

UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

MULTISTATIC PASSIVE WEATHER RADAR

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

By
ANDREW BYRD
Norman, Oklahoma
2020

MULTISTATIC PASSIVE WEATHER RADAR

A DISSERTATION APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Robert Palmer, Chair

Dr. Caleb Fulton, Co-Chair

Dr. Nathan Goodman

Dr. Pierre Kirstetter

Dr. Alan Shapiro

Dr. Dusan Zrnić

© Copyright by ANDREW BYRD 2020
All Rights Reserved.

*For my wife, Suzy. I can't imagine beginning this next great adventure with anyone
but you.*

Acknowledgments

I have to begin by thanking Robert Palmer and Caleb Fulton, who have been my advisors, my coauthors on the two journal papers that form the cornerstone of this dissertation, and my friends. Dr. Palmer extended my initial offer to work at the ARRC and has, since then, set an almost implausibly high standard for leadership. It is a very rare advisor or manager who may be so confidently depended upon for advocacy and invaluable insight (both technical and otherwise) by those who work for him. Dr. Fulton is a model of astonishing depth and breadth of technical brilliance, but also of how to take genuine joy in the work of engineering. I think most people at the ARRC have a great love for their work, but Dr. Fulton's is so incandescent that it leaves you little choice but to have fun. Working for and with these two has truly been a privilege.

I also owe a tremendous debt to the family that has supported me through this journey. My mother, Janet, has helped to support me in every way possible through the years, and no thanks I give her will ever be sufficient. I am extremely grateful for my wife, Suzy, who has been incredibly loving and patient throughout this process. Suzy's family, Fernando, Lupe, and Sylvia have also been wonderfully supportive and welcoming. They have truly given me a home in Oklahoma (quite literally in these last months).

Finally, there are some necessary acknowledgments specific to the work detailed in Chapters 3-5. I wish to extend special thanks to ARRC staff engineer Cody Piersall for his work on the radar control software, and to the administration and facilities staff at OUHSC (particularly Stuart Hall and Dustin Bozarth) for their extraordinary level of cooperation. Thanks to the administration and engineers at the Radar Operations Center (particularly Adam Heck, Steven Smith, and Terrance Clark) for their help in understanding the finer details of WSR-88D operations.

Thanks also to Alan Shapiro for his excellent suggestion of using the KTLX radial velocity comparison to validate the system and to David Bodine for several useful discussions of multistatic radar applications, as well furnishing the numerical weather prediction data used in the simulations for Chapter 5, and, in cooperation with Clarice Dyson, for operating the passive radars to collect the scattered convection observations in Chapter 4. Finally, thanks to Dusan Zrnić for some helpful discussion and references regarding sidelobe whitening.

This work was supported by the National Severe Storms Laboratory through cooperative agreement NA15OAR4320115.

IEEE Copyright Notice

The contents of two journal papers (for which I was the first author) were reused substantially in Chapters 3-6 of this work. Chapters 3-4 are composed primarily of material drawn from [1] (©2019 IEEE). Chapter 5 is composed primarily of material from [2] (under review, IEEE). Chapter 6 contains substantial reprinted material from [1] and [2]. There are additional figures, datasets, and explanations added to the material present in these papers. Some figures may appear slightly different from the corresponding figures in these works due to updates and improvements to the processing software. The paper structure has also been rearranged to accommodate the dissertation format. However, these changes are not substantial enough to distinguish these chapters as new works independent from these papers.

Table of Contents

1	Introduction	1
1.1	Multiple Doppler Observations With Monostatic Radar	2
1.2	Multistatic Weather Radar	5
1.3	Motivation	9
1.4	Outline	11
2	Multistatic Weather Radar Fundamentals	13
2.1	Bistatic Radar Range Equation	14
2.2	Spatial Resolution	16
2.3	Echo localization	22
2.4	Velocity Estimation	27
2.5	Polarimetric Considerations	30
2.6	Wind Field Estimation	32
3	Design of a Multistatic Weather Radar Receiver	37
3.1	Passive Radar Signal Processing	40
3.2	Coarse Time Alignment and CFAR	42
3.3	Standard PRF Fitting	46
3.4	Batch PRF Fitting	50
3.4.1	Original Method	50

3.4.2	Improved Method for Precise Time Alignment	54
3.5	Phase Decoding and Frequency Offset Removal	55
3.6	Transmit Antenna Pattern Phase Effects	63
3.7	Receiver Module Hardware	66
4	Weather Observations With a Multistatic Network	71
4.1	Experimental Validation of Velocity Measurements	71
4.2	Convective Observations	78
4.3	Discussion of Sidelobe Contamination in Collected Data	85
5	Sidelobe Whitening Simulation and Analysis	91
5.1	Pattern Synthesis for Sidelobe Whitening	93
5.1.1	Phase and Amplitude Constraints	96
5.1.2	Least-Squares Excitation Retrieval	100
5.1.3	Modified Two-Dimensional Implementation	101
5.2	Whitening Algorithm Results	103
5.3	Weather Radar System Simulations	106
5.3.1	Simulator Description	106
5.3.2	Results	119
6	Conclusions and Future Work	122
6.1	Conclusions	122
6.2	Future Work	124
Appendix A Derivation of the Relationship Between Scatterer Velocity and Multistatic Velocity Measurements		140
Appendix B List of Symbols		147

List of Figures

2.1	The fundamental geometry of a bistatic radar observation including the transmitter range R_T , receiver range R_R , baseline length L , and bistatic angle β . The arrows indicate the direction of radiation propagation along each path.	14
2.2	An example of the transmitter and receiver locations as the focii of an ellipse of constant range (shown in purple). This ellipse is a cross-section of the three-dimensional constant-range ellipsoid along the bistatic plane. For a scatterer located at any point on this ellipse, the bistatic range $R_B = R_T + R_R$ will be constant.	17
2.3	The range and beamwidth boundaries that form a resolution volume in a typical bistatic weather radar. The concentric ellipses correspond to constant bistatic range ellipses. The translucent cones represent the main beam widths of the transmitter and receiver. The resolution volume formed by the intersection of the transmit beam with the ellipses is highlighted by a black dashed line.	18
2.4	Example of a degenerate resolution volume geometry in which the near bistatic range boundary is defined by the system baseline and the far bistatic range volume does not come into play, as the exterior volume edges are defined by the transmit and receive beampatterns.	22

2.5	Diagram of the geometric quantities used to calculate Cartesian scatterer coordinates in (2.12)-(2.14).	23
2.6	Echo localization through multilateration in a network consisting of a single transmitter and three receivers. The purple ellipses represent isorange contours corresponding to detection of the actual scatterer by each receiver. If only two receivers are used, there will be two candidate detection locations where the isorange contours intersect: one at the location of the actual scatterer and a “ghost” represented by a white triangle. The third receiver is necessary for disambiguation.	25
2.7	Complications of multilateration when multiple targets are present. Ghost targets are represented by white triangles. The especially problematic ghost target at which three ellipses intersect is indicated in red.	26
2.8	Relevant geometric quantities for modeling Doppler velocity measurements using a bistatic radar.	28
2.9	Three-dimensional bistatic RCS of raindrop with impingent radiation polarized along (a) the z-axis (V polarization) and (b) the y-axis (H polarization).	31
2.10	One-dimensional Cressman weighting function for $R_i = 1$	33
3.1	This satellite image [44] shows the locations of the two passive receivers at the Radar Innovations Lab (RIL) and University of Oklahoma Health Science center (OUHSC), as well as the location of KTLX, the WSR-88D being utilized as a transmitter of opportunity and also operating as an independent monostatic radar.	38

3.2	Flowchart depicting an overview of the signal processing scheme used to translate the passive radar time-series into wind field estimates. CFAR detection is followed by a quality control process based on PRF-fitting. The frequency offset between transmitter and receiver is removed, and then estimates of power and velocity are calculated and localized to Cartesian coordinates. Finally the estimates are interpolated to a common grid and wind-field estimation is performed. Special WSR-88D transmit modes (batch PRF and SZ-2 coding) are detected and handled appropriately.	41
3.3	The approximate matched filter used to process the received time series data. It was derived from experimental measurements of WSR-88D pulses. The half-power width of the approximate filter is consistent with the $1.57\mu s$ nominal pulse length. The index n assumes a sampling rate of 5 MSPS.	42
3.4	In this segment of data collected from a transmitted WSR-88D signal, the rising edge of a direct-path pulse is located near sample 400. Note the drastically higher power levels due to ground clutter and weather after the pulse in contrast to the samples collected prior. This severity of the issue is more striking given the context that the pulse itself is only about 8 samples long ($1.57\mu s$ at a 5 MSPS).	43
3.5	Measurements of direct-path pulse power over time. These pulses form a repeating measurement of the WSR-88D beam pattern at changing elevation cuts based on the mechanical tilt of the transmitter.	45

3.6	This 0.44° elevation beampattern cut is taken from the same dataset used to produce Figure 3.5. Note the saturation of the main beam as well as the apparent irregularities in the sidelobe topography caused by strong ground clutter.	45
3.7	High-level outline of the PRF fitting process used to quality control the pulse locations estimated by the CFAR detector.	46
3.8	Example of the “noisy sawtooth” form exhibited by the error function $\delta[n]$, as calculated from actual recorded data. Also shown is the error function recalculated after systematic error in samples-per-PRT has been corrected using the algorithm described in Section 3.1.	48
3.9	Example of measured pulse-to-pulse phase rotation distributions after decoding using each of the 8 possible delays. A delay of 2 samples is clearly the correct option, as shown by the low variance of the resulting phases.	53
3.10	Example of a simulated pulse train spectrum, both with and without a strong multipath interferer. Note that the peak of the macroscopic spectrum shown in the top plot shifts significantly due to the change in mainlobe shape induced by the multipath signal, but the locations of the individual “samples” visible at the finer scales shown in the lower two plots are virtually unaffected. This compares favorably with the measured data in Figure 3.11.	58

3.11	Example of two consecutive measured pulse train spectra with varying multipath. It shows the same effect as in the simulations of Figure 3.10: significant peak movement at a macroscopic scale but extreme consistency at finer scales. Note that the smoothing and rapid rolloff of the expected sidelobes in the spectrum are due to low-pass filtering within the transceiver.	59
3.12	Loss in matched filter sensitivity due to errors in frequency estimation.	61
3.13	Example azimuthal pattern of the X-band proxy antenna used to study the effects of pattern phase on velocity estimation biases. . . .	63
3.14	Results of the pattern phase error analysis across a range of rotation rates and transmit antenna mechanical elevations. These results reflect a 0.5 s data collection interval, and each plotted point corresponds to the worst possible result across the measured antenna pattern range.	65
3.15	Block diagram of the passive radar hardware.	67
3.16	Normalized radiation patterns of the receiver antenna measured in the ARRC far-field chamber. As the antenna is V-polarized, the E-plane corresponds to elevation, and the H-plane corresponds to azimuth.	68
3.17	An installed system on the roof of the Radar Innovations Lab at the University of Oklahoma.	69

4.1	Raw estimates of range-corrected power (not reflectivity) and bistatic Doppler velocity obtained by the passive receivers, along with reflectivity and radial velocity estimates from KTLX for comparison. The black contour illustrates the region selected for analysis based on the censoring criteria described in Section 4.1. The data shown were collected using a 4° elevation KTLX scan on April 6, 2019 at 16:11 UTC. Note that a ground clutter filter has not yet been applied to these data, in order to more clearly show the logic behind the censoring boundaries.	72
4.2	Retrieved horizontal wind vectors over the analysis region.	73
4.3	Actual KTLX radial velocity field and the estimate of the KTLX radial velocity obtained by projecting the retrieved horizontal wind field onto the vectors representing the pointing direction of KTLX at each point.	74
4.4	Results of the KTLX radial velocity retrieval are depicted here as (a) spatial map of the root squared error in the retrieved estimate and (b) scatterplot of the relationship between the measured and retrieved values. The blue line in (b) represents a theoretical exact match between retrieved and measured values.	77

4.5	Raw estimates of range-corrected power (not reflectivity) and bistatic Doppler velocity obtained by the passive receivers, along with reflectivity and radial velocity estimates from KTLX for comparison. The black contour illustrates the region selected for analysis based on the censoring criteria described in Section 4.1. The data shown were collected using a 4° elevation KTLX scan on May 25, 2019 at 2:58 UTC. Note that a ground clutter filter has not yet been applied to these data, in order to more clearly show the logic behind the censoring boundaries.	79
4.6	Retrieved horizontal wind vectors over the analysis region.	80
4.7	These plots show the actual KTLX radial velocity field as well as the estimate of the KTLX radial velocity obtained by projecting the retrieved horizontal wind field onto the vectors representing the pointing direction of KTLX at each point.	82
4.8	Results of the KTLX radial velocity retrieval are depicted here as (a) spatial map of the root squared error in the retrieved estimate and (b) scatterplot of the relationship between the measured and retrieved values. The blue line in (b) represents a theoretical exact match between retrieved and measured values.	83

4.9	a) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the RIL receiver b) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the OUHSC receiver c) RIL receiver range-corrected power with corresponding bistatic range contours and analysis region outline d) OUHSC receiver range-corrected power with corresponding bistatic range contours and analysis region outline. See Figure 4.1 for a wider view of the same observation set.	88
4.10	a) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the RIL receiver b) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the OUHSC receiver c) RIL receiver range-corrected power with corresponding bistatic range contours and analysis region outline d) OUHSC receiver range-corrected power with corresponding bistatic range contours and analysis region outline. See Figure 4.5 for a wider view of the same observation set.	90
5.1	High-level summary of the proposed pattern synthesis technique. . .	93
5.2	Example whitening pattern synthesis results for a 152 element $\lambda/2$ -spaced ULA. The mask used is an approximation of the WSR-88D azimuth pattern at 0° elevation. Coding phase is the phase difference between the two synthesized patterns to be used in the binary whitening scheme. Note that there is some deviation from ideal results (0 in the mainlobe and π elsewhere), particularly within the first sidelobe and in the pattern nulls.	96

5.3	Evolution of the convergence criterion ϵ during the synthesis of the pattern shown in Figure 5.2. It crosses the threshold value of 1×10^{-12} after 80 iterations.	101
5.4	Map showing which of the three possible whitening codes is applied to each sidelobe region of a sample two-dimensional array pattern. .	102
5.5	Results of a 64 point whitening code used in conjunction with the synthesized pattern. Doppler spectrum peak attenuation is the reduction in the maximum Doppler spectrum value for an impinging clutter signal at each angle. As expected, near zero attenuation occurs in the mainlobe.	103
5.6	Several examples of the clutter signal Doppler spectra produced using the 64 point code (calculated with a Hamming window). As anticipated based on the peak attenuations in Figure 5.5, the mainlobe spectrum receives virtually no spectral spreading, although the sidelobes are slightly perturbed relative to an ordinary Hamming window spectrum. The far sidelobe is extremely well whitened, with DC no longer having a dominant peak, and the first sidelobe spectrum lies between these extremes.	105
5.7	High-level view of the simulation process. The receiver-related portion of the radar range equation calculation, as well as the sorting of scattering centers into range bins, are done prior to calculation of the transmitter-related weighting contributions. This structure maximizes the efficiency of the simulation process.	107

5.8	Gain patterns for each simulated transmitter scenario. Since the one-dimensional array factors of the two binary whitening patterns are identical, this means that the two-dimensional power pattern will remain constant during the whitening process and is also identical to that of the unwhitened array pattern.	108
5.9	Constant-altitude slice of the NWP data grid used to produce the simulation results. Reflectivity is shown in (a) while (b) and (c) show the zonal and meridional wind field components, respectively. The receiver and transmitter locations are also indicated, as well as a dashed line indicating the boundary of the observation region shown in Figures 5.10-5.11.	114
5.10	Doppler velocity fields measured by each receiver for each simulated transmitter. There are significant errors due to sidelobe contamination in each image, primarily along areas of sharp reflectivity and velocity gradients. Note also regions of degraded spatial resolution along the transmitter/receiver baselines.	115
5.11	Differences between the dish, unwhitened array, and whitened array simulations and the sidelobe-free ideal simulation. There are significant reductions in bias prevalence and magnitude between the dish and either of the arrays. Whitening provides a noticeable improvement compared to the unwhitened array.	116
5.12	Examples of the unwhitened, whitened, and ideal spectra from the simulation results shown in Figure 5.10. The sidelobe leakage visible in the unwhitened results is spread throughout the whitened spectrum, resulting in a closer match to the ideal results.	121

Abstract

Practical and accurate estimation of three-dimensional wind fields is an ongoing challenge in radar meteorology. Multistatic (single transmitter / multiple receivers) radar architectures offer a cost effective solution for obtaining the multiple Doppler measurements necessary to achieve such estimates. In this work, the history and fundamental concepts of multistatic weather radar are reviewed. Several developments in multistatic weather radar enabled by recent technological progress, such as the widespread availability of high performance single-chip RF transceivers and the proliferation of phased array weather radars, are then presented. First, a network of compact, low-cost passive receiver prototypes is used to demonstrate a set of signal processing techniques that have been developed to enable transmitter / receiver synchronization through sidelobe radiation. Next, a pattern synthesis technique is developed which allows for the use of sidelobe whitening to mitigate velocity biases in multistatic radar systems. The efficacy of this technique is then demonstrated using a multistatic weather radar system simulator.

Chapter 1

Introduction

Radar is among the most powerful and widely utilized technologies for remote observation of the atmosphere. However, it is not without limitations. Prominent among these is the fact a single radar receiver is only capable of measuring the component of scatterer motion projected onto a single vector in three-dimensional space. For the most common case of a monostatic radar (a system in which the transmitter and receiver are colocated), this single vector extends radially along the pointing direction of the transmitter, giving rise to the familiar “radial velocity” product. Atmospheric motion, however, occurs in three dimensions, all of which are critical for understanding and predicting the weather. In order to completely reconstruct the motion of a given scatterer, it is necessary to measure its velocity along a minimum of three linearly independent vectors or to impose some additional constraints on scatterer motion using fluid dynamics. The desire for complete three-dimensional wind field reconstruction has, for this reason, given rise to the concept of multiple Doppler observations, in which two or more radar systems observe a volume of weather as simultaneously as possible in order to measure wind fields

more directly and reduce dependence on theoretical constraints. One method of obtaining multiple Doppler observations is the use of a multistatic (single transmitter / multiple receiver) radar network. The ultimate goal of the work presented here is to significantly decrease the expense and improve the practicality of these multistatic networks through the development of new signal processing techniques. This introduction seeks to contextualize that effort within the existing body of work on multiple Doppler data and methods, and on the radar systems used to obtain these types of measurements.

1.1 Multiple Doppler Observations With Monostatic Radar

The concept of using a pulsed-Doppler radar for the purpose of measuring wind speed was first documented in a brief 1960 *Nature* article by J.R. Probert-Jones [3]. This work merely notes the possibility that the radial component of wind velocity can be measured. Work over the subsequent decade gradually generalized single-radar wind field estimation using the concepts of Velocity-Azimuth Display (VAD) [4] and Volume Velocity Processing (VVP) [5], which can be useful for determining the wind fields in the immediate surroundings of a single radar given certain assumptions on the homogeneity of both the wind field and hydrometeor size / fall speed in that area [6]. However, the problem of estimating the wind field in a region *remote* to a radar system inevitably required the coordination of two or more systems. The first attempt at this problem was made by Lhermitte [7] for the simplified case limited to two radars and an assumed negligible contribution of hydrometeor fall speed to observed radial velocities. Armijo [8] expanded on this work by incorporating the continuity equation into a multiple Doppler analysis. This allowed for

the assumption of known, rather than negligible, hydrometeor terminal velocities in the two radar case, and imposed no such assumptions for the three radar case, for which he also outlined an analysis procedure. Multiple Doppler wind field estimation has continued to evolve in the years since the Armijo paper, however the basic process remains recognizable. The radial velocities observed by a set of radars are combined with information about hydrometeor terminal velocity (typically derived from reflectivity) and/or the continuity equation. Innovation has come primarily in the form of improvements in how the resulting systems of equations are solved and in additional assumptions and constraints that can be added in order to reduce error. The analytic solution developed by Armijo requires an integration of the continuity equation that is susceptible to undesirable levels of error accumulation [9]. In order to combat this effect, variational (least squares) methods have been employed. These allow for the solution of an overdetermined problem, meaning that additional criteria such as smoothing constraints and multiple boundary conditions can be used to fit an optimal wind field estimate [10]–[12]. Iterative techniques have also been developed for the purpose of deriving solutions to these fitting problems that will satisfy the continuity equation in a Cartesian coordinate system [9]. While these methods have been useful in improving the accuracy of radar-based wind field estimation, they do not address the fundamental challenges of temporal variation and coordination between radars - both critical in the collection of useful multiple Doppler observations.

Temporal variation in observed wind fields is problematic for multiple Doppler analysis due to the fact that (typically) radial velocity measurements throughout the analysis domain cannot be collected simultaneously. The wind field is evolving while the measurements are being collected, which poses particular problems when attempting to apply the continuity equation or perform spatial interpolation

of sampled points onto a Cartesian grid. This effect is most dramatic between resolution volumes that are adjacent in elevation, due to the azimuth-first scanning strategies of typical weather radars. This is a challenge even for the single-radar variety of wind field estimation, but it is exacerbated in the case of multiple radars. Not only should the data collected by each radar be contemporaneous with each other, but simultaneity is also desired between the samples collected by each of the different systems. While it is possible (although not necessarily simple) to coordinate radar operation in such a way that differences in temporal synchronization are minimized, exact agreement is quite impossible. Conical scans by two spatially separated radars can only overlap to a limited extent, and the problem is only exacerbated for larger networks. For this reason, advection correction techniques are used to rectify mismatches in temporal sampling, both within and between radar systems.

Chong et al. [13] suggest that temporal wind field variation associated with a storm be thought of as having two components: intrinsic variation consisting of evolution in the internal structure of the storm's wind field, and advection consisting of the storm's translational motion due to some uniform background wind field. Correction of the advection component is the most common and most readily achievable way to reduce errors due to temporal variation in the wind field. However, sophisticated non-uniform advection correction methods have been developed in recent years [14], [15] that seek to address the intrinsic variation of the wind field. These techniques have been shown to improve advection correction performance compared to simple correction of uniform translational motion. However, they still suffer from limitations including sensitivity to the initial guess used to initialize the process of correction optimization, as well as the possibility of temporal aliasing/non-uniqueness. In light of these limitations it is natural that engineering

solutions to the temporal variation problem have been sought. The use of a single-transmitter / multiple receiver radar network is one potential way to mitigate the problem of inter-radar scan synchronization.

1.2 Multistatic Weather Radar

Bistatic radar systems, in which the distance between the transmitter and receiver have a separation on the order of the target distance, have a long history in defense applications [16]. Initially, a bistatic radar presented a less radical change in architecture from a monostatic system as it does today, since the duplexer was not invented until 1936 [16]. Transmitters and receivers were separate by necessity until that point, even when colocated. Bistatic systems saw renewed interest in the modern era due in no small part to the advantages they offer in the context of electronic warfare. Most obviously, a transmitter inherently reveals its location, making a colocated receiver vulnerable to directive jamming or physical attack. By contrast, a bistatic receiver would remain hidden. Passive radar takes this concept a step further. These systems consist solely of a receiver that observes targets using signals transmitted by other radars or communications devices outside of its control. Interestingly, these concepts were key to some of the earliest radar experiments, which also had close ties to atmospheric science. The earliest experiments in radio direction-finding conducted by Robert Watson-Watt utilized passive receivers to localize lightning strikes based on their associated microwave radiation [17]. When the British government requested a proof of concept demonstration before authorizing funding for the development of the first aircraft detection radar, Watson-Watt showed that the idea was feasible through passive radar observations using radio broadcast signals as an illumination source [18].

