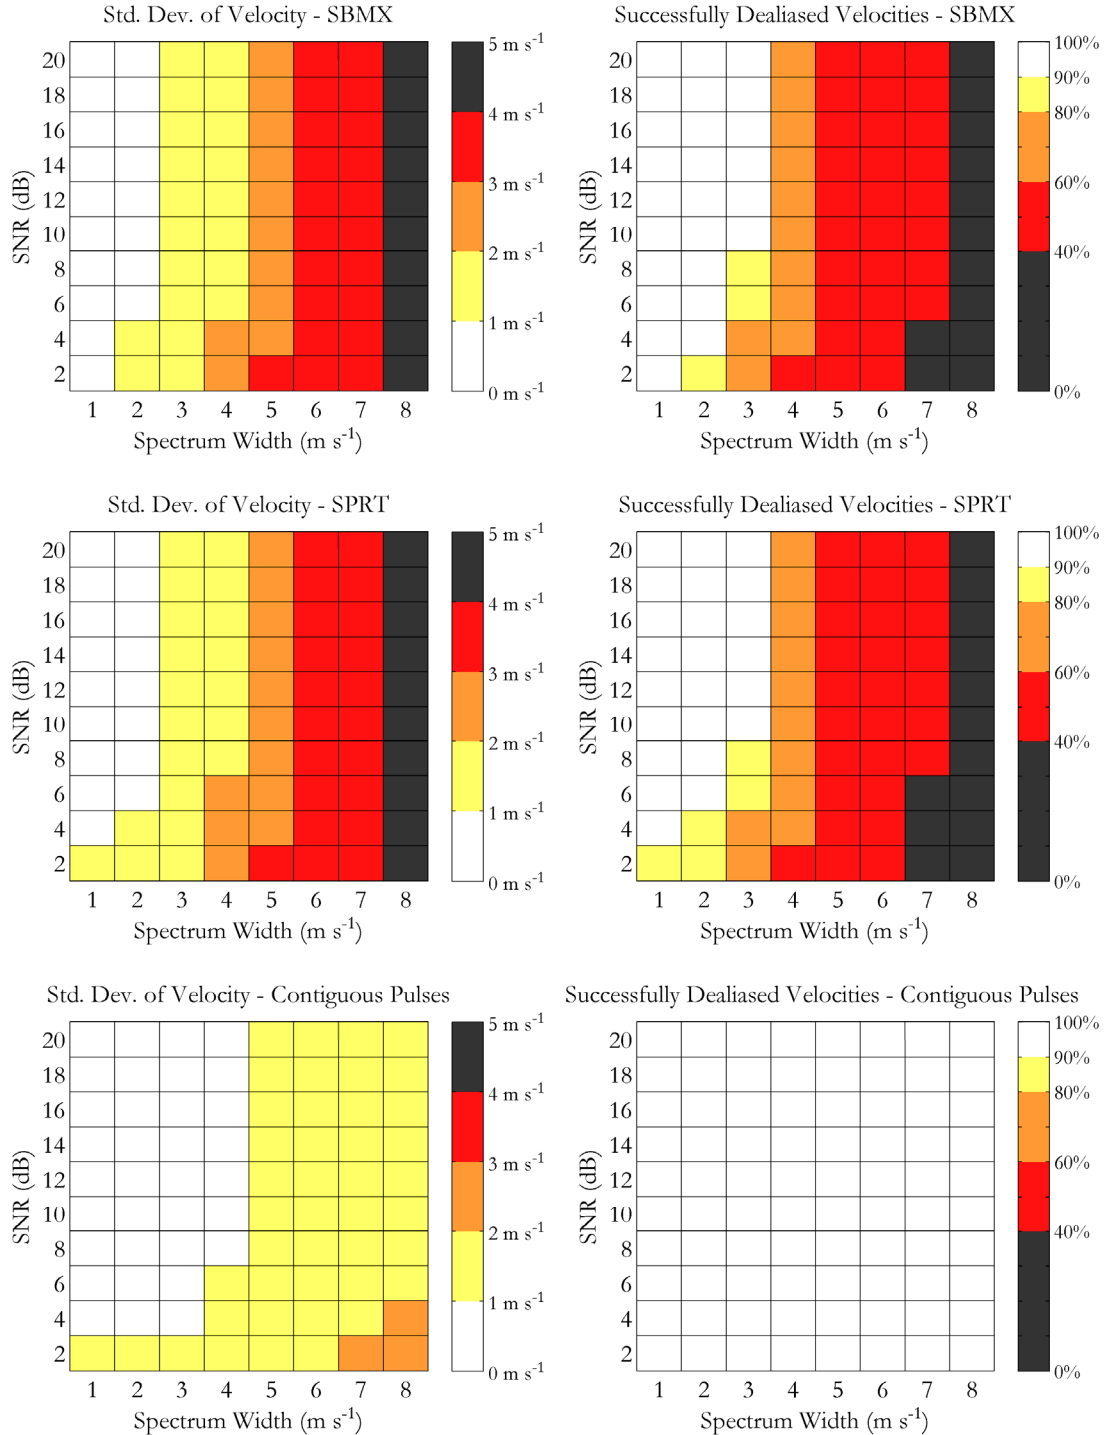


Figure 3.14. Plots of velocity standard deviations and percentage of successfully dealiased velocities for an array of true spectrum width and SNR values. The scanning strategies used in the simulations are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.

The staggered strategies have large errors at large spectrum widths, and the percentage of successfully dealiased velocities approaches what would be expected from noise (33%). Staggered strategies do not perform well in these situations when such a small number of sample pairs are utilized; better performance would require longer acquisition times. Large errors could possibly be tolerated if there was a relatively low incidence of large spectrum width values. Although the staggered strategies are competitive with the contiguous pulse strategy at narrower spectrum widths, SBMX only outperforms staggered PRT when estimating reflectivity over a small range of spectrum width values. SBMX also takes 40% longer than staggered PRT to achieve this similar performance. Based on these results, SBMX does not look like a viable candidate for a general weather collection strategy. If good performance at large true spectrum widths is required, the contiguous pulse strategy produces significantly lower errors than the staggered strategies. If performance at narrower spectrum widths is required, staggered PRT performs comparably to the contiguous pulse strategy with an acquisition time less than 50% of the contiguous pulse acquisition time.

One bright spot for BMX occurs when estimating reflectivity at narrow spectrum widths, but the overall performance is almost identical to staggered PRT while still having a longer acquisition time. The additional time between samples for SBMX produces some benefits, but those benefits do not outweigh the shorter acquisition time of staggered PRT. Apparently, the increase in the time between pairs used to compute velocity from 1.024 ms for the contiguous pulse strategy to 5.12 ms for the staggered PRT strategy provides most of the benefits of beam multiplexing along with a similar shape to the performance curves. If a staggered strategy is needed for general weather collection, staggered PRT should be a strong contender. Staggered PRT performs like SBMX but also has associated clutter filtering techniques that have reasonable performance (Sachidananda and Zrnić 2002).

3.6 Weather Data Results

In addition to studying the performance of these strategies using simulated data, weather data were also collected in July, 2007 using the NWRT to verify that the results from simulations can be applied to weather data. To avoid the difficulties of comparing scanning strategies collected at different times, a method similar to the simulation method was used for processing the weather data. The beam was pointed in a particular direction, and a few minutes of data were collected using a fixed PRT of 1024 μ s. Figure 3.15 shows two 90° sectors collected using the NWRT close to the times of the experiments. The red line shows the direction the beam was pointing when the experiment data were collected.




Figure 3.15. Two 90° sectors collected on July 5, 2007 at 21:30 UTC (left) and July 13, 2007 at 14:27 UTC (right) close to the times of the experiments.

Approximately 130 s of data consisting of 126976 samples were processed as 62 groups of 2048 samples. Each group of 2048 pulses took about 2.1 s to collect which satisfies the 2.3 s constraint introduced earlier for the total dwell time. This means that the stationarity assumption should hold for each of the groups of samples, but stationarity cannot be

assumed over the whole 130 s collection. Rather than attempting to remove a linear trend in the moment estimates because of changes in the moments over time (Yu et al. 2007), another approach was utilized. Both the mean and variance were estimated over the shorter 2.1 s periods where the stationarity assumption holds, and the 62 means and variances were then averaged to obtain representative values for the complete 130 s period. The steps for processing the data are discussed in more detail in the following paragraphs.

DC bias and noise values were computed from parts of the time series data without weather echoes. Point targets were used to calculate a simple digital matched filter of length four to produce range bins of length 240 m which is close to the pulse length of 250 m. The antenna was pointed in a direction so as to try to avoid second trip echoes, and the first 40 km of data were not processed to avoid ground clutter contamination.

The July 13, 2007 case is examined first. For the contiguous pulse case, 51 sets of 40 samples were utilized from each of the 2048-sample groups to estimate velocity and spectrum width, and 51 sets of 9 samples with an effective PRT of 3072 μ s were processed to estimate the signal power. Similarly, 73 sets of samples were used to estimate SPRT moments, and 56 sets of samples were used to estimate SBMX moments while employing the 30.72 ms and 43.01 ms versions of the collection strategies, respectively. The variances for each of the moments for each of these sets were calculated, and then the 62 variances were averaged to estimate the overall variance for each of the moments and for each scanning strategy. Although the 62 estimate means changed over the data collection time, the averaged variance should be representative of the variance of a weather signal with characteristics of the mean over the data collection period. To compute the mean value at each range bin, all 2048 pulses were used to estimate the moments and then averaged to find a “true” mean value. These “true” estimates of the signal power were then used to calculate

the SNR for each range bin. The estimates of the SNR and spectrum width were then used to predict the standard deviation for each of the scanning strategies using arrays of standard deviations derived through simulations (similar to the ones plotted in Figures 3.13 and 3.14 except the SNR was extended to 40 dB). Bilinear interpolation was used to find the standard deviation corresponding to the particular SNR and spectrum width.

Figure 3.16 shows the mean values for the scanning strategies in the left column and the standard deviations in the right column. Only every fifth standard deviation derived from the simulations is plotted so that the underlying estimates from the real data can be seen. The mean signal power from the three strategies agrees very well, and the calculated reflectivity standard deviations match the standard deviations predicted from the simulations. The SBMX strategy has the best performance which is expected because the spectrum width varies from a little less than 1 m s^{-1} to a little less than 2 m s^{-1} . The measured noise level is very close to 0 dB, so the mean signal power is also a good estimate for the SNR. The high SNR also leads to a strong performance for the SBMX strategy. The performances of the staggered PRT strategy and the contiguous pulse strategy are similar.

The mean velocity values from the three strategies also match up well, but some of the largest discrepancies between the standard deviations estimated from the simulation and the calculated standard deviations occur for velocity. The relative performance is well-predicted from the simulations, but the standard deviations calculated from the real data are larger for both the staggered PRT and contiguous pulse strategies. This is most likely due to deviations from the Gaussian spectrum assumption (Yu et al. 2009). Figure 3.17 shows a sample spectrum from 54.8 km in range; the spectrum was calculated using Welch's method from a length-2048 sequence using a spectrum length of 256 and a Von Hann window.

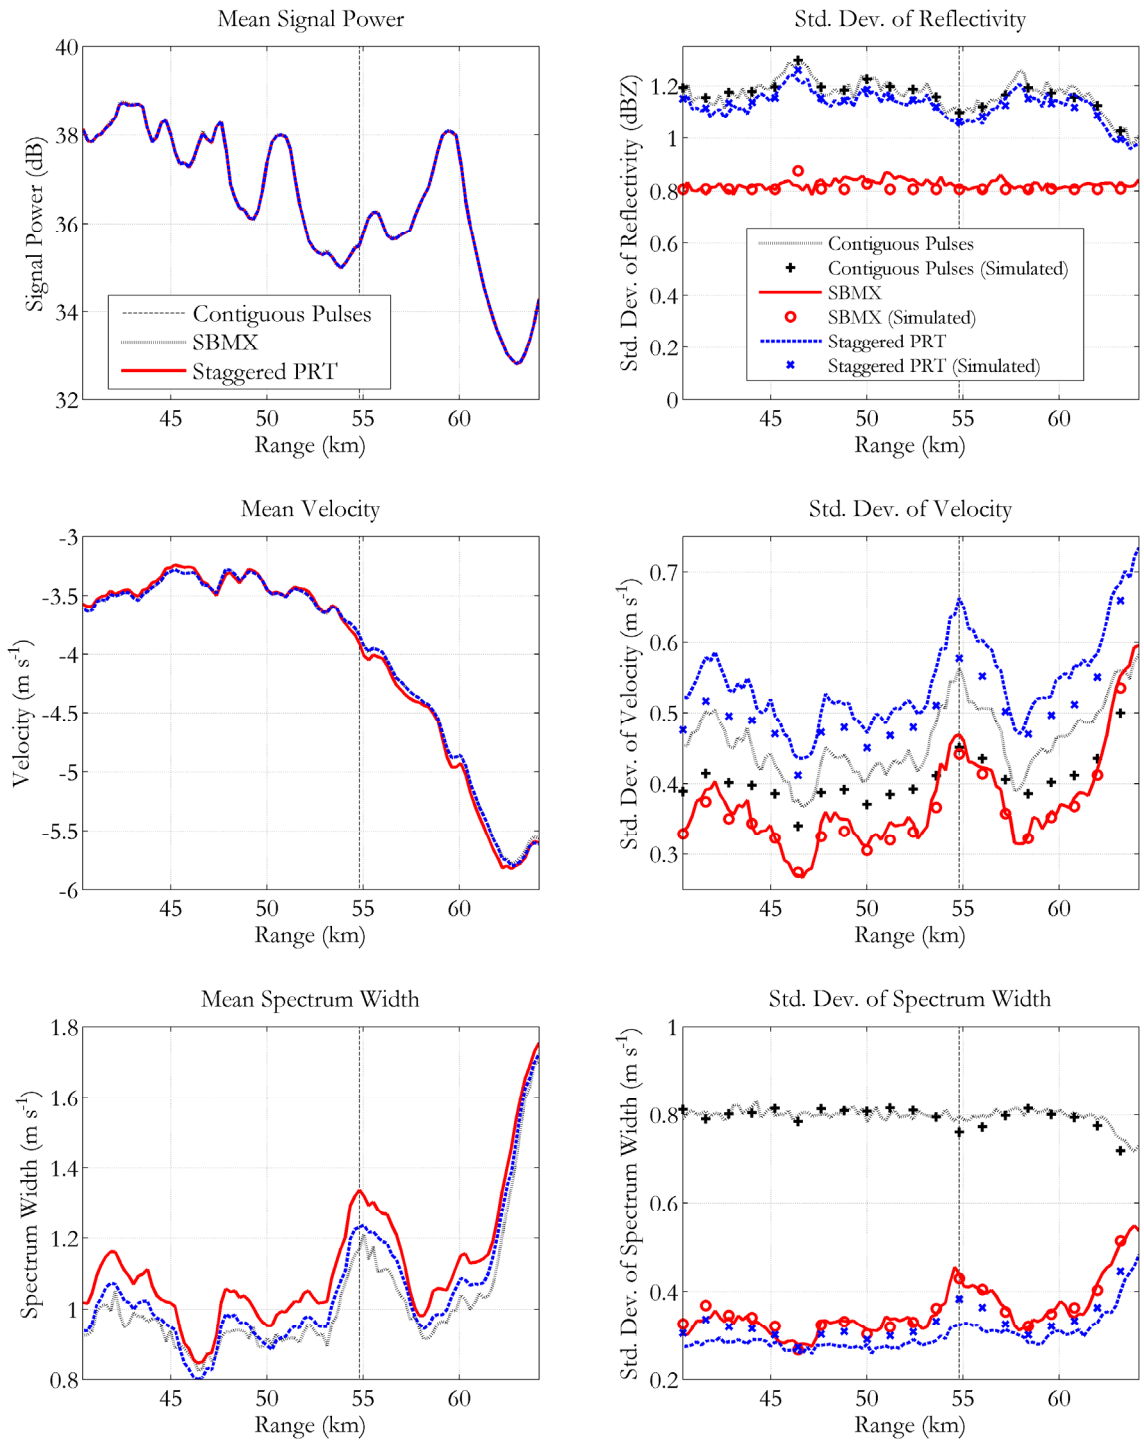


Figure 3.16. Means and standard deviations from weather data collected on July 13, 2007 at 14:23 UTC. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies. The dashed vertical line delineates the range of the spectrum plotted in Figure 3.17.

The spectrum is non-symmetric and departs from a Gaussian shape (a parabola in the log domain). It is interesting to note that the SBMX strategy does not show the same discrepancies between the calculated and simulated standard deviations. The biggest difference between the SBMX and staggered PRT strategies is the small number of samples collected at a time for the SBMX strategy which may explain the differing behavior.




Figure 3.17. Sample non-Gaussian spectrum from a range of 54.8 km for the July 13, 2007 experiment.

The mean spectrum width shows the largest discrepancies for any of the estimated mean values, but the differences are still less than 0.2 m s^{-1} . The spectrum width values for SBMX are puzzling since a negative bias was predicted from the simulations compared to a smaller or no negative bias for the contiguous pulse strategy. This result may also be due to a departure from the Gaussian shape used in the simulations. The standard deviations for both

the SBMX and contiguous pulse strategies closely match the values predicted from the simulations. The performance of the staggered PRT strategy is the best of the three as predicted from the simulations, but the actual estimated standard deviation is smaller than the simulation value in almost all of the cases.

