©Copyright 2013
Laura Vertatschitsch

High Bandwidth, Multi-Purpose Passive Radar Receiver Design For Aerospace and Geoscience Targets

Laura Vertatschitsch

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

John D. Sahr, Chair

Miguel F. Morales

Lih Y. Lin

Brian Nelson

James J. Clauss

Program Authorized to Offer Degree:
UW Electrical Engineering

University of Washington

Abstract

High Bandwidth, Multi-Purpose Passive Radar Receiver Design For Aerospace and Geoscience Targets

Laura Vertatschitsch

Chair of the Supervisory Committee:
Associate Dean, Undergraduate Academic Affairs and Professor John D. Sahr
UW Electrical Engineering

Passive radar permits inexpensive and stealthy detection and tracking of aerospace and geoscience targets. Transmitters of opportunity such as commercial FM broadcast, DTV broadcast, and cell phone towers are already illuminating many populated areas with continuous power. Passive radar receivers can be located at a distance from the transmitter, and can sense this direct transmission as well as any reflections from ground clutter, aircraft, ionospheric turbulence and meteor trails. The 100% duty cycle allows for long coherent integration, increasing the sensitivity of these instruments greatly. Traditional radar receivers employ analog front end downconverters to translate the radio frequency spectrum to an intermediate frequency (IF) for sampling and signal processing. Such downconverters limit the spectrum available for study, and can introduce nonlinearities which limit the detectability of weak signals in the presence of strong signals. With suitably fast digitizers one can bypass the downconversion stage completely. Very fast digitizers may have relatively few bits, but precision is recovered in subsequent signal processing. We present a new passive radar receiver designed to utilize a broad spectrum of commercial transmitters without the use of a front end analog downconverter. The receiver centers around a Reconfigurable Open Architecture Computing Hardware (ROACH) board developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) group. Fast sampling rates (8-bit samples as fast as 3 GSps) combined with 640 multiply/addition

operations on the Virtex-5 FPGA centered on the ROACH allows for coherent processing of broad spectrum and dynamic decision-making on one device all while sharing a single front end, putting this device on the cutting edge of wideband receiver technology. The radar is also designed to support mobile operation. It fits within a 19" rack, it is equipped with solid state hard drives, and can run off an uninterruptible power supply (UPS) for up to 1 hour of continuous operation. In this document we provide technical details of the hardware, firmware, and software of the system and design strategies and decisions. We cover the topic of coherent processing for passive radar, specifically an overview of the cross-ambiguity function as a detection mechanism. While the applications of a system like this are incredibly broad, the initial validation and performance analysis was applied specifically to detection of aircraft using Digital Television (DTV) broadcast as an illuminator. We present results of both stationary and mobile operation. In stationary operation, the same helicopter has been detected using two different DTV transmissions. Early mobile operation results show the Doppler-spread ground clutter and possible detection of aircraft. In addition to the fully-functional aircraft detection signal chain, alternative FPGA designs are presented with modes for fast sampling on two antennas or four antennas, with access to an aggregate 240 MHz of spectrum, with 8-bit samples. At these extremely high data rates, moderate data loss occurs while saving this data to disk, but as detailed within this document, it can be accounted for and the effects minimalized, still allowing for detection of aircraft. With these modes, FM transmission and DTV transmission can be captured synchronously from a single antenna and digitizer feed, an exciting result that offers promise for both aerospace and geoscience applications.

TABLE OF CONTENTS

Page

List of Figures iv

List of Tables xiii

Glossary xvi

Chapter 1: Introduction, Background Research and Literature Review 2
1.1 Radar, Passive Radar Background 5
1.2 Bistatic Radar Geometry 9
1.3 Passive Radar Systems and Applications 23
1.4 Focus of this work 27

Chapter 2: Performance Prediction 31
2.1 The SCR and Dynamic Range 31
2.2 Self-Ambiguity of the Radar Waveform 34
2.3 Cross-Ambiguity Calculations 42

Chapter 3: System Design and Specifications 45
3.1 Typical Passive Radar System design: ISIS and the MRR 45
3.2 Current System Design 48
3.3 FPGA Processing 50
3.4 IF Processing 55
3.5 Precision Gain, Processing Gain, and Dynamic Range Considerations in System Design 59

Chapter 4: Usability 73
4.1 FPGA Design, Simulation, and Compilation 73
4.2 Programming the FPGA and interacting with the ROACH board 75
4.3 Initialization: Front End and ADC control 76
4.4 Control center 81

4.5 Diagnosing the FPGA Output 84
4.6 Mobile Operation Considerations 88

Chapter 5: FPGA Digital Receiver Modes 91
5.1 FPGA Resources 92
5.2 Polyphase Decimation, 1.6 GSamp/s, Two Antennas, Two 6 MHz spectral windows (2 DTV stations per antenna) 93
5.3 Polyphase Decimation, 800 MSamp/s, Two Antennas, One 6 MHz spectral windows (1 DTV station per antenna), 8-bit IQ 95
5.4 Digital mixer, 1.6 GSamp/s, two antennas, two 6 MHz spectral windows (2 DTV stations per antenna) 8-bit IQ 97
5.5 Four antennas, 800 MSamp/s, one 6 MHz spectral window, 8-bit IQ 99
5.6 Four antennas, 800 MSamp/s, one 60 MHz spectral window, 8-bit IQ 101
5.7 Two antennas, 1.6 GSamp/s, two 60 MHz spectral windows, 8-bit IQ 103

Chapter 6: High Aggregate Bandwidth Challenges 106
6.1 Introduction to current data rates and data transfer 107
6.2 Characterization Of Data Loss 114
6.3 Simulation and Analysis of Data Loss 120
6.4 Strategies for Data Loss at High Bandwidth 143

Chapter 7: Aircraft Detection 145
7.1 SCR Performance Specification 146
7.2 When SNR matters 150
7.3 DTV-based aircraft detection: Stationary Results with two transmissions . . 158
7.4 DTV-based aircraft detection: Mobile Results with two transmitters 164
7.5 High Aggregate Bandwidth: Four antennas, Two DTV stations 169
7.6 DTV and FM Broadcast on the Same Antenna 184

Chapter 8: Future Experimentation with this system 186
8.1 Hybrid Transmitter Detection of Aircraft 186
8.2 Testing New Coherent Processing Strategies 186
8.3 Auroral Observations 187
8.4 Considerations for Wi-Fi or GPS transmitter applications 188

Chapter 9: Conclusions 191

Bibliography 193

Appendix A: Polyphase Decimation 198

Appendix B: Animation of Detections 202

Appendix C: Signal Modelling and Cross-Ambiguity Calculations 203

LIST OF FIGURES

Figure Number Page

1.1 Example of monostatic radar, with transmitter and receiver located at A, and bistatic radar, with transmitter at A, receiver at B. The airplane will be detected at a certain range. For the monostatic case, this is anywhere on a sphere centered at A; the footprint of this sphere on land is shown as the blue circle. For the bistatic case, the plane is anywhere on an ellipsoid with one focus at A and the other at B. The footprint of this ellipsoid is shown as the red ellipse. The plane is located on the arc defined by the intersection of the two surfaces. 11

1.2 Details of the bistatic geometry associated with target detection. We define the velocity $\hat{(v)}$ of a target, as well as the bistatic angle β seen via the angle bisector value $\beta/2$, and the angle between the velocity and this angle bisector, δ. Also seen are the various ranges involved with the detection. 12

1.3 A 2-dimensional map displaying a bistatic radar (transmitter and receiver displayed as x's), with contours of constant bistatic path difference, a simple variant of "range" which is more intuitive for the bistatic geometry. Contours close to the focii appear elliptical, while contours further away are already approaching their circular limit. 13

1.4 The same 2-dimensional map and geometry of Figure 1.3, but with the specific and common case of a target located at a bistatic path difference of 10 km. The transmitter and receiver are 15 km apart. The region within 10 km bistatic path difference is indicated with purple. 14

1.5 Direction of maximum doppler shift, shown pointing in the *negative* direction, or "away" from the radar system overlayed with Figure 1.3, a contour plot of constant bistatic path difference. 16

1.6 Contour plot of received power for the same bistatic radar system plotted in Figure 1.3. These are Ovals of Cassini. 18

1.7 An example of relative received signal power as compared to noise (or clutter). This particular example illustrates a threshold ratio of -85 dB. Perhaps our particular receiver can be characterized by sensing every target with a signal-to-clutter (SCR) ratio greater than -85. The region in purple shows the location (projected down to 2 dimensions) of all visible target locations. . . . 21

1.8 This shows an aircraft flying at 200 m/s at a particular instance in time (purple filled triangle, past positions empty triangles) as seen on a 2-dimensional map in plot (a) and as measured by the bistatic radar on a range-Doppler plot in plot (b). In both cases, the full path of the aircraft over time is shown by a black dotted line. 22

1.9 Map of the stationary receiver location, marked RX, at Sieg hall at the University of Washington (UW) and various DTV transmitters, marked TX, in the greater Seattle area. 29

2.1 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth. We can note the long correlation time by the large power horizontal 0 doppler line. 36

2.2 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth. By removing a larget deterministic component we see an expected drop of the horizontal 0 doppler line, improving the self-ambiguity. 37

2.3 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth, isolating the 0 doppler slice (simple fft, no windowing). We see the correlation remain quite high here due to the deterministic component of the spectrum rather than legitimate ground clutter reflections. 38

2.4 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth, isolating the 0 doppler slice (simple fft, no windowing). We see the correlation drop significantly. It is important to note that what we expect to see here is the true self-ambiguity of the signal as well as any stationary ground clutter reflections or moving targets with 0 doppler shift. By removing a large deteministic component we can tell that this is truly an improvement that will help us detect 0 doppler targets. 39

2.5 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth, isolating the 0 range slice. We see the frequency correlation drop off rather quickly, showing a rather nice self-ambiguity property without modification. 40

2.6 Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth, isolating the 0 range slice. We see the frequency correlation changes little by using this subset. . . 41

3.1 An example radar receiver that receives data on an antenna, mixes the data down to an intermediate frequency with an image reject filter and analog mixer pair, applies an anti-aliasing filter and samples the signal at IF with the Analog-to-Digital converter (ADC), and finally applies a digital down-converter to mix the sampled IF signal down to IQ and minimize the data rate. Passive radars often do not require signal amplification, but other radar receivers may include a low-noise amplifier (LNA) behind the antenna. 46

3.2 A particular example: the ISIS passive radar system. DTV processing (top) is similar to typical radar receiver design explored in Figure 3.1, however the FM processing avoids the analog mixing stage and digitizes the filtered FM broadcast directly. 47

3.3 Detailed block diagram of the new receiver architecture for DTV signal reception. 49

3.4 Block diagram of the processing on the Xilinx Virtex-5 FPGA. The dashed rectangle indicates the processes of conversion to baseband IQ data. The ROACH board completes the first few operations, while our control computer Thorex absorbs the data over a 10GbE link. The two machines together comprise the mobile portion of the signal processing. 54

3.5 Spur at one fourth the sampling rate in the ADC interleaved sampling mode. With a sampling rate of 1.6 GHz, we display the FFT from 0-800 MHz. The FM spectrum is visible around 100 MHz, a large spur is visible at 400 MHz, and the DTV spectrum can be seen through the front end band-pass filter from 500-626 MHz. 61

3.6 Block Diagram noting important stages for dynamic range gains (A-E) in the FPGA processing. Stages B-D are discussed in Table 3.2 67

3.7 Block Diagram noting important stages (F-K) of floating point signal processing from the stored FPGA samples to IQ data. Stages F, G, and H are analyzed in Table 3.3. 69

4.1 Output of initialization script to connect to the ROACH board, program the FPGA, calibrate the ADCs and iteratively setting the programmable attenuators. 77

4.2 Histograms of 1.6 GHz samples in interleaved sampling mode for each of our two KATADC cards. The data shows the expected Gaussian amplitude distribution. Without attenuation we are achieving the optimal signal strength into the ADC when the antennas are located on the roof of Sieg Hall. 78

4.3 Histograms of 800 MHz samples in non-interleaved sampling mode for each of our two KATADC cards. Like the 1.6 GHZ interleaved sampling data in Figure 4.2, the data shows the expected Gaussian amplitude distribution. Without attenuation we are achieving the optimal signal strength into the ADC when the antennas are located on the roof of Sieg Hall. 79

4.4 Live, real-time data from radar operation. One-sided spectrum of 1.6 GHz samples in interleaved sampling mode. We see roughly 6 DTV stations falling within our 500-626 MHz front end band-pass filter. We can also see the presence of spurs, notably one at 400 MHz. 79

4.5 Live, real-time data from radar operation. One-sided spectrum of 800 MHz samples in non-interleaved sampling mode. We see the 6 DTV stations within the 500-626 MHz passband of the front end filter aliased due to the under-sampling. 80

4.6 Snapshot from Simulink model file showing digital gain and quantization on the FPGA. A 32-bit software register can be set any time after the FPGA is programmed to provide an integer between 0 and 15 and is presented to the `scale` input. After mixing and data rate processing, full precision samples (57 bits with binary point at 54) are processed in parallel with the Xilinx Scale operator (no hardware cost), applying a digital gain of 2^n where n is an integer between 0 and 15. The blocks marked with z^{-1} are delay blocks, inserting latency in our model to achieve timing closure. The Multiplexer block (shown as the sideways trapezoid labeled `Mux`) passes the scaled data selected by the user. Finally, a `cast` block will cast the data as an 8 bit number with binary point at 5, set to saturate the data on any overrun. . . . 83

4.7 Monitoring the raw ADC samples at 800 MSamp/s on 4 antennas. The FPGA is designed to isolate 60 MHz of this spectrum in 8-bit IQ samples centered between two DTV stations: KZJO at 539 MHz and KONG at 575 MHz. These stations show up here at 225 MHz and 261 MHz, respectively, from aliasing. 86

4.8 Diagnosis of 100 MHz (60 MHz passband) data saved to disk. The digital mixer on the FPGA is set to 557 MHz, centered between KZJO to the left at 539 MHz and KONG to the right at 575 MHz. The second row of graphs show histograms of the real part of data samples for diagnosing digital gain. . 87

5.1 2 antennas, 2 stations, 6 MHz with polyphase filtering 93

5.2 2 antennas, 1 stations, 800 MSamp/s, 6 MHz window with polyphase filtering 96

5.3 2 antennas, 1 stations, 1.6 GSamp/s, 6 MHz window with polyphase filtering 97

5.4 4 antennas, 1 stations, 800 MSamp/s, 6 MHz window with polyphase filtering 99

5.5 4 antennas, 800 MSamp/s, 4 60 MHz window with polyphase filtering 102

5.6 2 antennas, 1.6 GSamp/s, 4 60 MHz window with polyphase filtering 103

6.1 After channel selection processing on the FPGA, the IQ data streams from each of four antennas are packed together in a 64-bit word to enter the 10 GbE buffer. 109

6.2 A boolean clock signal, tx_valid, determines which 64-bit words (packed with IQ data from 4 antennas) enters the 10 GbE buffer for transmission. Another clock signal, tx_end_of_frame, signifies the last word in the UDP frame. 111

6.3 A snapshot of the Simulink environment showing the 10 GbE block in the various input signals, most notably the data input tx_data, and the clock signals tx_valid and tx_end_of_frame (eof) as noted in Figure 6.2. . . . 112

6.4 A Simulink scope plotting the simulated clock signals driving the 10 GbE buffer. The scope taps into the wires for tx_valid and tx_end_of_frame (eof) as noted in Figure 6.3. 113

6.5 Bursts of data loss between consecutive packets of standard rate (≈14 MHz) data. Three 10-second data sets are presented, with the set labeled "Data 2" (red) showing the valuable piece of information that occasionally no data loss will occur. "Data 1" (blue) and "Data 3" show a large burst of data loss after 6 seconds of continuous operation, allowing roughly 200,000 packets or 800 MB of data to be reliably saved. 115

6.6 Packet Loss determined by value of the 64-bit header received in the storage computer. UDP packet headers increment by one each sent packet, with the black line indicating the expected received header number. The blue, red, and green lines all represent real data taken on the same day, showing the actual header values received. The deviation is recorded as the percentage of data lost, and is traced for all three data sets at several instances in time, marked by vertical dashed lines. These numbers are presented as integers relative to the first header value received, for comparison. After 2.6 seconds, all three data sets show aggregate data loss around 35%. 117

6.7 Bursts of data loss between consecutive packets of fast rate (100 MHz) data. Three 2.6-second data sets are presented, showing the first large burst of data loss around 200,000 packets in. 118

6.8 Histogram of amount of consecutive data loss on a log scale, x-axis, with the number of occurances on the y-axis. Still most common is small, 1 or 2 packet loss. 119

6.9 Diagram of packet loss experiment. Data is lost in whole packets at the fast, 14 MHz rate. The data loss produces two streams of data: one where the packet loss is known and represented as packets of zeros, and one where the lost packets are removed entirely creating a scenario much like processing data with unknown losses. 121

6.10 This is the cross-ambiguity calculation for a simulated DTV signal (band-limited white noise) and scatter from target with SCR = -10 dB. We will compare the affects of processing data with packet loss to this baseline. Here we see spreading in range due to the auto-correlation of anti-aliasing filters used to process the data from the fast rate of the 10 GbE link down to final IQ data rates. The target at 3 km and 100 m/s shows vertical spreading in power due to a doppler shift not perfectly aligned with the doppler samples. . 131

6.11 This is the cross-ambiguity calculation for a simulated DTV signal (band-limited white noise) and scatter from target with SCR = -10 dB with the direct path removed by our simple adaptive nulling. Any detection, including the direct path, contributes power across the entire 2-dimensional plot. Removing the direct path successfully shows not only is the peak at 0,0 completely removed, but the range spreading features from the direct path and the contribution to the level of the detection floor are removed as well. Pushing the detection floor down results in a detection floor linked directly to the target returned power. This allows the doppler and range spreading features associated with the processing of the target return to be more visible than Figure 6.10. 132

6.12 Bernoulli trial for 500 packets at a rate of 1% packet loss. This shows that packets 13, 21, 176, and 257 will be lost. 133

6.13 Data streams with a lost packet. The data is identical until roughly 6,000 samples at which time a packet (512 samples) of data is lost. The top plot represents unknown data loss, where these samples are simply missing. The bottom plot shows a data stream with zeros inserted in place of the missing packet. 134

6.14 Cross-ambiguity calculation for 1 of 10 trials with 1% unknown data loss. The only difference between this and the control simulation of Figure 6.10 is that the measured Doppler shift may have increased, noted by higher power in the adjacent Doppler bin above the main target peak. 135

6.15 Cross-ambiguity calculation for 1 of 10 trials with 1% known data loss corrected with zeros. There are no major differences visible between this and the control simulation of Figure 6.10. 136

6.16 Cross-ambiguity calculation for 1 of 10 trials with 5% unknown data loss. Without knowing data was dropped, there is nothing suspicious about this plot. However, the target peak occurs at a Doppler shift 5% higher than the true value. 137

6.17 Cross-ambiguity calculation for 1 of 10 trials with 5% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data. . . 138

6.18 Cross-ambiguity calculation for 1 of 10 trials with 10% unknown data loss. Without knowing data loss has occurred, there is nothing particularly suspicious about this plot. However, the target peak occurs at a Doppler shift 10% higher than the true value! 139

6.19 Cross-ambiguity calculation for 1 of 10 trials with 10% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data. The effect is stronger than in Figure 6.17. 140

6.20 Cross-ambiguity calculation for 1 of 10 trials with 50% unknown data loss. The target has smeared in Doppler space, showing that while roughly centered around 200 m/s, there are now many peaks along the vertical Doppler slice, rather than the tight thumbtack-like formation of the control simulation with no data loss. 141

6.21 Cross-ambiguity calculation for 1 of 10 trials with 10% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data. This is a significant feature; the Doppler line is consistently 40 dB below the peak. This implies that reflection from the target must be less than 40 dB below any other scatterer at that range. 142

7.1 SCR determined for a 0 dBsm aircraft, with a distance of 3 km to the receiver. As the receiver and transmitter are separated, this shows the SCR approaches a limit of -81 dB. Any distance greater than 4.5 km separation (not close to the monostatic limit!) will have an SCR above -85 dB. 149

7.2 Cross-ambiguity output for stationary operation of radar using KZJO with no attempt to null the direct path. We cannot detect any targets with confidence. 159

7.3 Cross-ambiguity output for stationary operation of radar using KZJO with attempt to null direct path. We null the direct path at (0,0) by 20 dB, reducing the floor and allowing targets to be detected clearly. Three aircraft are seen here, with the most interesting being the helicopter (seen by the doppler spread blade modulation) at close range. 160

7.4 Cross-ambiguity output for stationary operation of radar using KONG with no attempt to null the direct path. We cannot detect any targets with confidence. 161

7.5 Cross-ambiguity output for stationary operation of radar using KONG with attempt to null direct path. We null the direct path at (0,0) by 32 dB, reducing the floor by 17 dB and allowing targets to be detected clearly. Three aircraft are seen here, with the most interesting being the helicopter (seen by the doppler spread blade modulation) at close range. 162

7.6 Bistatic ellipses associated with simultaneous detection of helicopter from KONG and KZJO on September 12, 2012. The radar was operating in stationary mode with receiver at Sieg Hall. 163

7.7 Mobile operation with KZJO, driving 50-60 mph. Doppler spread ground clutter is striking at this fast speed and makes detecting moving targets difficult without adaptive filtering. 165

7.8 Mobile operation with KONG, driving 50-60 mph. Doppler spread ground clutter is striking at this fast speed and makes detecting moving targets difficult without adaptive filtering. 166

7.9 Mobile operation with KZJO, stop-and-go traffic as seen by the ground clutter remaining close to 0 doppler. We may have a possible detection of aircraft at a bistatic path difference of 0.9 m and doppler velocity -40 m/s. 167

7.10 Mobile operation with KONG, stop-and-go traffic as seen by the ground clutter remaining close to 0 doppler. We may have a possible detection of aircraft with small bistatic path difference close to -40 m/s doppler velocity. . 168

7.11 Diagnosis of FPGA output for an aircraft detection experiment with fast rate data transfer. The FPGA mixer is set to 557 MHz, directly between two DTV stations: KZJO (536-542) and KONG (572-578). The passband is shown from -30 MHz to 30 MHz. 170

7.12 Cross-Correlations for each of four antennas for KONG using the fast rate data transfer mode, processed without taking into account packet loss. A target is visible with measured Doppler velocity around 36 m/s, at range of 170 meters. We expect that the aircraft may have experienced some range migration, however it appears to be a concise peak in range, so this may indicate minimal losses. 173

7.13 This shows the next second of data associated with Figure 7.12. The bistatic path difference shrinks as expected from the positive Doppler velocity in the previous plot. 174

7.14 Cross-Correlations for each of four antennas for KZJO using the fast rate data transfer mode, processed without taking into account packet loss. . . . 175

7.15 This shows the next second of data associated with Figure 7.12. The target advances in range by roughly the amount associated with the Doppler velocity estimate from the first second. 176

7.16 Packet Loss for the first second of fast rate data on a log 10 scale. Fairly minimal packet loss for most of the duration. 177

7.17 Packet Loss for the second second of fast rate data on a log 10 scale. Many instances of bursty packet loss indicate that the measured range and Doppler velocity in Figures 7.13 and 7.15 could be inaccurate. 178

7.18 First second of KONG data, with packet loss corrected by inserting zeros. Minimal packet loss makes these results very similar to Figure 7.12, showing the true range and Doppler velocity of the target is 170 m at 32 m/s. The difference in the true Doppler velocity and the Doppler velocity measured previously is consistent with 12% packet loss for most of the processing. . . . 180

7.19 Part two of KONG data, with large amounts of packet loss corrected with zeros. The target is completely masked by the effects of the correction scheme.181

7.20 First second of KZJO data, with packet loss corrected by inserting zeros. Minimal packet loss makes these results very similar to Figure 7.14, showing the true range and Doppler velocity of the target is ¡75 m at -15 m/s. The difference in the true Doppler velocity and the Doppler velocity measured previously is consistent with 12% packet loss for most of the processing. . . . 182

7.21 Part two of KZJO data, with large amounts of packet loss corrected with zeros. The target is completely masked by the effects of the correction scheme.183

7.22 FFT of data saved to disk for two antennas, each digitized at 1.6 GSamp/s and split into two streams that isolate 6 MHz of the spectrum. This shows we are observing both FM and DTV from a single antenna input. 184

A.1 Block diagram of decimation. 198

A.2 Example of polyphase decimation as seen in Simulink. 200

A.3 Polyphase computation comparison to high rate decimation. For 160 taps, polyphase filtering only takes 319 computations per clock whereas straightforward decimation would take over 2500. 201

LIST OF TABLES

Table Number Page

1.1 Various transmitters available for passive radar use and their properties. Original table displayed in [20], presented here with extensions for GPS and Wi-Fi, as well as a rough calculation of direct path power received, P_c. 24

2.1 An FM and DTV transmitter comparison. 34

3.1 Flop count for xambi, single range slice. Correlation is the multiplication of the two received signals, $z[n, n_r] = y[n]x^*[n-n_r]$. Decimation is the coherent integration and reduction in data rate, $z_M[m, n_r]|_{m=Mn} = \sum_{k=0}^{M-1} z[Mn+k, n_r]$. Finally the FFT is computed to generate a single range slice. Multiplied by the number of ranges (120) gives the total count for 1 coherent integration (1.416 M), and multiplying by the number of incoherent averages (4) gives the total for 1 second of data processing. Typically, to process 1s of data in 1s of time, we need 5.7 GFlops. 58

3.2 Fixed-point SNR gain on the FPGA. The bandwidth (BW) referred to here represents the band of the original sampled signal that has been preserved. In stage D, after the rate reduction, aliasing has occured but only in areas of the spectrum that will later get filtered out, explaining the 28.57 MHz data rate compared to the quoted 32 MHz "bandwitdh". This is by design to use fewer filter coefficients. CNR and SNR gain are related to the BW, rather than the data rate (as long as aliasing has not occured in the band of the direct path clutter or scattered signal). We gain 2.8 bits of precision in SNR (and CNR) for a total of 17 dB improvement. 68

3.3 Floating point processing from IF to IQ at baseband. The total reduction of bandwidth gives us an SNR gain of 8 dB. 69

6.1 Disk space required to store IQ data compared to simply storing the cross-ambiguity result. The savings is a factor of 228. 107

6.2 Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target. 122

6.3 Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 1% unknown data loss. 124

6.4 Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 1% data loss corrected with packets of zeros. 124

6.5 Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 5% unknown data loss. The measured Doppler shift is precise, but biased by roughly 5 m/s in the positive direction. 125

6.6 Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 5% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth. 126

6.7 Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 10% unknown data loss. The measured Doppler shift is consistent across the 10 trials, but is nonetheless wrong by close to 10%. 126

6.8 Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 10% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements, due to incorporating zero for 10% of the data. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth. 127

6.9 Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 50% unknown data loss. As anticipated from past experiments, the Doppler shift is measured at double the true value. The measured Doppler shift is less consistent across the 10 trials as previous data loss experiments shown in Tables 6.3, 6.5, and 6.7. The power in the main target peak has dropped by more than anticipated, reflected in the SCR increase from ≈12 to 20 dB. 128

6.10 Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 50% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce close quantities, however the overall power has dropped in both the direct path and the target measurements, due to half the data set having value zero. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth. 129

7.1 Various transmitters available for passive radar use and the resulting SCR estimate as pertaining to detection of an aircraft-like target 3 km from the receiver. 146

7.2 Various transmitters available for passive radar use and the resulting Precision Gain based on initial sampling rate. 151
7.3 Various transmitters available for passive radar use and the resulting Processing Gain for an aircraft-like target. 152
7.4 Various transmitters available for passive radar analyzed for detection of an aircraft-like target. We analyze the SNR and CNR gain through stages of the radar processing. Including both Precision and Processing Gain, we explore the SNR and SCR values at several points in our processing scheme. We assume the front end attenuation is equal to 0 dB for all cases. 154
7.5 Various transmitters available for passive radar use and their properties as pertaining to detection of an aircraft-like target. Including both precision and processing gain, we determine for each system whether the waveform would result in a clutter-limited or quantization noise-limited scenario. 155
7.6 Using the system with an estimated 25 dB clutter suppression (direct path and ground clutter if possible). Aircraft detection with GSM cellphone transmission or Wi-Fi is now limited by quantization noise. 156

GLOSSARY

R_{tr}: Range from transmitter to receiver

R_{ts}: Range from transmitter to scatterer (target)

R_{sr}: Range from scatterer (target) to receiver

$\widehat{tr}$: Unit vector pointing from transmitter to receiver

$\widehat{ts}$: Unit vector pointing from transmitter to scatterer (target)

$\widehat{sr}$: Unit vector pointing from scatterer (target) to receiver

RCS: Radar Cross Section (m^2, dBsm), measured in units of area, this complicated term describes how electrically large a target appears to a specific radar system and geometry

SCR: Signal-to-Clutter Ratio (dB), ratio of signal power to clutter or unwanted received power, usually presented on a log scale

SNR: Signal-to-Noise Ratio (dB), ratio of signal power to noise power, usually presented on a log scale

DR: Dynamic Range (dB), a property of the radar receiver describing the ratio between the largest and smallest detections that can be made

RX: Shorthand for receiver

TX: Shorthand for transmitter

ACKNOWLEDGMENTS

I would like to acknowledge The Boeing Company for funding and collaboration on this work. Additionally, many sincere thanks to the CASPER community for their generous help interacting with the hardware. I'd like to also mention my lab mate Weiwei Sun, who provided much energy and support for this work.

DEDICATION

My time at the University of Washington has been profound. I'm very thankful to have a chance at the culmination here to honor the many people who have put me on this path.

To my parents, Ed and Michelle Vertatschitsch, who told me to never stop going to school. They offered immense support in the physical world (love, money, rides, coaching) but they ingrained in me a mentality of valuing education, specifically graduate school. This started so early I cannot remember a time I was not planning to try for a PhD, whatever those letters meant.

To my high school Physics teacher, Mr. Thomas Haff, whose many teaching awards can't tell the story of how deeply and profoundly he inspires his students to a lifetime of science love and learning. He took a bunch of teenagers usually more interested in passing notes[1] and showed them an unrivaled love of nerd culture. We left calling Carl Sagan our "homeboy."[2] I reorganized my high school course schedule to take three years of science and physics courses with Mr. Haff.

To my professors at Whitman College who put Cygwin[3] on my laptop while I worked late in the math lab, all so I could do my C++ homework on basketball roadtrips. Who reached out to me when I struggled with the material, who taught me how to learn, study, and rediscover as an adult, with conviction, how much joy I take in mathematics, especially when it's hard. Dr. Bob Fontenot, Dr. Fred Moore, and Dr. Albert Schueller, to name a special few.

This work is also dedicated to the spectacular staff, faculty, and culture at the University of Washington. I could add 10 dBpages to this document if I were to recall for you, the

[1] Kids, this was the original "text message"

[2] Kids, again, this is ancient slang, but I promise you it was very cool at the time

[3] http://www.cygwin.com/

reader, the many blessings afforded me during my tenure here.

Work with Dr. Eve Riskin involving outreach to local young women and educating them about engineering, rewarding in its own right, brought me to Cape Canaveral to see one of NASA's final space shuttle launches. I explored the great Southwest while collaborating with Dr. Sigrid Close and Dr. Pat Colestock at the historic Los Alamos National Laboratory researching tiny meteoroids and their affect on the earth, satellites, and science. It offered me a brand new perspective on life, quite literally, from 100 km altitude where human nature and its daily problems shrink just a little bit more. I've travelled to the far south looking for meteor trails in Tierra del Fuego, Argentina and the far north to learn about Incoherent Scatter Radar in Kangerlussuaq, Greenland. For the record, both icy.

Dr. John Sahr has led me to many of these opportunities, including one of the most rewarding experiences of my graduate school days: teaching. The quarters I spent lecturing were sweat-filled, sleepless, and eager; easily some of the most productive and enthusiastic work I've ever accomplished, where 80 hours a week flew by in exhausted fulfillment. Not all graduate students get this chance, and not all advisors would have given me the space that John did to sink into it, heart and soul. It may have added time to my degree, but it was time that I will cherish. Seeing former students at graduation this Spring gave me a joy of contributing to something much bigger than myself; to University with a capital "U." Thanks to Dr. Evan Goldstein and Dr. John Sahr for being excellent teachers, examples, mentors, and for freely sharing your time and advice.

For John, I could cover pages with thank yous and it would not be enough to express my gratitude. Thank you for your patience, dedication, and encouragement. Thank you for having such a fun, intricate, and accessible research project. I learned so much just by being excited to share what we do with others. Passive radar has all the makings of a good action film: space, airplanes, music (FM), television (DTV), explosions[4], and invisible electromagnetic waves! For those of you with a more technical taste, it also has

[4]Not truthfully, thankfully, but plenty of accidentally broken hardware and software, more than one segmentation fault, and mistaken deletion of much work. Yay version control!

all the makings of an excellent engineering project: Statistics, signal processing, hardware, firmware, software, earth science, design, creativity and invisible electromagnetic waves! It's wonderful to be able to take so much pleasure in my work; I can only hope this for the rest of my days.

Through engineering I have made deep and supportive relationships with friends who will celebrate each others' victories with a beer, attend each others' weddings, babysit each others' kids, and quiz each other for the qualifying exam. Most important of these friendships has been with Evan Hanusa. Thank you so much for your understanding, teaching, brain-storming, patience, and love. These relationships have made Electrical Engineering at the University of Washington a home.

The sketches that lead each chapter of the book are from my own wobbly hand and are of pictures I have taken or of items that have been close to my heart over my stay at the UW. In these sketches the reader will find latin phrases all loosely related to the idea of being a "careful listener," a warm and thoughtful little description of passive radar. I have to thank Dr. Jim Clauss of the Classics Department at the UW (my thesis committee General School Representative) for thoughtfully contributing these phrases, which bring a touch of heart to this technical document.

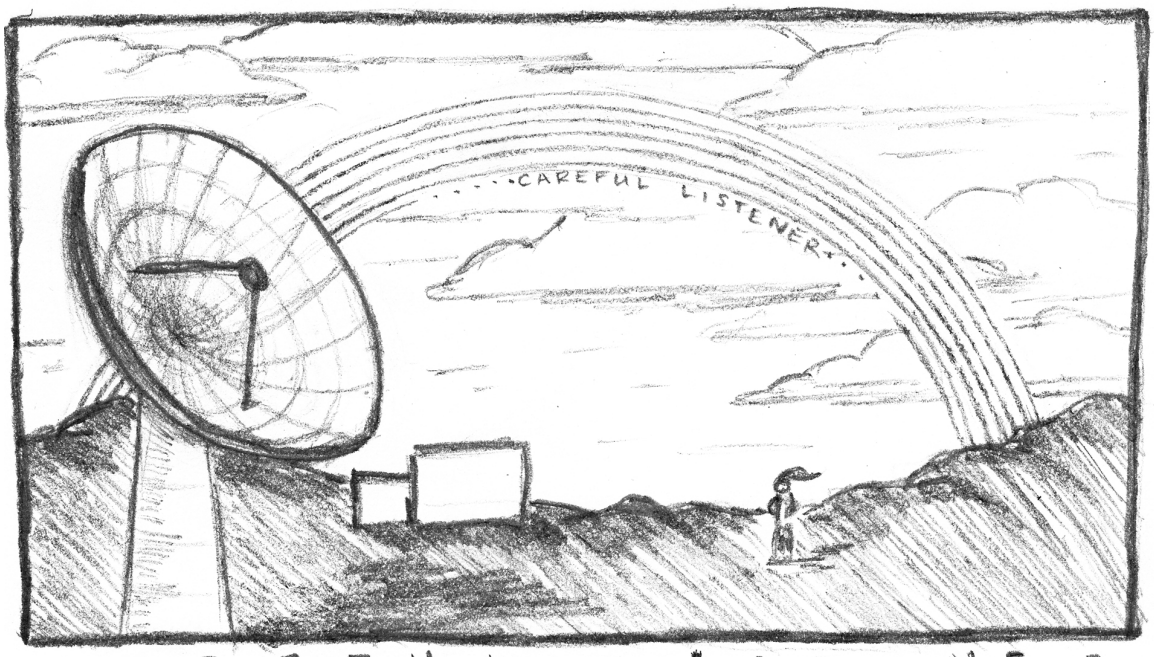
CAREFUL LISTENER
CAREFUL LISTENER

Chapter 1

INTRODUCTION, BACKGROUND RESEARCH AND LITERATURE REVIEW

Passive radar, also referred to as passive bistatic radar (PBR), permits inexpensive and stealthy detection and tracking of aerospace and geoscience targets. Unlike active radar, passive radar uses "transmitters of opportunity" or "non-cooperative transmitters" such as commercial FM broadcast, DTV broadcast, and cell phone transmission. The receiver and transmitter are not co-located and therfore form a bistatic geometry. This allows continuous illumination of targets and receivers can be designed with 100% duty cycle. While challenges arrise from bistatic geometry, this sensitivity enhancement is one of several advantages. There has been considerable interest in building a mobile passive radar for detection of moving targets. Passive radar receivers rely on much computation, and with current technology advancements such as fast digitizers, high bandwidth data links, and field-programmable gate arrays (FPGAs), a mobile passive radar is now feasible. We have created the next generation of the Manastash Ridge Radar (MRR), designed to look at a broad spectrum of transmitters (with the current analog-to-digital converters (ADCs) the Nyquist rate is as high as 1.5 GHz) and is designed for light weight, low cost, low power consumption, and mobile operation. This new system significantly increases the versatility of the MRR, increases the range resolution, and decreases the cost. The MRR has always been, and still is, a unique radar for the geoscience community, and this next generation receiver increases its usefulness to this community as well as advancing the state-of-the-art for aircraft detection with passive radar by exploring fast, direct sampling of the spectrum, omitting any use of an analog mixer. Presented here are technical details and considerations for the system as well as results from initial experimentation with DTV-illuminated aircraft detection. This introductory chapter will cover the literature, background, and basics of passive radar.

Chapter 2 discusses basic analysis of passive radar performance parameters and estimation. Unlike the monostatic, active radar systems with which the reader may be more familiar, single site passive radar systems must receive simultaneously both the very powerful transmitted signal as well as the small echoes from desirable targets. These systems are therefore limited by unwanted received power, rather than noise.

Chapter 3 describes the technical system details. Our key contribution is to demonstrate the efficacy of direct sampling through UHF without the use of analog down-conversion. The bulk of the initial signal processing is completed on an FPGA. The heart of the radar is the Reconfigurable Open Architecture Computing Hardware (ROACH) board[1] developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) group[2], which houses the Xilinx®Virtex-5 XC5VSX95T-1FF1136 FPGA[3]. Dynamic range is vital to these receivers and we discuss the design choices and system components that influence this parameter.

Chapter 4 covers the usability of this system. We describe the software environment used to design the FPGA-based receiver as well as how the user can interact, control and diagnose the radar at run-time or in the field.

Details of several useful designs are provided in Chapter 5. Polyphase decimation is compared to a more traditional digital down-conversion (DDC) scheme for DTV station isolation on the FPGA. Additionally, an FPGA mode is presented that isolates both DTV and FM broadcast from each RF input, a technological first.

In Chapter 6 we discuss the challenge of accessing and storing large windows of the spectrum. Data can be lost here, but if the loss is tracked accurately performance degradation can be kept to a minimum.

Chapter 7 details the initial successes of this receiver used as an aircraft detection system. The general parameters and strategies of our work are posed and specified for this particular application. We present aircraft detection for both mobile and stable operation. We also present the stationary detection of several different kinds of aircraft, including a helicopter

[1]https://casper.berkeley.edu/wiki/ROACH

[2]https://casper.berkeley.edu/

[3]http://www.xilinx.com/support/documentation/data_sheets/ds100.pdf

as noted by the blade modulation. Preliminary results of the road-vehicle mobile campaign, in addition to aircraft, show Doppler-spread ground clutter. Additionally we include early results of aircraft detection using fast data transfer and discuss the effects of packet loss on this detection.

In Chapter 8 we propose future experiments for the system to follow the success of the initial aircraft experiments. Suggested are further improvements to the aircraft detection processing including ground clutter suppression and tracking, improvements to mobile operation, and using the increased bandwidth of the MRR to enhance observations of auroral turbulence.

We begin with an in-depth background to passive radar, highlighting the attributes and defining the important features and terms that will govern the rest of our discussion.

1.1 Radar, Passive Radar Background

1.1.1 History

A nice background of passive radar is found in [32], and we use this source to summarize the history. Even in the early days of radar (1930s) it was clear that operation of the radar system with stealth was incredibly important. The first bistatic radar detection of aircraft was performed in 1935, just before WWII, in the UK by Sir Robert Watson-Watt using illumination from a shortwave BBC Empire transmitter. In 1943, the German 'Klein Heidelberg' receivers operated passively, using a series of British transmitters to detect aircraft. The stealth and resistance to hostile jamming made these radars incredibly useful, however the advent of the duplexer in 1936 allowed radars to operate a transmitter and receiver in one location, and the convenience of this caused monostatic, active radar technology to dominate until current day. Passive radar interest picked up in the 1980s, when European countries developed passive radar receivers that chased the rotating narrow-beam transmitters used at airports. Again in the 1990s a revival of passive radar technology grew from the ability of bistatic geometry to aid in detection of stealthy targets. Systems were developed in the UK, France, Germany and the US. Significant technology increases such as fast digitizers, fast processors, GPS, the growth and coverage of available transmitters, and the internet (for data transport) have allowed passive radar systems to become ever more useful.

Passive radar used for detection of geospace targets such as auroral turbulence began with the MRR, coming online in 1997 ([48], [34]), with a significant upgrade as part of the ISIS array in roughly 2006, first reported in [33].

1.1.2 Comparison to Active, Monostatic Radar

Passive radar maintains no control of authority of the transmitter, and for that reason there are disadvantages to this kind of radar system. Being bistatic in nature, detection of any target requires two clear propagation paths. Being passive, we are at the mercy of the transmitter. In the case of FM broadcast, the bandwidth changes as a function of the programming. Certain radio stations provide better radar waveforms due to sustained

high volume in the audio, which gives large instantaneous bandwidth to the transmitted waveform. For example, rock music outperforms talk radio (large pauses with low volume), shown in [19] and verified as well by our own observations in the Radar Remote Sensing Lab (RRSL). This same limitation extends to the radiation pattern of available transmitters. Commercial broadcasts concentrate their power toward the horizon and ground, and in the case of FM, DTV, and cell phones, are concentrated on land for their target customers. Not controlling the transmitter also means that we do not have access to a perfect matched filter, important for target detection in radar. Instead passive radar relies on making the cleanest recording of the transmitted signal possible, but this is always perturbed by atmospheric and propagation effects such as multipath. Recently, however, the advent of digital transmission such as DTV allows passive radars to reconstruct a more perfect matched filter through the error correction coding schemes as long as they have been provided to the public. Passive radars may require large instantaneous dynamic range in order to receive both the direct transmission and the small scatter from the target of interest.

There are several advantages resulting from the bistatic nature of passive radar. First, these receivers can be extremely sensitive. They operate with 100% duty cycle which gives high average power. As a radar for geoscience, this actually makes the MRR using VHF as sensitive as active instruments designed for similar purposes [48]. Without a transmit/receive switch, any preamp can be directly connected to the antenna, further increasing the sensitivity. Second, the receiver location is difficult to determine, and is therefore inherently resistant to hostile jamming, electronic counter measures, and physical attacks. Third, the bistatic nature allows for detection of stealth targets due to the forward scatter phenomenon which drastically increases the radar cross section (RCS) of a target. Finally, there is a lower cost of operation because it is unnecessary to shield the sensitive receiver electronics from a transmitter sharing the same antenna.

Geometry aside, operating a radar without a transmitter can offer its own advantages. First, these receivers can have access to powerful transmitters at important wavelengths without licensing. This gives access to 3 m wavelengths excellent for Bragg scatter from 1.5 m perturbations in the ionosphere. Stealth targets aim to reduce their RCS for higher frequencies, and thus the FM and DTV broadcasts at VHF and UHF may be able to detect

these targets more readily. Second, the entire radar system is stealthy because there is no transmitter to locate and destroy - unless, of course, war is waged on rock 'n' roll music and our enemies subject us to a one-two punch of talk radio (painful on the ears and on the range resolution!) Additionally, the long wavelengths associated with FM and DTV broadcast makes the transmitting antennas less vulnerable to "anti-radiation" missiles. Finally, with no transmitter operation of the radar is safe and much less costly.

1.1.3 A Passive Radar by Any Other Name...

Throughout its history, passive radar research has been done under many names. Quite recently, aerospace applications have made an effort to unify the research under a commonly used term passive bistatic radar (PBR) [39], [7].

Passive radars can also be referred to as passive coherent location (PCL) radar. Coherent location is not a technique that can be applied to overspread, underdense scatterers such as auroral echoes, thus the term "passive radar" is inclusive of these applications. Other names include hitchhikers, parasitic radar, and bistatic (or multistatic) radars utilizing transmitters of opportunity or non-cooperative transmitters [26].

Emphasizing their stealth features, passive covert radar (PCR) is yet another keyword used to describe these systems ([32]).

Jamming is the hostile and intentional broadcast of high-power noise in the receiver bandwidth to prevent a radar from efficient operation. Passive radars suffer from self-jamming, and the performance is limited not by the system noise but by this interference [53]. While not prominent in the literature, describing passive radar as a "self-jamming" radar is accurate. A more widely used term for this is direct path interference and will be addressed with detail in Chapter 2 and Chapter 3. Noise Radar is a technology that utilizes continuous band-limited noise as a waveform for its excellent range and Doppler resolution [35]. FM and DTV broadcast are similar to band-limited noise transmissions as seen in their ambiguity functions, and in fact we have described the Manastash Ridge Radar (MRR) as a distributed passive noise radar system as the fields share receiver signal processing technologies [38].

Passive radar is inherently bistatic, and can easily be implemented with multiple receivers (or transmitters) and requires appropriate fusion of measurements and tracking algorithms that will be found with multistatic sensor systems and data fusion research and conferences.

Passive radar is growing rapidly, and is a multi-faceted technology with dedicated and diverse research areas. The passive radar system designer must be a renaissance woman and must be prepared to collaborate and search diligently for improvements to the system that may be found in many diverse fields. Bistatic radars seem to have a 15-20 year cyclical resurgence in history [58], due to discovery of new applications or technological advancements that increase their performance. In 2000, the availability of GPS offered a way for multistatic receivers to synchronize. Today, our ability to deal with GSamp/s data rates offers a way of increasing the dynamic range and improve the frequency agility of the passive radar receiver, making it a more useful instrument.

1.2 Bistatic Radar Geometry

1.2.1 Bistatic Range

Bistatic radar involves a different detection geometry than monostatic radar. Figure 1.1 shows the same target detected by a monostatic system, location A, and a bistatic system using the transmitter at A and the receiver at B. A monostatic receiver can measure the range to a target, which places the target anywhere on a sphere of radius R centered at the combined transmitter/receiver location. A bistatic receiver also measures a range, placing the target anywhere on a bistatic ellipsoid with focii at the transmitter and receiver locations, as seen in Figure 1.2. This figure describes the range from the transmitter to the receiver, R_{tr}, the range from the transmitter to scatterer or target, R_{ts}, and the range from scatterer or target to the receiver, R_{sr}. The diagram also shows the velocity vector, $\hat{v}$, of a target. From the location and velocity of a given scatterer, we can also derive the bistatic angle, β, formed by the scatterer's corner of the bistatic triangle, and the angle δ between the velocity vector and the bistatic angle bisector.

In the limit that a target is very far away, $R_{ts} = R_{sr} >> R_{tr}$, or $\beta \approx 0$, the bistatic radar behaves as a monostatic radar, with receiver and transmitter close enough (relatively) to be considered co-located. The bistatic range ellipsoid approaches spherical. The closer the target, the more eccentric the bistatic range ellipsoid. Many radar systems are designed to produce an aggregate range-to-target measurement. It's clear from our discussion thus far that at close distances, the idea of "range" to a target in a bistatic system becomes physically unintuitive. For these cases, a measurement of bistatic path difference is perhaps a more accurate variable. As the name suggests, we measure the amount of extra distance radio waves travelled to reach the target and then the receiver, from the direct transmitter-receiver distance. Mathematically, this is simply $R_{ts} + R_{sr} - R_{tr}$. A 2-dimensional plot showing countours of constant range, or constant bistatic path difference, for a given transmitter and receiver 15 km apart can be seen as x's in Figure 1.3. Notice for smaller values of bistatic path difference (blue), the ellipsoid structure in the curve is apparent, while even a target located just 20 km north of the transmitter, receiver pair would already be placed on a contour looking much less eccentric and more like a circle (truly a sphere in 3 dimensions).

A common specification of monstatic radars is the maximum range they can reliably detect. Often, we see aircraft with a bistatic path difference of 10 km. This specific value of range is illustrated on the same map size map in Figure 1.4, highlighting for us the region within 10 km bistatic path difference as purple, perhaps a theoretical maximum range, with the shading representing range measurements for each location on the map. Unfortunately, the bistatic geometry may mean that maximum range is not an appropriate metric.

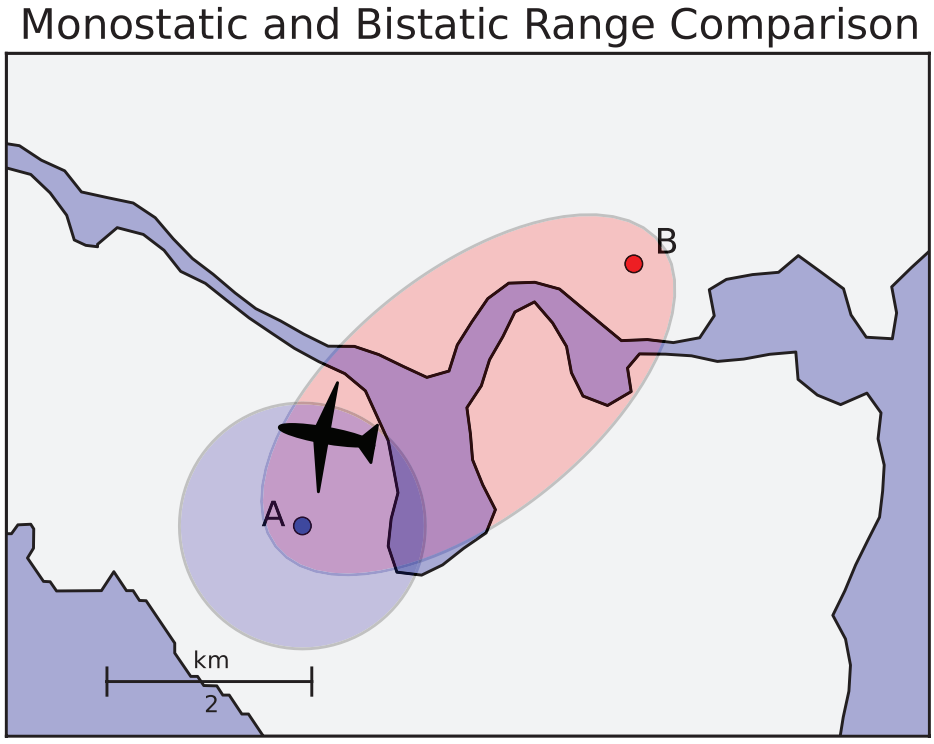

Figure 1.1: Example of monostatic radar, with transmitter and receiver located at A, and bistatic radar, with transmitter at A, receiver at B. The airplane will be detected at a certain range. For the monostatic case, this is anywhere on a sphere centered at A; the footprint of this sphere on land is shown as the blue circle. For the bistatic case, the plane is anywhere on an ellipsoid with one focus at A and the other at B. The footprint of this ellipsoid is shown as the red ellipse. The plane is located on the arc defined by the intersection of the two surfaces.

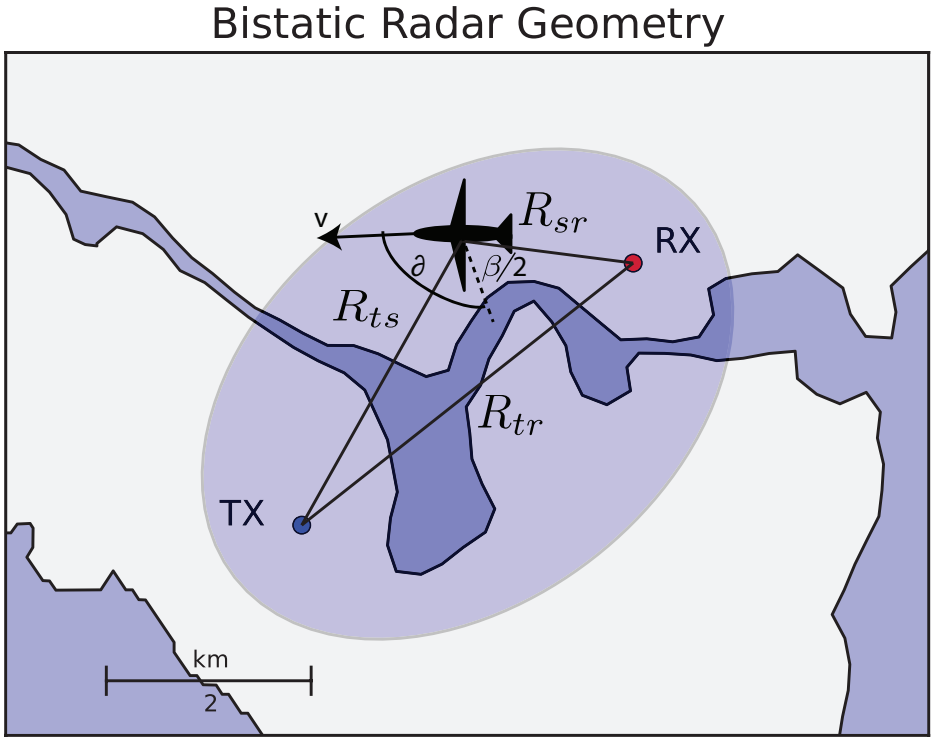

Figure 1.2: Details of the bistatic geometry associated with target detection. We define the velocity $\hat{(v)}$ of a target, as well as the bistatic angle β seen via the angle bisector value $\beta/2$, and the angle between the velocity and this angle bisector, δ. Also seen are the various ranges involved with the detection.