Interest in bistatic radar within the context of modern atmospheric radar be-

gan in earnest in the late 60s. The original focus of research in this area was on clear-air observations [19]. This is due to the fact that the reflectivity of a turbulent medium has a strong positive dependence on forward scatter angle. This allows clear-air observations using bistatic configurations at a much shorter wavelength than what would be realizable using a monostatic system. Subsequent experiments over the next few years produced bistatic measurements of scattering intensity from the melting layer [20], as well as successful bistatic Doppler measurements of precipitation [21]. The concept of using a single-transmitter / multiple receiver network for the collection of multiple Doppler weather observations was put forward in a 1993 paper by Wurman et al. [22]. The primary scientific advantage of this methodology is that it solves the inter-radar scan simultaneity problem, as the echoes from a given volume measured by each of the passive receivers originate from the illumination of that volume by the transmitter at a common instant. This significantly reduces the complexity and potential for error introduced by the advection correction schemes needed to perform wind field estimation. The second major advantage of this architecture is economic; the cost of a multistatic network is miniscule compared to that of a monostatic radar network with the same number of sites. This is attributable to both the fact that the receivers do not need the electronics necessary to produce a powerful transmit signal, and that they do not need the large antenna aperture necessary to form a highly directive beam.

The central problem of bistatic / multistatic radar design is how to achieve synchronization and coherence between the transmitting and receiving systems. This encompasses both carrier frequency synchronization, which must be extremely precise in order to obtain accurate Doppler information, as well as pulse timing synchronization, which is necessary in order to accurately determine scatterer locations. There is a spectrum of approaches available for achieving this kind of syn-

chronization. On one end of this spectrum is a set of cost and infrastructure intensive hardware-based solutions. Due to the fact that two free-running local oscillators (LOs), which govern the carrier frequency at the transmitter and receiver, will drift relative to each other at levels that can cause significant errors in Doppler velocity estimation, some technique must be used to lock the LOs at the hardware level or correct for this drift in post-processing. The more expensive hardware-based synchronization solutions typically use GPS-disciplined oscillators in order to maintain carrier frequency synchronization. This is often complemented by communication of pulse-timing information over a dedicated communications channel between the transmitter and receiver [22]. This type of approach has the advantage of simplicity, precision, and low signal processing overhead. However, it has the disadvantages of more expensive hardware, strict specifications on the transmitter and receiver hardware, and potentially the need to tie receivers to a specific transmitter due to the communications infrastructure. A less hardware intensive approach [23] may involve using inexpensive hardware, for instance a free-running but high-quality ovenized oscillator, that is capable of maintaining acceptable stability over some temporary interval. The example in Wurman also eschews the use of a dedicated communications channel and achieves pulse timing synchronization through a manual tuning process in which the transmitting radar is pointed directly at the receiver and a human operator adjusts both the pulse timing and frequency parameters. Used in combination with a static pulse repetition time (PRT) at the transmitter, such a system could have both the frequency and pulse timing synchronization manually retuned at periodic intervals in order to ensure continued synchronization. This system, while less restrictive than the first, still poses some non-trivial requirements on hardware design and also requires significant human intervention, as well as control over the transmitter for the purposes of conducting this manual tuning process.

Existing multistatic radars [22], [24]–[27] typically utilize the first technique which relies on GPS disciplined oscillators.

While the advantages offered by multistatic weather radar in terms of scan simultaneity and system cost are undeniable, the architecture also presents some significant challenges for the collection of scientifically useful data [22]. One of these challenges is low sensitivity compared to monostatic radar. This problem arises from the fact that multistatic weather radar receivers typically utilize low-directivity antennas in order to allow for signal reception over a wide area without the need for beam steering. In contrast, a monostatic weather radar will have the advantage of a highly directive beam on both transmit and receive. The consequence of this is that the area over which a multistatic receiver has adequate sensitivity for data collection is relatively small. Another problem of multistatic networks is spatial variability in both spatial resolution and Doppler resolution. For reasons that will be explained in detail in Chapter 2, resolution in both space and Doppler frequency is significantly degraded for each receiver near the baseline between itself and the transmitter. This places a further limitation on the area over which a given receiver can collect useful observations. These limitations are, however, not of particularly grave concern. They are offset significantly by the extreme cost advantages offered by a multistatic architecture. Receive modules are sufficiently inexpensive that both the sensitivity and resolution concerns can be significantly mitigated merely through the installation of additional receivers. Sensitivity concerns can be mitigated through the addition of receive modules at longer ranges from the transmitter, while resolution concerns can be mitigated by the strategic addition of receive modules within the coverage area so as to achieve redundant coverage with non-overlapping transmit/receive baselines.

A much more challenging and more frequently discussed problem is the issue

of sidelobe contamination. An additional effect of low-directivity receiver antennas is that the sidelobe levels will essentially be dictated by the one-way beam pattern of the transmitter. This can lead to extreme sidelobe contamination (particularly in areas of sharp reflectivity gradients) that can cause significant biases in velocity measurements. The complex geometries involved in this problem make compensation through an increased number of receivers impractical. Some research on this topic has focused on studying and quantifying the severity of the problem [28] in order to allow for censoring based on contamination levels. Alternatively, Chong et al. [29] has proposed a technique for variational correction of the measured velocities based on precise knowledge of the antenna patterns. While this technique appears to be reasonably effective, it would be preferable to remove the bias at its source rather than through an approximate correction (much like correcting for temporal variation in multiple Doppler observations). Kawamura et al. [30] suggest a method to mitigate this effect with a receive antenna that has closely spaced narrow grating lobes. This does reduce sidelobe effects, but greatly reduces the observable area of the receiver to discrete non-contiguous angles. A more promising avenue for mitigation of this issue is sidelobe whitening [28], [31] (which will be discussed later), but this technique has been little studied up to this point.

1.3 Motivation

The objective of this work is to advance the state of the art in multistatic weather radar through the development of signal processing techniques that allow receive modules to automatically synchronize with any in-band coherent transmitter using only the direct-path radiation and time-stamped pointing angles. This synchronization method has a number of advantages over existing techniques that have relied on some combination of expensive, highly stable oscillators and the ability to record

and/or communicate information about precise pulse timings and transmit phases. By removing restrictions on oscillator quality, it further reduces the cost associated with receiver modules. More importantly, it removes many prohibitive restrictions on what transmitter may be used. Most existing weather radar systems can serve as viable candidate transmitters without any modification. This allows both further cost savings, and the ability to utilize the receivers in conjunction with systems over which one has limited or no access. To demonstrate this capability, the prototype receivers developed over the course of the work detailed here utilize the WSR-88D as the transmitter of opportunity. This capability has a number of implications. It certainly offers an appealing avenue for multiple Doppler field research for institutions that may not have the considerable funding necessary to field multiple monostatic radars of their own. The interchangeability offered by the relaxation of requirements on transmitter hardware also makes this kind of system an intriguing choice for use in large field campaigns involving many monostatic systems. Receive units could theoretically switch seamlessly between transmitters involved in a given campaign based on need. The freedom from bulky oscillators and communications infrastructure also makes this type of receiver a good candidate for mounting on drones or for permanent/semi-permanent installation in elevated positions. Both of these strategies could be extremely useful for the collection of multiple Doppler data in rough or cluttered terrain (such as the southeastern United States). The practical realities of beam blockage and locating feasible deployment locations in this type of environment have historically made the collection of multiple Doppler data in these areas extremely challenging.

One exciting possible application for this type of flexible synchronization technique is to enable studies of the sidelobe whitening method for the mitigation of sidelobe contamination problem that has historically bedeviled multistatic weather

radar. This technique was first proposed by Sachidinanda and Zrnić [31] for use in monostatic systems and was subsequently identified as a promising avenue for mitigating the multistatic sidelobe problem by De Elia and Zawadski [28]. It involves perturbing the sidelobes of the transmitting antenna beam pattern on each pulse in order to decorrelate (and thereby mitigate) their contribution to Doppler velocity estimates. The challenge is, of course, that this kind of sidelobe perturbation is really only practical for a phased array weather radar. Suitable prototype systems for implementation of this type of technique are only now becoming available. Prominent examples of suitable phased array radars include the Advanced Technology Demonstrator (ATD) [32], Polarimetric Atmospheric Imaging Radar [33] (PAIR), and Horus [34]. If bistatic experiments demonstrating the efficacy of sidelobe whitening could be carried out using one or more of these systems, it would mark a potentially important development for multistatic weather radar. However, modification of the transmit electronics of these complex systems in order to make them suitable for use with more traditional multistatic weather radar hardware would be difficult or impossible. This makes the technology described in this dissertation an ideal tool for carrying out investigations into sidelobe whitening.

1.4 Outline

In Chapter 2, the fundamental theory of multistatic radars is outlined generally and then extended to the case of distributed scatterers in order to describe the operation of multistatic weather radars. The signal processing scheme developed to allow automated transmitter/receiver synchronization through sidelobe radiation is described in Chapter 3. The hardware that comprises the prototype receiver modules that have been constructed and deployed in the Oklahoma City, OK metropolitan area is also discussed in this chapter. Chapter 4 is dedicated to the presentation

and discussion of weather observations using this prototype system. First, a set of velocities measured by the prototype systems is checked for consistency against the radial velocities measured by KTLX, the nearest WSR-88D. Once this validation is presented, several additional cases of more significant meteorological interest are discussed. In Chapter 5, a new algorithm is described which makes the sidelobe whitening principle applicable to multistatic applications. Then a multistatic weather radar simulation framework is discussed and used to perform the first simulations of the possible benefits of sidelobe whitening for mitigation of sidelobe contamination in multistatic weather radar networks. Finally, in Chapter 6, the conclusions of this work are summarized and a number of possible avenues for further research are presented.

Chapter 2

Multistatic Weather Radar

Fundamentals

A bistatic radar can be defined as a radar system in which the distance between the transmitter and the receiver is on a similar order to the distance from a typical scatterer to the receiver. A multistatic radar system is an extension of this concept in which a group of at least two receivers with overlapping observation areas are deployed in conjunction with a single transmitter. Alternatively, a multistatic system can be implemented with a single receiver and multiple transmitters, but this is not common for weather radar applications. Typically, data from a multistatic radar network is processed by processing data from each transmitter / receiver pair, and then using a synthesis process to extract additional information beyond what is observable by a single bistatic radar (such as three-dimensional velocity) and/or to use the additional measurements to reduce error in estimates of scatterer properties. In this chapter, descriptions of the fundamental physical properties and mathematical




Figure 2.1: The fundamental geometry of a bistatic radar observation including the transmitter range R_T , receiver range R_R , baseline length L , and bistatic angle β . The arrows indicate the direction of radiation propagation along each path.

models of the multistatic radar observation process are provided.

2.1 Bistatic Radar Range Equation

In order to build an understanding of multistatic radar operation, it is useful to begin with a discussion of how an individual transmitter/receiver operates. Figure 2.1 shows a representation of bistatic radar geometry in the bistatic plane. It also introduces some key geometric variables that will be referenced frequently throughout this work. L is the baseline length, or the distance between the transmitter and

receiver. The transmitter range R_T is the distance between the transmitter and the scatterer under observation. Similarly, the receiver range R_R is the distance between the receiver and the scatterer. Finally, the bistatic angle β is defined as the angle between transmitter and receiver with a vertex located at the scatterer. The bistatic range, defined as $R_B = R_T + R_R$, is another fundamental geometric parameter which appears frequently in mathematical models of bistatic systems and signals.

In a perfect analog to monostatic radar systems [35], a bistatic radar range equation may be derived which relates the power emitted by the transceiver to that received by the receiver. The power density incident upon the scatterer is given by

$$S_s = \frac{P_T G_T f_T^2(\theta_T^s, \phi_T^s)}{4\pi R_T^2} \text{ (W/m}^2\text{)}, \quad (2.1)$$

Where P_T is the transmitted power, G_T is the transmit antenna gain in the relative direction of the scatterer, and $f_T^2(\theta_T^s, \phi_T^s)$ is the power pattern of the transmitter antenna evaluated at the angular position (θ_T^s, ϕ_T^s) of the scatterer with respect to the transmitter. The radiated power is then reflected from the scatterer toward the receiver. The total scattered power in the direction of the receiver is

$$P_s = S_s \sigma_{bi} \text{ W}, \quad (2.2)$$

Where σ_{bi} is the bistatic radar cross-section (RCS). This is a scatterer-specific property determined by the composition and geometry of the scatterer. For a given scatterer, this value will vary dependent on the relative angular positions of the transmitter and receiver, as well as the polarization of the incident radiation. The

power density at the receiver will then be

$$S_r = \frac{P_s}{4\pi R_R^2} \text{ (W/m}^2\text{)}. \quad (2.3)$$

Given an effective antenna area A_e , the power measured by the receiver from a scatterer located at its main beam peak is:

$$P_r = A_e S_r \text{ W}. \quad (2.4)$$

However, in the bistatic case, the relative position of the scatterer is unknown. This received power is therefore scaled by $f_R^2(\theta_R^s, \phi_R^s)$, the power pattern of the receiver antenna evaluated at the angular position (θ_R^s, ϕ_R^s) of the scatterer with respect to the receiver, yielding:

$$P_r = A_e S_r f_R^2(\theta_R^s, \phi_R^s) \text{ W}. \quad (2.5)$$

An antenna's effective aperture can be related to its gain G_R and operating wavelength λ by $A_e = \lambda^2 G_R / 4\pi$. Substituting that expression and equations (2.1)-(2.3) into (2.5), P_r can be expressed as:

$$P_r = \frac{P_T G_T G_R f_T^2(\theta_T^s, \phi_T^s) f_R^2(\theta_R^s, \phi_R^s) \lambda^2 \sigma_{bi}}{(4\pi)^3 R_T^2 R_R^2} \text{ W}. \quad (2.6)$$

2.2 Spatial Resolution

Because the focus of this work is on multistatic weather radar, it is desirable to generalize this radar range equation for point scatterers into a form appropriate for distributed scatterers, once again in a manner exactly analogous to the manner in which the same transformation is performed in the monostatic case. This is done by integrating the equation for point scatterers over some volume determined by




Figure 2.2: An example of the transmitter and receiver locations as the focii of an ellipse of constant range (shown in purple). This ellipse is a cross-section of the three-dimensional constant-range ellipsoid along the bistatic plane. For a scatterer located at any point on this ellipse, the bistatic range $R_B = R_T + R_R$ will be constant.

the transmitter and receiver antenna patterns and the range resolution corresponding to the transmitted waveform bandwidth [36]. However, due to the additional geometric complexity of a bistatic system, the bistatic weather radar equation does not lend itself to a simple closed form expression (even with such simplifications as Gaussian beam patterns). In order to understand why this is so, it is necessary to discuss the characteristics of bistatic radar resolution volumes. The range quantity directly measured by a bistatic receiver is R_B . This is obtained through the relationship $R_B = c\Delta t$ where Δt is the time delay between the instant a pulse is emitted by the transmitter and when the echo from the scatterer arrives at the receiver. The instant that the pulse is emitted is known through either dedicated communications infrastructure between the transmit and receive sites, or through monitoring of the direct path signal from from the transmitter. Using only the known value of R_B , it can be determined that the scatterer which produced the observed echo is located somewhere on an ellipsoid with its focii at the transmitter and receive locations and




Figure 2.3: The range and beamwidth boundaries that form a resolution volume in a typical bistatic weather radar. The concentric ellipses correspond to constant bistatic range ellipses. The translucent cones represent the main beam widths of the transmitter and receiver. The resolution volume formed by the intersection of the transmit beam with the ellipses is highlighted by a black dashed line.

a major axis of R_B . Figure 2.2 shows a cross-section of such an ellipsoid along the bistatic plane (the plane passing through the scatterer, transmitter, and receiver locations). The three-dimensional ellipsoid corresponds to a revolution of this ellipse about its major axis. Methods of determining the scatterer's precise location on that ellipsoid will be discussed in Section 2.3. However, precisely as in the monostatic case, a given waveform only offers finite range resolution. Generally, this resolution is

$$\Delta r = \frac{c}{2\beta_w} \text{ m}, \quad (2.7)$$

where β_w is the bandwidth of the transmitted pulse. For a single-frequency, unmodulated pulse, this resolution is

$$\Delta r = \frac{c\tau}{2} \text{ m}, \quad (2.8)$$

where τ is the temporal length of the pulse. This means that a resolution cell is bounded in range by two concentric ellipsoids representing two constant bistatic ranges with a difference of $c\tau/2$. One notable feature is that the “thickness” of the shell formed by region between the constant range ellipsoids (measured along the direction of the bistatic bisector) is not constant, but rather depends on bistatic angle. There is not a convenient analytic expression for this variation although an exact but implicit solution exists. That solution is derived in [35] along with the following useful approximation:

$$\Delta R = \frac{\Delta r}{\cos(\beta/2)} \text{ m}. \quad (2.9)$$

This varying resolution cell thickness can be observed in the constant-range ellipses shown in Figure 2.3.

The cross-range boundaries of a resolution cell are defined by the antenna patterns of the transmitter and receiver. Because a typical multistatic weather radar architecture (including the one described in this work) uses a highly directive transmit antenna and a receive antenna with a broad non-directive beam, the simplifying assumption that the cross-range boundaries are defined exclusively by the transmit antenna pattern is reasonable. An illustration of the geometry of such a resolution volume is shown in Figure 2.3. It is important to note again that this representation is merely a cross-section. In reality, the resolution volume is formed by the inter-

section of the conical transmit beam with the space between the two constant range ellipsoids determined by the range resolution associated with the transmitted waveform. At this point, it may become evident where the difficulty lies in converting the bistatic radar equation to a distributed scattering form. In order to perform this conversion, the radar cross section σ_{bi} is replaced with an integral of reflectivity η over the resolution volume weighted by the normalized transmit and receive antenna patterns $f_{\text{T}}^2(\theta_{\text{T}}, \phi_{\text{T}})$, $f_{\text{R}}^2(\theta_{\text{R}}, \phi_{\text{R}})$, as well as the range response $|W(r)|^2$ associated with the receiver transfer characteristic and transmitted waveform

$$\sigma_{\text{eff}} = \int_V f_{\text{T}}^2(\theta_{\text{T}}, \phi_{\text{T}}) f_{\text{R}}^2(\theta_{\text{R}}, \phi_{\text{R}}) |W(r)|^2 \eta dV \quad \text{m}^2, \quad (2.10)$$

where $\theta_{\text{T}}, \phi_{\text{T}}$ and $\theta_{\text{R}}, \phi_{\text{R}}$ are the angles associated with spherical coordinate systems centered at the transmitter and receiver, respectively. For the monostatic case it is simple to make some reasonable approximations and arrive at a closed form expression for this quantity in terms of η . This is because the resolution volume is simple to represent in terms of spherical coordinates with the origin located at the radar. Here, the antenna patterns are best represented by different spherical coordinate systems centered on the transmitter and receiver, while the range weighting function is best described by an elliptical coordinate system. The size and shape of the resolution volumes also changes dramatically depending on range and bistatic angle. This precludes the calculation of a convenient closed form solution for this integral in terms of η , so this quantity will be represented as $\sigma_{\text{eff}}(\eta, \theta_{\text{T}}, \phi_{\text{T}}, R_{\text{b}})$. This gives the bistatic weather radar equation the following form:

$$P_{\text{R}} = \frac{P_{\text{T}} G_{\text{T}} G_{\text{R}} \lambda^2 \sigma_{\text{eff}}(\eta, \theta_{\text{T}}, \phi_{\text{T}}, R_{\text{b}})}{(4\pi)^3 R_{\text{T}}^2 R_{\text{R}}^2} \quad \text{W}. \quad (2.11)$$

In the monostatic weather radar equation, the transmit portion of P_{R} dependency

on range is effectively canceled by the corresponding increase in resolution volume size (while incident power density at the scatterers decreases as r^2 , the volume size increases as r^2). The same general effect does, in fact, occur here; the solid angle of the isorange ellipsoidal shell intercepted by the transmitter beam is increasing with distance. However, the mathematical expression of this effect is not so clean as in the monostatic case. This is because the angle of incidence onto the shell is also dependent on bistatic range, which affects the size of the intercepted volume. However, the effect of this angular change approaches zero as the bistatic range goes to infinity. It is important to recognize that even though the R_T variable cannot be precisely cancelled, bistatic weather radars are not somehow exempt from the general effects of increasing volume size on distributed scatter.

In order to extract calibrated values of η , the meteorological quantity of interest, from received power measurements using this range equation, the most viable avenue would be numerical computation due to the complex resolution volume geometries. However, the remainder of this work focuses almost exclusively on Doppler measurements, so this will not be discussed in detail. However, this discussion of bistatic resolution volume variation and how it contributes to received power will be helpful in understanding some of the experimental measurements of received power shown in Chapters 3-4, as well as the simulation results in Chapter 5. A final aspect of this resolution volume variation that can be inferred from the discussion thus far, but deserves special mention, is the way that the resolution volume size degrades appreciably in regions near the baseline between the transmitter and receiver. This occurs due to loss of orthogonality between the beamwidth boundaries and the range boundaries. Along the baseline, this reaches its most extreme form. The entire baseline falls within a single resolution volume; this makes sense as there is no change in bistatic range regardless of where an object falls




Figure 2.4: Example of a degenerate resolution volume geometry in which the near bistatic range boundary is defined by the system baseline and the far bistatic range volume does not come into play, as the exterior volume edges are defined by the transmit and receive beam patterns.

along that line. An example of the kind of resulting degenerate geometry that can result from this effect is shown in Figure 2.4. In radar observations this will typically manifest as a roughly elliptical distorted “blob” caused by the presence of any scatterers in the region subject to severe resolution volume degradation.

2.3 Echo localization

The simplest method by which to localize a received echo in a multistatic radar system is to employ a highly directive antenna at either the transmitter or receiver, and combine knowledge of the pointing angle of that antenna with the measured bistatic range. The point at which the directive antenna’s beam intersects the ellipsoidal shell corresponding to the measured bistatic range is the target location. Assume a three-dimensional Cartesian coordinate system in which the scattering




Figure 2.5: Diagram of the geometric quantities used to calculate Cartesian scatterer coordinates in (2.12)-(2.14).

particle is located at (x, y, z) , a non-directive receiver is located at $(L, 0, 0)$, and a highly directive transmitter is located at the origin. In this scenario, the particle location can be expressed as [22]:

$$x = \frac{R_B^2 - L^2}{2[R_B - L \cos(p)]} \sin(a_T) \cos(e_T), \quad (2.12)$$

$$y = \frac{R_B^2 - L^2}{2[R_B - L \cos(p)]} \cos(a_T) \cos(e_T), \quad (2.13)$$

$$z = \frac{R_B^2 - L^2}{2[R_B - L \cos(p)]} \sin(e_T), \quad (2.14)$$

where a_T and e_T are the azimuth and elevation of the scatterer with respect to the transmitter and p is the angle between the transmitter pointing direction and the baseline between the transmitter and receiver. These quantities are illustrated in Figure 2.5. The drawback of this technique is that processing cannot be performed without timestamped pointing angles from the transmitter. Using the WSR-88D as a transmitter of opportunity eliminates the possibility of real-time processing; it must wait until the WSR-88D data are publicly released. There does exist a technique known as multilateration [35] which is used to localize scatterers observed by a multistatic network without any knowledge of the radiation pattern or pointing angle of the transmitter. This technique, however, is not viable for use in the observation of distributed scatterers.

To demonstrate why, first consider Figure 2.6, which depicts the observation of a single scatterer by a multistatic network. Each receiver is capable of measuring a constant-range ellipse corresponding to the scatterer location. The places where these ellipses intersect represent candidate target locations. However, each pair of ellipses intersects twice. One of these locations corresponds to the actual scatterer, whereas the other is a spurious “ghost” target location. This ambiguity is resolved through the use of three simultaneous ellipses. The only location at which all three




Figure 2.6: Echo localization through multilateration in a network consisting of a single transmitter and three receivers. The purple ellipses represent isorange contours corresponding to detection of the actual scatterer by each receiver. If only two receivers are used, there will be two candidate detection locations where the isorange contours intersect: one at the location of the actual scatterer and a “ghost” represented by a white triangle. The third receiver is necessary for disambiguation.




Figure 2.7: Complications of multilateration when multiple targets are present. Ghost targets are represented by white triangles. The especially problematic ghost target at which three ellipses intersect is indicated in red.

ellipses intersect is the true scatterer location. An alternative technique to resolve these ambiguities is by checking for consistency in the Doppler information measured by each transmitter / receiver pair. This type of technique can work well in the observation of a small number of discrete targets. However, there arises a problem of target association that renders this method increasingly impractical as the number of target-containing resolution volumes increases. A common weather scenario in which distributed scatterers filling an area of hundreds or thousands of contiguous resolution volumes represents an extreme case of this problem, which is why it is necessary to rely on information about transmitter pointing angle for the purposes of the system described in this work. Figure 2.7 depicts an illustration of this association problem in the simplest possible two-scatterer case.