Because of the high SNR and narrow spectrum width, the percentages of successfully dealiased velocities were not plotted for this case. Instead, time series data from July 5, 2007 were processed that had a wider variety of SNR and spectrum width values. The mean values and standard deviations are shown in Figure 3.18. The mean signal power values again match up very well, but standard deviations for reflectivity predicted from the simulations are underestimated from 45 to 52 km. The SBMX predictions are not affected over this range and match the predicted values closely as in the previous case. The SBMX strategy performs the best when the spectrum widths are narrow as predicted, and all three strategies perform similarly when the spectrum widths are large.

For velocity, the predicted standard deviations underestimate the calculated standard deviations for a range from 40 to 50 km. After looking at some spectra from this case, the underestimates occurred because of the non-Gaussian shape of the spectra. A sample spectrum from 45.5 km is shown in Figure 3.19. The non-Gaussian spectra also seem to affect the mean velocity values because the velocities for the staggered strategies extend from close to -3 m s^{-1} up to 9 m s^{-1} while the contiguous pulse velocities only extend from less than -1 m s^{-1} to less than 7 m s^{-1} . This could also be a result of the dealiasing algorithm for the staggered strategies and the effects of the catastrophic error methodology. This is one case where the SBMX standard deviations from the simulations do not match the calculated values, but this can be attributed to the “saturation” of the spectrum widths. At ranges above 50 km, both the mean velocities and the standard deviations match up well.

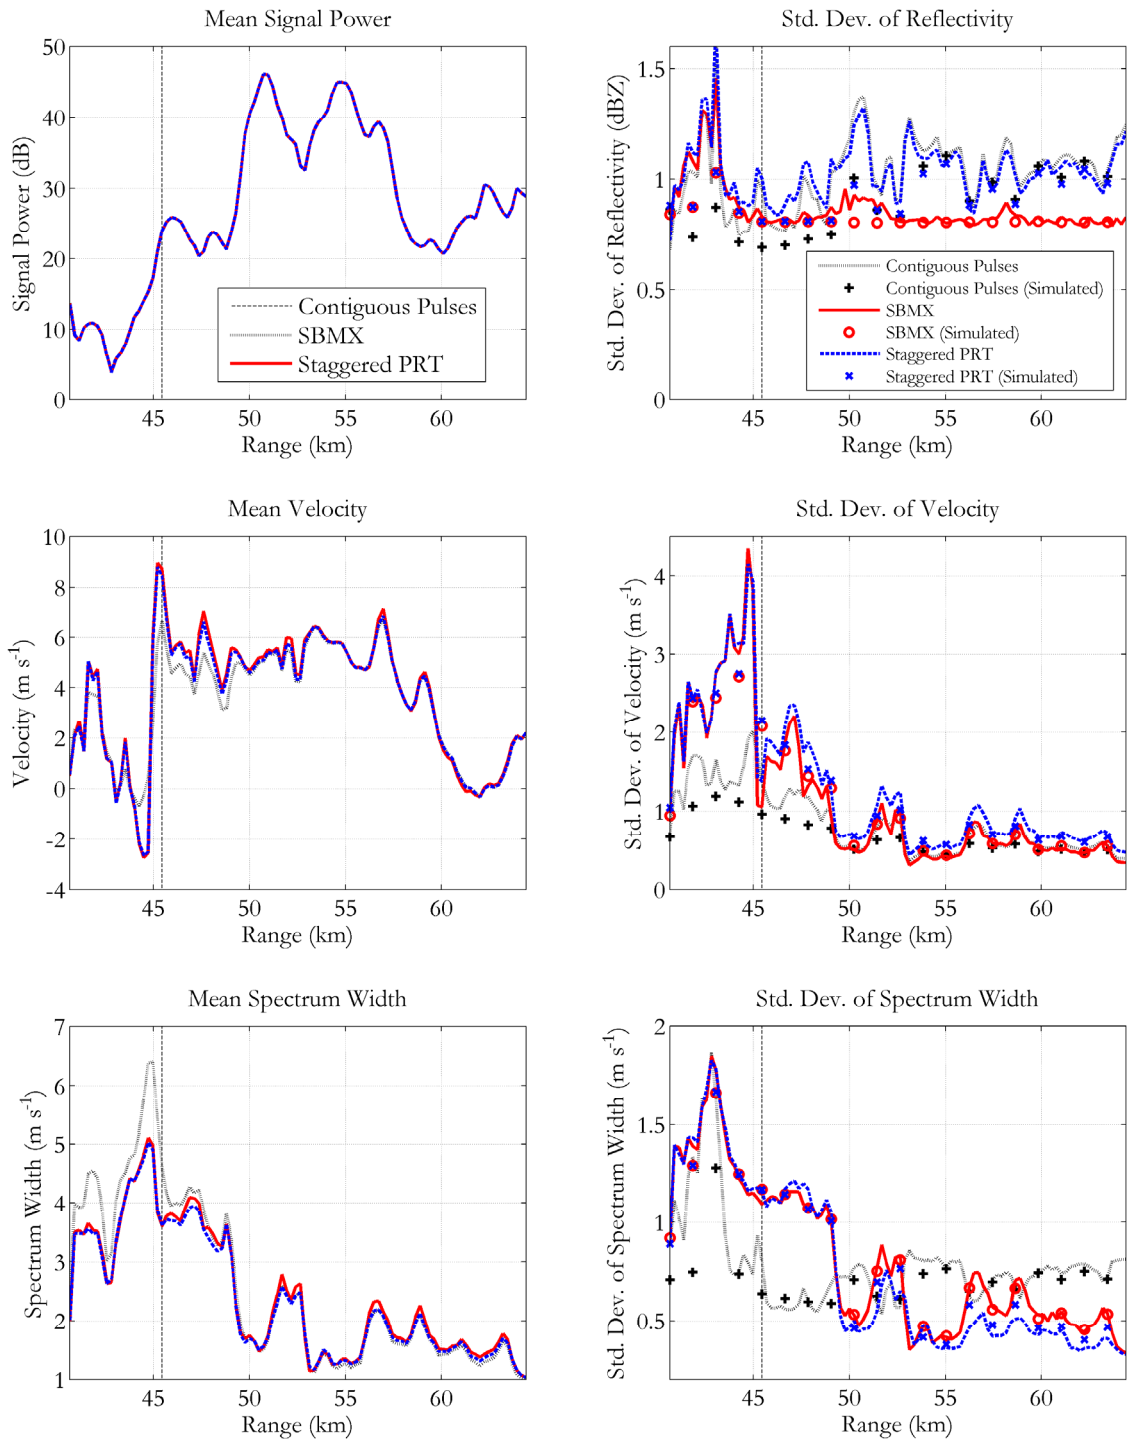


Figure 3.18. Means and standard deviations from weather data collected on July 5, 2007 at 21:34 UTC. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies. The dashed vertical line delineates the range of the spectrum plotted in Figure 3.19.




Figure 3.19. Sample non-Gaussian spectrum from a range of 45.5 km for the July 5, 2007 experiment.

The mean spectrum widths clearly show spectrum width “saturation” for the staggered strategies. The spectrum width for the staggered strategies only reaches a maximum of 5 m s^{-1} while the contiguous pulse strategy reaches a maximum close to 6.5 m s^{-1} . The mean spectrum widths for all three strategies match well beyond about 48 km. The standard deviation arrays produced by the simulations successfully predict the standard deviations calculated from the weather data. The crossover point near the 50 km mark shows up clearly in both the simulated and real data. As noted in most of the cases, the SMBX predicted values are closer to the calculated values. The SBMX strategy seems to be significantly less sensitive to spectral shape than the other scanning strategies at least in both of the cases examined in this section.

Because of the large spectrum widths, catastrophic errors should be much more common in this case than in the first case that was examined. The last figure, Figure 3.20, shows the percentage of successfully dealiased velocities for all three strategies both predicted from the simulations and calculated from the real data.




Figure 3.20. Percentage of successfully dealiased velocities computed from weather data collected on July 5, 2007 at 21:34 UTC along with percentage computed from simulations. The scanning strategies are the contiguous pulse (65.54 ms), staggered PRT (30.72 ms), and SBMX (43.01 ms) collection strategies.

For the real data, the velocity was considered to be successfully dealiased if it was in the region that was $\pm v_a/3$ from the velocity computed from all 640 samples. Also, every other simulated value was plotted for this figure instead of every fifth value. The results match up well with some minor discrepancies throughout. The deep valley in the percentage of

successfully dealiased velocities for the staggered strategies was well-predicted. This is an additional verification that short collection times for the staggered strategies lead to large numbers of catastrophic dealiasing errors when the spectrum width is large.

3.7 DC Removal for SBMX Data

This section combines the techniques for simulating and processing SBMX data with simulated data from the quadratic clutter model discussed in Chapter 2 to examine the effectiveness of the DC removal clutter filtering approach. DC removal should be more effective with a stationary antenna because the clutter width is narrower than for a rotating antenna, but the increase of clutter width with wind speed will limit the effectiveness of the filtering technique. Ground clutter data from the “Full Foliage, Prairie” case (which were already presented in Chapter 2) were processed to compare the results of real data to the simulation.

In the simulation, 4096 samples were utilized for each realization which is consistent with the number of samples collected at each beam position in Chapter 2. This allows the same processing to be used for both simulated and real data. To match the PRT of the real data, a simulation PRT of 800 μ s was used instead of 1024 μ s. The 43.01 ms SBMX strategy with six three-pulse collections utilized in the previous section will also be used for this simulation, but the acquisition time will decrease to 33.6 ms because of the change in PRT. For each 4096-sample realization, 112 simulated SBMX collections were processed. There was 268.8 ms between the three-pulse collections for a total dwell time of 1.69 s. For each SBMX collection, the DC component was computed by averaging all 18 IQ samples. The unfiltered power was computed from the original data, and the filtered power was computed after the DC component was subtracted from all 18 samples. Next, the ratio of the powers

was converted to dB to measure the amount of power removed. These values were averaged to produce a final estimate of the power removed for the realization.

To better understand the dependence of removed power on wind speed, 5000 realizations were simulated at each wind speed from 1 to 17 mph in 1 mph increments using the quadratic clutter model. The parameters for the “Full Foliage, Prairie” conditions were used because those conditions produce a wide range of clutter widths over the wind speed range. The data were simulated without additional noise so the spectral floor is determined by the simulated phase noise. This results in the widest possible clutter width and corresponds to strong ground clutter targets where the phase noise is larger than the system noise. After 5000 estimates of the removed power were computed based on the simulated data, the mean and standard deviation were calculated. Approximately 68% of the estimates should fall within a range that is plus or minus one standard deviation from the mean because the distribution of the 5000 estimates is roughly normal. This range was used to show whether the removed power values estimated from the real data were reasonable. For each of the three spectra shown in Appendix D for the “Full Foliage, Prairie” conditions, the 4096 samples of time series data used to produce the spectra were processed identically to the simulated data to produce three estimates of the power removed for the “Light,” “Breezy,” and “Windy” conditions, 11.9 dB, 5.8 dB, and 1.8 dB, respectively. Figure 3.21 displays the results of these computations for both real and simulated data.

The DC removal technique effectively removes over 50 dB of the power when the wind speeds are 1-2 mph. The amount of power removed is constant for 1-2 mph wind speeds because the model produces a single spike at DC. The amount of power removed quickly drops with wind speed so that only approximately 20 dB is removed at a wind speed of 4 mph.

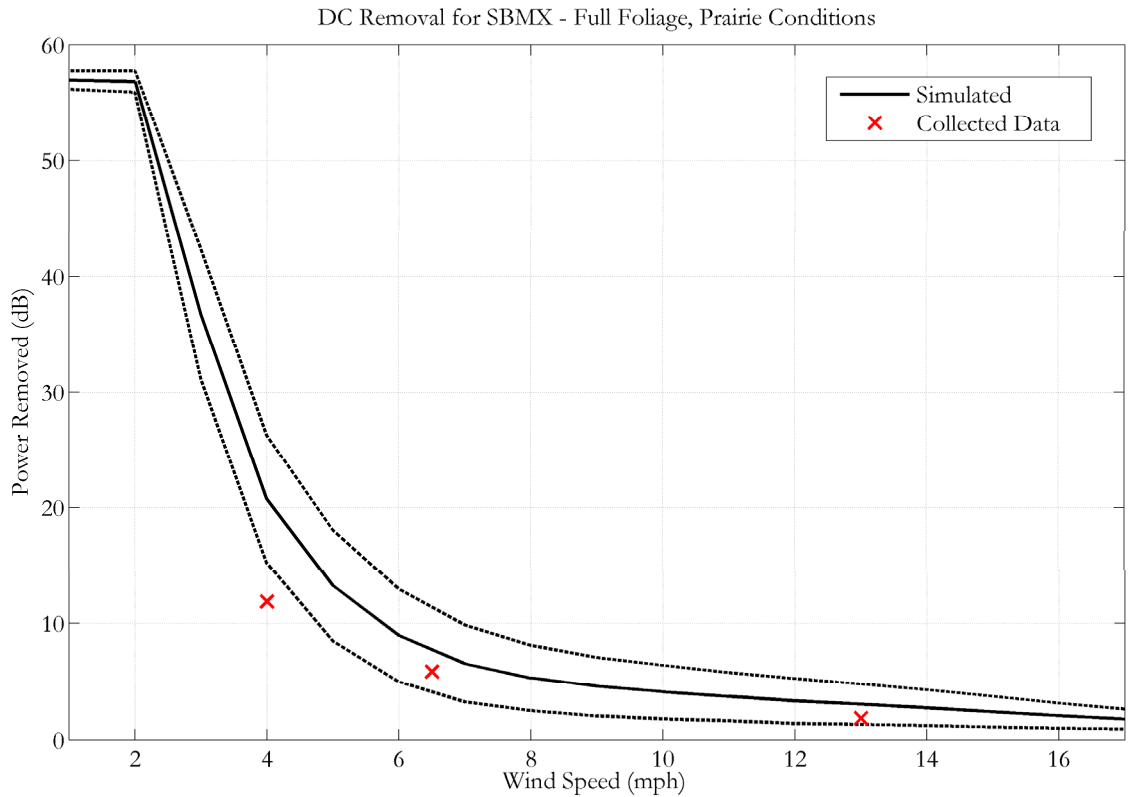


Figure 3.21. Dependence of amount of power removed using DC removal on wind speed. Data simulated using the quadratic clutter model to match collected data from the “Full Foliage, Prairie” conditions. The dotted lines are one standard deviation from the mean of the simulated data.

At wind speeds greater than about 6 mph, less than 10 dB of power is removed on average. The estimates corresponding to the three spectra from the “Full Foliage, Prairie” conditions are also plotted versus estimated wind speed. All three of the estimates fall below the predicted average power removed, and the “Full Foliage, Prairie, Light” case falls below the range that includes approximately 68% of the simulated data. The overall trend matches the simulated data because less power is removed as the wind speed increases. One possible explanation for the smaller amount of power removed in the “Light” case could come from additional power outside of the main clutter signature near zero velocity. Figure 3.22 shows the sample spectrum from the “Full Foliage, Prairie, Light” case over the whole Nyquist interval.




Figure 3.22. Sample spectrum from the “Full Foliage, Prairie, Light” case displayed over the whole Nyquist interval.

There is a significant amount of power from slightly over 0 m s^{-1} to approximately 18 m s^{-1} . A large portion of this power will not be removed by the DC removal technique which would lead to a lower than expected value for the removed power for this case. If this additional power is due to a weather signature, then retaining most of it after DC removal results in successful clutter filtering. If the additional power is due to some kind of moving clutter, the DC removal will not be effective, but other clutter filters would also have difficulty. Overall, the DC removal technique does not perform well enough for wind speeds over about 2 mph to be used as an effective clutter filter for SBMX data. As discussed in Section 3.2, the lack of an effective clutter filter is a significant limitation of BMX scanning strategies.

3.8 Conclusion

Beam multiplexing (BMX) was introduced to illustrate one way to take advantage of the electronic scanning capabilities of a phased array weather radar. A simple BMX scanning strategy called BMX pairs was described which has already been implemented on the NWRP. This strategy was shown to be feasible on a phased array, but several limitations were discussed. The main limitation specific to BMX pairs is that it is difficult to recover velocity and spectrum width estimates from the second trip. To address this limitation, a new strategy was developed called staggered PRT beam multiplexing or SBMX. This strategy combines the use of staggered PRT data collection, velocity dealiasing, and beam multiplexing to enable the recovery of velocity and spectrum width estimates at longer ranges compared to BMX pairs.

Simulations were then used to compare conventional contiguous pulse, staggered PRT, and SBMX scanning strategies. A requirements-based approach was implemented that looked at all of the scanning strategies in terms of the NTR requirements (NTR 1991). Using the NTR spectrum width value of 4 m s^{-1} , the staggered strategies needed significantly longer acquisition times to meet the requirements than the contiguous pulse strategy. More recent research shows that the median spectrum width for a large proportion of weather events is closer to 2 m s^{-1} (Fang et al. 2004). After running the simulations with the new spectrum width value, it was shown that the staggered strategies met the modified requirements with significantly shorter acquisition times. One key observation is that the velocity errors of estimates increase with spectrum width for all three strategies, but the errors increase faster for the staggered strategies than for the contiguous strategy. A limitation of NTR-type requirements is that the requirements are only defined for one set of conditions. When only one type of scanning strategy is utilized, the performance under other conditions can be

predicted relatively well based on the performance curve for that strategy. The addition of scanning strategies that have performance curves with different shapes (like the staggered strategies) causes the performance to vary at conditions other than the set used to define the requirements.