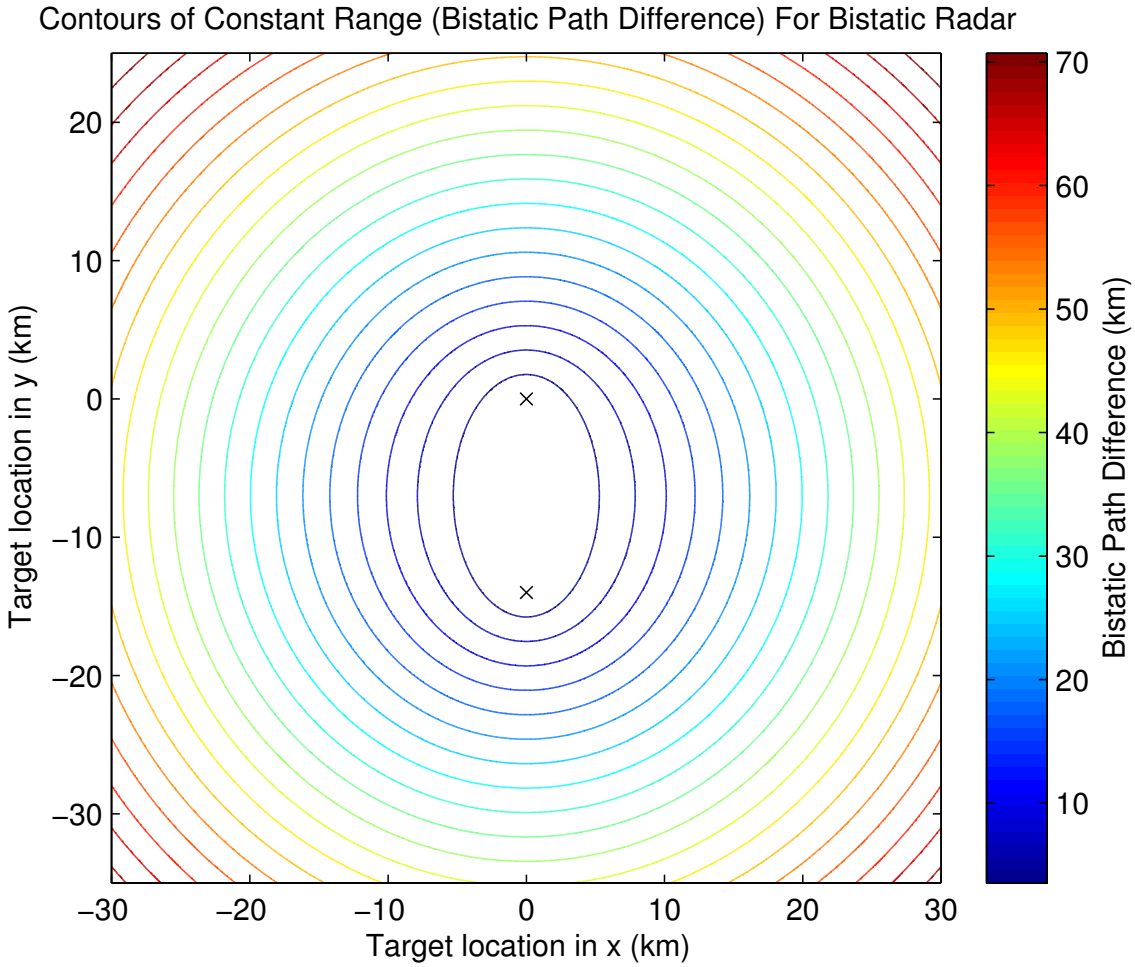

Figure 1.3: A 2-dimensional map displaying a bistatic radar (transmitter and receiver displayed as x's), with contours of constant bistatic path difference, a simple variant of "range" which is more intuitive for the bistatic geometry. Contours close to the focii appear elliptical, while contours further away are already approaching their circular limit.

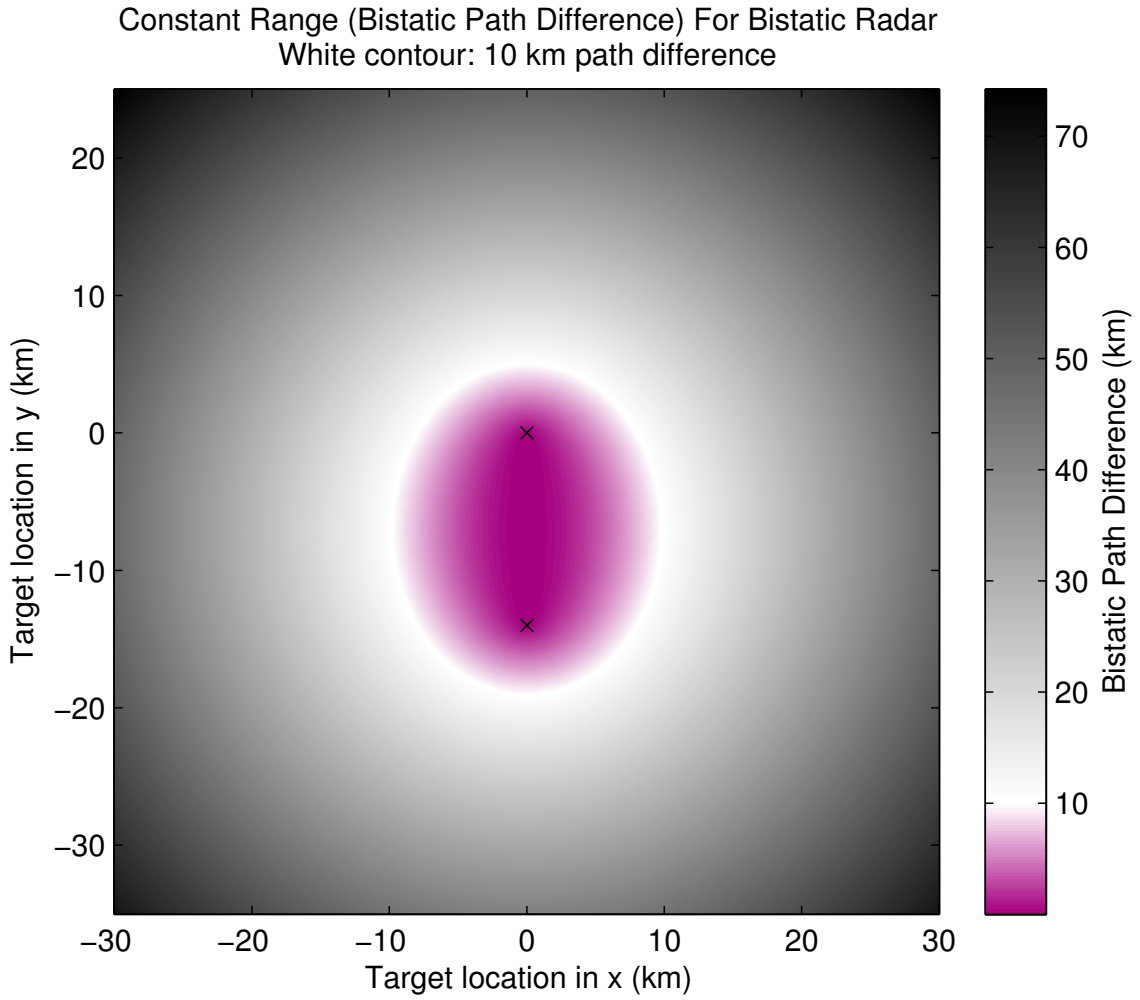

Figure 1.4: The same 2-dimensional map and geometry of Figure 1.3, but with the specific and common case of a target located at a bistatic path difference of 10 km. The transmitter and receiver are 15 km apart. The region within 10 km bistatic path difference is indicated with purple.

1.2.2 Bistatic Doppler Velocity

Doppler velocity in the monostatic case is the radial component of the velocity, with the radar located at the center of the range sphere. It is derived as $\delta R/\delta t$. If a target is traveling tangent to this sphere of constant range, it will have zero Doppler velocity. Refering to Figure 1.2, we can derive this for the bistatic case, the Doppler velocity is derived as $\frac{1}{2}\delta(R_{ts} + R_{sr})/\delta t$, the rate of change of the range from transmitter to scatterer plus scatterer to target over two. Note that in the monostatic limit, this equation agrees with the monostatic equation. A target has zero Doppler velocity as it travels tangent to the contour of constant range, the 3-dimensional bistatic ellipsoid. The equation for Doppler velocity is a similar "radial" component of velocity. It is the velocity, v, projected onto the angle bisector of the bistatic angle, β, with an additional scaling by $\cos(\beta/2)$.

$$v_d = |v| \cos(\delta) \cos(\beta/2) \tag{1.1}$$

In the limit of the target being very far away, $R_{ts} \approx R_{sr} >> R_{tr}$, the bistatic radar is effectively a monostatic radar and the Doppler equations are equal.

Doppler velocity is defined in units of m/s, however what we will measure is a Doppler shift f_d defined in Hz. The Doppler shift will depend on the wavelength of the transmit waveform.

$$f_d = \frac{2v_d}{\lambda} \tag{1.2}$$

In the monostatic sense, Doppler shift is positive when a target is approaching the radar and negative when the scatterer is moving away. For radars that are bistatic in nature, the idea of "approaching the radar" is ambiguous. Doppler velocity or Doppler shift is defined to be positive for a target whose bistatic range is decreasing and negative for a scatterer whose bistatic range is increasing.

Figure 1.5 overlays vectors of maximum doppler shift with the range contours from the previous section. Here every arrow is indicating the direction of maximum *negative* doppler shift of target at that location. As anticipated, the maximum doppler shift is in the direction normal to the curve of constant range. Continuous curves perpendicular to curves

of constant range are hyperboloids.

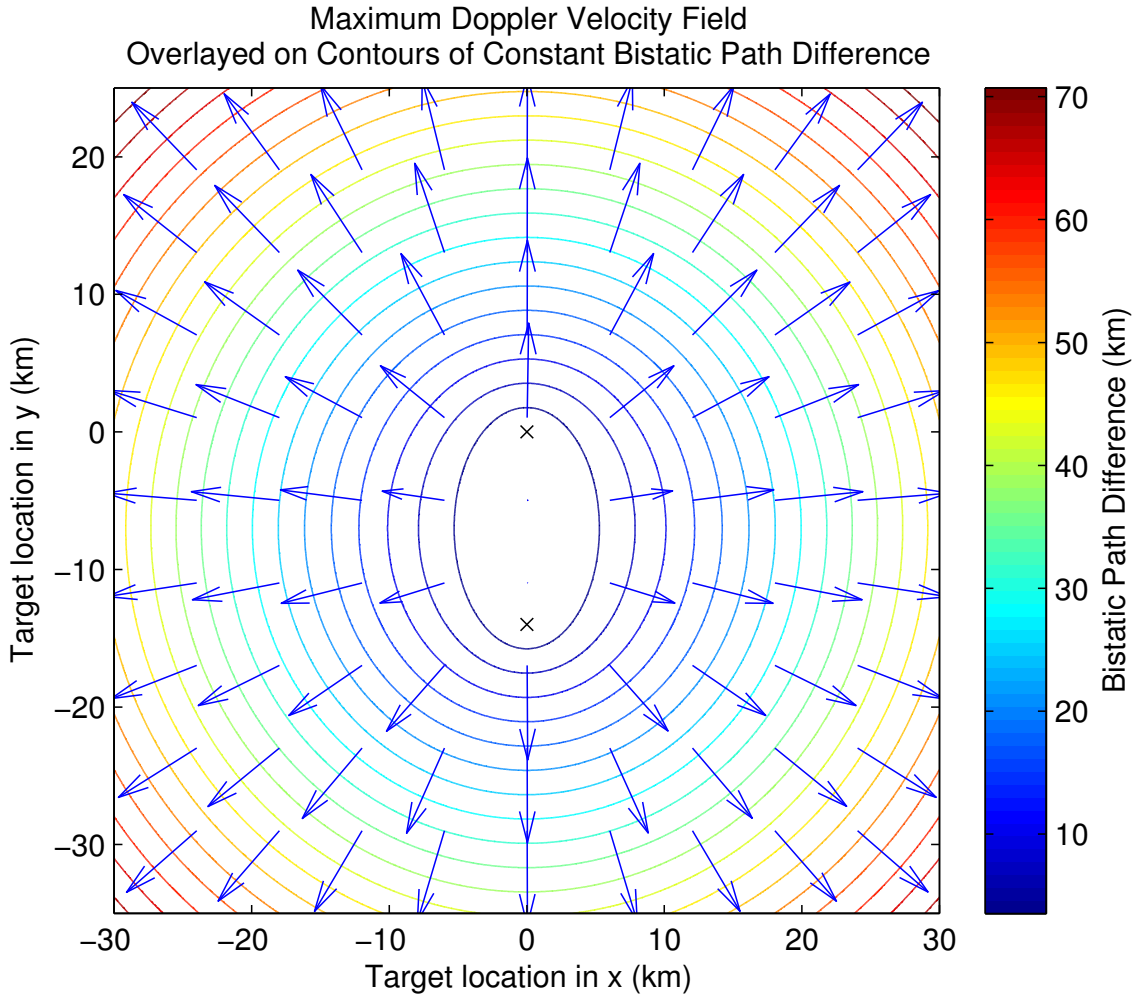

Figure 1.5: Direction of maximum doppler shift, shown pointing in the *negative* direction, or "away" from the radar system overlayed with Figure 1.3, a contour plot of constant bistatic path difference.

Like range, with a target close to the bistatic radar system the doppler velocity meaning can be unintuitive. In these cases we measure "bistatic doppler," where we simply do not divide the doppler velocity measurement by two.

1.2.3 Bistatic Signal Power

Monostatic radars can analyze their performance by the minimum SNR, or the smallest power received from a scatterer compared to the noise power. Lines of constant SNR or constant power in the monostatic radar case are spheres, the same as constant range contours. Because of this, these radars can invert their minimum SNR metric to a maximum Range. For bistatic radar, contours of constant power return are Ovals of Cassini. Contours of constant power for the bistatic case are shown in Figure 1.6. As expected, in the monostatic limit these curves approach spherical. If the minimum received signal power that the passive radar receiver can detect is still a large value, the cassini oval will be one of the curves close to the receiver and transmitter. In this area, the contours of constant range draw very different lines across the map, and the idea of a "maximum range" of a target to the radar cannot be derived to represent radar performance.

To illustrate this further, a passive radar may be characterized by a minimum signal-to-clutter (SCR) ratio. Our particular radar was built to detect targets with a minimum SCR of -85 dB (more on this in Chapter 2). This threshold is shown on our now-familiar 2-dimensional map, for a 0 dBsm target, in Figure 1.7. All 0 dBsm targets in the purple region are detectable. Note that an aircraft can be detected not only if it is close to the receiver, but also if it is close to the transmitter.

1.2.4 Bistatic Radar Cross-Section

The bistatic RCS of a target, like the monostatic RCS, depends on the angle of incidence, frequency, polarization, physical dimensions and electric properties of the target. The bistatic RCS expands this definition to include the scattering angle. For monostatic radar the angle of incidence and scattering angle are the same, and in that limit the definitions are equivalent. Bistatic RCS has an interesting advantage when the target is between the transmitter and receiver. A target with cross-sectional area A, (as seen by the transmitter) will have an enhanced RCS for this forward scatter geometry:

$$\sigma_b = \frac{4\pi A^2}{\lambda^2} \tag{1.3}$$

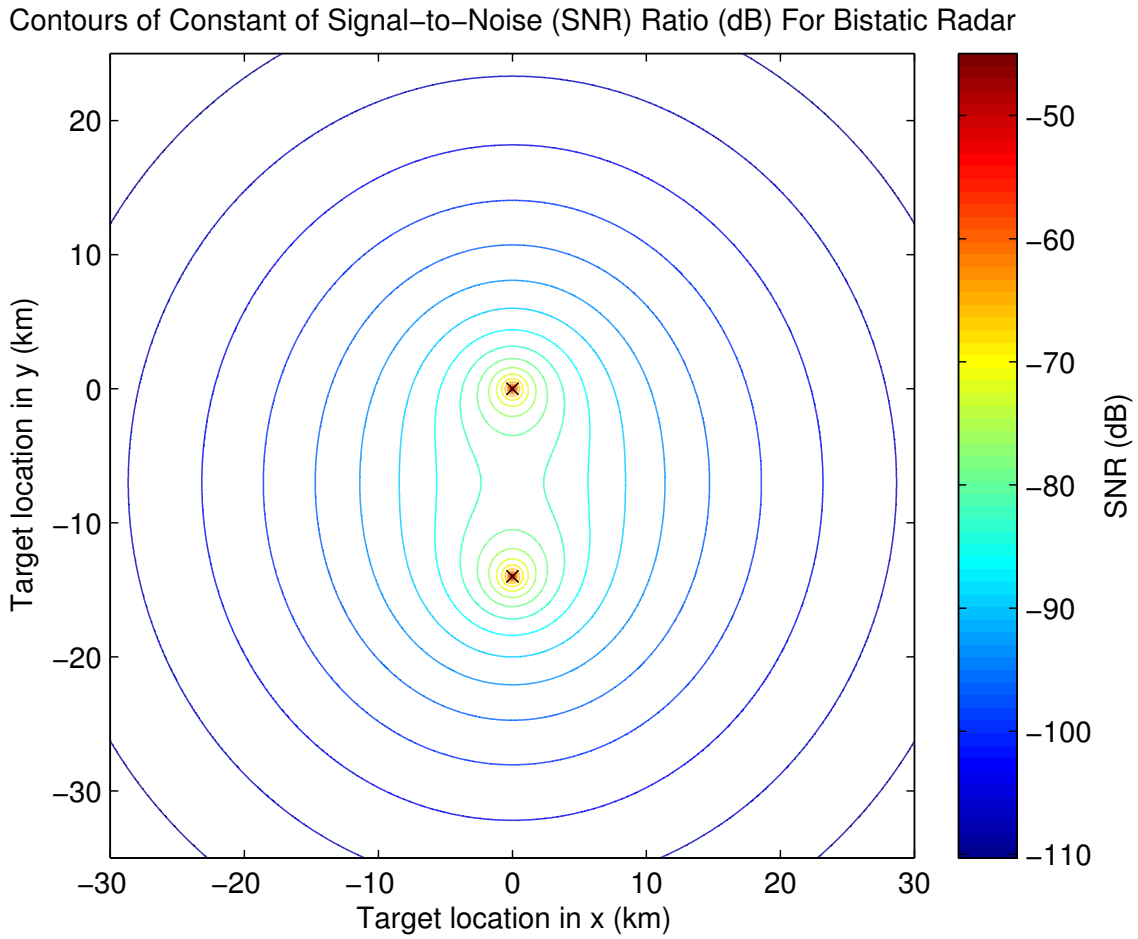

Figure 1.6: Contour plot of received power for the same bistatic radar system plotted in Figure 1.3. These are Ovals of Cassini.

This RCS will be valid for any receiver that lines up within angle $\Theta_b = \lambda/d$, where d is the linear dimension in the plane of the Θ_b. These parameters are shown in [58] and [20]. The latter paper shows that FM and DTV broadcast are frequencies particularly suited for this bump in detectability. For what it is worth, a target exactly on the line between the transmitter and receiver has 0 range and has nearly undetectable velocity.

The RCS of a target is quite complicated, and a full discussion is well beyond the scope of this work. The reader may wish to refer to [46] for an in-depth overview, [28] which is a thorough text on electromagnetic scattering, showing the complexity of the RCS, or [56] for

an example of scatter from a group of meteoroids, whose overall RCS varies in polarization and absolute value with the group spacing.

1.2.5 Common Data Product: Range-Doppler Plot

The range-Doppler plot is a common data product for radars of this type. The radar detection space is divided into cells of a matrix indexed by range (x-axis) and Doppler velocity (y-axis). The motion of a target, say an aircraft, in 3 dimensions does not always lead to an intuitive motion across the range-Doppler plot. A simple demonstration of an aircraft flying in a straight line from south to north near the bistatic radar can be used to demonstrate the mapping of simple motion to the range-doppler equivalent. Figure 1.8 shows the path of an aircraft doing such a fly-by at 200 m/s on both a 2-dimensional map in plot (a) and the equivalent range-Doppler plot in plot (b). The aircraft at current time is indicated by a filled purple triangle, with past instances indicated by empty triangles. The path of the aircraft on the map and in the range-Doppler plot is shown with a dotted black line.

Starting the simulation, the aircraft approaches from the south. As measured by the radar, the target has large bistatic path difference and has a bistatic Doppler velocity of 400 m/s. With the intention to measure targets in close proximity to the radar we are avoiding the convention of dividing by two in Doppler velocity. Positive Doppler velocity indicates decreasing bistatic range.

As the aircraft continues to approach the radar system, the bistatic path difference is shrinking, but the Doppler shift is decreasing. Eventually, the aircraft will come to a point of closest approach, where the Doppler velocity is zero. Notice that the bistatic path difference will never be zero unless the aircraft actually cross the receiver-transmitter line.

After the point of closest approach, the radar will measure the aircraft moving away (negative Doppler shift). Due to symmetry of the maneuver in the 2-dimensional map, the range-Doppler also turns out to be symmetric.

It is important to note that translating from 3-dimensional distance and velocity to range and Doppler is not one-to-one. To transform from range and Doppler back to 3-dimensional

location and velocity will require additional measurement such as angle-of-arrival or multiple simultaneous detections from diverse receivers or transmitters to name a few.

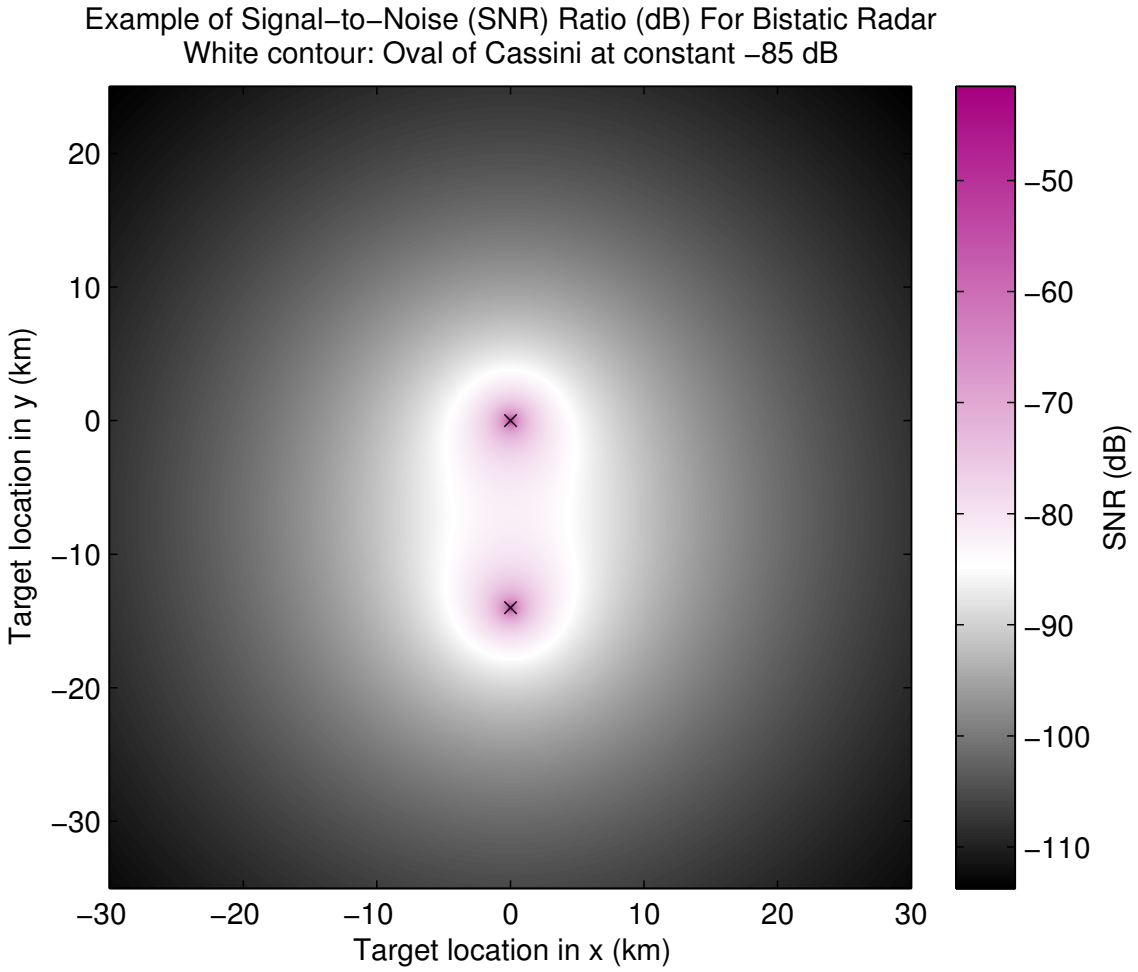

Figure 1.7: An example of relative received signal power as compared to noise (or clutter). This particular example illustrates a threshold ratio of -85 dB. Perhaps our particular receiver can be characterized by sensing every target with a signal-to-clutter (SCR) ratio greater than -85. The region in purple shows the location (projected down to 2 dimensions) of all visible target locations.

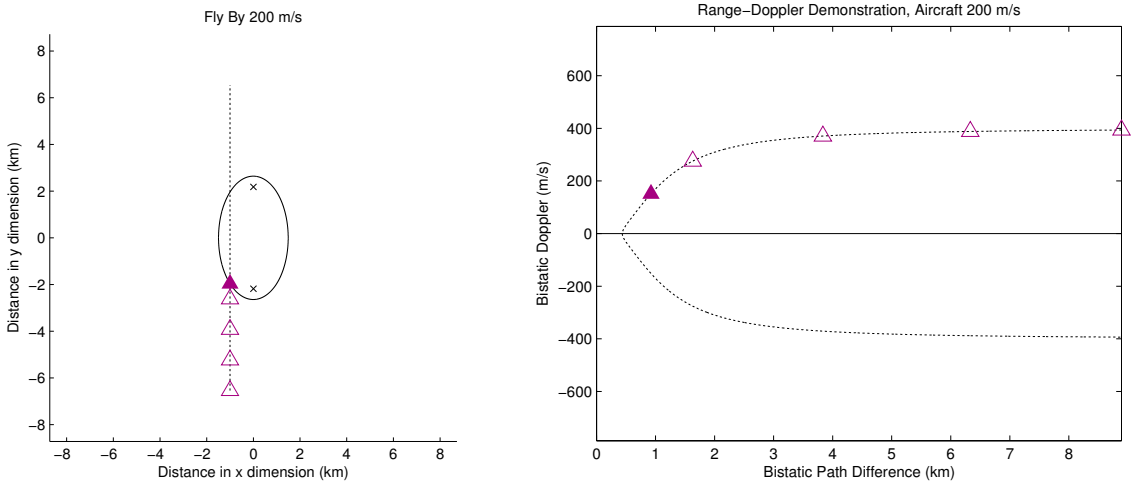

(a) (b)

Figure 1.8: This shows an aircraft flying at 200 m/s at a particular instance in time (purple filled triangle, past positions empty triangles) as seen on a 2-dimensional map in plot (a) and as measured by the bistatic radar on a range-Doppler plot in plot (b). In both cases, the full path of the aircraft over time is shown by a black dotted line.

1.3 Passive Radar Systems and Applications

Several transmitters can be utilized here in the Seattle area. We have focused our efforts on FM and most recently heavily focused on building up our DTV capabilities. However, across the globe there are various systems that also take advantage of cell phone basestations, WiFi, GPS, and satellites.

1.3.1 Transmitters

There are various transmitters readily available for passive radar. In [20], a table is presented that analyzes several transmitters available in the United Kingdom. In Table 1.1, we include the information previously reported as well as extend that table to include analysis of Wi-Fi and GPS[4] signals, as well as an estimate of power of the direct path transmission at the receiver. Since the direct path signal is in many ways the largest clutter signal in the radar, this quantity is referred to as the clutter power, P_c. In terms of radar waveforms, the instantaneous bandwidth can be used to estimate the range resolution of the radar. Here range resolution is defined in the monostatic sense. From this table we can see that the high bandwidth of DTV, GPS, Wi-Fi and cellphone (3G) coverage have range resolution on the order of tens of meters. This baseline precision makes a good argument for using these transmitters in aircraft detection systems. Another important metric of the transmitted waveform is how much power is radiated. By choosing a very rough estimate for the distance from the transmitter to receiver, R_{tr}, the power at the receiver can be estimated. HF, FM, and DTV are very strong signals. GPS, on the other hand, is quite weak by comparison. This will need to be accounted for with appropriate amplification as needed.

[4]GPS values are found from [29], and we quote P_tG_t as the EIRP for the L1 P-code transform quoted as 28.3 dBW or 240 W needed to supply the minimum received power at ground level standard, $P_c = -162$W.

Transmission	Frequency	BW	Range Res.	P_tG_t	R_{tr}	$\frac{P_tG_t}{(4\pi R_{tr}^2)}$	P_c
HF, 20 MHz	20 MHz	9 kHz	16.7 km	50 MW	1000 km	-54 dBWm^{-2}	-11 dBm
FM, KMTT Seattle	103.7 MHz	150 kHz	1 km	50 kW	30 km	-54 dBWm^{-2}	-25 dBm
DTV, KONG Seattle	575 MHz	6 MHz	25 m	700 kW	10 km	-33 dBWm^{-2}	-19 dBm
DAB (UK)	220 MHz	220 kHz	681.8 m	10 kW	100 km	-71 dBWm^{-2}	-49 dBm
Cellphone (GSM)	900 MHz	200 kHz	750 m	100 W	10 km	-71 dBWm^{-2}	-62 dBm
Cellphone (3G)	2 GHz	5 MHz	30 m	100 W	10 km	-71 dBWm^{-2}	-68 dBm
GPS (LP1 P-code)	1.575 GHz	20 MHz	7.5 m	240 W	20200 km	-117 dBWm^{-2}	-132 dBm
Wi-Fi (802.11b)	2.4 GHz	20 MHz	7.5 m	0 W	100 m	-61 dBWm^{-2}	-60 dBm

Table 1.1: Various transmitters available for passive radar use and their properties. Original table displayed in [20], presented here with extensions for GPS and Wi-Fi, as well as a rough calculation of direct path power received, P_c.

1.3.2 Applications

The application of a particular passive radar is determined by the type of transmitter, as suggested by our analysis of Table 1.1. If the particular application needs high range resolution, then large bandwidth (short correlation time) is important. If target detection at large range is needed then a transmitter with high average power may be necessary. Just as the active radar designer tailors the radar waveform for the application, the passive radar designer must choose the best from the available waveforms in the environment. FM and DTV broadcast provide significant terrestrial coverage, as they are designed for inefficient antennas, receivers, and non line-of-sight propagation. Radars using these transmitters are well-suited for detection of aircraft. DTV, with its large bandwidth, provides the excellent range resolution needed, but this can also be achieved with FM transmitters through tracking. The previous list of systems includes a popular interest in aircraft detection showing systems using FM, DTV, DVB-T, DAB, cell phone basestations and satellite communications for this means. Wifi-based radars are better suited for more localized applications due to their lower power and higher bandwidth (can achieve meter level range resolution) and have been used for indoor surveillance. Other applications of passive radar range from oceanography ([54]), detection of radioactive pollution [59], mapping lightning in thunderstorms ([18]) and ionospheric studies ([48]).

1.3.3 Systems

With the increasing interest in passive radar, recent articles have been published to summarize the history, breadth and depth of the field ([20], [32], and [57]). I summarize these works with the following list of passive radar systems:

- HA100 RM-based Passive Radar and SINBAD DVB-T based in France by Thales [37]
- CELLDAR by Roke Manor Research Limited, UK (commercial) [55]
- Science Applications International Corporation (SAIC) Passive, Multi-Static Radar System [4]
- GAMMA, GSM-based passive radar by Fraunhofer FKIE in Denmark
- CORA, DAB system for aircraft and ferry detection Fraunhofer FHR in Germany
- PETRA PCL radar from Fraunhofer FHR in Germany
- Bistatic SAR from Fraunhofer FHR in Germany
- Marauder's Map - WLAN radar network for surveillance from Fraunhofer FHR in Germany
- WiFi-based passive radar in Italy [15]
- Passive Radar from University College London, UK for TV [22] and Satellite illuminators [21]
- Silent Sentry from Lockheed Martin (commercial) [2]
- Air Target detection using LEOS satellites [6]
- Passive Radar Demonstrator (PaRaDe) in Poland [36]

- FM-based PCL radar by Dynetics [60]
- HF radar receivers for surface wave scatter [54]
- FAA traffic control (1.295 GHz, pulsed) Passive radar for lightning detection in thunderstorms by Severe Storms Research Center (SSRC) in Georga, USA [18]
- SELEX FM-based passive radar [11]
- VERA system by ERA (Czech Republic) detects aircraft location from their own secondary surveillance radar emmission [3]
- Cassidian, of European Aeronautic Defence Space Company N.V. (EADS)[52]
- Occiu FM-based passive radar (no longer in operation) [13]

This list is surely incomplete. With many of these systems developed since the 1990s, it is clear that passive radar system development is currently of great importance.

1.4 Focus of this work

Our key contribution is our experimental investigation of high speed sampling as an alternative to analog down-conversion. This dissertation serves as a recipe for end-to-end system design and considerations for modern passive radar. In addition to direct sampling experimentation, we explore high aggregate bandwidth with the intention of using DTV as an illuminator of opportunity. The DTV signal has 6 MHz of bandwidth, while the existing ISIS MRR can only handle about 2 MHz. With DTV in mind, we explored several modes of operation for the radar, some of which do not limit us to DTV at all. These are described in detail in Chapter 5. The predicted system performance generically derived in the document is brought to life by applying this receiver to DTV-illuminated aircraft detection, which we cover in detail in Chapter 7, but we keep this goal in mind throughout the document. Figure 1.9 shows the location of our receiver at the University of Washington (marked RX) and various DTV transmitters (marked as TX) in the greater Seattle area.

Surprisingly, the spectrum is fairly sparse in terms of channel occupancy here, and this is important to consider with mobile operation. We may enter an environment that relies on our adjacent channel suppression to be very good. This is one of the many factors to consider when redesigning the FPGA software of this receiver to fit the application.

The system offers many advantages over current technology:

- The system provides new capabilities for ionospheric studies; there is no other radar like it for that application

- The system provides an important technology study which allows the user to combine data from several transmitters spread far in frequency space (such as FM at 100 MHz and DTV at 540 MHz) without replication of the RF path, including the digitizer

- We will be able to run experiments to contribute to the fresh field of mobile passive radar receivers, especially in terms of ground clutter studies

- The system is low-cost, $24k

- The system is open source. There are (I'm sure) many classified systems and commercial systems that protect their IP. We have a chance to offer our technological developments to the open scientific community.

- The system uses the latest technologies, which are new developments for radar and will only become less expensive and more efficient, specifically: FPGA, 10 GbE data link, and graphical processing units (GPUs).

- The system provides high aggregate bandwidth, up to 6.4 Gbps, which we can allocate as desired. A current choice is to produce 8-bit IQ samples on 4 antennas, each observing 60 MHz of the spectrum. While data loss occur at these high rates, we demonstrate the ability to detect aircraft and retrieve accurate range and Doppler velocity measurements.

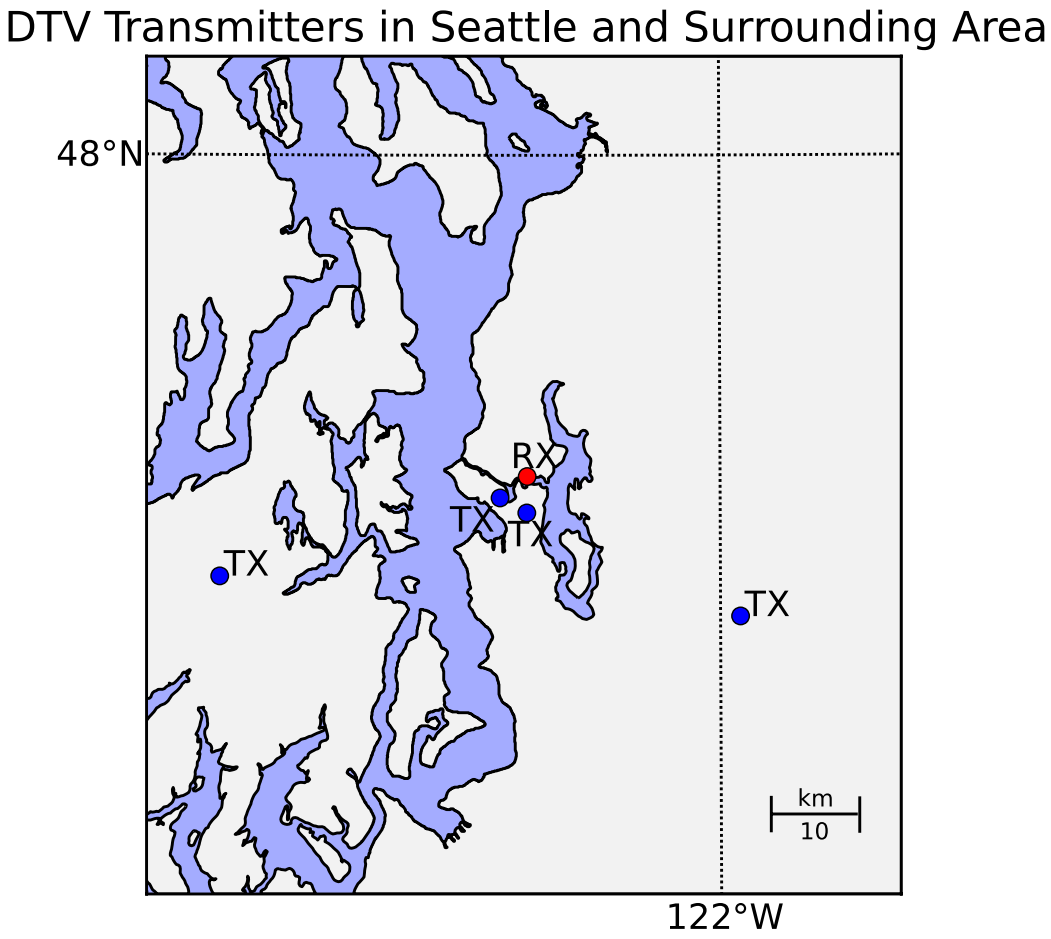

Figure 1.9: Map of the stationary receiver location, marked RX, at Sieg hall at the University of Washington (UW) and various DTV transmitters, marked TX, in the greater Seattle area.

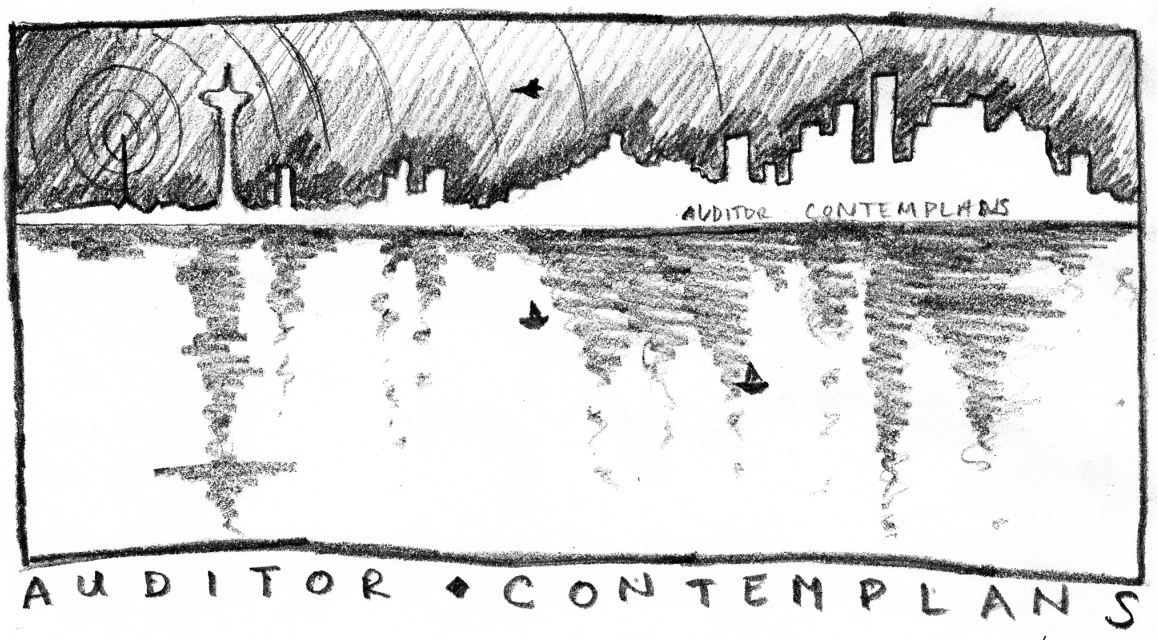
AUDITOR CONTEMPLANS
AUDITOR • CONTEMPLANS

Chapter 2

PERFORMANCE PREDICTION

2.1 The SCR and Dynamic Range

We will focus now on major components of passive radar performance, and give specific examples of our geometry and the current DTV transmitters in use. An in-depth discussion based on several types of transmitters can be found in [20]. All radars, active and passive, require sufficient transmitter power to enable detectable scatter from the desired target. Unlike active radars, passive radar sensitivity is usaully limited by the ability to detect scattered signals that are stronger than the noise, but much weaker than the dominant transmitter signal. These radars are in general not limited by noise power (although many other radars are limited by noise power, others by electronic counter measures (ECM) such as jamming, interference), but instead by the signal-to-clutter (SCR) ratio, also seen in literature as the signal-to-interference (SIR) ratio. This parameter compares the received power from a scatterer to the received power from the direct path transmission. In general, this parameter is quite small, even smaller than the SNR associated with these radars at the antenna. A thorough derivation of the SCR can be found in [47], and we will derive it here as well. P_c, the signal power at the receiver antenna associated with the direct transmission can be defined using the Friis free space transmission equation:

$$P_c = P_t \frac{\lambda^2}{(4\pi)^2} \frac{G_t(\widehat{tr}) G_r(\widehat{rt})}{R_{tr}^2} \tag{2.1}$$

This describes how the transmitted power spreads as it travels through free space. At the receive antenna, the transmitter power P_t is reduced by the square of distance from the transmitter to the receiver, R_{tr}, in addition to a factor of $(4\pi)^2$. We also consider the boost from the gain of the transmit antenna in the direction of the receiver, $G_t(\widehat{tr})$, the gain from the receive antenna in the direction of the transmitter, $G_r(\widehat{rt})$, and the square of the

wavelength, λ. This description, however, does not account for a variety of propagation phenomena such as ducting, troporefraction, or Lloyd's mirror, to name a few. The latter considers the destructive and constructive interference of this direct path signal reflecting off of the ground. Also not considered here are system losses such as transmitter impedence mismatch, transmitter illumination loss, transmission loss, etc. The reader can learn more about these terms in [53]. A typical value for the bistatic power received at the UW site from the DTV transmitter KONG (assuming $G_{\widehat{rt}} = 0$) is:

$$P_c = -11.5 \text{ dBm} \tag{2.2}$$

The power received due to the signal scattered from a target, P_s, can be shown as:

$$P_s = P_t \frac{\lambda^2 \sigma(\lambda, \widehat{ts}, \widehat{sr})}{(4\pi)^3} \frac{G_t(\widehat{ts}) G_r(\widehat{rs})}{R_{ts}^2 R_{sr}^2} \tag{2.3}$$

Again we do not consider Lloyd's mirror or ground clutter, which considers the new directional gains $G_t(\widehat{ts})$ and $G_r(\widehat{rs})$ pointing from transmitter to scatterer, and receiver to scatterer as well as the RCS of the target, $\sigma(\lambda, \widehat{ts}, \widehat{sr})$, which is dependent on angle of incidence as well as angle of reflection. RCS is dependent on a number of factors, but we have also included the wavelength here to remind ourselves that this radar, looking at transmitters across VHF and UHF, and potentially on the same antenna, can detect targets whose RCS is frequency dependent. If we choose an arbitrary 0 dBsm target seen at a distance of 3 km from the receiver, we can expect a value in the ballpark of:

$$P_s = -96.8 \text{ dBm} \tag{2.4}$$

Here we have used $R_{sr} = 3$ km, and estimated a worst case $R_{ts} = R_{tr} + 3$ km, forming a straight line from transmitter to receiver to scatterer. We use the very same R_{tr} value used to produce the result in Equation 2.2. We can therefore define the SCR (or SIR) as:

$$SCR = \frac{P_s}{P_c} = \frac{\sigma(\lambda, \widehat{ts}, \widehat{sr})}{4\pi} \frac{G_t(\widehat{ts}) G_r(\widehat{rs})}{G_t(\widehat{tr}) G_r(\widehat{rt})} \frac{R_{tr}^2}{R_{ts}^2 R_{sr}^2} \tag{2.5}$$

The transmitter power is absent, as well as any explicit dependence upon the wavelength. There is a dependence on the bistatic RCS, $\sigma(\lambda, \widehat{ts}, \widehat{sr})$, which is itself dependent on wavelength and angles of incidence (from transmitter to scatterer $\widehat{ts}$) and reflection (from scatterer to receiver $\widehat{sr}$). There is also dependence on the various gains of the antennas of the transmitter($G_t(\widehat{tr})$ and $G_t(\widehat{ts})$) and receiver ($G_r(\widehat{rt})$ and $G_r(\widehat{rs})$), as well as the distances from the transmitter-to-receiver, R_{tr}, from the transmitter-to-scatterer, R_{ts}, and from the scatterer-to-receiver, R_{sr}. These terms are illustrated in Figure 1.2 which shows bistatic scattering geometry.

If we assume isotropic antennas, the SCR becomes:

$$SCR = \frac{\sigma(\lambda, \widehat{ts}, \widehat{sr})}{4\pi} \frac{R_{tr}^2}{R_{ts}^2 R_{sr}^2} \tag{2.6}$$

For stationary operation with our receiver located at the UW we typically see airplanes landing on nearby Lake Union, $R_{sr} \approx 2.5$ km. Considering this geometry and the nearby DTV transmitters, the transmitters on Queen Anne Hill maximize the SCR at -70 dB. This is actually a convenient geometry, and in general, especially when mobile, we can expect to encounter a smaller SCR. Using our metric of detecting any 0 dBsm target 3 km away from the transmitter, we get an SCR of -85 dB, and we aim to consistently detect targets in these circumstances.

With a dynamic range of 85 dB, the scattered power from the target would be at the threshold of the smallest signal we could measure (assuming the front end signal gain can be set to maximize this - challenging in the mobile scenario). For reliable detection we need to impose a threshold of 15 dB signal strength above the noise floor, leaving us to achieve a dynamic range of 100 dB.

To meet this goal we can consider antenna patterns with good side-lobe rejection, coherent integration, and space-time adaptive processing (STAP) to suppress the direct path as well as ground clutter components.

Feature	FM Broadcast	DTV Broadcast
Call Sign	KMTT	KONG
ERP	50 kW	700 kW
Carrier Frequency	103.7 MHz	575 MHz
Bandwidth (BW)	± 75 kHz	± 3 MHz
Wavelength (λ)	2.89 m	0.52 m
Range Resolution (c/BW)	2 km	50 m
Maximum Doppler Shift (200 mph Target) ($f_d = v/\lambda$)	± 30 Hz	± 171 Hz

Table 2.1: An FM and DTV transmitter comparison.

2.2 Self-Ambiguity of the Radar Waveform

This radar is designed to use both VHF and UHF transmitters. Commercial FM (VHF) broadcast ranges from 88 to 108 MHz, with a bandwidth of 150 kHz. This leads to a range resolution of 2 km. The DTV (UHF) waveforms near Seattle range in center frequency from roughly 400 to 800 MHz, and with 6 MHz of bandwidth offer the exceptional range resolution of 50 m. A summary of these properties for particular transmitters can be found in Table 2.1. In the past few decades, several investigators have reported the utility of television broadcast for radar applications, including [22], [1], [45], [49], and [16]. In this work we will demonstrate this range resolution in results.

An important aspect of the transmitted waveform in all radar is the self-ambiguity:

$$|\chi(\tau,\nu)|^2 = \left| \int_{-\infty}^{+\infty} u(t)u^*(t+\tau)e^{j2\pi\nu t}dt \right|^2 \tag{2.7}$$

It describes how much the waveform $u(t)$ correlates with itself in time, τ, and frequency ν. The ideal radar waveform has a "thumbtack" self-ambiguity, infinitesimally short correlation in both time and frequency, with peak located at (0,0). The self-ambiguity of a waveform lets the radar engineer know how well she could expect to localize a target unambiguously in range and doppler space. With tracking, one can improve upon this "prior" for the resolution.

The width of this self-ambiguity peak determines the range and doppler resolution of the system. These metrics describe how accurately you can resolve two targets that are

near each other in range and doppler space. The range resolution can be approximated analytically as the inverse of the instantaneous modulation bandwidth (the deviation from the carrier frequency) or the correlation time of the modulation.

It may be the case that using a subset of the spectrum will produce a more ideal ambiguity function. In [49] the authors show that if deterministic components in the DVB-T waveform are removed, the ambiguity function improves greatly. We see similar improvement by notch filtering the line component associated with the DC offset of the 8VSB modulation used in the DTV standard. The result of this preconditioning is a sacrifice range resolution, increasing from 50 to 60 m, but as you can see in the full spectrum self-ambiguity (Figure 2.1) and in the modified self-ambiguity (Figures 2.2), the self-ambiguity is greatly improved. This is indeed not the true self-ambiguity of the transmitted waveform, but is the self-ambiguity of the received transmission. An obvious effect of this is the strong 0 doppler slice associated with multipath reflections from ground clutter. Of course, without knowing the waveform at the transmitting antenna a stationary receiver cannot tell the difference between range ambiguity inherent in the waveform and reflections of the transmission off non-moving clutter like buildings and mountains. Figures 2.3 and 2.4 compare just the 0 doppler slice of the self-ambiguities, or in other words, the auto-correlation of the two signals. This is the horizontal slice found halfway up the self-ambiguity plot. Figures 2.5 and 2.6 show the 0 range slice, the vertical left-most slice of the self-ambiguities. Also noted on these graphs is their height above the floor of the self-ambiguity. These slices show the range resolution improves by this modification of the spectrum, but the doppler resolution is more or less unchanged.

The final data rate need not exceed the signal bandwidth as it will not improve the range resolution. The measured range can be described as $R_{measured} = R_{true} + \Delta R_{noise,quant}$, where this quantization noise comes from the range resolution inherent in the waveform. Measures can be taken to reduce the range resolution below the waveform limitation. The measured range of a target can be better estimated by the data fusion of multiple measurements and tracking [17].

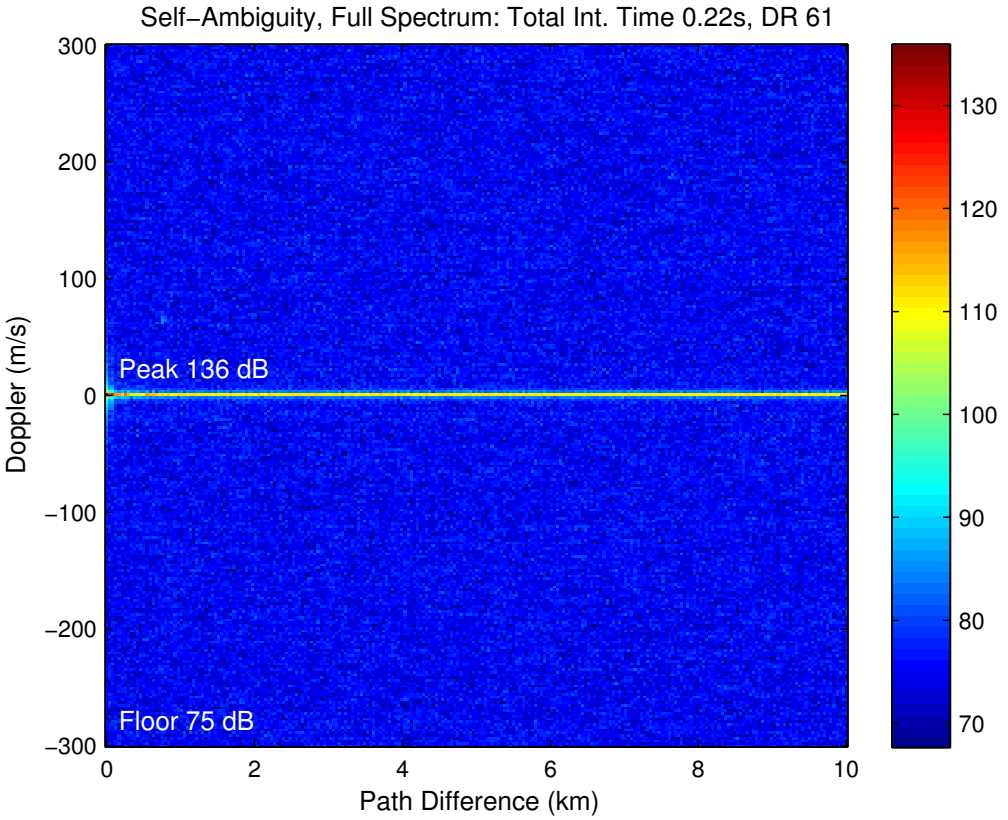

Figure 2.1: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth. We can note the long correlation time by the large power horizontal 0 doppler line.

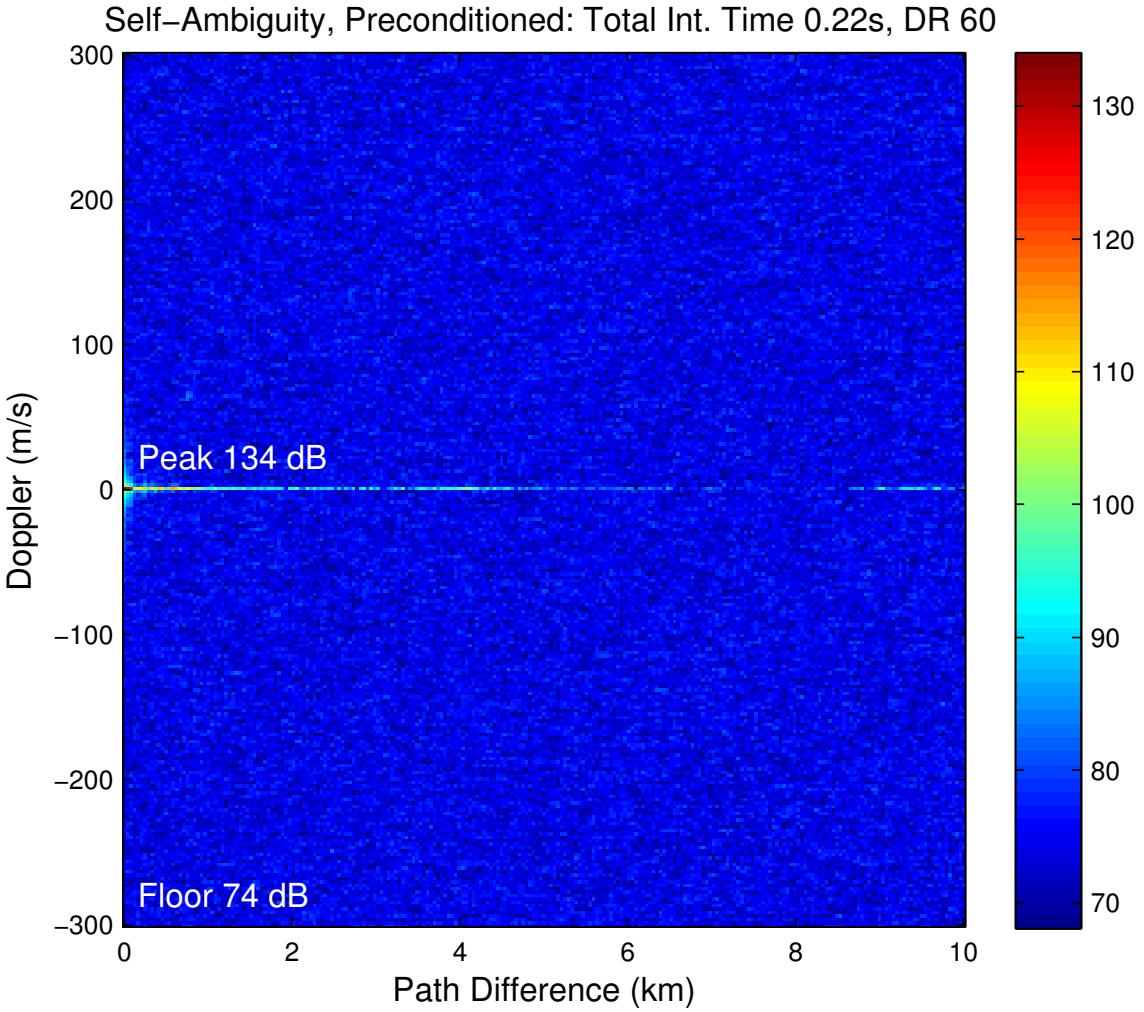

Figure 2.2: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth. By removing a larget deterministic component we see an expected drop of the horizontal 0 doppler line, improving the self-ambiguity.