Additional ghost targets are created by intersections between the isorange con-

tours corresponding to different scatterers. More problematically, it is even possible to have locations where three contours can intersect to create a ghost, making it significantly harder to distinguish from an actual scatterer. This issue is exacerbated by the fact that real systems have only finite range resolution, as discussed in Section 2.2, which means that the ellipses do not need to perfectly intersect to introduce ambiguity. The number of these ambiguities increases combinatorially with an increase in the number of scatterers present. Clearly, the presence of scatterers in tens of thousands of adjacent resolution volumes, as one might expect to see in weather radar applications, presents a completely intractable problem. Thus, we must rely on pointing angle knowledge for the localization of echoes from distributed scatterers.

2.4 Velocity Estimation

One feature that distinguishes bistatic from monostatic systems is the geometry involved with the measurement of Doppler velocities. It is well known that the Doppler velocity measured by monostatic systems is a radial velocity, or the projection of the scatterer velocity vector onto the vector which represents the pointing angle of the radar system. Rather than measuring velocity along a radial corresponding to either the transmitter or receiver, bistatic systems measure velocity along the bistatic bisector. This is the vector bisecting the angle formed by the transmit and receive lines of sight to the scatterer. The Doppler frequency shift f_d measured by a bistatic system is, as in the monostatic case, the first time derivative of the path length traveled by a pulse divided by its wavelength (assuming that acceleration over the observation interval is negligible). However, the separate transmit and receive paths must be accounted for in developing an expression for




Figure 2.8: Relevant geometric quantities for modeling Doppler velocity measurements using a bistatic radar.

that quantity:

$$f_d = \frac{1}{\lambda} \left[\frac{d}{dt} (R_R + R_T) \right] \text{ Hz} \quad (2.15)$$

$$= \frac{1}{\lambda} \left[\frac{dR_R}{dt} + \frac{dR_T}{dt} \right] \text{ Hz}, \quad (2.16)$$

where λ is the operating wavelength of the transmitting radar. Consider then the geometry defined by Figure 2.8. Here, v_s represents the total velocity of the scatterer, v_{bi} is the projection of that velocity onto the bistatic bisector, and δ is the angle between the scatterer's direction of motion and the bistatic bisector. The time derivatives of R_R and R_T can be calculated by projecting v_s onto the transmitter and

receiver lines of sight [35]:

$$\frac{dR_{\text{R}}}{dt} = v_s \cos\left(\delta + \frac{\beta}{2}\right) \text{ ms}^{-1}, \quad (2.17)$$

$$\frac{dR_{\text{T}}}{dt} = v_s \cos\left(\delta - \frac{\beta}{2}\right) \text{ ms}^{-1}. \quad (2.18)$$

Substituting these quantities into (2.16) yields

$$f_d = \frac{v_s}{\lambda} \left[\cos\left(\delta - \frac{\beta}{2}\right) + \cos\left(\delta + \frac{\beta}{2}\right) \right] \text{ Hz}, \quad (2.19)$$

$$= \frac{2v_s}{\lambda} \cos(\delta) \cos\left(\frac{\beta}{2}\right) \text{ Hz}. \quad (2.20)$$


To describe what this equation represents conceptually, it is useful to consider it in three pieces. $v_s \cos(\delta)$ is the projection of the scatterer velocity onto the bistatic bisector. $2/\lambda$ accounts for the two-way path length change and converts the velocity in m/s to a frequency in rad/s through the division by λ . Finally, the $\cos(\beta/2)$ is a scaling factor to account for spatial variation in the amount of path change produced by a given movement along the bistatic bisector. The necessity of this term is best illustrated through a consideration of the limiting cases. For a bistatic angle of 0° in which the scatterer is colinear with the transmitter and receiver (but does not lie between them), an infinitesimal scatterer movement by a distance dr along the bistatic bisector results in a path length change of $2dr$. However, given a bistatic angle of 180° in which the scatterer lies on the baseline between transmitter and receiver, the bistatic bisector is orthogonal to the transmit and receive lines of sight. This means that a movement of dr along the bisector will produce no path length change whatsoever. This fact can be generally summarized by stating that the larger the bistatic angle becomes, the smaller the frequency shift induced

by a given v_{bi} . By extension, the larger the bistatic angle becomes, the larger the unambiguous Doppler velocity becomes, as the unambiguous Doppler frequency remains constant. However, because the frequency resolution of the receiver also stays constant, this means that the velocity resolution of the system degrades for large bistatic angles. This will occur in the region near the baseline, which is also the area subject to severe loss of spatial resolution due to loss of orthogonality between the isorange contours and transmitted beam.


2.5 Polarimetric Considerations

Another aspect of multistatic weather radar that must be considered is how polarization affects scattering behavior. While there has been some interest in bistatic weather radar polarimetry [37] for the purpose of discriminating and characterizing hail, the system described herein is single-polarization and designed specifically for wind field estimation. However, this aspect of bistatic weather radar is still important to the design of the system as it is desirable to design the polarization of the receive module antenna so as to optimize system performance. It has been shown [38] that the scattering behavior of liquid precipitation can be accurately modeled by assuming that each raindrop behaves as an electrically small dipole with a moment vector oriented along the polarization direction of the impinging radiation. The consequences of this behavior are illustrated in Figure 2.9.

If the radiation is polarized along the z-axis, the RCS will be proportional to $\sin^2(\theta)$ in a spherical coordinate system. In other words, the raindrop will have nulls in its bistatic RCS on the axis corresponding to the polarization of the impinging wave, and it will be omnidirectional about that axis. For the shallow elevation angles at which most weather observations are carried out, this means that H-polarized radiation will produce nulls in the scattered wavefront at 90° in azimuth



(a)



(b)

Figure 2.9: Three-dimensional bistatic RCS of raindrop with impinging radiation polarized along (a) the z-axis (V polarization) and (b) the y-axis (H polarization).

relative to the transmitter direction, as illustrated in Figure 2.9b. Any receivers located near those null angles will have degraded signal-to-noise ratio (SNR) due to the weak scattering. This is undesirable, as such geometries produce excellent spatial and good Doppler resolutions, so it is not desirable to create blindness at those angles. By contrast, V-polarized radiation creates nulls directly above and below the raindrop. For most applications this is much more desirable, as scattering from the raindrop can be observed well by a receiver at any azimuthal position.

2.6 Wind Field Estimation

Once measurements of v_{bi} have been obtained, additional processing must be performed to obtain an estimate of the observed wind fields in Cartesian coordinates. The basic process is to utilize some objective analysis method to interpolate the observations obtained by each radar system to a common grid and then to solve a system of equations to obtain each component of the wind vector at every point. While sophisticated techniques for wind field estimation such as those discussed in Chapter 1.1 can be applied to multistatic observations as readily as they may be applied to monostatic multiple Doppler data, such methods are beyond the scope of this work, which uses essentially the same simple techniques outlined in [22].

Objective analyses are used to take data obtained by multiple instruments on differing and irregularly spaced grids and estimate the state of those variables at common, regularly spaced points. This is a key stage in the pre-processing of data not only for wind field estimates such as those of interest here, but also for numerical weather prediction and corresponding data assimilation processes. While modern numerical weather prediction techniques typically use complex methods based in optimum control theory [39] to perform these estimates, several of the earliest and simplest interpolation-based techniques, such as the Barnes [40] and Cressman




Figure 2.10: One-dimensional Cressman weighting function for $R_i = 1$.

[41] techniques, are still widely used for multiple-Doppler analyses. The processing scheme used to obtain wind fields from multistatic observations collected for this work utilizes the Cressman technique. In this method, the value interpolated to each desired grid point is a weighted average of the available observations within some radius of influence R_i . The Cressman weighting function is

$$W(d) = \frac{R_i^2 - d^2}{R_i^2 + d^2}, \quad (2.21)$$

where d is the distance between the desired grid point and the observation being weighted. The unit-radius Cressman weighting function in one dimension is shown in Figure 2.10.

Once the Cressman method or another objective analysis scheme has been used to interpolate the observations from each receiver to a common grid, the next step is to solve for the Cartesian velocity components at each point. For a monostatic radar system, the measured radial velocity v_r can be expressed in terms of the zonal velocity component u , the meridional component v , and the vertical component w

as follows [22]:

$$v_r = u \sin(a_T) \cos(e_T) + v \cos(a_T) \cos(e_T) + w_p \sin(e_T) \text{ ms}^{-1}. \quad (2.22)$$

A similar relationship exists for a bistatic velocity:

$$v_{bi} = \frac{\sin(a_R) \cos(e_R) + \sin(a_T) \cos(e_T)}{2 \cos(\beta/2)} u + \frac{\cos(a_R) \cos(e_R) + \cos(a_T) \cos(e_T)}{2 \cos(\beta/2)} v + \frac{\sin(e_R) + \sin(e_T)}{2 \cos(\beta/2)} w \text{ ms}^{-1}. \quad (2.23)$$

It is clear that in either the monostatic or multistatic case (or a combination), at least three independent velocity measurements are necessary to solve for the three components of the wind field. For a multistatic network consisting of a single transmitter (with accompanying receiver allowing it to function as an independent monostatic radar) and multiple receivers, the corresponding system of equations can be represented in matrix format as:

$$\begin{bmatrix} \frac{\sin(a_1) \cos(e_1) + \sin(a_T) \cos(e_T)}{2 \cos(\beta_1/2)} & \frac{\cos(a_1) \cos(e_1) + \cos(a_T) \cos(e_T)}{2 \cos(\beta_1/2)} & \frac{\sin(e_1) + \sin(e_T)}{2 \cos(\beta_1/2)} \\ \frac{\sin(a_2) \cos(e_2) + \sin(a_T) \cos(e_T)}{2 \cos(\beta_2/2)} & \frac{\cos(a_2) \cos(e_2) + \cos(a_T) \cos(e_T)}{2 \cos(\beta_2/2)} & \frac{\sin(e_2) + \sin(e_T)}{2 \cos(\beta_2/2)} \\ \vdots & \vdots & \vdots \\ \frac{\sin(a_N) \cos(e_N) + \sin(a_T) \cos(e_T)}{2 \cos(\beta_N/2)} & \frac{\cos(a_N) \cos(e_N) + \cos(a_T) \cos(e_T)}{2 \cos(\beta_N/2)} & \frac{\sin(e_N) + \sin(e_T)}{2 \cos(\beta_N/2)} \\ \sin(a_T) \cos(e_T) & \cos(a_T) \cos(e_T) & w_p \sin(e_T) \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} v_{bi1} \\ v_{bi2} \\ \vdots \\ v_{biN} \\ v_r \end{bmatrix}, \quad (2.24)$$

where a_n and e_n are the azimuth and elevation relative to the n th receiver, v_{bin} is the bistatic velocity measured by each of the receivers, β_n is the bistatic angle corresponding to the n th receiver, and N is the total number of receivers in the network. Note that Equation 3 in [22] conflicts with this result, erroneously omitting the β_n dependency. A proof of the equation presented here is given in Appendix A. This system of equations can be solved through direct matrix inversion for the two-receiver / one-transmitter or three-receiver cases; however, the overdetermined system in which there are more than three available observations can also be solved in the least-squares sense using the Moore-Penrose pseudoinverse. Wurman et al. [22] showed that additional observations beyond the minimum of three help to reduce error levels in the resulting wind field estimates. For typical network geometries in which the receivers are distributed near ground level and have their antennas oriented at shallow elevation angles, this error level reduction not only improves data quality generally, but extends to lower altitudes the height at which vertical velocity estimates will meet a given quality metric. As in the monostatic case, the vertical velocity estimates are typically lower quality than the horizontal velocities, as the shallow observation angles make them more sensitive to error. This is due to the fact that at low altitudes, the vertical motion is nearly orthogonal to the bistatic bisector at most azimuth angles (those not near to the system baseline).

The fundamental aspects of multistatic radar operation and performance laid out in this chapter form the foundation for the remainder of this work. They are necessary to understand and interpret the observations, both experimental and simulated, presented to demonstrate the system designs and signal processing techniques discussed here. Furthermore, they add context in the sense that they provide an understanding of the limitations and challenges that still remain for multistatic weather radar data quality, even provided that such problems as sidelobe contami-

nation can be significantly mitigated. Having provided this groundwork, it is now possible to discuss a prototype multistatic weather radar network that has been designed and deployed to demonstrate the feasibility of passive transmitter / receiver synchronization.

Chapter 3

Design of a Multistatic Weather Radar Receiver

The drawbacks in cost and logistics associated with the use of a monostatic radar network for multiple-Doppler data collection can be mitigated through the adoption of a multistatic network architecture. However, these systems must address the key challenge in the engineering of any multistatic radar: precise synchronization in both carrier frequency and pulse timing between the transmitter and the distributed receivers. Frequency synchronization is critical for accurate Doppler velocity estimation, while the pulse timing synchronization is needed to accurately localize received echoes. Existing systems [24], [27], [28], [42], [43] use a combination of GPS-disciplined oscillators and communications infrastructure linking the transmitter and receivers to achieve this synchronization. While this is a simple and effective way to achieve precise synchronization, it has disadvantages in the form of expense and inflexibility. Byrd et al. [45], [46] proposed a passive radar module constructed



Figure 3.1: This satellite image [44] shows the locations of the two passive receivers at the Radar Innovations Lab (RIL) and University of Oklahoma Health Science center (OUHSC), as well as the location of KTLX, the WSR-88D being utilized as a transmitter of opportunity and also operating as an independent monostatic radar.

using inexpensive components that achieves synchronization through monitoring of the direct-path signal between the transmitting radar and the receiver. This type of architecture significantly reduces the cost of a multistatic network by obviating the need for a highly accurate oscillator and allowing the use of any in-band transmitter as a radiator without modification or the establishment of any communications infrastructure. The only cooperation necessary from the transmitting radar is the availability of time-stamped pointing angle information. This feature is essentially necessary even for monostatic weather radars. In addition to its favorable impact on system cost, this also allows for an exceptional degree of flexibility. Receive modules of this type can be installed virtually anywhere and could be easily redeployed to new locations or used with new transmitters. A two-receiver network comprised of the modules described by [46] has been constructed by the Advanced Radar Research Center (ARRC) at the University of Oklahoma and deployed in the southern Oklahoma City, OK metropolitan area at the locations shown in Figure 3.1. These receivers utilize KTLX, a nearby WSR-88D, as a transmitter of opportunity. The operating parameters for KTLX are summarized in Table 3.1.

In order to obtain high-quality measurements using synchronization through sidelobe radiation there were three signal processing challenges that had to be overcome:

- Accurate frequency estimation of a small duty cycle signal in a strong multi-

Table 3.1: KTLX operating parameters.

Carrier Frequency	2.91 GHz
Peak Transmit Power	700 kW
Pulse Widths	1.57 and 4.71 μ s
Bandwidth	637 and 212 kHz
Beamwidth	$\approx 1^\circ$
Peak Sidelobe Level	-29 dB

path environment

- Robust quality control of estimated direct-path pulse locations in the presence of both clutter and weather
- Adaptation to special transmit schemes used by the WSR-88D, specifically Sachidinanda/Zrnić (SZ-2) phase coding and batch pulse repetition frequencies (PRFs)

Section 3.1 reviews the signal processing methods that resolve these issues. Section 3.7 describes the hardware design of the receiver module. Section 4.1 discusses the radar network and presents a validation of its velocity measurements by evaluating the consistency of a retrieved two-dimensional wind field with the radial velocity measurements of the WSR-88D.

3.1 Passive Radar Signal Processing

The process of synchronizing each receiver with the transmitting radar takes place at three levels. First, there is the coarse synchronization needed to correctly calculate the transmitter pointing angle associated with any detected echoes. This is achieved by combining time-stamped pointing angle information from the transmitter with measurements of direct-path pulse power collected by the receiver. The next level of synchronization is pulse timing. The timing of each transmitted pulse's arrival at the receiver must be known to separate the data into intervals corresponding to a single PRT and correctly determine the bistatic range associated with each detected echo. Finally, carrier frequency synchronization must be achieved. This is necessary to accurately calculate the Doppler velocities associated with observed scatterers. The synchronization process is complicated by several special transmit modes utilized by WSR-88D, the SZ-2 phase coding and batch PRF modes. The




Figure 3.2: Flowchart depicting an overview of the signal processing scheme used to translate the passive radar time-series into wind field estimates. CFAR detection is followed by a quality control process based on PRF-fitting. The frequency offset between transmitter and receiver is removed, and then estimates of power and velocity are calculated and localized to Cartesian coordinates. Finally the estimates are interpolated to a common grid and wind-field estimation is performed. Special WSR-88D transmit modes (batch PRF and SZ-2 coding) are detected and handled appropriately.

basic workings of this synchronization process were outlined in [46], but they will be discussed in more exhaustive detail here. Once the synchronization process is complete, bistatic velocity estimates are calculated and the data are localized to Cartesian coordinates. Finally, the data are censored based on SNR and bistatic angle and then interpolated to a uniform grid using Cressman analysis [41], at which point it may be utilized for wind-field retrieval. Figure 3.2 shows a basic outline of the complete signal processing scheme. It should be noted that the existing radar network carries out this signal processing scheme offline, due to lack of availability of real-time timestamped pointing angles from the WSR-88D, as well as the fact that the software currently exists as a relatively slow MATLAB prototype.




Figure 3.3: The approximate matched filter used to process the received time series data. It was derived from experimental measurements of WSR-88D pulses. The half-power width of the approximate filter is consistent with the $1.57\mu s$ nominal pulse length. The index n assumes a sampling rate of 5 MSPS.

3.2 Coarse Time Alignment and CFAR

The first step in the signal processing chain is a coarse synchronization of the timestamped pointing angles provided by the transmitting radar with the timestamped beam pattern measured by the receiver using the direct-path signal. The system time available to the internet-connected receivers is typically within about a second of the time indicated by the WSR-88D timestamps. However, that fraction of a second of potential inaccuracy can result in significant errors in localizing echoes and complicates the quality control process used for pulse-timing synchronization. In order to mitigate this source of error an initial constant false alarm rate (CFAR) detection is carried out on the time-series data collected by the receiver. Prior to CFAR detection the data are passed through an approximate matched filter shown




Figure 3.4: In this segment of data collected from a transmitted WSR-88D signal, the rising edge of a direct-path pulse is located near sample 400. Note the drastically higher power levels due to ground clutter and weather after the pulse in contrast to the samples collected prior. This severity of the issue is more striking given the context that the pulse itself is only about 8 samples long ($1.57 \mu\text{s}$ at a 5 MSPS).

Parameter	Value
Reference Samples	500
Guard Samples	10
Nominal PFA	1e-6
Minimum Pulse Samples	9
Window Type	Lagging
Algorithm	Cell-Averaged

Table 3.2: CFAR Parameters

in Figure 3.3.

The CFAR detector (see Table 3.2 for key parameters) uses a lagging window to favor the detection direct-path pulses over strong close-in ground clutter that will likely follow it in the received time-series, as illustrated in Figure 3.4. In short, this is because any echoes originating from the transmitted pulse must arrive *after* the direct path signal as they have a longer propagation path. Therefore, the samples collected immediately before the arrival of a direct path pulse correspond to distant bistatic ranges from the preceding pulse. At those distances, the main beam /

near-sidelobe heights will be large, and therefore ground clutter influence will be minimal, and significant weather contributions are also less likely. Combined with the reduced sensitivity of the multistatic receiver due to its low-gain antenna, this means that these samples will typically represent near-pristine thermal/atmospheric noise. In contrast, the samples immediately following the pulse represent low beam heights at which ground clutter signals will be relatively strong. This will strongly bias any estimate to determine a noise level using these samples. Therefore, to avoid missed direct-path pulse detections, it is strongly desirable to use only the samples preceding the sample-under-test to estimate noise. As the data collected are significantly oversampled for the purposes of the frequency estimation process, each direct-path pulse (as well as clutter and weather) will result in chains of detections. The next processing step is to sift through the detector output and group each chain of consecutive detections into a “pulse” with an associated time and peak power level. A minimum length threshold is set to add a layer of protection against false detections. Plotting the peak power vs. pulse time results in a series of consecutive approximate beam pattern cuts taken across varying elevations, as shown in Figure 3.5. These patterns have not yet been quality controlled, so they contain some missed pulses and false detection. The missed pulses are generally due to insufficient received energy for detection due to nulls in the transmitter antenna pattern, while “false” detections are due to ground clutter or even weather. Additionally, these pattern measurements are subject to saturation at the highest power levels, which is why no main beam is clearly visible, even at the lowest elevation cuts. Note that this saturation will destroy the phase information needed for frequency offset removal, but this happens only near the transmitter / receiver baseline, where any recorded data would be of dubious value anyway due to the degraded spatial and Doppler resolution associated with that region. The saturation and ground clutter-




Figure 3.5: Measurements of direct-path pulse power over time. These pulses form a repeating measurement of the WSR-88D beam pattern at changing elevation cuts based on the mechanical tilt of the transmitter.




Figure 3.6: This 0.44° elevation beampattern cut is taken from the same dataset used to produce Figure 3.5. Note the saturation of the main beam as well as the apparent irregularities in the sidelobe topography caused by strong ground clutter.

ter effects are clearly visible in the detailed beampattern cut shown in Figure 3.6. Despite these flaws, this raw CFAR output is, however, sufficiently accurate for use in the coarse synchronization step.

Each repetition of the measured beam pattern represents one azimuthal rotation of the WSR-88D. The time-stamped pointing angles provided by the WSR-88D are used to construct an approximation of the beam pattern that should be measured at the receiver over the observation period. The cross-correlation between this theoretical pattern and the realized pattern is then used to estimate the delay between the timestamps of the transmitter and the receiver. This delay is then corrected.




Figure 3.7: High-level outline of the PRF fitting process used to quality control the pulse locations estimated by the CFAR detector.

3.3 Standard PRF Fitting

As mentioned above, the output of the CFAR detector will contain false direct-path pulses due to clutter and weather, as well as missed pulses due to nulls in the transmitting radar beam pattern. In order to produce final products of acceptable quality, it is essential to correct these errors. This is achieved through a process (outlined in Figure 3.7) by leveraging some minimal knowledge about the operation of the WSR-88D. The only real assumption needed is that the WSR-88D transmits using either 1 or 2 PRFs at each elevation angle. It takes some additional knowledge to deal with the two-PRF case, which will be discussed later; the single-PRF case will be addressed presently. The basic guiding principle is to seek the sequence of regularly spaced pulse locations that best fits the observed data. Once this is done, detections that do not fit the sequence can be discarded as false, while locations in the sequence that are not associated with nearby detections can be identified as misses. The first step in this process is to make some initial guess at the number

of receiver samples M' in each transmitter PRT. This is done by looking at the observed distribution of delays between successive pulses and taking the peak of that distribution (if there are 2 peaks, the data must be processed using the multiple PRF method). This assumes that a significant proportion of the PRT intervals will contain no false detections. This could potentially be violated by a sufficiently cluttered environment, but such an instance has not yet been identified within the several hours of weather observations collected by the network. Error in M' can be measured by the function $\delta[n] = t[n] \bmod M'$, where $t[n]$ is the vector of detected pulse timestamps and n is the pulse index. A perfect estimate of M' would result in $\delta[n]$ manifesting as a constant value (with some noise due to false detections and misses), representing the scalar delay between the actual pulse train with period M' and a pulse train with period M' that has its first pulse located at $n = 0$. An error in M' will give $\delta[n]$ the form of a noisy sawtooth wave with the following equation:

$$\delta[n] = fn - M' \left\lfloor \frac{fn}{M'} \right\rfloor + v[n], \quad (3.1)$$

$$f = \frac{M}{M'} - 1, \quad (3.2)$$

where M is the actual number of samples-per-PRT, and $v[n]$ is a function representing the errors introduced by misses and false detections. An actual measured example of such a function is shown in Figure 3.8. It is evident through inspection that the value of M could be obtained if the slope f of this sawtooth wave could be accurately measured. However, this is rendered difficult by both the noise component $v[n]$ and the wrapping behavior. The task is rendered much easier through conversion to a problem of frequency estimation. This is achieved by scaling $\delta[n]$ to an amplitude of 2π , yielding a phase function $\phi[n]$. This, in turn, can be used to generate a complex exponential with frequency $2\pi f/M'$ and some amount of phase




Figure 3.8: Example of the “noisy sawtooth” form exhibited by the error function $\delta[n]$, as calculated from actual recorded data. Also shown is the error function recalculated after systematic error in samples-per-PRT has been corrected using the algorithm described in Section 3.1.

noise:

$$\phi[n] = \frac{2\pi}{M'}\delta[n], \quad (3.3)$$

$$\begin{aligned} g[n] &= \exp(j\phi[n]) \\ &= \exp\left(j\frac{2\pi}{M'}(fn + v[n])\right). \end{aligned} \quad (3.4)$$

The function $g[n]$ can then be resampled to a uniform interval so that its fast Fourier transform (FFT) may be calculated. The peak frequency of the resulting spectrum corresponds to the phase function slope $2\pi f/M'$, from which the value of M can be extracted. This is used to then recalculate a set of theoretical pulse locations with the first pulse located at $n = 0$. An offset is then applied to minimize the distance between each theoretical pulse location and the nearest detected pulse. This yields the final set of estimated pulse locations. False detections may now be discarded and missed pulses flagged. This information is then used to break the recorded stream of time-series data into a matrix in the familiar fast time / slow time format in preparation for the next step of frequency synchronization. An attentive reader may note that the slope of the sawtooth wave could also be calculated by directly taking its Fourier transform. This would, however, yield a spectrum symmetric about zero with peaks at both the positive and negative frequencies corresponding to the slope. Determining the sign of the slope would require disentangling the effects of time delay on the phase of the obtained spectrum. It is simpler (and adequate) to convert the wave to a complex sinusoid.