After looking at performance over a wide range of SNRs and spectrum widths, it was determined that the staggered PRT strategy performs comparably to or better than SBMX under a wide range of conditions with acquisition times shorter than SBMX. The only clear strength of beam multiplexing came at narrow spectrum widths especially when estimating reflectivity. In most cases, staggered PRT is the scanning strategy of choice over SBMX if a staggered strategy is being employed. The contiguous pulse strategy performs better at larger spectrum widths when utilizing an acquisition time over twice as long as the staggered PRT strategy used in most of the simulations. The contiguous pulse strategy avoids the “saturation” problem for spectrum width because of the use of a shorter PRT, but the range unfolding used to mitigate ambiguities results in much more obscuration or “purple haze” compared to the staggered strategies.

A simple DC removal clutter filter was implemented for the SBMX strategy, but it is only effective at very low wind speeds. There are no known clutter filters for SBMX that are as effective as the well-established approaches used for contiguous pulses. Fortunately, there is an effective clutter filtering technique for the staggered PRT strategy (Sachidananda and Zrnić 2002). An issue that effects the implementation of a clutter filter is the number of samples that need to be collected for the filter to successfully be applied. For example, the contiguous pulse strategy implemented for the simulations only used 9 samples for the reflectivity computations. In the operational strategies, at least 15 or 16 pulses are collected at low elevation angles when clutter filtering is being utilized. It would be difficult to apply a

spectral clutter filter such as GMAP to only 9 samples. Before implementing a staggered PRT scanning strategy, the minimum number of samples needs to be determined to ensure that the clutter filter is effective. Clutter filtering can be an additional constraint on the number of samples required in addition to requirements on the errors of estimates of the moments.

Although SMBX has several limitations that keep it from being optimal for general weather collection, the limitations should not preclude the search for other possible BMX strategies. As mentioned previously, BMX performs especially well when measuring reflectivity at small spectrum widths. This could be very useful in a scanning strategy for detecting weather signatures which utilizes only a few pulses to compute reflectivity. The capability of a phased array radar to use different scanning strategies at different beam positions also allows a lot of flexibility. Beam multiplexing should be seen as one of many tools when using a phased array radar and its utility should be assessed when new problems arise.

4. Strategies for Weather Detection Using a Phased Array Radar

4.1 Introduction

Phased array technology opens up several new possibilities for scanning compared to conventional parabolic dish antennas. Beam multiplexing has already been discussed in Chapter 3 as a technique that is feasible with a phased array antenna, but one of the most promising possible advances using the phased array comes from “adaptive scanning.” In this context, adaptive scanning is divided into two major areas: the first is the ability to collect or not collect data at particular beam positions, and the second is tailoring the waveform at each beam position in an optimal way. This chapter will focus on strategies that address the first of these areas. Although some techniques are already being implemented on conventional radars to attempt to reap the benefits of adaptively changing the scanning strategy (Ray and Chrisman 2008), the limitations of moving a dish antenna preclude true adaptive scanning. There are many aspects of adaptive scanning to study because stationary phased arrays have only recently been utilized to collect weather data (Heinselman et al. 2008). A large amount of work has been done using adaptive scanning for point targets on military phased array radars (Lee et al. 2003), but detecting and tracking weather scatterers is different. This chapter will address some of the issues that are encountered when developing detection scanning strategies for weather scatterers.

The first section will investigate detecting weather signatures in terms of probability of detection (POD) and probability of false alarm (PFA). Probability of detection is the probability of detecting a signal in noise. Probability of false alarm is the probability of noise being falsely detected as a signal. These are the standard tools of detection theory and can be applied in a straightforward way for weather detection. The second section will examine

issues with sine space scanning for the phased array. The broadening of the beam as it is steered away from broadside is unique to phased array antennas and has an effect on detection scanning strategies. The third section addresses the overall feasibility of detection scanning strategies and some additional related issues.

4.2 Detecting Weather Signatures

One way that phased array weather radars can be utilized to determine where to spend most of the data collection time is to quickly scan the volume at regular intervals and only collect additional data at beam positions which are thought to contain weather. Additional data may need to be collected at nearby beam positions to ensure that the weather has not advected since the original detection scan. These detection scans need to be very fast so that the time spent on them is balanced by a reduction in data collection time in areas that do not require additional data collection. In contrast, conventional weather radars tend to scan a volume while transmitting a predetermined number of pulses to ensure that the estimates of the spectral moments meet certain error requirements. Deciding which weather returns are significant for each processed radial is also a type of detection problem. A possible approach for a phased array antenna is to quickly scan the volume using a detection scan and then collect additional data to meet the error requirements. In this section, some of the methods used with conventional antennas to determine if weather returns are significant (or detected) will be extended to detection scans that only transmit a small number of pulses at each beam position. The detection scans and conventional scans have a different purpose, but the work on conventional weather scanning can still give insight into some of the basic detection issues involving weather radars.

The NEXRAD network of WSR-88D radars uses a simple power threshold to determine whether the radar return in a particular range bin is significant or not. One of the nice features of the power threshold is that the threshold for a particular PFA can be computed analytically if the noise is assumed to be white and Gaussian. For a conventional radar such as the WSR-88D, contiguous pulses are processed because of the limitations of the dish antenna and its associated mechanical pedestal. For a phased array antenna, BMX can be used to collect samples that are nearly independent. As shown in the previous chapter, BMX is especially effective at larger signal-to-noise ratios (SNRs) and at smaller spectral widths. These properties could prove to be very useful for a detection scan. When collecting a small number of pulses for a detection scan, using BMX to collect data that are nearly independent should increase the POD compared to contiguous pulses. The following analysis will examine the use of power thresholds for both independent and contiguous pulses to detect weather signatures.

The standard power threshold used for reflectivity on the WSR-88D is actually a SNR threshold and is most commonly set to 2 dB. Some recent work has been done on using other variables in addition to the SNR to enhance the detection of weather on dual polarization radars (Ivić et al. 2009). Attempts were made to utilize the lag-1 autocorrelation along with the SNR for detection, but the combination did not significantly improve the POD. For this study, the standard SNR threshold was employed. A simple simulation was developed to examine the POD for independent and contiguous pulses for a given signal power. For the first simulation, the signal power in the simulation was set to 2 dB above the noise to see how well weather with a relatively low SNR is detected with the WSR-88D. Although a small number of pulses can be used to compute reflectivity in batch mode, at least 15 pulses are collected in the surveillance mode at lower elevations to compute

reflectivity (Brown et al. 2005). Therefore, the simulation was performed with a range of pulses from 1 to 15. The threshold was set to a fixed value of 2 dB to match the most common setting for WSR-88D radars. In order to compute the POD, 1000 trials were run with 10,000 time series realizations for each trial. The POD was computed for each trial, and all 1000 estimates of the POD were averaged to get a final estimate. The simulation was implemented with a true spectrum width, $\sigma_v = 2 \text{ m s}^{-1}$, and a PRT = 3 ms. The radar frequency was set to the frequency of the PAR, 3200 MHz. The simulation of the BMX strategy was based on a scanning strategy where a single pulse is transmitted at each beam position. The samples used to simulate BMX were spaced 24 ms apart and should be almost completely uncorrelated.

Because the threshold is fixed, the PFA varies with the number of pulses. To see this change in the PFA, it was computed analytically using a formulation similar to that found in Ivić et al. (2009):

$$\text{PFA}(M, \text{THR}_{\text{dB}}) = \mathcal{Q}\left(M, M\left(1 + 10^{\frac{\text{THR}_{\text{dB}}}{10}}\right)\right) \quad (4.1)$$

where M is the number of pulses, THR_{dB} is the SNR threshold in dB and $\mathcal{Q}(a, x)$ is the regularized incomplete gamma function. The regularized incomplete gamma function is defined as follows:

$$\mathcal{Q}(a, x) = \frac{\Gamma(a, x)}{\Gamma(a)} \quad \text{where} \quad \Gamma(a, x) = \int_x^{\infty} t^{a-1} e^{-t} dt \quad (4.2)$$

and $\Gamma(a)$ is the standard gamma function. Because the beam multiplexed samples are nearly independent, the POD can also be computed analytically. The formula for the POD is closely related to the PFA formula and can be written using both THR_{dB} and SNR_{dB} (the SNR in dB):

$$POD(M, THR_{dB}, SNR_{dB}) = \mathcal{Q} \left(M, \frac{M}{\left(1 + 10^{\frac{SNR_{dB}}{10}}\right)} \left(1 + 10^{\frac{THR_{dB}}{10}}\right) \right). \quad (4.3)$$

This analytical expression can be used to check the accuracy of the simulation by comparing the calculated values of the POD to the simulated values for the beam multiplexed pulses.

Figure 4.1 shows the POD for both independent and contiguous pulses in one plot with the PFA in the second plot. The POD does not vary significantly but does increase with the number of samples as expected. The POD for the independent (beam multiplexed) pulses is higher for all $M > 1$. Of course, the values for $M = 1$ are the same for both cases. The analytically computed values also match the simulated values very well. The POD is less than 0.5 over the whole range of M ; this is relatively low but gives an idea of the POD for signals with $SNR = 2$ dB on the WSR-88D. The PFA varies significantly from a value of about 7.5×10^{-2} for $M = 1$ to 4.4×10^{-6} for $M = 15$. The large PFA of 0.0754 for $M = 1$ would cause a significant number of false detections. With a range bin size of 240 m which is the value used on the PAR, this would result in approximately 140 range bins being indicated as weather even if the whole radial was noise (assuming an unambiguous range of about 450 km with the PRT = 3 ms). Clearly, a higher threshold needs to be used or some other method needs to be employed if a small number of pulses will be utilized for detection scanning.

The problem is to find a way to have a relatively high POD when using a small number of pulses without having an extremely high PFA. The nearly independent samples from BMX help the POD, but the improvement is not especially significant for this case. Part of it is the low SNR (2 dB) which does not fit into the small spectrum width, high SNR cases where BMX is most effective.

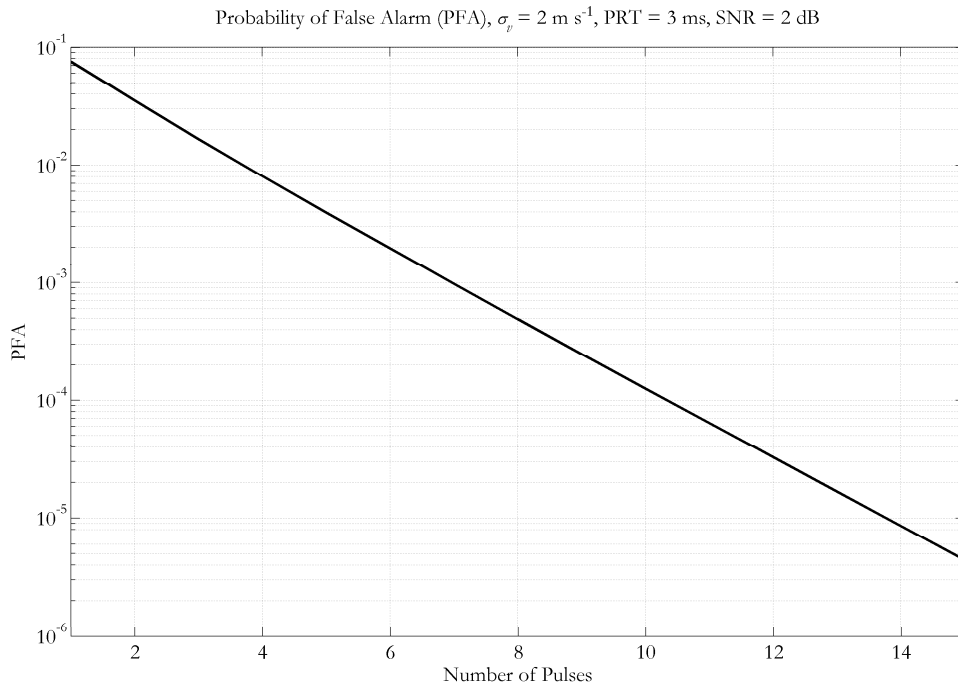



Figure 4.1. Probability of detection (POD) and probability of false alarm (PFA) for independent and contiguous pulses using a spectrum width of 2 m s^{-1} , SNR of 2 dB, and PRT of 3 ms.

It is difficult to reliably detect weather signatures with a small number of pulses when the SNR is so low. One answer is to relax the constraint on SNR and look at weather with an SNR of 5 dB or larger for detection purposes. In order to keep the PFA the same for different numbers of pulses, the threshold should change instead of staying fixed. If equation (4.1) is solved for THR_{dB} given a particular PFA, a unique threshold can be found for each value of M :

$$\text{THR}_{dB} = 10 \log_{10} \left(\frac{Q^{-1}(M, \text{PFA})}{M} - 1 \right) \quad (4.4)$$

where $Q^{-1}(M, \text{PFA})$ is the inverse regularized incomplete gamma function which is calculated using a zero-finding routine.

The next decision is to select an appropriate value for the PFA. There are many ways to select a PFA value, but a conservative value of 1×10^{-6} will be used for the rest of this section. This value is smaller than the PFA of 4.4×10^{-6} for $M = 15$ found earlier and is close to the PFA for $M = 17$ when a 2 dB fixed threshold is used. To see the effects of these changes, the POD is plotted in Figure 4.2 for these new parameters. All of the other parameters are unchanged. The POD of detection varies significantly more in this plot than in Figure 4.1 which used a fixed threshold. Increasing the SNR to 5 dB increases the POD above 0.9 for 14 or 15 independent pulses. Unfortunately, the POD is below 0.3 for 1 to 4 pulses because of the PFA constraint.

The constraint on the PFA significantly decreases the POD for small numbers of pulses. Although this is expected, quickly detecting weather signatures with SNR values less than 5 dB seems to be a difficult problem. In order to achieve a POD of 0.5, seven independent pulses need to be collected.




Figure 4.2. Probability of detection (POD) for independent and contiguous pulses using a spectrum width of 2 m s^{-1} , SNR of 5 dB, and PRT of 3 ms. The SNR threshold is computed so that the PFA has a constant value of 1×10^{-6} .

Increasing the SNR of the signal even more will be necessary to increase the POD. Another possibility is to impose additional constraints on the detection of weather signatures. If it is assumed that samples in range are independent, a continuity constraint could increase the POD while keeping the same PFA. For example, the probability of two consecutive false alarms would be the square of the PFA for a single range bin assuming independent bins. This would allow the use of a PFA of 1×10^{-3} for each individual range bin because the combined PFA for two consecutive bins would still be 1×10^{-6} . One drawback is that the POD is also squared, but the POD should still be larger than the POD for one range bin assuming the same aggregate PFA. This idea could also be extended for 3 or more

contiguous range bins. Imposing this constraint for a small number of bins seems reasonable since the weather will extend beyond one range bin in most cases.

Rather than guessing at possible values of the SNR which lead to a particular POD, it would be helpful to select the desired POD and PFA and determine the SNR needed to achieve it. Because there are analytical expressions for both the POD and PFA for independent pulses, an equation for the required SNR can be derived for given values of the POD and PFA. By combining equations (4.1) and (4.4), the corresponding SNR_{dB} can be found:

$$\text{SNR}_{dB} = 10 \log_{10} \left(\frac{Q^{-1}(M, \text{PFA})}{Q^{-1}(M, \text{POD})} - 1 \right). \quad (4.5)$$

Although this formula is only valid for a single range bin, it is straightforward to compute the proper SNR_{dB} value if a continuity constraint is imposed. If n consecutive range bins are required for detecting a weather signature, then $\text{PFA}^{1/n}$ and $\text{POD}^{1/n}$ are used in equation (4.5) where PFA and POD are the probability values desired after the continuity constraint is enforced. Figure 4.3 shows the SNR_{dB} values for a $\text{PFA} = 1 \times 10^{-6}$ and two different probabilities of detection, $\text{POD} = 0.5$ and $\text{POD} = 0.9$. For each POD, four lines are plotted with continuity constraint, n , from 1 to 4 range bins. The SNR_{dB} value is computed for 1 to 4 pulses.

For a POD of 0.5, it takes four pulses and a continuity constraint of four consecutive bins in order to detect a weather signature with $\text{SNR} = 5$ dB, but two pulses are sufficient to detect a weather signal with $\text{SNR} = 10$ dB. The continuity constraint provides the largest jump in improvement when going from $n = 1$ to $n = 2$, and the improvement increases with larger numbers of pulses. In the two pulse case, the SNR drops from 9.52 dB to 8.79 dB which is an improvement of 0.73 dB.





Figure 4.3. SNR_{dB} for weather signatures having a POD of 0.5 and 0.9 including the effects of a continuity constraint.

For the four pulse case, the improvement increases to 1.05 dB. Even for a relatively low POD of 0.5, the weather signals need to be significantly stronger than the originally examined SNR of 2 dB when using a small number of pulses. As expected, the higher POD of 0.9 requires even higher SNR values. A signal with an SNR of almost exactly 10 dB can be detected with three pulses using a continuity constraint of $n = 4$ while two pulses are needed for SNR values close to 15 dB. Although this analysis provides some information about the power levels needed in order to reliably detect weather signatures, reflectivity values are highly dependent on range so a signal with 5 dB SNR at a close range has a vastly different reflectivity value than a signal with 5 dB SNR at a distant range. Just as with the SNR threshold used on the WSR-88D, weather signatures at distant ranges are much less likely to be detected.