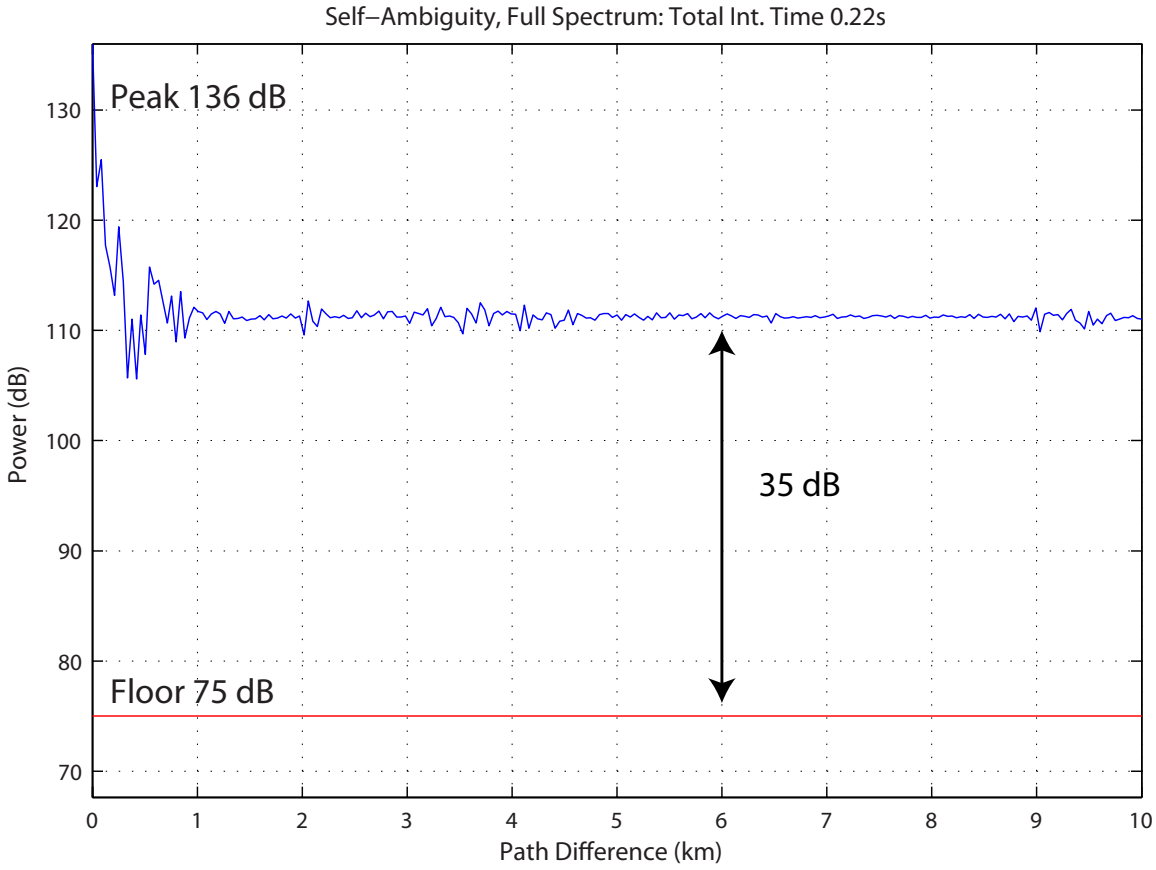

Figure 2.3: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth, isolating the 0 doppler slice (simple fft, no windowing). We see the correlation remain quite high here due to the deterministic component of the spectrum rather than legitimate ground clutter reflections.

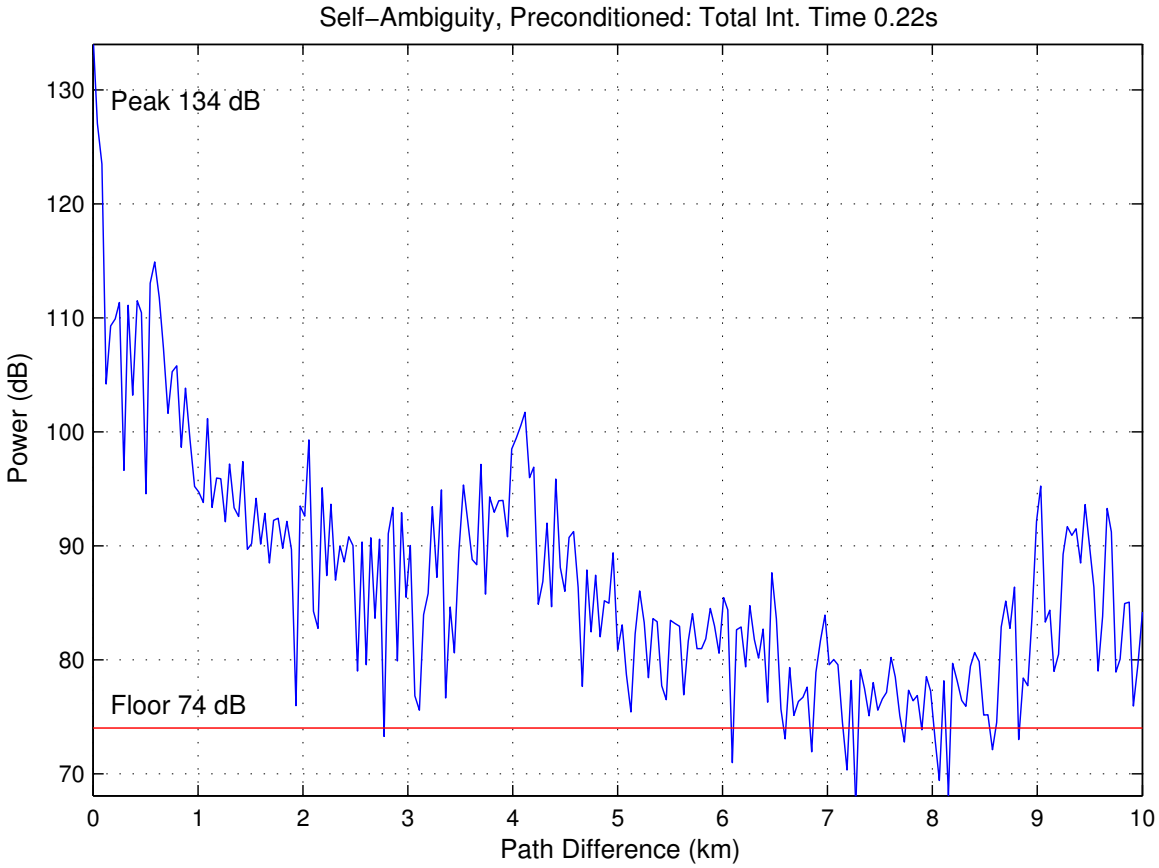

Figure 2.4: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth, isolating the 0 doppler slice (simple fft, no windowing). We see the correlation drop significantly. It is important to note that what we expect to see here is the true self-ambiguity of the signal as well as any stationary ground clutter reflections or moving targets with 0 doppler shift. By removing a large deteministic component we can tell that this is truly an improvement that will help us detect 0 doppler targets.

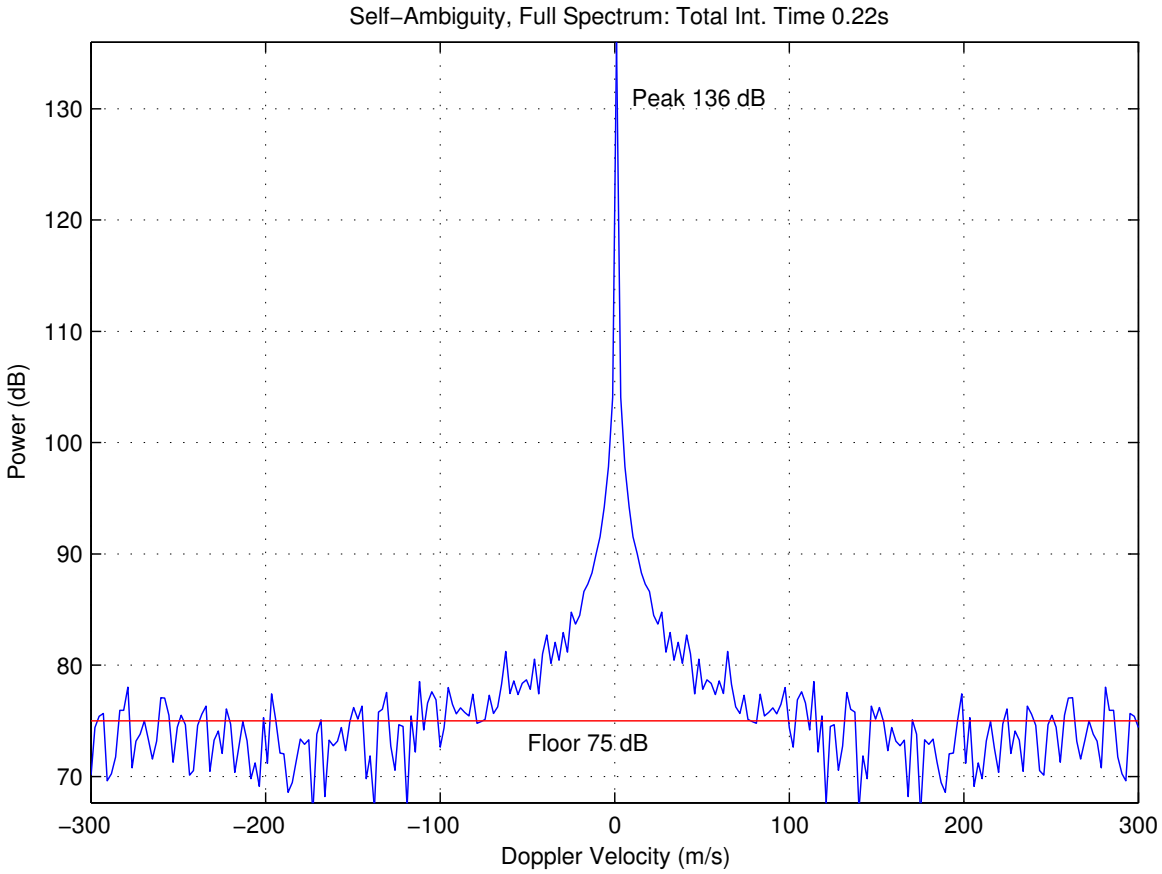

Figure 2.5: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using the full 6 MHz bandwidth, isolating the 0 range slice. We see the frequency correlation drop off rather quickly, showing a rather nice self-ambiguity property without modification.

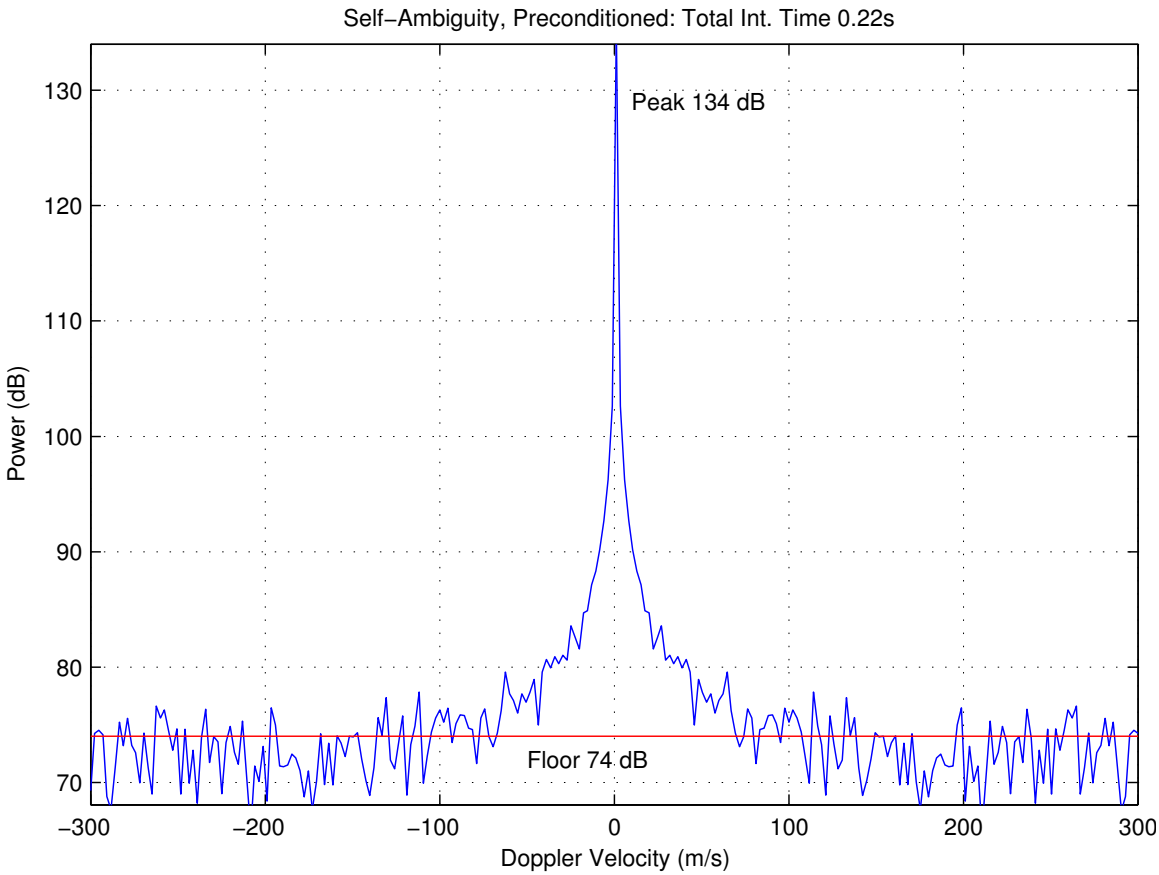

Figure 2.6: Self-ambiguity function of DTV transmitter KONG as received at Sieg Hall at the UW, using a 5 MHz subset of the full bandwidth, isolating the 0 range slice. We see the frequency correlation changes little by using this subset.

2.3 Cross-Ambiguity Calculations

Passive radar detection is performed by computing the cross-ambiguity estimate of the received signals. This is essentially performing a matched filter operation on a signal where the direct path has been suppressed. This can be done by distance in the case of the Manastash Ridge Radar ([48]) and the Intercepted Signals for Ionospheric Science (ISIS) Array ([33]), or by means of directional antennas, beamforming, or STAP.

At baseband, the digitized matched filter $x[n]$ and the signal with the suppressed direct path $y[n]$ are combined to form the product $z[n, n_r]$:

$$z[n, n_r] = y[n]x^*[n - n_r] \tag{2.8}$$

This product is now a two-dimensional signal in time index n, and delay n_r. The signal $z[n, n_r]$ is strongly oversampled, thus we can perform a coherent integration. This provides a significant increase in the SCR, by reducing the effective bandwidth of the receiver.

$$z_M[m, n_r]|_{m=Mn} = \sum_{k=0}^{M-1} z[Mn + k, n_r] \tag{2.9}$$

This new signal $z_M[m, n_r]$ has been reduced in data rate in one dimension. In the case of DTV observation of aircraft, the decimation factor M can exceed 6000. The range-oriented time signals now evolve at 1 kHz rather than 6 MHz. This is an integration time of one one-thousandth of a second for an increase in dynamic range of 38 dB.

Finally we can compute the Fourier Transform to get:

$$\chi[n_r, \nu] = \sum_{m=0}^{255} z_M[m, n_r]e^{-j\frac{2\pi\nu m}{256}} \tag{2.10}$$

With a 256-point FFT the total coherent integration is one quarter of a second, for a total dynamic range gain of 62 dB. We can additionally compute four incoherent averages to lower the variance of our estimate, and create a 1s frame of data.

This cross-ambiguity shows us detectable moving targets and ground clutter with various ranges and doppler shifts. With monostatic radar, the range indicates a location on a

sphere centered at the transmitter/receiver. For the case of bistatic radar, the location of the scatterer is anywhere on an ellipsoid with the transmitter at one focus and the receiver locations at the other. For stationary operation, unmoving structures and targets moving tangential to the surface of this ellipsoid of constant range will have 0 doppler velocity. For this work, we define range to be $R_{ts} + R_{sr} - R_{tr}$. With radars that primarily receive backscatter, "range" is typically half of this quantity. The doppler velocity is the bistatic range rate of change. Just as range alone does not determine position in 3D space, this measurement of doppler shift presents just a single point of view, the radial velocity, of the true 3D velocity of the target.

These cross-ambiguity plots are the data product we shall present in Chapter 7.

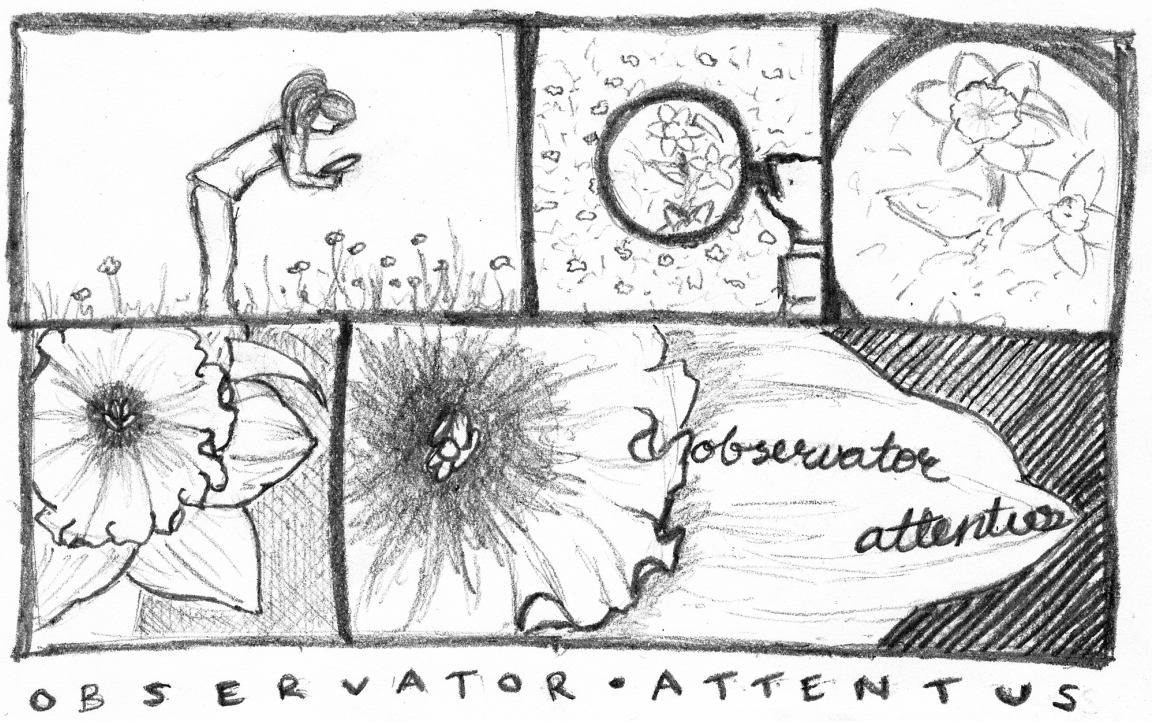
observator
attentus
OBSERVATOR • ATTENTUS

Chapter 3

SYSTEM DESIGN AND SPECIFICATIONS

In this chapter we outline the main features of a receiver for passive radar application.

3.1 Typical Passive Radar System design: ISIS and the MRR

Unlike most receiver designs, our work emphasizes receivers which have no analog down-conversion stage. Instead the received signals are passed directly from the antenna, through some filtering, and then are directly digitized with little or no amplification. The filter is quite important, as it suppresses undesired signals as well as limits the total bandwidth presented to the digitizer, in order to minimize aliasing. By avoiding analog downconversion, we expect to accomplish the following:

1. reduced number of active electronic components which contribute nonlinearity, and thus limit the dynamic range detectable,

2. reduced analog complexity and cost

3. increased flexibility of operation

4. increased simplicity of operation

A standard radar receiver is shown in Figure 3.1. Data is received on an antenna, then passes through an image reject filter and analog mixer to shift the data frequency to an intermediate frequency (IF). To properly digitize the data, first an anti-aliasing (low-pass or band-pass) filter is applied and the resulting data is streamed into an Analog-to-Digital converter (ADC). Finally, a digital downconverter (DDC) is used to mix the sampled IF signal down to baseband and minimize the data rate. This example shows that the digitized IF signal has many channels of interest, and the DDC processes each to deliver four different baseband

channels from the same data. Passive radars often do not require signal amplification, but other radar receivers may include a low-noise amplifier (LNA) behind the antenna. This style of radar relies on extremely linear and costly analog components. The first generation MRR receiver was slightly modified from this, with a single antenna having an analog IQ mixer (one input, two outputs representing in-phase and quadrature components), two digitizers, and was capable of observing one FM station after the DDC.

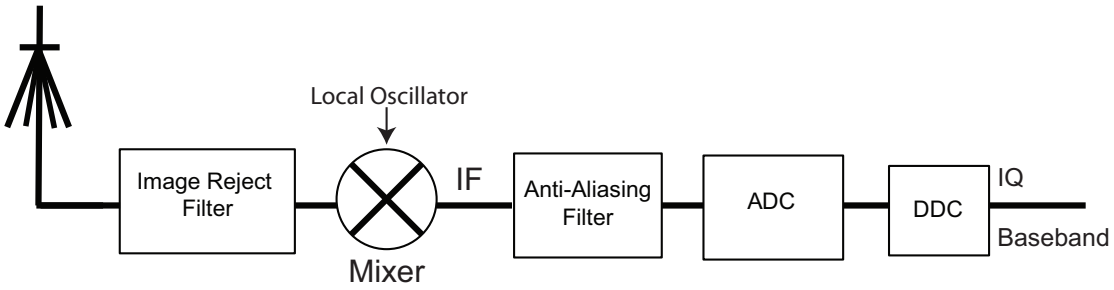

Figure 3.1: An example radar receiver that receives data on an antenna, mixes the data down to an intermediate frequency with an image reject filter and analog mixer pair, applies an anti-aliasing filter and samples the signal at IF with the Analog-to-Digital converter (ADC), and finally applies a digital downconverter to mix the sampled IF signal down to IQ and minimize the data rate. Passive radars often do not require signal amplification, but other radar receivers may include a low-noise amplifier (LNA) behind the antenna.

Newer generations of the MRR, such as that of the ISIS receiver had greatly improved performance capabilities. Figure 3.2 shows the process for the ISIS receiver to observe both DTV signals and FM signals. The top path shows DTV processing, which follows the typical radar design discussed earlier. The FM path below shows the mixer is eliminated entirely. With regards to receiver design, as time progressed digitizers became faster and cheaper while digital signal processing (DSP) became a well-established field. Digitizers began to move to the left of our diagram, getting closer to the antenna. The FM processing on ISIS is an example of this trend, and begins the RRSL's investigation of direct sampling. ISIS is an important upgrade from the previous radars as the final aggregate bandwidth is much greater than the 150 kHz FM broadcast. For a single antenna, just one digitizer is needed and the DDC can produce four different FM stations simultaneously. In DTV mode, each antenna input passes through an analog mixer, a single digitizer, however it can only look

at a 1 MHz from the 6 MHz signal. In either case, the ADC operates at 72 MSamp/s, while the signal content is centered at 126 MHz (DTV) or 100 MHz (FM). This is an example of undersampling. While the Nyquist criteria is not strictly met, the signal has been filtered so the spectrum is sparse, and aliasing occurs but does not have adverse effects on signal integrity.

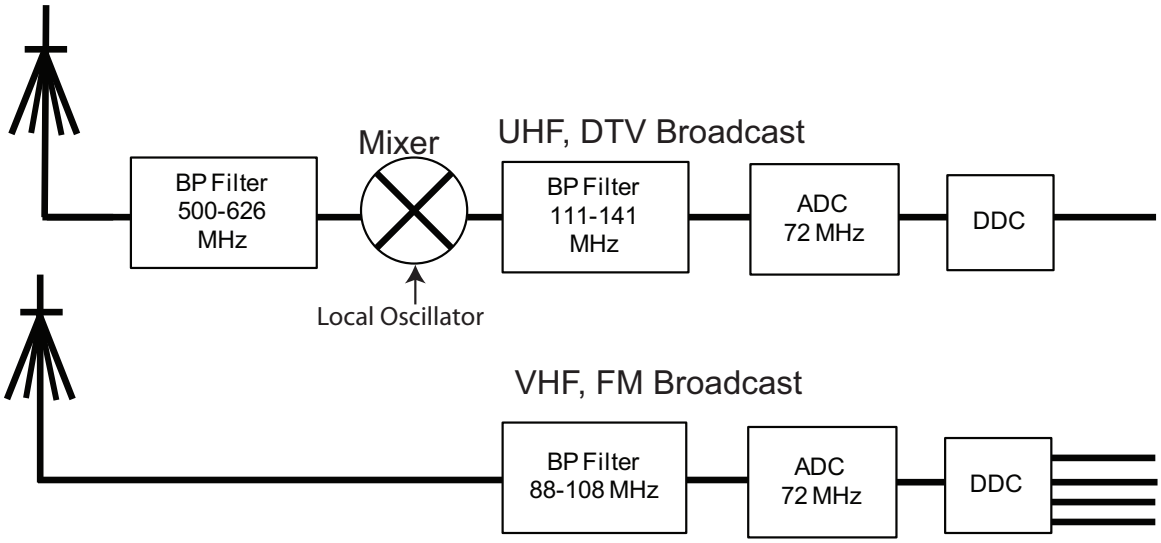

Figure 3.2: A particular example: the ISIS passive radar system. DTV processing (top) is similar to typical radar receiver design explored in Figure 3.1, however the FM processing avoids the analog mixing stage and digitizes the filtered FM broadcast directly.

3.2 Current System Design

Figure 3.3 shows a block diagram of the new receiver architecture, specifically for one of our DTV modes. For this configuration, we use two omnidirectional wideband discone antennas. One of the few analog components is a band-pass filter that isolates a portion of the DTV spectrum from 500-626 MHz. The outputs of these filters connect directly to the two KATADC boards on the ROACH. Before the digitizer, the KATADC board has a fixed front end amplifier (Sirenza Microdevices SBB-2089Z (50 - 850 MHz) or SBB-5089Z (50 MHz - 6.0 GHz)) with gain of 20 dB followed by a software-controlled attenuator that varies by 0.5 dB increments from 0 to -31.5 dB. The KATADC uses the Texas Instruments ADC08D1520, an 8 bit sampler that can operate with two channels up to 1.5 GSamp/s or in single channel mode up to 3 GSamp/s. For initial experimentation, we take advantage of the fast rate of the single channel interleaved mode and sample at 1.6 GSamp/s using an 800MHz clock input. The FPGA clock is derived from the ADC sample clock. The 8-bit data from a single ADC is transfered in 8 parallel streams to the Xilinx Virtex-5 FPGA. In interleaved mode the FPGA runs at one eighth the ADC sample rate, in this case 200 MHz, and receives 8 samples in parallel every FPGA clock cycle. 200 MHz is within the FPGA 350 MHz threshold, which allows us to avoid timing closure issues for complicated designs. The antenna, filter, ROACH board, and the data storage computer Thorex, together comprise the mobile components of our system. Additional signal conditioning and downconversion is done offline, and with the mobile components comprises receiver designs comparable to Figures 3.1, 3.2, and 3.3. Included in this diagram is the processing after the DDC. The IQ, baseband data can be used to do beamforming (in our case simply nulling the direct path transmission), cross-ambiguity calculations, and in the future detections can be automated, combined, and allow for tracking of targets of interest. The following sections will break down the major processing components for this particular design in detail. Our initial design for the channelization of the data relied on polyphase decimation rather than digital mixing on the FPGA. While we have also explored modes using digital mixing and found these superior for our needs, they are rather straightforward. The polyphase mode is arguably more complex, and therefore allows us to go into subtle details of design choices here.

Other modes of operation are described in Chapter 5 with results of detections presented in Chapter 7.

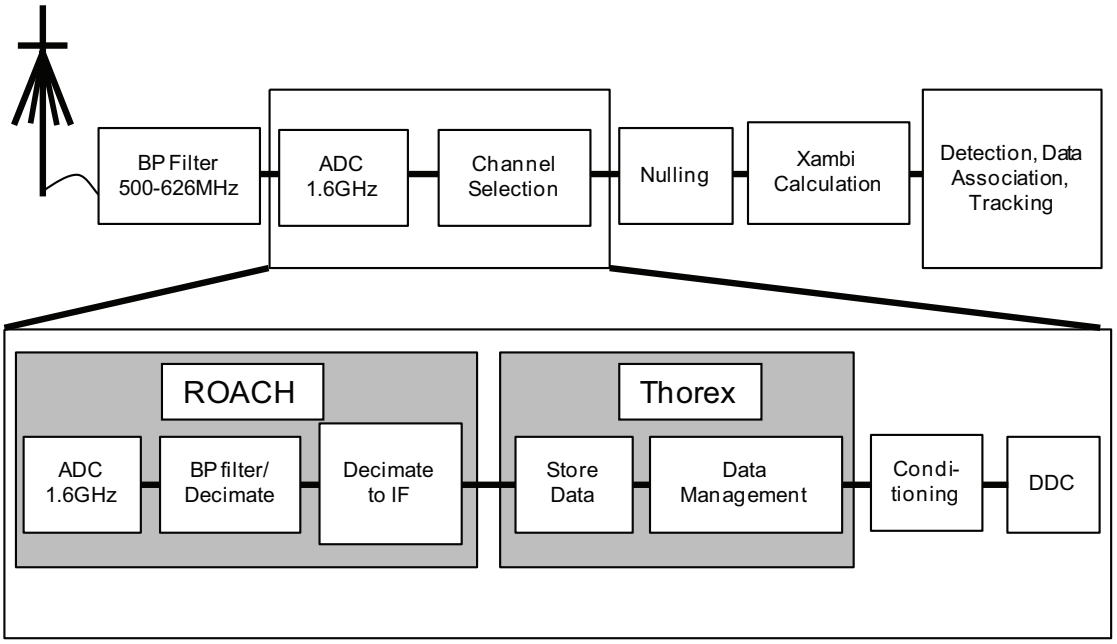

Figure 3.3: Detailed block diagram of the new receiver architecture for DTV signal reception.

3.3 FPGA Processing

The main purpose of the FPGA processing is to translate the desired channel to an intermediate frequency and reduce the data rate. There are two basic schemes for converting one frequency to another. The first is based on Fourier Transform Theory. Let $x[n]$ be a signal with Discrete Time Fourier Transform(DTFT) $X(e^{j\omega})$. Then

$$e^{-j\omega_0 n}x[n] \overset{DTFT}{\Longleftrightarrow} X(e^{j(\omega+\omega_0)}) \tag{3.1}$$

Multiplying the incoming data samples by a complex exponential is referred to as downconversion.

The other method is resampling or aliasing, a pure DSP function. If $x[n]$ is a digital signal and $y[n]$ is the resampled signal at new sampling rate f_0 such that $\omega_0 = 2\pi f_0$, each with respective DTFT $X(e^{j\omega})$ and $Y(e^{j\omega})$. With proper anti-aliasing filtering, these can be related by

$$Y(e^{j\omega}) = X(e^{j(\omega+N\omega_0)}) \tag{3.2}$$

showing that at baseband, the resampled signal has the frequency content of the original but shifted by an integer multiple of the new sample rate.

Both techniques require something like an anti-aliasing or image-reject filter.

3.3.1 An Example: Polyphase Decimation

The automatic 8-stream parallelization of the ADC samples into the FPGA lends itself directly to polyphase decimation. The polyphase filter bank and its use in downconversion is discussed in [23]. For this particular design, polyphase decimation is used instead of a traditional digital downconversion (DDC) for channelization of the data. Polyphase decimation is a computationally efficient way of performing filtering followed by downsampling. The first task is to downconvert the DTV channel desired from the fast 1.6 GHz rate to the 200 MHz FPGA clock speed. Using theoretical decimation, a stream of 1.6 GHz data, $x[n]$, passes through a full rate (1.6 GHz) digital band-pass filter, $h[n]$, to isolate the particular channel desired, and then downsample the output $y[n]$ by 8 (only save one of

every 8 samples!) to get the resulting $y[8n]$. It is computationally inefficient to compute all the intermediate values that simply do not get used. Polyphase decimation processes the data in parallel with low rate filters of different phase, allowing just the final low rate data to be calculated directly. However, since the FPGA receives data already parallelized from the ADCs to keep clock speeds low. Since the data is already parallelized, we can implement a polyphase filter directly by sending the 8 parallel streams of incoming data to the 8 sub-filters with varying phase. These are created by downsampling the full rate filter:

$$h_k[n] = h[8n + k] \tag{3.3}$$

By rerouting the input streams to their matching filter, we can add the outputs to directly create $y[8n]$, without wasting computation on the full rate output samples we would have thrown away. More detail on polyphase decimation is found in Appendix A.

Effectively we have band-pass filtered and resampled the DTV channel at a lower rate, aliasing our real-valued data down to 200 MHz. The data rate must still be reduced in order for us to transmit data off the FPGA on the 10 GbE link with the desired precision. If the initial filter was narrow enough, the data may simply be resampled again without the need for a second stage anti-aliasing filter. This is shown in Figure 3.4. This particular design samples two antennas at 1.6 GHz and isolates two different DTV stations with different data rates for transmission off the FPGA. This is one of our more complicated modes of operation.

Beginning with a single ADC (and single antenna) the interleaved mode produces 8 parallel streams of data which are routed into two separate polyphase decimating filters. These filters are very narrow, using many taps and in turn many of the FPGA resources, the intention being to perform all filtering in one stage. This is less efficient than creating two filtering stages, each with an incremental rate reduction. Performing filtering in a single stage requires a narrow filter with very steep edges. By breaking this into multiple stages, wider filters requiring fewer taps followed by small rate reductions can be cascaded resulting in fewer overall filter taps. The multi-stage model is used in subsequent designs.

Upon completion of the polyphase decimation, the data from the first DTV station is

truncated to 16 bits[1] and packed into a 64-bit word. One of every 14 samples is sent into the 10 GbE buffer for transmission, effectively performing the final data rate reduction. This means from DTV station 1, we save 14.2857 MHz real-valued 16-bit samples to disk on Thorex, perhaps while mobile. The DTV station is only 6 MHz wide, so there is still processing to be done to prepare it for cross-ambiguity calculation, however the FPGA does not have resources available for further computation. The data can be loaded into another computer for the final signal conditioning (removing the pilot tone as discussed in Chapter 2) and digital downconversion. Further details of this mode are covered in a late section on preparation of data for the 10 GbE link. First, a quick note about polyphase decimation versus digital downconversion.

3.3.2 Polyphase Decimation Technique Trade-Off

It is important to note that downconversion using decimation rather than a digital downconverter (DDC) requires careful attention to aliasing. This method will not work for all portions of the spectrum, and we chose our ADC sampling rate so several DTV stations of interest in the Seattle area would be feasible for this configuration. The goal of target localization in three dimensions would require at least three transmitters per antenna. The goal is to access as many different channels on as many antennas as possible with the existing 640 math operations per clock cycle of the FPGA, and push the rest of the downconversion to IQ onto another machine. The choice of the polyphase filter implementation was the first attempt, and it relied heavily on the spectrum being sparse, so filters could be wider and the affects of aliasing were small due to the lack of presence of transmitters in certain portions of the spectrum. This method will not work in all locations, and with mobility in mind resources on the FPGA must be devoted to sharper filtering. By implementing digital downconversion, we can easily access different parts of the spectrum, free from aliasing limitations due to polyphase resampling, and have been able to use resources more widely to guarantee 40 dB adjacent channel suppression. We have currently moved to modes that allow us one station on each of four antennas in this way, and have explored high aggregate

[1]The precision of the data samples grow throughout the filtering stages. Samples as large as 50-60 bits are possible.

bandwidth modes allowing up to 10 DTV stations on each of four antennas. Discussion of these modes is found in Chapters 5 and 6.

3.3.3 Data Preparation and Transfer to Disk: the 10 GbE Link

The amount of total decimation possible is limited by the amount of data we can send over the 10 GbE link[2]. The two DTV channels investigated are KZJO (536-542 MHz) and KONG (572-578 MHz). In the case of KZJO, we can simply decimate again by a factor of 14, to give a final data rate of 14.286 MHz. This is real-valued data, so the spectrum contains the 6 MHz DTV channel with only 1.286 MHz to spare - a very efficient use of bandwidth. Paying careful attention to the precision of the data, we end with a 16 bit sample from KZJO from each antenna. In the case of KONG, there was another DTV broadcast close by in frequency that was difficult to remove with the initial polyphase filter. To avoid more computation, we were able to relax our filter constraints to absorb this nearby station into the passband, and chose instead of 14 to decimate by 28 to avoid the aliasing of these two stations. This stream now has a final data rate of 28.572 MHz. By carefully limiting the data precision to 8 bits, we could send two streams of this data, each 7 clock cycles apart, into the 64-bit word required by the 10 GbE link. Each 64-bit word now contains two 16-bit samples of KZJO, and four 8-bit samples from KONG. We actually stream the data into the 64-bit word at the 200 MHz clock rate, and can control the final downsampling (anti-aliasing filter already applied) by only allowing every 14th sample into the 10 GbE buffer for transmission.

The data connection to 10 GbE port can be programmed onto the FPGA within the Simulink environment. The data to transmit must be stacked into a 64-bit word, which is sent to an internal buffer, which is drained over the link using a separate 156.25 MHz clock. This asynchronous transmission and buffering for streaming data makes it difficult to reach the full 10 GbE. Transmitting 64-bits at 14.286 MHz rate is a comfortable 0.914 Gbps, which the 10 GbE link handles easily. In-depth discussion of high aggregate bandwidth (6.4 Gbps) is found in Chapter 6.

---

[2]CX4 ports, IEEE 802.3ak

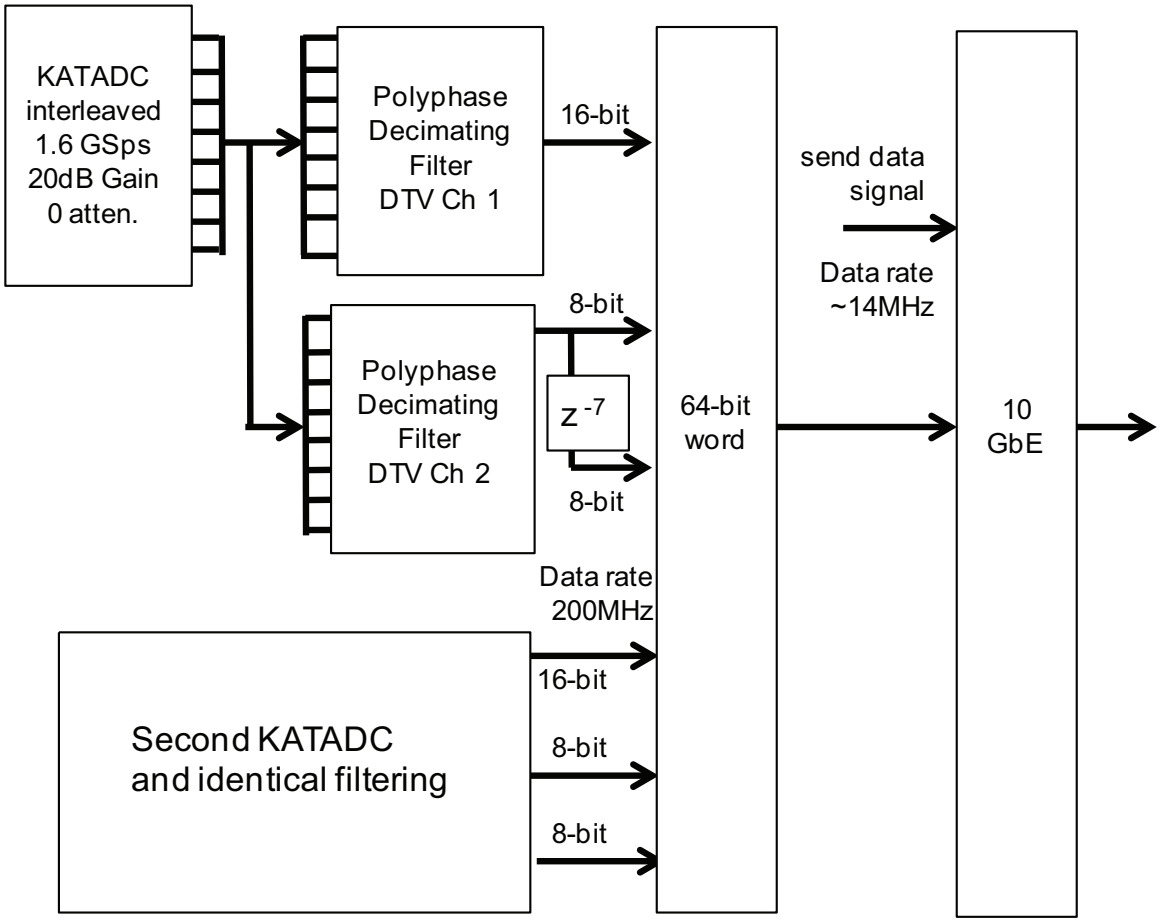

Figure 3.4: Block diagram of the processing on the Xilinx Virtex-5 FPGA. The dashed rectangle indicates the processes of conversion to baseband IQ data. The ROACH board completes the first few operations, while our control computer Thorex absorbs the data over a 10GbE link. The two machines together comprise the mobile portion of the signal processing.

3.4 IF Processing

3.4.1 IF Data Storage For Mobile Operation

The ROACH board connects to a second computer, which will be referred to as Thorex, via Ethernet and 10 GbE connections. This second computer is necessary for two purposes: 1) controlling and interacting with the ROACH board, and 2) storing the enormous amount of data streaming over the 10 GbE link. For this last purpose, we equipped Thorex with 1 TB of solid state hard drive space. Through Thorex we can program the ROACH board for a particular mode of operation, set variable parameters through software registers, and run a suite of monitoring and diagnostic tools.

The antennas, front end band-pass filters, ROACH board, Thorex, and a uninterruptible power supply (UPS), the radar can be used for up to 1 hour of continuous mobile operation. The entire set-up uses roughly 8U of a standard 19" rack space.

3.4.2 Signal Conditioning and Conversion to Baseband In-Phase and Quadrature

After data has been saved, we can download the files of real-valued IF samples to a computer for final processing into IQ form. Loading the data into Matlab, the first thing we do is convert the data from fixed to double precision floating point numbers. We first continue our channel selection by aggressively filtering out any adjacent noise and aliased channel stations to isolate the transmitter of choice. Filters are designed to suppress out-of-band signals by 40 dB. In this same step we perform the signal conditioning for ambiguity improvement by filtering a 1 MHz portion 6 MHz station, we can remove the pilot tone that causes long auto-correlation. To convert to IQ format we multiply the signal by $e^{-j2\pi f_0 t}$, where f_0 is the center frequency of the DTV station. The resulting data will be complex valued with the center of the DTV station at zero frequency. This frequency translation can be similarly done by resampling, as mentioned in Section 3.3. The ultimate sample rate with this conditioning will be 5 MHz.

3.4.3 Direct Path Suppression

There are several techniques involved with suppression of the direct path, to increase the detectability of our 0 dBsm target. One solution would be analog phased arrays and directional antennas to null the direct path before the voltage is sampled. This is an elegant solution, and would be straight forward for any stationary operation. For mobile operation, this front end SCR improvement may be difficult to manage as the antenna boresite will need to be mechanically steered in coordination with the changing geometry. Alternatively the MRR traditionally takes a different approach, which is to geographically separate the "reference" and "scatter" receivers, using distance/topography to suppress the direct path. MRR has superb, inherent direct path suppression. However the technique does not work for short range applications.

In order for the receiver to detect aircraft at any azimuthal angle, omnidirectional antennas are used. This leaves us with combining multiple antennas in post processing to do basic beamforming. The goal is to suppress the direct path interference by nulling the angle of the direct transmission. Attempts have been made here ([50]) using adaptive filtering. Even with our very basic two-antenna adaptive transmitter nulling, we see great improvement in the dynamic range of our receiver. Besides this additional gain from suppression, much gain is achieved simply through the correlation process. The math is detailed below in the Coherent Integration discussion, and the results of this nulling on real detections are shown in Chapter 7. We will now discuss the details of our basic beamforming.

With the direct path being the largest component in each, we simply assume that one one antenna feed will have a phase delay (if the antennas are close enough), and perhaps a magnitude difference due to some differences in the signal paths. We can model this in the IQ data as:

$$x_0[n] = \alpha x_1[n - \tau] \tag{3.4}$$

where $x_0[n]$ and $x_1[n]$ are the IQ signals corresponding to the same transmitter from antennas 0 and 1. α is a complex valued constant that contains the phase and magnitude difference between the two signals. Currently there is a random time offset between the two ADCs, and that can also be taken care of in this step with the τ variable, the lag

that maximizes the cross-correlation. If we use our two degrees of freedom, we can create a combined signal with the direct path removed to some degree.

$$y[n] = x_0[n] - \alpha x_1[n - \tau] \tag{3.5}$$

From here we can use this signal with the larger power of the scattered signal from the target and either of our original IQ signals $x_0[n]$ or $x_1[n]$ as the matched filter in our cross-ambiguity calculation. The removal of the direct path can be thought of as peeling its self-ambiguity off of the cross-ambiguity result. The floor will now be associated with the next largest peak on the cross-ambiguity plot.

Further adaptive-processing can be done to remove other large peaks in the cross-ambiguity, specifically large ground clutter. The overall drop of the cross-ambiguity floor, call this the Gain due to Nulling (G_N) in dB, allows the scattered signal-to-floor (SFR) ratio, and the dynamic range of the system, to increase by G_N. The goal is to be able to detect our 0 dBsm target at a range of 3 km with an SFR of 15 dB.

3.4.4 Cross-Ambiguity Operation Count

Chapter 2 detailed the cross-ambiguity function and its utility in passive radar target detection. Here we will discuss some of the complexities of the calculation. The first thing to note is that the cross-ambiguity calculation is extremely parallelizable in the range delay variable. The operation count broken down by stage of xambi processing for a single range slice of coherent integration is shown in Table 3.1. For a given range delay, the first step is a massive multiply of several numbers. Using DTV broadcast and looking particularly for airplanes, we use a very large decimation ($N_{dec} = 6626$ samples in a particular example) and an FFT size of 256 ($N_{fft} = 256$). Typically we look out to 120 range lags and compute 4 incoherent averages in a second. This gives a total count of 5.7 GFlop to process 1 second of data. Currently, this takes about 40 seconds in Matlab which could be reduced to 5-10 seconds in custom C code.

	Correlation	Decimation	FFT of $z_M[m, n_r]$(assume Radix 2)
Complex Adds		$N_{fft}(N_{dec}-1)$	$N_{fft}log_2N_{fft}$
Complex Mults	$N_{fft}N_{dec}$		$\frac{N_{fft}}{2}log_2N_{fft}$
Real Adds	$2N_{fft}N_{dec}$	$2N_{fft}(N_{dec}-1)$	$2N_{fft}log_2N_{fft}+2\frac{N_{fft}}{2}log_2N_{fft}$
Real Mults	$3N_{fft}N_{dec}$		$3\frac{N_{fft}}{2}log_2N_{fft}$
Total Float Operations	$5N_{fft}N_{dec}$	$2N_{fft}(N_{dec}-1)$	$2N_{fft}log_2N_{fft}+5\frac{N_{fft}}{2}log_2N_{fft}$
Total with numbers	8,481,280	3,392,000	9,216

Table 3.1: Flop count for xambi, single range slice. Correlation is the multiplication of the two received signals, $z[n, n_r] = y[n]x^*[n - n_r]$. Decimation is the coherent integration and reduction in data rate, $z_M[m, n_r]|_{m=Mn} = \sum_{k=0}^{M-1} z[Mn+k, n_r]$. Finally the FFT is computed to generate a single range slice. Multiplied by the number of ranges (120) gives the total count for 1 coherent integration (1.416 M), and multiplying by the number of incoherent averages (4) gives the total for 1 second of data processing. Typically, to process 1s of data in 1s of time, we need 5.7 GFlops.

3.5 Precision Gain, Processing Gain, and Dynamic Range Considerations in System Design

3.5.1 Front End Components

As we will see from the processing gain later in the system, the system is incredibly robust to front end performance limitations. This is an extremely nice feature of a system as high performance RF components are costly and can require human attention in terms of tuning and calibration. The current front end components of the system are the passive components (antenna, cables, a band-pass filter, attenuator) and the active components (amplifier, mixer, ADC). The band-pass filters are wide enough (126 MHz) to capture 5 powerful DTV stations in the Seattle Area.

Passive analog components are very linear and only become nonlinear in extreme conditions such as a lightning strike. Active analog components are inherently nonlinear, and care must be taken to insure radar operation within the their limits.

The antennas are connected to a lightning arrester, band-pass filter, and then the KATADC board which contains (in order) a 20 dB amplifier, variable attenuator, conversion to a differential signal, and finally the ADC.

The amplifier (RFMD SBB-2089Z[3]) is the first active component, with a passband from 50 MHz to 850 MHz. The 1 dB compression point is where the output power deviates from linearity by 1 dB, and for this particular amplifier is specified as typically 21 dBm with 400 MHz input. Nonlinearity can cause distortion due to third order terms. As input power increases, the power of this 3rd order harmonic rivals the power of the linear component (in log scale the slope is 3 compared to the linear slope of 1). The output power where this third order harmonic equals the linear term is called the output third order intercept, and is quoted at 41 dBm with 400 MHz input. In other applications these terms can be specified by their input power rather than output power. Dynamic range is maximized at the point where the power of the 3rd order harmonic just passes the power of the noise.

[3]The RFMD SBB-2089Z amplifier with passband 50-850 MHz has data sheet: http://www.rfmd.com/CS/Documents/SBB2089ZDS.pdf. Alternately, the KATADC can be ordered with the RFMD SBB-5089Z amplifier with passband 50 Hz - 6.0 GHz, with data sheet: http://www.rfmd.com/CS/Documents/SBB5089ZDS.pdf

If front end dynamic range is important, you want to operate in this sweet spot. While our system requires large dynamic range, the limitation on this is not the noise, but the difference between the desired scatter from a target and the powerful direct transmission. We can operate below this third order intercept without sacrificing dynamic range, especially because the noise will be greatly suppressed through digital signal processing. While weak signals will still be dominated by the quantization noise of the ADC, coherent integration allows us to drive this noise power down, increasing the available dynamic range. With the variable attenuator set to 0 dB, we achieve an input peak-to-peak voltage into the ADC, same as the output voltage from the amplifier, close to 50% of full scale peak-to-peak or 500 V out of 870 mV. This tells me our output power from the amplifier is about 7 dBm, and we can safely assume close to linear operation. Additionally this amplifier has a good noise figure, typically 2.6 to 3.1 dB up to 850 MHz.

The ADC can operate in single-edged sampling (non-DES) mode or dual-edged sampling mode (DES). Given the sample clock signal, the ADC will sample at every rising edge of the signal (former), allowing a maximum sample rate of 1.5 GHz or rising and falling edges (latter), extending the maximum sample rate to 3.0 GHz.

Analog-to-digital converters have many non-idealities, and the effects of the non-idealities are described by important metrics found in the data sheet. For our application, we care mostly about the resolution and frequency domain characteristics like harmonic distortion and spurious free dynamic range (SFDR). Due to nonlinearity, harmonic distortion occurs and the ADC output will contain spurious signals (spurs) at harmonics of the input frequency. The KATADC specifies the total harmonic distortion (THD) out to 9 harmonics, and shows that the total power in these harmonics compared to the power of the input tone ranges from -57 to -47 dB in the case for all input frequencies below 1.5 GHz. This is for single channel non-interleaved mode. For interleaved mode, a single typical value of -47 dB is given for interleaved mode sampling again at 1.5 GHz. The THD increases with clock frequency, and so far we are operating with a clock at 800 MHz in both interleaved and non-interleaved mode, well below this 1.5 GHz example they provide, although we may move to a 1.4 GHz clock.

Spurs can also be generated within the ADC, unrelated to the input signal. In interleaved

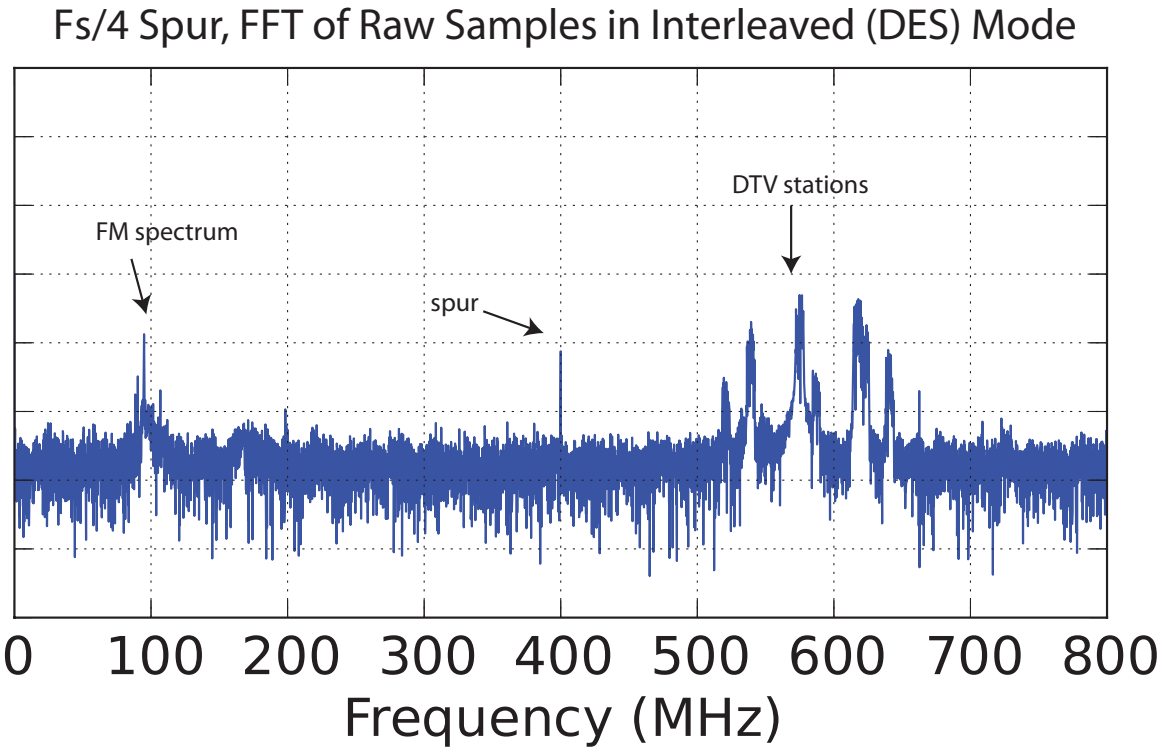

Figure 3.5: Spur at one fourth the sampling rate in the ADC interleaved sampling mode. With a sampling rate of 1.6 GHz, we display the FFT from 0-800 MHz. The FM spectrum is visible around 100 MHz, a large spur is visible at 400 MHz, and the DTV spectrum can be seen through the front end band-pass filter from 500-626 MHz.

mode, we have observed a strong spur at $F_s/4$, whose power is consistent regardless of input signal power. Figure 3.5 shows this spur for a single antenna with an ADC sampling rate of 1.6 GHz. The spurious free dynamic range (SFDR) is quoted at 55 dB in non-DES mode for a sample rate of 1.5 GHz with input frequency 748 MHz and in DES mode 47 dB with sample rate 3 GHz with input frequency 748 MHz. The poorer performance of the DES mode can be attributed to jitter between the even and odd samples, as they correspond with the rising edge and falling edge of the sample clock. This explains a feature occuring at $F_s/4$.

An important and encompassing metric is the effective number of bits (ENOB). The ENOB takes into account the ratio of signal to noise and distortion (SINAD), a more complete metric than SNR or SFDR. The KATADC has 7.4 ENOB, saying the ADC op-

erates equivalent to a perfect ADC of 7.4 bits rather than 8. Resulting in an SINAD of $6.02(7.4) + 1.76 = 46$ dB. This is a very high performance figure, and is plenty for our current application.