3.4 Batch PRF Fitting

Batch PRF mode is utilized to gain both long-range surveillance data and close-range unambiguous Doppler data within the same scan. Within each dwell, the WSR-88D transmits a fixed number of slow PRF pulses followed by a larger, flexible number of fast-PRF pulses. The number of fast pulses is flexible due to the fact that the rotation rate of the WSR-88D's pedestal is not perfectly constant. Therefore, to keep the angle traversed by the radar within each dwell constant, the WSR-88D will add or remove pulses from the dwell as necessary according to the realized rotation rate measured by its angular position sensor. This adds a unique challenge to the problem of quality control through PRF estimation. This section will detail two different solutions to this challenge, the version used for initial system validation using stratiform precipitation observations, and a more robust but computationally intensive method developed later.

3.4.1 Original Method

The first step in overcoming this challenge is to precisely align the receiver and transmitter timing beyond what is achievable using only timestamped pointing angles and the measured radiation pattern of the transmitter. This is done by identifying the instants at which the received pulses transition from the slow to the fast PRF. This metric is then converted to dwell length by taking the first difference of the time samples in this vector. This measure of dwell length over time is then resampled to a uniform time axis. This process is then repeated using the time stamps provided by the WSR-88D, which are used to estimate dwell lengths and then resampled to the same time axis as that used to process the receiver data. Thus, we have a mea-

surement of the dwell lengths detected by the receiver $x_d(t)$, and those transmitted by the WSR-88D $x_{\text{ref}}(t)$, measured on a common time axis. A fine time alignment can then be achieved by calculating the cross-correlation function of these two measurements and using the peak of the resulting signal to identify the delay necessary to align the signals. However, there is an issue with this approach; gaps in the received data caused by missed pulses can cause large measured dwell lengths that severely bias the estimated cross-correlation. Therefore, rather than utilize a traditional cross-correlation, the weighted normalized cross-correlation (WNCC) method is used [47]. This allows the contribution of outliers in dwell length estimate to be censored. A weight vector $w(t)$ is assigned a value of 0 for all dwell lengths above some threshold and 1 for all dwell lengths below the threshold. The WNCC $C(\tau)$ is calculated as follows:

$$C(\tau) = \frac{[w(t) \cdot \hat{x}_d(t)] * \hat{x}_{\text{ref}}(t)}{\sqrt{w(t) * \hat{x}_{\text{ref}}^2(t) [w(t) \cdot \hat{x}_d^2(t)]}}, \quad (3.5)$$

$$\hat{x}_d(t) = x_d(t) - \bar{x}_d, \quad (3.6)$$

$$\hat{x}_{\text{ref}}(t) = x_{\text{ref}}(t) - \bar{x}_{\text{ref}}, \quad (3.7)$$

where $*$ is the convolution operator and \bar{x}_d and \bar{x}_{ref} are the sample means of $x_d(t)$ and $x_{\text{ref}}(t)$, respectively. The offset between the receiver time axis and that recorded by the WSR-88D is estimated by finding the lag corresponding to the peak of $C(\tau)$. This offset is then removed such that the axes are aligned. Once this process is complete, a theoretical pulse train $T_{\text{ref}}[n]$ is constructed based on the WSR-88D timestamps. First, the number of slow pulses in each dwell S (which does not change over the observation period) is estimated based on the peaks of the observed distribution of delays between pulses. Starting with the first timestamp, S

theoretical pulse timestamps are added to $T_{\text{ref}}[n]$ at intervals corresponding to the longer of the two PRTs detected in the observed pulses. At this juncture, pulses are added at an interval corresponding to the shorter of the two PRTs are added until the next WSR-88D timestamp is exceeded. Then slow pulses are again added and the process repeats. This continues until the pulses corresponding to all available WSR-88D timestamps have been incorporated into $T_{\text{ref}}[n]$. Because we are interested in the collected data primarily for Doppler analysis, the long PRT pulses are of little use, and are therefore discarded. Due to limitations on the precision of the WSR-88D timestamps, the first and last short PRT pulses within each dwell are unreliable and are also discarded. The remaining short PRT pulses are then used to calculate an error vector $\delta_{\text{batch}}[n]$ as follows:

$$\delta_{\text{batch}}[n] = T_{\text{d}}[n] - \mathcal{N}_{\text{ref}}\{T_{\text{d}}[n]\}, \quad (3.8)$$

where $T_{\text{d}}[n]$ is the vector containing timestamps of all detected pulses and the operator \mathcal{N}_{ref} denotes finding the nearest neighbor for all input values within the vector T_{ref} . This will produce a sawtooth wave structure similar to that produced in the single PRT case, but in the time domain rather than the sample index domain. This sawtooth wave will have the following form:

$$\delta_{\text{batch}}[n] = fT_{\text{d}}[n] - T_{\text{s}}' \left\lfloor \frac{fn}{T_{\text{s}}} \right\rfloor + v[n], \quad (3.9)$$

$$f = \frac{T_{\text{s}}}{T_{\text{s}}'} - 1, \quad (3.10)$$




Figure 3.9: Rexample of measured pulse-to-pulse phase rotation distributions after decoding using each of the 8 possible delays. A delay of 2 samples is clearly the correct option, as shown by the low variance of the resulting phases.

where T_s is the actual fast PRT and T'_s is the current estimate of the fast PRT. This signal is converted into a noisy complex sinusoid as follows:

$$\phi[n] = \frac{2\pi}{T'_s} \delta[n], \quad (3.11)$$

$$\begin{aligned} g[n] &= \exp(j\phi[n]) \\ &= \exp\left(j\frac{2\pi}{T'_s}(g[n] + v[n])\right), \end{aligned} \quad (3.12)$$

The estimated true value of T_s is then retrieved using Fourier analysis and the received sample indices are corrected accordingly.

3.4.2 Improved Method for Precise Time Alignment

Although the precise time alignment technique using the WNCC method is adequate in many cases, it can break down when there are particularly large numbers of false detections due to weather and/or clutter. An elevated number of false detections can lead to failures in identifying the intervals when the transmitter transitions from fast to slow PRF. If enough of these transitions are missed, the WNCC technique will have an insufficient amount of accurately measured dwell lengths available, and will produce inaccurate results.

In order to rectify this problem, a more computationally expensive but robust technique can be used. In this method, the theoretical timestamp vector $T_{\text{ref}}[n]$ is constructed prior to any time alignment. The proper time delay Δt with which to correct $T_{\text{ref}}[n]$ is then estimated through the following optimization problem:

$$\min_{\Delta t} \sum_{n=1}^N [(T_{\text{ref}}[n] - \Delta t) - \mathcal{N}_d\{T_{\text{ref}}[n] - \Delta t\}]^2. \quad (3.13)$$

In other words, a time delay Δt is sought which minimizes the sum of the squared

distances between each theoretical timestamp and its nearest detected counterpart. It should be noted that as error in the initial PRF estimates has not yet been removed, the total error will most likely not closely approach zero. Once this optimal time offset has been determined, the process of batch PRF processing may proceed as previously specified from (3.8).

3.5 Phase Decoding and Frequency Offset Removal

The other special WSR-88D transmit case is SZ-2 phase coding [48]. In this technique, a pulse-to-pulse phase code is applied to the transmitted signal to decorrelate second-trip echoes, mitigating contamination of radial velocity estimates. SZ phase codes are specified by two parameters M and n . The phase rotation applied to the k th pulse is given as follows:

$$\phi_k = \sum_{m=0}^k \frac{n\pi m^2}{M}, \quad (3.14)$$

$$k = 0, 1, 2, \dots, M - 1. \quad (3.15)$$

In the case of the WSR-88D, $n = 8$ and $M = 64$. It is simple to replicate this phase code for use in the decoding process, but before it can be removed from the received pulses the correct delay for the code (i.e., what value of k corresponds to each pulse) must be determined. This is accomplished by testing the set of possible delays and finding the solution that minimizes pulse-to-pulse phase variance in the decoded signal. Fortunately, this code has the property that when shifted by M/n pulses, it will produce a constant phase shift across all pulses. As we do not care about the absolute phase shift between the transmitter and receiver, this is an acceptable result. Therefore, we only need to test 8, rather than 64, possible delays. Figure 3.9

shows some representative results of this decoding process. Once the signal has been successfully decoded, frequency synchronization may proceed.

Accurate frequency synchronization is critical: a 1-ppb error in synchronization between the transmitter and receiver corresponds to a 0.15 m/s error in Doppler velocity estimation for a monostatic radar operating at 3 GHz. The most direct method for estimation of the frequency offset between the transmitter and the receiver is to take the Fourier transform of the received signal and estimate the peak frequency. However, due to multipath effects, this will result in intolerable errors in frequency estimation. Fortunately, by exploiting our knowledge of the signal properties, the effects of aliasing, and some minimal assumptions on the stability and accuracy of the transmitter and receiver LOs, we find that the use of pulse-pair processing to estimate the frequency offset produces acceptably accurate results through a process that is massively more computationally efficient than FFT processing.

To understand this conclusion we must first develop a mathematical model for our received signal spectrum. First assume a transmitter sending a pulsed signal to a receiver over an infinite time interval. The signal amplitude and link gain is unity and there is an angular carrier frequency offset of ω' between the two systems. This yields the following model for the received baseband signal:

$$f(t) = \exp(-j\omega't) \sum_{n=-\infty}^{\infty} \text{rect}\left(\frac{t - nT_s}{\tau}\right), \quad (3.16)$$

where T_s is the transmitter PRT and τ is the pulse width. The Fourier series repre-

sentation of this signal is as follows:

$$f(t) = D \sum_{n=-\infty}^{\infty} \text{sinc} \left[\frac{\tau}{2\pi} (n\omega_0 + \omega') \right] \exp(n\omega_0 t), \quad (3.17)$$

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}, \quad (3.18)$$

$$D = \frac{\tau}{T_s}, \quad (3.19)$$

$$\omega_0 = \frac{2\pi}{T_s}. \quad (3.20)$$

This yields the following amplitude spectrum:

$$|F(\omega)| = D \sum_{n=-\infty}^{\infty} \text{sinc} \left[\frac{\tau}{2\pi} (n\omega_0 + \omega') \right] \delta(\omega - n\omega_0). \quad (3.21)$$

This spectrum has the form of a sinc function, centered at ω' , with a bandwidth corresponding to the reciprocal of τ and sampled at intervals of 2π times the PRF of the signal. However, any real signal we measure will be of finite length L . Assuming we use a unit amplitude rectangular window, the final amplitude spectrum will have this form:

$$|F(\omega)| = DL \sum_{n=-\infty}^{\infty} \text{sinc} \left[\frac{\tau}{2\pi} (n\omega_0 + \omega') \right] \text{sinc} \left[\frac{L}{2} (\omega - n\omega_0) \right]. \quad (3.22)$$

This results in each of the sample impulses being replaced by a sinc function with a width inversely proportional to the length of the observation window (generally on the order of 10^{-6} times the width of the sinc function corresponding to the pulse bandwidth). Observe that the frequency offset could be accurately estimated by calculating the position of the spectrum's peak. Now, consider the effect of a catastrophic multipath scenario. Figure 3.10 shows a simulated pulse train signal both with and without a strong multipath component at a range delay of 450 m. The mul-




Figure 3.10: Example of a simulated pulse train spectrum, both with and without a strong multipath interferer. Note that the peak of the macroscopic spectrum shown in the top plot shifts significantly due to the change in mainlobe shape induced by the multipath signal, but the locations of the individual “samples” visible at the finer scales shown in the lower two plots are virtually unaffected. This compares favorably with the measured data in Figure 3.11.




Figure 3.11: Example of two consecutive measured pulse train spectra with varying multipath. It shows the same effect as in the simulations of Figure 3.10: significant peak movement at a macroscopic scale but extreme consistency at finer scales. Note that the smoothing and rapid rolloff of the expected sidelobes in the spectrum are due to low-pass filtering within the transceiver.

tipath changes the shape of the spectrum main lobe, causing it to become bimodal with peaks that no longer have a known relationship to the underlying frequency offset. However, an examination of the spectrum on a finer scale shows that the locations of the “sample” peaks have shifted by only a negligible amount due to their extremely small width relative to the gradient of the distortion induced by the multipath, as shown in the bottom two plots of Figure 3.10. Figure 3.11 demonstrates this same phenomenon using data actually measured by one of the receivers.

This property, that the “samples” remain essentially stationary in the presence of multipath, is exploited to achieve robust frequency offset estimation. The samples exist at the frequencies $\omega' \pm m\omega_0$ where ω_0 is the PRF in radians per second. Accurately estimating the location of any one of these samples is equivalent to estimating the correct frequency offset plus or minus some integer multiple of the PRF. This error carries over directly to estimation of the Doppler frequency. However, we know due to aliasing that an error of some exact integer multiple of the PRF will result in all Doppler velocities aliasing back to their correct values. Therefore, the location of any spectrum sample could be estimated and used to obtain correct Doppler velocities. The only drawback to choosing the “wrong” sample is that the matched filter will now be mismatched by some frequency offset. Fortunately, the actual offsets between the receiver and transmitter LOs are in practice small (on the order of 10 kHz) compared to the bandwidth of the pulse (on the order of 1 MHz). Figure 3.12 shows the relationship between sensitivity loss and frequency estimation error for a $1.57 \mu\text{s}$ pulse, the most common pulse length utilized by the WSR-88D. It is evident that errors of the magnitude that are observed in practice lead to negligible sensitivity loss. This allows us to estimate the location of a sample within the frequency range of $\pm\omega_0/2$ and use that for processing. This can be accomplished through pulse pair processing, using effectively the same expression




Figure 3.12: Loss in matched filter sensitivity due to errors in frequency estimation.

conventionally used to calculate Doppler frequency:

$$\omega' = \frac{1}{T_s} \arg \hat{R}(T_s), \quad (3.23)$$

$$\hat{R}(T_s) = \frac{1}{N} \sum_{n=0}^{N-1} V^*[n]V[n+1], \quad (3.24)$$

where $V[n]$ is the IQ signal measured at the receiver over a single dwell, N is the number of pulses in the dwell, and $\hat{R}(T_s)$ represents the estimated autocorrelation of the received signal at a lag of T_s .

There is one final step to obtain an accurate estimate of the frequency offset between the transmitter and the receiver. The frequency estimate that has been obtained at this point contains an undesirable component that must be removed before the correction is applied to the time-series data: the Doppler shift caused by the rotation of the antenna. This motion will cause a sinusoidally varying apparent Doppler shift in the direct-path signal. However, it does not affect the weather echo frequency (as the main beam is always orthogonal to the motion of the antenna), so it should not be part of our correction. To correct for this effect, the rotational Doppler shift must be calculated at each azimuth and removed from our estimate of the frequency offset. The distance from the WSR-88D phase center to its axis of rotation is approximately 4.8 m [49], [50], and its angular velocity and position are determined by the time stamped pointing angles used elsewhere in the processing scheme. Once the rotational Doppler shift has been compensated, the frequency offset is estimated and removed from the time series data corresponding to each pulse. At this point, the data are ready for pulse pair processing to obtain Doppler velocity estimates.

Once the synchronization process has successfully been carried out, the time series is organized using the estimated pulse locations into slow time/fast time for-




Figure 3.13: Example azimuthal pattern of the X-band proxy antenna used to study the effects of pattern phase on velocity estimation biases.

mat. Using the WSR-88D timestamped pointing angles, the fast-time vectors are then grouped into dwells corresponding to a traversal of the 1° beamwidth of the transmitter. Estimates of received power and Doppler frequency are retrieved from each range bin within a given dwell using basic pulse-pair processing [36]. These estimates are then localized in Cartesian space (relative to the transmitter) using (2.12)-(2.14).

3.6 Transmit Antenna Pattern Phase Effects

One theoretically possible source of frequency error is the phase of the antenna pattern. The significance of this effect has been noted in other phase-sensitive weather

applications such as the estimation of refractivity gradients using ground clutter observations [51]. Antenna radiation patterns do not merely change in amplitude with angle, but also change in phase. This will introduce some unknown error into the signal measured by the receiver that could only be corrected with relatively precise knowledge of the antenna pattern characteristics. This frequency estimate error will take some combination of two possible forms. The first possibility is frequency bias, caused by some non-zero average rate of antenna pattern phase change over the frequency estimation interval. The second possibility is spectrum broadening caused by zero-mean variation in the rate of antenna pattern phase change. This will not bias frequency drift estimates, but rather will increase their variance. In addition to the characteristics of the antenna pattern, the exact nature of the phase errors will also be influenced by the mechanical rotation rate of the transmitting antenna. For instance, in the case of a zero-mean, but slowly varying transmit antenna phase-modulation, the rotation rate could be slow enough that even though phase is zero-mean over the entire antenna pattern, the particular subsection of the pattern captured in a single observation could introduce a significant frequency estimation bias.

The favorable results of the system validation using weather measurements described in Chapter 4.1 suggest that this is not a major source of error. Ideally, this could be verified through measurements of an actual WSR-88D antenna. However, this is difficult in practice. Accurate antenna phase measurements require coherence between the transmitter and receiver. Such measurements for the WSR-88D are not publicly available, and there is no way to separate antenna phase effects from the LO frequency drift that we are trying to measure using passive receiver. As a proxy for a study of the actual WSR-88D antenna, measurements of an X-band parabolic dish antenna with a similar electrical size to that of a WSR-88D were col-




Figure 3.14: Results of the pattern phase error analysis across a range of rotation rates and transmit antenna mechanical elevations. These results reflect a 0.5 s data collection interval, and each plotted point corresponds to the worst possible result across the measured antenna pattern range.

lected at the RIL. An example pattern measurement from this antenna is shown in Figure 3.13. These measurements were then used to simulate the errors in frequency estimation that would result from antenna pattern phase over a range of rotation rates and elevation angles used operationally by the WSR-88D. The results of this study are shown in Figure 3.14.

The study assumed an 0.5 s data collection interval and the velocities shown correspond to a bistatic angle of 0° . The plotted points correspond to the worst-performing measurement interval across the available antenna pattern range. The obtained values are well within acceptable error levels for multiple Doppler measurements. While these velocity error values will scale inversely with $\cos(\beta/2)$, the bistatic angle would have to become quite large before these errors are seriously detrimental. At such large bistatic angles, the spatial resolution degradation would likely be a larger concern for measurement accuracy than these antenna pattern phase errors. This finding is consistent with the favorable system validation results presented in Chapter 4.1

3.7 Receiver Module Hardware

The presented bistatic radar hardware places an emphasis on simplicity, low cost, and a compact form factor. The primary components of the radar system are an antenna, a radio frequency (RF) front-end, a transceiver, and a small processing computer. As in existing multistatic weather radars, this system uses a fixed, relatively non-directive antenna. Currently, the systems use a series-fed V-polarized microstrip patch array. The beam pattern, shown in Figure 3.16, has an azimuthal 3 dB width of approximately 47.5° and a width in elevation of approximately 17.5° . It should be noted that the 3 dB beamwidth is not any sort of practical limit on the region in which useful data can be measured, particularly in azimuth where the




Figure 3.15: Block diagram of the passive radar hardware.

beam pattern rolloff is gradual. Echoes from well beyond the 3 dB points can be measured with sufficient SNR for high quality velocity measurements. Affixed to the back of the antenna is a weatherproofed enclosure containing the low-noise amplifier (LNA) and a bias tee. The LNA is separated from the rest of the RF front-end to be as close as possible to the antenna, maximizing SNR. The bias tee allows the LNA to receive power over the RF cable used to carry the received signal back to the main electronics enclosure.

The main enclosure is weatherproofed and climate controlled. It is outfitted with an ethernet port for system control and data transfer, as well as a 120 VAC power input, and two type-N ports that lead to the two channels of the RF transceiver. Within the box is the remainder of the RF front-end, consisting of a second bias tee used to transmit direct current to the LNA, a bandpass filter, a limiter, and an attenuator (to ensure that the maximum power output of the limiter could not damage the digital transceiver). After the limiter, the signal passes into the two-channel




Figure 3.16: Normalized radiation patterns of the receiver antenna measured in the ARRC far-field chamber. As the antenna is V-polarized, the E-plane corresponds to elevation, and the H-plane corresponds to azimuth.

digital transceiver. The transceiver is based on the Analog Devices 9361 integrated circuit and was developed at the Advanced Radar Research Center. This transceiver is, in turn, controlled by an Udoo x86 Ultra single board computer. A simplified schematic of this system is shown in Figure 3.15 and an example is shown in Figure 3.17 as installed on the roof of the Radar Innovations Laboratory at the University of Oklahoma. The total cost of this system, even when produced in small



Figure 3.17: An installed system on the roof of the Radar Innovations Lab at the University of Oklahoma.

quantities, is less than \$7,000. There is significant margin for improvement on this figure. \$5,000 of the total cost is the RF transceiver, which was chosen based on the facts that it was developed in-house, removing the need to deal with vendors for purchasing or support, and that suitable software for performing data collections had already been developed. Use of a commercially available software defined radio, while necessitating some significant software development efforts, could significantly drive down the price (as well as the size) of the receiver. One promising alternative, for example, is the bladeRF 2.0 micro [52] at a cost of as little as \$480. In its current form the main electronics enclosure has a size of 14"x11"x5", which could be reduced not only by a more compact transceiver, but by integration of the discrete RF and power supply components into purpose-built printed circuit boards.

Chapter 4

Weather Observations With a Multistatic Network

4.1 Experimental Validation of Velocity Measurements

On April 6, 2019 the prototype passive receiver network was used to collect observations of stratiform precipitation as it passed through the viable observation region of the network. Figure 4.1 shows the SNR and bistatic radial velocities measured by each of the receivers, as well as the radial velocity and reflectivity measured by KTLX. A few important features of these data should be noted. First, both spatial resolution and velocity resolution degrade with proximity to the baseline between the transmitter and the receiver. This leads to the high reflectivity “blobs” along these baselines with accompanying unstable velocity estimates. Clutter near these baselines is particularly troublesome, as it spreads out across ellipses of constant bistatic range, and is harder to filter without destroying the weather signal due to the degraded velocity resolution. This effect is particularly salient in the OUHSC data. Additionally, it should be noted that returns are visible over what might be a




Figure 4.1: Raw estimates of range-corrected power (not reflectivity) and bistatic Doppler velocity obtained by the passive receivers, along with reflectivity and radial velocity estimates from KTLX for comparison. The black contour illustrates the region selected for analysis based on the censoring criteria described in Section 4.1. The data shown were collected using a 4° elevation KTLX scan on April 6, 2019 at 16:11 UTC. Note that a ground clutter filter has not yet been applied to these data, in order to more clearly show the logic behind the censoring boundaries.




Figure 4.2: Retrieved horizontal wind vectors over the analysis region.

surprisingly large area despite antennas with fixed pointing directions (north from the RIL and south from OUHSC). This is due to the extremely gentle rolloff of the antenna pattern beyond the nominal 47.5° beamwidth. The poor front-to-back isolation of the antenna even allows clear echoes to be received from the backlobes (although much of this, particularly near the radar, will be badly contaminated by ground clutter as the antenna has a mechanical elevation tilt of about 9° that points the backlobes into the ground).