The analytical expressions for computing the SNR are straightforward when looking at independent samples, but to see the effect on contiguous samples, simulations are needed. The biggest difference between independent and contiguous samples comes from changing spectrum width. Small spectrum widths result in higher correlations between consecutive samples which ends up lowering the probability of detection compared to independent samples. To see this, another simulation was run while varying the spectrum width from 1 to 8 m s^{-1} . In addition to looking at a PRT of 3 ms which was used for the earlier simulations, a PRT of 0.8 ms was also studied. This is the shortest PRT that is allowed on the current PAR. Because of the shorter time between samples, the correlation is even higher than for the 3 ms PRT. Figure 4.4 shows the relationship between spectrum width and POD for both independent and contiguous samples. Two cases were simulated with a continuity constraint of $n = 2$, one with $M = 2$ and the other with $M = 4$. The SNR was set to achieve a POD of 0.5 in the two-pulse case and 0.9 in the four-pulse case.





Figure 4.4. The relationship between POD and spectrum width for both independent and contiguous samples at PRTs of 0.8 and 3 ms. For the top plot, $M = 2$ and SNR = 8.79 dB. For the bottom plot, $M = 4$ and SNR = 9.28 dB.

As mentioned earlier, the SNR needed to achieve a POD of 0.5 using two pulses is 8.79 dB. For a POD of 0.9 and four pulses, the SNR is 9.28 dB. These values of the SNR were used in the simulation to check the validity of equation (4.5).

As long as there is enough time between the BMX pulses, the samples will be nearly independent. In both of the BMX cases, the time between pulses was 24 ms. Because of this, the POD for the independent samples is constant. The POD for the contiguous samples drops significantly for small spectrum widths. For the two-pulse case with a POD of 0.5 and a 3 ms PRT, the POD drops from 0.5 to 0.38 and for the shorter 0.8 ms PRT, the POD drops all the way to 0.36. As the spectrum width increases, the POD for the 3 ms PRT reaches 0.5 around a spectrum width of 6 m s^{-1} . The POD for the 0.8 ms PRT never reaches 0.5 even at a spectrum width of 8 m s^{-1} . The results for the four-pulse case with a POD of 0.9 are similar, but the POD for contiguous pulses and narrow spectrum widths drops even further below the BMX cases. Using BMX for detecting weather reflectivity signatures provides two major benefits. The first is that the POD is higher than the POD when using contiguous pulse data collection. The second is that the POD is constant over this particular range of spectrum width values. This avoids the drop in POD compared to contiguous pulses and ensures the same POD over a wide range of conditions. BMX is a valuable tool for detection scanning and can aid in the detection of weather signatures especially at small spectrum widths and short PRTs.

4.3 Detection Scanning Strategies

One unique aspect of phased array antennas is the change in the size and shape of the beam as it is scanned off broadside. As mentioned in Chapter 1, the beamwidth increases by roughly $1/\cos(\theta_s)$ where θ_s is the angle that the beam is steered off broadside. For example, the beamwidth increases by a factor of about $\sqrt{2} \approx 1.4$ when steering 45° away from broadside. This can be a significant issue for detection scanning. If the beams are equally spaced in azimuth and elevation, there will be significant beam overlap away from broadside. The overlap can result in longer scanning times. Another issue is the decrease in antenna gain; this may need to be taken into account when determining scan times at different steering angles.

To examine these issues, several scanning strategies will be compared. A natural way to study phased array scanning is in terms of direction cosines. The direction cosines are determined by a unit vector, $(\cos(\alpha), \cos(\gamma), \cos(\beta))$, where the array is coincident with the x - z -plane of a right-handed coordinate system with the y -axis perpendicular to the array. The z -axis points out of the top of the array, and the $\cos(\alpha)$ and $\cos(\beta)$ terms completely determine the unit vector with α representing the angle between the x -axis and the unit vector and β the angle between the z -axis and the unit vector. One of the important advantages of using direction cosines is mentioned in Skolnik's *Radar Handbook* (2008): "A feature of $\sin \theta$ space is that the antenna pattern shape is invariant to the direction of the scan." If the beam is represented as a circle in what will be referred to as "sine space," the shape of the beam in sine space does not change as it is steered away from broadside. By converting from sine space to angle space, the shape of the beam in angle space can be examined. For the NWRT, the array is tilted back at an angle of 10° from vertical in order to

minimize the effects of beam broadening over a wider range of elevations. The equations that relate $\cos(\alpha)$ and $\cos(\beta)$ to the azimuth and elevation for an array with a tilt back angle are shown below (Katz 2003):

$$\begin{aligned}
 \cos(\alpha) &= \cos(EL)\sin(AZ) \\
 \cos(\beta) &= \sin(EL)\cos(E_0) - \cos(EL)\cos(AZ)\sin(E_0) \\
 EL &= \sin^{-1}\left(\cos(E_0)\cos(\beta) + \sin(E_0)\sqrt{1 - \cos^2(\alpha) - \cos^2(\beta)}\right) \\
 AZ &= \sin^{-1}\left(\frac{\cos(\alpha)}{\cos(EL)}\right)
 \end{aligned} \tag{4.6}$$

where AZ and EL are azimuth and elevation, respectively, and E_0 is the tilt back angle of the array (measured from vertical). For example, antenna broadside corresponds to $\cos(\alpha) = 0$ and $\cos(\beta) = 0$ which results in $AZ = 0^\circ$ and $EL = 10^\circ$ for the NWRT because E_0 is 10° .

Although the broadside beamwidth of the current NWRT is about 1.5° , the beamwidth of the WSR-88D antenna is less than 1° . For the following scanning strategies, a broadside beamwidth of 1° will be assumed. This is a nice round number and is also close to the maximum beamwidth that may be implemented on a future MPAR. The first strategy that is considered is one that has equal spaced beam positions in azimuth and elevation. All of the scanning strategies span angles from -45° to 45° in azimuth and from 0.5° to 80° in elevation. The lowest elevation is limited to 0.5° because this is the lower elevation limit for both the WSR-88D and the NWRT. The azimuth limits were set assuming the use of a four-face phased array to cover most of the hemispherical scanning volume. Various elevation extents for the scans will be considered, but 80° is the maximum for this study. Figure 4.5 shows a detection scan using a 1° square pattern for the beams in angle space. The complete scan is shown in the bottom left of the figure while close-up plots of the top middle, top right, and bottom right portions of the scan are also included. In the bottom right portion of

the scan, there is a small amount of beam overlap, but at higher elevation angles (towards the top) there is significant overlap. This overlap causes more time to be spent than necessary over large portions of the scan.

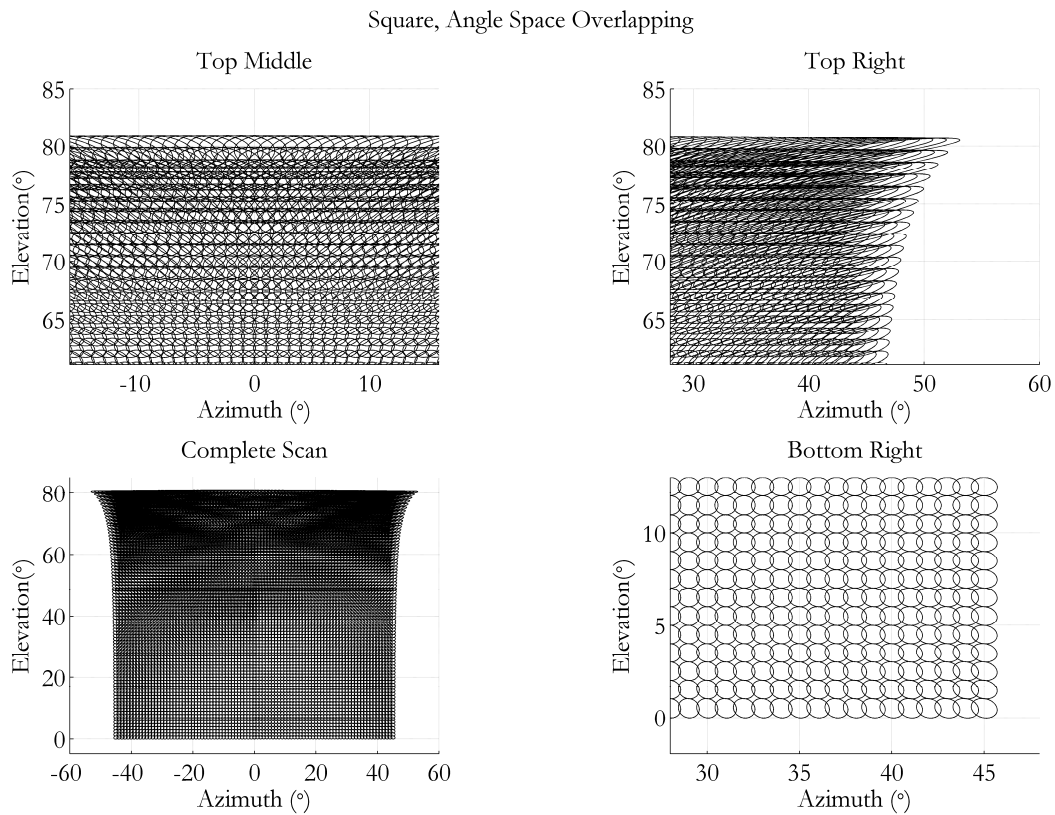


Figure 4.5. Detection scan with square pattern for beam positions in angle space.

To estimate the amount of time for this scan, the PRT for each elevation needs to be determined. The standard method for setting the unambiguous range is to constrain the volume of interest using an altitude limit. In this case, 18 km was chosen which matches the limit used for staggered PRT scanning in Torres et al. (2004). The formula for computing the height of the beam is the same one used by Wood and Brown (1999), and the slant range can also be found if the height is given:

$$\begin{aligned}
b &= r_s \sin EL_k + \frac{r_s^2}{2r_e} \\
r_s &= \sqrt{r_e^2 \sin^2 EL_k + 2r_e b} - r_e \sin EL_k
\end{aligned}
\tag{4.7}$$

where b is the height above the earth, r_s is the slant range, EL_k is the k^{th} elevation angle, and r_e is the effective earth radius. Following Wood and Brown, a 6/5 earth radius is used. After determining the PRT based on the slant range and the number of beam positions for each elevation, a time can be computed for the total scan assuming one transmitted pulse per beam position. This particular scan has 91 beam positions per elevation based on the 1° spacing in azimuth. In this case, the total time for the scan is 2.74 s using PRTs that range from 122 μs to 3083 μs . If four pulses were collected at each beam position, the time would be 10.98 s which is relatively large compared to the overall weather scan time. Based on current WSR-88D VCPs, a single face of a four-face phased array weather radar should take a minute or less to collect weather data. For the current NWRT, the minimum PRT is 800 μs which increases the scan time using one pulse per beam position to 6.45 s. The ability to use PRTs shorter than 800 μs with shorter pulse lengths is important if scanning at high elevations is implemented.

Because of the changing shape of the beam in angle space, the beam positions can be distributed evenly in sine space where the shape of the beam does not change. Figure 4.6 shows a detection scan where the beam positions are distributed in sine space in a square pattern so that the beams do not overlap. Because the beam is 1° wide at broadside, a beamwidth of $b_s = 2 \sin(0.5^\circ)$ was used in sine space. The first beam is positioned at an azimuth of 0° and an elevation of 0.5° , and all of the other beams are placed with respect to that beam. This distribution of beam positions can also be referred to as a square packing in sine space.

Square, Sine Space Packing

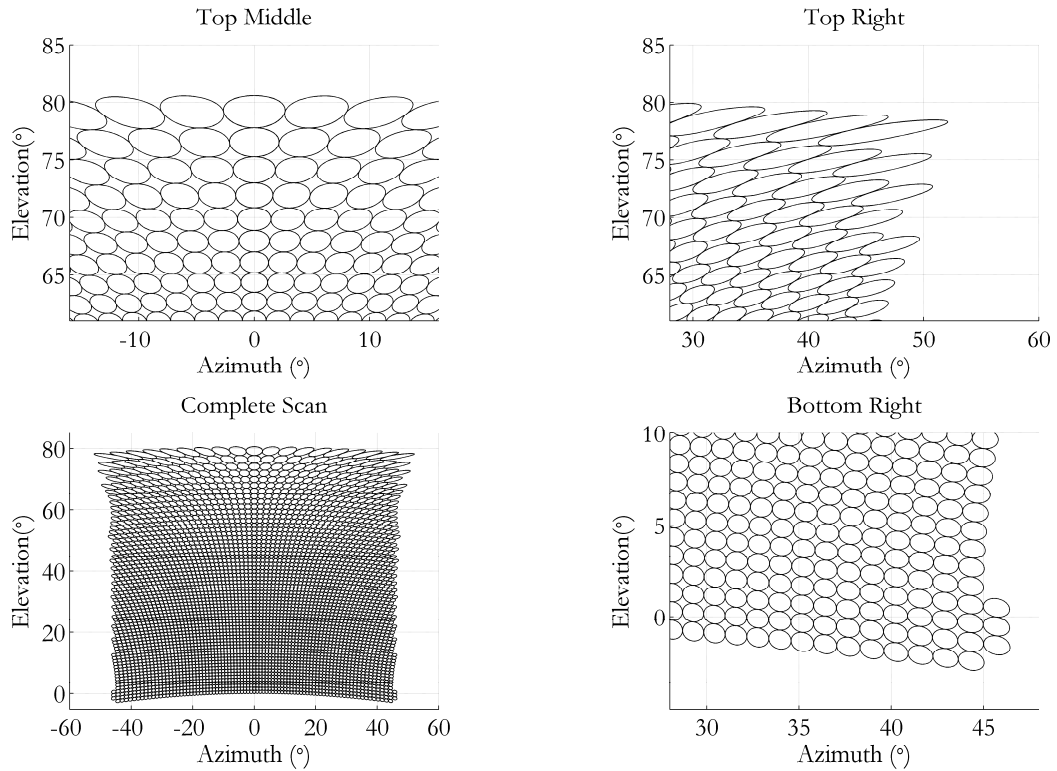


Figure 4.6. Detection scan with square pattern for beam positions in sine space.

As expected, the beams no longer overlap (based on the one-way half-power beamwidth) which leads to significantly fewer beams at high elevations compared to the first scan. The beams do extend beyond the calculated beamwidth, but the gain is lower. The biggest drawback for this strategy is that a constant $\cos(\beta)$ value does not correspond to a constant elevation angle. This occurs because of the array tilt back angle and causes negative elevation angles and a curved shape to the scan. A constant elevation constraint can be imposed to address these problems.

For the next scan, each beam position was adjusted to match the elevation of the previous beam position to ensure a constant elevation. The next beam position needs to be a distance b_s from the previous position in sine space. If the original beam position is (x_ϕ, y_θ) ,

then the following equations ensure that the next beam position (x, y) is at the same elevation:

$$\begin{aligned} x &= x_0 + b_s \cos \varphi \\ y &= y_0 + b_s \sin \varphi \\ \sin(EL) &= \left(y \cos(E_0) + \sin(E_0) \sqrt{1 - x^2 - y^2} \right) \end{aligned} \quad (4.8)$$

Then, a zero-finding method can be used to find φ assuming EL to be the elevation of the previous beam position. Figure 4.7 shows a detection scan with a square pattern in sine space corrected using the constant elevation correction.

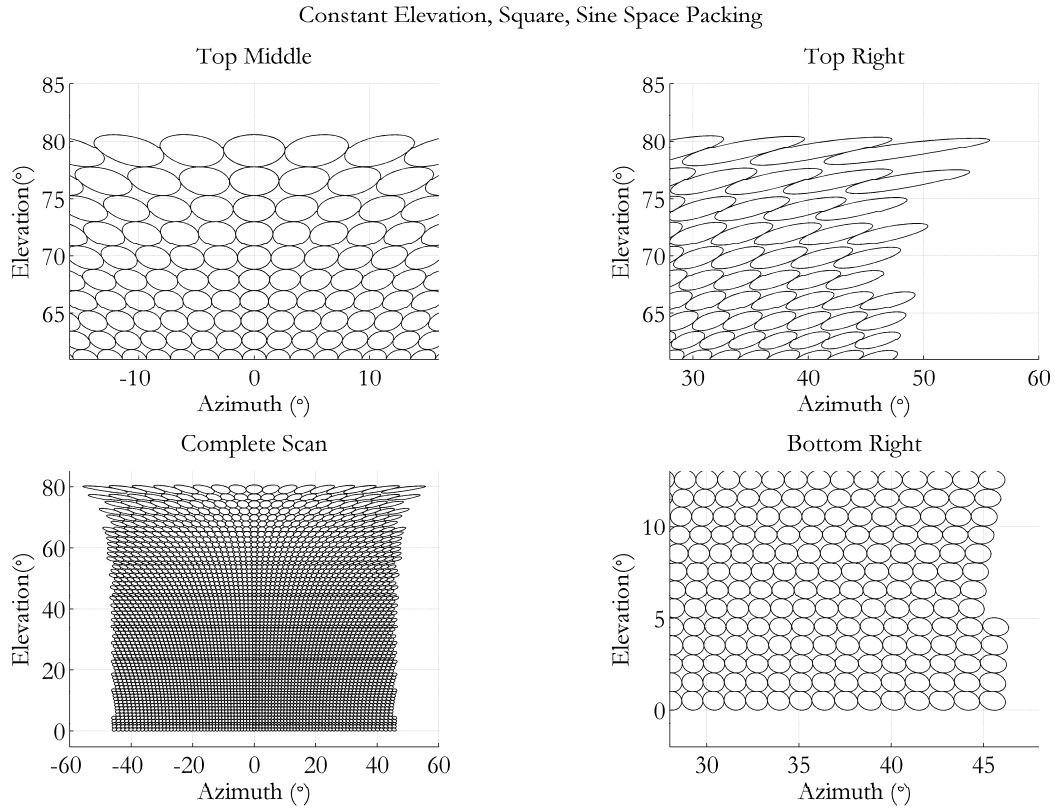


Figure 4.7. Detection scan with square pattern for beam positions in sine space with constant elevation constraint.