Experimentally, the KATADC has a nice calibration feature that is run on every power up and also by user command. With our software-controlled startup, we calibrate every time we turn the radar on to take data. This calibration sets the offset error, the full-scale error, the integration nonlinearity (INL), and differential nonlinearity (DNL) to maximize the SINAD and ENOB. In interleaved mode, spurs can still occur that require external adjustments beyond the self-calibration of the ADCs. The CASPER community has issued details on how to reduce these spurs even further [5].

We should do further testing to determine the effects of these spurs with our experimental conditions, although as you will see from the results in Chapter 7 we are achieving excellent dynamic range with our current operation. If in the future more sensitivity is desired with the same front end, we would have to put in some significant work in the post-processing algorithms before we reach this noise floor limitation.

The practice of sampling much faster than the frequency content of the signal of interest is called oversampling. A nice explanation can be found in a tech report by Atmel [14]. This allows us to choose an ADC with less resolution (fewer bits), because we actually gain bits through signal processing. A quick way of looking at this is that regardless of the sample rate, the quantization noise power stays the same. To compute an FFT using 1 second of data, however, an ADC with a faster sample rate will have a lower noise floor in the frequency domain, just by nature of having more data. For identical resolution ADCs operating at different sample rates, the power level of the signal of interest to the noise floor in the frequency domain will be greater for the faster sampler. If we choose to oversample, reduce the data rate, we have a chance to filter out some of this noise. With decimation (filtering and resampling), the faster rate samples can be converted to a lower rate data, but maintain that initial SNR savings. This can be quantified as literally adding more bits to this 7.4 ENOB of our ADC. In the case of our radar, we sample at 1.6 GSamp/s but we are only interesting in a 12 MHz band (the minimum we could get away with sampling our 6 MHz DTV station of interest), oversampling by a factor of 130. This adds an effective 3.5

bits (in the ideal case), meaning we operate our 7.4 ENOB ADC as a 10.9 bit ADC. This is the ideal case where we sample nothing but our signal of interest and white noise, so we expect the adjacent DTV transmitters and other received signals to degrade this slightly. In the following section we will detail the precision gain due to ADC oversampling.

This means, for a given antenna, there is a single RF path with a single digitizer. We can than examine as many transmitters as the digital data path and computational power permit. This makes our passive radar not an FM radar or a DTV radar as many of the current systems, but truly a versatile system that simply through software design can utilize any transmitter up to 1.5 GHz. Currently we have emphasized DTV, and we plan to run an experiment that looks at the entire FM spectrum (10 stations), and an experiment that looks at DTV and FM on a single antenna. However, the door is truly open to design modes including GPS and satellite communication systems as well with no hardware upgrades, just software development time.

It is actually optimal for the input signal to use less than the full scale range when considering saturation and quantization noise. Considering Gaussian input, the RMS signal voltage should be the full-scale voltage divided by k, and for an 8 bit ADC $k = 3.9$, about 25% of full-scale [51]. This means we hope to be utilizing on the order of 65 voltage bins, about 6 bits (5 bits amplitude).

An important note about using a fast digitizer is that sample clock phase noise or jitter is incredibly important. The total amount of jitter tolerated by the ADC before any jitter-induced degradation of performance is given by the ADC spec sheet [30]:

$$t_{J(max)} = \left(\frac{V_{INFSR}}{V_{IN(P-P)}}\right) \times \left(\frac{1}{2^{8.4} \times \pi \times f_{in}}\right) \tag{3.6}$$

For a conservative estimate, we can let the first term go to 1, with a maximum input frequency of 800 MHz and we get $t_{J(max)} = 1.2$ ps. The aperture jitter, or the noise associated with sampling at the correct time is a property of the ADC, quoted at 0.4 ps RMS. This allows for external sources to provide an RMS jitter of 1.1 ps. The jitter of the phase-locked oscillator we are currently using is 0.84 ps RMS, showing that we are not in danger of clock jitter affecting our ADC performance.

ADC linearity can be degraded by data that does not require use of all bit levels. The particular radar waveform associated with DTV has Gaussian structure, and more importantly, the aggregate spectrum up to 800 MHz (the signal we are sampling) is close to Gaussian. This inherently maintains ADC linearity similar to dithering [10].

Alternatives to this fast sample rate include band-pass filtering and undersampling, employed at the MRR as early as 1996 for FM broadcast [48]. To see a diverse range of transmitters on a single antenna, the RF path must be replicated with different band-pass filters and different ADCs. If not replicated, to select a different transmitter for a single RF path, analog band-pass filters may have to be swapped on the front end in order to observe illuminators across a wide frequency range in a single antenna feed. Alternatively, a tuneable band-pass filter technology could be used, but still the problem remains of simultaneously observing multiple transmitters. This band-pass filter method will only allow you to observe transmitters that fall entirely within frequency bands with boundaries of integer multiples of half your sample rate $f_s/2$, else aliasing will occur. To avoid this criteria (which becomes claustrophobic the lower the sample rate), a system will need to use an analog downconverter.

Mixers, as active RF components, are inherently nonlinear and inject spurs into the system. Spurs will occur at linear combinations of the input frequency f_{in} and the local oscillator frequency f_{lo}. Often most troublesome are the 3rd order nonlinearities. These spurious signals limit the dynamic range of the receiver by raising the effective noise floor, from spectral power scattered by the nonlinearity. Adding a mixer to the system can greatly increase the price (a nice programmable downconverter currently costs *half* the *total* cost of our system). The programmable downconverter for the ISIS system cost upwards of $10,000 and will handle only two RF inputs, provides two IF outputs, meaning a single antenna again may only be set up to observe transmitters in the band to which the mixer has been tuned, unless the RF path is split before the mixer. If more antennas *or* transmitters are desired, then a more expensive downconverter or duplicate downconverters must be procured.

Another downside to a mixer is simply the fact that it is another knob for the engineer to control. It's not a large downside, but can prove to be inconvenient, especially if the radar is designed to be operated by a non-expert end user or operate autonomously. It

would limit the ability for the radar to make "smart" decisions, such as switch to another transmitter given measurements it has taken. This sort of operation would be a possibility with this new system, however it will be beyond the scope of this PhD.

We can conservatively estimate the analog system temperature to be 2000 K based on the noise figure quoted earlier, and at the front of the ADC, we get a noise power of:

$$P_{n,sys} = kT_{sys}B \equiv -84 \text{ dBm} \tag{3.7}$$

Where k is Boltzmann's constant, B is 126 MHz, the passband of the analog filters. With full-scale voltage of the 8-bit ADC being 870 mV, and assuming a 50 Ohm load, we can estimate the quantization noise power to be:

$$P_{n,adc} = 1.92 \times 10^{-8} \text{ W} = -47.2 \text{ dBm} \tag{3.8}$$

This shows that the largest noise source is the ADC quantization noise rather than the system noise. A detailed analysis of the overall noise figure of the system is still needed, but it is important to note that a system temperature of $\approx$ 11,000,000 K would be required for the noise power to surpass the quantization noise power from the ADC. The surface temperature of the sun is estimated at 5,778 K.

We can come up with an estimate of the important signals and their ratio to noise power just after the ADC. We began Chapter 2 with a discussion of the SCR, where we refered to "signal" as the scattered power from the theoretical 0 dBsm target at 3 km range from the receiver and its ratio to the power of the clutter, the power from the direct path. At the ADC, both of these signals have been amplified by 20 dB and potentially attenuated. 20 dB amplification puts the peak-to-peak voltage of the nearly gaussian signal outside of the range of the ADC. Recalling that at the (isotropic, lossless) antenna, the power of the direct path clutter was $P_c = -11.5$ dBm and the power from the target was $P_s = -96.8$ dBm. Let us assume we set the attenuator to 20 dB, so the signal receives a total gain of 0 dB, putting the direct path safely within the limits of the ADC full-scale range. Therefore the CNR $\approx$ 35 dB and SNR ≈ -50 dB.

3.5.2 Precision Gain

Here we detail how the precision gain in the SNR, and identical gains in the CNR, grow at each stage of fixed point and floating point processing down to IQ data.

Precision Gain on the FPGA

Throughout the FPGA, the data undergoes stages of band-pass filtering. From a frequency domain perspective, this is equivalent to rejecting unwanted out-of-band signals such as adjacent transmissions and noise. The time-domain analysis, however, provides insight into how this filtering actually increases the SNR, and equally the CNR. Filtering is the act of scaling and adding several samples together. If signals in the selected band are coherent, the time domain samples added together will constructively add, suppressing any non-coherent signals. Low-pass filtering is an averaging operation. It preserves signals with DC offsets while suppressing higher frequency components. Band-pass filtering can be considered an averaging operation that preserves signals that in the desired frequency band. In this sense, coherent signals in the frequency range specified by the band-pass filter will get the same benefit of increasing in magnitude compared to the out-of-band noise and transmissions.

The overall CNR and SNR gain in the FPGA can be computed by considering an initial sampling rate of 1.6 GHz, that is reduced in bandwidth via band-pass filters to 32 MHz before transmission, a total reduction in bandwidth of a factor of 50, a total gain of 2.8 bits and an increase of 17 dB in CNR and SNR. The FPGA processing is diagramed in Figure 3.6, and the CNR and SNR are tracked through these stages in Table 3.2.

The filters must have deep enough side lobes that any undesired transmissions are suppressed down to the noise floor. When the rate change is applied after this anti-aliasing step, these transmissions will not affect the SNR and CNR growth.

First we note that when bandwidth is quoted in Table 3.2, we refer to bandwidth of the original sampled signal that has been preserved. Aliasing has occurred (only in areas of the spectrum that will later get filtered out) in several stages to reduce the data rate. This leads to the initially puzzling quote of 28.57 MSamp/s and 32 MHz of bandwidth in stage D.

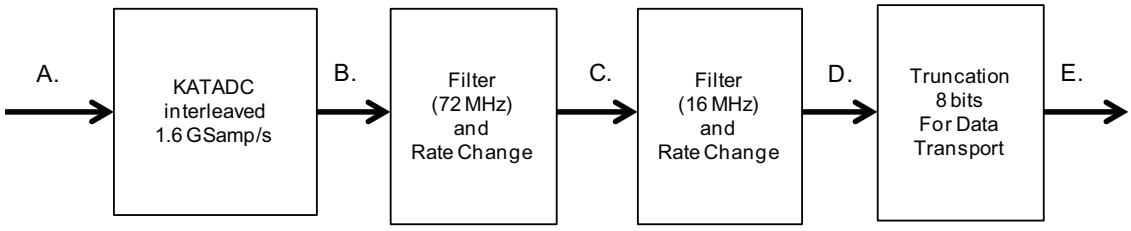

Figure 3.6: Block Diagram noting important stages for dynamic range gains (A-E) in the FPGA processing. Stages B-D are discussed in Table 3.2

The growth in CNR and SNR is also growth applied to all coherent signals within the final 32 MHz of bandwidth. It is important to note that there are limitations to the size of data that can be pushed off the board, hence the truncation in stage E. The growth of the individual data samples (which will include other transmitters not yet removed by filtering) must be truncated. Currently we limit these samples to 8 bits, and it is therefore important to consider the bit level growth of the sample itself as well the component of the sample associated with the DTV transmitter of interest and the received power from the target. As we filter, we remove noise as well as other powerful transmitters from the voltage level samples. With truncation to 8 bits, introducing a new quantization noise, the CNR and SNR have an upper limit of $6.02 \times 8 + 1.76 = 50$ dB, likely slightly less due to yet unfiltered high power transmitters.

Truncation can also be tricky to implement. It may not be optimal to slice the 8 MSBs from the 28.57 MHz samples. With small front end voltage into the ADC, the space held in the FPGA for bit-growth of the initial 8 bit samples may be largely unused. It is possible dynamically select the correct 8 bits which can remove the guess work and help in mobile conditions (changing signal power at the ADC) we always get the optimal 8 bits. These strategies are covered more in Chapter 4.

A front end low-noise amplifier would allow us to make use of the variable attenuator which currently is not needed. We could recalibrate the attenuators whenever our environment changes in mobile operation to maintain the desired voltage level into the ADC. This

	Stage		
	B.	C.	D.
Data Rate (Samp/s)	1.6 G	200 M	28.57 M
BW (Hz)	1.6 G	144 M	32 M
Reduction Factor		11.11	4.5
CNR (dB)	35	45.5	52
SNR (dB)	-50	-39.5	-33
Δ SNR (dB)		10.5	6.5
Δ SNR (bits)		1.7	1.1

Table 3.2: Fixed-point SNR gain on the FPGA. The bandwidth (BW) referred to here represents the band of the original sampled signal that has been preserved. In stage D, after the rate reduction, aliasing has occured but only in areas of the spectrum that will later get filtered out, explaining the 28.57 MHz data rate compared to the quoted 32 MHz "bandwitdh". This is by design to use fewer filter coefficients. CNR and SNR gain are related to the BW, rather than the data rate (as long as aliasing has not occured in the band of the direct path clutter or scattered signal). We gain 2.8 bits of precision in SNR (and CNR) for a total of 17 dB improvement.

would allow us to maintain a standard truncation method.

An interesting note: the cases during mobile operation where we are "shielded" from the direct path by a building may in fact reduce our dynamic range needs as the SCR will be much lower for targets with which line-of-sight is maintained.

Floating point precision gain: from FPGA to IQ baseband samples

After downloading the saved data, we must still convert the signal to IQ baseband samples, perform any conditioning such as removing parts of the DTV broadcast spectrum that have poor auto-correlation features, attempt to null the direct path, and then compute the cross-ambiguity. This is done currently via Matlab®(The Mathworks, Inc.) on a computer with 16 GB of memory. The major stages of processing are shown as a block diagram in Figure 3.7 and summarized in Table 3.3. For this we can use a "worst-case scenario" that our truncation gave us only a single bit of precision, reducing the CNR and SNR. The reduction in bandwidth in these stages gives us total gain in CNR and SNR of 8 dB or 1.3 bits.

We begin by assuming no loss in CNR or SNR by converting the fixed-point 8 bit

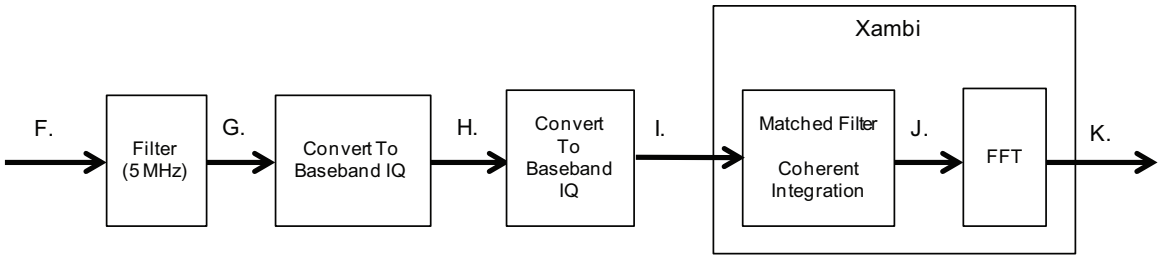

Figure 3.7: Block Diagram noting important stages (F-K) of floating point signal processing from the stored FPGA samples to IQ data. Stages F, G, and H are analyzed in Table 3.3.

	Stage		
	F.	G.	H.
Data Rate (Samp/s)	28.57 M	28.57 M	7.14 M
BW (Hz)	32 M	10 M	5 M
Reduction Factor		3.2	2
CNR (dB)	7.8	13	16
SNR (dB)	-77	-72	-69
Δ SNR (dB)		5	3
Δ SNR (bits)		0.8	0.5

Table 3.3: Floating point processing from IF to IQ at baseband. The total reduction of bandwidth gives us an SNR gain of 8 dB.

numbers to double precision floating point.

3.5.3 IQ to Xambi: Processing Gain

The IQ processing is shown in stages I, J, and K of Figure 3.7. First let us consider the case of no nulling, just a self-ambiguity calculation. We can define the process gain (PG) as the peak of the self-ambiguity of the transmitted waveform to the floor, for us we are limited to using the direct path reception. We can define the clutter-peak-to-floor ratio (CFR) as the ratio of the peak of the self-ambiguity of the received waveform to the floor. Similarly, we define our desired target scattered power peak-to-floor ratio (SFR). We also define N as the amount of decimation multiplied by the FFT size, $N = N_{fft}N_{dec}$. If we had a copy of the transmitted signal, we would be executing matched filtering. An echo of particular

range delay and doppler shift will grow coherently in voltage (N^2), while all other signals including the direct path, scatter from ground clutter, scatter from other moving targets and the quantization noise all add incoherently in voltage (N). The CFR is not limited by the quantization noise level as before, but by this total contribution of incoherent clutter noise. We expect the CFR, therefore to be on the order of $10log_{10}(N) = 62$ dB. In fact, every scattered component of the signal will contribute a peak at its particular range delay and doppler shift (for point targets), and its detectability will be determined by the height of its peak relative to the floor. It is quite difficult to detect moving targets on the zero doppler line as all of the ground clutter peaks reside there. The peak of our signal of interest in cross-ambiguity space will be 85 dB (the SCR) below the clutter peak. However, if the target is doppler shifted, we should be able to detect a positive SFR. We roughly state that an SFR of 15 dB will be enough for reliable detections.

With no nulling, if our 0 dBsm target is doppler shifted it will have an SFR of:

$$SFR = SCR + CFR = -85 + 62 = -23 \text{ dB} \tag{3.9}$$

In order to make this detection, we must remove some of these large and unwanted echoes that are causing this incoherent numerical noise in the IQ domain before be compute the cross-ambiguity. Every strong signal echo we remove can be thought of as peeling its two dimensional matched filter output off of the final result. This means we may reduce the power in the range, doppler cell where the peak of the signal occured by quite a bit, but the floor level will only change by the small amount of its incoherent contribution. We begin by simply attempting to remove the strongest return, the direct path signal. In the results section, you will see that we can reduce the peak at the 0 range and 0 doppler by up to 32 dB, however the floor only drops by 17 dB. The amount of suppression of the floor, G_N directly adds to the SFR.

In the case of nulling, the SFR can be estimated as:

$$SFR = SCR + CFR + G_N = -6 \text{ dB} \tag{3.10}$$

At the moment our floor suppression is simple, and only involves removing the first large 0 doppler component, the direct path. This is achieving a total $G_N \approx 17$ dB, as shown in Chapter 7. To drop the floor as much as possible, we must continue to remove any and all peaks larger than our target peak. If we can remove as much from the floor as we have from the 0 range 0 doppler component ($G_N = 32$ dB), we can hope to achieve an SFR of

$$SFR = SCR + CFR + G_N = +9 \text{ dB} \tag{3.11}$$

Given the addition of 4 antennas, we expect to do much better here.

We can basically see 3 important terms that lead up to our SFR = +15 dB goal: the SCR, the CFR or processing gain, and the gain from dropping the noise floor, G_N. The processing gain is more or less fixed, however we can turn to more antennas to help improve G_N, and we can still improve the SCR at the tip of the receiver antenna by using directional antennas, or by shielding ourself with buildings. This SCR was also a worst case scenario, and as you will see in the following chapter, our radar is currently detecting airplanes at closer ranges and demonstrates the dramatic effect of dropping the floor by even just a few dB.

The dynamic range of the system at this point can be estimated as:

$$DR = CFR + G_N = 79 \text{ dB} \tag{3.12}$$

and we have identified ways of moving this number closer to 100 dB.

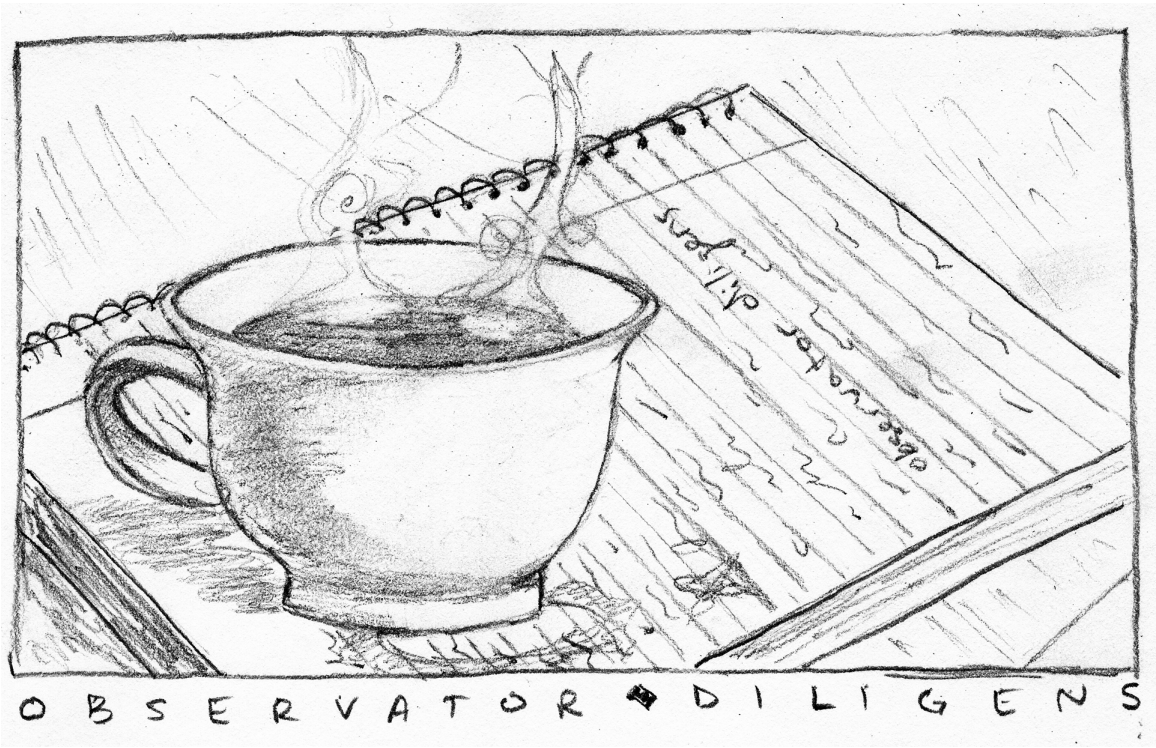
OBSERVATOR DILIGENS

Chapter 4

USABILITY

4.1 FPGA Design, Simulation, and Compilation

Programming the FPGA was facilitated with the Xilinx System Generator for DSP^{TM}, a tool that adds high-level blocks to Simulink®(The Mathworks, Inc.). The CASPER community has created their own libraries of Simulink blocks that use these Xilinx tools to create common functions for astronomy instrumentation. The ROACH board itself has a PowerPC running a stripped-down version of Linux called BusyBox[1]. The FPGA can be programmed entirely in the Simulink environment with the appropriate blocks, and then compiled to VHDL and ultimately to a binary image that can be loaded directly from the PowerPC. In this sense we can control and change the signal processing algorithm while mobile. Typical compile times are 2 hours, so it is important to save desirable configurations.

Simulink is a helpful tool for testing. Since much of our code is written in Matlab®(The Mathworks, Inc.), we can simulate the entire radar with an artificial set of ADC samples. The benefit of this is being able to export data from any point in the signal processing chain into a Matlab variable. We can isolate any component designed for the FPGA and work with the input and output data in the familiar Matlab workspace. Although our design is fairly simple, we encountered difficulties testing the polyphase filters, and with data truncation and management The Simulink environment allows us to specify a sequence of 8 bit 1.6 GSamp/s as if they were just emitted from the ADCs. We have found several helpful test sequences:

- Band-limited white noise
- Frequency sweep

[1] www.busybox.net

- Real Data

Band-limited white noise allows us to study the shape of the filters we have designed, and see how that response propagates through the FPGA. In all modes of operation, anti-aliasing filters are employed in conjunction with changing data rates. Tracking the effect of these changes and any aliasing can be done through simulation with white noise or frequency sweep input. This will also let us understand the effects of any truncation of the data as well as bit precision decisions for filter coefficients. The frequency sweep specifically allows us to study the overall transfer function of the entire FPGA processing. Using a small amount of fast rate ADC samples from our real system and propogating those through the simulation gives us a very important sanity check. Having much experience with the DTV spectrum by now, we can tell quickly that the mixers and filters are operating as expected by verifying the operation on known stations. From this processing and simulation, we are able to design a mode that will isolate a DTV station in a dense transmitter environment with a guarantee of 35 dB suppression of any similar-powered transmitters in the received spectrum, including frequency-adjacent transmitters. All this can be done without the need of analog equipment.

4.2 Programming the FPGA and interacting with the ROACH board

The ROACH board can be accessed via TCP or RS232 via the Karoo Array Telescope Control Protocol (KATCP) [9]. Using this package to interact with the PowerPC on the ROACH board, we control the radar via an external computer, Thorex, using Python scripts. This tool allows us to do various things such as program the FPGA, monitor the board temperature, monitor the FPGA clock rate, set the front end attenuator gain, read or write to any software registers created in the design, and access memory on the FPGA to read and plot data. These tasks can be split into initialization, a control center stage where software registers are adjusted for the particular experimental needs, and diagnosis of FPGA output.

4.3 Initialization: Front End and ADC control

4.3.1 Automatic Gain Adjustments and calibration on the KATADC

The KATCP protocol allows us to calibrate and set the ADC modes. This process is found in several of the Tutorial files, and thus we have only made minor adjustments to open source code provided to us for control of the KATADCs. The author of the original code is Jason Manley, at the time of writing is at the SKA-KAT (The Square Kilometre Array project in South Africa, using the Karoo Array Telescope) and the University of Capetown. The initial code and models provided from the Casper community have served as the skeleton of all of our designs. While the data processing we do is quite different, much of the control, monitoring, and testing code provided to us has been placed in our designs and relied upon. It is important again to mention a very hearty thanks to the open-source code and help of the community.

The initialization script begins with connecting to the ROACH board, programming the FPGA, and setting the value for the variable attenuators. This can be done iteratively by calculating the rms voltage value of raw ADC samples. This value is calculated constantly and stored in memory on the FPGA. By reading this value, code will increment or decrement the attenuation until desired rms value is reached. Quoted in Figure 4.1 is the script output for this process, noting that when the term "Gain" is refered to, it includes the fixed 20 dB front end amplification with the variable 0 to -31.5 dB attenuator, giving the overall gain a range from -11.5 to +20 dB.

4.3.2 ADC fast rate samples for diagnostics

Using the Python programming language, we can build a real-time display for the radar. For example, it is possible to extract ≈16,000 full rate samples from the ADCs, plot a histogram of these sample values, and generate a spectrum estimate. In interleaved (dual-edged) mode, we sample two antennas at 1.6 GSamp/s, and in non-interleaved mode we sample two antennas at 800 MSamp/s. The histograms from the two ADCs are shown in Figure 4.2 and Figure 4.3. These figures verify the voltage into the ADC is Gaussian. The calibration performed automatically at the start of the radar mode determined that the

```
Connecting to ROACH... done.

------------------------
Programming FPGA with fourant_100mhz_2013_Jun_13_1748.bof
________________________

Initializing Cores...

Checking clocks... ok, 200 MHz.
_________________________________
Auto-calibrating ANT A... done
Attempting automatic RF gain adjustment...
         Gain was -11.5dB, resulting in an ADC input level of
              -36.19dB. Trying gain of 4.69dB...
         Gain was 5.0dB, resulting in an ADC input level of
              -24.11dB. Trying gain of 9.11dB...
         Gain was 9.5dB, resulting in an ADC input level of
              -19.80dB. Trying gain of 9.30dB...
done!
Resetting counters for adc0... done
Current status: all ok for ANT A ADC0
_________________________________
```

Figure 4.1: Output of initialization script to connect to the ROACH board, program the FPGA, calibrate the ADCs and iteratively setting the programmable attenuators.

variable attenuator is not needed (and is set to 0 dB) to achieve optimal signal level into the ADC. The RF signal is unattenuated (perhaps 1 dB of insertion loss in the attenuator itself), then passes into the amplifier for a fixed gain of 20 dB. With this configuration, the RMS value of the input voltage is close to 25% of the full-scale amplitude. The only software control on this value is the attenuation. If the signal level drops we must live with this smaller SNR, however as seen in Chapter 3, sample precision is gained through signal processing, and a single bit suffices for detection of aircraft where no clutter rejection necessary.

In addition to the histograms we can compute a quick FFT in Python and display this to the screen. Figure 4.4 shows the spectrum from 0 to 800 MHz. We can see the effects of

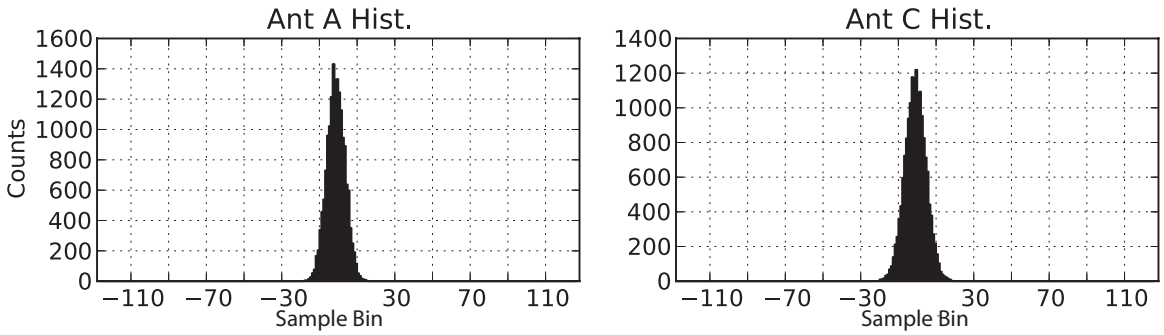

Figure 4.2: Histograms of 1.6 GHz samples in interleaved sampling mode for each of our two KATADC cards. The data shows the expected Gaussian amplitude distribution. Without attenuation we are achieving the optimal signal strength into the ADC when the antennas are located on the roof of Sieg Hall.

the analog band-pass filter from 500-626 MHz. They do a nice job of suppressing the rather large FM broadcast signals around 100 MHz. We can count about 6 DTV stations. Figure 4.5 shows the spectrum up to 400 MHz of the non-interleaved data sampled at 800 MHz. Here we see the DTV stations aliased due to the lower sampling rate.

As predicted in Chapter 3 and shown in Figure 3.5, spurs occur in interleaved sampling mode. Figure 4.4 shows a spur that we see consistently at 400 MHz in interleaved sampling mode, with power unrelated to input signal. With the polyphase processing, the data is immediately converted to 200 MHz and this spur aliases to DC. It is removed with low-pass filtering before the data is converted to IQ. Alternatively, with digital down-conversion implementation on the FPGA, this spur is filtered out immediately as it is outside the passband centered on the DTV station or FM station of choice. Figure 4.5 shows that sampling in non-interleaved mode reduces these spurs. This lower sampling rate will reduce our estimate of precision gain and SNR by a factor of 2 in power.

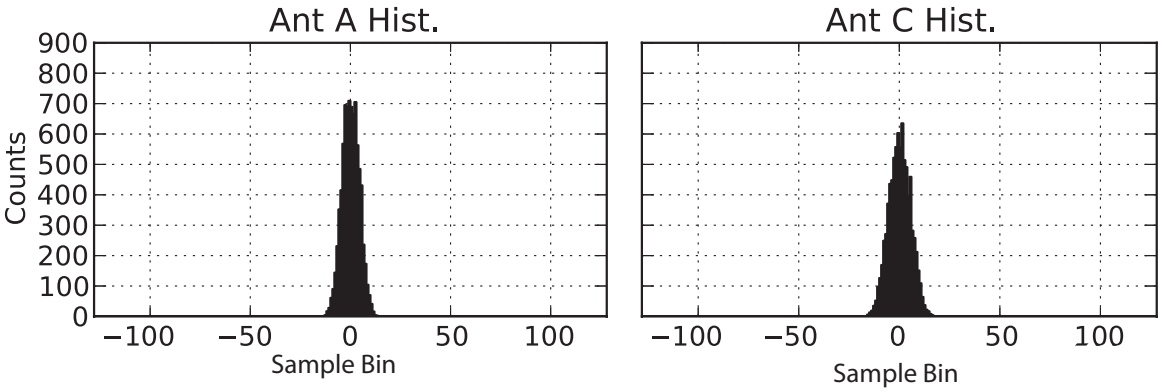

Figure 4.3: Histograms of 800 MHz samples in non-interleaved sampling mode for each of our two KATADC cards. Like the 1.6 GHZ interleaved sampling data in Figure 4.2, the data shows the expected Gaussian amplitude distribution. Without attenuation we are achieving the optimal signal strength into the ADC when the antennas are located on the roof of Sieg Hall.

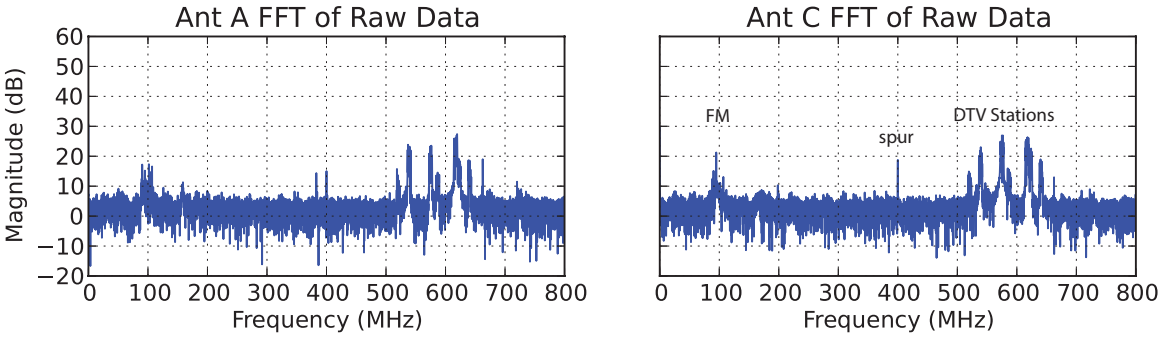

Figure 4.4: Live, real-time data from radar operation. One-sided spectrum of 1.6 GHz samples in interleaved sampling mode. We see roughly 6 DTV stations falling within our 500-626 MHz front end band-pass filter. We can also see the presence of spurs, notably one at 400 MHz.

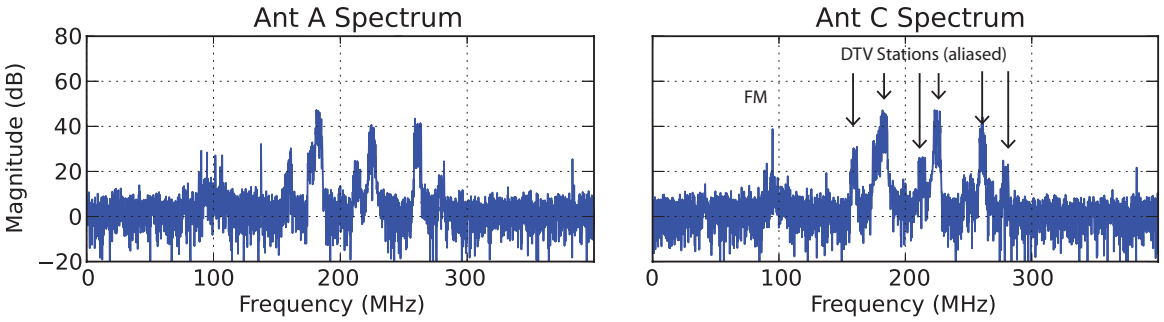

Figure 4.5: Live, real-time data from radar operation. One-sided spectrum of 800 MHz samples in non-interleaved sampling mode. We see the 6 DTV stations within the 500-626 MHz passband of the front end filter aliased due to the undersampling.

4.4 Control center

After an initialization script programs the FPGA, establishes ideal front end gain settings, and initializes the 10 GbE cores, a software register in the FPGA is set to finally allow the data to stream off the 10 GbE link.

While monitoring the front end, we can now access software registers on the FPGA. This allows the user to separate initialization of the radar (time consuming) from other decision making, such as setting the digital gain and the mixer frequency.

Setting the mixer frequency is fairly straightforward, however setting the digital gain is an iterative process. Through each stage of the FPGA processing, regardless of the mode, the strategy has been to let the sample size grow as needed, and truncate down to 8-bit samples as the very last step. The data can grow as large as a 57-bit sample, however the most significant 8 bits may very well only contain the sign of the sample value. Blind truncation of the data is not optimal. Used in previous MRR implementations, the Graychip GC4016 ([27]) digital receiver is an ASIC implementation of a DDC, and it requires considerable care in setting the gain and controlling the precision.

To optimize data quantization, we first simulate our design in Simulink and view the precision of the samples into this final quantization stage. After all processing, the full precision samples are signed (two's complement) fixed point numbers, 57 bits wide with a binary point at 54. The values will therefore cover a range from -4 to 4, with the smallest possible increment being 2^{-54}. The digital gain and quantization process from our Simulink model is shown in Figure 4.6. We process each full-precision sample into 16 parallel streams, each stream applying the Xilinx `scale` block, which will adjust the binary point of the data. The data receives a digital gain of 2^n where n is the stream number from 0 to 15. After scaling, each stream is passed through a block labeled with z^{-1}, a delay block used throughout the model to achieve timing closure. The data is then sent into a multiplexer. We take advantage of Casper's software register block to select the correct digital gain stream at the time of radar operation. The output of the multiplexer will be fed to a block that will cast the data properly as a signed 8 bit fixed point number with binary point at 5, thus still covering from -8 to 8 (saturation on any overflow), but quantizing so the minimum

increment between values is 2^{-5}.

Proper digital gain is set iteratively. Upon initialization, all software registers are set to the value zero, passing the top stream of parallel digital gain processing to the output. The data initially streaming off the FPGA contains 8-bit samples with a gain of $2^0 = 1$. For roughly gaussian data output, as is the case with our standard 14 MHz FPGA output, the user can simply save some of this streaming data and view the histogram of either the I or Q samples from each antenna. By examining the maximum and minimum values, we can measure the width of the peak and determine if digital gain is needed. This process can be automated with a `for` loop to take data, determine optimal digital gain, set the correct software register, and test the result. The expertise of a human operator to adjust and determine the digital gain by hand is always an option.

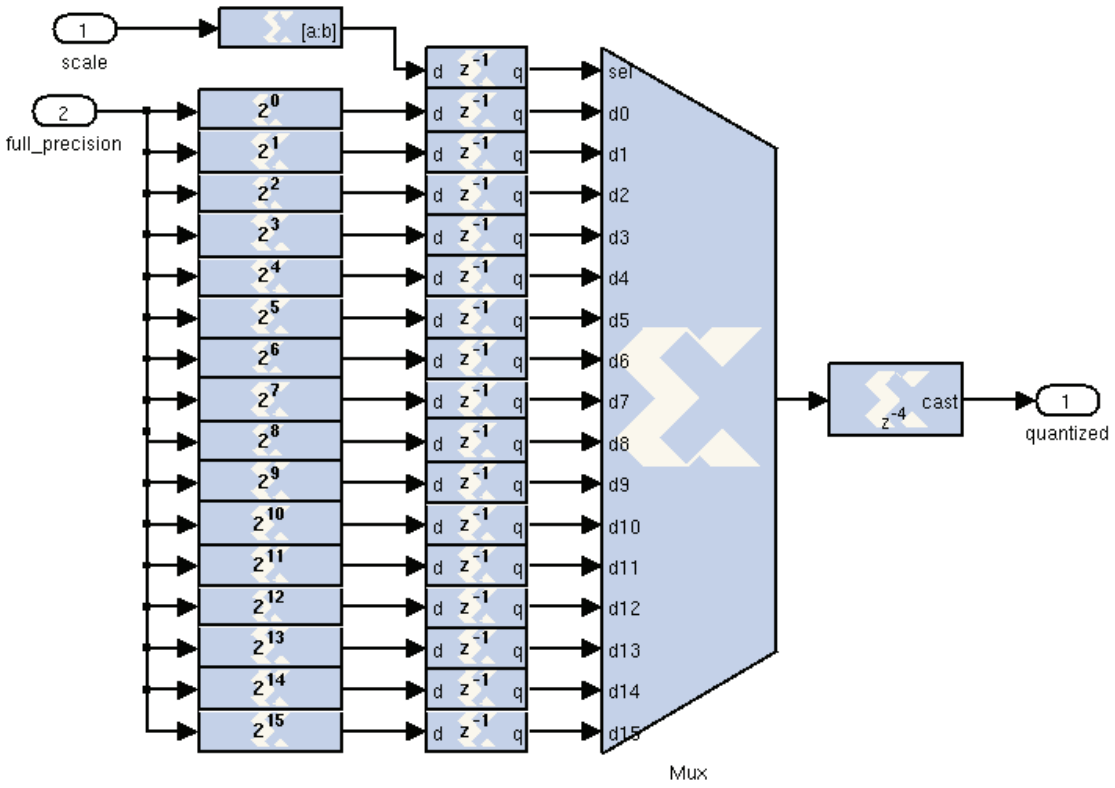

Figure 4.6: Snapshot from Simulink model file showing digital gain and quantization on the FPGA. A 32-bit software register can be set any time after the FPGA is programmed to provide an integer between 0 and 15 and is presented to the `scale` input. After mixing and data rate processing, full precision samples (57 bits with binary point at 54) are processed in parallel with the Xilinx Scale operator (no hardware cost), applying a digital gain of 2^n where n is an integer between 0 and 15. The blocks marked with z^{-1} are delay blocks, inserting latency in our model to achieve timing closure. The Multiplexer block (shown as the sideways trapezoid labeled `Mux`) passes the scaled data selected by the user. Finally, a `cast` block will cast the data as an 8 bit number with binary point at 5, set to saturate the data on any overrun.

4.5 Diagnosing the FPGA Output

After the the FPGA is programmed, the controls for digital gain and/or mixer frequency are set, the final stage of testing is to process the spectra of the data streaming off the FPGA and examine how well the radar is working. Dubbed simply `diagnose.py`, we use this code to check the data streaming off the FPGA to guarantee everything is operating as expected. Disk space and processing time for these data rates is significant; to know exactly what the data looks like and to be given a chance to correct before committing to a large data recording is important. Using Python, Numpy, and Scipy, many signal processing commands can be executed from the code. Common operations we use on the fly are data resampling and power spectral density calculations.

When operating in fast rate mode and transmitting 6.4 Gbps off the 10 GbE link, we isolate 8-bit samples from a 60 MHz window of the spectrum. By centering the mixer frequency at 557 MHz, both KZJO and KONG DTV stations fall within the passband. If the spectrum is full of DTV stations, then any one station may only represent 1/10th the total power in the samples off the FPGA. This alleviates our need to choose the appropriate 8 bits of the full precision samples at the end of the FPGA processing; as long as we do not saturate the data, the noise can be fairly large compared to the signal strength since so much of our precision gain will come entirely from processing after this stage. Monitoring a few of the fast, 800 MSamp/s data upon initialization of the radar is shown for all four antennas in Figure 4.7. KZJO at 539 MHz shows here at 261 MHz, and KONG at 575 MHz aliases to 225 MHz. When diagnosing the data streaming off the FPGA, it can be helpful to reference the front end data if anything looks suspicious.

Output of the diagnose program for this situation is shown in Figure 4.8. The DTV spectrum in Seattle is fairly sparse, so in this 60 MHz band of the spectrum only two DTV stations are isolated. On the left, centered at 539 MHz is KZJO. To the right, at 575 MHz is KONG.; note the 6 MHz bandwidth of each. We use 0.1 seconds of data to compute the power spectral densities shown, a typically small amount of data for these quick diagnoses. The second row of plots show histograms of the real part of the samples. Just read as 8-bit integers, the values range from -128 to +127. Starting with a shift of 0 bits or a digital

gain of 1, the histograms showed the data contained mostly 0s and 1s, hardly spanning the full 8 bits. Incrementing by a few bits we see a result here of a shift of 3 bits, or a factor of 8 in gain. The gain could be set more aggressively, but again so much precision gain will take place after this and the signal to ADC quantization noise ratio is quite small, that additional digital gain will not improve the results.

In Figure 4.7, the noise floor in the spectrum from antenna B looks different from the others. This may be an occasion where we examine the samples from the ADC to see if there is anything suspicious there. From exploring different digital gains, we can see that eventually the floor reaches the original quantization noise error from the ADC. The ADC samples confirm that antenna B receives KONG and KZJO with lower power, so this must be the effect of the ADC quantization. The spectrum on antenna C also shows interesting features, as if there are other stations next to KZJO that are not received on the other antennas. Upon close examination of the front end ADC samples, although not terribly visible in the still snap shot of Figure 4.7, these features appear to be there from the beginning. If we were not satisfied that the FPGA was behaving as expected, this stage would be an opportunity to exchange antenna cables and note any spectrum changes; coaxial cables fail remarkably frequently.

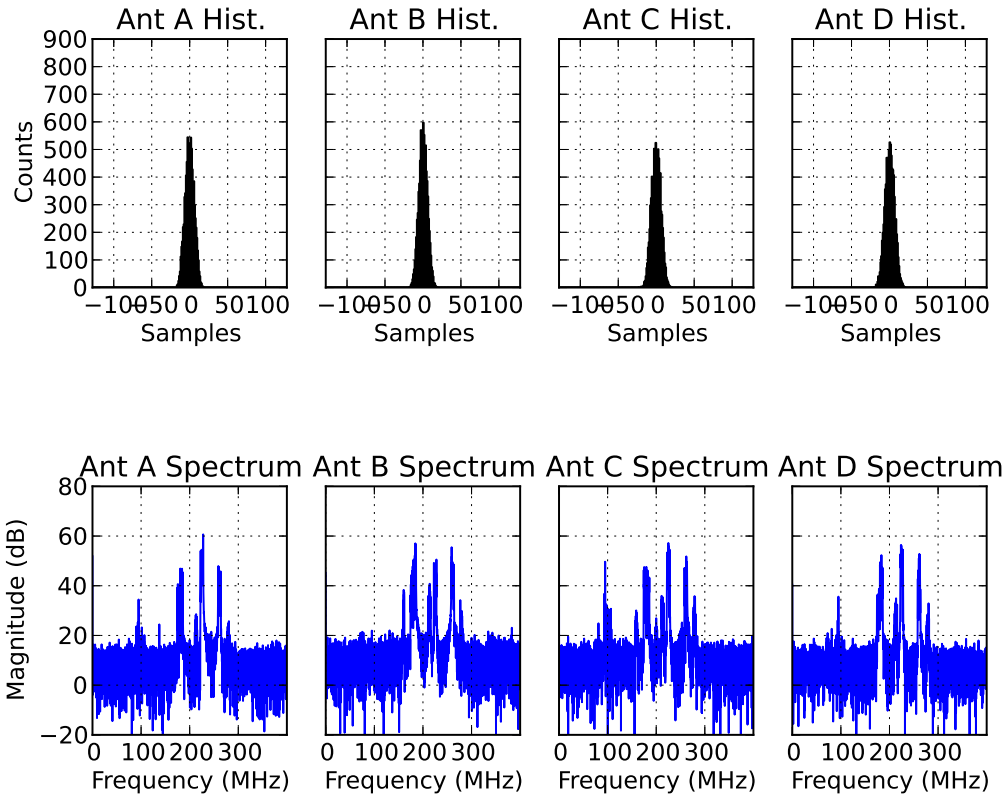

Figure 4.7: Monitoring the raw ADC samples at 800 MSamp/s on 4 antennas. The FPGA is designed to isolate 60 MHz of this spectrum in 8-bit IQ samples centered between two DTV stations: KZJO at 539 MHz and KONG at 575 MHz. These stations show up here at 225 MHz and 261 MHz, respectively, from aliasing.

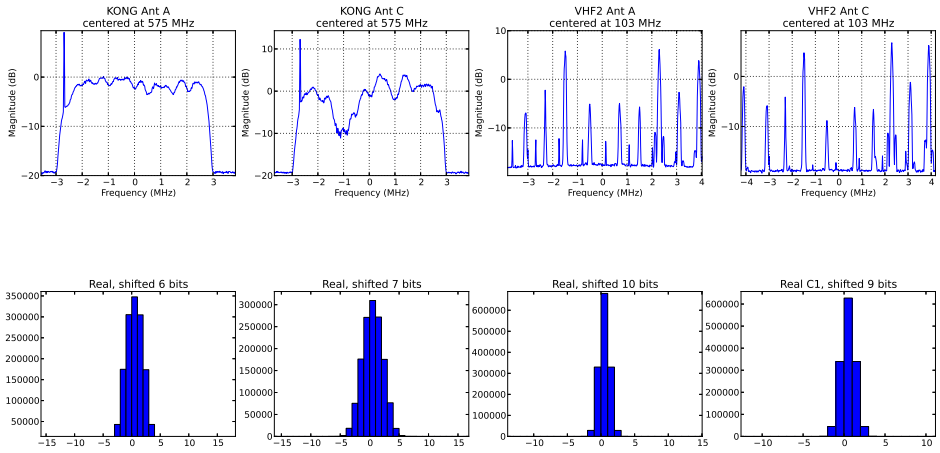

Figure 4.8: Diagnosis of 100 MHz (60 MHz passband) data saved to disk. The digital mixer on the FPGA is set to 557 MHz, centered between KZJO to the left at 539 MHz and KONG to the right at 575 MHz. The second row of graphs show histograms of the real part of data samples for diagnosing digital gain.

4.6 Mobile Operation Considerations

All of this code is run on the same computer, named Thorex, that stores the data from the 10 GbE link from the ROACH board. Additionally, Thorex connects to ROACH over dedicated standard ethernet for these communications. Since the radar is controlled and monitored with Python code with its open source license[2], there are no complications running Thorex and ROACH mobile in terms of accessing or verifying software licenses.

In addition to this open source convenience, software registers allow us to build in variations in operation that can accommodate a changing multipath environment. This is easily utilized while operating at a fixed location but offline, perhaps taking the radar to a new location to gather data for an hour off battery. All of the calibrations mentioned above can be re-done to account for these new surroundings. It may even be necessary to spend some time and care choosing different mixer frequencies to find transmitters in an unknown environment. Digital gain selection is then performed as usual.

In a truly mobile application it would be important to adjust the gain and frequency settings as the radio environment changes. At the moment, setting the front end attenuator gains on all four antennas in turn takes a total of 6 seconds. In a slowly changing multipath environment such as driving down a highway, it may be reasonable to occasionally trigger this function to optimize the transmitter strengh on the front end. In a fast changing environment, such as driving through a city surrounded by tall buildings, this may not be an option. If the user is carefully monitoring the front end gain and attenuator values, she can set the value specifically, rather than the iterative, automatic process, for a single antenna about as fast as she can type (actual time to execute command is a few milliseconds).

Adjusting the digital gain is another option when considering changing signal strength. To set the digital gain to a specified value, similar to front end gain is just a few milliseconds. It may be useful to stream more data from the FPGA without increasing the aggregate spectrum, in order to increase the bit depth of the final samples. Recall the maximum precision gain on the FPGA with 14 MHz final data rate is only 3 bits, meaning the data has at most 11 significant bits. If we increase the final data precision to 16-bit samples, the

[2]http://docs.python.org/2/license.html

system will be more tolerant of gain variation expected in mobile operation, without over- or underflowing the available precision.

These diagnostic tools, and the registers that can be set "instantly" during operation of the radar are large improvements in usability over the existing MRR system.

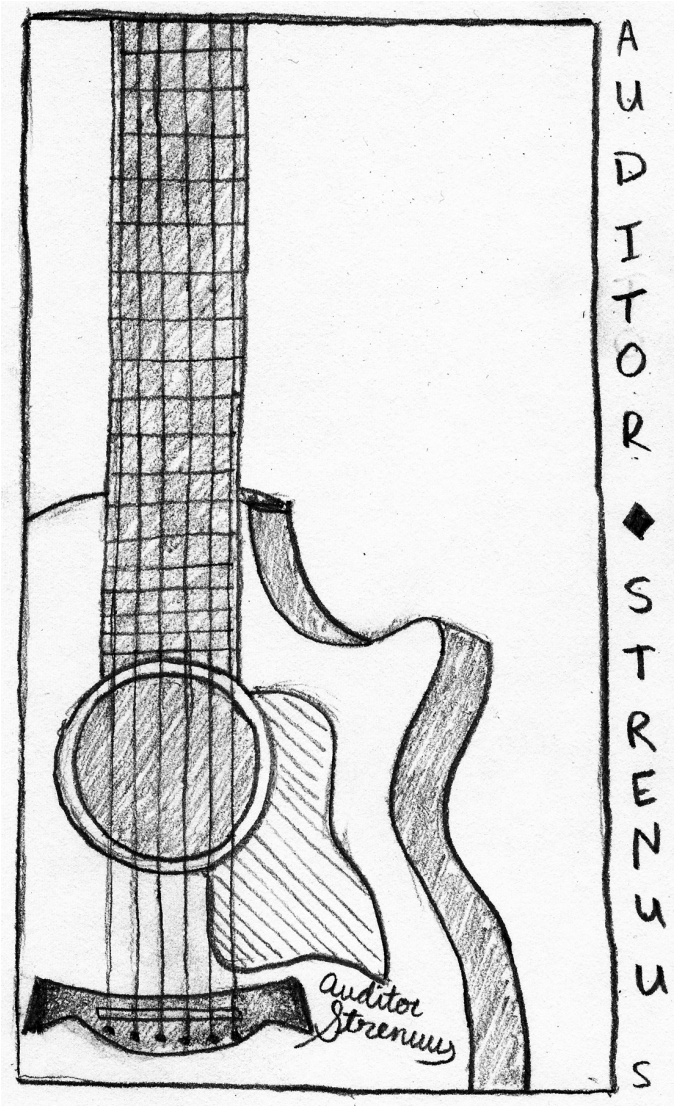
AUDITOR ◆ STRENUUS
auditor strenuus

Chapter 5

FPGA DIGITAL RECEIVER MODES

The FPGA is a tremendous complement to high speed digitizers in modern radar systems. This combination allows users to build several different test systems all with a one-time hardware purchase. With this freedom, we can use one to four antennas, understanding that extra resources can be put toward deeper filter sidelobes, more diagnostics, or more processing. We present briefly several radar modes that have been expressed in the FPGA. Some of these are not completely polished for use today, but could be quickly fit with the most recent developments in digital gain selection, ADC control, and UDP packetization.

This chapter includes block diagrams of the modes as well as a print out of the FPGA resource usage. While the block diagrams may indicate small changes between modes, in truth several small changes to the actual implementation occur on the FPGA, making it difficult to compare resources across the modes. Of particular imporance is the use of DSP48E slices[1] (highlighted in red). The DSP38E slice is available in all Virtex-5 devices, and allows for a 25-bit by 18-bit, two's complement multiply followed by an adder/subtractor stage.

This particular usage parameter can give the designer an idea of how many filter taps can be used in the design. These slices are not used one per filter tap, however, due to the fixed point nature of the slice and the current strategy of our designs to allow the data to grow in precision and quantize at the end for data transfer. A single large data value may need to be distributed across several of these slices to complete a filtering operation, putting strain on timing closure in addition to DSP resources. Filters can also be employed using distributed logic, and even in our most complex designs we still have many logic resources free to use. Re-quantizing the data earlier in the stream and implementing filters in logic as well as DSP slices are both strategies to fit a larger, more complex design on the FPGA.

---

[1] http://www.xilinx.com/technology/dsp/xtremedsp.htm

5.1 FPGA Resources

The KATADC cards, Virtex 5 FPGA (XC5VSX95T), and 10 GbE output link place several constraints on the digital receiver design. Some of these constraints are listed as follows:

- input data rate: maximum 3 GSamp/s, 8 bit samples for a total of 24 Gbps
- output data rate: maximum 10 Gbps
- parallel data samples input, 8x or 4x demultiplex
- 8.7 MB of RAM
- 640 DSP48E blocks consisting of a 25 x 18 precision multiplier, adder, accumulator

The following sections will provide block diagrams and present resource usage for several modes of radar operation expressed on the FPGA.