In order to validate the velocity measurements collected by the passive radars, an experiment similar to those carried out in [53] was conducted. The two passive radars were used to carry out a horizontal dual-Doppler analysis, with the assumption that at low elevations, the contribution of vertical velocity to measured radial velocities would be small. The resulting horizontal wind vectors were then projected onto the vectors representing the radial velocities measurable by KTLX. This




Figure 4.3: Actual KTLX radial velocity field and the estimate of the KTLX radial velocity obtained by projecting the retrieved horizontal wind field onto the vectors representing the pointing direction of KTLX at each point.

was then compared to the radial velocities that KTLX actually measured.

The first step in this process was to select a region of the available data suitable for analysis, as shown by the black contours in Figure 4.1. This was determined by three primary criteria: SNR (minimum of 3 dB), bistatic range (minimum of 25 km for the RIL receiver and 34 km for the OUHSC receiver), and antenna orientation (all points in the backplanes of the receiver antennas are censored). The reason for the SNR criterion (more specifically the minimum SNR at a given point between the two passive receivers) is straightforward. Low SNR would preclude accurate velocity estimation. This criterion most strongly influenced the western and northern boundaries of the censored region. The bistatic range criterion is used to eliminate the areas of poor spatial and Doppler resolution near the baselines, as well as to assist in rejecting problematic ground clutter. It should be noted that the OUHSC baseline was censored out to a larger bistatic range; this is due to the fact that its region of problematic ground clutter happens to extend further (the final ring of ground clutter used to determine this boundary is visible primarily in the velocity data). The bistatic range criterion is responsible for the ellipsoidal boundaries along the eastern edge of the analysis region. Finally, data in the backplane of each antenna was excluded to avoid the necessity of sorting what areas are too contaminated by the ground to be salvageable. The influence of this criterion is seen in the relatively straight southern border of the analysis region.

After the data were censored, the points within the analysis region were pre-processed. The data from KTLX and each receiver were first passed through a two-dimensional median filter in order to remove any speckle-like noise in the velocity images. The data were then interpolated to a uniform grid using a Cressman average with a radius of influence of 0.8 km. At this point the two-dimensional velocity retrieval was performed by solving the following the two dimensional ana-

log of (2.24) for u and v at each point in the analysis region:

$$\begin{bmatrix} \frac{\sin(a_1) \cos(e_1) + \sin(a_t) \cos(e_t)}{2 \cos(\beta_1/2)} & \frac{\cos(a_1) \cos(e_1) + \cos(a_t) \cos(e_t)}{2 \cos(\beta_1/2)} \\ \frac{\sin(a_2) \cos(e_2) + \sin(a_t) \cos(e_t)}{2 \cos(\beta_2/2)} & \frac{\cos(a_2) \cos(e_2) + \cos(a_t) \cos(e_t)}{2 \cos(\beta_2/2)} \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} v_{bi1} \\ v_{bi2} \end{bmatrix} \quad (4.1)$$

This calculation yields the horizontal wind field depicted by the quiver plot in Figure 4.2. While qualitatively consistent with the observed storm motion during the data collection period, the wind field does not provide any quantitative insight into the accuracy of the retrieved velocities. In order to obtain a real metric of accuracy, the estimated wind vector at each point is projected onto the along-beam direction of KTLX to obtain an estimate of the radial velocity that the WSR-88D *should* have measured. That estimated radial velocity is shown in Figure 4.3 alongside the measured KTLX radial velocity. Again, these results look qualitatively reasonable, but it is difficult to assess the magnitude of the differences between the two systems, particularly along the zero isodop. These error magnitudes are shown clearly in Figure 4.4. The root squared error (Figure 4.4a) falls below 1 ms^{-1} through much of the analysis domain, rising to $3\text{-}3.5 \text{ ms}^{-1}$ in some regions. Within the main body of the analysis region, these errors appear to correlate with gradients in reflectivity, which is consistent with expectations; it is in these areas that varying spatial resolution and higher sidelobes in the passive systems are likely to cause the largest differences in velocity estimates. The root mean squared error over the analysis domain is 1.25 ms^{-1} . The scatterplot in Figure 4.8b offers an alternative visualization of the error. This error has non-negligible components of both bias (0.78 ms^{-1}) and variance ($0.94 \text{ m}^2\text{s}^{-2}$). The bias is most likely driven primarily by sidelobe contamination effects as discussed in Section 4.3. These error levels are similar to those obtained in [54] using a multistatic weather radar synchronized through direct




Figure 4.4: Results of the KTLX radial velocity retrieval are depicted here as (a) spatial map of the root squared error in the retrieved estimate and (b) scatterplot of the relationship between the measured and retrieved values. The blue line in (b) represents a theoretical exact match between retrieved and measured values.

communication between transmitter and receiver. Note that if one looks closely, it is possible to notice some evidence of clustering or coherence within the scatterplot. This is attributable partially to the fact that the spatially sampled points are not independent; the data were oversampled initially (to varying degrees due to varying bistatic resolution volume size), and the Cressman averaging process uses overlapping regions of influence. Even amongst samples that are far enough apart to have a reasonable degree of independence, but are still located in the same general area, they are likely to have relatively similar radial velocities and similar biases due to experiencing sidelobe contamination originating from the same areas. This will, again, lead to the appearance of coherence within the scatter plot. It also appears that there is some dependence of bias on the measured KTLX velocity. Note, for example, the region of positive bias between -5 and 0 m/s. This is, however, misleading. What is actually occurring is that both the measured velocity and the level of bias depend on the spatial region being observed. Therefore, if a given spatial region is suffering from some sort of bias due to sidelobe contamination, for instance, then the range of radial velocities associated with that region will be biased.

4.2 Convective Observations

The observations discussed above were extremely useful for their intended purpose of verifying the validity of the experimental transmit/receive synchronization method used in the multistatic network. Stratiform rain subject to relatively uniform advection represents a near-best-case scenario in terms of avoiding serious sidelobe effects. This is because large gradients in both reflectivity and radial velocity are largely absent from the phenomena under observation. The observations presented in this section place that aspect of system performance under a higher degree of stress.




Figure 4.5: Raw estimates of range-corrected power (not reflectivity) and bistatic Doppler velocity obtained by the passive receivers, along with reflectivity and radial velocity estimates from KTLX for comparison. The black contour illustrates the region selected for analysis based on the censoring criteria described in Section 4.1. The data shown were collected using a 4° elevation KTLX scan on May 25, 2019 at 2:58 UTC. Note that a ground clutter filter has not yet been applied to these data, in order to more clearly show the logic behind the censoring boundaries.




Figure 4.6: Retrieved horizontal wind vectors over the analysis region.

On May 25, 2019 observations of a system of scattered convection were collected as it passed through the multistatic network observation area. Figure 4.5 shows data collected by KTLX and by each of the passive receivers during a 4° elevation scan at 2:58:05 UTC. This elevation was chosen to match that of the dataset used for the system validation discussed in Section 4.1, so as to facilitate as direct a comparison as possible. As in Figure 4.1, the black contours in Figure 4.5 represent the analysis region selected for study, based on the exact same geometric and SNR-based criteria utilized for the previous dataset. The differences in the weather within the two observation regions are most easily discerned through examination of the KTLX data, which lends itself to much more intuitive interpretation than the raw bistatic observations from the two passive receivers. The relevant differences between the two scenarios are increased variation in both reflectivity and velocity in the convective case.

On a large scale, both the stratiform and convective data sets contain moderate reflectivity gradients. For the stratiform data, these gradients are likely produced primarily by the melting layer, as evidenced by the distribution of an elevated reflectivity region in an arc about the transmitter location (with an interior edge distance that would correspond to a reasonable melting layer height of approximately 2 km). The western edge of the analysis region captures some of this bright band. One feature of particular interest is the area of weak echoes within this band centered at approximately -5 km meridional / -30 km zonal distance in the KTLX data shown in Figure 4.1. This represents the area of the analysis region most susceptible to sidelobe contamination, which will be discussed in more detail later in the section. Within the convective data, by contrast, reflectivity variation is produced by the distribution of convective cells and the space between. This, in turn, has produced more widespread reflectivity gradients, although it is not as intuitively obvious what areas are likely to be subject to particularly severe contamination problems.

Large scale examination of the KTLX radial velocity data suggests roughly uniform advection, as evidenced by the approximately linear geometry of the 0 m/s isodop. The greatest deviation from this structure is apparent in the bight-like northward incursion of the negative velocity region toward the western edge of the analysis region (which happens to roughly coincide with the aforementioned area of weak echoes). The non-uniformity of the velocity field captured in the convective dataset is immediately obvious through the extreme non-linearity of the 0 m/s isodop. Significant non-uniformity is also apparent within the analysis region, where radial velocities tend to become significantly more negative with decreasing radial distance to the transmitter in the southern portion of the analysis region and exhibit a roughly north/south gradient in the northern portion.

Precisely the same analysis procedure was carried out on the convective case






Figure 4.7: These plots show the actual KTLX radial velocity field as well as the estimate of the KTLX radial velocity obtained by projecting the retrieved horizontal wind field onto the vectors representing the pointing direction of KTLX at each point.



(a)



(b)

Figure 4.8: Results of the KTLX radial velocity retrieval are depicted here as (a) spatial map of the root squared error in the retrieved estimate and (b) scatterplot of the relationship between the measured and retrieved values. The blue line in (b) represents a theoretical exact match between retrieved and measured values.

as was used for the stratiform data. After median filtering and Cressman averaging, the simplified wind field estimation procedure was carried out, the results of which are shown in Figure 4.6. As predicted based on the raw KTLX radial velocity data, this wind field is much less heterogeneous, making it generally more likely that any sidelobe contamination will actually result in velocity biases. The estimated and actual KTLX radial velocities, are shown in Figure 4.7. Again, good general agreement is demonstrated. As in the stratiform case, however, the color scale makes it difficult to quantify differences visually. The differences (errors) between the estimated and actual KTLX radial velocities are shown in Figure 4.8. As expected, the estimated and actual velocities are highly correlated, with a root mean squared error of approximately 1.8 m/s. Note that the scatterplot shown in Figure 4.4 exhibits more obvious evidence of local error coherence, likely due to sidelobe contamination playing a more dominant role as a root cause of error. This, as previously mentioned, can produce correlated velocity errors even between grid points that should represent statistically independent samples based on the Cressman average radius. While still within an acceptable range, this does represent an increase in RMSE of approximately 50% over the stratiform case. Higher error levels in this instance are concentrated in the northern part of the analysis region, which is somewhat counterintuitive, as more convective activity and associated gradients are found in the southern portion of the analysis region. However, a closer analysis of the KTLX reflectivity observations and the range-corrected received power from the passive receivers makes a convincing case that the primary culprit for the local regions of elevated error in both analyses is, in fact, sidelobe contamination.

4.3 Discussion of Sidelobe Contamination in Collected Data

The simplest technique for the identification of regions of potential sidelobe contamination in bistatic receiver data is a comparison between the range-corrected power structures observed by the receiver to the reflectivity structures observed by the transmitting radar. It is obvious from a cursory examination of these data in Figures 4.1 and 4.5 that differences between the two passive systems and KTLX are significant. The first and most obvious reason for this is that range-corrected power and reflectivity are not the same quantity. To convert range-corrected power to reflectivity at the bistatic receiver would require correction for the receiver antenna pattern and for the resolution volume size at each point in the observed weather field. While this would be possible in principle, it would require detailed measurements of each individual receiver antenna to perform accurately. As accurate reflectivity measurements were not a primary interest of this work, that task was not undertaken. However, the effects of the receiver beam shape should be small within the central portion of the main beam, where its gradients are small. The next factor is varying spatial resolution, which is significant primarily near the transmitter receiver baseline. At favorable bistatic angles, this factor should not have a large effect on the shape of observed reflectivity structures. The final factor that can significantly affect the observed reflectivity fields is sidelobe contamination.

Based on these facts about the possible causes of reflectivity / received power structure differences between the transmitter and the passive receivers, it can be reasonably stated that major differences in structure that occur a) reasonably far from the system baseline and b) well within the main beam of the receiver can confidently be attributed to sidelobe contamination. As such, these regions will

therefore be *susceptible* to velocity measurement biases. It bears repeating that whether and to what degree these biases actually materialize is entirely dependent on how different the Doppler spectrum of the integrated sidelobe signal is from that of the main beam signal. Both the convective and stratiform datasets serve to illustrate these effects.

Figure 4.9 shows magnified versions of the reflectivity and range-corrected power observations from the stratiform dataset with an overlay of the constant-range ellipses corresponding to each transmitter/receiver pair. Notice that the primary structural difference between the KTLX images and each of the passive receiver images is the absence of the weak echo region in the western half of the analysis region. This area meets the general criteria of having a small-moderate bistatic angle and being located in the central portion of the observation region. It is therefore reasonable to assume that the “filling in” of this area in the passive data is due to sidelobe contamination. This idea can be further tested for plausibility through the examination of the bistatic range contours for each system. Sidelobe contamination in a given volume must originate from the same constant-range ellipsoidal shell as the resolution volume being contaminated. It can be readily observed that the bistatic range contours for both systems which pass through the weak echo region continue on through the bright band / melting layer to both the north and south. This creates strong reflectivity gradients within the constant range shell, leading to significant sidelobe contamination. There are also significant structural differences along the southern edge of the analysis region in the data collected by the RIL receiver. This likely reflects the fact that this area is at an extremely steep elevation angle relative to the receiver. Therefore, signals from this area are attenuated somewhat due to the receiver pattern. The RIL receiver contours which run through this region are largely filled with weather, which is sufficient to significantly contam-

inate the attenuated signal collected from the area of interest by the RIL receiver. The most significant areas of sidelobe-induced velocity error are therefore expected to occur through the western and southern regions of the analysis area. Referring back to Figure 4.4, it is evident that this is, in fact, the case. Note, however, that even though the sidelobe contamination is quite severe, the consequences for velocity estimation are relatively mild thanks to the homogeneity of the observed wind field.

Similarly, Figure 4.10 shows magnified versions of the reflectivity and range-corrected power observations from the convective dataset with an overlay of the constant-range ellipses corresponding to each transmitter/receiver pair. Here there is actually relatively little structural difference throughout the image corresponding to the OUHSC receiver. The RIL receiver, however, has some significant differences. The most significant of these occurs at the northern end of the analysis region. Here, any reflectivity structure is masked by contamination from the convection to the west of the analysis region. Note that a similar effect in the OUHSC data was likely mitigated by the fact that the contaminating weather region lies at an extreme elevation angle and partially in the backplane relative to the receiver, and its bistatic range contours do not pass through as much high-reflectivity area as those of the RIL receiver. The RIL receiver suffers from another contamination region at the southeastern edge of the analysis area. Here, the region of interest is again attenuated due to its extreme elevation angle relative to the receiver, while the contaminating weather lies closer to the center of the main beam. Based on these observations, there are possibilities of significant velocity error at both the northern and southeastern edges of the analysis area. Referring to Figure 4.8, there is heavy error at the northern end, but the southwestern region gave relatively unblemished wind retrieval results. Figure 4.6 can shed some light as to why this area escaped




Figure 4.9: a) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the RIL receiver b) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the OUHSC receiver c) RIL receiver range-corrected power with corresponding bistatic range contours and analysis region outline d) OUHSC receiver range-corrected power with corresponding bistatic range contours and analysis region outline. See Figure 4.1 for a wider view of the same observation set.

significant consequences of contamination. Note that the retrieved wind vectors are near orthogonal to the RIL/KTLX bistatic bisector, aligning much more closely with the OUHSC/KTLX bisector. This is supported by the much larger velocity measured in this area by the OUHSC receiver, as shown in Figure 4.5. Therefore, the wind vector estimate here was dominated by the relatively pristine data collected at the OUHSC receiver.

In summation, the data collected by this prototype network certainly demonstrates that the experimental transmit/receive synchronization method outlined in Chapter 3 works effectively, as even small synchronization errors would result in catastrophic errors in velocity measurements that are clearly not present here. However, as is expected, observations of both stratiform and convective weather systems exhibit convincing evidence of the presence of significant sidelobe contamination. The next chapter looks to the future and examine a potential technique for mitigating this effect in future multistatic networks.




Figure 4.10: a) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the RIL receiver b) KTLX H-polarized reflectivity with analysis region outline and bistatic range contours corresponding to the OUHSC receiver c) RIL receiver range-corrected power with corresponding bistatic range contours and analysis region outline d) OUHSC receiver range-corrected power with corresponding bistatic range contours and analysis region outline. See Figure 4.5 for a wider view of the same observation set.

Chapter 5

Sidelobe Whitening Simulation and Analysis

As documented in prior work on multistatic weather radar [28], [42], as well as in Chapter 4 of this work, sidelobe contamination is a major contributor to velocity estimation error using multistatic systems. One potential method of mitigating this error is sidelobe whitening, which was proposed for use in multistatic applications in [28], but was first developed for monostatic systems by Sachidananda and Zrnić [31], [55]. The basic idea of the Sachidananda/Zrnić (SZ) sidelobe whitening algorithm is to synthesize a pair of array patterns with equal mainlobe phases, but with sidelobe phases that are 180° out of phase with each other. By then switching between these array patterns according to a pseudorandom phase code, the Doppler spectrum contribution of sidelobe signals is whitened, but signals from the mainlobe retain their original degree of coherence. Thus, while sidelobe leakage will still make a significant contribution to the level of received power, the induced

Doppler velocity bias will be greatly reduced. It is worth noting that this concept bears some similarities to sidelobe manipulation techniques developed later for the purpose of secure communications, which seek to send a desired sequence of bits in the direction of one or more intended receivers while broadcasting a distorted or scrambled sequence elsewhere. One general technique for achieving this goal is to transmit two independently weighted signals simultaneously: an “information” signal with a traditional directive beam pattern pointed toward the intended receiver, and an “interference” signal with a pattern having a null in the intended receiver direction and a roughly uniform power level elsewhere [56], [57]. This technique does not offer the kind of precise sidelobe phase control offered by the technique presented in this chapter. More importantly for the multistatic radar application, it also relies on an array architecture which allows for the injection of two signals with independent spatial weighting, which is not a typical feature in PAWRs. Another similar method is directional modulation [58], [59], which manipulates array weight phases in order to produce a desired bit sequence in one or more directions of interest while maximizing bit error rate in some number of undesirable directions. Not only does this technique not offer the precise sidelobe control given by the sidelobe whitening algorithm, but it does not guarantee any particular shape for the transmitted power pattern. While this is not necessarily a problem in many communications applications, it is not acceptable for most radar use-cases. A variation on this idea is presented by Snow and Chappell [60], who perform directional modulation through the transmission of independent waveforms on each element of an array, producing a decorrelation of sidelobe contributions at the matched filter output. The use of such a system assumes the ability to transmit independent waveforms at each element, which, like the ability to inject an independently weighted signal, is not a common feature (although at least one forthcoming PAWR, Horus




Figure 5.1: High-level summary of the proposed pattern synthesis technique.

[34], would be capable of implementing this method). It also would result in sensitivity issues, as again the power pattern is not controlled (varying in this case over the pulse duration).

However, the SZ algorithm has limits to the achievable level of sidelobe phase control, requiring a two-way pattern to achieve the necessary 180° phase modulation for a binary-coded whitening scheme. For a multistatic application, only the transmit pattern is available for modulation, presenting a problem to which this chapter offers a solution. The details of the new pattern synthesis process are laid out in Section 5.1.

5.1 Pattern Synthesis for Sidelobe Whitening

The additive beamforming weight perturbation calculated using the SZ technique will shift the sidelobe phases on each side of the mainlobe by $\pm 45^\circ$. The two-way sidelobe phases are therefore shifted by $\pm 90^\circ$ relative to the main lobe. A

pulse-to-pulse phase shift of 180° in the two-way sidelobes can then be achieved by reversing the order of the weights across the array on successive pulses. It would, in principal, be possible to follow a similar approach to the monostatic problem. One could seek an additive component to one particular set of conventional weights that would cause a 180° difference between the sidelobes on each side of the mainlobe. However, a major drawback of this approach is that it is unnecessarily overconstrained. By seeking an additive perturbation to a single fixed set of weights, one is specifying not only the desired sidelobe phases and amplitudes, but their precise locations and shapes (as dictated by the initial set of weights). These unnecessary constraints lead to significant degradation in pattern synthesis performance. In [31] this problem was mitigated by using non-uniform array spacings to create initial array factors with sidelobe topography more conducive to the whitening process. However, using such a technique experimentally would require a purpose-built array, which is often not practical. This concern motivated a search for a new, less constrained approach that would allow satisfactory whitening results using a uniform array.

The pattern synthesis technique proposed here for multistatic sidelobe whitening is a variation on the method of alternating projections [61]–[63]. The basic concept of this method is to take an initial set of element excitations, and transform it to a beampattern. The beampattern is then minimally adjusted such that every point falls within some desired maximum amplitude envelope. A new set of excitations is then calculated by finding the values that best approximate the pattern in a least-squares sense. If there are any constraints on the element excitations, such as a maximum phase difference between adjacent excitations, the excitation vector is adjusted to meet those requirements. This process is then repeated iteratively until converging to a solution that meets all constraints (insofar as possible)

in both the excitation and beampattern domains. What differentiates the technique proposed here from a typical application of this method is an additional set of phase constraints in the beampattern domain. The modified technique is outlined in Figure 5.1 and described in detail throughout the remainder of this section.

The first step in implementing the alternating projections technique is to establish a set of linear equations linking the element excitations to the beampattern. For the purposes of this chapter, it is assumed that the desired beampattern (and therefore the corresponding excitation vector) is symmetric, which has the convenient effect of reducing the number of independent variables in the system of equations by a factor of two. Given this symmetry, the excitations and beampattern for a uniform linear array (ULA) with $2N$ elements can be completely described as follows:

$$\mathbf{F}^{(k)} = \mathbf{C}\mathbf{x}^{(k)}, \quad (5.1)$$

$$\mathbf{C} = \begin{bmatrix} \cos(\psi_1 u_1) & \dots & \cos(\psi_N u_1) \\ \vdots & \ddots & \vdots \\ \cos(\psi_1 u_L) & \dots & \cos(\psi_N u_L) \end{bmatrix}, \quad (5.2)$$

$$\psi_n = \frac{\pi(2n-1)d}{\lambda}, \quad (5.3)$$

$$u_l = \sin(\theta_l), \quad (5.4)$$

where \mathbf{F} is an $L \times 1$ vector containing sampled beampattern values, \mathbf{x} is an $N \times 1$ vector containing element excitations, L is the number of sampled angles in the beampattern, n is the element index, l is the sample angle index, θ is the sampled angle of incidence measured from broadside, d is the element spacing, λ is the operating wavelength of the array, and (k) denotes the iteration index within the alternating projections algorithm. Due again to the symmetry assumption, sample values of u for all results produced in this paper were chosen as dense uniformly




Figure 5.2: Example whitening pattern synthesis results for a 152 element $\lambda/2$ -spaced ULA. The mask used is an approximation of the WSR-88D azimuth pattern at 0° elevation. Coding phase is the phase difference between the two synthesized patterns to be used in the binary whitening scheme. Note that there is some deviation from ideal results (0 in the mainlobe and π elsewhere), particularly within the first sidelobe and in the pattern nulls.

spaced points on the interval $[0, 1]$.