This gives a scan that covers the same volume as the original overlapping angle-space scan without any overlapping beams. The constant elevation correction does cause a small

amount of extra space between beams in the upper right portion of the scan. Of course, empty space does not guarantee that scatterers in that area would not be detected. It just means that the gain is below the level used to determine the beamwidth. The constant elevation, sine space scan has 4068 beam positions compared to the $91 \times 80 = 7280$ beam positions for the original angle space scan. The time for the constant elevation, sine space scan is 2.09 s with one pulse per beam position if there are no PRT constraints and 3.82 s if the shortest PRT is limited to 800 μ s. Compared to the angle space scan, the scan time drops by nearly 25% without PRT constraints and drops slightly over 40% when the shortest PRT is constrained to be 800 μ s. This is a significant time savings with most of the difference coming at higher elevation angles. The angle space scan has so much overlap at the higher elevation angles that the 800 μ s at each beam position really adds up. The effect is less pronounced when shorter PRTs are allowed.

The constant elevation, sine space scan does a relatively good job covering the volume with no overlap and with a relatively small amount of space in between the beams. The modified square packing used in this scan is a natural way to collect data but is not the most efficient packing in sine space. A hexagonal packing of the circular beams in sine space results in less space between the beam positions which leads to a higher chance of detection. The downside is that the increase in the number of beam positions causes the time for the scan to increase. A constant elevation, hexagonal, sine space scan was developed to look at the effects on scan time. An estimate of the time increase can be computed based on the different spacing in elevation. For a square packing, the beam centers are b_s apart in sine space between elevations but are $\sqrt{3}/2b_s$ apart when using a hexagonal packing. This would

result in a predicted increase of about 15% in the scan time compared to the square packing.

Figure 4.8 shows the constant elevation, hexagonal, sine space packing scan.

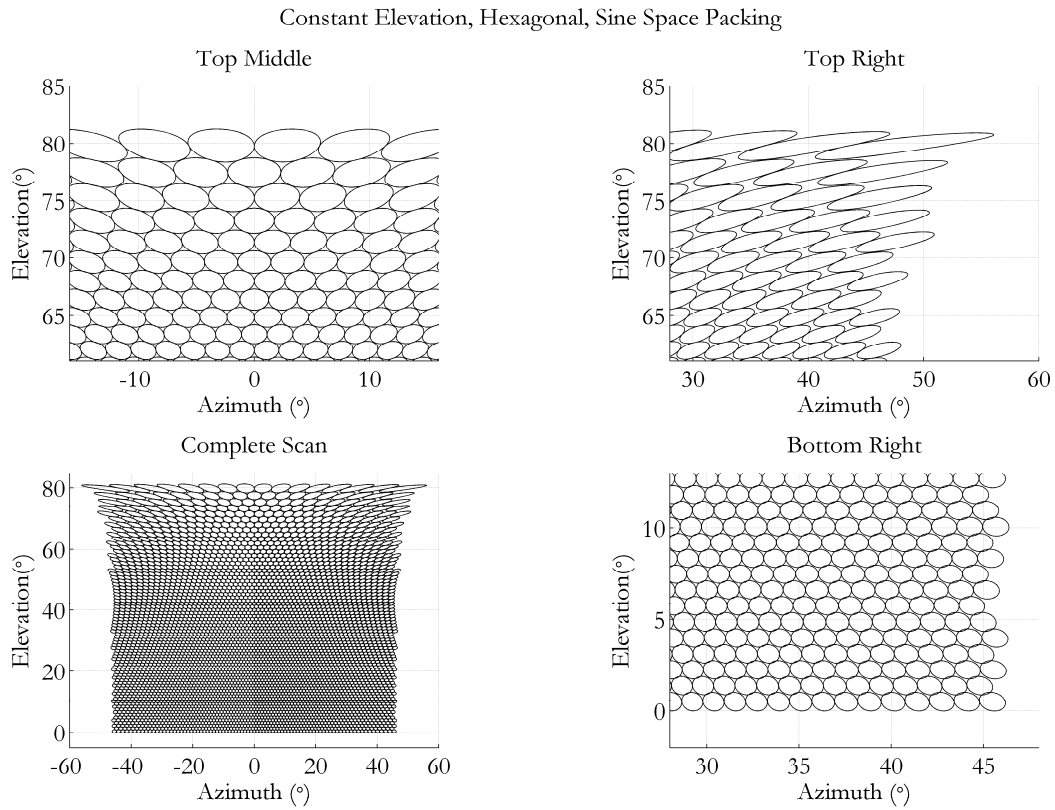


Figure 4.8. Detection scan with hexagonal pattern for beam positions in sine space with constant elevation constraint.

This scan does a better job of decreasing the space between the beam positions without introducing any overlap. The extra space between elevations in the upper right portion of the scan is also minimized. The total number of beam positions for the hexagonal packing increases to 4,685 compared to 4,068 with the square packing. The time also increases to 2.39 s and 4.38 s without and with the PRT constraint, respectively. This ends up being very close to the 15% increase predicted using the change in distance between elevations.

The last detection scan to be considered is a modification of the hexagonal, sine space packing. As mentioned previously, the empty space between beams corresponds to a part of the antenna beam that falls below the gain threshold that defines the beamwidth. To ensure that the gain threshold is met over the whole volume, the beams need to overlap a small amount. The most efficient strategy to overlap the beams as little as possible is to use a hexagonal pattern with less space between beams. This new hexagonal spacing decreases the spacing between elevations in sine space from $\sqrt{3}/2b_s$ to $3/2b_s$, and decreases spacing between azimuths in sine space from $2b_s$ to $\sqrt{3}b_s$. The combination of decreasing the spacing in both dimensions should result in a predicted increase in the time compared to the square packing of about 54%. Figure 4.9 shows the constant elevation, hexagonal, sine space, overlapping scan.

This scan has beam overlap similar to the original angle space scan except that the overlap is spread out evenly instead of occurring mostly at higher elevations or at the edges. The number of beam positions increases to 6,251 which is just about 54% more than the 4,068 from the square packing but is still less than the 7,280 from the original angle space scan. The scan times increase to 3.16 s and 5.82 s without and with the PRT constraint, respectively. The scan time with the 800 μ s PRT constraint is still less than the 6.45 s for the angle space scan, but the time without the PRT constraint is now higher than the angle space scan time of 2.74 s. The closer spacing of beams at lower elevations leads to the increased time. The empty space in the upper right portion of the scan nearly disappears with this new scan with only some very small spaces that occur because of the constant elevation constraint.

Constant Elevation, Hexagonal, Sine Space Overlapping

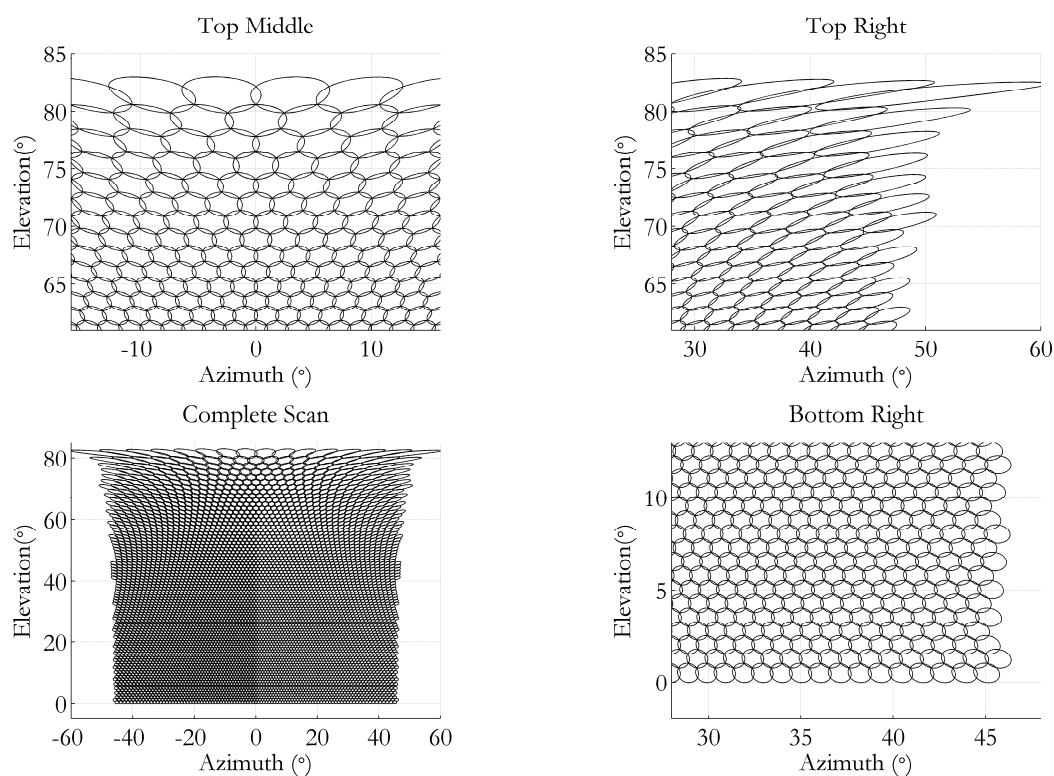


Figure 4.9. Detection scan with hexagonal pattern for beam positions in sine space with constant elevation constraint and overlap.

The initial comparisons of these scans were based on the complete coverage area from -45° to 45° and 0° to 80° , but lower elevation extents could also be considered depending on the acceptable “cone of silence” for the radar. In Table 4.1, the number of elevations, number of beams, and times are given for elevation extents of 20° , 40° , 60° , and 80° . The information in the table clearly shows that most of the scan time is spent at the lowest elevation angles. The most significant consequence of the $800 \mu\text{s}$ constraint is to increase the scan times at higher elevation angles. The rest of this discussion will focus on the scan times without any PRT constraints. For the angle space scan, scanning the lowest 20° takes 65% of the time to scan the whole 80° . In contrast, the lowest 20° takes about 76% of the time to scan the full 80° for all three of the sine space scans. The sine space scans are much more

efficient scanning the top 60° of the 80° scan, but the times for scanning the lowest 20° are comparable to the original, overlapping angle space scan. In general, although sine space scanning is more efficient at higher elevation angles, the PRTs are so much shorter at higher elevation angles that the time savings is limited.

Table 4.1. Detection scanning strategies, scan times, and number of beams

Scanning Strategy	Elev. Extent	Number of Elevations	Number of Beams	Time (s)	Time (s) ($\geq 800 \mu\text{s}$)
Angle Space, Square	20°	20	1820	1.77	2.08
	40°	40	3640	2.22	3.54
	60°	60	5460	2.51	4.99
	80°	80	7280	2.74	6.45
Sine Space, Square, Constant Elevation	20°	20	1616	1.59	1.86
	40°	39	2967	1.92	2.94
	60°	54	3768	2.05	3.58
	80°	64	4068	2.09	3.82
Sine Space, Hexagonal, Constant Elevation	20°	23	1856	1.81	2.12
	40°	45	3421	2.20	3.37
	60°	63	4368	2.35	4.13
	80°	74	4685	2.39	4.38
Sine Space, Hexagonal, Constant Elevation, Overlapping	20°	27	2514	2.40	2.83
	40°	52	4556	2.90	4.47
	60°	72	5786	3.10	5.45
	80°	86	6251	3.16	5.82

For example, look at the constant elevation, hexagonal, sine space scan. The time to collect one pulse at all beam positions up to 20° is 1.81 s which is very similar to the 1.77 s for the angle space scan. If two pulses were collected at each beam position, the scan time would be 3.62 s and would increase to 7.24 s if four pulses were collected. Based on a weather update time of one minute, the two and four pulse detection scans would take about 6% and 12% of the total update time. If the scan is extended to cover all elevations up to 80°, the time for one pulse per beam position increases from 1.81 s to 2.39 s or an increase of about 32%. This corresponds to scan times of 4.78 s and 9.56 s for two and four pulse

scans, respectively. Extending the scan up to 80° increases the percentages to 8% and 16% again assuming a one minute weather update time. The increase is relatively small when only two pulses are collected but becomes significant if four or more pulses are collected. Data collection at the lower elevation angles is the biggest driver of scan time, but a relatively large minimum PRT can greatly increase the scan time at higher elevations.

4.4 Additional Considerations

Even after determining a detection scanning strategy, there are several additional considerations that need to be taken into account. The first to be addressed is the effect of ground clutter on the detection scan. Ground clutter could be detected during the detection scan, be misidentified as weather, and needlessly lead to additional data collection. Filtering the ground clutter is difficult because of the small number of pulses collected and is complicated even more if BMX is used. In the BMX case, the ground clutter from pulse to pulse will be uncorrelated or very weakly correlated which diminishes the effectiveness of two-pulse cancellers and similar clutter mitigation strategies. In the contiguous pulse case, these types of strategies are still not especially effective even if they are more effective than in the BMX case. This is another case where alternative methods of mitigating clutter such as the use auxiliary receive channels could be useful (Le et al. 2009). Another approach would be the use of a dynamic clutter map. If the detection scan is run continuously, ground clutter power levels could be tracked over time, and weather signals that caused deviations in the clutter power could be detected. The main benefit to this approach is that it does not rely on clutter filtering. The main drawback is that there will be quite a bit of variance in the powers since only a few pulses are being collected. A strategy similar to the contiguous range gate constraint could be implemented that only declared a detection if a certain number of

consecutive scans showed an increase in the power over the tracked clutter level. If anomalous propagation (AP) clutter was detected, it could be identified using successive scans. The AP clutter could then be tracked the same way as normal ground clutter.

If detection scans were continuously interleaved with longer scans to more accurately retrieve all of the spectral moments, additional information could be obtained from the detection scans. One possibility would be continuous measurement of refractivity in both clear air and precipitation situations. The feasibility of measuring refractivity on the NWRT has already been shown, and as few as two pulses are needed for usable measurements (Cheong et al. 2008). It is necessary to collect at least two pulses at each beam position because the radial velocity is used to determine the suitability of ground clutter targets for refractivity retrieval. Instead of collecting two beam multiplexed pulses at low elevation angles for the detection scan, either one or both of the separated collections could be extended to collect a pair of pulses instead of a single pulse. A shorter PRT could also be used for one of the pulses in the pair since ground clutter does not extend out beyond a certain range. For example, at the lowest elevation angle a 1 ms PRT could be used for the first pulse followed immediately by 3.2 ms PRT pulse. This pair could be used to measure the change in refractivity and an additional 3.2 ms PRT pulse that was separated from the pair could be transmitted for both the detection scan and for refractivity. The additional pulse for the radial velocity measurement would only add 90 ms to the detection scan, and usable refractivity data could be collected continuously with a very fast update rate. Other useful information could conceivably be extracted from the detection scan, and relatively fast updates from a phased array radar could increase the benefits.

The basic structure of a possible detection scan and some applications have been explored, but there are also considerations on how to use the information from the detection

scan to collect additional data. This is the other aspect of adaptive scanning that was mentioned in Section 4.1. One way would be to collect weather data at the same beam positions used for the detection scan. Information from the detection scan could then be used to optimize the weather collection scan. For example, shorter PRTs could be used if the weather was limited in extent compared to the detection scan PRT. Scanning strategies could also be chosen based on the detected weather; uniform sampling, staggered PRT, or other multiple PRT strategies could be employed. If the weather collection scan covered the full range of the detection scan, the collected data could be used instead of the standard detection scan. If the weather collection scan did not cover the full range, the detection scan would still need to be implemented to possibly detect weather beyond the unambiguous range limit of the weather collection scan. The simpler approach would be to collect a full detection scan instead of trying to selectively use longer weather collection scans for weather detection.

In some cases, the pattern used for the detection scan might not match the desired pattern for the weather data collection. Even when a hexagonal pattern might be optimal for the detection scan, a more conventional square pattern might be optimal for the weather scan. A range height indicator (RHI) scan lends itself more to the square collection pattern than the hexagonal pattern. Interpolation could be used to produce an RHI from a hexagonal scan pattern, but the results may be suboptimal. If the scanning patterns do not match directly, an algorithm will need to be developed to determine how to map the beam positions from the detection scan to the weather collection scan. A simple approach would be to use the beam position center closest to the weather collection beam position center to determine whether to execute the weather collection scan. If the weather collection beam position center was close to halfway between two detection scan beam positions, a more

complicated algorithm could be used. Maybe the weather data collection would be executed if either one of the nearest two detection scans detected weather. These are the kinds of issues that are unique to adaptive scanning using a phased array weather radar and need to be studied further.

Another way to speed up detection scanning is to collect data from multiple beam positions simultaneously using multiple receive beams with a single spoiled transmit beam. This is a capability of certain active phased arrays and has been suggested as a possible future capability of the MPAR (Weber et al. 2005). The particular design mentioned in the paper enables 9 simultaneous receive beams for each spoiled transmit beam. Although this approach could be used to speed up detection scans, the technique can only be applied at higher elevation angles. The spoiled transmit beam reduces the antenna gain and can only provide adequate detection performance at limited ranges. It is not completely clear at what elevation angle this could be used, but based on the terminal aircraft area surveillance scan timeline, 20° seems to be a reasonable estimate. Applying this to the previously discussed constant elevation, hexagonal, sine space scan, the time to collect a single pulse at all the beam positions in the detection scan would drop from 2.39 s to about 1.87 s. This is a reduction of nearly 23% in the scan time for the one pulse case. The times for the two and four pulse cases drop to 3.74 s and 7.48 s from the previously computed values of 4.78 s and 9.56 s, respectively. Although this technique is useful in reducing scan times, it still does not address the most time intensive portion of the scan between 0° and 20° .