5.2 Polyphase Decimation, 1.6 GSamp/s, Two Antennas, Two 6 MHz spectral windows (2 DTV stations per antenna)

As described in Chapter 3, the first design for this digital receiver takes advantage of the sparsity of the DTV spectrum in Seattle and utilizes polyphase decimation, a simple re-sampling strategy to isolate two DTV stations on each of two antennas. A detailed block diagram is found in Figure 5.1. Despite the fact that the total number of filter taps in this design is on the order of 600, only 46 DSP48E blocks are used. The Xilinx DA_FIR filter blocks used in this Simulink design used distributed arithmetic as opposed to the DSP48E blocks. The resource utilization is catalogued below.

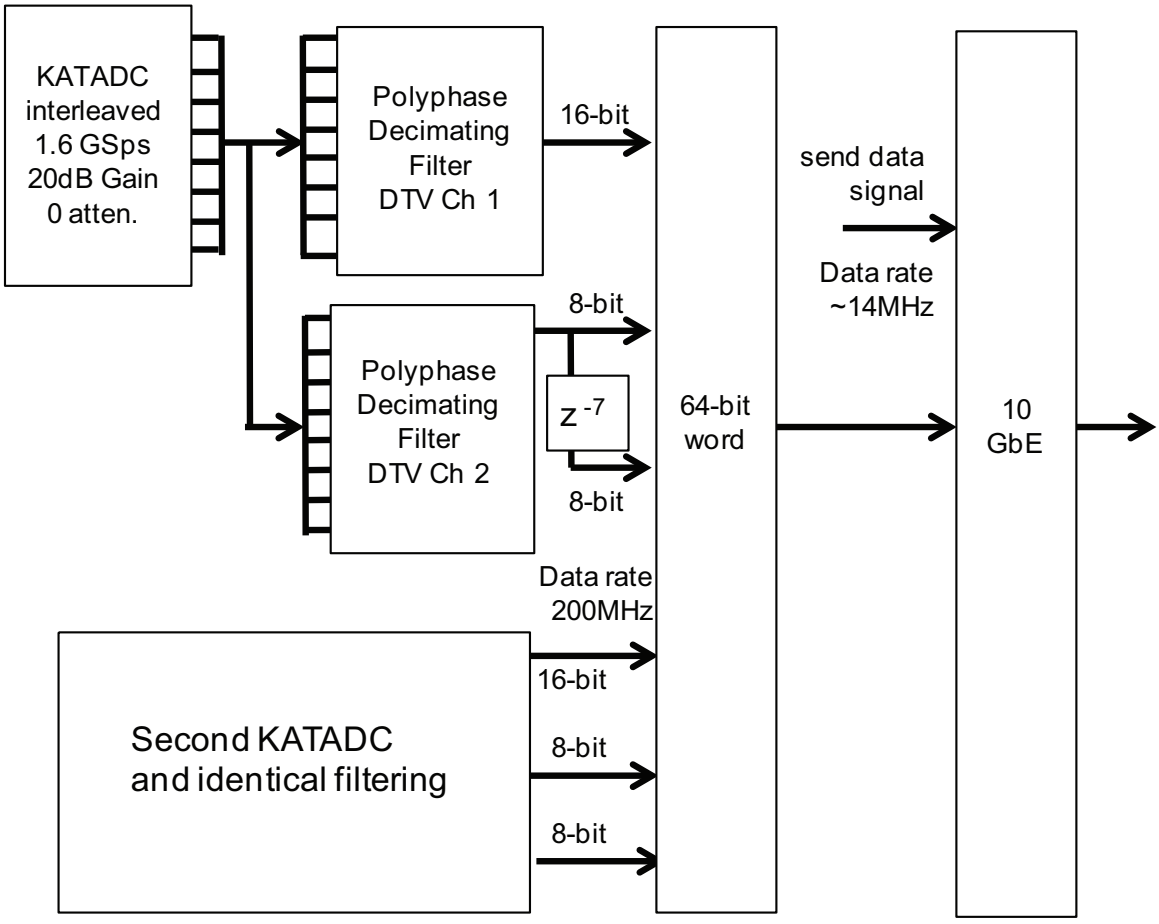

Figure 5.1: 2 antennas, 2 stations, 6 MHz with polyphase filtering

```
Slice Logic Utilization:
  Number of Slice Registers:                46,311 out of  58,880   78%
    Number used as Flip Flops:              46,310
    Number used as Latch-thrus:                  1
  Number of Slice LUTs:                     37,689 out of  58,880   64%
    Number used as logic:                   34,602 out of  58,880   58%
      Number using O6 output only:          33,261
      Number using O5 output only:             448
      Number using O5 and O6:                  893
    Number used as Memory:                   2,972 out of  24,320   12%
      Number used as Dual Port RAM:             780
        Number using O6 output only:            728
        Number using O5 and O6:                  52
      Number used as Shift Register:         2,192
        Number using O6 output only:         2,192
    Number used as exclusive route-thru:        115
  Number of route-thrus:                        798
    Number using O6 output only:                549
    Number using O5 output only:                249
Slice Logic Distribution:
  Number of occupied Slices:                13,141 out of  14,720   89%
  Number of LUT Flip Flop pairs used:       48,917
    Number with an unused Flip Flop:         2,606 out of  48,917    5%
    Number with an unused LUT:              11,228 out of  48,917   22%
    Number of fully used LUT-FF pairs:      35,083 out of  48,917   71%
    Number of unique control sets:              383
    Number of slice register sites lost
      to control set restrictions:              754 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                       16 out of     244    6%
    Number using BlockRAM only:                  16
    Total primitives used:
      Number of 36k BlockRAM used:               15
      Number of 18k BlockRAM used:                1
    Total Memory used (KB):                     558 out of   8,784    6%
  Number of BUFG/BUFGCTRLs:                      10 out of      32   31%
    Number used as BUFGs:                        10
  Number of BUFDSs:                               2 out of       8   25%
  Number of CRC64s:                               2 out of      16   12%
  Number of DCM_ADVs:                             3 out of      12   25%
  Number of DSP48Es:                             46 out of     640    7%
  Number of GTP_DUALs:                            8 out of       8  100%
  Number of PLL_ADVs:                             2 out of       6   33%
```

In fact, a single FIR subfilter of one of the polyphase decimating filters contains 20 coefficients. With the DA_FIR block using 16 bit integers with a binary point at 14, this breaks down into:

- Slices: 377
- Flip Flops: 1408
- BRAMs: 0
- Look Up Tables: 1003
- IOBs: 0
- Embedded Multiplies: 0
- TBUFs: 0

5.3 *Polyphase Decimation, 800 MSamp/s, Two Antennas, One 6 MHz spectral windows (1 DTV station per antenna), 8-bit IQ*

With mobility in mind, a redesign of the digital receiver is necessary to accommodate environments with a crowded DTV spectrum. This necessarily means diverting resources used to process a second DTV station per antenna to more aggressive filtering on one DTV station. The decision to move to a lower sampling rate, 800 MSamp/s, is made as moving to four antennas is a goal. The ADCs are now sampling in non-interleaved mode (or single-edged sampling mode), with two of the RF inputs left unprocessed. As can be seen in Figure 5.2, filtering is split into 2 stages: one operating on the fast 800 MHz data in parallel to reduce the data rate properly to 200 MHz, the second anti-aliasing filter operating on the 200 MHz data preparing for the final decimation for the 10 GbE link. Note that for this design, we chose to pack the data identically to the previous mode of operation. By producing the same data structure at the output of the ROACH board, this new design could be tested and integrated into our existing processing scheme while making no changes to further processing down the signal chain.

This design includes a more aggressive filters to guarantee 40 dB adjacent channel suppression. It is clear from the number of DSP48E slices (330 out of 640), the design cannot be doubled to include a second spectral window. This design specifically isolates the DTV station KONG.

```
Slice Logic Utilization:
  Number of Slice Registers:                17,534 out of  58,880   29%
    Number used as Flip Flops:              17,533
    Number used as Latch-thrus:                  1
  Number of Slice LUTs:                     22,258 out of  58,880   37%
    Number used as logic:                   11,001 out of  58,880   18%
      Number using O6 output only:           9,638
      Number using O5 output only:             338
      Number using O5 and O6:                1,025
    Number used as Memory:                  11,233 out of  24,320   46%
      Number used as Dual Port RAM:            780
        Number using O6 output only:           728
        Number using O5 and O6:                 52
      Number used as Shift Register:        10,453
        Number using O6 output only:         9,211
        Number using O5 output only:         1,242
    Number used as exclusive route-thru:        24
  Number of route-thrus:                       400
    Number using O6 output only:               361
    Number using O5 output only:                38
    Number using O5 and O6:                      1
Slice Logic Distribution:
  Number of occupied Slices:                 7,692 out of  14,720   52%
  Number of LUT Flip Flop pairs used:       25,278
    Number with an unused Flip Flop:         7,744 out of  25,278   30%
    Number with an unused LUT:               3,020 out of  25,278   11%
    Number of fully used LUT-FF pairs:      14,514 out of  25,278   57%
    Number of unique control sets:             439
    Number of slice register sites lost
      to control set restrictions:             806 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                      16 out of     244    6%
```

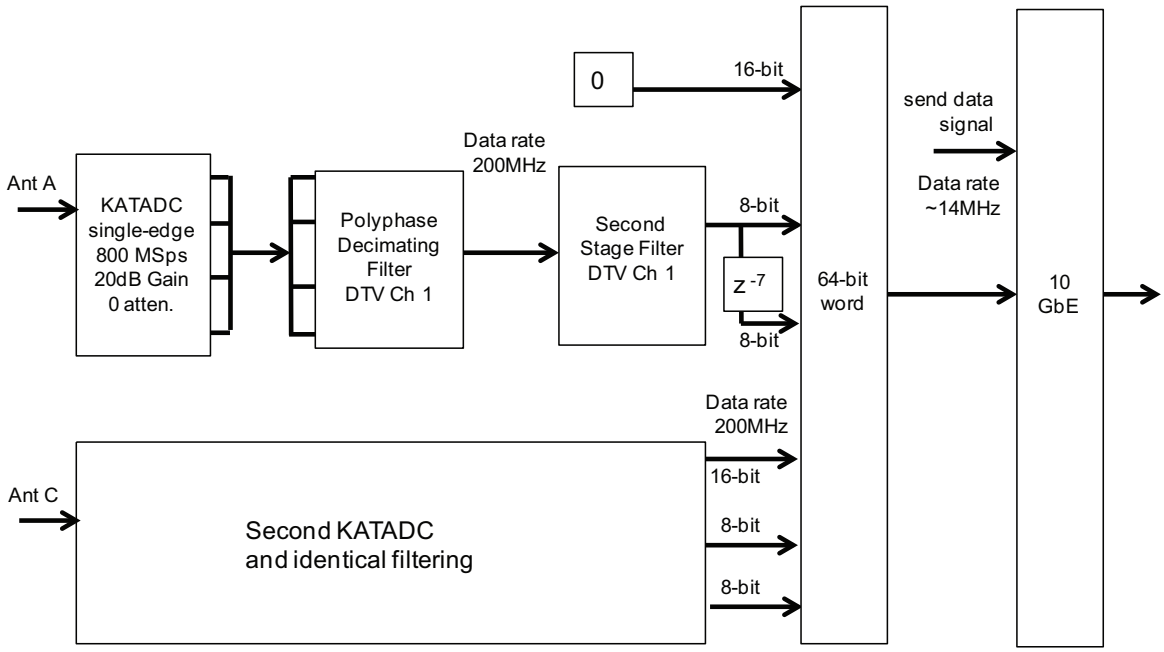

Figure 5.2: 2 antennas, 1 stations, 800 MSamp/s, 6 MHz window with polyphase filtering

```
   Number using BlockRAM only:                16
   Total primitives used:
     Number of 36k BlockRAM used:             15
     Number of 18k BlockRAM used:              1
   Total Memory used (KB):                   558 out of   8,784    6%
Number of BUFG/BUFGCTRLs:                     10 out of      32   31%
   Number used as BUFGs:                      10
Number of BUFDSs:                              2 out of       8   25%
Number of CRC64s:                              2 out of      16   12%
Number of DCM_ADVs:                            3 out of      12   25%
Number of DSP48Es:                           330 out of     640   51%
Number of GTP_DUALs:                           8 out of       8  100%
Number of PLL_ADVs:                            2 out of       6   33%
                  2 out of       6    33%
```

5.4 *Digital mixer, 1.6 GSamp/s, two antennas, two 6 MHz spectral windows (2 DTV stations per antenna) 8-bit IQ*

When comparing the mixer to the existing polyphase system, it was important to produce a similar 64-bit word. This guarantees that our processing after the 10 GbE link will not have to be changed to perform the test. In that sense, we begin with 1.6 GSamp/s data on two antennas. The signal is then split and mixed using two different frequencies, each centered on the DTV station of choice. The two-stage filters now become low pass with purely real coefficients that are identical for each DTV station, unlike the polyphase decimating strategy which required carefully designed filters for each station.

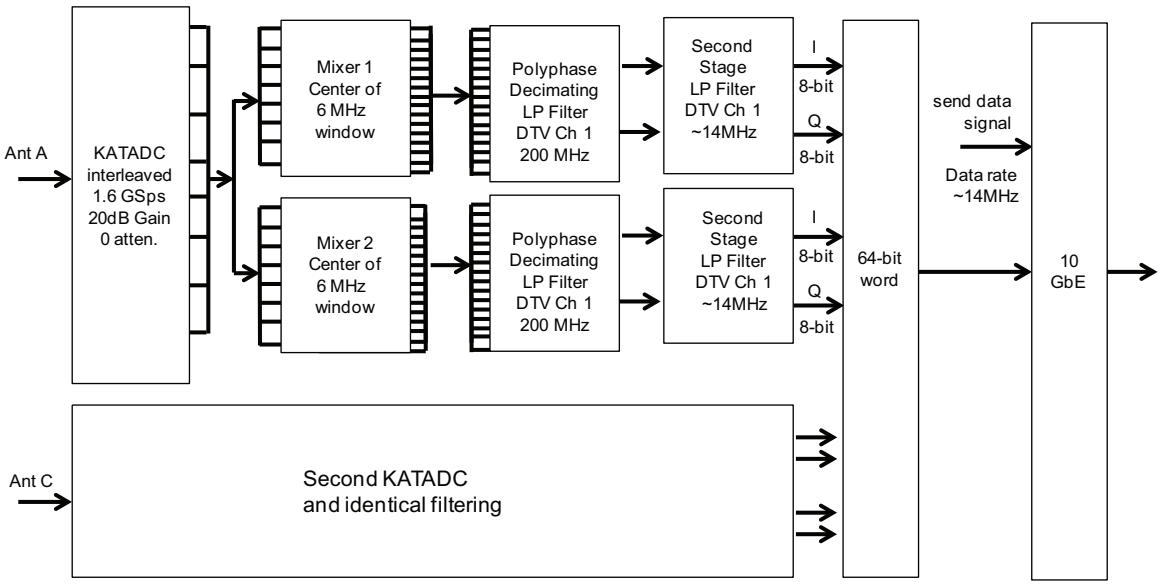

Figure 5.3: 2 antennas, 1 stations, 1.6 GSamp/s, 6 MHz window with polyphase filtering

```
Slice Logic Utilization:
  Number of Slice Registers:              17,769 out of  58,880   30%
    Number used as Flip Flops:            17,768
    Number used as Latch-thrus:                1
  Number of Slice LUTs:                   16,464 out of  58,880   27%
    Number used as logic:                 10,493 out of  58,880   17%
      Number using O6 output only:         9,207
      Number using O5 output only:           378
      Number using O5 and O6:                908
    Number used as Memory:                 5,788 out of  24,320   23%
      Number used as Dual Port RAM:            952
        Number using O6 output only:           900
        Number using O5 and O6:                 52
      Number used as Shift Register:       4,836
        Number using O6 output only:       4,836
    Number used as exclusive route-thru:       183
  Number of route-thrus:                       663
```

```
    Number using O6 output only:              557
    Number using O5 output only:              102
    Number using O5 and O6:                     4
Slice Logic Distribution:
  Number of occupied Slices:                6,963 out of  14,720   47%
  Number of LUT Flip Flop pairs used:      21,906
    Number with an unused Flip Flop:        4,137 out of  21,906   18%
    Number with an unused LUT:              5,442 out of  21,906   24%
    Number of fully used LUT-FF pairs:     12,327 out of  21,906   56%
    Number of unique control sets:            433
    Number of slice register sites lost
      to control set restrictions:            796 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                     20 out of     244    8%
    Number using BlockRAM only:                20
    Total primitives used:
      Number of 36k BlockRAM used:             15
      Number of 18k BlockRAM used:              5
    Total Memory used (KB):                   630 out of   8,784    7%
  Number of BUFG/BUFGCTRLs:                    10 out of      32   31%
    Number used as BUFGs:                      10
  Number of BUFDSs:                             2 out of       8   25%
  Number of CRC64s:                             2 out of      16   12%
  Number of DCM_ADVs:                           3 out of      12   25%
  Number of DSP48Es:                          282 out of     640   44%
  Number of GTP_DUALs:                          8 out of       8  100%
  Number of PLL_ADVs:                           2 out of       6   33%
```

5.5 *Four antennas, 800 MSamp/s, one 6 MHz spectral window, 8-bit IQ*

With the digital down-conversion test producing similar results to the previously successful polyphase resampling strategy, we focus on observing the same DTV station on each of four antennas.

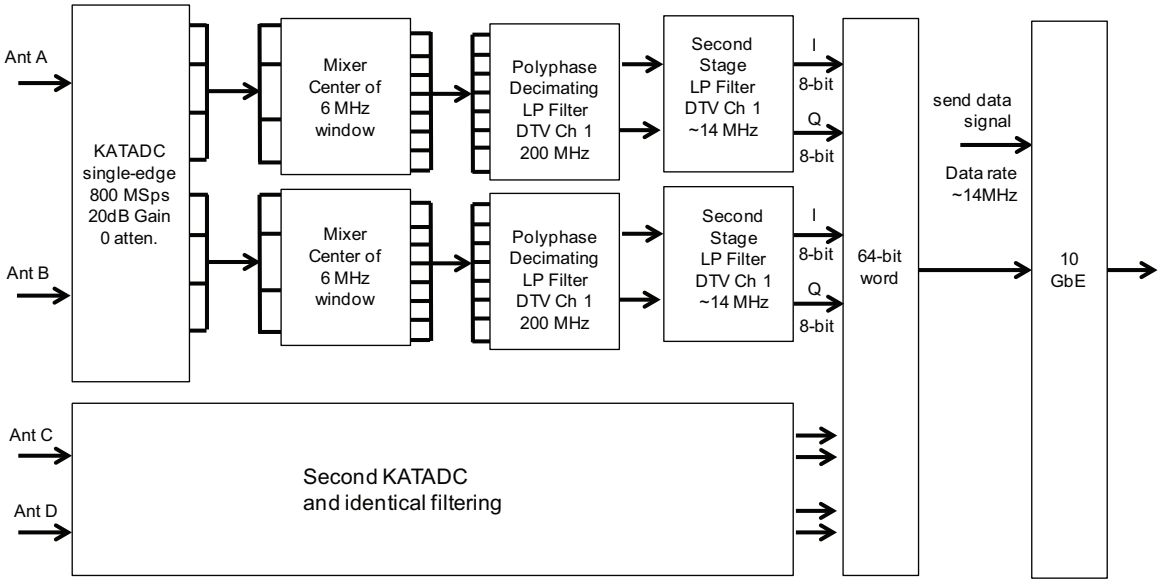

Figure 5.4: 4 antennas, 1 stations, 800 MSamp/s, 6 MHz window with polyphase filtering

```
Slice Logic Utilization:
  Number of Slice Registers:                38,599 out of  58,880   65%
    Number used as Flip Flops:              38,598
    Number used as Latch-thrus:                  1
  Number of Slice LUTs:                     30,337 out of  58,880   51%
    Number used as logic:                   19,061 out of  58,880   32%
      Number using O6 output only:          17,506
      Number using O5 output only:             351
      Number using O5 and O6:                1,204
    Number used as Memory:                  10,435 out of  24,320   42%
      Number used as Dual Port RAM:            952
        Number using O6 output only:           900
        Number using O5 and O6:                 52
      Number used as Shift Register:         9,483
        Number using O6 output only:         9,483
    Number used as exclusive route-thru:       841
  Number of route-thrus:                     1,685
    Number using O6 output only:             1,028
    Number using O5 output only:               510
    Number using O5 and O6:                    147
Slice Logic Distribution:
  Number of occupied Slices:                12,728 out of  14,720   86%
  Number of LUT Flip Flop pairs used:       44,274
    Number with an unused Flip Flop:         5,675 out of  44,274   12%
    Number with an unused LUT:              13,937 out of  44,274   31%
    Number of fully used LUT-FF pairs:      24,662 out of  44,274   55%
    Number of unique control sets:             445
    Number of slice register sites lost
```

```
      to control set restrictions:             819 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                      35 out of     244   14%
    Number using BlockRAM only:                 35
    Total primitives used:
      Number of 36k BlockRAM used:              15
      Number of 18k BlockRAM used:              33
    Total Memory used (KB):                  1,134 out of   8,784   12%
  Number of BUFG/BUFGCTRLs:                     10 out of      32   31%
    Number used as BUFGs:                       10
  Number of BUFDSs:                              2 out of       8   25%
  Number of CRC64s:                              2 out of      16   12%
  Number of DCM_ADVs:                            3 out of      12   25%
  Number of DSP48Es:                           558 out of     640   87%
  Number of GTP_DUALs:                           8 out of       8  100%
  Number of PLL_ADVs:                            2 out of       6   33%
```

5.6 Four antennas, 800 MSamp/s, one 60 MHz spectral window, 8-bit IQ

One way to access more of the spectrum without straining the FPGA resources is simply to push higher data rates off the FPGA. This shifts part of the processing load to other machines. If real-time is not an immediate concern, then this is a good strategy. Instead of focusing on a 6 MHz spectral window and transmitting 0.8 Gbps across the 10 GbE link, we increase the data rate to 6.4 Gbps, allowing us to look at a 60 MHz spectral window. Note: the data rate off the FPGA is really 100 MSamp/s, but we can relax the filtering constraints on the FPGA (to use fewer taps) by only preserving 60 MHz of the original spectrum and allowing some aliasing into the remaining 40 MHz. Careful spectrum management, knowing where you can tolerate aliasing, is critical to relaxing constraints on filters and maintaining high aggregate bandwidth. The block diagram shown in Figure 5.5 differs only from Figure 5.4 in the final decimation. This change, while small, was not without technical challenges. For one, creating the clock signal to drive data into the 10 GbE link must be done with care. We provide elaborate details on generating this signal in Chapter 6.

With 60 MHz of the spectrum isolated on 4 antennas, and even with the sparsity of the spectrum as seen at the UW in Seattle, the mixer can be set to include both KONG and KZJO, the original two DTV stations of interest in the band. To use this data rate, we did not have to sacrifice precision on the 8-bit IQ samples. Depending on the application, more of the original spectrum can be preserved by sacrificing sample precision, as we saw in our original polyphase filtering technique where we transmitting 16-bit IQ samples of one DTV station off the board at 7 MHz, while simultaneously transmitting 8-bit IQ samples of another DTV station at 14 MHz.

```
Slice Logic Utilization:
  Number of Slice Registers:                28,581 out of  58,880   48%
    Number used as Flip Flops:              28,580
    Number used as Latch-thrus:                  1
  Number of Slice LUTs:                     20,660 out of  58,880   35%
    Number used as logic:                   15,915 out of  58,880   27%
      Number using O6 output only:          14,529
      Number using O5 output only:             303
      Number using O5 and O6:                1,083
    Number used as Memory:                   3,909 out of  24,320   16%
      Number used as Dual Port RAM:            952
        Number using O6 output only:           900
        Number using O5 and O6:                 52
      Number used as Single Port RAM:          288
        Number using O6 output only:           288
      Number used as Shift Register:         2,669
        Number using O6 output only:         2,669
    Number used as exclusive route-thru:       836
  Number of route-thrus:                     1,479
```

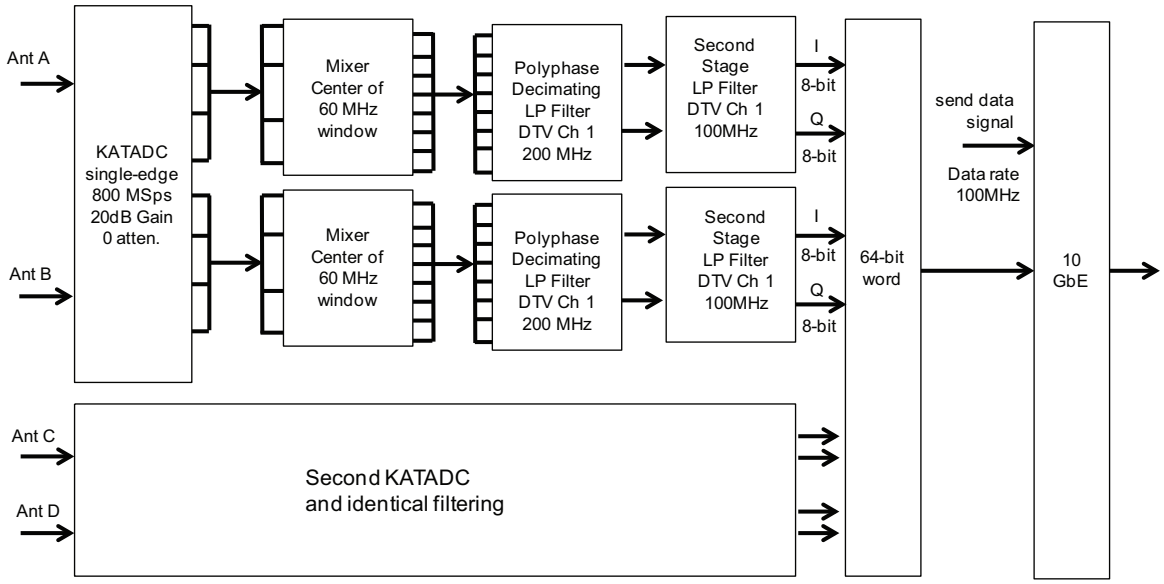

Figure 5.5: 4 antennas, 800 MSamp/s, 4 60 MHz window with polyphase filtering

```
    Number using O6 output only:             1,024
    Number using O5 output only:               342
    Number using O5 and O6:                    113
Slice Logic Distribution:
  Number of occupied Slices:                10,868 out of  14,720   73%
  Number of LUT Flip Flop pairs used:       35,032
    Number with an unused Flip Flop:         6,451 out of  35,032   18%
    Number with an unused LUT:              14,372 out of  35,032   41%
    Number of fully used LUT-FF pairs:      14,209 out of  35,032   40%
    Number of unique control sets:             454
    Number of slice register sites lost
      to control set restrictions:             855 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                      35 out of     244   14%
    Number using BlockRAM only:                 35
    Total primitives used:
      Number of 36k BlockRAM used:              15
      Number of 18k BlockRAM used:              33
    Total Memory used (KB):                  1,134 out of   8,784   12%
  Number of BUFG/BUFGCTRLs:                     10 out of      32   31%
    Number used as BUFGs:                       10
  Number of BUFDSs:                              2 out of       8   25%
  Number of CRC64s:                              2 out of      16   12%
  Number of DCM_ADVs:                            3 out of      12   25%
  Number of DSP48Es:                           230 out of     640   35%
  Number of GTP_DUALs:                           8 out of       8  100%
  Number of PLL_ADVs:                            2 out of       6   33%
```

5.7 Two antennas, 1.6 GSamp/s, two 60 MHz spectral windows, 8-bit IQ

By increasing the ADC sampling rate to 1.6 GHz, we observe the spectrum up to 800 MHz unaliased. This allows us to isolate both FM and DTV broadcast from the same antenna. We show here the results of fast data transfer, allowing us to observe the entire FM band from 88-108 while simultaneously observing 10 frequency-adjacent DTV stations. This can be simply adapted to a standard data rate, allowing for 1 DTV station isolated at the same time from the same antennas as 6 MHz of the FM band.

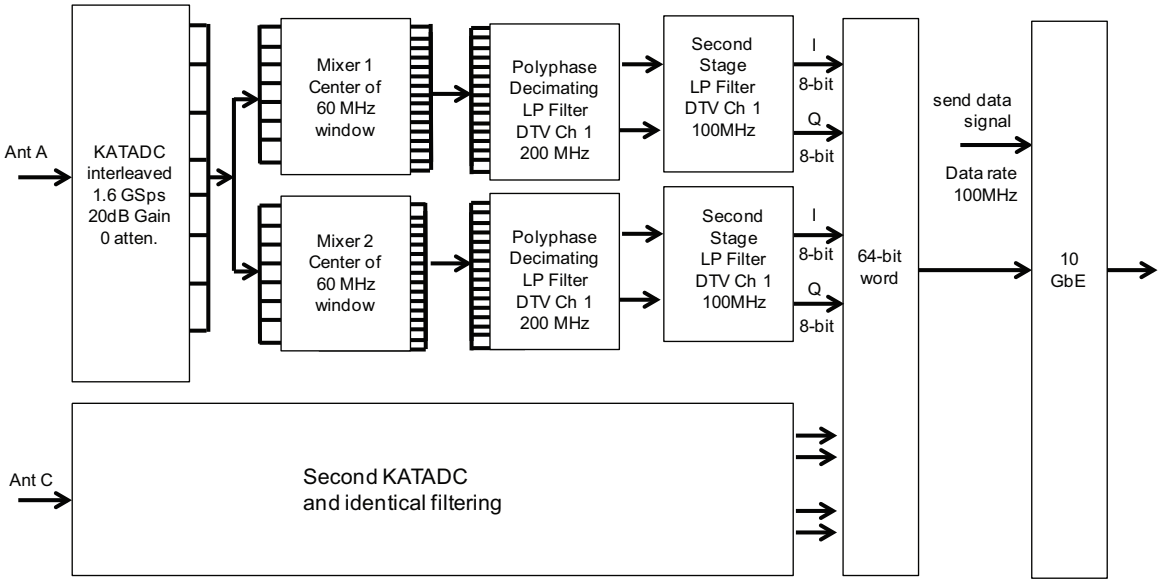

Figure 5.6: 2 antennas, 1.6 GSamp/s, 4 60 MHz window with polyphase filtering

```
Slice Logic Utilization:
  Number of Slice Registers:                35,871 out of  58,880   60%
    Number used as Flip Flops:              35,869
    Number used as Latch-thrus:                  2
  Number of Slice LUTs:                     23,345 out of  58,880   39%
    Number used as logic:                   17,155 out of  58,880   29%
      Number using O6 output only:          15,939
      Number using O5 output only:             308
      Number using O5 and O6:                  908
    Number used as Memory:                   4,613 out of  24,320   18%
      Number used as Dual Port RAM:            952
        Number using O6 output only:           900
        Number using O5 and O6:                 52
      Number used as Single Port RAM:          296
        Number using O6 output only:           296
      Number used as Shift Register:         3,365
        Number using O6 output only:         3,365
    Number used as exclusive route-thru:     1,577
  Number of route-thrus:                     2,687
    Number using O6 output only:             1,576
    Number using O5 output only:               805
```

```
    Number using O5 and O6:                  306
Slice Logic Distribution:
  Number of occupied Slices:              12,116 out of  14,720   82%
  Number of LUT Flip Flop pairs used:     41,700
    Number with an unused Flip Flop:       5,829 out of  41,700   13%
    Number with an unused LUT:            18,355 out of  41,700   44%
    Number of fully used LUT-FF pairs:    17,516 out of  41,700   42%
    Number of unique control sets:           453
    Number of slice register sites lost
      to control set restrictions:           862 out of  58,880    1%
Specific Feature Utilization:
  Number of BlockRAM/FIFO:                    53 out of     244   21%
    Number using BlockRAM only:               53
    Total primitives used:
      Number of 36k BlockRAM used:            15
      Number of 18k BlockRAM used:            65
    Total Memory used (KB):                1,710 out of   8,784   19%
  Number of BUFG/BUFGCTRLs:                   10 out of      32   31%
    Number used as BUFGs:                     10
  Number of BUFDSs:                            2 out of       8   25%
  Number of CRC64s:                            2 out of      16   12%
  Number of DCM_ADVs:                          3 out of      12   25%
  Number of DSP48Es:                         370 out of     640   57%
  Number of GTP_DUALs:                         8 out of       8  100%
  Number of PLL_ADVs:                          2 out of       6   33%
```

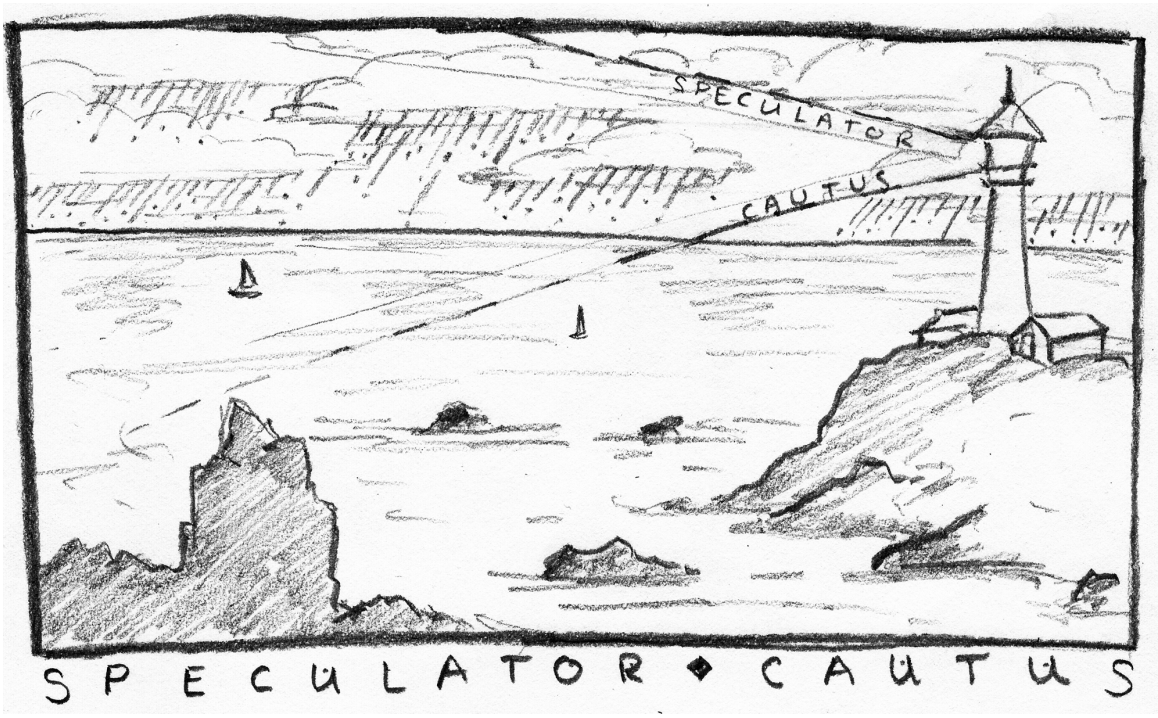
SPECULATOR
CAUTUS
SPECULATOR ◆ CAUTUS

Chapter 6

HIGH AGGREGATE BANDWIDTH CHALLENGES

Digitizing a large portion of the spectrum increases the frequency agility of the radar and increases the number and diversity of usable transmitters. For example, there can be redundancy in target location estimation by processing many DTV transmitters simultaneously. With a 6.4 Gbps data rate on the 10 GbE link and 640 DSP48E slices, we can digitize and store 60 MHz of spectrum on each of 4 antennas with 8-bit IQ samples. That is up to 10 different DTV stations on each antenna. This gives us flexibility to operate in a geographic location with a sparse DTV spectrum. As engineers, we want to push the bandwidth limitations. A radar that can save the raw ADC samples at 3 GHz would be ideal; we could processes every transmitter under 1.5 GHz illuminating the target. At this time, however, the FPGA is needed to reduce the ADC data rate to a manageable rate for real time processing and data storage. Creating FPGA designs to accommodate fast data transfer off the 10 GbE port is only half the problem; the receiving computer must be able to handle that fast rate data. There are challenges associated with high aggregate bandwidth, and this chapter is designed to illustrate and analyze these issues.

6.1 Introduction to current data rates and data transfer

6.1.1 Data Rates and disk storage

Currently, the radar operates in two modes:

- standard 0.91 Gbps (114 MBps), which can accommodate the reception of a single DTV station on each of 4 antennas with 8 bit samples

- fast rate 6.4 Gbps (800 MBps), which can accommodate up to 10 frequency-adjacent DTV stations on each of 4 antennas with 8-bit samples.

At these data rates the recording media will fill quickly. Without careful attention to disk space and moving data to external storage in between experiments, data will not be able to be saved to disk. This would be a frustrating hitch in a carefully coordinated field experiment, and we take care to ensure this does not happen. In our experimental setting it is possible for graduate students to monitor and move data with to the available storage on Thorex. However, other arrangements need to be made before the system could operate continuously in support of geoscience experiments. During magnetic storms, it is common for MRR to operate continuously for 12 hours. Dramatic improvements in the cost and capacity of hard disk drives make it possible to store more data. However the data must still be processed eventually. If we were able to process the data in real time, the total data volume reduction would be substantial, as cross ambiguity arrays are much smaller than the data used to compute them. Table 6.1 shows that storing the IQ data takes 228 times the volume as simply storing the cross-ambiguity. The ability to perform real-time processing, even with latency, would enable us to support long duration experiments.

Data Type	Format	Storage (MB)
14.57 MHz IQ data	1s 8-bit streams, 2 antennas	456 MB
Cross-Ambiuity	256 x 1000 matrix, double precision floating point	2 MB
	Storage savings:	$\approx$ 228x

Table 6.1: Disk space required to store IQ data compared to simply storing the cross-ambiguity result. The savings is a factor of 228.

For the challenge of aircraft detection in particular, even more computational burden can be placed on the GPU. After the cross-ambiguity calculations, automatic detection, data association, feature-recognition, feature recognition and 3D tracking are still performed offline with this system. The near future holds the potential for real-time operation of all these tasks. In this case, the final data product could be as small as a handful of numbers such as number of aircraft, features of each detection such as range, angle, 3D location, 3D velocity, and perhaps classification, such as labeling the detection as a helicopter or float plane.

Besides data storage, another large caveat to wideband system design is reliable transfer of fast rate data. It is fair to assume that in order for the radar to operate, the fundamental task of reliably receiving the fast rate data over the 10 GbE link is vital to the operation of the system. While for many systems this may be true, passive radar is incredibly robust to data loss. In that sense it is a great fit for bleeding edge fast data transfer technologies. In the following sections, we will characterize our particular packet loss scenario, and identify solutions, and demonstrate just how robust our system is to packet loss.

6.1.2 Data Stream Structure

The ROACH board utilizes User Datagram Protocol (UDP) to packetize and transmit the data over a point-to-point, 10 GbE connection to the storage computer. The reader may be familiar with the common protocol for internet access, Transmission Control Protocol (TCP), that guarantees data packets arrive reliably, and in order. This is done by two-way communication between the sender and receiver, acknowledging packets sent. The cost of this reliability is extra bytes spent on more header information and lower throughput. UDP is in general unreliable, with no guarantee packets will be received at all, and even if received there is no notion of order. It is a poor choice for a general internet protocol, where data packets may travel through many nodes, and based on congestion may be redirected and arrive out of order. For a point-to-point connection, such as the dedicated 10 GbE link between the ROACH board and the receiving computer, these sorts of errors are less likely, and UDP is a good choice given less header overhead.

The 10 GbE core on the ROACH has an internal buffer that can store up to 8 kB of data. The network card is not clocked off the FPGA, but by a fixed 156.25 MHz oscillator. This makes it inherently asynchronous to the FPGA clock, and in order to avoid buffer overflow even with much lower data rates than the 10 Gbps limit of the link, we maintain a payload size of 4 kB, or 512 64-bit words. Each 64-bit word is made up of an 8-bit I and Q sample from each of 4 antennas.

6.1.3 FPGA interaction with 10 GbE interface

The FPGA interface with the 10 GbE link can be programmed from the Simulink environment using the Casper tools. The internal buffer stores data in 64-bit words with a maximum storage capacity of 8 kB. For a standard mode of operation, the 64-bit data word is packed with 8-bit I and Q samples from each of four antennas, as shown in Figure 6.1.

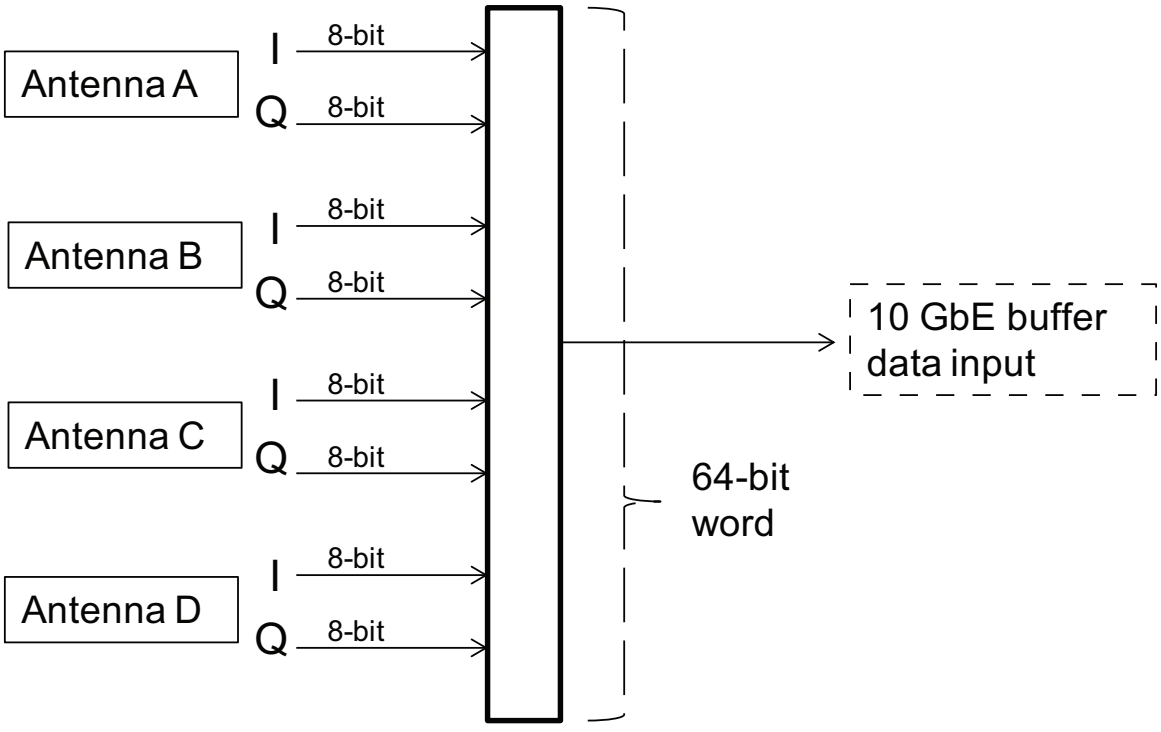

Figure 6.1: After channel selection processing on the FPGA, the IQ data streams from each of four antennas are packed together in a 64-bit word to enter the 10 GbE buffer.

While the 64-bit word stream is connected directly to the 10 GbE buffer, words do not enter the buffer every clock cycle. Depending on the design, clocking signals are devised to reduce the data rate down to manageable data rates for the 10 GbE. At this point, the data has already gone through an anti-aliasing filter. Then we can select one of every few samples to reduce the data rate for loading the buffer. The 10 GbE buffer needs a signal to indicate on when to pass data to the buffer. When this signal is high for a single FPGA clock cycle, a single 64-bit word is passed into the buffer. When the signal is held low, the 64-bit word associated with that FPGA clock period is discarded. With an FPGA clock rate of 200 MHz, we create a clock signal that holds high for one FPGA clock period, and low for 13 FPGA clock periods. This passes every 14th data sample into the buffer, in other words we downsample the data by a factor of 14, reducing the data rate from 200 MHz to roughly 14.2857 MHz (200/14 MHz). The 10 GbE buffer also requires a signal to indicate the final sample to be packed in a UDP frame. To maintain a UDP data payload of 4096 bytes, 512 64-bit samples are sent in a frame. One of every 7168 FPGA clock periods (512*14) are held high, and this is coordinated in time with the final data sample sent into the buffer for the frame. Figure 6.2 shows how the 64-bit word data stream and these clock signals connect to the 10 GbE buffer.

Shown in Figure 6.3 is the actual Simulink programming interface with the 10 GbE block, labeled `gbe0`. Note `tx_data` is the data input for the 64-bit words, `tx_valid` is the boolean clock signal choosing samples to pass into the buffer, and `tx_end_of_frame` or abbreviated `eof` is the clock signal indicating the final sample in the UDP packet. In addition to transmitting data off the FPGA, the core can also be configured to receive data, perhaps from another ROACH board. This functionality is not currently used in our system. Simulink can show the data types of all signals involved, a helpful debugging tool. Simulation is very helpful to ensure the clock signals are generated properly and coordinated. Figure 6.4 is a Simulink scope tapped into the `tx_valid` and `eof` wires, varifying the programmer designed these signals correctly.

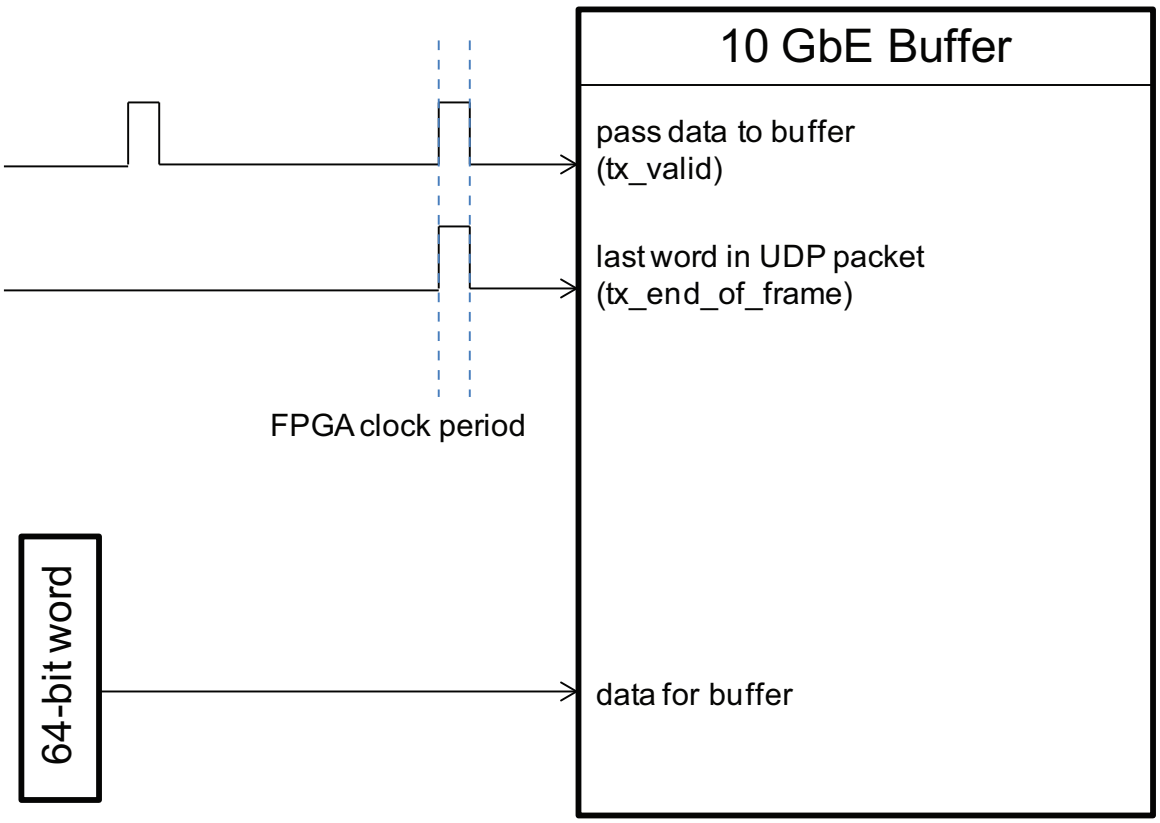

Figure 6.2: A boolean clock signal, `tx_valid`, determines which 64-bit words (packed with IQ data from 4 antennas) enters the 10 GbE buffer for transmission. Another clock signal, `tx_end_of_frame`, signifies the last word in the UDP frame.

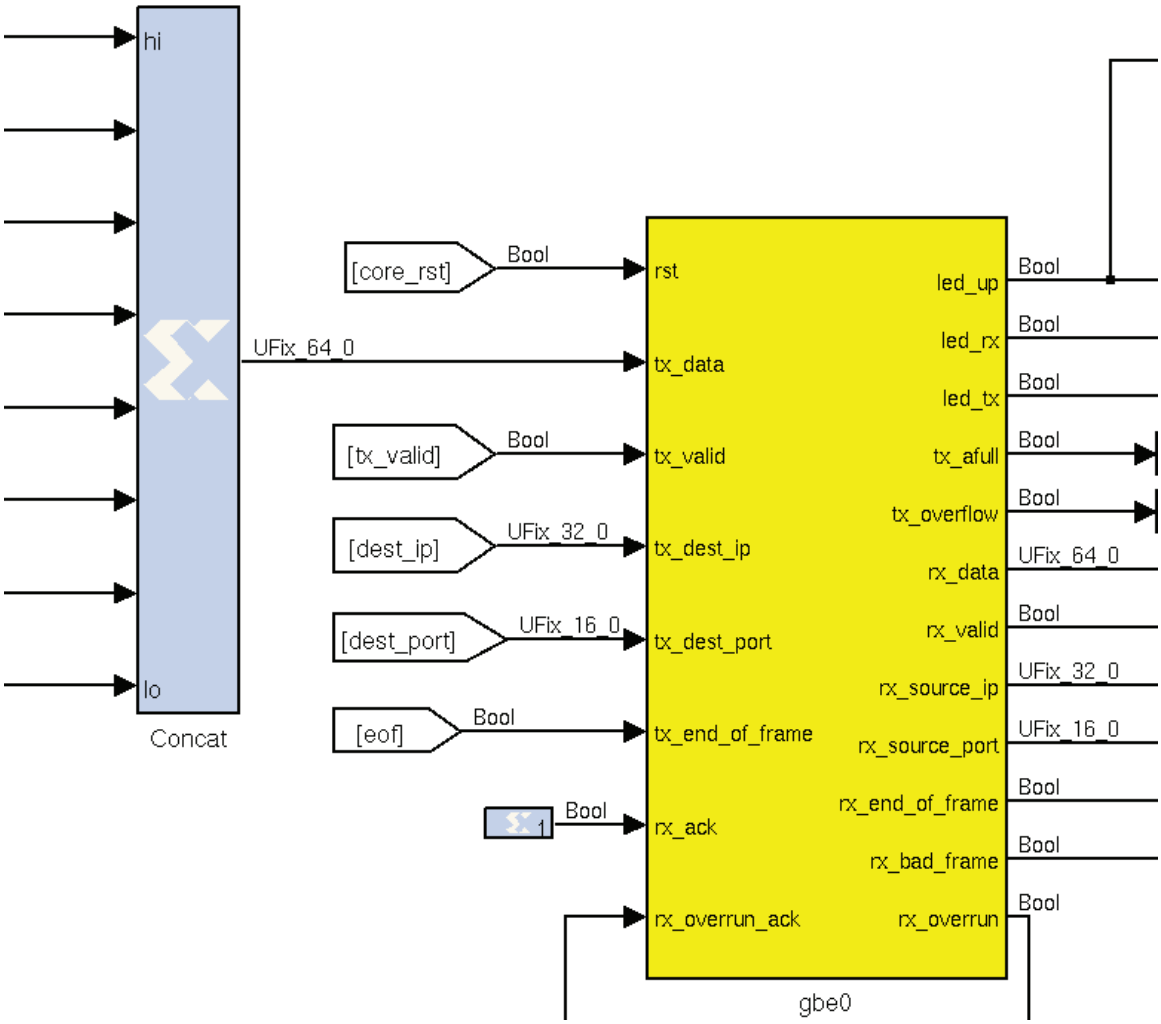

Figure 6.3: A snapshot of the Simulink environment showing the 10 GbE block in the various input signals, most notably the data input `tx_data`, and the clock signals `tx_valid` and `tx_end_of_frame` (`eof`) as noted in Figure 6.2.

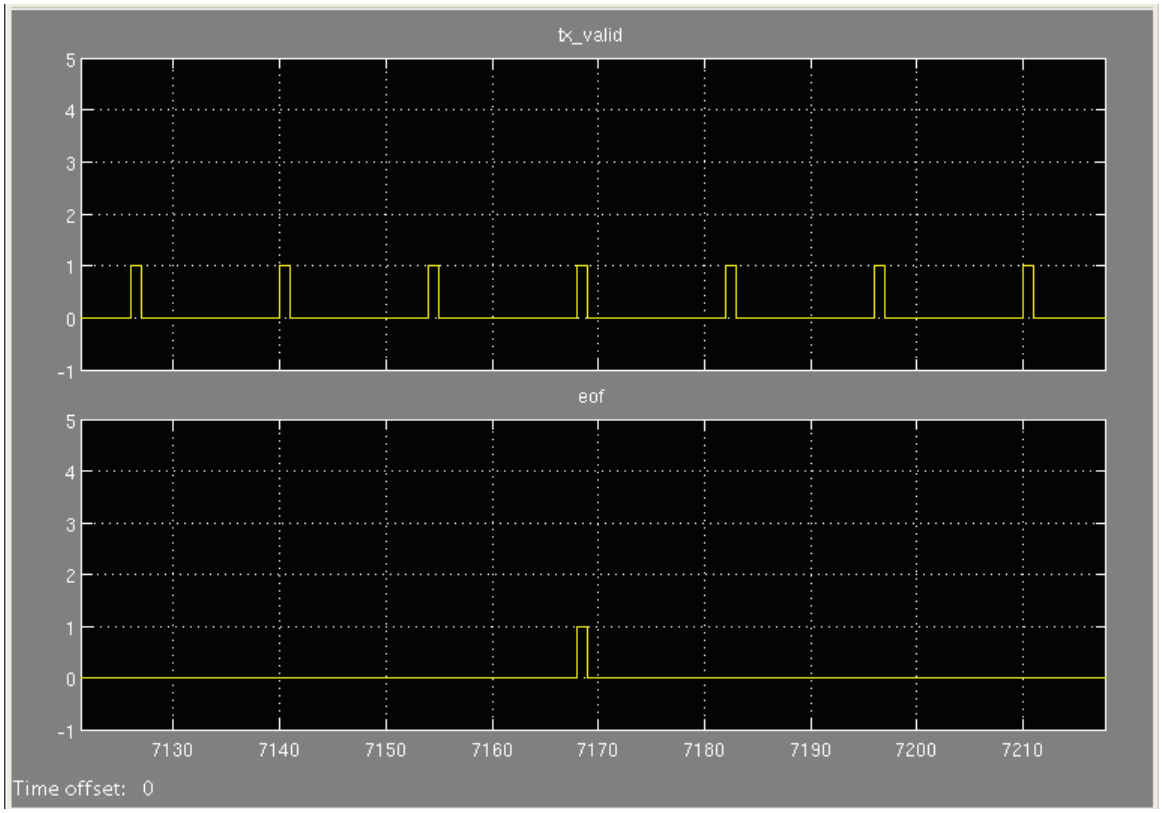

Figure 6.4: A Simulink scope plotting the simulated clock signals driving the 10 GbE buffer. The scope taps into the wires for `tx_valid` and `tx_end_of_frame` (`eof`) as noted in Figure 6.3.