5.1.1 Phase and Amplitude Constraints

Once the excitation vector has been transformed to a beampattern, that pattern must then be modified to fit the phase constraints necessary for whitening, and the amplitude constraints imposed by some specified radiation pattern envelope. The phase constraints are imposed first. In considering what these phase constraints should be in order to achieve the desired whitening effect, first note that the i^{th} sidelobe phases of a typical (e.g., real-valued, symmetric, and non-negative) array taper are:

$$\phi(i) = \pi[1 + (-1)^{|i|}]. \quad (5.5)$$

In other words, it alternates symmetrically between 0 and π , beginning with a phase of π in the first sidelobe. The SZ binary algorithm synthesizes a pattern in which the

sidelobe phases of the one-way pattern in the positive u direction are rotated by $\pi/4$ radians, while those in the negative u direction are rotated by $-\pi/4$ radians. The sidelobe amplitude remains unaffected. This yields the following sidelobe phases:

$$\phi'_a(i) = \begin{cases} \pi[1 + (-1)^{|i|}] + \frac{\pi}{4}, & i > 0 \\ \pi[1 + (-1)^{|i|}] - \frac{\pi}{4}, & i < 0 \end{cases}, \quad (5.6)$$

where the sign of i indicates whether the sidelobe lies in the negative or positive u direction relative to the mainlobe. The two-way sidelobe phases are then given as:

$$\phi'_{a2}(i) = \begin{cases} 2\pi[1 + (-1)^{|i|}] + \frac{\pi}{2}, & i > 0 \\ 2\pi[1 + (-1)^{|i|}] - \frac{\pi}{2}, & i < 0 \end{cases} \quad (5.7)$$

$$= \begin{cases} \frac{\pi}{2}, & i > 0 \\ -\frac{\pi}{2}, & i < 0 \end{cases}. \quad (5.8)$$

The excitation function corresponding to this beampattern will be designated $x_a(n)$, while the beampattern will be designated $F_a(u)$. Consider the excitation function $x_b(n) = x_a(-n)$. By the time-reversal property of the Fourier transform, we know that its frequency response will be $F_b(u) = F_a(-u)$. The phases of its two-way sidelobes will then be:

$$\phi'_{b2}(i) = \begin{cases} -\frac{\pi}{2}, & i > 0 \\ \frac{\pi}{2}, & i < 0 \end{cases}. \quad (5.9)$$

Thus, the phase difference between the two-way sidelobes of $F_a(u)$ and $F_b(u)$ at any point has a magnitude of π , as is desired for a binary whitening scheme. In order to implement a similar algorithm for the purposes of multistatic whitening,

a pattern pair must be synthesized that has a one-way sidelobe phase difference magnitude of π . This will be achieved through the synthesis of a beampattern with the following sidelobe phases:

$$\phi_c(i) = (-1)^{|i|} \frac{\pi}{2}. \quad (5.10)$$

From the corresponding excitation function $x_c(n)$ and beampattern $F_b(n)$, and utilizing time reversal and complex conjugation properties of the Fourier transform, we can create a new excitation function $x_d(n) = x_c^*(-n)$ which will have frequency response $F_d(n) = F_c^*(n)$. This corresponds to sidelobe phases of:

$$\phi_d(i) = (-1)^{|i|-1} \frac{\pi}{2}, \quad (5.11)$$

which produces a phase difference of magnitude π at each point in the sidelobe regions. As we have already constrained the excitation vectors to be even-symmetric, the calculation of the second excitation vector can be simplified even further to $x_d(n) = x_c^*(n)$. The selection of an even-symmetric rather than an odd-symmetric (as in the SZ method) sidelobe phase modulation bears some discussion. Firstly, the odd-symmetric phase modulation would not allow symmetry assumptions on the excitation vector and would therefore increase the dimension of the alternating projections optimization problem. Secondly, the even-symmetric modulation results, in practice, in smaller residual phase gradients across the central portion of the main beam compared to the odd-symmetric modulation. These gradients result in some small, but nevertheless undesirable, spectral spreading effect on signals within the main beam.

In order to impose the phase constraint described by (5.10), the beampattern samples in vector \mathbf{F} are broken into segments with boundaries defined by local min-

ima in amplitude. Given sufficiently dense pattern sampling, these minima should correspond to beampattern nulls. The first segment (assuming samples of u range from zero to 1) will contain the mainlobe region and should be discarded, as its phase should not be modulated. For each of the remaining segments, every individual sample in the i^{th} segment is adjusted in phase to $\phi_c(i)$ while preserving its amplitude. The amplitude constraints are defined by some $L \times 1$ mask vector \mathbf{M} , containing the maximum beampattern amplitude for each value of u_i . This defines both the desired main beam shape and sidelobe amplitude. This can be specified arbitrarily; however, specifications that exceed the limits of what is physically realizable (e.g., a main beam narrower than that of a uniformly weighted array) will yield poor results. Each element of \mathbf{F} with an amplitude greater than the corresponding element of \mathbf{M} has its amplitude adjusted to the maximum specified by \mathbf{M} . Together, application of these constraints allows us to construct a modified beampattern vector $\mathbf{G}^{(k)}$:

$$G_l^{(k)} = \begin{cases} F_l^{(k)}, & |F_l^{(k)}| \leq |M_l| \text{ and } i = 0 \text{ (mainlobe)} \\ F_l^{(k)} \frac{|M_l|}{|F_l^{(k)}|}, & |F_l^{(k)}| > |M_l| \text{ and } i = 0 \\ F_l^{(k)} \exp[j(\phi_c(i) - \angle F_l^{(k)})], & |F_l^{(k)}| \leq |M_l| \text{ and } i \neq 0 \\ F_l^{(k)} \exp[j(\phi_c(i) - \angle F_l^{(k)})] \frac{|M_l|}{|F_l^{(k)}|}, & |F_l^{(k)}| > |M_l| \text{ and } i \neq 0 \end{cases} \quad (5.12)$$

where $G_l^{(k)}$, $F_l^{(k)}$, and M_l represent the l^{th} elements of $\mathbf{G}^{(k)}$, $\mathbf{F}^{(k)}$, and \mathbf{M} respectively.

5.1.2 Least-Squares Excitation Retrieval

Once the adjusted beampattern vector $\mathbf{G}^{(k)}$ has been generated, it is next necessary to retrieve the excitation vector that most nearly reproduces the desired pattern in a least-squares sense. However, an unweighted least-squares solution will generally yield unsatisfactory results. Typical beampatterns for radar applications have sufficiently low sidelobe levels relative to the mainbeam that an unweighted solution will tend to focus on accurate reproduction of the mainbeam shape at the expense of unacceptable inaccuracy in meeting target sidelobe levels. It is desirable to be able to tune this performance tradeoff, which can be accomplished through the use of a weighting matrix:

$$\mathbf{W} = \text{diag}(\mathbf{M})^\alpha. \quad (5.13)$$

The mainbeam/sidelobe tradeoff is adjusted through the scaling parameter α , where more negative values of α place more emphasis on meeting amplitude and phase specifications in the sidelobes at the expense of main beam shape fidelity. For the examples provided here, an α of -0.485 was used. Using this weighting matrix, our least squares solution for the excitation vector can be formulated as:

$$\mathbf{x}^{(k+1)} = (\mathbf{C}'\mathbf{W}\mathbf{C})^{-1}\mathbf{W}\mathbf{G}^{(k)}. \quad (5.14)$$

The new excitation vector can then be converted to its corresponding beampattern using (5.1), and this process can be repeated until \mathbf{x} converges to a stable solution. In this implementation of the algorithm, convergence is declared when the criterion ϵ drops below some specified threshold:

$$\epsilon = \frac{(\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)})^H (\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)})}{\mathbf{x}^{(k)H} \mathbf{x}^{(k)}} \quad (5.15)$$




Figure 5.3: Evolution of the convergence criterion ϵ during the synthesis of the pattern shown in Figure 5.2. It crosses the threshold value of 1×10^{-12} after 80 iterations.

5.1.3 Modified Two-Dimensional Implementation

The final modification that will be proposed here for the SZ algorithm is for the extension of the whitening technique to two dimensions. The method suggested in [31] is to simply implement the coding along a single dimension of the array, while leaving the taper along the orthogonal direction constant from pulse to pulse. This whitens the principal sidelobes along the dimension with the coded taper, and also whitens the non-principal sidelobes in the four quadrants of the array pattern, but leaves the principal sidelobes along the dimension of the unperturbed taper unwhitened. This is suboptimal, as large gradients in reflectivity are possible across both azimuth and elevation. It is possible, however, to achieve whitening through-

out the entire sidelobe region of a two dimensional array. This can be achieved by encoding the tapers along the two array dimensions using two uncorrelated pseudorandom whitening codes c_u and c_v . For the simulations that follow, Gold codes [64] were used, as they have excellent auto-correlation properties, as well as the property that the exclusive OR of two Gold codes is also a Gold code. In this way, every area in the sidelobe region of the two-dimensional array will experience one of three possible codes ($c_u, c_v, c_u c_v$), each of which is an effective whitening code. This is illustrated in Figure 5.4.




Figure 5.4: Map showing which of the three possible whitening codes is applied to each sidelobe region of a sample two-dimensional array pattern.




Figure 5.5: Results of a 64 point whitening code used in conjunction with the synthesized pattern. Doppler spectrum peak attenuation is the reduction in the maximum Doppler spectrum value for an impinging clutter signal at each angle. As expected, near zero attenuation occurs in the mainlobe.

5.2 Whitening Algorithm Results

The algorithm described in Section 5.1 was used to synthesize a pair of whitening patterns using a mask based on an approximation of a 0° elevation cut of the WSR-88D radiation pattern. The array manifold used to synthesize the pattern is a 152-element $\lambda/2$ spaced ULA. This number of elements was chosen because a 152×152 element square array has a total aperture of 57.76 m^2 . This is the closest possible approximation to the WSR-88D aperture of 57.16 m^2 achievable using an even number of elements, an attribute that becomes relevant for the simulations discussed in Section 5.3. This mask, the resulting pattern amplitude, and the achieved phase difference between the two synthesized patterns are all shown in Figure 5.2. The behavior of the convergence criterion ϵ throughout the synthesis process is shown in Figure 5.3.

While the synthesized pattern fits the specified envelope reasonably well, the fit

is not perfect. These imperfections are primarily due to compromises in beamwidth and sidelobe level made in order to meet the phase constraints. This fact can be demonstrated through comparison with a set of pattern synthesis results that use the same envelope but do not impose any phase constraints. The pattern synthesized without phase constraints stays within the envelope to within 0.01 dB at all points, and has a 3 dB beamwidth of 0.898° . By contrast, the whitened pattern shown in Figure 5.2 has a beamwidth of 0.997° . Figure 5.2 also shows the imperfections in the final phase behavior. Ideally, the phase difference between the two patterns synthesized for the binary whitening process (indicated in the figure as “Coding Phase”) would be 0 within the mainlobe and π elsewhere. However, this is not realized perfectly. There is some deviation from a phase of 0 within the mainlobe, the magnitude of which is less than 0.038 radians within the 3 dB beamwidth. The phase at the first sidelobe peak is -2.2 radians, while the phases of the remaining sidelobe peaks vary about a median value of approximately 2.8 radians. There are phase gradients across each sidelobe and erratic phase behavior near nulls where the pattern amplitude is extremely low, and therefore an insignificant contributor to system performance. Figure 5.5 shows the achieved whitening performance across the synthesized pattern. The metric used to characterize performance is Doppler spectrum peak attenuation. To calculate this metric, a 64-point Doppler power spectrum is calculated corresponding to a stationary point target (similar to ground clutter) at every angle in the beam pattern, both with and without application of a whitening code. Doppler spectrum peak attenuation is the ratio of the maximum value of the unwhitened spectrum to the maximum value of the whitened spectrum. As expected, signals within the mainlobe are virtually unattenuated, with measured attenuations of less than 0.001 dB across the 3 dB beamwidth. Attenuation in the first sidelobe is approximately -6.4 dB, while the sidelobes beyond that have attenua-




Figure 5.6: Several examples of the clutter signal Doppler spectra produced using the 64 point code (calculated with a Hamming window). As anticipated based on the peak attenuations in Figure 5.5, the mainlobe spectrum receives virtually no spectral spreading, although the sidelobes are slightly perturbed relative to an ordinary Hamming window spectrum. The far sidelobe is extremely well whitened, with DC no longer having a dominant peak, and the first sidelobe spectrum lies between these extremes.

tions of approximately -12.9 dB. Attenuation near pattern nulls is naturally erratic, but this is both expected based on the coding results and of little consequence.

Figure 5.6 shows examples of the Doppler spectra corresponding to a DC (clutter-like) signal incident through a sidelobe in each of these regions. The spectra have been normalized such that their total power is equal (i.e. the effects of relative beam pattern power have been eliminated). A Hamming window was applied to the time series samples for sidelobe reduction prior to spectrum calculation. The spectrum peak due to the stationary point target is distinguishable in all three spectra at varying levels. As expected, the peak is strongly dominant for the mainlobe case, with the only clear evidence of the main beam phase imperfections being the irregular perturbation of the sidelobes from the topography that would be expected from an unmodulated Hamming window. A significant whitening effect is achieved within the first sidelobe, with the spectrum peak only exceeding the next highest spectrum peak by approximately 5.5 dB. In the far sidelobes, sufficient whitening has been achieved such that the DC peak is no longer the global maximum.

5.3 Weather Radar System Simulations

5.3.1 Simulator Description

In order to examine what effects sidelobe whitening might have in a realistic weather observation scenario, a simple multistatic weather radar time-series simulator was developed. The simulator uses a Monte Carlo framework similar to the simulation architectures described in [65], [66]. A high level outline of the simulator structure is shown in Figure 5.7. The first step in the simulation process is the population of the observation domain with scattering centers (SCs). These are point targets, each of which represents some small volume of precipitation. These point targets are




Figure 5.7: High-level view of the simulation process. The receiver-related portion of the radar range equation calculation, as well as the sorting of scattering centers into range bins, are done prior to calculation of the transmitter-related weighting contributions. This structure maximizes the efficiency of the simulation process.




Figure 5.8: Gain patterns for each simulated transmitter scenario. Since the one-dimensional array factors of the two binary whitening patterns are identical, this means that the two-dimensional power pattern will remain constant during the whitening process and is also identical to that of the unwhitened array pattern.

placed throughout the domain according to a uniform random distribution. Based on the position of each SC, it is assigned a reflectivity, and a three-dimensional velocity, using quad-linear interpolation from a set of regularly gridded numerical weather prediction data as described in [65].

The eventual goal of the simulation process is to calculate the signal formed by summing the (complex) radar range equation over all the SCs for every sample instant, transmitter pointing angle, range bin and receiver, requiring a set of nested for-loops. The order of these loops, as well as which calculations are made at each level of the nested structure, can affect the computational efficiency of the simulation process by orders of magnitude. This is particularly true when a few simplifying conditions are applied to the simulation architecture. The most significant of these simplifications is that advection between transmitter dwells (i.e. transmitter pointing angles) is ignored. While this assumption does somewhat reduce the realism of the simulator output, and would make it unsuitable for certain studies (evaluation of advection correction techniques, for example) it has no bearing on the study carried out here. This simplification allows the portion of the radar range equation dictated by receiver location and antenna pattern to be pre-calculated such that it does not need to be calculated separately for each individual transmitter pointing angle. This portion of the equation can be written (for the m th scattering center) as:

$$w_{rx}^{(m)} = \frac{f_{rx}(u_{rx}^{(m)}, v_{rx}^{(m)})\lambda(\sigma^{(m)})^{1/2}}{(4\pi)^{1/2}r_{rx}^{(m)}} \exp\left(-2\pi j\frac{r_{rx}^{(m)}}{\lambda}\right), \quad (5.16)$$

where $f_{rx}(\theta, \phi)$ is the complex radiation pattern of the receiver evaluated at the uv coordinate position corresponding to each scatterer and $\sigma^{(m)}$ is the effective radar cross-section (RCS) of the scattering center. Antenna patterns within the simulator exist as lookup tables, from which the values corresponding to each scatterer are

determined through two-dimensional interpolation.

The RCS σ can be calculated as:

$$\sigma^{(m)} = Z_v^{(m)} \gamma^{(m)} \bar{v}, \quad (5.17)$$

where $Z_v^{(m)}$ is V-polarized reflectivity, \bar{v} is the average spatial volume represented by each SC, calculated by dividing the total simulation domain volume by the number of SCs, and γ represents an angle-dependent loss factor induced by the bistatic scattering behavior of a raindrop. It has been shown that raindrops can be accurately modeled as reradiating incident energy as infinitesimal dipoles with a moment vector aligned along the polarization direction of the incident wave [67], an effect described in greater detail in Section 2.5. This means that in a coordinate system in which the scatterer system lies at the origin, with the incident polarization oriented along the z -axis, $\gamma = \sin^2(\theta_{rx})$, where θ_{rx} indicates the location of the receiver in spherical coordinates. However, the simulator uses a coordinate system with the transmitter located at the origin, and an xy -plane that lies tangent to the earth's surface at the transmitter location. It is therefore necessary to translate between these two coordinate systems in order to calculate γ .

Making the simplifying assumption that the transmitting antenna exhibits ideal polarization behavior in the Ludwig II sense [68], the polarization direction of the incident wave at each scatterer will lie along the $\hat{\theta}$ direction. The coordinate system can be rotated and translated to position the scatterer at the origin and the polarization of the incident wave along the z -axis using the following steps:

1. Perform a rotation about the z -axis by $-\phi_s$ to position the scatterer in the xz -plane (in the positive- x half-plane).
2. Perform a rotation about the y -axis by $90^\circ - \theta_s$ to position the scatterer on the

positive x -axis. The polarization vector is now aligned along the z -axis.

3. Translate the scatterer along the x -axis to the origin.

The calculations corresponding to this process are as follows:

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \mathbf{R}_y(90^\circ - \theta_s) \mathbf{R}_z(-\phi_s) \begin{bmatrix} x \\ y \\ z \end{bmatrix} - \begin{bmatrix} r_{\text{tx}} \\ 0 \\ 0 \end{bmatrix}, \quad (5.18)$$

$$\mathbf{R}_z(-\phi_s) = \begin{bmatrix} \cos(-\phi_s) & -\sin(-\phi_s) & 0 \\ \sin(-\phi_s) & \cos(-\phi_s) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (5.19)$$

$$\mathbf{R}_y(90^\circ - \theta_s) = \begin{bmatrix} \cos(90^\circ - \theta_s) & 0 & \sin(90^\circ - \theta_s) \\ 0 & 1 & 0 \\ -\sin(90^\circ - \theta_s) & 0 & \cos(90^\circ - \theta_s) \end{bmatrix}, \quad (5.20)$$

where xyz and $x'y'z'$ represent the receiver position in the transmitter-centered and scatterer-centered Cartesian coordinate systems, respectively, and $\theta_s \phi_s$ represent the angular position of the scatterer in transmitter-centered spherical coordinates. The loss factor γ can then be calculated as:

$$\gamma = \sin^2(\theta_{\text{rx}}) \quad (5.21)$$

$$= \sin^2 \left[\cos^{-1} \left(\frac{z'}{\sqrt{x'^2 + y'^2 + z'^2}} \right) \right] \quad (5.22)$$

$$= 1 - \cos^2 \left[\cos^{-1} \left(\frac{z'}{\sqrt{x'^2 + y'^2 + z'^2}} \right) \right] \quad (5.23)$$

$$= 1 - \frac{z'^2}{x'^2 + y'^2 + z'^2}. \quad (5.24)$$

In addition to the calculation of the receiver pattern weights, all of the scattering

centers are pre-sorted into range bins for each receiver, prior to the transmitter-dependent portion of the radar range equation calculations. Assuming that the sampling time t_0 relative to pulse transmission is measured from the pulse center, the bounds of the range bin measured at t_0 are given by:

$$\frac{ct_0}{2} - \frac{c\tau}{4} \leq r_{\text{bi}} \leq \frac{ct_0}{2} + \frac{c\tau}{4}. \quad (5.25)$$

After this sorting process is carried out across all scattering centers for each receiver, final radar range equation calculations begin. For each transmitter pointing angle, the following transmitter-dependent weight is calculated across all SCs:

$$w_{tx}^{(m)} = \frac{f_{\text{tx}}(u_{\text{tx}}^{(m)}, v_{\text{tx}}^{(m)})}{4\pi r_{\text{tx}}^{(m)}} \exp\left(-2\pi j \frac{r_{\text{tx}}^{(m)}}{\lambda}\right), \quad (5.26)$$

where $f_{\text{tx}}(u_{\text{tx}}^{(m)}, v_{\text{tx}}^{(m)})$ is the complex transmitter radiation pattern value evaluated at the uv coordinate position corresponding to each scatterer. Then, for each range bin / receiver combination, the receive weights corresponding to the SCs within that range bin are retrieved, and the final calculation to obtain a time series point at the receiver is carried out:

$$V(t_0) = \sum_{m=1}^M w_{tx}^{(m)} w_{rx}^{(m)}, \quad (5.27)$$

where M is the total number of SCs in the range bin.

Once the receiver voltage for every range bin / receiver / transmit angle combination has been calculated, the simulator advances to the next time series point. The scatterer positions are updated based on their velocities, and new time series points are calculated. However, the scatterers' parameters (velocity/reflectivity) are not recalculated at their new positions. This simplifies the necessary calculations for the scatterer update, and is reasonable for the simulator's intended application of providing single frames of radar data with dwell times spanning only 10s of

milliseconds. If the simulator was intended to accurately capture the evolution of weather over time, these parameters would need to be updated at each time step. Also, as realistic spectrum widths are not particularly important for this application, the scatterer positions are deterministic. Thus, the only significant contributor to spectrum width is wind shear.

Once the time-series $V(t)$ has been simulated, Gaussian white noise of power N_0 is added, where:

$$N_0 = k(T_{\text{rx}} + T_{\text{A}})B. \quad (5.28)$$

Here k is Boltzmann's constant, T_{rx} is receiver noise temperature, T_{A} is antenna noise temperature, and B is receiver bandwidth.

Simulated Scenarios

In order to test the proposed whitening scheme, and provide some reference to levels of sidelobe contamination that would be expected from currently available transmitters, four different systems were simulated. The radiation patterns of each of these simulated systems are shown in Figure 5.8. The first of these systems is designed to be a parabolic dish with characteristics similar to the WSR-88D. This pattern was constructed by first using the theoretical expression [69]:

$$S(u) = \frac{\frac{48J_3(u)}{u^3} + \frac{0.32J_1(u)}{u}}{1.16}, \quad (5.29)$$

where J_n denotes a Bessel function of the first kind with parameter n . This expression, however, is highly idealized and does not capture the effects of the feed structure or diffraction and spillover around the dish edges. It provides a fairly accurate model for the main beam shape. However, the sidelobes are much too optimistic. Therefore, the sidelobe amplitudes of the pattern produced by the theo-




Figure 5.9: Constant-altitude slice of the NWP data grid used to produce the simulation results. Reflectivity is shown in (a) while (b) and (c) show the zonal and meridional wind field components, respectively. The receiver and transmitter locations are also indicated, as well as a dashed line indicating the boundary of the observation region shown in Figures 5.10-5.11.




Figure 5.10: Doppler velocity fields measured by each receiver for each simulated transmitter. There are significant errors due to sidelobe contamination in each image, primarily along areas of sharp reflectivity and velocity gradients. Note also regions of degraded spatial resolution along the transmitter/receiver baselines.




Figure 5.11: Differences between the dish, unwhitened array, and whitened array simulations and the sidelobe-free ideal simulation. There are significant reductions in bias prevalence and magnitude between the dish and either of the arrays. Whiten- ing provides a noticeable improvement compared to the unwhitened array.

retical expression were adjusted to levels which approximate those observed in the 0° elevation pattern cuts in [69]. In order to produce a two-dimensional pattern it is then assumed that these sidelobe levels are radially symmetric. This property of radial symmetry is a simplifying approximation. Sidelobe levels (particularly for a center-fed dish like the WSR-88D) are not uniform. The most pronounced non-uniform property is that there will be several ridges of raised sidelobes corresponding to the radial positions of the spars supporting the feed. For the WSR-88D, the 0 degree cut actually lies along one of these ridges, so the sidelobe envelope which is radially symmetric here is actually “several dB” [69] higher than the sidelobes of the WSR-88D away from the elevated ridges. This cut was used primarily due to a dearth of available high-quality published measurements of WSR-88D patterns along cuts less affected by the feed spars.