4.5 Conclusion

Detection scanning of some sort will be important to future adaptive scanning schemes used on phased array weather radars. If the goal is to avoid weather collection scanning over

the whole volume, some type of detection scanning will be necessary to identify newly forming weather. The tradeoff comes in deciding how much time to spend detection scanning so that the weather scanning does not suffer. All of different scanning strategies have to be scheduled in an optimal way. In the section on detecting weather signatures, some results related to the probability of detection of weather signatures were examined. To keep low false alarm rates and reasonable scan times, SNR threshold values had to be increased compared to the standard SNR thresholds used on WSR-88D radars. BMX increases the probability of detection at smaller spectrum widths but may complicate clutter mitigation strategies. With both BMX and conventional scanning, dealing with clutter will be a challenge when only small numbers of pulses are collected. Future research will be needed to determine how sensitive the detection scans need to be to detect weather quickly without missing significant early formation.

The section on detection scanning strategies introduced several angle-space and sine-space strategies and showed that a considerable amount of time can be wasted when using angle-space strategies especially at higher elevation angles. Sine space strategies can cut down significantly on overlapping beams but may not be as natural as angle-space strategies for producing RHI-type scans. One of the problems with all of the detection strategies is that most of the scanning time is spent at lower elevation angles, and there are no simple answers for speeding up this part of the scan. Balancing volume coverage with the amount of time spent scanning is another key area for future research.

The final section addressed some possible ways to deal with ground clutter when utilizing detection scans. Filtering strategies may be difficult to apply and new avenues for mitigating ground clutter will need to be explored. If detection scans were run continuously on the radar, refractivity measurements could be updated continuously at a rapid rate. This

information could also be used as input for future numerical models to aid in short-term forecasts. The next issue addressed was how to use detection scans to determine the best way to utilize the radar for weather scanning. Scanning strategies could be tailored on a beam position basis to improve scanning times and lower errors. Finally, a technique for collecting data from multiple beam positions simultaneously was discussed. The technique takes advantage of the unique capabilities of the type of active phased array antenna that could be implemented on the MPAR. Adaptive scanning is one of the most exciting and promising areas of research on phased array weather radars, and detection scanning will be instrumental in achieving its full potential.

5. Conclusions and Future Work

With phased array antennas becoming more affordable, the prospect of widespread use of phased array weather radars is increasing. Phased array antennas are being considered for use in both a future replacement for the NEXRAD network of radars and a network of smaller, gap-filling weather radars (National Research Council 2002). As the WSR-88D radars age and with the possibility of replacing them with an MPAR that may perform multiple functions in addition to weather surveillance, it is a critical time for phased array research. This includes using the ideas that have been percolating in the weather research community for years and adding new ideas based on the capabilities of active phased array antennas and the challenges of multifunction radars. In this dissertation, both long-discussed research areas and new research areas related to adaptive scanning were examined. This multi-pronged research approach needs to be extended over the next several years to address the challenges and opportunities associated with phased array weather radars and a possible future multi-function radar system.

Chapter 2 on spectral clutter characterization examined the shape of ground clutter spectra collected with a stationary phased array antenna. The electronic scanning capability leads to the possibility of using fixed array faces for future phased array weather radars. The prospect of using a stationary antenna for weather data collection eliminates the smearing that occurs when using rotating antennas. In this chapter, the main objectives included finding a clutter model for the collected data, examining different interpretations of the clutter width, looking at the effect of beam broadening on the clutter model, and observing the effects of time series length on the spectral shape. Narrower spectra and the associated

longer decorrelation times should be considered when designing new ground clutter filters for phased array weather radars.

In Chapter 3, a new beam multiplexing technique, SBMX, was introduced to mitigate some of the limitations of an earlier technique, BMX pairs. Although SBMX successfully addresses some of those limitations, conventional staggered PRT processing performs similarly to SBMX with the additional advantage of having proven ground clutter filtering strategies. The conditions where beam multiplexing outperforms both staggered PRT and conventional contiguous pulse scanning strategies is at narrow spectrum widths and relatively high SNRs. This performance was explored in more detail in Chapter 4 when discussing the detection of weather signatures with small numbers of pulses. Even though beam multiplexing faces several challenges, it should still be considered as a tool when encountering new obstacles with phased array weather radars.

Chapter 4 examined some of the issues concerning the detection of weather signatures when using fast surveillance scans. These weather detection scans may play a large role in future adaptive scanning strategies. The ability to rapidly jump from one beam position to another is one of the most significant advantages of phased array antennas, and a large portion of future research should look at the best ways to take advantage of this capability. Any adaptive scanning strategy that attempts to scan only where weather is actually occurring needs to have a way to periodically sample the entire volume to ensure that newly developing weather is not overlooked. This chapter addressed two parts of this detection scanning process: the probability of detection of weather signatures and the best way to sample the volume using a phased array antenna. Beam multiplexing was shown to be a viable scanning strategy for keeping the probability of detection constant as the spectrum width varies; contiguous pulses strategies have lower probabilities of detection at narrow

spectrum widths. Several strategies for sampling a large portion of the total volume were introduced based on the invariance of the shape of the beam in sine space. These strategies trade-off time for volume coverage but do give an idea of the amount of time needed for detection scanning.

As has been mentioned previously, adaptive scanning will play a large role in future phased array research. If the multifunction capabilities of a future MPAR are to be fully realized, adaptive scanning for both weather and aircraft targets will need to be studied along with the interaction between the two. Even if two or more frequencies are implemented on an MPAR, the dead zones during transmission will need to be addressed. For weather scanning, finding the optimal scanning strategy at each beam position will become very important. If the total volume can be scanned quickly using a fast detection strategy, the information from that scan can also be used to determine the optimal scanning strategies at each beam position. The PRT could be adjusted based on coverage or to minimize range velocity ambiguities. The type of scanning could be changed on a beam position-to-beam position basis from a contiguous pulse to a staggered PRT or beam multiplexing strategy or any other yet-to-be-invented collection strategy. Some approaches to adaptive waveform selection are already being implemented on conventional weather radars (Cho et al. 2004), but the capabilities of phased array antennas provide even more opportunities for extending and improving these early attempts at adaptively selecting optimal waveforms. Haykin (2006) extends the idea of adaptive scanning even further in his discussion of “cognitive radar” which depends on intelligent signal processing, feedback from the receiver to the transmitter, and the preservation of the information content of radar returns. Following a path toward a true “cognitive radar” could lead to an enormous amount of research in the future.

In addition to the areas already touched on, the capabilities of future active phased array antennas need to be considered. The possibility of using a spoiled transmit beam with multiple, simultaneous receive beams was mentioned in Chapter 4 (Weber et al. 2005). Other ideas that could speed up data collection by taking advantage of digital beamforming should also be considered. With multiple channels, it may be possible to implement space-time adaptive processing (STAP) techniques for mitigating ground clutter or to solve other potential problems. The cost of phased array antennas is significantly higher than conventional antennas, and research looking at the approaches that give the best return on hardware investment will be important in the future. The cost-benefit analysis for these types of techniques will play a large role in determining the requirements for future phased array antennas.

The benefits of phased array antennas are numerous, but taking advantage of these new capabilities will require a significant amount of research. Radar engineers have been thinking about how to capitalize on the unique capabilities of phased array antennas for years, and some of those ideas are finally being implemented and tested on the NWRT. Innovations in both hardware and algorithms are needed to fully realize the potential of phased array antennas for weather data collection. On the hardware side, the cost of phased array antennas needs to be reduced to make them feasible. Other obstacles such as dual polarization need to be overcome for a possible future network of phased array radars to provide the same capabilities of the dual polarized WSR-88D. Future algorithms need to address the challenges of adaptive waveform selection and agile beam scanning to achieve the full benefits of the electronic scanning capabilities of the phased array. Matching the performance of the current WSR-88D radar is important, but looking beyond the limitations of conventional antennas is also essential to improve the performance of future phased array

weather radars. The promise of phased array antenna technology is exciting, but a lot of work still needs to be done before this technology becomes the new standard for weather data collection.

Some ideas for possible future research are described below:

- Collect data at higher wind speeds (above 20 mph) under “Full Foliage” conditions to better capture the shape of the AC part of the spectrum. Study the changes in both the f and b parameters especially where f and b are nearly equal and a is close to zero or negative.
- Explore alternative ways to model the DC portion of the clutter spectrum. This could include the use of the autocorrelation function to aid in characterizing the change in the decorrelation time under varying conditions.
- Examine the possible benefits of narrower clutter spectra on the mitigation of ground clutter. Apply the results to improve spectral clutter filters by minimizing the effects of clutter filtering on spectral moment estimation.
- Find novel ways to utilize beam multiplexing that take advantage of its strengths and minimize its limitations.
- Study new ways to apply the capabilities of active phased array antennas to improve the performance of beam multiplexing. For example, the use of multiple, simultaneous receive beams could enable the collection of second trip data for BMX pairs.
- Investigate clutter mitigation techniques that could be used for beam multiplexing to address one of its major limitations.

- Collect and analyze weather data to optimize the methods used to detect weather. The results could be used to determine appropriate SNR thresholds and continuity constraints for the practical implementation of rapid weather detection scanning.
- Optimize weather detection strategies to balance scan time and probability of detection. Investigate sine space strategies and quantify their performance.

6. References

- Billingsley, J. B., 2002: *Low Angle Radar Land Clutter: Measurement and Empirical Models*, William Andrew Publishing, 703 pp.
- Brown, R.A., R.M. Steadham, B.A. Flickinger, R.R. Lee, D. Sirmans, and V.T. Wood, 2005: New WSR-88D Volume Coverage Pattern 12: Results of Field Tests. *Wea. Forecasting*, **20**, 385–393.
- Cheong, B.L., R. D. Palmer, C. D. Curtis, T.-Y. Yu, D. S. Zrníc, and D. Forsyth, 2008: Refractivity Retrieval Using the Phased-Array Radar: First Results and Potential for Multimission Operation, *IEEE Transactions on Geoscience and Remote Sensing*, **46**, 2527–2537.
- Cho, J. Y. N., N. G. Parker, and G. R. Elkin: 2004: Improved Range-Velocity Ambiguity Mitigation for the Terminal Doppler Weather Radar, *11th Conference on Aviation, Range, and Aerospace Meteorology (ARAM)*, Hyannis, MA, *Amer. Meteor. Soc.*, P4R.8.
- Crum, T. D., and R. L. Alberty, 1993: The WSR-88D and the WSR-88D operational support facility. *Bull. Amer. Meteor. Soc.*, **74**, 1669–1687.
- Dennis, A. S., and P. L. Smith, 1970: Next Generation of Weather Radars, unpublished memorandum.
- Doviak, R. J., and D. S. Zrníc, 1993: *Doppler Radar and Weather Observations*. 2nd ed. Academic Press, 562 pp.
- Fabry, F., and R. Jeffrey Keeler, 2003: Innovative Signal Utilization and Processing. *Meteor. Monogr.*, **30**, 199.
- Fang, M., R. J. Doviak, and V. Melnikov, 2004: Spectrum width measured by WSR-88D: Error sources and statistics of various weather phenomena. *J. Atmos. Oceanic Technol.*, **21**, 888–904.
- Forsyth, D. E., J. F. Kimpel, D. S. Zrníc, R. Ferek, J. F. Heimmer, T. J. McNellis, J. E. Crain, A. M. Shapiro, R. J. Vogt, W. Benner, 2005: The National Weather Radar Testbed (Phased-Array). Extended Abstracts, *32nd Conference on Radar Meteorology*, Albuquerque, NM, USA, American Meteorological Society, CD-ROM, 12R.3.
- Frasco, B. and S. Katz, 2003: National Weather Radar Testbed User Manual, Lockheed Martin.
- Free, A. and N. K. Patel, 2005: Clutter Censoring Theory and Application for the WSR-88D, *32nd Conference on Radar Meteorology*, P4.R15.
- Gradshteyn, I. S., I. M. Ryzhik, A. Jeffrey, and D. Zwillinger, 2000: *Table of Integrals, Series, and Products*, Academic Press, 1216 pp.

- Gumbel, E. J., 1960: Multivariate Extremal Distributions. *Bull. Inst. Internat. de Statistique*, **37**, 471-475.
- Haykin, S., 2006: Cognitive radar: a way of the future. *Signal Processing Magazine, IEEE*, **23**, 30-40.
- Heinselman, P. L., D. L. Priegnitz, K. L. Manross, T. M. Smith, R. W. Adams, 2008: Rapid sampling of severe storms by the National Weather Radar Testbed Phased Array Radar. *Weather and Forecasting*, **23**, 808-824.
- Hössjer, O. and C. Croux, 1995: Generalizing Univariate Signed Rank Statistics for Testing and Estimating a Multivariate Location Parameter, *Non-parametric Statistics*, **4**, 293-308.
- Ice, R. L., D. A. Warde, A. D. Free, R. D. Rhoton, D. S. Saxion, C. A. Ray, N. K. Patel, O. E. Boydston, D. S. Berkowitz, J. N. Chrisman, J. C. Hubbert, C. J. Kessinger, M. Dixon, S. M. Torres, 2007: Optimizing clutter filtering in the WSR-88D. Preprints, *23rd International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, San Antonio, TX, USA, AMS, CD-ROM, P2.11.
- Ivić, I. R., D. S. Zrnić, and T.-Y. Yu, 2009: Use of Coherency to Improve Signal Detection in Dual-Polarization Weather Radars. *J. Atmos. Oceanic Technol.*, in publication.
- Katz, S., T. Maese, and W. Mazur, 2003: NWRT/EP Training Presentation Material, Lockheed Martin.
- Le, K. D., R.D. Palmer, B. L. Cheong; T.-Y. Yu, G. Zhang and S. M. Torres, 2009: On the Use of Auxiliary Receive Channels for Clutter Mitigation With Phased Array Weather Radars. *IEEE Transactions on Geoscience and Remote Sensing*, **47**, 272-284.
- Lee, C.-G., P.-S. Kang, C.-S. Shih, L. Sha 2003: Radar dwell scheduling considering physical characteristics of phased array antenna. *Real-Time Systems Symposium, 2003. RTSS 2003. 24th IEEE* , 14-24.
- National Research Council, 2002: *Weather Radar Technology beyond NEXRAD*. National Academy Press, 81 pp.
- NEXRAD Technical Requirements (NTR), R400-SP401A, 1 November 1991, Joint Systems Program Office, Silver Spring, MD.
- Orescanin, M. B., T. Y. Yu, C. D. Curtis, D. S. Zrnic, and D. E. Forsyth, 2005: Signal processing of Beam-multiplexed data for Phased-Array weather radar. *Extended Abstracts, 32nd Conference on Radar Meteorology*, Albuquerque, NM, USA, American Meteorological Society, CD-ROM, 4R.6.
- Raghavan, S., 2003: *Radar Meteorology*. Springer, 549 pp.

- Ray, T. and J. Chrisman, 2008: Automating the WSR-88D. *NEXRAD Now*, Radar Operations Center, Issue 18.
[Available online at <http://www.roc.noaa.gov/news/NNautumn08d1.pdf>.]
- Rouseeuw, P. J., and K. Van Driessen, 1999: A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, **41**, 212-223.
- Sachidananda, M., and D.S. Zrnić, 2002: An Improved Clutter Filtering and Spectral Moment Estimation Algorithm for Staggered PRT Sequences. *J. Atmos. Oceanic Technol.*, **19**, 2009–2019.
- Siggia, A. D., and R. E. Passarelli, Jr., 2004: Gaussian Model adaptive processing (GMAP) for improved ground clutter cancellation and moment calculation. *Proceedings of ERAD*, 67-73.
- Skolnik, M. I., 2008: *Radar Handbook*. 3rd Ed. McGraw Hill Professional, 1328 pp.
- Torres, S., 2005: Range and velocity ambiguity mitigation on the WSR-88D: Performance of the SZ-2 phase coding algorithm. *Preprints, 21st International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, San Diego, CA, USA, American Meteorological Society, 19.2.
- Torres, S., 2006: Range and velocity ambiguity mitigation on the WSR-88D: Performance of the staggered PRT algorithm. *Extended Abstracts, 22nd International Conference on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*, Atlanta, GA, USA, AMS, CD-ROM, 9.9.
- Torres, S. M., Y. Dubel, and D. S. Zrnic, 2004: Design, implementation, and demonstration of a staggered PRT algorithm for the WSR-88D. *J. Atmos. Oceanic Technol.*, **21**, 1389-1399.
- Torres, S. and D. Zrnić, 1999: Ground Clutter Canceling with a Regression Filter. *J. Atmos. Oceanic Technol.*, **16**, 1364-1372.
- Torres, S., D. Zrnic, 2004: Range and velocity ambiguity mitigation techniques for the WSR-88D weather radar. Preprints, *International Geoscience and Remote Sensing Symposium 2004*, Anchorage, AK, USA, Institute of Electrical and Electronics Engineers, III-1727-III-1729.
- Weber, M., Cho, J. Y. N., Flavin, J., Herd, J, Vai, M., 2005: Multi-function Phased Array Radar for U.S. Civil-Sector Surveillance Needs, Preprints, *32nd Conf. on Radar Meteorology*, Albuquerque, NM, Amer. Meteor. Soc., 12R.2.
- White, R. G., 1994: A model for MTI radar clutter. *Proc. SEE Int. Radar Conf.*, Paris, 540-545.
- Wood, V., R. Brown, 1999: The optimization of WSR-88D scanning strategies for convective storms. Preprints, *29th International Conference on Radar Meteorology*, Montreal, QC, Canada, American Meteorological Society, P2.21.