6.2 Characterization Of Data Loss

Data loss symptoms first presented themselves with the high-data rate mode. When saving data, a particular volume in bytes in requested, and data is received for as long as it takes to fill that volume. We noticed that saving 10 seconds worth of data really took over 30 seconds of time, indicating data loss. Also of important note is that by changing parameters in the code used to save the data from the 10 GbE link to the storage computer, the time "dilation" that first alerted us to the data loss problem could be significantly altered. This points to the issue lying not in the ROACH board but with the receiving computer. It also tells us that a majority of the data loss was not occuring at the network card level.

To characterize this data loss, the FPGA was reprogrammed to insert a 64-bit header into the front of each packet. In this case, the header is simply an unsigned 64 bit integer incrementing by one every time a UDP packet is sent. By stripping the received headers from the received data, the loss of data can be categorized. The data is saved as a long stream, concatenating all UDP packets for the desired amount of data saved. With a UDP packet size of 512 64-bit words (1 header, 511 data words, for a total of 4 kB), every 512th 64-bit word from the resulting stream can be read, and forms the expected header counter, albeit with missing integers where data is dropped. This shows us that data loss occurs in full-packet increments.

By adjusting parameters in the data-writing code, the standard mode of operation drops no data for the first six seconds of operation. During the seventh second, there is a single, large burst of data loss of over 1000 packets. In Figure 6.5, we plot the amount of data lost between consecutive received packets for three 10-second data sets. The three data sets were taken at different times during the same day, and for completion we include the set labeled "Data 2" (red) although no packets were dropped, illustrating that data loss may not even occur during normal operation. The first large data loss for Data 1 (blue) and Data 3 (green) occurs around 200,000 packets, or 800 MB of data. Six seconds of clear data is plenty for reliable target detection. By knowing where data is lost, we can simply choose to omit that portion of data in our analysis. In fact, we may be able to achieve reliable operation by simply imposing a duty-cycle such as five seconds on, two seconds off (a rough

example). At any rate, the overall data loss in 10 seconds is kept to under 2%. The distance between large bursts of data varies.

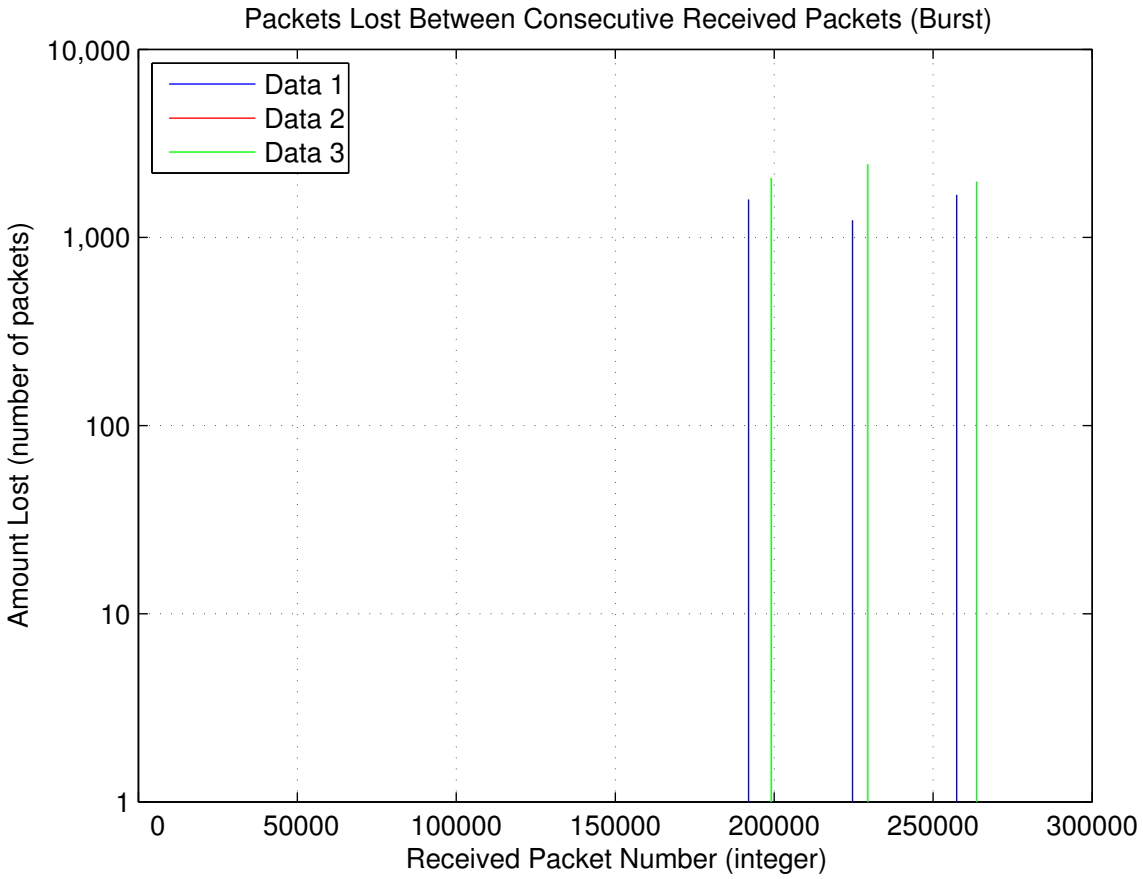

Figure 6.5: Bursts of data loss between consecutive packets of standard rate (≈14 MHz) data. Three 10-second data sets are presented, with the set labeled "Data 2" (red) showing the valuable piece of information that occasionally no data loss will occur. "Data 1" (blue) and "Data 3" show a large burst of data loss after 6 seconds of continuous operation, allowing roughly 200,000 packets or 800 MB of data to be reliably saved.

The fast rate data loss is more interesting to study. For three different fast rate data sets, we plot the received header and the expected header in Figure 6.6. The x-axis indicates expected header value. Meaning if we recorded 500,000 packets of data (2.5 seconds) we expect to see consecutive headers from 1 to 500,000. This expected header value is noted by the black line, with a linear relationship to the y-axis of actual received header value. As

packet loss occurs, we expect to read header values greater than this linear line, indicating that the actual value of the header between consecutive packets is more than the expected counter increment of 1. Noted here is the percentage of data lost, relative to the first received data packet in each stream, at several instances in time denoted by dashed vertical lines. Considering the first and last packet received, roughly 35% of the data was lost, implying that when we intended to save 2.6 seconds of data, we received the same amount of data in bytes but due to losses contains samples from over 3.8 seconds of real time.

All three data sets seem to be roughly piecewise linear, meaning minor data losses with intermittent bursts of data loss contributing largely to the overall total. To explore the bursty nature of the data loss, we look at a plot of the amount of data lost between consecutive received packets in Figure 6.7. Similar to the standard rate data, all three data sets experience their first large data loss around 200,000 packets in (≈800 MB). If we need to write our own software to save the data off the 10 GbE link, this suspiciously consistent location of large data drop would lead me to investigate a buffer or memory overflow. While the distance between large bursts varies, it seems regular enough to suggest a deterministic component such as a buffer overrun. A similar duty-cycle to the standard rate operating suggestion could be imposed here, however the time scale would lead to running the radar for under one second, then off for a fraction. The green and red data sets show many smaller amounts of packet loss over short stretches of time. We can explore data set 1 with a histogram showing the number of times loss at that level occurs in Figure 6.8. Recall that the data lost is only 35% of the total data run. It is still more common NOT to lose data than to lose it. The histogram shows it is stil more likely to lose data in small bursts of less than 10 packets than large bursts off 100, 1,000, or 10,000 or more packets.

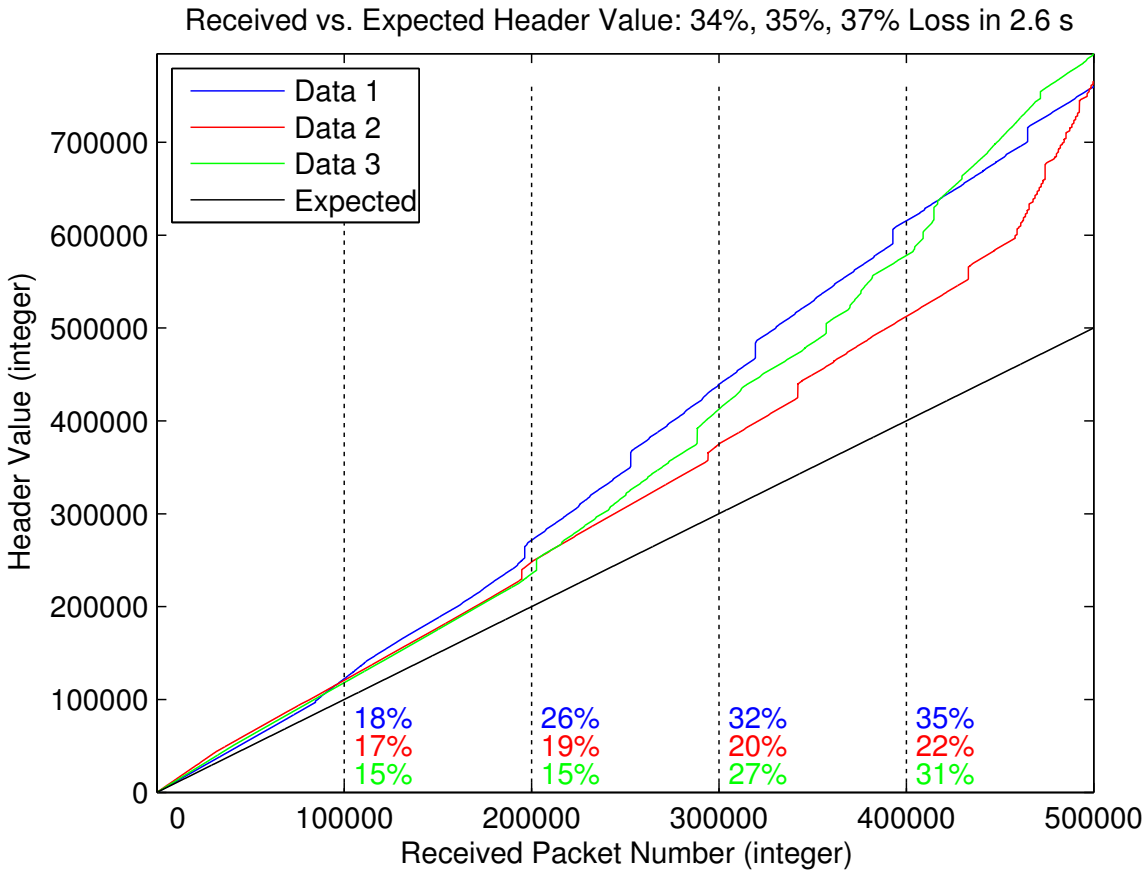

Figure 6.6: Packet Loss determined by value of the 64-bit header received in the storage computer. UDP packet headers increment by one each sent packet, with the black line indicating the expected received header number. The blue, red, and green lines all represent real data taken on the same day, showing the actual header values received. The deviation is recorded as the percentage of data lost, and is traced for all three data sets at several instances in time, marked by vertical dashed lines. These numbers are presented as integers relative to the first header value received, for comparison. After 2.6 seconds, all three data sets show aggregate data loss around 35%.

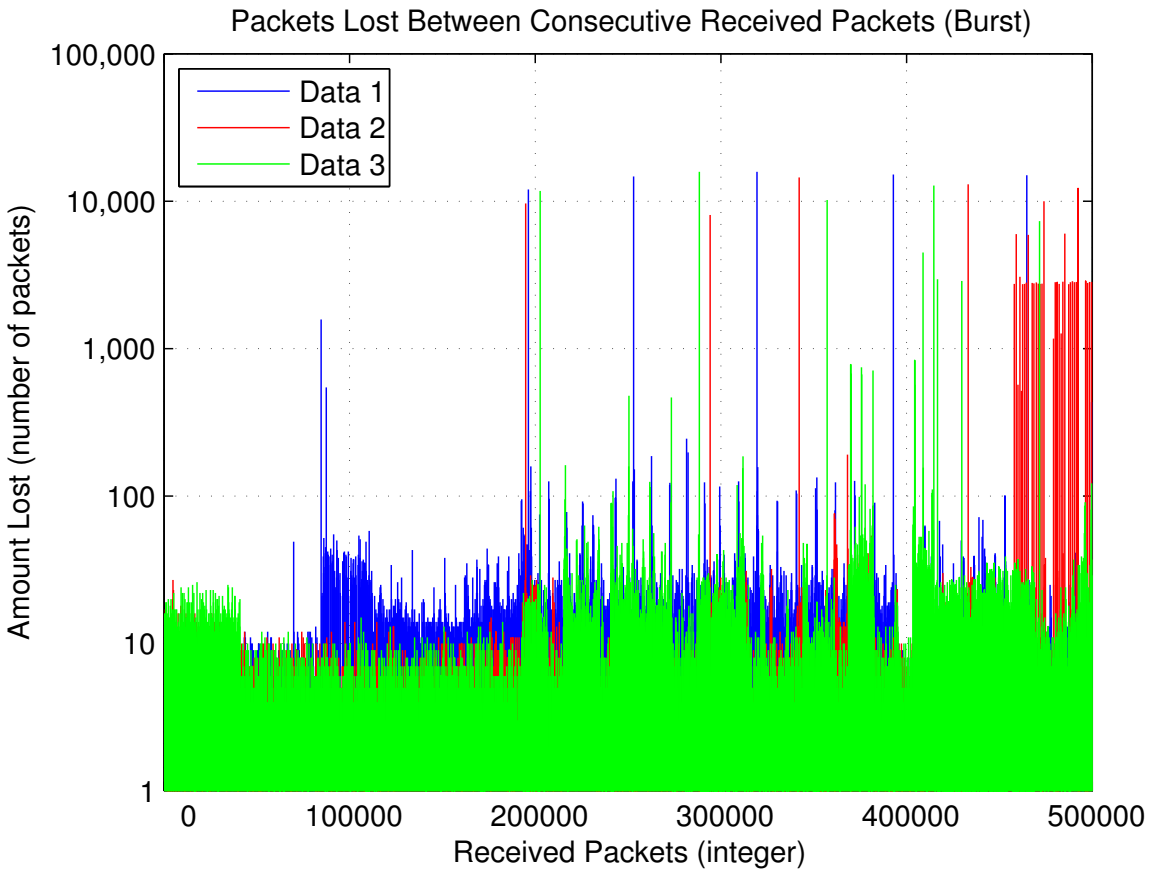

Figure 6.7: Bursts of data loss between consecutive packets of fast rate (100 MHz) data. Three 2.6-second data sets are presented, showing the first large burst of data loss around 200,000 packets in.

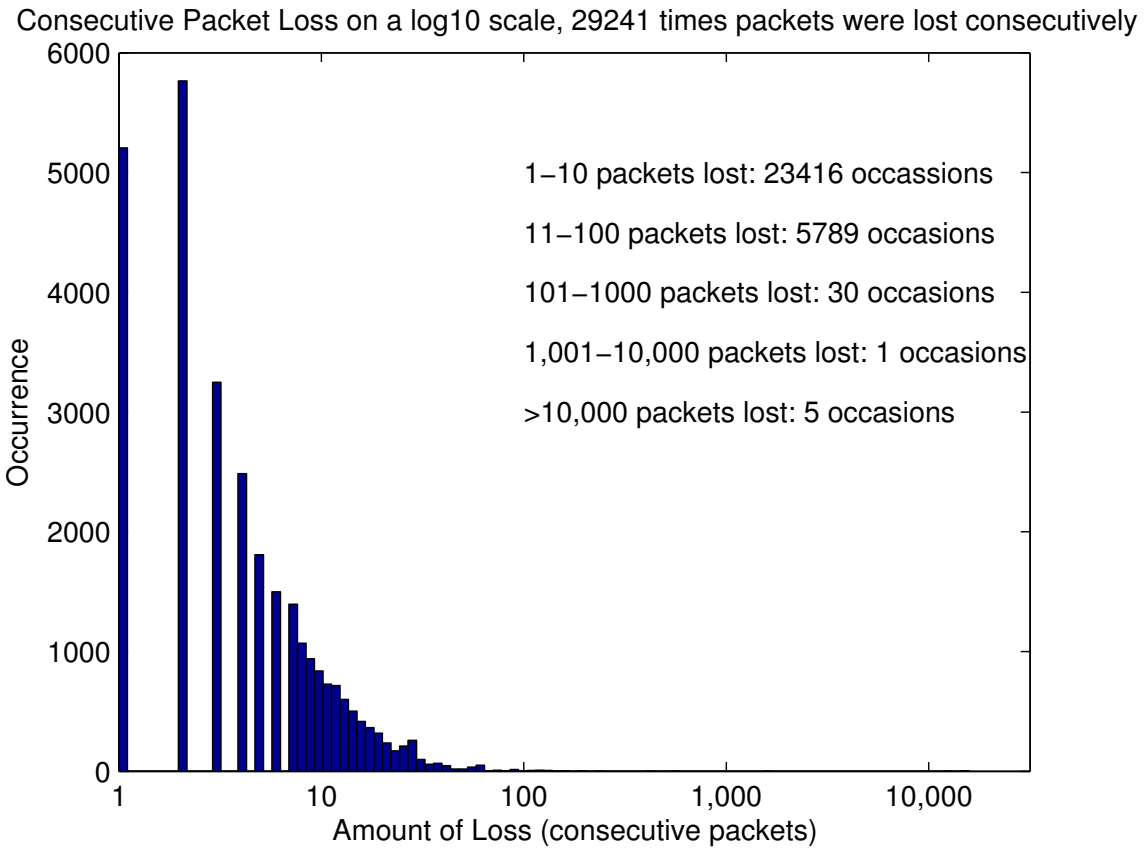

Figure 6.8: Histogram of amount of consecutive data loss on a log scale, x-axis, with the number of occurances on the y-axis. Still most common is small, 1 or 2 packet loss.

6.3 Simulation and Analysis of Data Loss

To understand exactly how data loss affects the target detectability, range, and Doppler estimtes, we can simulate a DTV signal and reflection from a target and observe the resulting transformation.

6.3.1 Signal Modeling

We model the signal received on antennas A and B in the following way:

$$x_A[n] = x[n] + \alpha e^{j2\pi f_0 n} x[n - n_0] \tag{6.1}$$

$$x_B[n] = e^{j\theta} x[n] + \alpha e^{j\phi} e^{j2\pi f_0 n} x[n - n_0] \tag{6.2}$$

where $x[n]$ is the direct path transmission, arriving on both antennas at time zero, but with a relative phase difference of θ. The second component in each antenna signal is the reflection from a target. The direct path is scaled by α and Doppler shifted by f_0. This indirect path arrives at both antennas at time n_0 with relative phase difference of ϕ.

The parameters are chosen for a target at a 3 km range, 102.4 m/s Doppler velocity, and a SCR of -10 dB. This target will be informative to our understanding of packet loss effects, but is unrealistically large in a physical sense.

We expect the cross-correlation to yield a peak at 0, 0 with a second peak at τ_0, f_0 and with power roughly 10 dB below the 0,0 peak. A more complete discussion of the mathematics is found in Appendix C, although we may expect more than 10 dB difference here because the assumption that $|\alpha|^2 \ll 1$ may not hold for such a strong target.

6.3.2 Packet Loss Model

Packet loss occurs on the 10 GbE link between the ROACH board and the storage computer. This data rate is in general larger than the final IQ rate, since subsequent processing is performed after this data transfer. For this simulation, we drop packets by Bernoulli trial, selecting a packet loss probability of 1%, 5%, 10% and even 50%. This simulation loses data at the 14 MHz rate, close to our standard 900 Mbps on the 10 GbE link. The data is then

processed two ways: as if the packet loss is undetected and subsequently we process data that has been artificially sped up in time, and again replacing these unknown packets with packets containing data with value zero. This process is presented in Figure 6.9. For each regime of packet loss percentage, 10 independent trials will be executed. We will examine the affect of this packet loss by features of the cross-ambiguity plot including range, doppler, and power estimates of the target and target return. To model the direct path transmission, band-limited white noise is used.

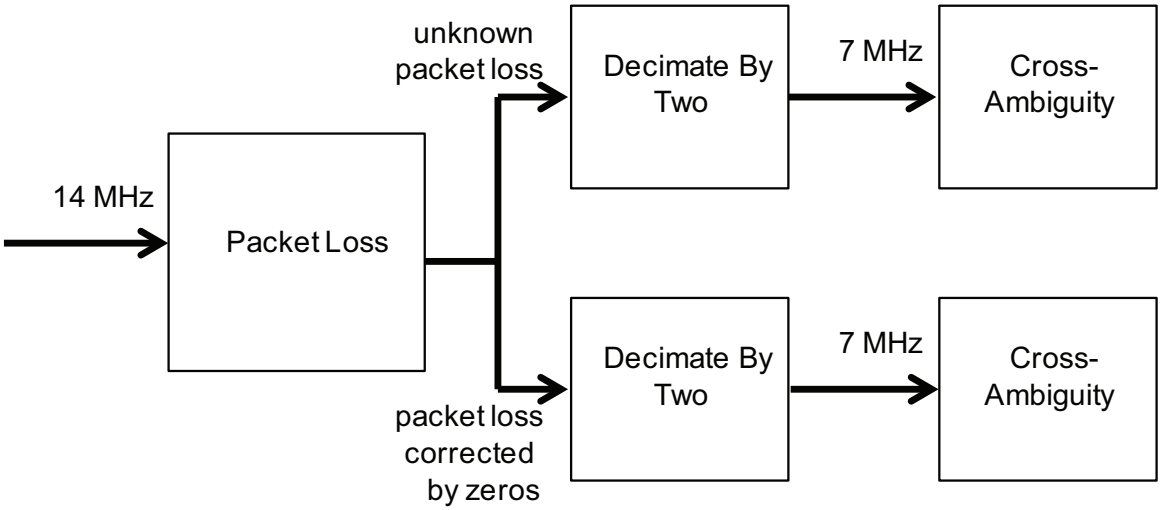

Figure 6.9: Diagram of packet loss experiment. Data is lost in whole packets at the fast, 14 MHz rate. The data loss produces two streams of data: one where the packet loss is known and represented as packets of zeros, and one where the lost packets are removed entirely creating a scenario much like processing data with unknown losses.

6.3.3 Control Experiment: 0% Packet Loss Cross-Ambiguity

We create a fairly powerful scattered signal, so that the SCR is roughly -10 dB. This is unrealistic for an aircraft, but given the processing we wanted to see both the peak of the cross-ambiguity at 0,0 as well as the peak at the target location. The initial doppler shift of the target is not an exact doppler sample, thus we see the target power spread vertically, showing the spreading of power associated with sampling a periodized sinc function, as expected. If the perceived doppler shift of the target changes to align perfectly with a

doppler bin, this vertical spreading will diminish. Also noted here is that with no ground clutter or noise injected into the simulation, we see spreading along the range axis from both the direct path and from the target return. The rate reduction from the 10 GbE link speed to the final IQ data used in the cross-ambiguity requires decimation (anti-aliasing filtering followed by rate reduction), and here we see the sidelobes of the autocorrelation of our filters. These sidelobes, however, are 30 dB down. These affects are shown in Figure 6.10 for a target at 3 km path difference and 100 m/s doppler shift. The peak of the cross-ambiguity at the target (or at the closest doppler sample to the true doppler shift) is roughly 13 dB below the peak at 0,0, showing that at 0,0 there is some contribution from the artificially large target scattering cross section.

We can observe the affects of our direct path adaptive nulling on simulated data in Figure 6.11. As expected, the entire 2-dimensional cross-ambiguity contribution from the direct path is removed. The median value of the plot, which we use to estimate the average value of the floor has changed from 60 dB in Figure 6.10 to 46 dB in this Figure 6.11. While the radar is not calibrated, thus we present unitless dB values, the units in both plots are identical and can be directly compared. The power return at 0,0 and any range spreading due to filtering is gone. This includes removing power from the "floor" as well as removing all of the range-spread features associated with the direct path. What we are left with is simply the return at every range and doppler bin entirely associated with the target. The difference between the peak value at the target and the floor of this cross-ambiguity is 62 dB, the expected peak-to-floor ratio for roughly a quarter second of coherent integration of 7 MHz data. With the floor pressed downward, the range and doppler spreading features from the processing are more visible. These features are summarized in Table 6.2. Now that the features of the control experiment are well understood, we can explore the affects of packet loss.

Peak Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
122.7	111.3	-11.4	3.0	103.1

Table 6.2: Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target.

6.3.4 1% Packet Loss

We can plot the bernoulli trial for 1% packet loss for the first five hundred packets in Figure 6.12. Packets 13, 21, 176, and 257 will be lost. With each packet of data containing 512 samples of I and Q data from four antennas, we can zoom into the data streams and compare the data with unknown packet loss and with packets loss corrections in Figure 6.13. The data in both plots are identical up until roughly 6,000 samples at which time the first packet is lost. The top plot shows the signal with unknown data loss (512 samples missing), the bottom shows the lost packets replaced with zeros (512 zeros). On the bottom figure, there is a particularly large sample just under 7,000 samples into the data set. It is clear that this data point shows up 512 samples earlier in the top data set, showing that the unknown data set is artificially "sped up" in time by each of these lost packets.

Unknown Signal Loss

Using the exact same simulated data as in the control case, we can run 10 simulations of bernoulli trials where the probability any given packet is lost is 1%. First we process the stream of data representing unknown data loss. Statistics from these trials are shown in Table 6.3. Any change in range and Doppler shift in the peak of the target remains within a single bin. The direct path power does not change from trial to trial, which is expected. The same packets are lost on every antenna, so the peak at 0,0 will always be consistent. The power associated with the target, while maintaining the same average as the control case, shows variability on the order of 0.5 dB.

To consider details of the packet loss not as easily quantified, we can examine cross-ambiguity shown for one of these trials in Figure 6.14. With the scope of the cross-ambiguity zoomed on the target, we confirm that there is not a noticable difference between this example and the control case. The one difference is that this particular plot shows that power increases in the Doppler bin directly above the main target peak, showing that perhaps the measured Doppler shift is changing ever so slightly.

	Direct Path Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.7	110.7	-12.0	3.0	103.1
Std Dev:	0	0.5	0.5	<0.05 (bin)	<3.4 (bin)

Table 6.3: Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 1% unknown data loss.

Corrected Signal Loss

The same 10 bernoulli sample paths are used again, but instead of simply dropping the lost data, we assume the system can detect where packets are lost and can correct for it by inserting 512 zeros for each lost packet. Statistics for the results are summarized in 6.4. The values correspond exactly with the control experiment, and the variance of the target power is zero. We plot one of the cross-ambiguities in Figure 6.15, but there are no notable differences between this and the control case.

	Peak Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.7	111.2	-11.5	3.0	103.1
Std Dev:	0	0	0	<0.05 (bin)	<3.4 (bin)

Table 6.4: Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 1% data loss corrected with packets of zeros.

6.3.5 5% Packet Loss

Unknown Signal Loss

The same experiment is repeated with 5% data loss. Statistics are presented in Table 6.5. There is a similar variance in the target power as we saw in the 1% unknown data loss experiments. However, the measured target Doppler velocity has shifted positive by an entire Doppler bin, ≈ 4.5 m/s. The cross-ambiguity of one trial is presented in Figure 6.16. This, again, shows little affect of the data loss. In fact, without measuring it is hard to tell

that the Doppler value of the peak has shifted exactly one bin. If the system is processing data where losses occur and go undetected, bias in the Doppler estimate is guaranteed.

	Direct Path Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.7	110.9	-11.8	3.0	107.8
Std Dev:	0	0.9	0.9	<0.05 (bin)	<3.4 (bin)

Table 6.5: Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 5% unknown data loss. The measured Doppler shift is precise, but biased by roughly 5 m/s in the positive direction.

Corrected Signal Loss

Statistics for the results of the 5% data loss corrected by zeros are summarized in 6.6. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth. These values show that knowing when data is lost and simply replacing it with zeros ensures that the correct Doppler shift can be recovered.

Figure 6.17 shows one of the cross-ambiguities. The plot confirms the measurements for range and Doppler are correct, however interesting Doppler-spread power is present along the vertical zero range slice. This feature is in all likelihood present near the target detection as well, but is obfuscated the Doppler features already present before the loss of data. The data stream being processed is not simply the truth data, but is the truth data multiplied by a function with rectangular windows. Multiplying by rectangular windows in the time domain causes Doppler spreading, and in the future we may want to explore other windowing functions to reduce these sidelobes. This problem is common to spectral estimation, and texts such as [41] cover methods to taper the data in a way that will result in a more accurate spectral density estimate. While the affect does not alter the detectability of this particular target, it may cause problems for weaker targets in the same range as large ground clutter.

	Peak Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.3	110.7	-11.6	3.0	103.1
Std Dev:	0	0	0	<0.05 (bin)	<3.4 (bin)

Table 6.6: Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 5% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth.

6.3.6 10% Packet Loss

Unknown Signal Loss

The same experiment is repeated with 10% data loss. Statistics are presented in Table 6.7. Even more clear here is that the measured Doppler shift of the target has increased by even more. Surprisingly, even losing this much data does not reduce the measured target power significantly, on the order of 1 dB. Figure 6.18 shows a cross-ambiguity. Without knowing about the data loss, there is nothing particularly suspicious about this plot. Perhaps noting that the peak of the target does not seem to be centered in the doppler domain may clue the engineer into a non-physical behavior of the target. This may be indicative of the irregular way the signal is being sped up in time.

	Direct Path Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.7	110.1	-12.6	3.0	114.9
Std Dev:	0	0.9	0.9	<0.05 (bin)	<3.4 (bin)

Table 6.7: Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 10% unknown data loss. The measured Doppler shift is consistent across the 10 trials, but is nonetheless wrong by close to 10%.

Corrected Signal Loss

Statistics for the results of the 10% data loss corrected by zeros are summarized in 6.8. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth, perhaps unsurprising at this point even though the data loss is not small.

Figure 6.19 shows the recovery of the proper range and Doppler shift of the target, and again we see the Doppler-spreading due to the stretches of zero in the data stream.

	Peak Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	121.8	110.1	-11.7	3.0	103.1
Std Dev:	0	0	0	<0.05 (bin)	<3.4 (bin)

Table 6.8: Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 10% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce identical quantities, however the overall power has dropped in both the direct path and the target measurements, due to incorporating zero for 10% of the data. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth.

6.3.7 50% Packet Loss

Unknown Signal Loss

Finally we consider the extreme case of 50% packet loss. This is larger than what we have seen in the data. Statistics are presented in Table 6.9. As anticipated from past experiments, the Doppler shift is measured at double the true value. The measured Doppler shift is less consistent across the 10 trials as previous data loss experiments shown in Tables 6.3, 6.5, and 6.7. The power in the main target peak has dropped by more than anticipated, reflected in the SCR increased from ≈12 to 20 dB. This means targets can lose 8 dB in detectability, which may be enough to bury weaker reflections in the cross-ambiguity floor. The effect of this extreme packet loss is clear in Figure 6.20, where the cross-ambiguity shows the target having many peaks in Doppler shift, rather than one clear peak in the control simulation.

The packet loss caused smearing in the vertical direction. We suspect this feature increases with intensity, as the peak of the direct path at 0,0 remains focused in Doppler space. Again if we do not realize that loss has occured, we may mistakenly think this aircraft has a spinning feature, much like a helicopter, causing significant reflection at the same range but inducing different Doppler shifts on the reflected radio waves.

	Direct Path Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	122.7	102.4	-20.3	3.0	205.6
Std Dev:	0.1	1.3	1.3	<0.05 (bin)	3.3

Table 6.9: Summary of features of cross-ambiguity based on simulated DTV data and scatter from aircraft-like target with 50% unknown data loss. As anticipated from past experiments, the Doppler shift is measured at double the true value. The measured Doppler shift is less consistent across the 10 trials as previous data loss experiments shown in Tables 6.3, 6.5, and 6.7. The power in the main target peak has dropped by more than anticipated, reflected in the SCR increase from ≈12 to 20 dB.

Corrected Signal Loss

Statistics for the results of the 50% data loss corrected by zeros are summarized in 6.10. Even losing half of the data, the Doppler shift and Range can be measured accurately with this correction method. Power has been lost due to half of the data processed being equal to zero. We can explore the importance of this by examining Figure 6.20. The dynamic range of the system without any nulling can be defined as the power of the peak at 0,0 to the power of the floor of the cross-ambiguity. In general, for 0.22 seconds of integration at 7 MHz, this should be 62 dB. This tells us that targets with SCR less than -62 dB will not be detected as the peak of the correlation at their particular range and Doppler shift will not exceed the floor. Replacing half the data with the value zero says we should expect to achieve half as much integration, a hit of only 3 dB. The average floor value in Figure 6.20 is 57, showing a dynamic range of 59 dB, exactly 3 dB less than the control simulation. The value of the peak of the target above the floor here is 47 dB. In the control experiment, the value of the peak of the target above the floor was roughly 51 dB. We can estimate how

detectable a target is by this value. This simulation shows that by losing 50 % of the data and adding a simple correction of zeros for lost packets, the target range and Doppler shift can be recovered while losing only 3-4 dB of detectability. While the simulation is not a perfect match for reality, this is an incredibly promising result. For all the advantages that fast rate data transfer affords the passive radar system, the disadvantage of data loss leads to a small performance degradation.

	Peak Power (dB)	Target Power (dB)	SCR (dB)	Range (km)	Doppler (m/s)
Average:	116.8	104.0	-12.7	3.0	103.1
Std Dev:	0.1	0.2	0.1	<0.05 (bin)	<3.4 (bin)

Table 6.10: Summary of features of cross-ambiguity calculation based on simulated DTV data and scatter from aircraft-like target with 50% data loss corrected with packets of zeros. The values indicate that all 10 simulations produce close quantities, however the overall power has dropped in both the direct path and the target measurements, due to half the data set having value zero. Otherwise, the signal to clutter ratio and the range and Doppler measurements agree with the truth.

6.3.8 Conclusions from Simulation

The simulation is by no means complete. One feature that this simulation lacked to show is range migration. A physical target with Doppler shift is changing range, by definition. If integration by 0.22 seconds guarantees the target does not move range bins in that time, what if unknown data loss causes an integration of 0.44 seconds, such as the 50% data loss case? We expect not only smearing in Doppler, but perhaps smearing in range as well. I suspect this range smearing will occur in the direction toward the radar, as partial returns from the target are processed sooner in time than they should be. For real data, we expect that losses will cause the target to smear diagonally both up and left in the cross-ambiguity plot.

Additionally, while bernoulli trial is fine for simulation, we actually know our data loss is quite bursty, shown previously in Figures 6.5 and 6.7. If we can know when data loss has occured, for a particular set of data we have collected we can run this simulation with

the detected packet loss to observe the effects on a point target and perhaps invert a more accurate estimate of range and Doppler shift.

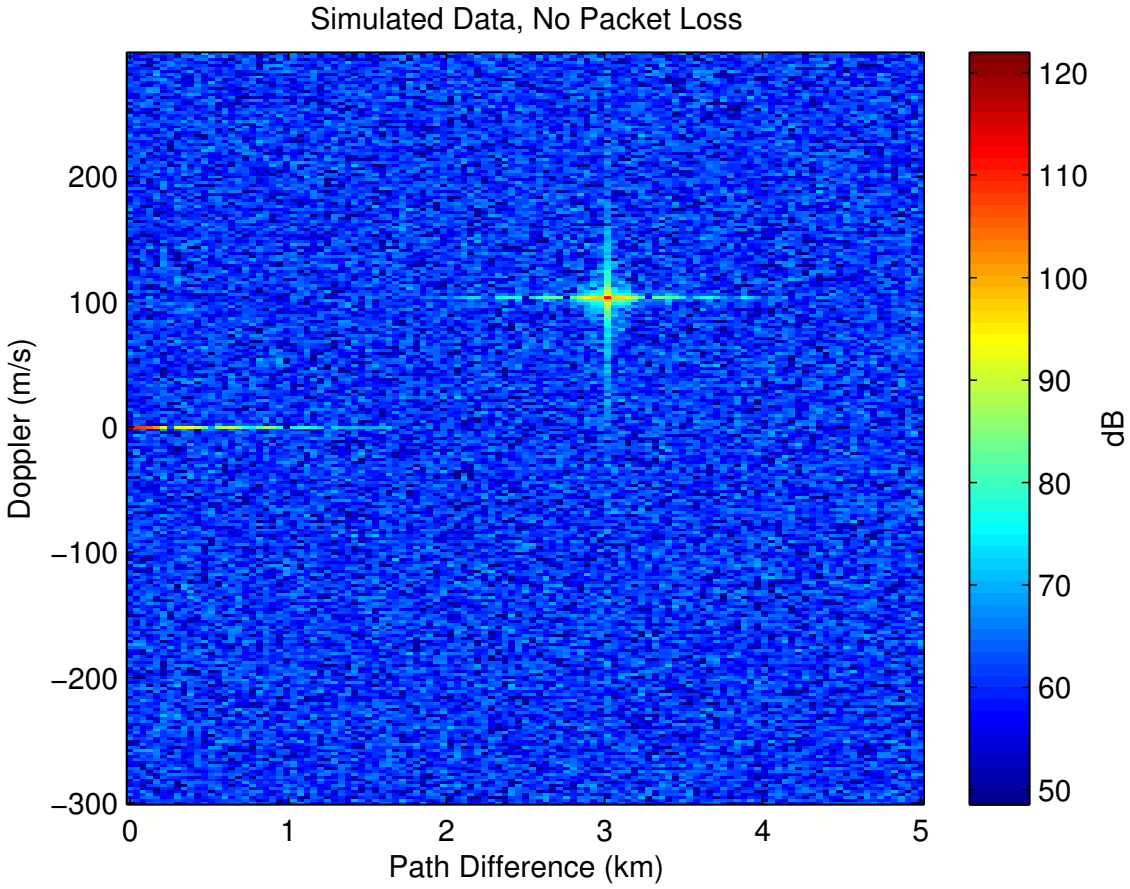

Figure 6.10: This is the cross-ambiguity calculation for a simulated DTV signal (band-limited white noise) and scatter from target with SCR = -10 dB. We will compare the affects of processing data with packet loss to this baseline. Here we see spreading in range due to the auto-correlation of anti-aliasing filters used to process the data from the fast rate of the 10 GbE link down to final IQ data rates. The target at 3 km and 100 m/s shows vertical spreading in power due to a doppler shift not perfectly aligned with the doppler samples.

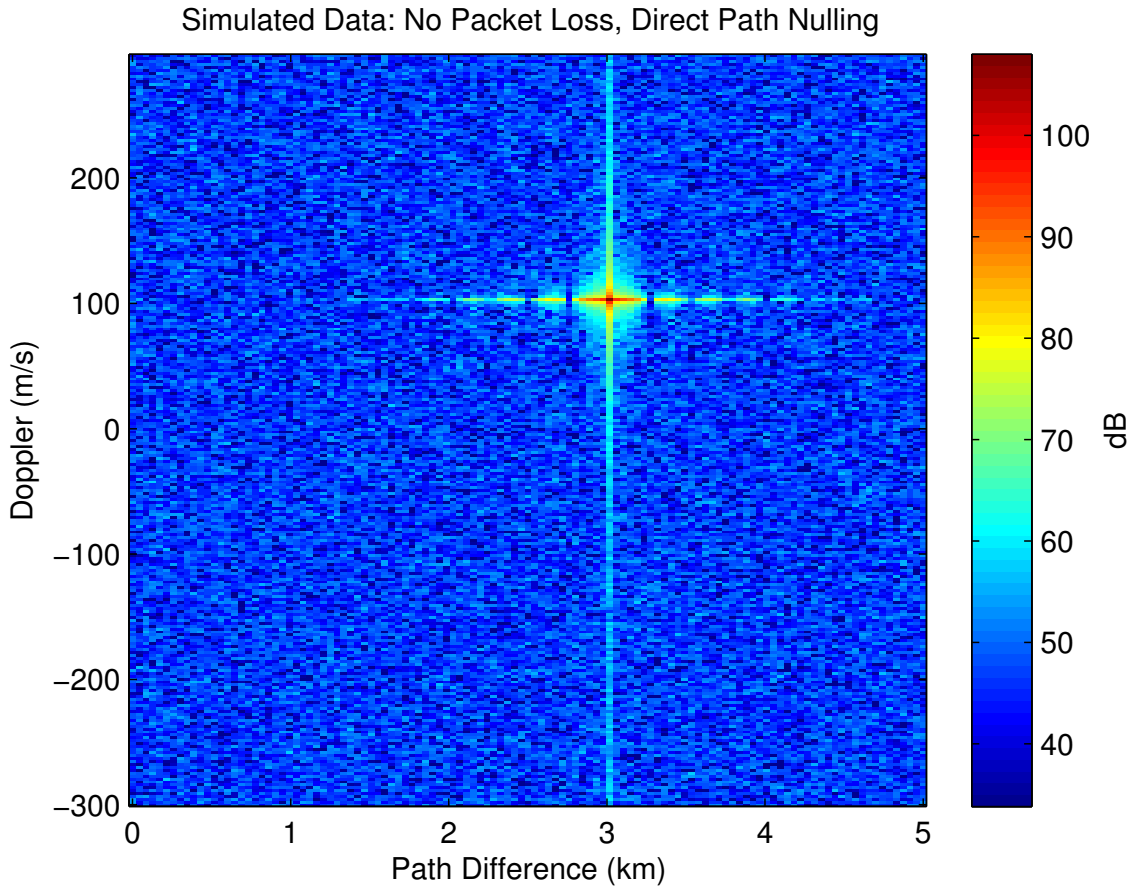

Figure 6.11: This is the cross-ambiguity calculation for a simulated DTV signal (band-limited white noise) and scatter from target with SCR = -10 dB with the direct path removed by our simple adaptive nulling. Any detection, including the direct path, contributes power across the entire 2-dimensional plot. Removing the direct path successfully shows not only is the peak at 0,0 completely removed, but the range spreading features from the direct path and the contribution to the level of the detection floor are removed as well. Pushing the detection floor down results in a detection floor linked directly to the target returned power. This allows the doppler and range spreading features associated with the processing of the target return to be more visible than Figure 6.10.

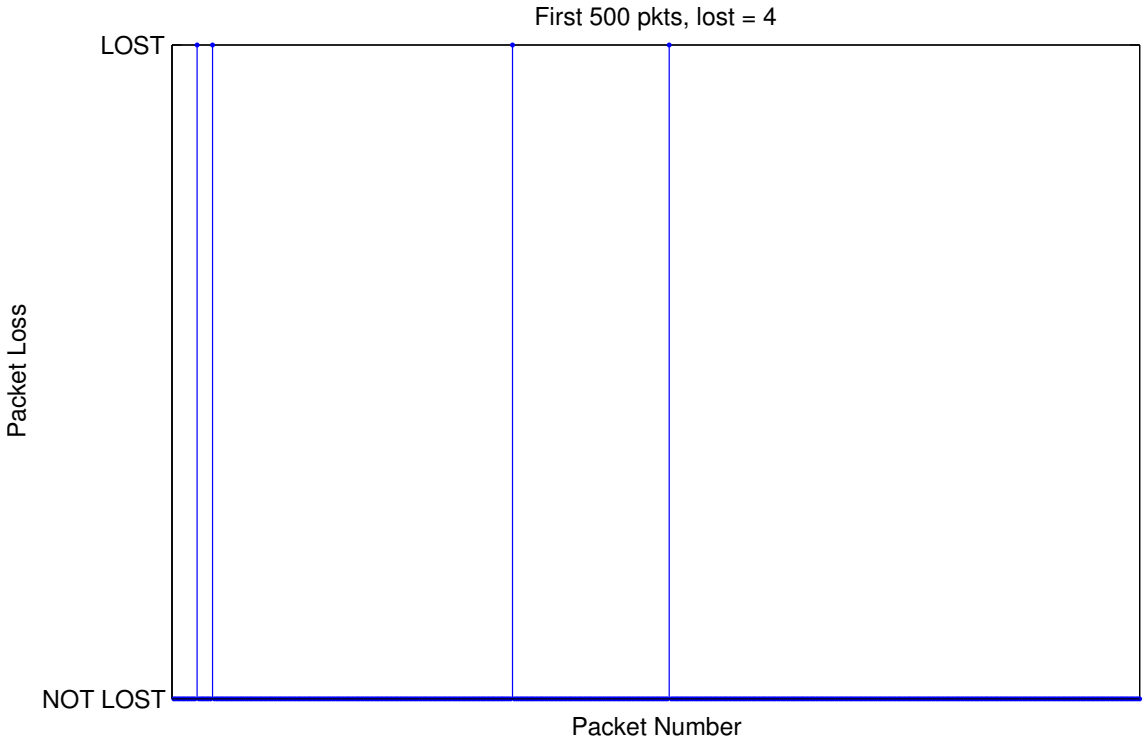

Figure 6.12: Bernoulli trial for 500 packets at a rate of 1% packet loss. This shows that packets 13, 21, 176, and 257 will be lost.

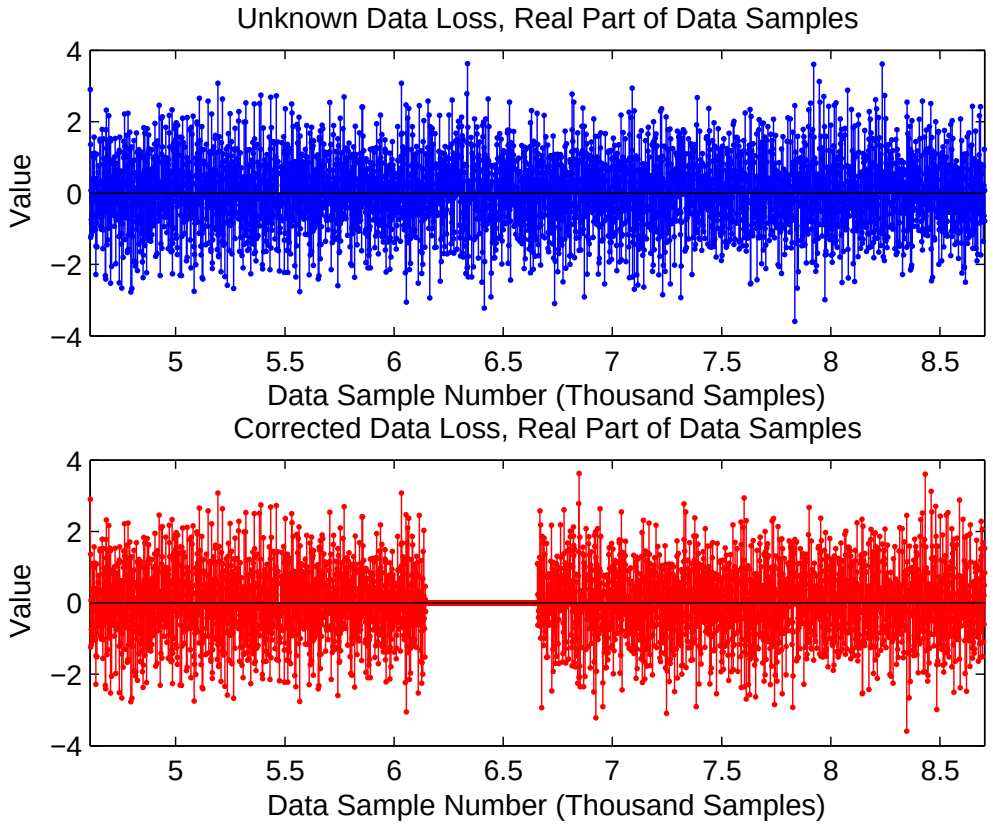

Figure 6.13: Data streams with a lost packet. The data is identical until roughly 6,000 samples at which time a packet (512 samples) of data is lost. The top plot represents unknown data loss, where these samples are simply missing. The bottom plot shows a data stream with zeros inserted in place of the missing packet.

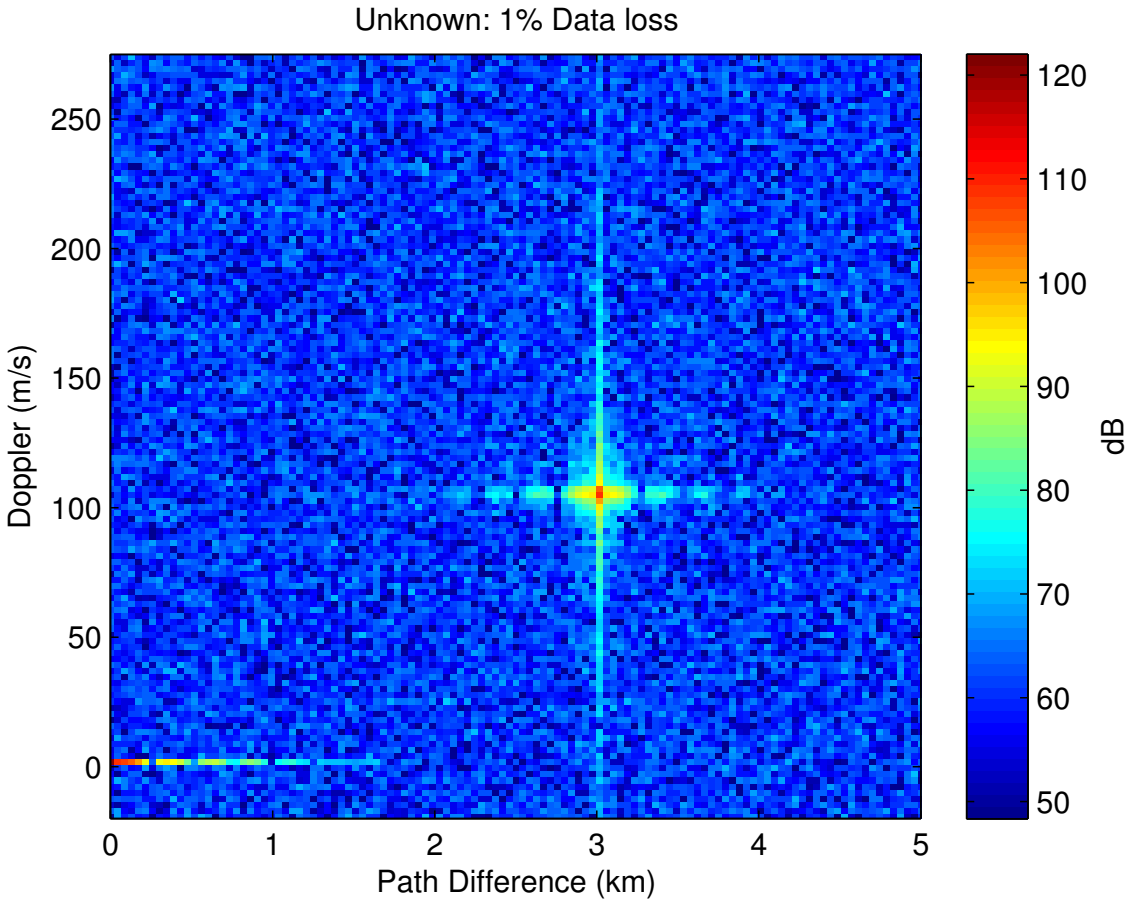

Figure 6.14: Cross-ambiguity calculation for 1 of 10 trials with 1% unknown data loss. The only difference between this and the control simulation of Figure 6.10 is that the measured Doppler shift may have increased, noted by higher power in the adjacent Doppler bin above the main target peak.

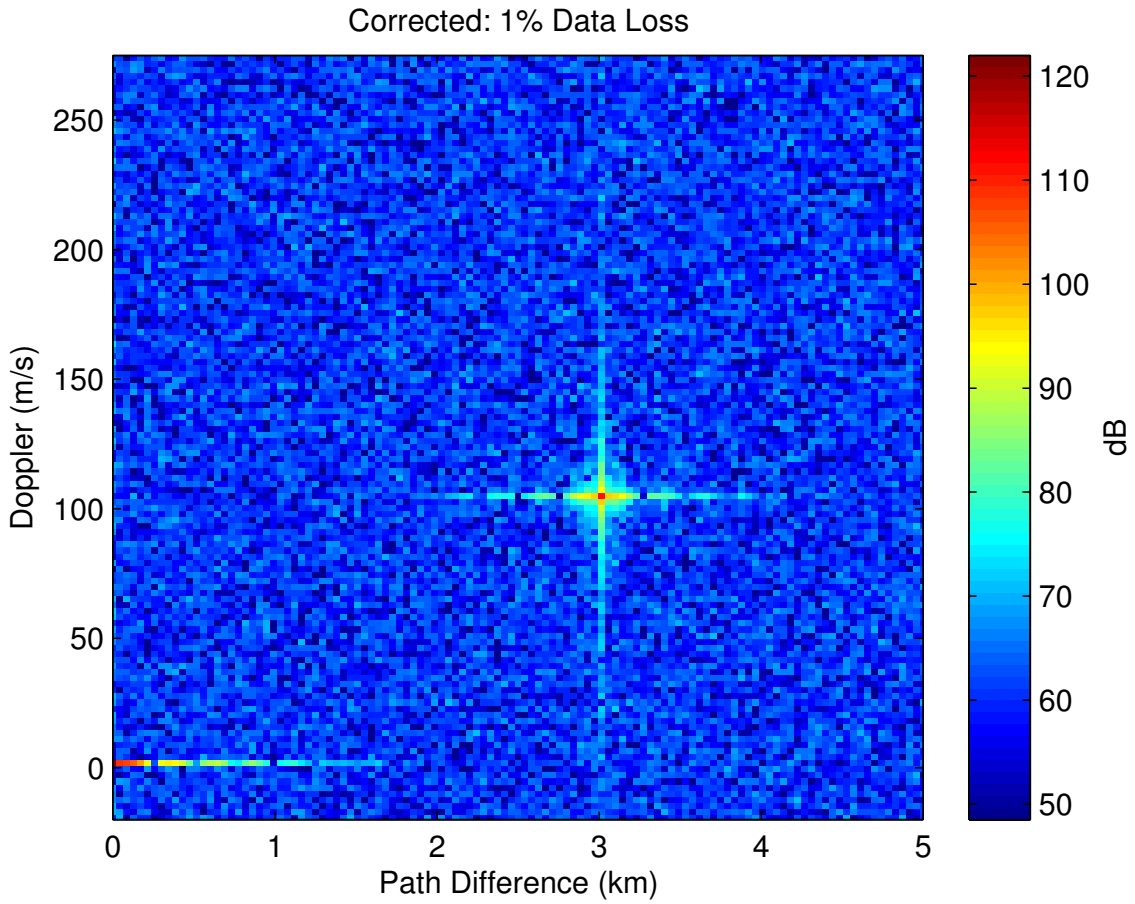

Figure 6.15: Cross-ambiguity calculation for 1 of 10 trials with 1% known data loss corrected with zeros. There are no major differences visible between this and the control simulation of Figure 6.10.