The second simulated system is a 7.6 m square array as described in Section 5.2. Here, the array taper was designed using the method described in Section 5.1. The resulting taper was then applied across both dimensions of the array. However, no whitening codes are applied to the array. It is immediately noticeable that this pattern has drastically lower sidelobes than the simulated dish away from the principal planes of the array (much lower than the “several dB” by which WSR-88D patterns may improve away from the elevated sidelobe ridges). One might ask why this array factor has total sidelobe energy that is so much lower than that of the simulated dish, which has a nearly equal total aperture area. This is a result of the fact that we are assuming a well-calibrated array, but the WSR-88D pattern includes all of the effects of the feed structure and dish edges. It is possible, however, to construct dish antennas using strategies that greatly mitigate these effects, which can then achieve levels of total sidelobe energy that approach those of a ideal values much more closely. One of these strategies is the use of an offset feed to reduce sidelobe per-

turbations due to the feed structure and its supports [70], which has actually already been demonstrated with significant success at CSU-CHILL [71]. Another is the use of edge treatments to reduce the effects of diffraction [70]. As such, a comparison between the results obtained by the unwhitened array and those obtained by the simulated dish should not be viewed primarily as a comparison between dishes and phased arrays, but rather a measure of what might be achieved with more efficient use of the aperture sizes already being used for long-range weather radars, which might be achieved with an upgraded dish design as well as a phased array. The simulated array is mechanically scanned in azimuth, a configuration that is gaining some recognition as a possible future configuration for a phased array replacement [72].

The third pattern has an identical power pattern to the unwhitened array, but binary pattern switching is enabled during the simulation as described in Section 5.1.3. This means that four different two-dimensional antenna patterns are actually used, however they all have precisely the same amplitudes, with the only difference being in phase. The final pattern is an “ideal” antenna, which has the same main beam shape and amplitude as the dish, but has no sidelobes. This is used as a reference in order to isolate the effects of the sidelobes in the simulation results.

The NWP data used are from a WRF simulation of the May 31, 2013 tornadic storm in El Reno, Oklahoma using the parameterizations described in [73]. This provides an excellent test case due to the extreme velocity and reflectivity gradients present in severe weather. A visualization of the section of WRF data used is shown in Figure 5.9. Figure 5.9 also shows the locations of the simulated transmitter and receivers. Table 5.1 provides a list of the key simulation parameters used.

Table 5.1: Simulation Parameters

Parameter	Value
Transmit Gain	45.5 dBi
Receive Gain	18 dBi
Transmit Power	375 kW
Pulse Length	1.57 μ s
PRT	0.78 ms
Transmit Elevation	1°
Receiver Bandwidth	570 kHz
Receiver Noise Figure	3 dB
Receive Antenna Noise Temp.	63.36 ° K

5.3.2 Results

The simulated Doppler velocity observations are shown in Figure 5.10. There are several unique aspects of bistatic Doppler velocity measurements that are important to understand these results. First, bistatic Doppler velocities do not measure the velocity along radials relative to the receiver, but rather the velocity component along the bistatic bisector, the vector bisecting the angle formed by the transmitter, scatterer, and receiver with the scatterer at the vertex. Second, the spatial resolution of bistatic systems degrades along the baselines between the transmitter and each receiver. This effect causes the distorted “blob” artifacts along the northern edge of each image for receiver 1, and along the eastern edge of each image for receiver 2.

There are two conditions that must be met for velocity biases due to sidelobe contamination to occur. First, there must be a significant velocity gradient tangent to the transmitted beam within a given bistatic range bin. If no velocity gradient exists, then sidelobe power will still contaminate the measurement, but it will not induce biases. Second, the scatterers contributing to the sidelobe contamination must have high enough reflectivity, and be distributed over a large enough solid angle that their total power contribution is able to overcome the sidelobe attenuation and contribute a significant amount of power to the received signal compared to the power incident

through the main beam. There are several locations in the simulated measurements where these conditions are obviously met. Specifically, this occurs in and around the two weak reflectivity “holes” in the central part of the simulation domain, as well as along the southwestern fringes of the storm. Anomalies in these areas are visible in the velocity images as differences between the “ideal” images (which suffer no sidelobe contamination) and the other simulated systems. However, due to the large dynamic range of the velocity measurements, it is difficult to observe these differences with accuracy in the velocity images. Figure 5.11 shows these differences much more clearly. Each of the plots in this figure shows the difference in estimated velocity between the simulated system indicated in the title and the “ideal” mainbeam-only system.

It is evident that the highest error levels relative to the “ideal” system occur for the simulated dish. Significant errors here occur not only in the aforementioned areas where severe reflectivity and velocity gradients are evident, but are spread throughout large areas of the simulation domain. This is to be expected, as the dish has an integrated sidelobe power that is dramatically higher than either of the simulated arrays. It has sidelobe levels equal to those along the worst-case cuts of the array pattern distributed across a much larger area. As mentioned in Section 5.3.1, this does not reflect a particular advantage of phased arrays, but rather a more efficient use of a given aperture size.

The errors in both the whitened and unwhitened array data are much more tightly concentrated near areas of severe gradients. As expected, the whitening process causes clear reductions in both the area affected by velocity biases and in the bias magnitudes within those areas. An example of the Doppler spectra within a contaminated resolution volume for the unwhitened, whitened, and ideal arrays is shown in Figure 5.12. All three simulations show significant peaks near -300




Figure 5.12: Examples of the unwhitened, whitened, and ideal spectra from the simulation results shown in Figure 5.10. The sidelobe leakage visible in the unwhitened results is spread throughout the whitened spectrum, resulting in a closer match to the ideal results.

Hz. The unwhitened array spectrum shows a spurious peak due to sidelobe leakage near 100 Hz. The whitening process successfully attenuates this peak, spreading its energy throughout the spectrum and resulting in a mean frequency estimate that more closely matches that obtained from the ideal array. There are limits to the potential improvements in Doppler frequency estimates available through whitening. Most significantly, if the peaks of the whitened sidelobe power spectrum are still large relative to the spectral peaks corresponding to power incident through the main beam, then estimates will still exhibit poor performance. This is the reason for the residual frequency estimation error in the central low-reflectivity region and along the southwest border of the storm.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

The use of multistatic radar networks for the collection of multiple-Doppler data has unquestionable advantages over the use of multiple monostatic systems in terms of cost and inter-radar synchronization. However, research into this technology has slowed over the last two decades. The primary goal of this work is to reexamine the potential of this system architecture in the context of recent technological developments. The research presented in Chapters 3-4 is dedicated to the development of a compact, low-cost passive receiver with accompanying signal processing techniques that allow for automated synchronization with a minimally cooperative transmitter. This was enabled by the recent widespread availability of highly integrated commercial off-the-shelf transceivers (such as the AD9361 used in the prototype systems) brought about principally by developments in the telecommunications industry. Systems like these have the potential to significantly decrease the barrier of entry into multistatic weather radar research through reduced cost and improved flexibility in choice of transmitter. The greatly reduced size and weight compared to past multistatic receivers also opens up a wide range of deployment

possibilities.

These receivers were used as platforms for the development of a set of techniques that allow for accurate pulse-timing and frequency synchronization via side-lobe radiation between a radar transmitting a typical weather radar signal and a remote receiver. These methods offer effective solutions to the problems of frequency estimation of a short duty cycle signal in strong multipath environments and of robust direct-path signal detection and quality control. Additionally, modifications to these techniques are shown which allow compatibility with specialized transmit schemes utilized by the WSR-88D. The validity of these techniques is shown through the presentation of a set of actual weather measurements, with accuracy verified through their consistency with the radial velocity measurements obtained by KTLX. This set of signal processing methods offers advantages in terms of both the cost and flexibility of the receiver modules used in a multistatic network. It eliminates the need for GPS-disciplined oscillators as well as the need for the transmitting radar to measure, store, or transmit any information on pulse timing or phase. This means that almost any existing weather radar could be used as a transmitter with little or no modification, as virtually every existing weather radar is capable of recording timestamped pointing angle information.

The second portion of this work, presented in Chapter 5, was devoted to an exploration into the use of pattern synthesis methods to mitigate the problem of velocity bias due to sidelobe contamination, possibly the most vexing problem facing practical applications of multistatic weather radar technology. While the theoretical / simulation studies carried out for this work could have been carried out at any point, they are more relevant now than ever due to the proliferation of PAWR systems within the research community, and the strong possibility that the eventual replacement for the WSR-88D will be a PAWR.

The pattern synthesis technique developed is a modified version of the alternating projections algorithm capable of producing near-180° sidelobe phase modulation in the one-way beam pattern, as required for the multistatic application. It also allows for the specification of an arbitrary beamwidth and sidelobe topography. A weather radar system simulation was used to demonstrate that this whitening technique is capable of achieving significant improvements in Doppler velocity estimation accuracy in problem areas of observed weather fields (those with large cross-beam gradients in velocity and reflectivity within a bistatic range bin).

An additional illuminating result of the simulation process was the considerable improvement in velocity bias achievable merely through the more efficient use of antenna apertures with sizes similar to those in many existing long-range weather radar systems such as the WSR-88D. In the simulations carried out here, this was achieved through the use of a well-calibrated array, but similar results should be achievable through the use of an offset feed and / or improved edge treatments. The possibility of improved data quality merely through improved transmitter antenna design has received little attention in prior research into the bistatic sidelobe contamination problem.

6.2 Future Work

The signal processing techniques developed for transmitter / receiver synchronization, as well as the prototype multistatic network itself, offer several interesting application possibilities, including the ability to flexibly use these radars with multiple transmitters as part of field campaigns or simply to augment WSR-88D installations as was done in this study. Both of these applications offer particularly exciting possibilities in environments where setting up multiple Doppler measurements with monostatic mobile radars is difficult or impossible due to terrain and beam block-

age issues (typified by the southeastern United States). The extreme low cost of these systems may also make them particularly attractive to research institutions with limited budgets that may not have the resources to obtain multiple monostatic radars in order to carry out multiple Doppler studies.

One particularly intriguing use-case for this technology is installation on airborne platforms. While the current form factor of the system is not yet suitable for this application, there is still considerable room for size and weight reduction as discussed in Chapter 3. This would provide a number of advantages. Airborne platforms would have the advantage mentioned above of virtually eliminating problems of beam blockage. Further, they would enable greater vertical spatial diversity in the network, which could prove helpful in establishing favorable geometries for the collection of high-quality vertical velocities at low altitudes. Finally, one challenge of a fixed multistatic network is the limited observation area, due in part to the reduced sensitivity of multistatic observations and fixed receiver antenna pointing directions. Airborne platforms, by contrast, would be able to relocate to track the advection of phenomena of interest. Furthermore, the positions of each receiver could be easily optimized to create the most favorable multistatic geometries for specific regions of interest within a storm. In theory, these last objectives could also be attempted through the use of ground-based mobile platforms. However, practical difficulties of navigation, coordination, and availability of suitable observation locations would make this an extreme challenge.

Future research plans also involve the expansion of the Oklahoma City network to include a greater number of receivers for the purpose of improving coverage and reducing error through the use of redundant observations from multiple systems. Such an expansion would likely also entail some improvements to the receiver design, such as the installation of a positioner to allow for remote adjustment of the an-

tenna pointing direction. The network will be used in conjunction with WSR-88Ds and the ARRC's mobile radars in order to conduct some of the first long-duration, long-domain, multiple-Doppler observations of supercells and MCSs.

The presented sidelobe whitening technique also offers promising avenues for future research. In addition to the primary motivating application of multistatic radar, the ability to perform sidelobe whitening using only the one-way beam pattern makes possible the use of this technique for receive-side beamforming in data collected using imaging radars, such as the Atmospheric Imaging Radar [74] or its forthcoming successor the Polarimetric Atmospheric Imaging Radar [33]. Velocity bias due to sidelobe contamination, particularly from ground clutter, is also a significant issue for these systems that may be improved through the use of this method.

While this work establishes the potential of the sidelobe whitening technique to effect improvements in Doppler velocity estimation quality, it is limited to a small number of promising examples. Future research should seek to explore how to best optimize that potential and explore its limits. In other words, such studies should seek to determine what combinations of beamwidth, sidelobe envelope, and presence / absence of whitening produce the highest quality Doppler velocity estimates across a range of weather and ground-clutter scenarios. This kind of work would be complemented by more rigorous simulation work that considers the practical limits of what kind of sidelobe performance can be achieved using improved parabolic dish designs and what levels of phased array calibration accuracy might be necessary to achieve parity with those designs.

These two lines of research come together with the possibility of future experiments in the practical application of multistatic sidelobe whitening using any of the several existing or forthcoming phased array weather radars operated by the

University of Oklahoma (e.g., Horus [34] and PAIR [33]) or the National Weather Service (e.g., the ATD [32]). The prototype network would be ideal for this purpose as its sidelobe synchronization capabilities will avoid any need for burdensome hardware modifications to these systems. Such experiments would be invaluable in experimentally verifying the potential of sidelobe whitening suggested by the simulation results presented here, as well as allowing for experimental exploration of the general multistatic transmitter optimization problem previously mentioned in this section.

Bibliography

- [1] A. D. Byrd, R. D. Palmer, and C. J. Fulton, “Development of a Low-Cost Multistatic Passive Weather Radar Network”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. In press, 2020. DOI: 10.1109/TGRS.2019.2955606.
- [2] A. Byrd, R. D. Palmer, and C. J. Fulton, “Doppler Velocity Bias Mitigation Through Sidelobe Whitening for Multistatic Weather Radar”, *Geoscience and Remote Sensing, IEEE Transactions on*, vol. In Review, 2020.
- [3] J. R. Probert-Jones, “Meteorological Use of Pulsed Doppler Radar”, *Nature*, vol. 186, no. 4721, pp. 271–273, 1960.
- [4] R. Lhermitte and D. Atlas, “Precipitation motion by pulse Doppler”, in *Proceedings, 9th Weather Radar Conference*, Boston, MA: Amer. Meteor. Soc., 1961, p. 223.
- [5] D. Boccippio, “A Diagnostic Analysis of the VVP Single-Doppler Retrieval Technique”, *Journal of Atmospheric and Oceanic Technology*, vol. 12, no. April, pp. 230–248, 1995.
- [6] K. A. Browning and R. Wexler, “The Determination of Kinematic Properties of a Wind Field Using Doppler Radar”, *Journal of Applied Meteorology*, no. February, pp. 105–113, 1968.

- [7] R. Lhermitte, “New Developments in Doppler Radar Methods”, in *Proc. 13th Radar Meteor. Conf.*, 1968, pp. 14–17.
- [8] L. Armijo, *A Theory for the Determination of Wind and Precipitation Velocities with Doppler Radars*, 1969. DOI: 10.1175/1520-0469(1969)026<0570:atftdo>2.0.co;2.
- [9] D. C. Dowell and A. Shapiro, “Stability of an iterative dual-Doppler wind synthesis in Cartesian coordinates”, *Journal of Atmospheric and Oceanic Technology*, vol. 20, no. 11, pp. 1552–1559, 2003, ISSN: 07390572. DOI: 10.1175/1520-0426(2003)020<1552:SOAIDW>2.0.CO;2.
- [10] P. S. Ray, C. Ziegler, W. C. Bumgarner, and R. J. Serafin, “Single- and Multiple-Doppler Radar Observations of Tornadoic Storms”, *Mon. Wea. Rev.*, vol. 108, no. October, pp. 1607–1625, 1980.
- [11] J. Testud and M. Chong, *Three-dimensional wind field analysis from dual-Doppler radar data. Part I: Filtering, interpolating and differentiating the raw data*. 1983. DOI: 10.1175/1520-0450(1983)022<1204:TDWFAF>2.0.CO;2.
- [12] A. Shapiro and J. J. Mewes, “New formulations of dual-Doppler wind analysis”, *Journal of Atmospheric and Oceanic Technology*, vol. 16, no. 6, pp. 782–792, 1999, ISSN: 07390572. DOI: 10.1175/1520-0426(1999)016<0782:NFODDW>2.0.CO;2.
- [13] M. Chong, J. Testud, and F. Roux, “Three-Dimensional Wind Field Analysis from Dual-Doppler Radar Data Part II: Minimizing the Error due to Temporal Variation”, *Journal of Climate and Applied Meteorology*, vol. 22, no. July, pp. 1216–1226, 1983.

- [14] A. Shapiro, K. M. Willingham, and C. K. Potvin, “Spatially variable advection correction of radar data. Part I: Theoretical considerations”, *Journal of the Atmospheric Sciences*, vol. 67, no. 11, pp. 3445–3456, 2010, ISSN: 00224928. DOI: 10.1175/2010JAS3465.1.
- [15] A. Shapiro, K. M. Willingham, and C. K. Potvin, “Spatially variable advection correction of radar data. Part II: Test results”, *Journal of the Atmospheric Sciences*, vol. 67, no. 11, pp. 3457–3470, 2010, ISSN: 00224928. DOI: 10.1175/2010JAS3466.1.
- [16] N. J. Willis, “Bistatic Radar”, in *Radar Handbook*, M. I. Skolnik, Ed., Second, Boston, MA: McGraw Hill, 1990, ch. 25, ISBN: 978-0070579132.
- [17] L. Brown, *Technical and Military imperatives: A Radar History of World War 2*. Bristol, 1999, p. 45.
- [18] B. A. Austin, “Precursors to radar—the Watson-Watt Memorandum and the Daventry experiment.”, *International Journal of Electrical Engineering Education*, vol. 36, no. 4, pp. 365–372, 1999.
- [19] D. Atlas, K. Naito, and R. E. Carbone, “Bistatic Microwave Probing of a Refractively Perturbed Clear Atmosphere”, *Journal of the Atmospheric Sciences*, vol. 25, no. 2, pp. 257–268, Mar. 1968, ISSN: 0022-4928. DOI: 10.1175/1520-0469(1968)025<0257:BMPOAR>2.0.CO;2. [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/1520-0469%7B%5C%7D281968%7B%5C%7D29025%7B%5C%7D3C0257%7B%5C%7D3ABMPOAR%7B%5C%7D3E2.0.CO%7B%5C%7D3B2>.
- [20] R. J. Doviak and C. M. Weil, “Bistatic Radar Detection of the Melting Layer”, *Journal of Applied Meteorology*, vol. 11, no. 6, pp. 1012–1016, Sep. 1972,

ISSN: 0021-8952. DOI: 10.1175/1520-0450(1972)011<1012:BRDOTM>2.0.CO;2. [Online]. Available: [http://journals.ametsoc.org/doi/abs/10.1175/1520-0450\(1972\)011<1012:BRDOTM>2.0.CO;2](http://journals.ametsoc.org/doi/abs/10.1175/1520-0450%7B%5C%7D281972%7B%5C%7D29011%7B%5C%7D3C1012%7B%5C%7D3ABRDOTM%7B%5C%7D3E2.0.CO%7B%5C%7D3B2). [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/1520-0450%7B%5C%7D281972%7B%5C%7D29011%7B%5C%7D3C1012%7B%5C%7D3ABRDOTM%7B%5C%7D3E2.0.CO%7B%5C%7D3B2>.

- [21] R. Crane, “Bistatic scatter from rain”, *IEEE Transactions on Antennas and Propagation*, vol. 22, no. 2, pp. 312–320, Mar. 1974, ISSN: 0096-1973. DOI: 10.1109/TAP.1974.1140766. [Online]. Available: <http://ieeexplore.ieee.org/document/1140766/>.
- [22] J. Wurman, S. Heckman, and D. Boccippio, “A Bistatic Multiple-Doppler Radar Network”, *Journal of Applied Meteorology*, vol. 32, no. 12, pp. 1802–1814, Dec. 1993, ISSN: 0894-8763. DOI: 10.1175/1520-0450(1993)032<1802:ABMDRN>2.0.CO;2. [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/1520-0450%7B%5C%7D281993%7B%5C%7D29032%7B%5C%7D3C1802%7B%5C%7D3AABMDRN%7B%5C%7D3E2.0.CO%7B%5C%7D3B2>.
- [23] J. Wurman, M. Randall, C. Frush, E. Loew, and C. Holloway, “Design of a bistatic dual-Doppler radar for retrieving vector winds using one transmitter and a remote low-gain passive receiver”, *Proceedings of the IEEE*, vol. 82, no. 12, pp. 1861–1872, 1994, ISSN: 00189219. DOI: 10.1109/5.338075. [Online]. Available: <http://ieeexplore.ieee.org/document/338075/>.
- [24] M. Hagen, P. Meischner, J. Wurman, M. Randall, and C. Burghart, “A C-Band Bistatic Doppler Radar System at DLR Oberpfaffenhofen”, *29th International Conference on Radar Meteorology*, no. July, pp. 813–814, 1999.

- [25] A. Protat and I. Zawadzki, “A variational method for real-time retrieval of three-dimensional wind field from multiple-Doppler bistatic radar network data”, *Journal of Atmospheric and Oceanic Technology*, vol. 16, no. 4, pp. 432–449, 1999, ISSN: 07390572. DOI: 10.1175/1520-0426(1999)016<0432:AVMFRT>2.0.CO;2.
- [26] K. Nakagawa, H. Hanado, S. Satoh, and T. Iguchi, “Development of the CRL Okinawa Bistatic Polarimetric Radar”, *Journal of the Communications Research Laboratory*, vol. 49, no. 2, pp. 225–231, 2002.
- [27] K. Nishimura, E. Gotoh, and T. Sato, “Fine Scale 3D Wind Field Observation With a Multistatic Equatorial Atmosphere Radar”, *Journal of the Meteorological Society of Japan*, vol. 84A, pp. 227–238, 2006, ISSN: 0026-1165. DOI: 10.2151/jmsj.84A.227. [Online]. Available: <http://joi.jlc.jst.go.jp/JST.JSTAGE/jmsj/84A.227?from=CrossRef>.
- [28] R. De Elia and I. Zawadzki, “Sidelobe contamination in bistatic radars”, *Journal of Atmospheric and Oceanic Technology*, vol. 17, no. 10, pp. 1313–1329, 2000, ISSN: 07390572. DOI: 10.1175/1520-0426(2000)017<1313:SCIBR>2.0.CO;2.
- [29] M. Chong, N. Lamrani, and M. Hagen, “A Variational Correction Method as an Alternative to Forced Rejection of Sidelobe-Contaminated Bistatic Doppler Measurements”, *Journal of Atmospheric and Oceanic Technology*, vol. 25, no. 11, pp. 1939–1954, Nov. 2008, ISSN: 0739-0572. DOI: 10.1175/2008JTECHA1117.1. [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/2008JTECHA1117.1>.
- [30] S. Kawamura, H. Hanado, S. Sugitani, and K. Nakagawa, “A new bistatic doppler measurement system with reduced contamination by sidelobe echoes”,

- in *2010 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, Jul. 2010, pp. 4150–4153, ISBN: 978-1-4244-9565-8. DOI: 10.1109/IGARSS.2010.5653081. [Online]. Available: <http://ieeexplore.ieee.org/document/5653081/>.
- [31] M. Sachidananda, R. J. Doviak, and D. S. Zrnić, “Whitening of Sidelobe Powers by Pattern Switching in Radar Array Antenna”, *IEEE Transactions on Antennas and Propagation*, vol. 33, no. 7, pp. 727–735, 1985, ISSN: 15582221. DOI: 10.1109/TAP.1985.1143661.
- [32] J. E. Stailey and K. D. Hondl, “Multifunction Phased Array Radar for Aircraft and Weather Surveillance”, *Proceedings of the IEEE*, vol. 104, no. 3, pp. 649–659, 2016, ISSN: 00189219. DOI: 10.1109/JPROC.2015.2491179.
- [33] J. L. Salazar, T.-Y. Yu, C. J. Fulton, M. McCord, R. D. Palmer, H. B. Bluestein, B.-L. Cheong, M. I. Biggerstaff, B. M. Isom, J. M. Kurdzo, R. J. Doviak, X. Wang, and M. Yearly, “Development of a Mobile C-band Polarimetric Atmospheric Imaging Radar (PAIR)”, in *AMS Special Symposium on Meteorological Observations and Instrumentation*, 2017.
- [34] R. D. Palmer, C. J. Fulton, J. L. Salazar, H. Sigmarsson, and M. Yearly, “The ”Horus” Radar - An All-Digital Polarimetric Phased Array Radar for Multi-Mission Surveillance”, in *35th Conference on Environmental Information Processing Technology*, Phoenix, AZ: Amer. Meteor. Soc., 2019.
- [35] N. J. Willis, *Bistatic Radar*, Second edi. Raleigh, NC: Scitech Publishing Inc, 2005, ISBN: 1891121456. DOI: 10.1049/sbra003e.
- [36] R. J. Doviak and D. S. Zrnić, *Doppler radar and weather observations*, Second edi. Mineola, NY: Dover Publications, Inc.