- Yu, T.-Y., M. B. Orescanin, C. D. Curtis, D. S. Zrnic, D. E. Forsyth, 2007: Beam Multiplexing Using the Phased-Array Weather Radar. *J. Atmos. Oceanic Technol.*, **24**, 616-626.
- Yu, T.Y., R.R. Rondinel, and R.D. Palmer, 2009: Investigation of Non-Gaussian Doppler Spectra Observed by Weather Radar in a Tornadoic Supercell. *J. Atmos. Oceanic Technol.*, **26**, 444-461.
- Zhang, G., and R.J. Doviak, 2007: Spaced antenna interferometry to measure crossbeam wind, shear, and turbulence: Theory and formulation, *J. Atmos. Oceanic Technol.* **24**(5), 791-805.
- Zrnić, D. S., 1975: Simulation of weatherlike Doppler spectra and signals. *J. Appl. Meteor.*, **14**, 619-620.
- Zrnić D. S., 1979: Spectrum width estimates for weather echoes. *IEEE Trans. Aerosp. Electron. Syst.*, **AES-15**, 613-619.
- Zrnić, D. S., and S. Hamidi, 1981: Considerations for the Design of Ground Clutter Cancelers for Weather Radar. Interim Report for the Department of Transportation.
- Zrnić, D. S., J. F. Kimpel, D. F. Forsyth, A. Shapiro, G. Crain, R. Ferek, J. Heimmer, W. Benner, T. J. McNellis, 2007: Agile beam phased array radar for weather observations. *Bulletin of the American Meteorological Society*, **88**, 1753-1766.

Appendix A. Estimating the Spectral Floor and Removing the Effect of System Noise

If a portion of the spectrum does not contain any extraneous signals, those spectral components can be averaged to find the spectral floor. Unfortunately, it is difficult to determine which spectral components are not contaminated. For this study, a technique similar to the “Dynamic noise power” technique used in the GMAP (Gaussian Model Adaptive Processing) clutter filter was utilized (Siggia and Passarelli 2004). Several time series realizations of 4096 complex, Gaussian noise samples with a power of 0 dB were simulated by properly scaling complex, pseudorandom, normal variates. Spectra were then computed using the same processing that was used for the clutter spectra. The spectral coefficients were sorted for each realization and then averaged. The smallest 1024 sorted values of this average curve (in dB) can then be compared to the smallest 1024 sorted values of each spectrum. The mean difference between the curves is an estimate of the spectral floor, F , in dB. Tests were run using simulated noise data with different average powers, and the results from this technique matched very closely.

To examine the characteristics of the spectral floor without system noise, the system noise floor, N_s , is estimated and subtracted from the measured spectral floor. F and N_s are in dB so they need to be converted to the linear scale to be subtracted. The final adjusted spectral floor, F_A , is calculated using the following equation:

$$F_A = 10 \log_{10} \left(10^{\frac{F}{10}} - 10^{\frac{N_s}{10}} \right) \quad (\text{A.1})$$

The system noise floor is not subtracted from the spectrum itself; it is only used to adjust how much the spectrum is shifted before fitting. This isolates the component of the spectral floor that depends only on the ground clutter from the system noise component.

The system noise floor for each radial, N_s , is estimated by finding the median of the noise values computed using the spectral noise measurement technique described earlier. Only noise values measured beyond a certain range (about 50 km) are included in the calculation. The ranges should be relatively clear of clutter or other contamination, but the use of the median ensures that small amounts of contamination will not effect the measurement. The distribution of these noise values has a symmetrical shape so the median estimator is unbiased. After subtracting the noise floor, only spectra that have an adjusted spectral floor greater than the noise floor are fit with the model. This ensures that the spectral floor has at least the same power as the system noise on average. The log spectrum is then shifted by the adjusted spectral floor, F_{adj} , before fitting the AC portion of the spectrum. In the worst case, the spectrum will be fit to within 12 dB of the original measured spectral floor. This technique was used instead of directly subtracting the noise floor from the power spectrum to avoid negative power values.

Appendix B. Mean Correction for Log Spectra

The correction value for the log spectra (or dB spectra) can be found using the properties of the Gumbel distribution (Gumbel 1960). The pdf of the Gumbel distribution is

$$f(x) = \frac{1}{\beta} \exp\left(-\frac{x-\alpha}{\beta}\right) \exp\left(-\exp\left(-\frac{x-\alpha}{\beta}\right)\right) \quad (\text{B.1})$$

where α is the location parameter and β is the scale parameter. If the variates from an exponential distribution with mean μ are converted to dB, the variates will be Gumbel distributed with $\alpha = 10 \log_{10}(\mu)$ and $\beta = 10/\log(10)$. The mode of this distribution is α , but the mean is $\alpha - \gamma\beta$ where γ is the Euler-Mascheroni constant. The fixed value that needs to be added to each spectrum in the log domain is given by $\omega = \gamma\beta \approx 2.5068$ dB. The variance of the Gumbel distribution is given by $\beta^2 \pi^2/6$.

Appendix C. Mean Squared Error (MSE) Threshold

One simple way to measure the difference between the model and the clutter spectra is to use the MSE (mean squared error). Because the distribution of each of the spectral components has the same shape with a possibly different mean, the error should be consistent for power values with different magnitudes. Thus, the MSE should work well over a varying range of powers and with a varying number of samples. If the MSE is used, then a threshold needs to be determined to differentiate between spectra that fit the model well and spectra that need to be excluded. Figure C.1 shows a histogram of the MSE values that were calculated when fitting the quadratic model to simulated data. After looking at histograms of several other cases, a value of 45 dB^2 was chosen for the threshold. The MSE is a slightly biased estimate of the variance, and the variance of the Gumbel distribution is known and is given by $\beta^2 \pi^2 / 6 \approx 31.03 \text{ dB}^2$ (see Appendix B). The histogram in Figure C.1 is consistent with this variance. Using 45 dB^2 as the threshold only excludes a few of the simulated spectra. For real data, this should allow most spectra that fit the model reasonably well to be included. This type of error test does self-select for spectra that fit the model, but after processing several cases, a large proportion of the spectra are still included. Most of the spectra that are excluded are contaminated in some way or have a strange shape that could not be characterized easily.




Figure C.1. Histogram of MSE values computed using a time series simulation.

Appendix D. Sample Spectra and MSE Histograms

Sample spectra and MSE histograms are presented for all of the sets of conditions except for “Light Foliage, Prairie” which can be found in Section 2.6. This includes the “Full Foliage, Urban,” “Full Foliage, Prairie,” and “Light Foliage, Urban” sets of conditions. The sample spectra show the expected increase in the clutter width with wind speed, and several spectra show other signals in addition to the clutter signal near zero velocity. The MSE histograms also follow the same pattern as the ones for the “Light Foliage, Prairie” conditions with the “Breezy” cases showing the largest discrepancies between the mean errors for the quadratic and linear fits.

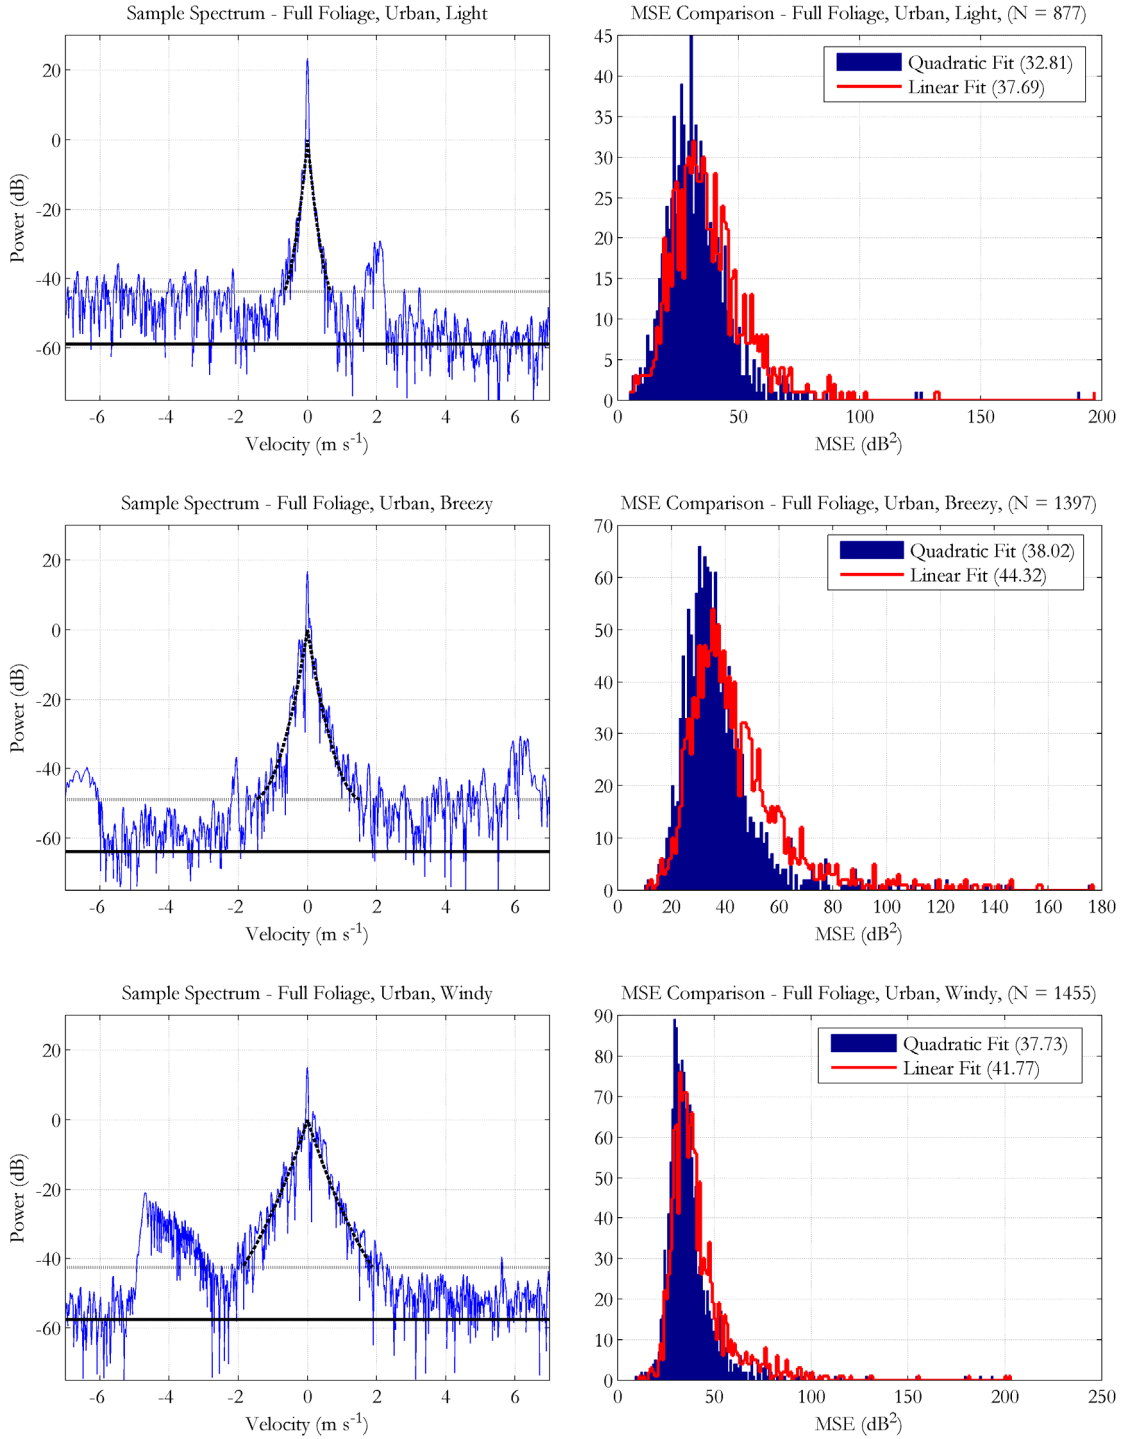


Figure D.1. Sample ground clutter spectra and MSE histograms for the “Full Foliage, Urban” conditions. (Estimated model parameters: Light - $\hat{a} = 121.0 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 106.2 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 58.9 \text{ dB}$, $\hat{d} = 23.3 \text{ dB}$, Breezy - $\hat{a} = 38.5 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 61.5 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 64.0 \text{ dB}$, $\hat{d} = 16.8 \text{ dB}$, Windy - $\hat{a} = 8.2 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 30.3 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 57.6 \text{ dB}$, $\hat{d} = 15.2 \text{ dB}$.)

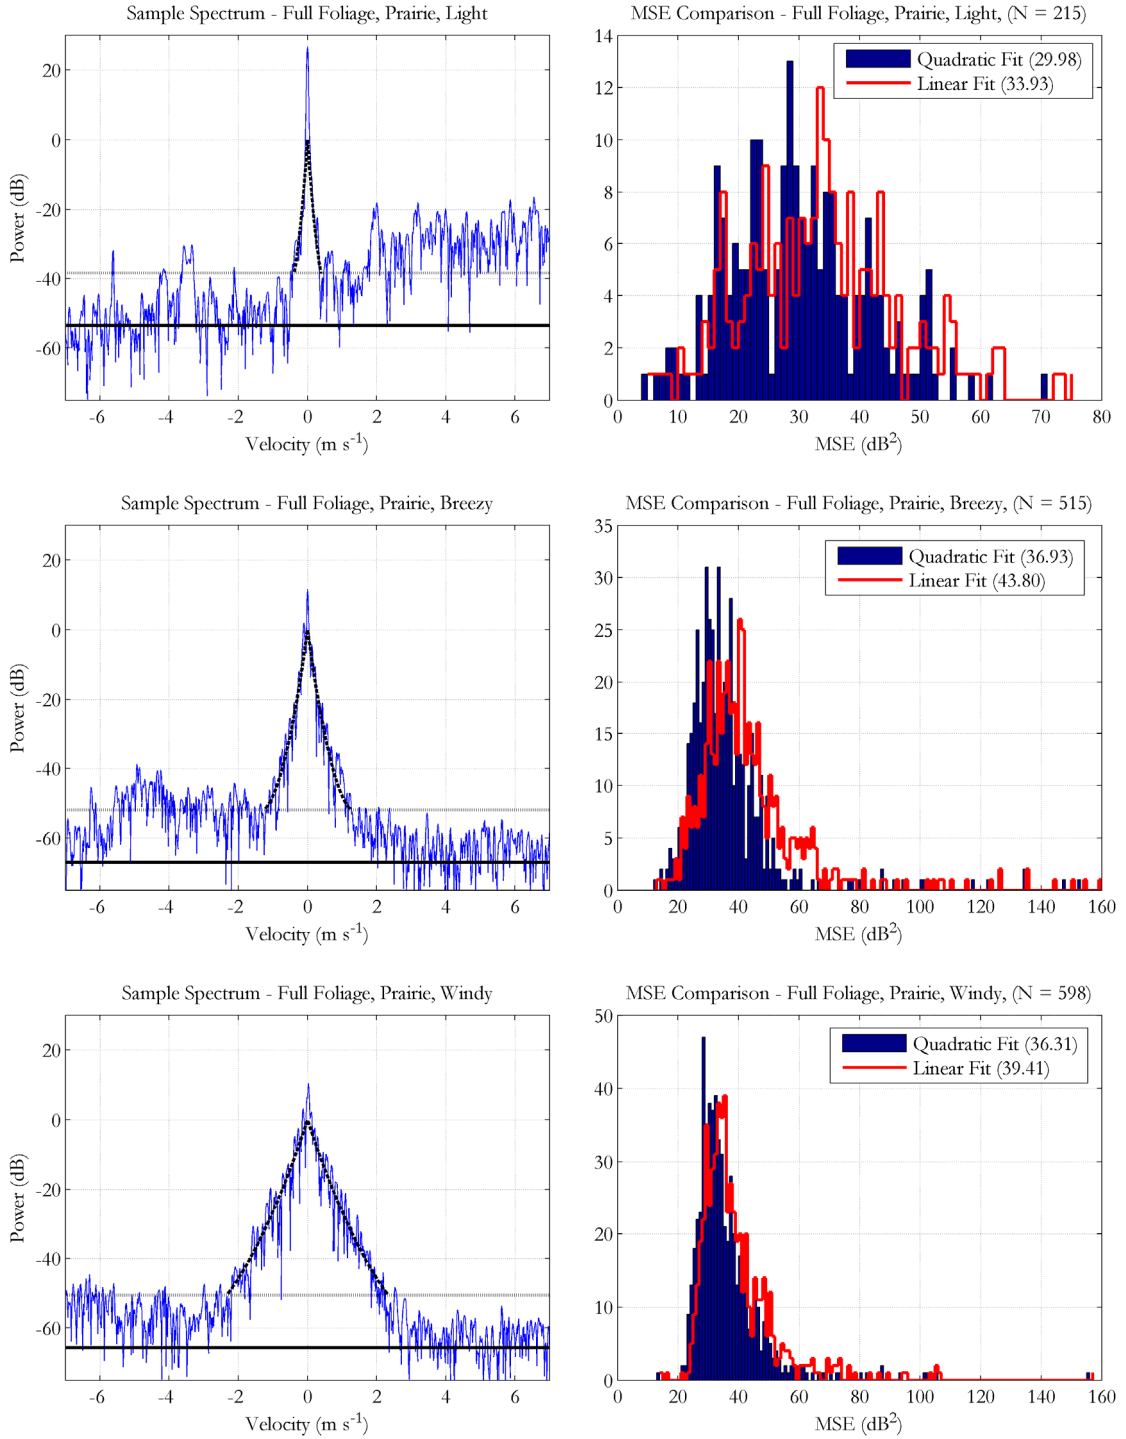


Figure D.2. Sample ground clutter spectra and MSE histograms for the “Full Foliage, Prairie” conditions. (Estimated model parameters: Light - $\hat{a} = 238.4 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 145.8 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 53.6 \text{ dB}$, $\hat{d} = 26.8 \text{ dB}$, Breezy - $\hat{a} = 51.8 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 73.8 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 67.0 \text{ dB}$, $\hat{d} = 11.8 \text{ dB}$, Windy - $\hat{a} = 6.8 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 30.0 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 65.3 \text{ dB}$, $\hat{d} = 7.0 \text{ dB}$.)