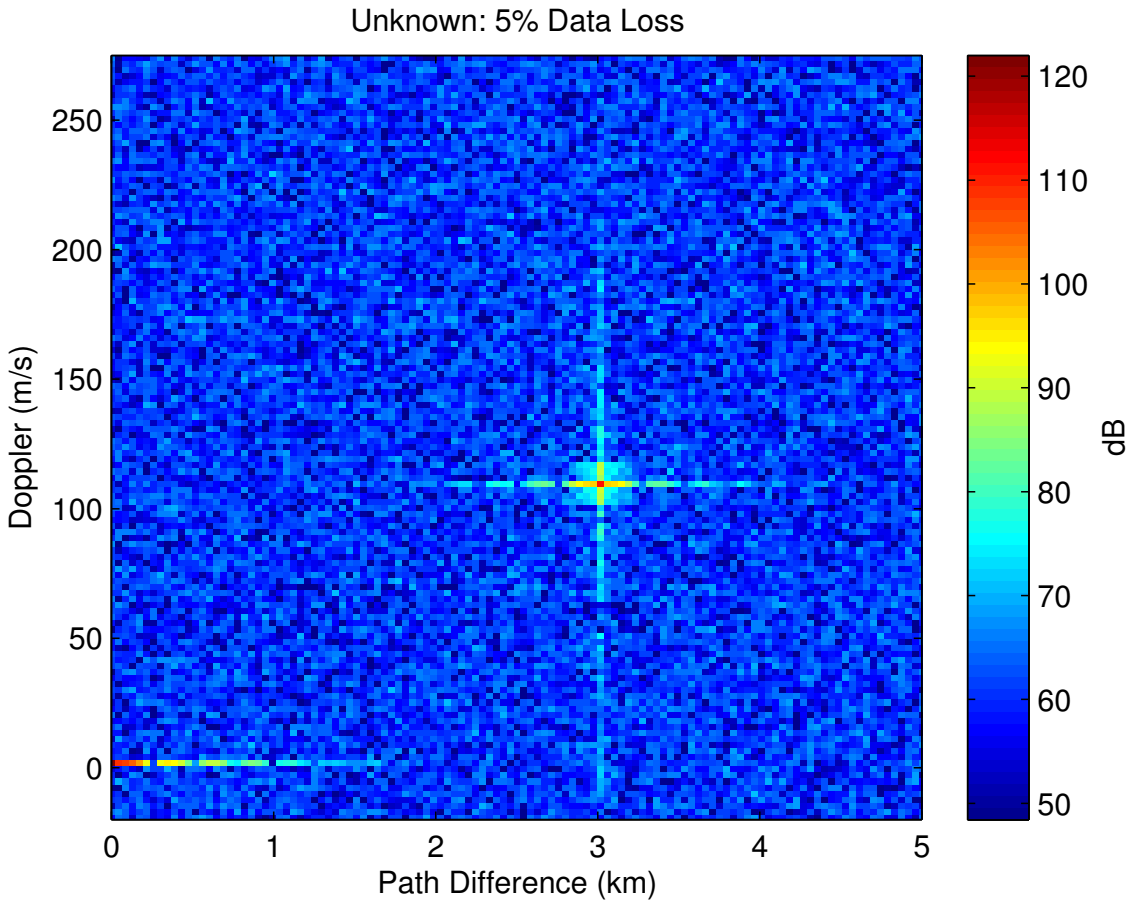

Figure 6.16: Cross-ambiguity calculation for 1 of 10 trials with 5% unknown data loss. Without knowing data was dropped, there is nothing suspicious about this plot. However, the target peak occurs at a Doppler shift 5% higher than the true value.

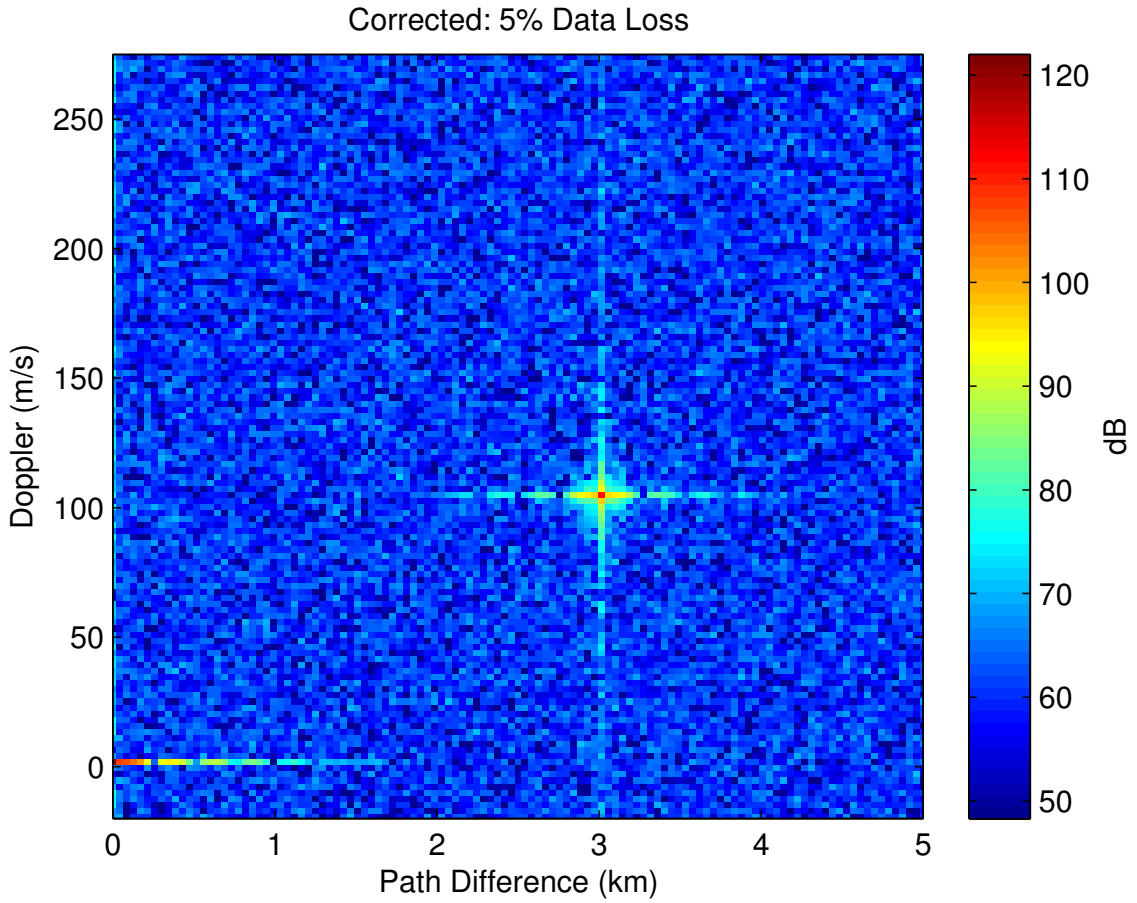

Figure 6.17: Cross-ambiguity calculation for 1 of 10 trials with 5% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data.

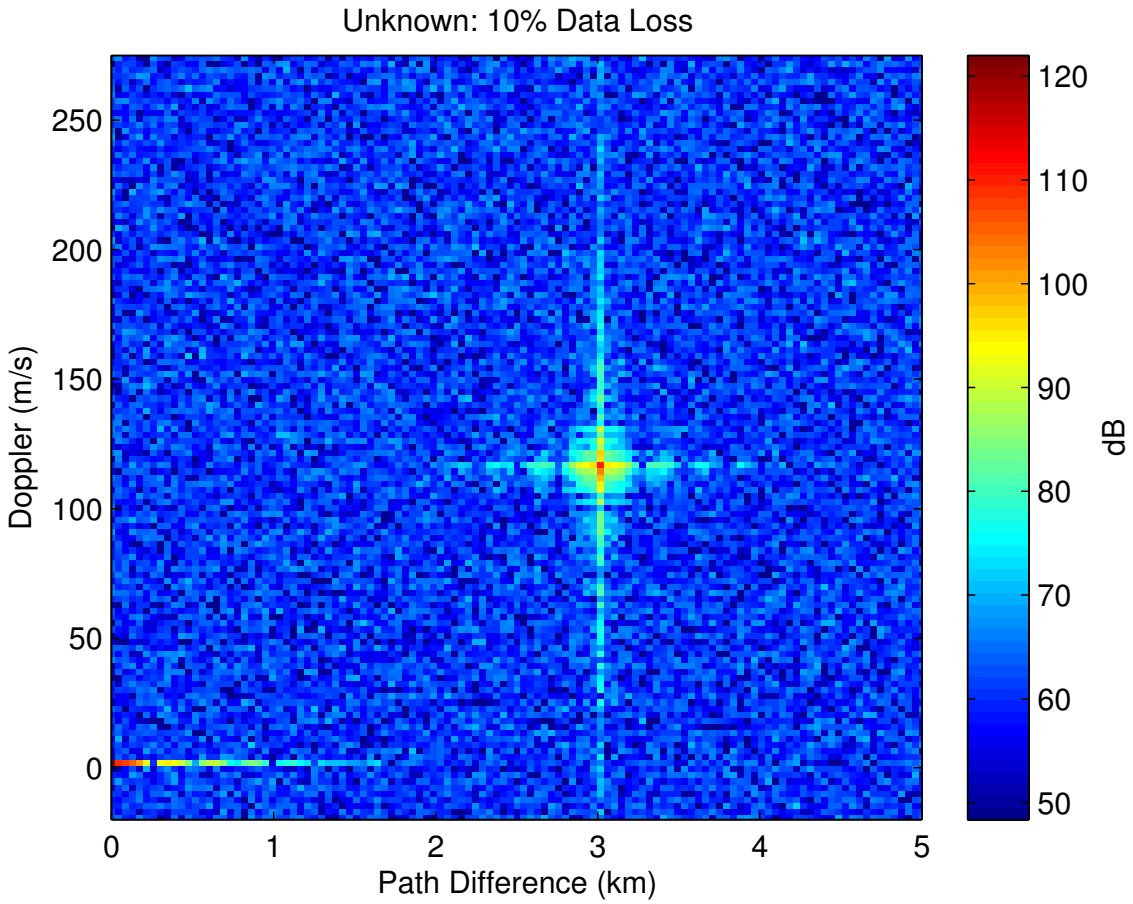

Figure 6.18: Cross-ambiguity calculation for 1 of 10 trials with 10% unknown data loss. Without knowing data loss has occurred, there is nothing particularly suspicious about this plot. However, the target peak occurs at a Doppler shift 10% higher than the true value!

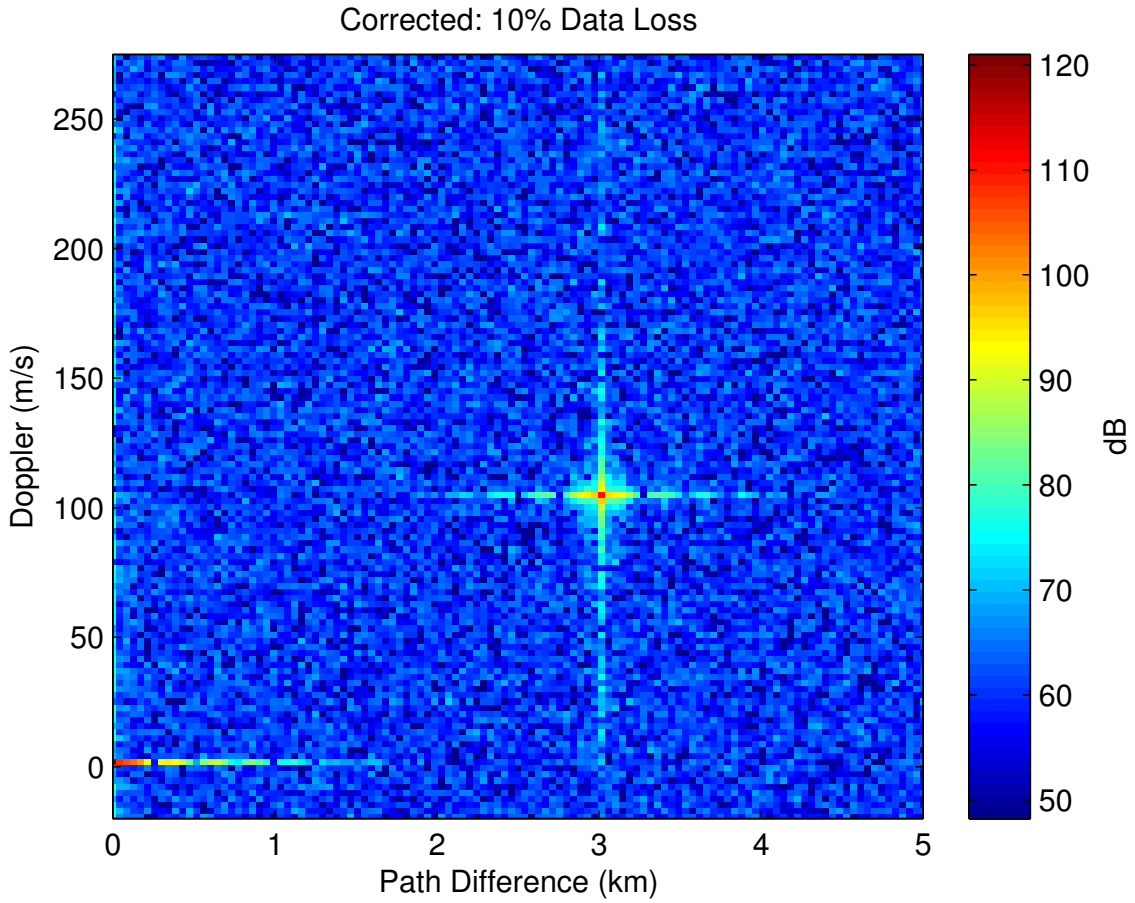

Figure 6.19: Cross-ambiguity calculation for 1 of 10 trials with 10% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data. The effect is stronger than in Figure 6.17.

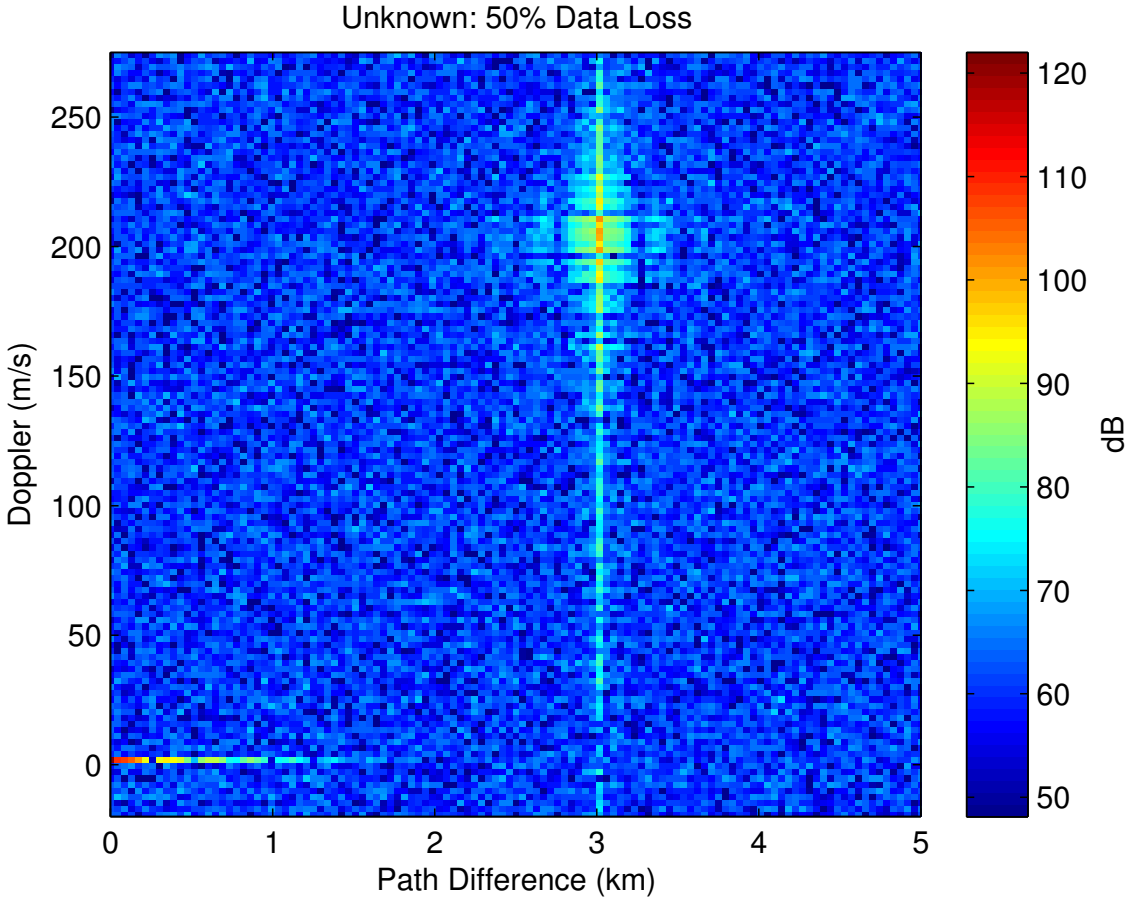

Figure 6.20: Cross-ambiguity calculation for 1 of 10 trials with 50% unknown data loss. The target has smeared in Doppler space, showing that while roughly centered around 200 m/s, there are now many peaks along the vertical Doppler slice, rather than the tight thumbtack-like formation of the control simulation with no data loss.

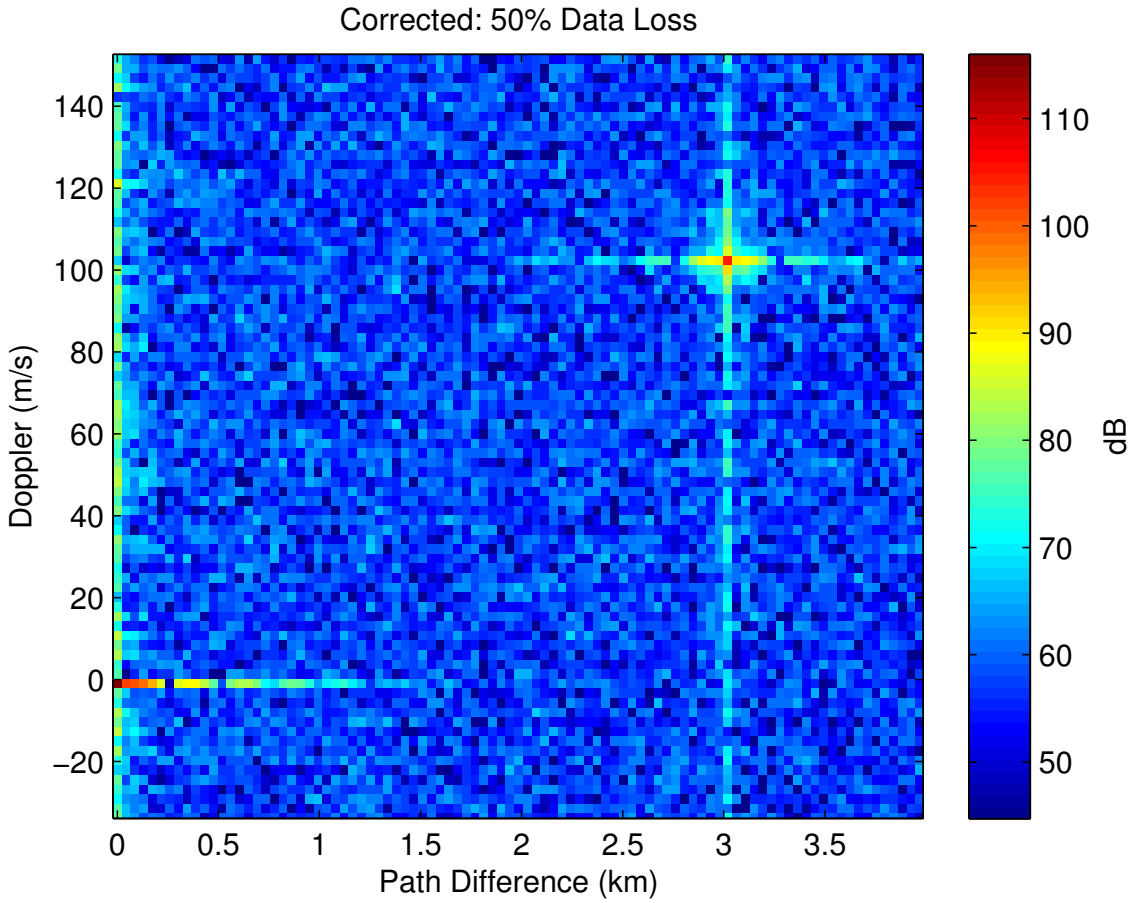

Figure 6.21: Cross-ambiguity calculation for 1 of 10 trials with 10% known data loss corrected with zeros. While the correction allows the target to be reliably detected at its true range and Doppler shift, Doppler spread (seen especially at zero range) occurs due to the nature of simply stuffing zeros in the data. This is a significant feature; the Doppler line is consistently 40 dB below the peak. This implies that reflection from the target must be less than 40 dB below any other scatterer at that range.

6.4 Strategies for Data Loss at High Bandwidth

Doppler and range estimates of targets will be inaccurate if the data loss is not detected. One obvious result is artificial increase in measured Doppler shift. However, recall that even from the IQ processing to the final cross-ambiguity, 1.5 million samples are combined to create a single range, doppler cell. With so much processing gain, a passive radar system such as this is incredibly robust to noise elements, including dropped data. By only reliably receiving 50% of the data and correcting for it with the appropriate zero-data packet substitution, the processing gain is reduced by as little as 3 dB. This is a very modest performance degradation.

Initial fixes for the data loss problem included adding header information to each UDP packet. This allows for the correct data to be properly time aligned, perhaps by inserting zeros for confirmed lost packets or something more sophisticated, which in turn allows for recovery of accurate range and Doppler measurements even in situations with much data loss.

To minimize data loss in the first place, an exhaustive rewrite of data-buffering and data-saving code should be explored. Additionally, purchase of a second ROACH board would allow the two be connected via the 10 GbE. This would allow us to operate one as usual, transmitting fast rate data out the 10 GbE connection to the second board. The second board can receive and process this data in real-time with minimal buffering, and could be used to further reduce the data rate and perhaps begin our lengthy cross-ambiguity calculations. The second board could then stream data out another 10 GbE port to a storage computer at data rates we know empiracally are reliable (below 14 MHz).

Every 64-bit word of the UDP packet is made up of time-synchronized IQ samples from each of four antennas. Data loss occuring in whole packet increments guarantees that our antennas see very similar data loss. Large bursts of the direct path filter may be lost, while at the particular time delay of a target of interest, completely different parts of the transmitted signal are dropped. This seems sub-optimal. By adding a known and different delay to each of the antennas, we could try and recreate the matched filter more precisely than simply stuffing zeros in for missing information.

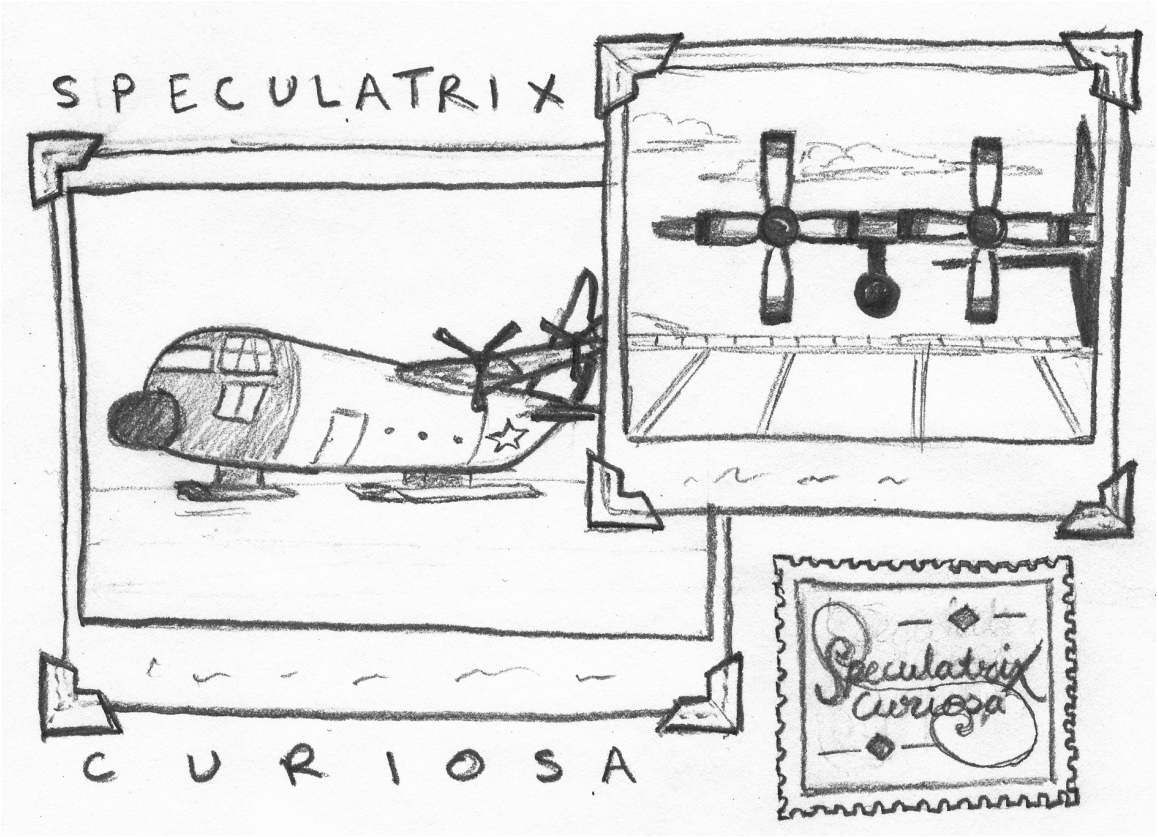
SPECULATRIX
CURIOSA
Speculatrix
curiosa

Chapter 7

AIRCRAFT DETECTION

With the fast digitizers and software-defined functionality, this radar is adaptable to several applications. The initial testing and performance evaluation can be done for the problem of detecting aircraft using DTV waveforms. In Chapter 2 we used this specific experiment to evaluate performance prediction. Here we will expand on that analysis to motivate our engineering choice of DTV as a transmitter of opportunity. The arguments and calculation we make are estimates intended to inform the engineer whether certain transmitters may be *plausible* or *preposterous* to use for aircraft detection with our system. While we come up with reasonable estimates for SNR, CNR, SCR and system gains, it is important to note that these are rough numbers and could be greatly improved with an elaborate and detailed model of system and propagation losses, the scope of which is beyond this work.

After initial discussion, we present results of aircraft detection, both with stationary receiver positioning and mobile receiver operation and preliminary results of the high aggregate bandwidth mode of operation.

7.1 SCR Performance Specification

As analyzed in Chapter 2, the SCR is fundamentally a function of geometry, rather than transmitted waveform properties like wavelength and average power. We consider the problem of detecting an aircraft 3 km from the receiver. For the worst-case scenario, we consider a maximum R_{ts}, which occurs when the aircraft is located on an extension of the transmitter/receiver line. This makes $R_{sr} = 3$ km and $R_{ts} = R_{tr} + 3$ km. Many transmitters[1] available to the passive radar system designer can be considered, as in Chapter 1. For a 0 dBsm aircraft, we can estimate the SCR at the tip of the receive antenna just using the rough estimate of distance to the transmitter, R_{tr}. These results, along with the range resolution in the monostatic sense and the power of the direct path transmission at the tip of the receiver P_c, are shown in Table 7.1. Most of these transmitters provide a similar SCR. One reason for this is that when the target is close to the transmitter or the receiver compared to the distance between transmitter and receiver (in other words, when $R_{tr} >> R_{sr}$ or $R_{tr} >> R_{ts}$), the SCR equation is entirely dependent on the single distance from target to receiver (or transmitter, whichever is the small distance). Using the calculated SCR we derive the received signal strength from the scatterer at the tip of the antenna.

Transmission	Range Res.	R_{tr} (Roughly)	P_c	P_s	SCR
HF, 20 MHz	16.7 km	1000 km	-11 dBm	-92 dBm	-81 dB
FM, KMTT Seattle	1 km	30 km	-25 dBm	-107 dBm	-81 dB
DTV, KONG Seattle	25 m	10 km	-19 dBm	-102 dBm	-83 dB
DAB (UK)	681.8 m	100 km	-49 dBm	-130 dBm	-81 dB
Cellphone (GSM)	750 m	10 km	-62 dBm	-144 dBm	-83 dB
Cellphone (3G)	30 m	10 km	-68 dBm	-151 dBm	-83 dB
GPS (LP1 P-code)	7.5 m	20200 km	-132 dBm	-212 dBm	-81 dB
Wi-Fi (802.11b)	7.5 m	100 m	-60 dBm	-170 dBm	-110 dB

Table 7.1: Various transmitters available for passive radar use and the resulting SCR estimate as pertaining to detection of an aircraft-like target 3 km from the receiver.

[1]GPS values are found from [29], and we quote $P_c = -162$W, the standard for minimum received power at ground level.

In the limit that R_{tr} gets large, R_{tr}^2/R_{ts}^2 goes to 1.

$$SCR = \frac{\sigma}{4\pi}\frac{R_{tr}^2}{R_{ts}^2 R_{sr}^2} \approx \frac{\sigma}{4\pi}\frac{1}{R_{sr}^2} \tag{7.1}$$

The SCR will approach a fixed limit (with fixed distance from scatterer to receiver, R_{sr}) as the transmitter to receiver distance increases, and the SCR settles on a minimum value of -81 dB.

A continuous curve of SCR for differing transmitter-receiver distance is shown in Figure 7.1. With the transmitter and receiver as close as 4.5 km, the bistatic radar is nowhere near the monostatic limit, however the SCR hits -85 dB, fairly close to its -81 dB upper limit. The fact that this plot is fairly uninteresting is actually quite excellent. It means that our performance metric holds for a wide range of geometries.

This asymptote is actually approached quickly, as indicated in the table. Almost all of these transmitters for their rough distance estimate show us similar SCR measurements. This also tells us that the R_{tr} estimate can be rough and the SCR is not terribly affected. This is helpful to know, as it reassures us that if our receiver is mobile, as long as it is far enough from the transmitter, our SCR performance does not change quickly.

It is important to note that a transmitter far away results in lower SCR than one nearby. This is drastically illustrated by the final row of the table, where we estimated the position of the receiver to be fairly close to a Wi-Fi transmitter, while the target still remains a large 3 km away. This places a strain of over 110 dB dynamic range on the receiver design, arguably making Wi-Fi a more difficult choice of transmitter for aircraft detection. Just looking at SCR, however, does not necessarily point to DTV as the superior waveform.

Also shown in Table 7.1 is the range resolution in the monostatic sense. One consideration for detecting a target such as an aircraft is how precisely its range can be determined. Without the need for additional processing or tracking algorithms to reduce the range resolution, DTV, 3G Cellphone transmission, GPS[2], and Wi-Fi all have native range resolution roughly the size of a large aircraft or smaller, an ideal range resolution for locating the

[2]GPS values are found from [29], and we quote P_c simply using the minimum received power at ground level standard, rather than make any assumptions about atmospheric or propagation losses.

target. This range resolution metric is not the final word, however. Several systems take advantage of the prolific nature of FM broadcast (1 km range resolution, natively) to detect aircraft with improved range resolution from beam-forming and processing of multiple stations, such as [36].

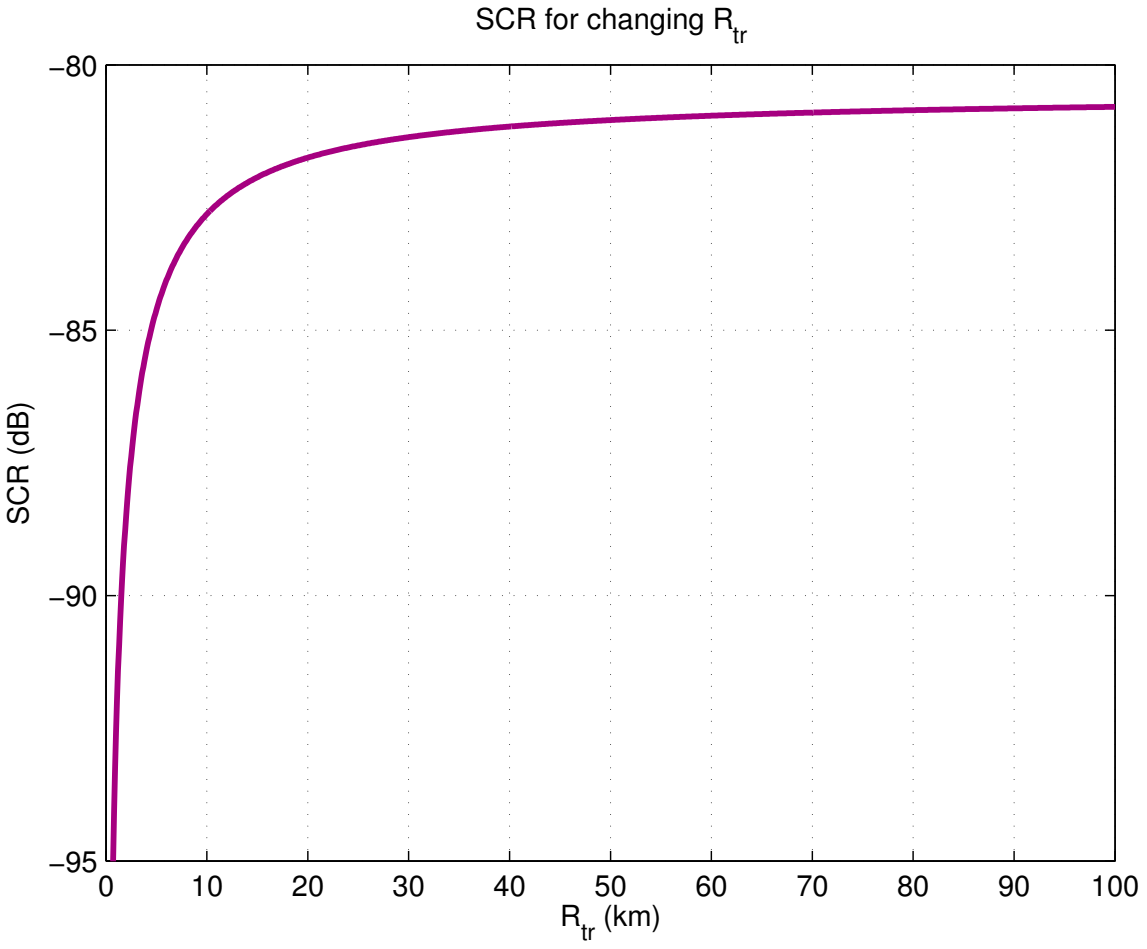

Figure 7.1: SCR determined for a 0 dBsm aircraft, with a distance of 3 km to the receiver. As the receiver and transmitter are separated, this shows the SCR approaches a limit of -81 dB. Any distance greater than 4.5 km separation (not close to the monostatic limit!) will have an SCR above -85 dB.

7.2 When SNR matters

7.2.1 Calculating Precision Gain

Chapter 2 showed that DTV-based aircraft detection is fundamentally limited by SCR rather than SNR. Is that the case for all types of transmission? As intuition would have you guess, that while moving the receiver farther away from the transmitter improves SCR, the received power drops off as R^2. To analyze how problematic this is, we begin by exploring how the largest noise source, quantization noise from the 8-bit digitizer, is suppressed through processing. With a full-scale voltage of 870 mV, the noise power is -47.2 dBm. Given the ADC parameters of bit-depth and sampling rate, we can determine the amount of Precision Gain, or simply gain over the truncation noise due to oversampling, for each waveform. Precision Gain can also be thought of as the effective increase of the ENOB by signal processing means.

By setting the ADC to different sampling rates, we can explore the amount of precision gain we achieve due to oversampling. Here we explore the limit of 3 Gsamp/s sampling, as well as our standard modes of 1.6 GSamp/s and 800 MSamp/s for the 6 MHz bandwidth of a DTV transmission.

$$\text{Precision Gain} = \begin{cases} (3GS/s) & : & 500 & \equiv & 27 \text{ dB} & \equiv & 4.5 \text{ bits} \\ (1.6GS/s) & : & 266.67 & \equiv & 24 \text{ dB} & \equiv & 4 \text{ bits} \\ (800MS/s) & : & 133.33 & \equiv & 21 \text{ dB} & \equiv & 3.5 \text{ bits} \end{cases} \tag{7.2}$$

The radar operates as if it had a 12.5 bit ADC, a 12 bit ADC, or an 11.5 bit ADC, respectively. Every factor of 2 in sampling rate provides 3 dB in gain, or equivalently half a bit to the ADC precision. We can explore this oversampling gain for several transmitters in Table 7.2. To receive these other transmitters we may condsider a large low-pass filter or band-pass filter on the front end. If we choose a low-pass filter with cutoff at 1.5 GHz (for sampling at 3 GHz), then we note that the assumption of quantization noise exceeding analog system noise is still good, as the system temperature would have to approach $\approx$ 929,000 K to compare.

		3 GSamp/s		1.6 GSamp/s		800 MSamp/s	
Transmission	Bandwidth	Gain (dB)	bits	Gain (dB)	bits	Gain (dB)	bits
HF, 20 MHz	9 kHz	55	9.2	52	8.7	49	8.2
FM, KMTT Seattle	150 kHz	43	7.1	40	6.7	37	6.2
DTV, KONG Seattle	6 MHz	27	4.5	24	4.0	21	3.5
DAB (UK)	220 kHz	41	6.9	39	6.4	36	5.9
Cellphone (GSM)	200 kHz	42	6.9	39	6.5	36	6.0
Cellphone (3G)	5 MHz	28	4.6	25	4.2	22	3.7
GPS (LP1 P-code)	20 MHz	22	3.6	19	3.2	16	2.7
Wi-Fi (802.11b)	20 MHz	22	3.6	19	3.2	16	2.7

Table 7.2: Various transmitters available for passive radar use and the resulting Precision Gain based on initial sampling rate.

Transmissions with large bandwidth have good range resolution, but lower Precision Gain. 3G Cellphone transmission has similar Precision Gain characteristics. Note if HF broadcast is of interest (10-30 MHz, depending on time of day), the Precision Gain is a large 55 dB, effectively turning our KATADC into a 17.2 bit digitizer operating at 3 GSamp/s! The relatively low bandwidth of FM broadcast offers 43 dB precision gain to our system, and the lare range resolution inherent in the signal can be improved through tracking of aircraft.

Oversampling gain, or precision gain, is not restricted to the case of aircraft detection. It is simply the advantage all coherent signals gain over quantization noise due to precision of the ADC. Through signal processing, fast rate data can be filtered (think a type of averaging) and reduced down to the minimum Nyquist rate for the transmitter bandwidth. In that sense uncorrelated noise is filtered out in a measurable way. To sample directly at this minimum Nyquist rate the 8-bit digitizer SNR cannot be improved.

7.2.2 Processing Gain

Processing Gain is applied by coherent integration. This gain is realized by any signal that correlates well in a specific range and doppler cell, where other components of power scattered from ground clutter, targets, and the direct path transmission will be uncorrelated. For example, the direct path transmission will correlate well at zero range and zero doppler, where scattered signal from ground clutter farther away, scattered signal from doppler-

shifted targets, and system and truncation noise will be uncorrelated with the matched filter at zero, zero. Similar to pulse compression, matched filtering allows for large gains in power as well as compression of radar echoes into a narrow spike in range and Doppler space, allowing computational separation of these echoes. The processing gain is defined as the coherent integration time multiplied by the bandwidth of the transmitted signal. The coherent integration time can be calculated by several means. For the sake of argument, we use the amount of time the signal scattered from an aircraft-like target is coherent, as described in [20]. It is a function of bandwidth, center frequency, and properties of target motion.

Coherent integration time, τ is described in [20], and can be determined by:

$$\tau = \sqrt{\frac{\lambda}{a_r}} \tag{7.3}$$

Based on our DTV data, we can estimate a maximum radial acceleration, a_r for aircraft of roughly 8 m/s[23], a data-driven estimate as we have found integration time of 0.25 seconds to provide good results. Using this target radial acceleration estimate, we can determine coherent integration times for other transmitters, and in turn determine the processing gain for each. The processing gain is therefore a function of both center frequency and bandwidth, and is presented in Table 7.3.

Transmission	BW	Int. Time, τ	Process Gain (BW*τ)
HF, 20 MHz	9 kHz	1.34 s	41 dB
FM, KMTT Seattle	150 kHz	0.59 s	49 dB
DTV, KONG Seattle	6 MHz	0.25 s	62 dB
DAB (UK)	220 kHz	0.41 s	50 dB
Cellphone (GSM)	200 kHz	0.2 s	46 dB
Cellphone (3G)	5 MHz	0.13 s	58 dB
GPS (LP1 P-code)	20 MHz	0.15 s	65 dB
Wi-Fi (802.11b)	20 MHz	0.12 s	64 dB

Table 7.3: Various transmitters available for passive radar use and the resulting Processing Gain for an aircraft-like target.

[3]This is an incredibly conservative estimate. Note 8 m/s^2 is close to 1.0g, a very aggressive maneuver for commercial or civilian aircraft.

By these metrics, DTV and Wi-Fi both have very similar performance in terms of processing gain.

$$\begin{gathered} \tau_{\text{Wi-Fi}} = 0.125 \text{ s} \\ \text{Processing Gain} = \tau \times BW \equiv 64 \text{ dB} \end{gathered} \tag{7.4}$$

The length of time a target is coherent is based on the wavelength of the radar waveform and the maximum radial acceleration of the target. The Doppler velocity, however, can also limit the integration time through an effect called range migration. There is a time for which the target dwells in a single range bin based its velocity. A fast-moving target, say 300 m/s Doppler velocity, illuminated by a 6 MHz DTV signal may pass through a single range bin in 0.16 s. If integration time is long, this range migration will result in the target power smearing across adjacent range bins. This effect should be calculated for a particular passive radar application to settle on a standard integration time for the system. With the IQ signals saved to disk, however, we are always afforded a chance to reprocess the data if these effects are noticed. Again, this processing gain can be considered an estimate to inform us if aircraft detection with each transmitter may be *plausible* or *preposterous*.

Choosing a rough distance from the DTV transmitter, perhaps choosing a distance that reflects the density of these types of transmitters in the area, we can analyze the power received at the tip of the antenna from the direct path as well as the aircraft.

This analysis is carried out for the potential transmitters and is summarized here in Table 7.4. For aircraft detection based on higher frequency transmitters, we would upgrade our front end amplifier to the SBB-5089Z (50 MHz - 6.0 GHz). To explore how agile our system is in terms of frequency, we assume a front end gain of 0 dB for each transmitter. The attenuator can be programmed in milliseconds to values from 0 to -31.5, making the total front end gain variable from +20 dB to -11.5 dB. If a particular transmitter is desired, the system designer can explore the affect of these limits on the calculations in Table 7.4.

	After ADC		Baseband IQ		Final	
Transmission	SCR	SNR	SCR	SNR	SCR	SNR
HF, 20 MHz	-81 dB	-45 dB	-81 dB	10 dB	-40 dB	51 dB
FM, KMTT Seattle	-81 dB	-59 dB	-81 dB	-16 dB	-32 dB	33 dB
DTV, KONG Seattle	-83 dB	-55 dB	-83 dB	-28 dB	-21 dB	34 dB
DAB (UK)	-81 dB	-83 dB	-81 dB	-42 dB	-31 dB	8 dB
Cellphone (GSM)	-83 dB	-97 dB	-83 dB	-55 dB	-37 dB	-9 dB
Cellphone (3G)	-83 dB	-104 dB	-83 dB	-76 dB	-25 dB	-18 dB
GPS (LP1 P-code)	-81 dB	-165 dB	-81 dB	-143 dB	-16 dB	-78 dB
Wi-Fi (802.11b)	-110 dB	-123 dB	-110 dB	-101 dB	-46 dB	-38 dB

Table 7.4: Various transmitters available for passive radar analyzed for detection of an aircraft-like target. We analyze the SNR and CNR gain through stages of the radar processing. Including both Precision and Processing Gain, we explore the SNR and SCR values at several points in our processing scheme. We assume the front end attenuation is equal to 0 dB for all cases.

Analysis of the SCR and SNR after the final stage of cross-ambiguity calculation shows whether our theoretical aircraft is detectable. If both quantities are positive, then the signal power from the target will just be at the threshold of the cross-ambiguity floor. Other systems may even wish to impose a detection threshold, in which case both the SNR and SCR would have to exceed this threshold for confirmed detection. In Table 7.5 compare these values to determine if aircraft detection is clutter-limited or quantization noise-limited, and we calculate the difference between these limits in the CNR. Note this reference to CNR is the clutter power to noise power in the entire cross-ambiguity plot except the value at range zero and Doppler zero. For example, in the case of GSM cellphone transmission as an illuminator of opportunity, if the clutter can be suppressed by 7 dB the aircraft detection will be limited by the quantization noise, but no amount of clutter suppression alone will make the target detectable (SNR still negative). All transmitters but GPS fall into the clutter-limited refime. The SCR for GPS is roughly the same as many of the other transmitters because R_{tr} is so large, however this distance causes power receieved from the direct path and scattered from the aircraft is to be small compared to the quantization noise.

Transmission	Final SCR	Final SNR	Limit?	CNR
HF, 20 MHz	-40 dB	51 dB	Clutter	+91 dB
FM, KMTT Seattle	-32 dB	33 dB	Clutter	+65 dB
DTV, KONG Seattle	-21 dB	34 dB	Clutter	+55 dB
DAB (UK)	-31 dB	8 dB	Clutter	+39 dB
Cellphone (GSM)	-37 dB	-9 dB	Clutter	+27 dB
Cellphone (3G)	-25 dB	-18 dB	Clutter	+7 dB
GPS (LP1 P-code)	-16 dB	-78 dB	Q Noise	-63 dB
Wi-Fi (802.11b)	-46 dB	-38 dB	Clutter	+9 dB

Table 7.5: Various transmitters available for passive radar use and their properties as pertaining to detection of an aircraft-like target. Including both precision and processing gain, we determine for each system whether the waveform would result in a clutter-limited or quantization noise-limited scenario.

7.2.3 Clutter Suppression

Computational suppression of the direct path and other ground clutter is performed on the baseband IQ data. For argument's sake, we assume the system can accomplish 25 dB clutter suppression. After matched filtering this computational scattered signal (a weighted subtraction of one IQ signal from one derived from an adjacent antenna), the final SCR increases by 25 dB, while the SNR remains unchanged. These results are presented in Table 7.6. With this gain, the GSM Cellphone transmission and Wi-Fi used with this system for aircraft detection would now fall into the quantization noise-limited category. Additionally, here we see that DTV is the first and only transmission to achieve positive SNR and SCR, meaning we expect this theoretical target to begin to rise out of the detection floor of the cross-ambiguity. This verifies that DTV is a good choice for this type of aircraft detection. From this analysis, we learn that if we can squeeze 15 dB more digital clutter suppression (a total of 40 dB), HF, DAB and FM could all be used to detect aircraft.

Transmission	Final SCR	Final SNR	Limit?	CNR
HF, 20 MHz	-15 dB	51 dB	Clutter	+66 dB
FM, KMTT Seattle	-7 dB	33 dB	Clutter	+40 dB
DTV, KONG Seattle	**4 dB**	**34 dB**	Clutter	+30 dB
DAB (UK)	-6 dB	8 dB	Clutter	+14 dB
Cellphone (GSM)	-12 dB	-9 dB	Clutter	+2 dB
Cellphone (3G)	0 dB	-18 dB	Q Noise	-18 dB
GPS (LP1 P-code)	9 dB	-78 dB	Q Noise	-88 dB
Wi-Fi (802.11b)	-21 dB	-38 dB	Q Noise	-16 dB

Table 7.6: Using the system with an estimated 25 dB clutter suppression (direct path and ground clutter if possible). Aircraft detection with GSM cellphone transmission or Wi-Fi is now limited by quantization noise.

7.2.4 Noise Suppression

If we take our 25 dB clutter suppression as consistently achievable, then 3G cellphone and GPS may be able to detect aircraft if the quantization noise is suppressed (SCR positive, SNR negative). If we desire to keep the front end gain identical for each all transmitters, maintaining full frequency agility in the DSP realm, then there are two changes to the system that will accomplish this SNR gain for every transmitter: increasing the bit-precision of the ADC or adding a front end analog amplifier.

An 18 dB increase in SNR is needed to make the aircraft illuminated by GSM cellphone transmission barely detectable. GPS, on the other hand, requires close to 80 dB noise suppression. A 3 GSamp/s ADC with 11 bits will get us the extra 18 dB (3 extra bits, roughly 6 dB per bit), but a 21-bit ADC (!) is needed to achieve the 80 dB of noise level suppression necessary for both. Oversampling can equivalently achieve higher bit-precision. That means we could operate an 8-bit ADC at 36 GSamp/s to incorporate our GSM Cellphone signal (3 extra bits is a factor of 12 in frequency), or at 156 GSamp/s (!) (13 extra bits is a factor of 52 in frequency) to incorporate all of the transmitters we have analyzed. With faster sampling, all transmitter may fit into a single Nyquist zone allowing the channel selection or radar application decisions entirely digital and most importantly software-programmable. These values may make aircraft detection with GPS-illumination with our system tough today, but in the future this may not be preposterous.

Amplification is another choice to increase the signal power relative to the quantization

noise. Amplifiers will also increase any external noise and clutter elements. Amplifiers may limit frequency agility as they usually cover a particular band. To detect the GPS signal, therefore, could mean sacrificing the ability to observe FM broadcast on the same antenna. That may be a fine sacrifice depending on the particular application.

It may not be feasible to observe all subsets of these transmitters from a single digitizer as we have done with FM and DTV due to the fact that the entire set spans several Nyquist zones (two zones in the case of 3 GHz sampling, up to eight for 800 MHz sampling). Another technical difficulty to simultaneous reception of any subset of this list is that all of these transmitters are not available in the same geographic location.

From this exercise we see that DTV is an excellent candidate for aircraft detection. We learn that FM, HF, and DAB may be within our reach. With a deeper precision or faster ADC (12-bits at 3 GHz and or 8-bit at 36 GHz, affordable in the near future) we someday may be able to incorporate GSM and 3G Cellphone transmitters as aircraft illuminators. It appears that Wi-Fi and GPS would be far-fetched.

Of course, this analysis is just for aircraft detection. Other applications would lead to different metrics. As we noted in Chapter 1, Wi-Fi is in passive radar receivers designed for mapping people in buildings.

7.3 DTV-based aircraft detection: Stationary Results with two transmissions

When plotting the cross-ambiguity, we use the bistatic path difference, $R_{ts} + R_{sr} - R_{tr}$, and doppler velocity defined by the rate of change of this quantity, rather than range and doppler velocity as defined previously for the limit of monostatic operation. These terms are more informative for the geometries of near-target detection, especially when forward scatter is possible. Figure 7.2 shows the basic cross-ambiguity calculation with the receiver stationary at Sieg Hall, using the KZJO transmitter on Capital Hill, without any attempts to null the direct path. In this figure we can see the strong returns along the horizontal 0 doppler line. This line contains any auto-correlation features of the transmitted waveform as well as reflections from ground clutter. The vertical axis to the left of the figure shows the frequency structure in the self-ambiguity of the waveform. The data rate is ≈7 MHz, and we can see the resulting 42 m range bins. With a coherent integration of 0.24 s and 4 incoherent averages, with the cross-ambiguity uses 1 full second of data. We can also measure the dynamic range of the system to be 59 dB. This particular data collection was 10 seconds long, but can be as long as 1 minute, with the limit only being large file sizes. The color scale is in arbitrary units of dB power. There are no obvious detections in this view. Figure 7.3 shows the same data with some initial attempts to null the direct path. We are able to suppress the direct path at (0,0) by 20 dB and the floor by 1-2 dB. This only improves our estimate of dynamic range to 60 dB. Even with G_N only a few dB, the results are dramatic. The SFR improved quite a bit (more than G_N) which tells us we are doing better than predicted, but that we still have more learning to make a more accurate estimate of performance. This view shows clearly 3 different aircraft: a helicopter (22 dB above the floor) at a range of 0.3 km with positive doppler velocity 43 m/s; an aircraft (13 dB above the floor) at 2.4 km range with negative doppler velocity -70 m/s; and an aircraft (19 dB above the floor) at 3.6 km with positive doppler velocity 60 m/s. The detections are even more compelling seeing the progress across the full 10 s of data, shown in Appendix B. The helicopter is particularly fascinating in this frame as we can see the doppler spread associated with the blade modulation.

Figure 7.4 shows the same moment in time using the KONG transmitter. We see a

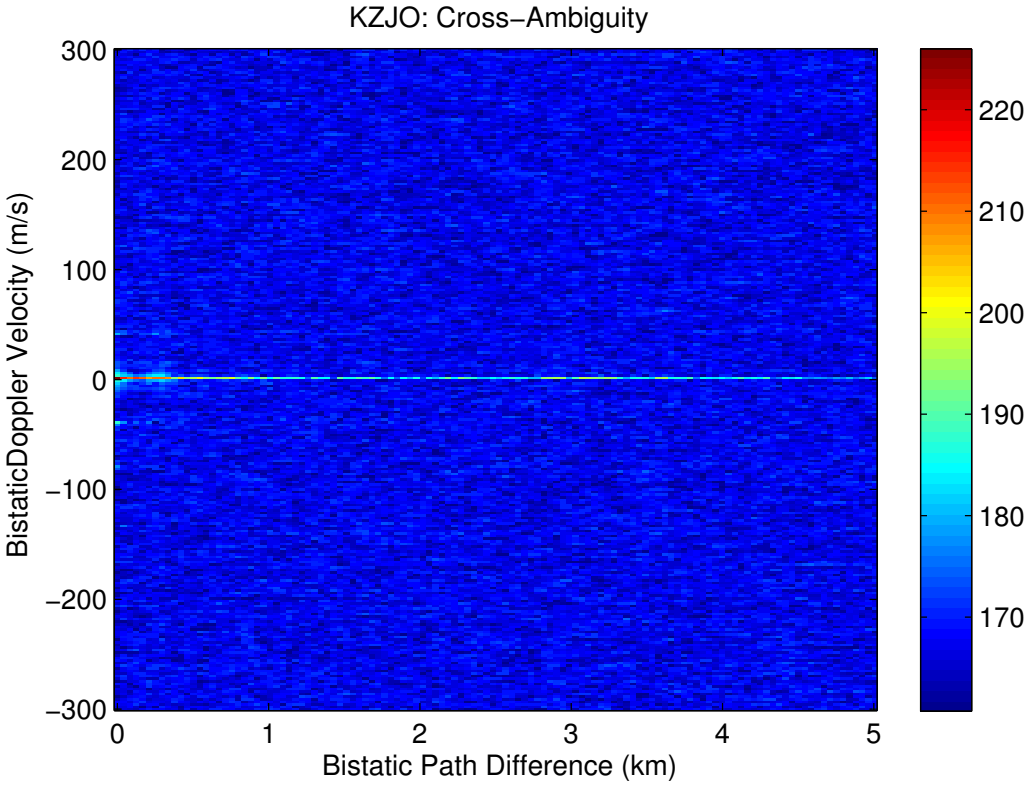

Figure 7.2: Cross-ambiguity output for stationary operation of radar using KZJO with no attempt to null the direct path. We cannot detect any targets with confidence.

similar dynamic range of 60 dB. Again, no detections are obvious here. For KONG we were able to reduce the direct path at (0,0) by 32 dB and the floor by 17 dB! These excellent improvements in nulling are due to the fact that we have line of sight from the transmitter to the receiver. This means there is a lot of direct path signal to remove. The transmitter at KZJO, however, is blocked by buildings. This should improve our SCR at the tip of the antenna but limits the amount of dynamic range we gain from nulling. With G_N=17 dB, we improve our dynamic range to 77 dB. Figure 7.5 shows the effect of the nulling, and the detection of the helicopter (19 dB above the floor) at a range of 0.7 km and positive doppler velocity 46 m/s. We can possibly make out a faint hint of the blade modulation, again more obvious as one looks at the full 10 s of data. Another aircraft is visible (17 dB above the floor) at a range of 1 km with a negative doppler velocity 8 m/s.