- [37] K. Aydin, S. H. Park, and T. M. Walsh, “Bistatic dual-polarization scattering from rain and hail at S- and C-band frequencies”, *Journal of Atmospheric and Oceanic Technology*, vol. 15, no. 5, pp. 1110–1121, 1998, ISSN: 07390572. DOI: 10.1175/1520-0426(1998)015<1110:BDPSFR>2.0.CO;2.
- [38] S. Sakagami, “Some experimental results on bistatic scatter from rain”, *IEEE Transactions on Antennas and Propagation*, vol. 28, no. 2, pp. 161–165, 1980.
- [39] P. Bauer, A. Thorpe, and G. Brunet, “The quiet revolution of numerical weather prediction”, *Nature*, vol. 525, no. 7567, pp. 47–55, 2015, ISSN: 14764687. DOI: 10.1038/nature14956.
- [40] S. L. Barnes, *A Technique for Maximizing Details in Numerical Weather Map Analysis*, 1964. DOI: 10.1175/1520-0450(1964)003<0396:atfmdi>2.0.co;2.
- [41] G. P. Cressman, “An Operational Objective Analysis System”, *Monthly Weather Review*, vol. 87, no. 10, pp. 367–374, 1959, ISSN: 0027-0644. DOI: 10.1175/1520-0493(1959)087<0367:A00AS>2.0.CO;2.
- [42] S. Satoh and J. Wurman, “Accuracy of Wind Fields Observed by a Bistatic Doppler Radar Network”, *Journal of Atmospheric and Oceanic Technology*, vol. 20, no. 8, pp. 1077–1091, Aug. 2003, ISSN: 0739-0572. DOI: 10.1175/1520-0426(2003)020<1077:AOWFOB>2.0.CO;2. [Online]. Available: <http://journals.ametsoc.org/doi/abs/10.1175/1520-0426%7B%5C%%7D282003%7B%5C%%7D29020%7B%5C%%7D3C1077%7B%5C%%7D3AAOWFOB%7B%5C%%7D3E2.0.CO%7B%5C%%7D3B2>.

- [43] J. George, D. Brunkow, and V. Chandrasekar, “Networking CSU-CHILL and CSU-Pawnee to form a bistatic radar system”, *International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 656–659, 2006. DOI: 10.1109/IGARSS.2006.172.
- [44] Google, *Google Earth*, Oklahoma City, OK. [Online]. Available: <http://www.earth.google.com>.
- [45] A. Byrd, R. Palmer, and C. Fulton, “Concept for a Passive Multistatic UAV-borne Weather Radar”, *2018 IEEE Radar Conference (RadarConf18)*, pp. 845–850, 2018. DOI: 10.1109/RADAR.2018.8378670.
- [46] —, “Implementation of a Low-Cost Passive Weather Radar and First Weather Observations”, in *2019 IEEE Radar Conference (RadarConf19)*, Boston: IEEE, 2019, Accepted.
- [47] E. D. B. Almeida, “Revisiting Normalized Cross-Correlation for Accurate Camera Pose Estimation and Accurate Real-Time Multiple View Stereo © Copyright 2015 by Eduardo de Brito Almeida”, PhD thesis, Brown University, 2015, pp. 34–38.
- [48] M. Sachidananda and D. S. Zrnić, “Systematic Phase Codes for Resolving Range Overlaid Signals in a Doppler Weather Radar”, *Journal of Atmospheric and Oceanic Technology*, vol. 16, pp. 1351–1363, 1999.
- [49] National Weather Service, *System Description: Doppler Meteorological Radar WSR-88D*, 2018.
- [50] R. Ice, A. Heck, J. Cunningham, W. Zittel, and R. Lee, “Polarimetric Weather Radar Antenna Calibration Using Solar Scans”, *AMTA 36th Annual Meeting and Symposium*, pp. 4–9, 2014.

- [51] Y.-C. Feng and F. Fabry, “The Imperfect Phase Pattern of Real Parabolic Radar Antenna and Data Quality”, *Journal of Atmospheric and Oceanic Technology*, pp. 2655–2661, 2016. DOI: 10.1175/JTECH-D-16-0143.1.
- [52] Nuand, *bladeRF 2.0 micro Datasheet*. [Online]. Available: <https://www.nuand.com/bladerf-2-0-micro/> (visited on 02/21/2019).
- [53] K. Friedrich and O. Caumont, “Dealiasing doppler velocities measured by a bistatic radar network during a downburst-producing thunderstorm”, *Journal of Atmospheric and Oceanic Technology*, vol. 21, no. 5, pp. 717–729, 2004, ISSN: 07390572. DOI: 10.1175/1520-0426(2004)021<0717:DDVMBA>2.0.CO;2.
- [54] K. Friedrich and M. Hagen, “Evaluation of wind vectors measured by a bistatic Doppler radar network”, *Journal of Atmospheric and Oceanic Technology*, vol. 21, no. 12, pp. 1840–1854, 2004, ISSN: 07390572. DOI: 10.1175/JTECH-1679.1.
- [55] D. S. Zrnić and M. Sachidananda, “A Switched Pattern Radar Antenna Array”, *IEEE Transactions on Antennas and Propagation*, vol. 35, no. 10, pp. 1104–1110, 1987, ISSN: 15582221. DOI: 10.1109/TAP.1987.1143995.
- [56] Y. Ding and V. Fusco, “Sidelobe manipulation using Butler matrix for 60 GHz physical layer secure wireless communication”, *2013 Loughborough Antennas and Propagation Conference, LAPC 2013*, no. November, pp. 61–65, 2013. DOI: 10.1109/LAPC.2013.6711852.
- [57] Y. Zhang, Y. Ding, and V. Fusco, “Sidelobe modulation scrambling transmitter using fourier rotman lens”, *IEEE Transactions on Antennas and Prop-*

- agation, vol. 61, no. 7, pp. 3900–3904, 2013, ISSN: 0018926X. DOI: 10 . 1109/TAP.2013.2254453.
- [58] M. P. Daly and J. T. Bernhard, “Directional modulation technique for phased arrays”, *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 9, pp. 2633–2640, 2009, ISSN: 0018926X. DOI: 10 . 1109 / TAP . 2009 . 2027047.
- [59] M. P. Daly, E. L. Daly, and J. T. Bernhard, “Demonstration of Directional Modulation Using a Phased Array”, *IEEE Transactions on Antennas and Propagation*, vol. 68, no. 5, pp. 1545–1550, 2010, ISSN: 00255327.
- [60] T. Snow and W. J. Chappell, “Directional modulation using independent waveforms at each element of a phased array radar”, *IEEE MTT-S International Microwave Symposium Digest*, pp. 1–4, 2013, ISSN: 0149645X. DOI: 10.1109/MWSYM.2013.6697796.
- [61] O. Bucci, G. Franceschetti, G. Mazzarella, and G. Panariello, “Intersection approach to array pattern synthesis”, *IEE Proceedings H - Microwaves, Antennas and Propagation*, vol. 137, no. 6, pp. 349–357, 1990.
- [62] O. Bucci and G. Mazzarella, “Reconfigurable arrays by phase-only control”, *IEEE Transactions on Antennas and Propagation*, vol. 39, no. 7, pp. 919–925, 1991.
- [63] C. Fulton, J. L. Salazar, Y. Zhang, G. Zhang, R. Kelly, J. Meier, M. McCord, D. Schmidt, A. D. Byrd, L. M. Bhowmik, S. Karimkashi, D. S. Zrnic, R. J. Doviak, A. Zahrai, M. Yearly, and R. D. Palmer, “Cylindrical Polarimetric Phased Array Radar: Beamforming and Calibration for Weather Applications”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 5, pp. 2827–2841, 2017. DOI: 10.1109/TGRS.2017.2655023.

- [64] R. Gold, “Optimal binary sequences for spread spectrum multiplexing (Corresp.)”, *IEEE Transactions on Information Theory*, vol. 13, no. 4, pp. 619–621, 1967. DOI: 10.1109/TIT.1967.1054048.
- [65] B. L. Cheong, R. D. Palmer, and M. Xue, “A time series weather radar simulator based on high-resolution atmospheric models”, *Journal of Atmospheric and Oceanic Technology*, vol. 25, no. 2, pp. 230–243, 2008.
- [66] A. D. Byrd, I. R. Ivic, R. D. Palmer, B. M. Isom, B. L. Cheong, A. D. Schenkman, and M. Xue, “A Weather Radar Simulator for the Evaluation of Polarimetric Phased Array Performance”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 7, pp. 4178–4189, 2016, ISSN: 0196-2892. DOI: 10.1109/TGRS.2016.2538179.
- [67] L. H. Doherty, *Z-R Relationships Deduced from Forward Scatter Doppler Measurements*, 1964. DOI: 10.1175/1520-0469(1964)021<0683:rrdffs>2.0.co;2.
- [68] A. Ludwig, “The definition of cross polarization”, *IEEE Transactions on Antennas and Propagation*, vol. 21, no. 1, pp. 116–119, 1973.
- [69] R. J. Doviak, D. S. Zrnica, J. Carter, A. V. Ryzhkov, S. M. Torres, and A. Zahrai, *Polarimetric Upgrades to Improve Rainfall Measurements*, 1998. [Online]. Available: https://www.nssl.noaa.gov/publications/wsr88d%7B%5C_%7Dreports/2pol%7B%5C_%7Dupgrades.pdf.
- [70] W. Imbriale, “Reflector Antennas”, in *Modern Antenna Handbook*, John Wiley & Sons, Inc., 2008, ch. 5, pp. 201–263.
- [71] D. A. Brunkow, J. George, V. N. Bringi, and V. Chandrasekar, “Recent data system and antenna upgrades to the CSU-CHILL radar”, *32nd Conference on Radar Meteorology*, pp. 1989–1994, 2005.

- [72] M. Weber and J. Herd, “Aperture Size Considerations for Future Operational Phased Array Weather Radar”, in *35th Conference on Environmental Information Processing Technology*, Phoenix, AZ: Amer. Meteor. Soc., 2019.
- [73] D. J. Bodine and K. L. Rasmussen, “Evolution of mesoscale convective system organizational structure and convective line propagation”, *Monthly Weather Review*, vol. 145, no. 9, pp. 3419–3440, 2017, ISSN: 15200493. DOI: 10.1175/MWR-D-16-0406.1.
- [74] B. Isom, R. Palmer, R. Kelley, J. Meier, D. Bodine, M. Yearly, B. L. Cheong, Y. Zhang, T. Y. Yu, and M. I. Biggerstaff, “The Atmospheric Imaging Radar: Simultaneous volumetric observations using a phased array weather radar”, *Journal of Atmospheric and Oceanic Technology*, vol. 30, no. 4, pp. 655–675, 2013.

Appendix A

Derivation of the Relationship Between Scatterer Velocity and Multistatic Velocity Measurements

Begin by assuming a three-dimensional Cartesian coordinate system with a scatterer positioned at the origin, and a bistatic transmitter and receiver each positioned at arbitrary points. The transmitter and receiver locations can be described by the following set of position vectors:

$$\mathbf{r}_r = \begin{bmatrix} x_r \\ y_r \\ z_r \end{bmatrix}, \quad (\text{A.1})$$

$$\mathbf{r}_t = \begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix}, \quad (\text{A.2})$$

where \mathbf{r}_r corresponds to the receiver and \mathbf{r}_t corresponds to the transmitter. The scatterer has a velocity vector \mathbf{v} :

$$\mathbf{v} = \begin{bmatrix} u \\ v \\ w \end{bmatrix}, \quad (\text{A.3})$$

where u, v , and w are the components of the scatterer velocity along the x, y , and z directions respectively. The Doppler velocity measured at the receiver v_{bi} is the negative scalar projection of \mathbf{v} onto a bistatic bisector, which is some vector bisecting the angle between \mathbf{r}_r and \mathbf{r}_t . This can be expressed as:

$$v_{bi} = -\hat{\mathbf{b}} \cdot \mathbf{v}, \quad (\text{A.4})$$

where $\hat{\mathbf{b}}$ is the unit magnitude bistatic bisector. Note that the negative sign is necessary to make the result consistent with the standard sign convention for velocity in weather radar (a positive sign corresponding to a positive change in propagation path length). $\hat{\mathbf{b}}$ can be calculated as follows:

$$\hat{\mathbf{b}} = \frac{\mathbf{b}}{\|\mathbf{b}\|}, \quad (\text{A.5})$$

$$\mathbf{b} = R_T \mathbf{r}_r + R_R \mathbf{r}_t, \quad (\text{A.6})$$

$$R_T = \|\mathbf{r}_t\|, \quad (\text{A.7})$$

$$R_R = \|\mathbf{r}_r\|. \quad (\text{A.8})$$

The expression for \mathbf{b} can be expanded:

$$\mathbf{b} = \begin{bmatrix} R_T x_r + R_R x_t \\ R_T y_r + R_R y_t \\ R_T z_r + R_R z_t \end{bmatrix}, \quad (\text{A.9})$$

which, in turn, gives the following expansion for $\|\mathbf{b}\|$:

$$\|\mathbf{b}\| = \sqrt{R_T^2(x_r^2 + y_r^2 + z_r^2) + 2R_T R_R(x_r x_t + y_r y_t + z_r z_t) + R_R^2(x_t^2 + y_t^2 + z_t^2)} \quad (\text{A.10})$$

$$= \sqrt{R_T^2 R_R^2 + 2R_T R_R(\mathbf{r}_r \cdot \mathbf{r}_t) + R_R^2 R_T^2} \quad (\text{A.11})$$

$$= \sqrt{2R_T^2 R_R^2 + 2R_T^2 R_R^2 \cos \beta} \quad (\text{A.12})$$

$$= 2R_T R_R \sqrt{\frac{1 + \cos \beta}{2}} \quad (\text{A.13})$$

$$= 2R_T R_R \cos\left(\frac{\beta}{2}\right), \quad (\text{A.14})$$

$$\beta = \cos^{-1}\left(\frac{\mathbf{r}_t \cdot \mathbf{r}_r}{R_T R_R}\right). \quad (\text{A.15})$$

Note that β is the angle between \mathbf{r}_t and \mathbf{r}_r , commonly referred to as the bistatic angle. Substituting (A.9) and (A.14) into (A.5) gives:

$$\hat{\mathbf{b}} = \frac{1}{2R_T R_R \cos(\beta/2)} \begin{bmatrix} R_T x_r + R_R x_t \\ R_T y_r + R_R y_t \\ R_T z_r + R_R z_t \end{bmatrix}, \quad (\text{A.16})$$

$$= \frac{1}{2 \cos(\beta/2)} \begin{bmatrix} \frac{x_r}{R_R} + \frac{x_t}{R_T} \\ \frac{y_r}{R_R} + \frac{y_t}{R_T} \\ \frac{z_r}{R_R} + \frac{z_t}{R_T} \end{bmatrix}. \quad (\text{A.17})$$

Next, consider the relationship between the Cartesian coordinates that have been used thus far and the corresponding spherical coordinate system:

$$x = r \sin \theta \cos \phi, \quad (\text{A.18})$$

$$y = r \sin \theta \sin \phi, \quad (\text{A.19})$$

$$z = r \cos \theta, \quad (\text{A.20})$$

$$r = \|\mathbf{r}\|, \quad (\text{A.21})$$

$$\mathbf{r} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}. \quad (\text{A.22})$$

Minor rearrangement yields:

$$\frac{x}{r} = \sin \theta \cos \phi, \quad (\text{A.23})$$

$$\frac{y}{r} = \sin \theta \sin \phi, \quad (\text{A.24})$$

$$\frac{z}{r} = \cos \theta, \quad (\text{A.25})$$

$$(\text{A.26})$$

which, when substituted into (A.17), gives:

$$\hat{\mathbf{b}} = \frac{1}{2 \cos(\beta/2)} \begin{bmatrix} \sin \theta_R \cos \phi_R + \sin \theta_T \cos \phi_T \\ \sin \theta_R \sin \phi_R + \sin \theta_T \sin \phi_T \\ \cos \theta_R + \cos \theta_T \end{bmatrix}, \quad (\text{A.27})$$

where θ_R, ϕ_R are the spherical coordinate angles corresponding to the receiver position and θ_T, ϕ_T are the spherical coordinate angles corresponding to the transmitter position.

It is desirable to next perform a change of variables such that the equation is written in terms of the scatterer position relative to the transmitter and receiver, rather than in terms of the transmitter and receiver positions relative to the scatterer. This conversion has two steps. The first is conversion to spherical coordinates referenced to the transmitter / receiver. The second is conversion of the result from spherical into azimuth / elevation format. Given the scatterer-referenced spherical coordinates θ, ϕ of some object, the object-referenced spherical coordinates of the scatterer θ', ϕ' are calculated:

$$\theta' = \pi - \theta, \quad (\text{A.28})$$

$$\phi' = \phi + \pi. \quad (\text{A.29})$$

This is obtained by assuming a linear translation of the coordinate system such that its origin is translated from the scatterer to the object in question. These coordinates must then be converted to azimuth / elevation angles. The elevation angle e is defined as the angle from the xy plane in the object-centered coordinate system to the scatterer. The azimuth angle a is the angle moving clockwise from the y axis to the scatterer's xy position (corresponding to the scatterer's compass bearing, assuming the y axis points due north). These angles are calculated:

$$e = \frac{\pi}{2} - \theta' \quad (\text{A.30})$$

$$= \theta - \frac{\pi}{2}, \quad (\text{A.31})$$

$$a = \frac{\pi}{2} - \phi' \quad (\text{A.32})$$

$$= -\phi - \frac{\pi}{2}. \quad (\text{A.33})$$

$$(\text{A.34})$$

This allows the following substitutions for the trigonometric functions in (A.27):

$$\cos \theta = -\sin(e), \quad (\text{A.35})$$

$$\sin \theta = \cos(e), \quad (\text{A.36})$$

$$\cos \phi = -\sin(a), \quad (\text{A.37})$$

$$\sin \phi = -\cos(a), \quad (\text{A.38})$$

yielding:

$$\hat{\mathbf{b}} = \frac{1}{2 \cos(\beta/2)} \begin{bmatrix} -\cos(e_R) \sin(a_R) - \cos(e_T) \sin(a_T) \\ -\cos(e_R) \cos(a_R) - \cos(e_T) \cos(a_T) \\ -\sin(e_R) - \sin(e_T) \end{bmatrix}, \quad (\text{A.39})$$

where e_R, a_R are the elevation and azimuth angles of the scatterer with respect to the receiver and e_T, a_T are the elevation and azimuth angles of the scatterer with respect to the transmitter.

Substitution of (A.3) and (A.39) into (A.5), and reformatting of the dot product as a matrix product yields:

$$\begin{bmatrix} \frac{\cos(e_R) \sin(a_R) + \cos(e_T) \sin(a_T)}{2 \cos(\beta/2)} & \frac{\cos(e_R) \cos(a_R) + \cos(e_T) \cos(a_T)}{2 \cos(\beta/2)} & \frac{\sin(e_R) + \sin(e_T)}{2 \cos(\beta/2)} \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = v_{bi} \quad (\text{A.40})$$

Using this expression and the fact that the corresponding result for the monostatic case can be obtained by setting $a_R = a_T$, $e_R = e_T$, and $\beta = 0$, the system of equations relating scatterer velocity to measured velocities for a multistatic network

is:

$$\begin{bmatrix} \frac{\cos(e_1) \sin(a_1) + \cos(e_T) \sin(a_T)}{2 \cos(\beta_1/2)} & \frac{\cos(e_1) \cos(a_1) + \cos(e_T) \cos(a_T)}{2 \cos(\beta_1/2)} & \frac{\sin(e_1) + \sin(e_T)}{2 \cos(\beta_1/2)} \\ \frac{\cos(e_2) \sin(a_2) + \cos(e_T) \sin(a_T)}{2 \cos(\beta_2/2)} & \frac{\cos(e_2) \cos(a_2) + \cos(e_T) \cos(a_T)}{2 \cos(\beta_2/2)} & \frac{\sin(e_2) + \sin(e_T)}{2 \cos(\beta_2/2)} \\ \vdots & \vdots & \vdots \\ \frac{\sin(a_N) \cos(e_N) + \sin(a_T) \cos(e_T)}{2 \cos(\beta_N/2)} & \frac{\cos(a_N) \cos(e_N) + \cos(a_T) \cos(e_T)}{2 \cos(\beta_N/2)} & \frac{\sin(e_N) + \sin(e_T)}{2 \cos(\beta_N/2)} \\ \cos(e_T) \sin(a_T) & \cos(e_T) \cos(a_T) & \sin(e_T) \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} v_{\text{bi}1} \\ v_{\text{bi}2} \\ \vdots \\ v_{\text{bi}N} \\ v_r \end{bmatrix}, \quad (\text{A.41})$$

where $v_{\text{bi}n}$ are the velocities measured by the n th receiver along its bistatic bisector, a_n and e_n are the azimuth and elevation of the observation location relative to the n th receiver, β_n is the bistatic angle corresponding to the n th receiver, v_r is the measured radial velocity at the transmitting radar, and N is the total number of receivers in the network. Note that this result differs consequentially from that presented by Wurman et al. [22]; (3) of that paper has erroneously omitted all $\cos(\beta_n/2)$ terms.

Appendix B

List of Symbols

A_e	Effective antenna area
a_R	Azimuth relative to receiver
a_T	Azimuth relative to transmitter
B	Bandwidth
C	Matrix relating excitations to beampattern
c_u	Horizontal whitening code
c_v	Vertical whitening code
F	Vector of beampattern values
d	Distance from grid point
D	Transmitter Duty Cycle
e_R	Elevation relative to receiver
e_T	Elevation relative to transmitter
f	Error function slope
f_d	Doppler frequency
f_R	Normalized receive antenna pattern
f_T	Normalized transmit antenna pattern
F	Vector of beampattern values

G_R	Receive antenna gain
G_T	Transmit antenna gain
\mathbf{G}	Adjusted beampattern vector
j	$\sqrt{-1}$
L	Transmitter-to-receiver baseline length
\mathbf{M}	Pattern synthesis mask vector
N_0	Receiver Noise Power
p	Angle from transmit beam to baseline
P_R	Received power
P_S	Scattered power
P_T	Transmit power
S_R	Receive power density
S_T	Transmit power density
R_i	Radius of influence
$\hat{R}(T_s)$	Estimate of time series autocorrelation
R_B	Bistatic range
R_R	Transmitter-to-receiver range
R_T	Transmitter-to-scatterer range
r	Range
t	Time
T_s	PRT
T_{rx}	Receiver noise temperature
T_a	Antenna noise temperature
u	Zonal component of velocity
V	Volume
$V[n]$	Received time series
v	Meridional component of velocity

v_{bi}	Bistatic velocity
v_r	Radial velocity
w	Vertical component of velocity
w_{rx}	Rx component of SC weight
w_{tx}	Tx component of SC weight
w	Vertical component of velocity
$W(d)$	Cressman weighting function
$W(r)$	Range weighting function
W	Weighting vector
x_p	Scatterer x coordinate
x	Vector of excitation values
y_p	Scatterer y coordinate
z_p	Scatterer z coordinate
α	Weight tuning factor
β	Bistatic angle
β_w	Pulse bandwidth
γ	Polarization loss factor
ΔR	Range resolution
Δr	Minimum/monostatic range resolution
Δt	Time delay from pulse transmission to reception
δ	Angle between scatterer velocity vector and bistatic bisector
$\delta[n]$	Pulse position error function
ϵ	Convergence criterion
η	Reflectivity

θ_R	Spherical coordinate θ relative to receiver
θ_T	Spherical coordinate θ relative to transmitter
ϕ_R	Spherical coordinate ϕ relative to receiver
ϕ_T	Spherical coordinate ϕ relative to transmitter
λ	Operating wavelength
σ	RCS
τ	Temporal pulse length
ω_0	PRF

Appendix C

List of Acronyms and Abbreviations

ARRC	Advanced Radar Research Center
ATD	Advanced Technology Demonstrator
CFAR	Constant False Alarm Rate
GPS	Global Positioning System
LO	Local Oscillator
MCS	Mesoscale Convective System
NWP	Numerical Weather Prediction
OUHSC	University of Oklahoma Health Science Center
PAIR	Polarimetric Atmospheric Imaging Radar
PAWR	Phased Array Weather Radar
PRF	Pulse Repetition Frequency
PRT	Pulse Repetition Time
RCS	Radar Cross-Section
RIL	Radar Innovations Laboratory
SNR	Signal-to-Noise Ratio
SZ	Sachidananda/Zrnić

ULA

Uniform Linear Array

VAD

Velocity Azimuth Display

WSR-88D

Weather Surveillance Radar - 1988

Doppler