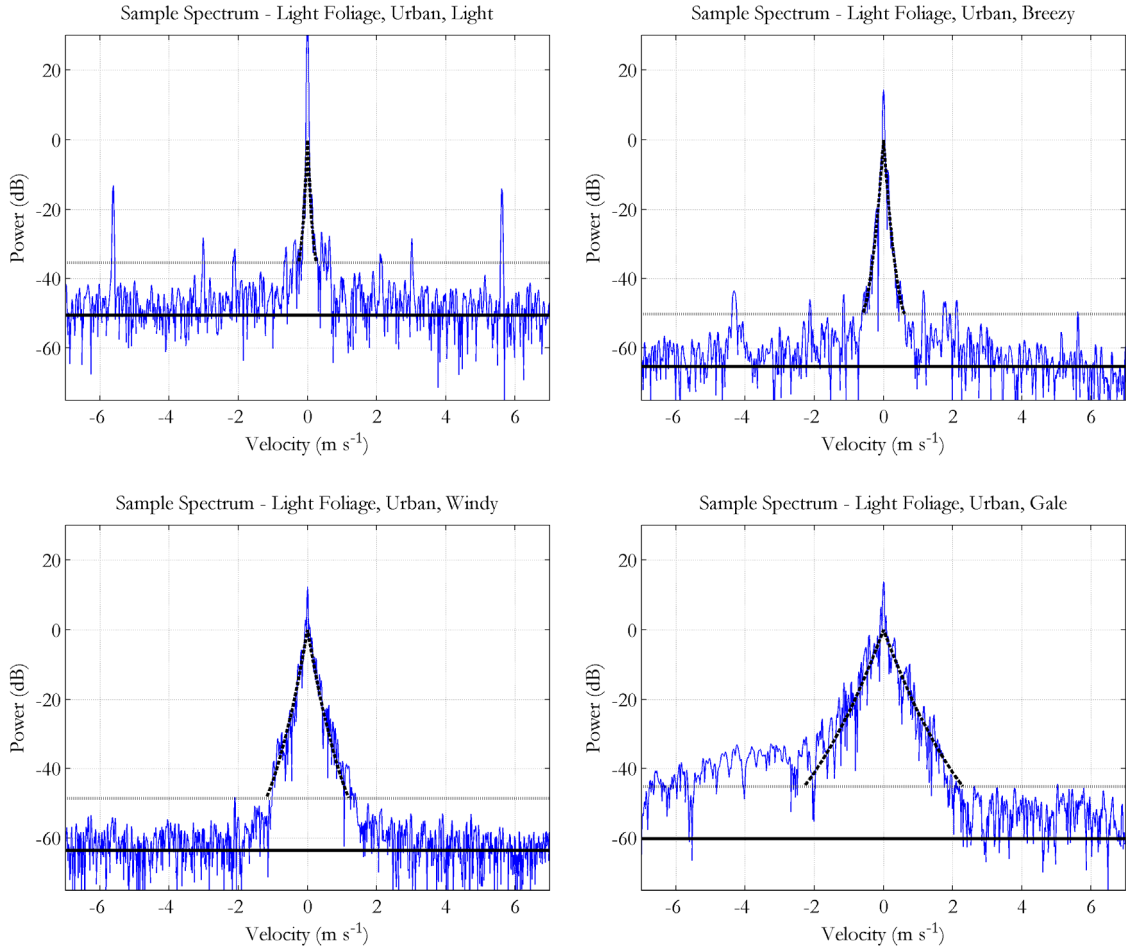


Figure D.3. Sample ground clutter spectra for the “Light Foliage, Urban” conditions.

(Estimated model parameters: Light - $\hat{a} = 785.7 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 239.0 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 50.6 \text{ dB}$, $\hat{d} = 32.9 \text{ dB}$, Breezy - $\hat{a} = 196.4 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 142.6 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 65.3 \text{ dB}$, $\hat{d} = 14.5 \text{ dB}$, Windy - $\hat{a} = 28.8 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 57.9 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 63.6 \text{ dB}$, $\hat{d} = 12.4 \text{ dB}$, Gale - $\hat{a} = 6.2 /(\text{m}^2 \text{ s}^{-2})$, $\hat{b} = 26.8 \text{ dB}/(\text{m s}^{-1})$, $\hat{c} = 60.3 \text{ dB}$, $\hat{d} = 13.9 \text{ dB}$.)

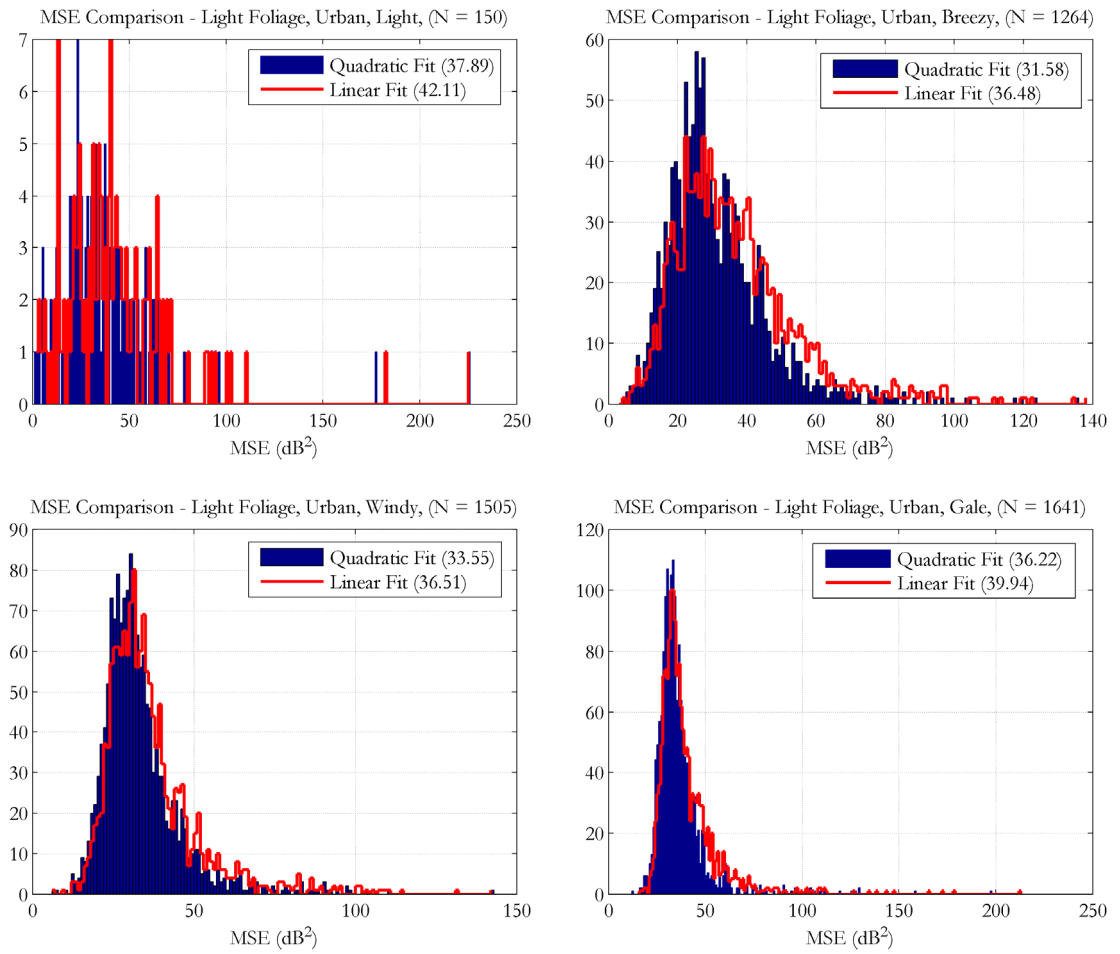


Figure D.4. Histograms of the MSE for both the quadratic and linear fits for the “Light Foliage, Urban” conditions.

Appendix E. The d Parameter: Normalization and Window Correction

The value of the DC parameter depends on the number of elements, N , in the spectrum. The d values that are dependent on the number of elements will be denoted as d_N . Because d_N reflects the amount of DC power in terms of the AC power in the zero-velocity bin, d_N depends on the width of the Doppler bins. If $N = 4096$ and there is a constant AC power of one across the zero-velocity bin, the amount of AC power in that bin is $(2 \cdot v_a)/4096$ where v_a is the unambiguous velocity. If the number of bins is doubled to 8192, the value of d_{8192} will increase by about 3 dB compared to d_{4096} because the power from the AC part of the spectrum will be reduced by about one half while the DC power stays the same.

Finding a value of d that does not depend on the number of elements in the spectrum provides an independent way for specifying the DC parameter. In Billingsley's two-part exponential model (2002), the DC parameter is scaled so that the sum of the DC and AC power is one. Although this approach could be used, it would complicate comparisons of the DC component when the underlying shape of the AC part changes. Instead, a normalized d_n will be defined independently of the power in the AC portion by relating the DC power to a fixed value of one. Any fixed value could be used, but this leads to a simple expression for the power, $10^{\frac{d_n}{10}} - 1$, where d_n is given in dB. The power in terms of d_N when the number of elements in the spectrum is N is given by $P_N \left(10^{\frac{d_N}{10}} - 1 \right)$ where P_N is the AC power in the zero-velocity bin. If d_N and P_N are known, the normalized value can be computed as follows:

$$d_n = 10 \log_{10} \left(P_N \left(10^{\frac{d_N}{10}} - 1 \right) + 1 \right) \quad (\text{E.1})$$

The key is to find the value of P_N . As mentioned previously, P_N should be very close to $(2 \cdot v_a)/N$ because the value at the center of the zero-velocity bin is one. For more accuracy, the integral of the AC portion of the spectrum can be computed.

The integral of the quadratic section of the AC portion of the spectrum given in equation (2.12) can be found using the imaginary error function based on a result from Gradshteyn et al. (2000):

$$\int 10^{\frac{-bv + \frac{a}{2}v^2}{10}} dv = \sqrt{\frac{5\pi}{a \log(10)}} \exp\left(-\frac{b^2 \log(10)}{20a}\right) \operatorname{erfi}\left(\frac{1}{2} \sqrt{\frac{a \log(10)}{5}} v - \frac{1}{2} \sqrt{\frac{\log(10)}{5a}} b\right). \quad (\text{E.2})$$

This equation is valid for $a \neq 0$. If the imaginary error function is not available, numerical integration can be used. Using this equation, the value of P_N can be found by evaluating the following definite integral.

$$P_N = 2 \int_0^{\frac{v_a}{N}} 10^{\frac{-bv + \frac{a}{2}v^2}{10}} dv \quad (\text{E.3})$$

For example, assume a model with the following parameters: $a = 20$ dB/(m² s²), $b = 50$ dB/(m s⁻¹), $c = 60$ dB, and $d_{4096} = 20$ dB. Using equation (E.3), the normalized value can be computed which results in $d_n = 3.7$ dB. If the power was assumed to be one across the whole bin, the estimate would be 3.8 dB which is very close the value based on the integral. Although these two values are within 0.1 dB, the more accurate calculation using the integral becomes more important when N is smaller and the Doppler bins are wider.

A plot showing the quadratic clutter model with $a = 20$ dB/(m² s²), $b = 50$ dB/(m s⁻¹), $c = 60$ dB, and $d_{4096} = 20$ dB (or $d_n = 3.7$ dB) is shown in Figure E.1. This shows how the

spectrum would look if $N = 4096$. A plot of the spectrum with another value of N would have a different value of d_N . This difference in d_N values needs to be taken into account when simulating time series data especially when additional Doppler bins are used to incorporate windowing effects. By solving for d_N instead of d_n in equation (E.1), an expression for any d_N can be found given the normalized d_n value.

$$d_N = 10 \log_{10} \left(\frac{1}{P_N} \left(10^{\frac{d_n}{10}} - 1 \right) + 1 \right) \quad (\text{E.4})$$

The d_N values can then be used in time series simulations depending on the number of elements in the spectrum.

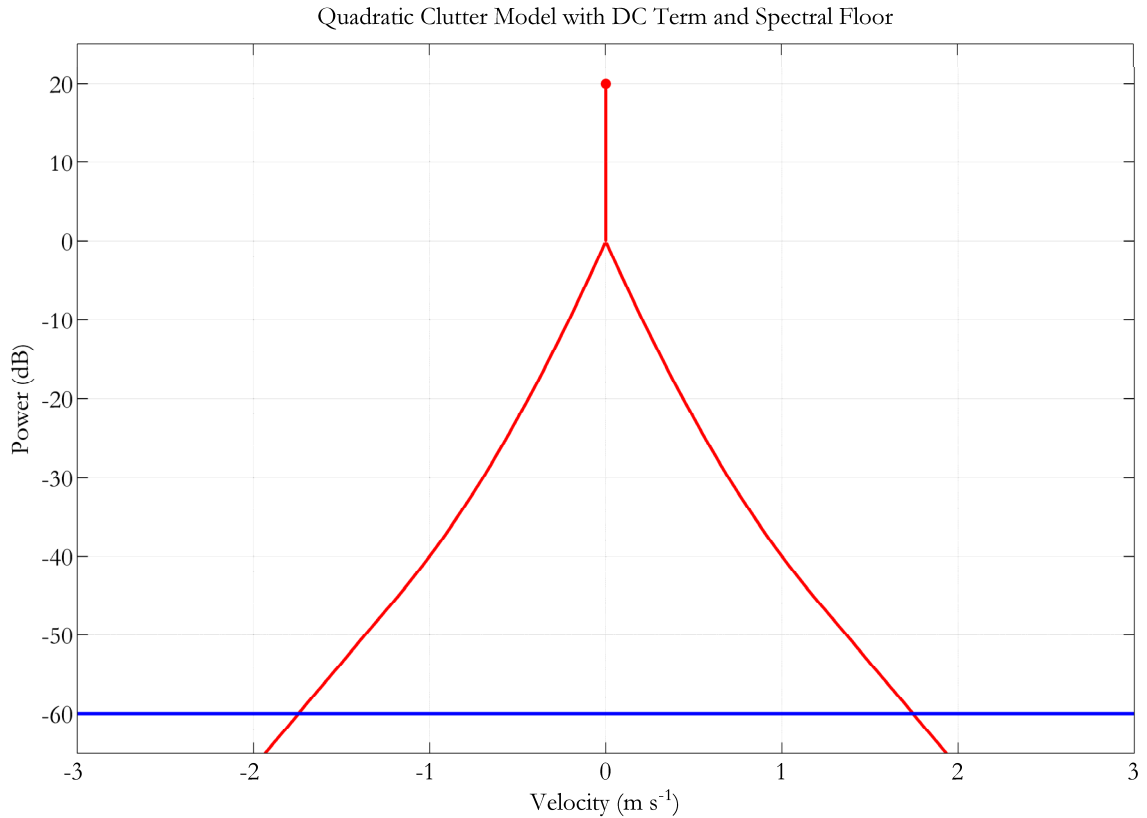


Figure E.1. Quadratic log-domain model with $a = 20$, $b = 50$, $c = 60$, and $d_{4096} = 20$ (or $d = 3.7$ db).

As stated earlier, the DC peak is broadened because of the application of the Chebyshev window. If d is estimated from a spectrum of length 4096 (d is a windowed version of d_{4096}), a normalized version of the estimated parameter, d_w , is computed first using equation (E.1). Then, the non-windowed version of the normalized parameter, d_n , is found. The relationship between d_w and d_n was determined through simulations. Several sets of time series data were simulated with the parameters used earlier ($a = 20 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $b = 50 \text{ dB}/(\text{m s}^{-1})$, $c = 60 \text{ dB}$), and several additional sets were generated for another set of parameters ($a = 0 \text{ dB}/(\text{m}^2 \text{ s}^{-2})$, $b = 20 \text{ dB}/(\text{m s}^{-1})$, $c = 60 \text{ dB}$). This second set of parameters corresponds to a linear shape for the AC part of the model. The value of d_n for the simulations was varied from 0 to 25 dB in 1 dB increments. The data were fit to a function that is zero when d_w is zero and approaches a straight line for larger d_w . The results are shown in Figure E.2. The approximate relationship between the estimated d_w and d_n is given by the following formula:

$$d_n = \left(0.4d_w^{1.3} + 2d_w\right)^{\frac{1}{1.3}} \quad (\text{E.5})$$

For example, this formula returns a value of 3.7 dB for d_n when d_w is estimated to be 2.2 dB.




Figure E.2. Data and fit showing relationship between d_n used in the simulations and the estimated d_w .