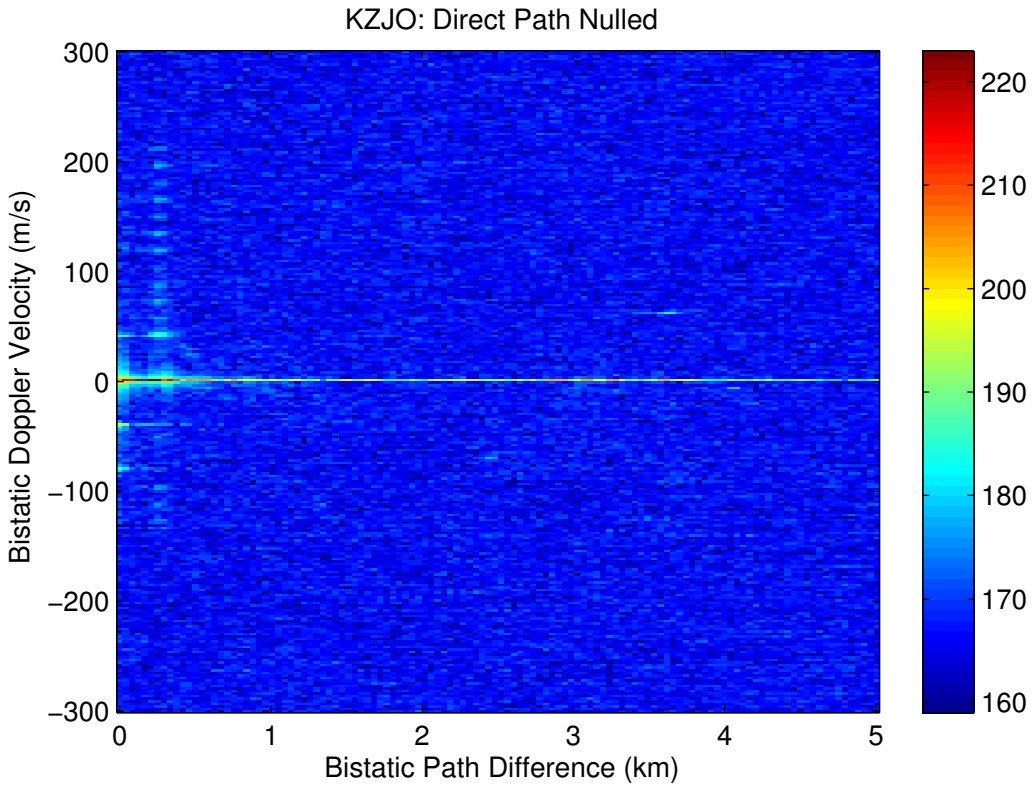

Figure 7.3: Cross-ambiguity output for stationary operation of radar using KZJO with attempt to null direct path. We null the direct path at (0,0) by 20 dB, reducing the floor and allowing targets to be detected clearly. Three aircraft are seen here, with the most interesting being the helicopter (seen by the doppler spread blade modulation) at close range.

Data association and tracking for this radar are not in place at this time, however we can plot the bistatic ellipse associated with the simultaneous helicopter detection from KZJO and KONG together in Figure 7.6. The helicopter will fall on the arc where these shapes intersect, which corresponds to our notes of the helicopter flying close to the UW from south to north. The evolution of these ellipses can be seen in the animations, linked in Appendix B.

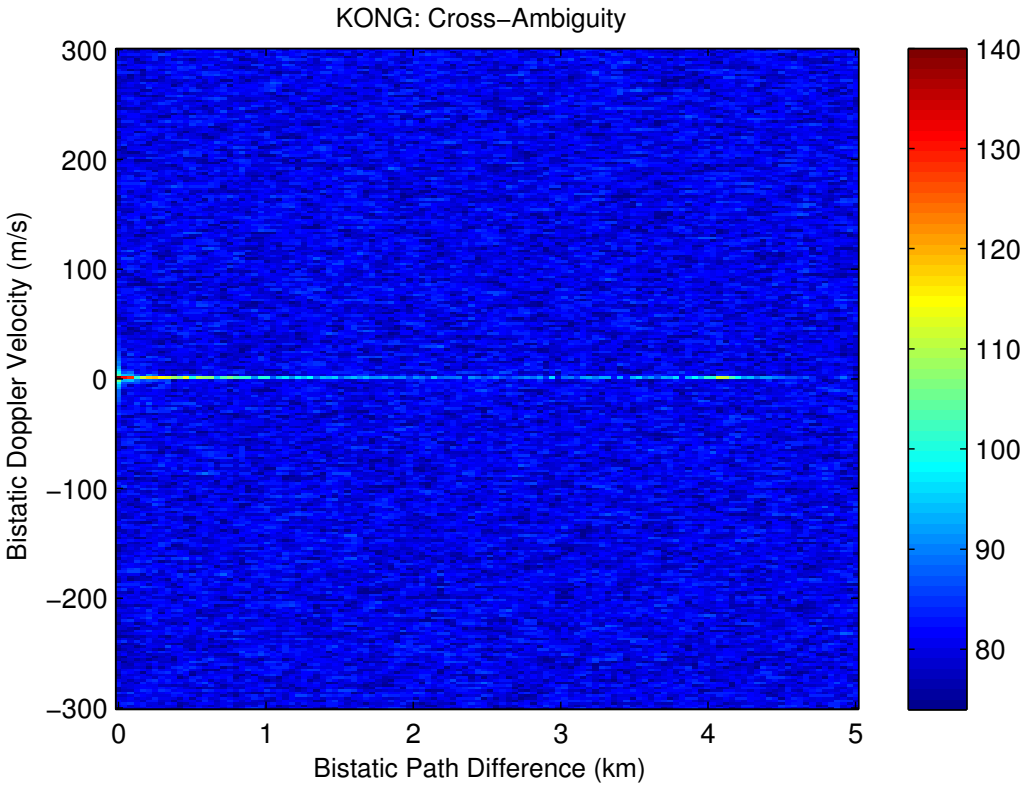

Figure 7.4: Cross-ambiguity output for stationary operation of radar using KONG with no attempt to null the direct path. We cannot detect any targets with confidence.

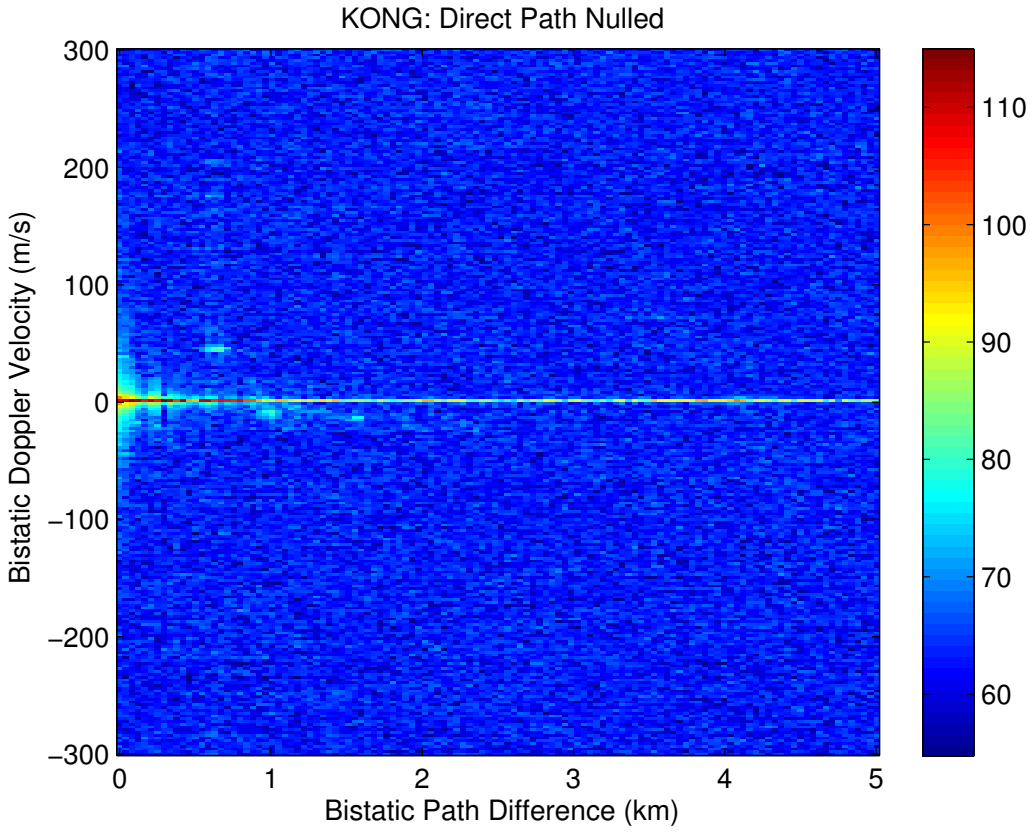

Figure 7.5: Cross-ambiguity output for stationary operation of radar using KONG with attempt to null direct path. We null the direct path at (0,0) by 32 dB, reducing the floor by 17 dB and allowing targets to be detected clearly. Three aircraft are seen here, with the most interesting being the helicopter (seen by the doppler spread blade modulation) at close range.

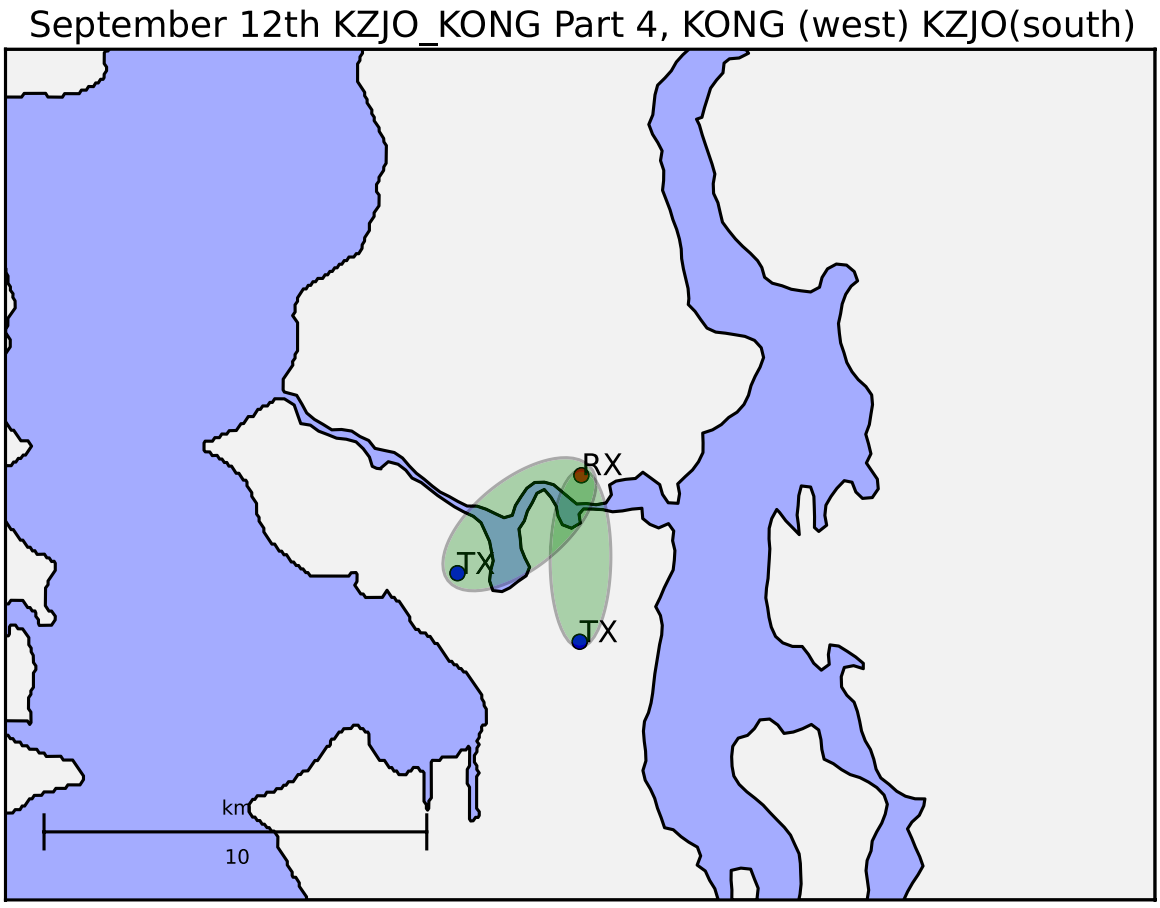

Figure 7.6: Bistatic ellipses associated with simultaneous detection of helicopter from KONG and KZJO on September 12, 2012. The radar was operating in stationary mode with receiver at Sieg Hall.

7.4 DTV-based aircraft detection: Mobile Results with two transmitters

Figures 7.7 and 7.8 show the results of the mobile operation of the radar in a vehicle driving 60 mph on the highway southbound from 10 km north of the center of Seattle. Each plot displays the doppler spreading of the ground clutter due to the motion of the car. In each plot, we see 3 distinct horizontal lines. The received direct path signal will itself have a doppler shift that reflects the speed of the car. Since our matched filter is based on the received direct path signal, we still expect the largest return to be at (0,0) and our nulling methods can remain the same. We measure doppler velocity of all other returns here relative to the direct path. If the car is traveling away from the transmitter, all stationary ground clutter behind the car will show up with the same doppler shift as the transmitter. So in fact, we still expect to see ground clutter at 0 doppler. Ground clutter ahead of the car will have measureable doppler shift that reflects the motion of the car. Far enough away, the doppler shift associated with this ground clutter will simply reflect the velocity of the car. As ground clutter transitions from near to far relative to the receiver, we expect it to trace a tangent curve on the cross-ambiguity plot, and this is visible as the arcs from 0 to 500 m range in both plots. A similar argument holds for the case of traveling toward the transmitter. Ground clutter behind you contributes a measureable doppler shift, while ground clutter in front arrives with the same doppler shift as the transmitted waveform. It appears this is the case in these figures: the car is moving toward the transmitters (toward seattle from the north) at 60 mph producing a bistatic path difference measurement of -55 m/s from the ground clutter behind us. The faint line at positive doppler shift is most likely associated with the conjugate symmetry of the cross-ambiguity function. We are seeing inherent autocorrelation from the transmitter signal. The correlation at "negative ranges" for the strong doppler shifted ground clutter contributes power in positive ranges at the equal and opposite doppler shift. While this ground clutter is understandable, it causes problems for reliable detection of moving targets in these areas of the cross-ambiguity plots without the use of adaptive filtering methods.

Figures 7.9 and 7.10 show a possible detection. This data is from stop-and-go traffic, as can be seen by the ground clutter remaining close to 0 doppler velocity. In Figure 7.9,

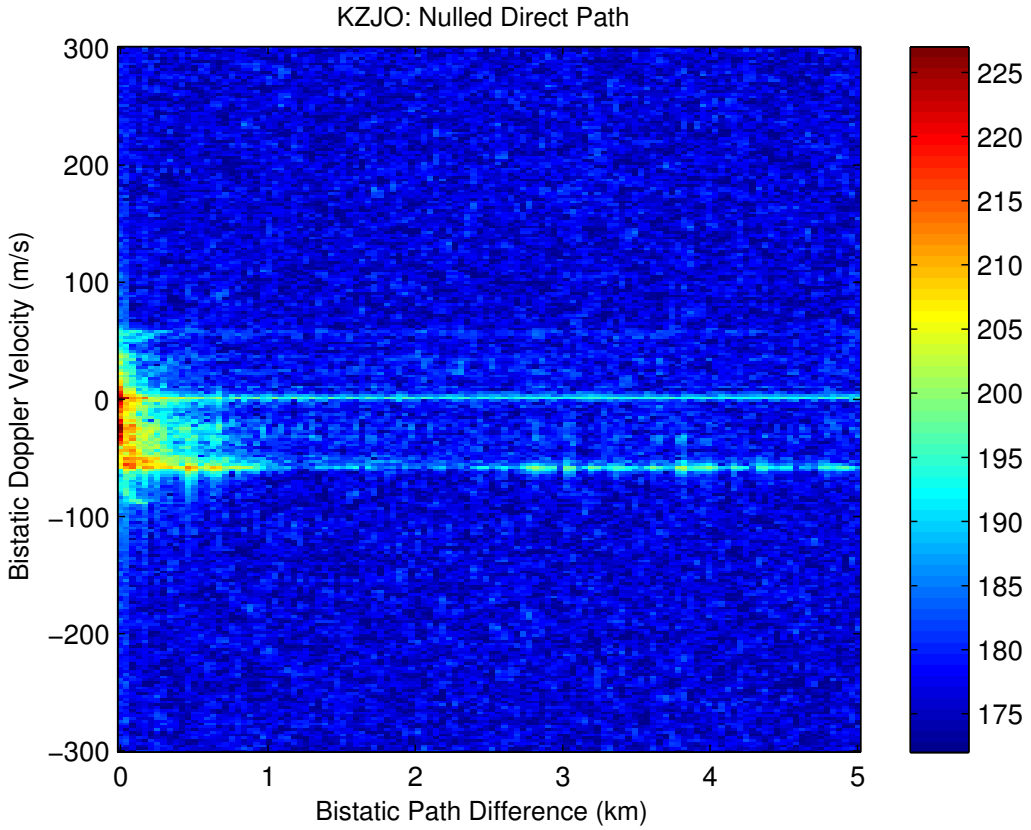

Figure 7.7: Mobile operation with KZJO, driving 50-60 mph. Doppler spread ground clutter is striking at this fast speed and makes detecting moving targets difficult without adaptive filtering.

the possible moving target (23 dB above the floor) is at range 0.9 km with negative doppler velocity -43 m/s. In Figure 7.10, it is not as clear how many targets are truly there. The strongest of the three doppler shifted "detections" (23 dB above the floor) is at range 0.15 km and negative doppler velocity -36 m/s. It is clear that better adaptive filtering, data association and tracking will clarify these preliminary findings.

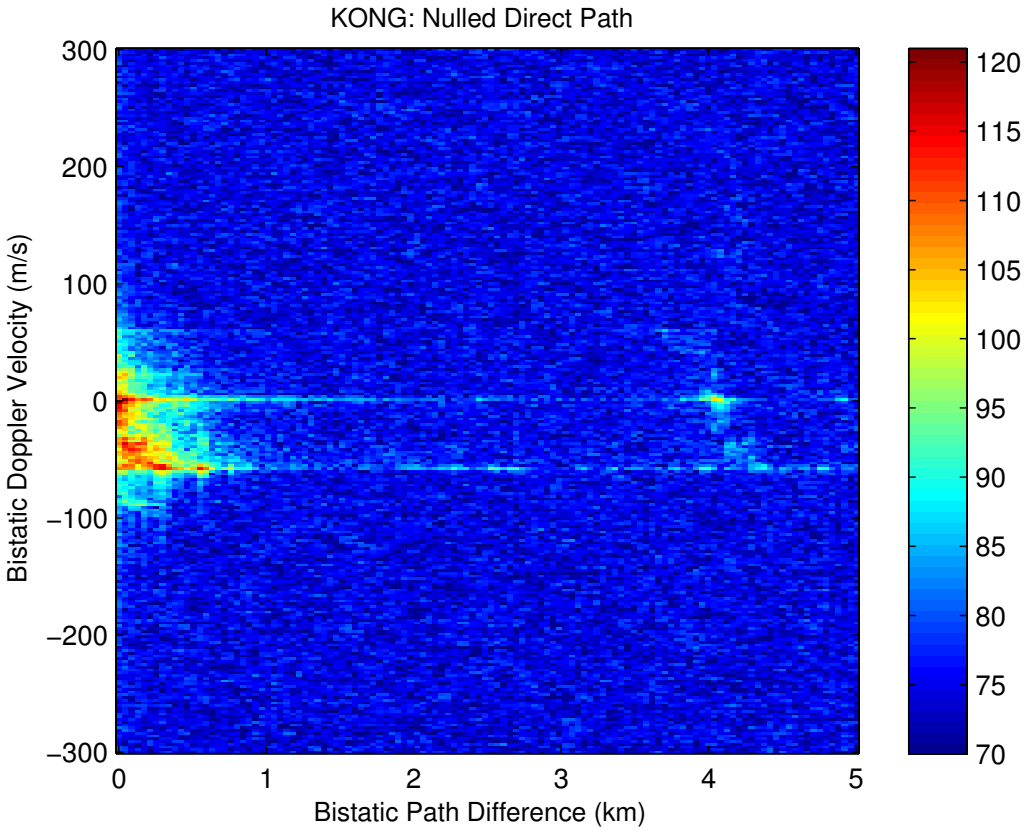

Figure 7.8: Mobile operation with KONG, driving 50-60 mph. Doppler spread ground clutter is striking at this fast speed and makes detecting moving targets difficult without adaptive filtering.

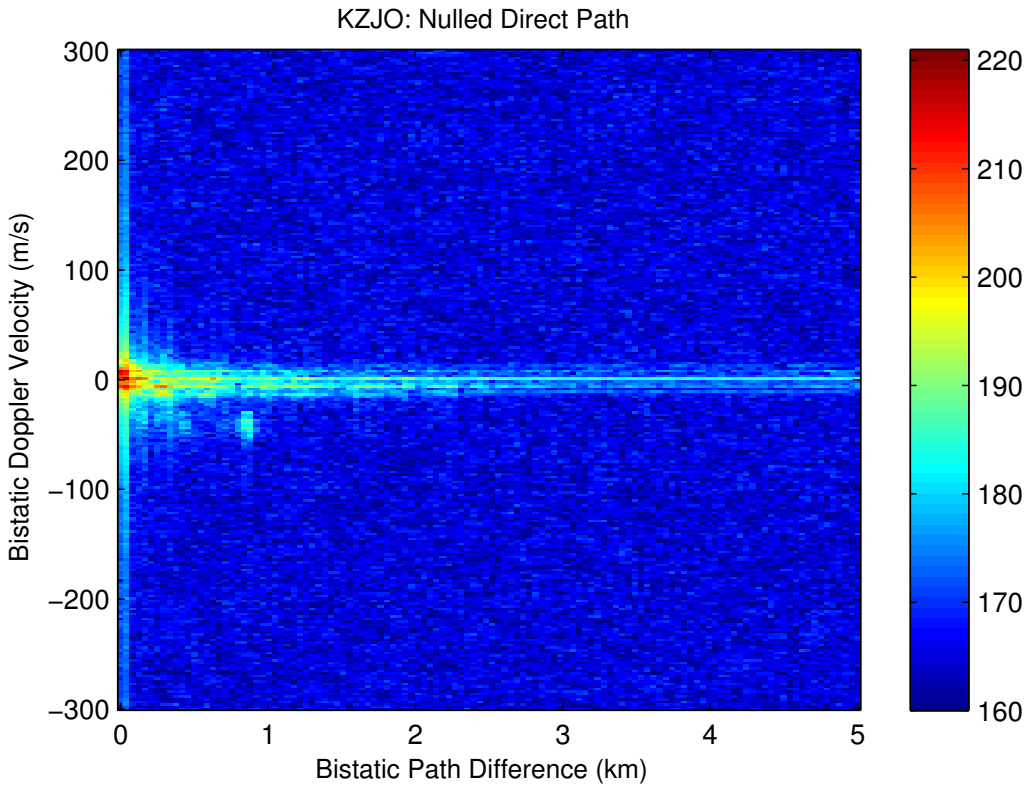

Figure 7.9: Mobile operation with KZJO, stop-and-go traffic as seen by the ground clutter remaining close to 0 doppler. We may have a possible detection of aircraft at a bistatic path difference of 0.9 m and doppler velocity -40 m/s.

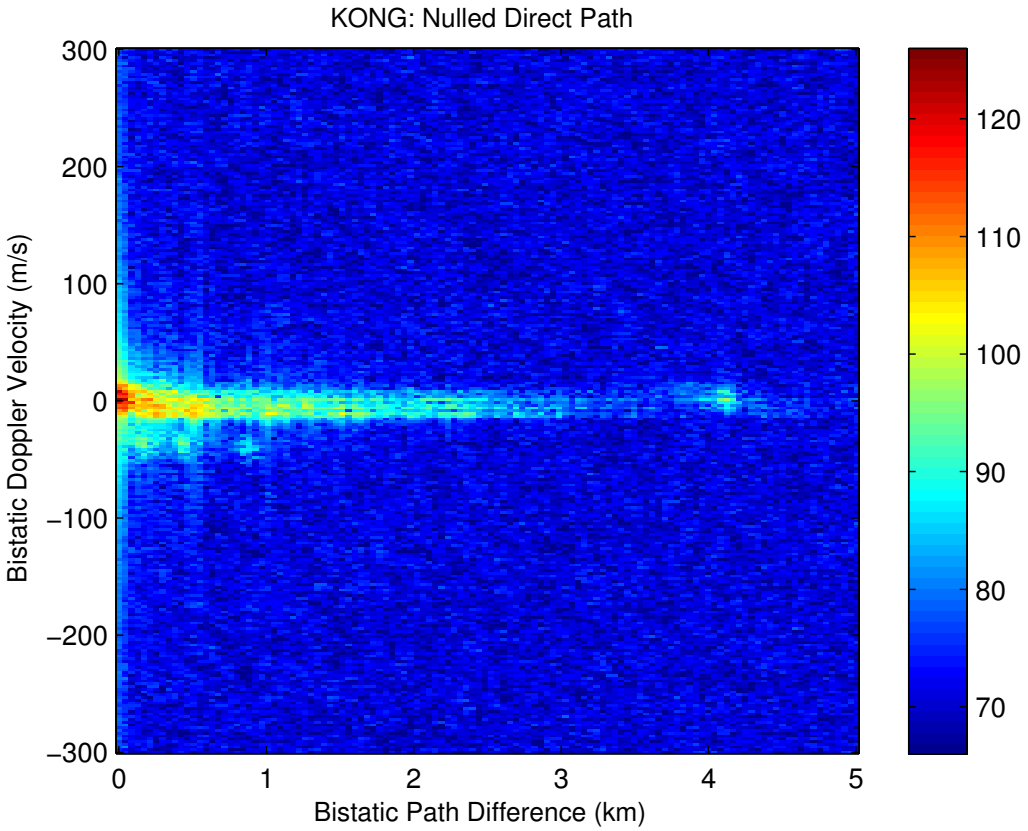

Figure 7.10: Mobile operation with KONG, stop-and-go traffic as seen by the ground clutter remaining close to 0 doppler. We may have a possible detection of aircraft with small bistatic path difference close to -40 m/s doppler velocity.

7.5 High Aggregate Bandwidth: Four antennas, Two DTV stations

Discussion of packet loss and other issues with fast data transfer are described in Chapter 6. Here we show results from real experiments on this mode. These results are fresh and should be considered preliminary; there is much for us to learn about processing this data. This mode of operation samples four antennas at 800 MSamp/s, and we save 100 MHz 8-bit IQ samples to disk. Of this 100 MHz, the clean passband is 60 MHz. The digital mixer can be set to center on several DTV stations, two locations in between DTV stations allowing 2 or 3 stations to fall within the passband, or 2 locations near FM broadcast. While 800 MHz sampling rate causes DTV and FM broadcast to alias into the same 0-400 MHz band, it is best to use band-pass filters on the front end to suppress FM and look for DTV stations. A block diagram of this processing is shown in Figure 5.5 from Chapter 5. For this particular experiment, we would like to observe aircraft on two DTV transmitters on all four antennas. After programming the radar and setting the front end attenuation, we can set the mixer frequency to 557 MHz, directly between KZJO (536-542 MHz) and KONG (572-578 MHz), the two stations used in the results presented earlier. After setting the digital gain, the data being transfered off the FPGA is shown in Figure 7.11 using our `diagnose.py` program. The 60 MHz passband is shown here, but the data is transfered off the FPGA, 8 bits I and Q on 4 antennas at 100 MHz, a total of 6.4 Gbps. These spectra are calculated using 0.1 seconds of data, and this data is saved to disk, so may contain discontinuity where packets have been lost (more on packet loss and high aggregate bandwidth in Chapter 6). It is clear here that the DTV spectrum in Seattle is sparse. In 60 MHz here, only 2 DTV stations are clear, thus we compute cross-correlations for only these two. However, this mode would allow us to process up to 10 stations were they all present and transmitting. With the radar mobile, we may very well encounter a situation like this.

With high aggregate bandwidth off the FPGA, packet loss will occur. This particular design has header information inserted into each packet, allowing us to precisely identify periods of missing data. Recall from Chapter 6, if we process the data with no corrections, we expect to see range and doppler migration toward the upper left (closer in range, positive in doppler). This data was taken on June 17th, at 1:33 pm. Processing the first second of

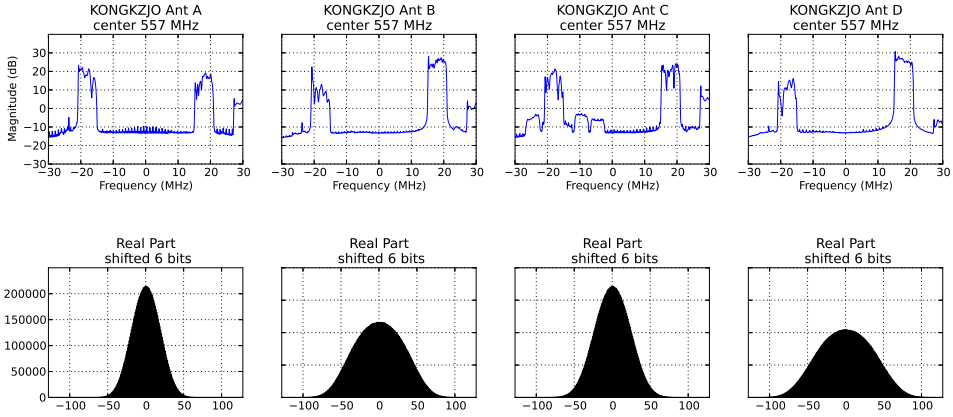

Figure 7.11: Diagnosis of FPGA output for an aircraft detection experiment with fast rate data transfer. The FPGA mixer is set to 557 MHz, directly between two DTV stations: KZJO (536-542) and KONG (572-578). The passband is shown from -30 MHz to 30 MHz.

data without taking into account dropped packets, we can compute the cross-correlations between all four antennas. These results are shown in Figure 7.12. A target is visible with measured Doppler velocity 36 ±2 m/s, at range of 170 ±25 meters. We expect that the aircraft may have experienced some range migration, however it appears to be a concise peak in range. I would expect this to be the accurate range measurement. Processing for the next second in time is shown in Figure 7.13. We can see the target's range reduce as predicted by the Doppler velocity measurement in Figure 7.12.

Similarly, we can process this detection for KONG for both seconds of data. These results are shown in Figures 7.14 and 7.15. We make the claim that the aircraft is represented in the plot as the "blob" in the lower left part of the plot at range <50 m with Doppler velocity -17 m/s. First of all, moving from second one to second two, this reflection moves correctly for a physical target. Recall that the range-Doppler plot shows the range and range rate. Therefore if we look at two consecutive range-Doppler plots, the first plot's Doppler velocity and range can predict the range location on the following plot. In that sense, all physical targets that show up with negative Doppler velocity, are moving away from the radar, and thus should predictably move to the right of the figure.

We can call the blob reflected above the range = 0 axis to positive Doppler shift a ghost

detection, for lack of a better term. This moves unphysically from second one to second two as it appears to move with positive Doppler velocity further away from the radar. In fact, the aircraft detection and it's ghost detection are linked and will always be opposite each other. The explanation for this feature is complex. First of all, it is important to note that we only plot positive range values. If the target detection has extent in range, it may extend into the negative range realm we do not see. The mathematics of the cross-ambiguity calculation guarantee the plot to be 180 degrees symmetric.

$$|\chi(\tau,\nu)|^2 = |\chi(-\tau,-\nu)|^2$$

With this being true, any negative range extent associated with our target will show up in the positive range axis with opposite Doppler shift. This explains the present of that ghost detection, and explains its relative motion to the true target detection.

All of this presumes that the target has extent in negative range. An explanation for why a target would show up multiple times and would appear to show up *earlier* in time (which leads to a measurement earlier in range) is multipath in the matched filter signal. Our matched filter is not perfect and in fact with the case of KZJO, we know that there is a large pine tree that suppresses much of the direct path. This may bring the direct path power down enough so that a secondary reflection from a large building or something is comparable. Without cleaning the matched filter, we expect to see these sorts of features, and they can even be present in the negative range realm, which is more mathematical than physically meaningful.

This argument is not conclusive, but plausible. Much more data is needed to confirm these effects, and confirm the usefulness of the fast data rate mode. While this isn't a confirmed aircraft detection, it behaves in a manner consistent with aircraft detection with features we can explain.

Unfortunately with a detection so small in Doppler, when applying the simple zero-stuffing packet loss correction described in Chapter 6, the elevated Doppler line due to the rectangular windowing nature of our correction overpowers the target strength, making detection nearly impossible.

Packet loss for second one and second two are shown in Figures 7.16 and 7.17. Recall that packet loss is fortunately identical for all antennas, and thus all stations. As we can see, the data loss is fairly small for almost the entire second. In total there is roughly 30% data loss and from our simulations of packet loss on a point target, we could expect that the true Doppler velocity of the target is 30% below the measured value in Figure 7.12 and 7.14. However, most of the data loss occurs at the end of this second and the true Doppler measurement may be quite a bit closer to our estimate. The packet loss indicated in Figure 7.17 is much worse for the second two. Overall packet loss is quoted around 60%, with bursts around 10,000-15,000, with a maximum over 25,000. From Figure 7.12 it is clear that not much distortion occured. It is not as clear from Figures 7.13 and 7.15 that a lot of distortion occured. As we know from our simulations, Doppler velocity bias is a function of actual Doppler velocity. With such a slow target (radially), even up to 60% distortion does not change our estimate drastically.

Inserting a packet of zeros where data is lost, we can process the data aligned correctly in time. Figure 7.18 shows the first second cross-correlation from KONG. We can compare this to Figure 7.12. The target is still very clear, indicating the effects of packet loss are small. In fact, this correction should give us the true velocity and range, which we measure to be 32 m/s at 170 m. Not much range migration was expected for the packet loss seen in the first second, but the Doppler velocity change from 36 m/s to 32 m/s matches simulation results with roughly 12% data loss, rather than the 30%. This may mean that the rather consistent packet loss occuring for the major portion of the first second was probably around 12%, and the one second mark fell in the middle of many zeros with little data tacked on to the end. The second part of the data is processed and shown in Figure 7.19. The Doppler power spread associated with our crude, zero-stuffing, correction scheme completely masks the target here since this second experienced much more data loss. In this case, processing the data without correcting for packet loss actually produced a result that, while inaccurate in range and Doppler, had a detectable target.

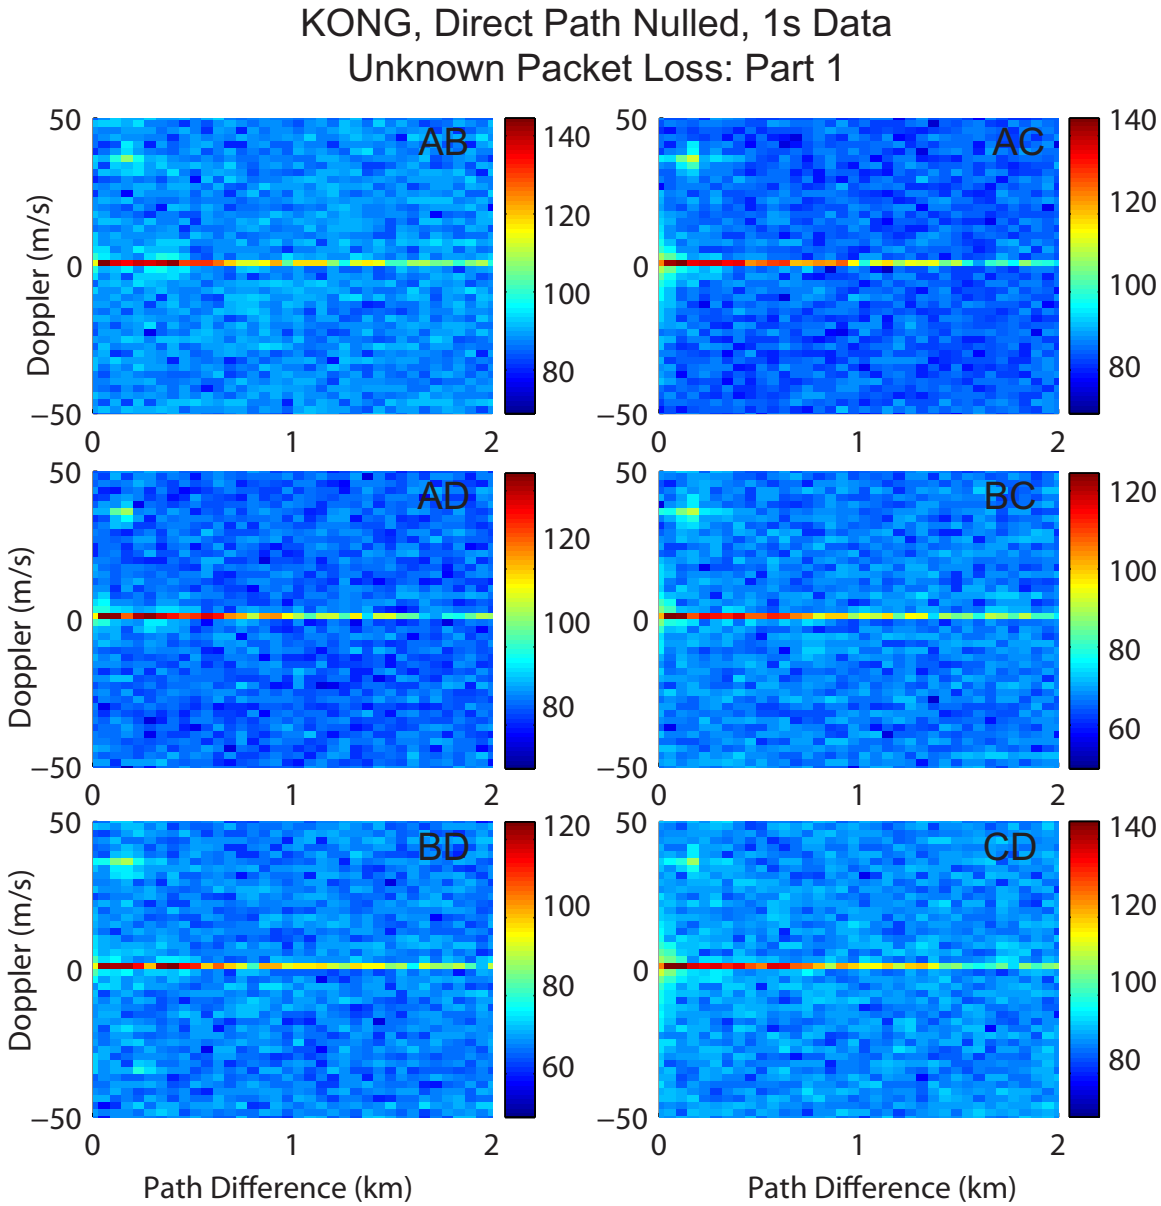

Figure 7.12: Cross-Correlations for each of four antennas for KONG using the fast rate data transfer mode, processed without taking into account packet loss. A target is visible with measured Doppler velocity around 36 m/s, at range of 170 meters. We expect that the aircraft may have experienced some range migration, however it appears to be a concise peak in range, so this may indicate minimal losses.

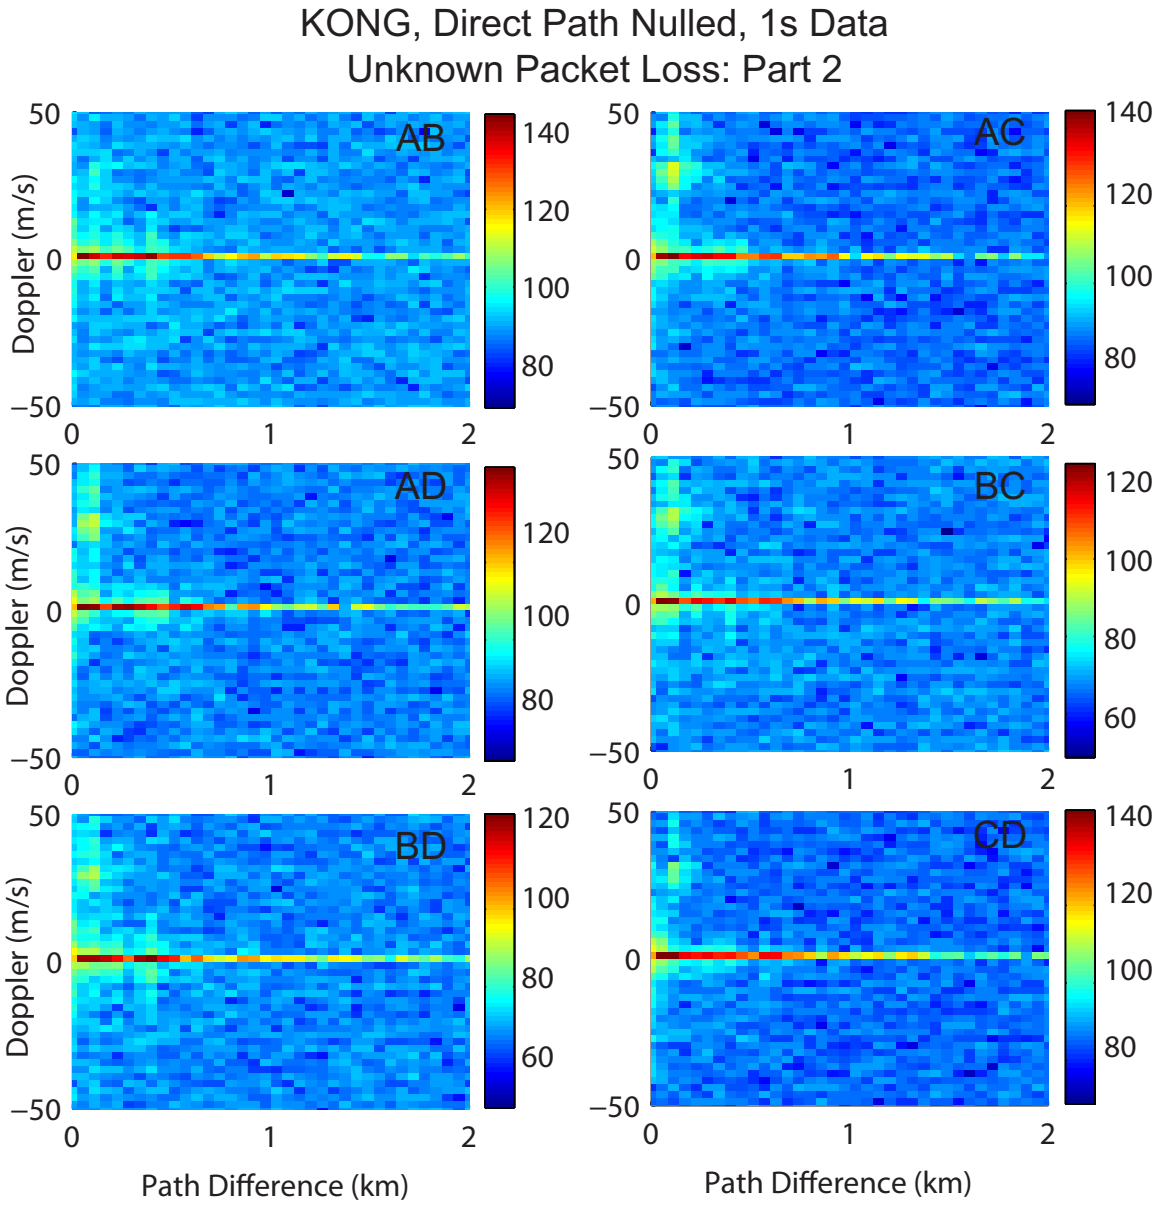

Figure 7.13: This shows the next second of data associated with Figure 7.12. The bistatic path difference shrinks as expected from the positive Doppler velocity in the previous plot.

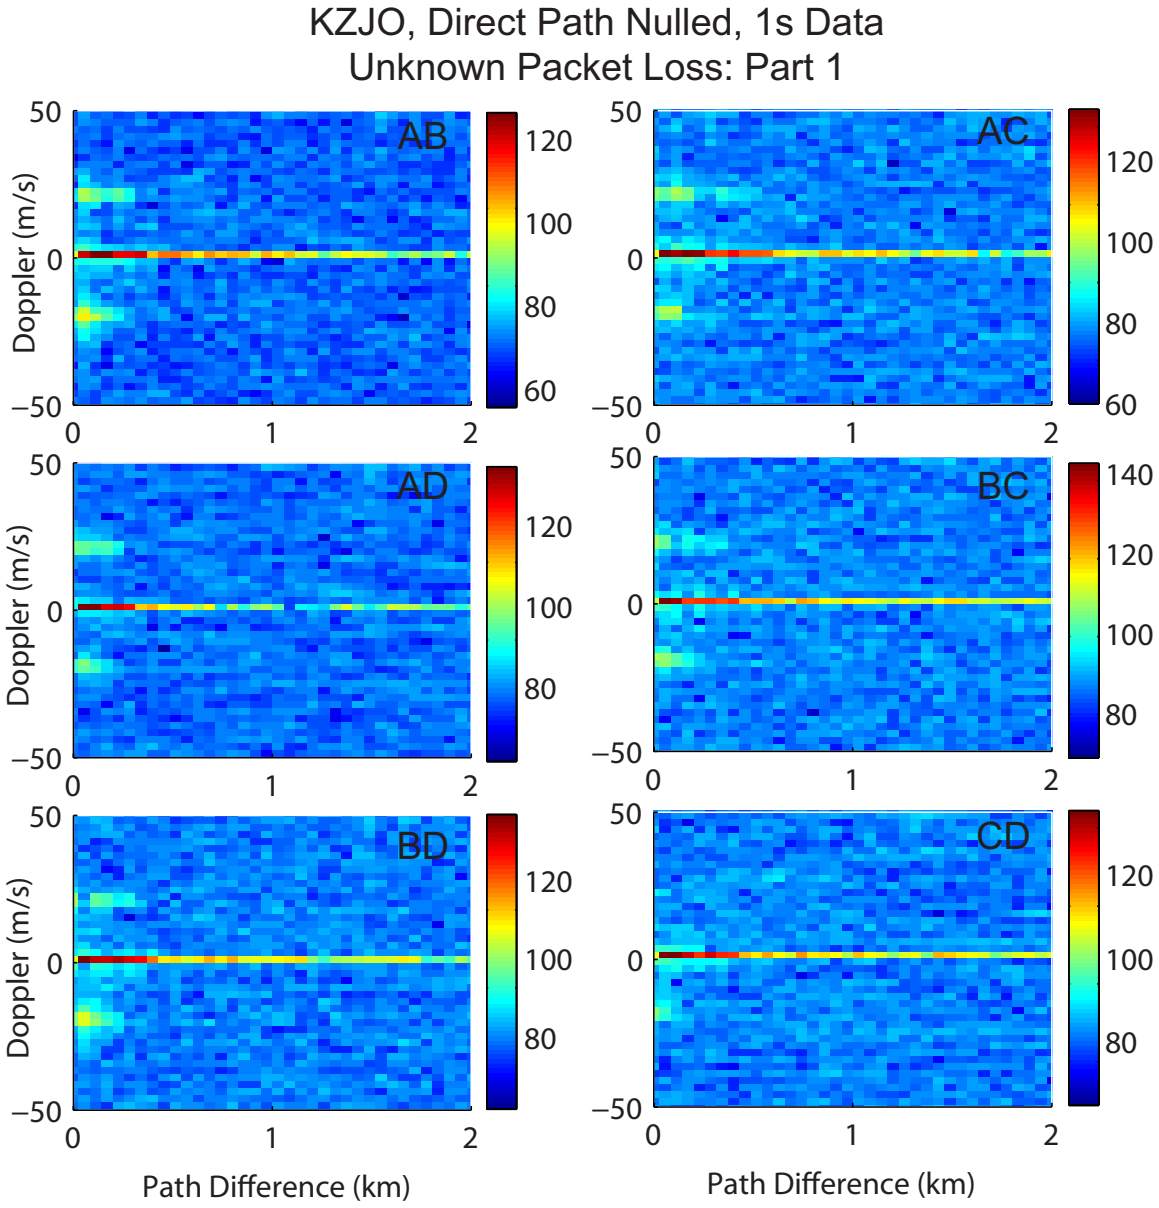

Figure 7.14: Cross-Correlations for each of four antennas for KZJO using the fast rate data transfer mode, processed without taking into account packet loss.

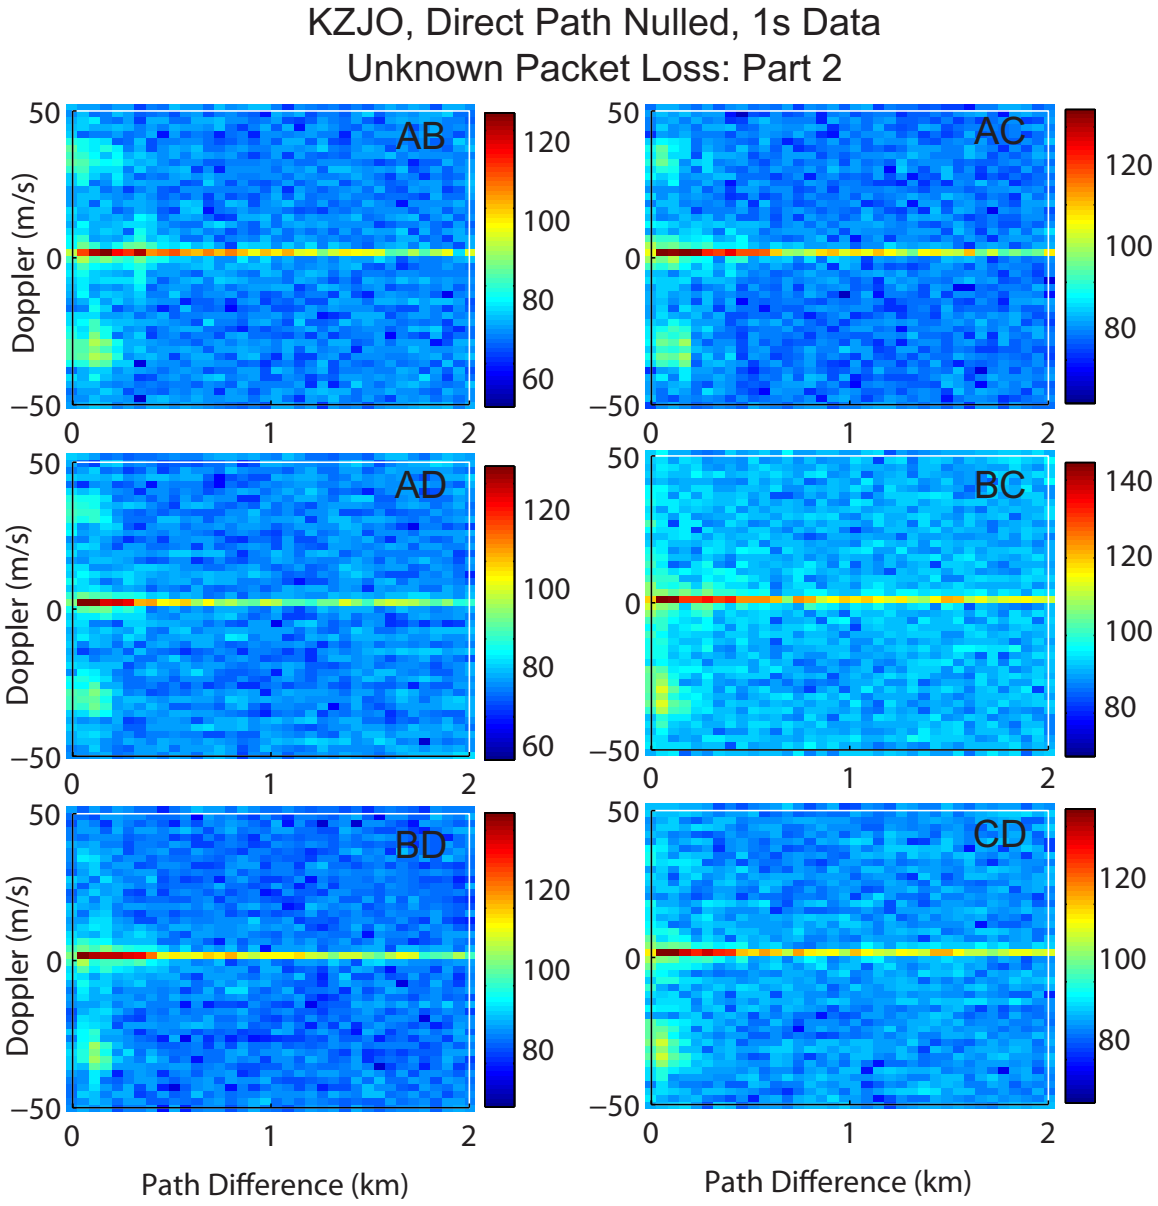

Figure 7.15: This shows the next second of data associated with Figure 7.12. The target advances in range by roughly the amount associated with the Doppler velocity estimate from the first second.

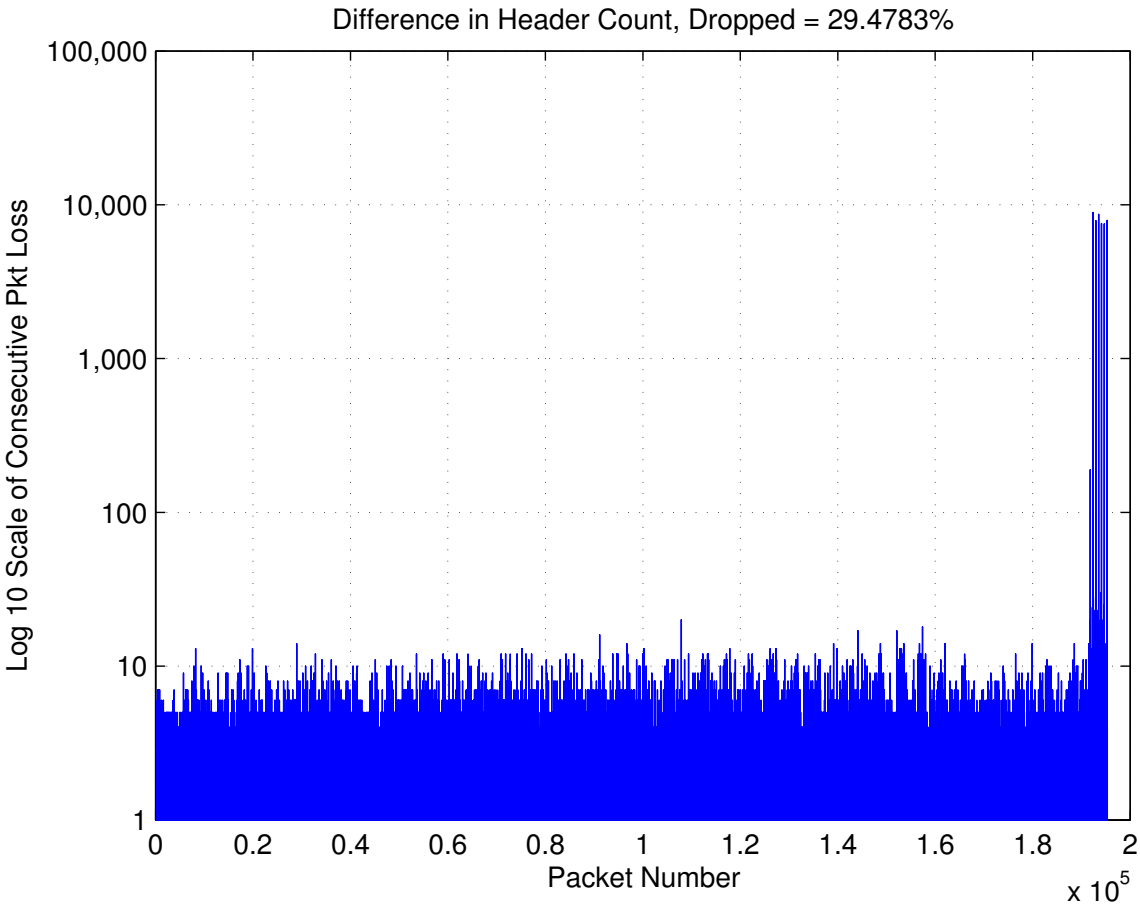

Figure 7.16: Packet Loss for the first second of fast rate data on a log 10 scale. Fairly minimal packet loss for most of the duration.

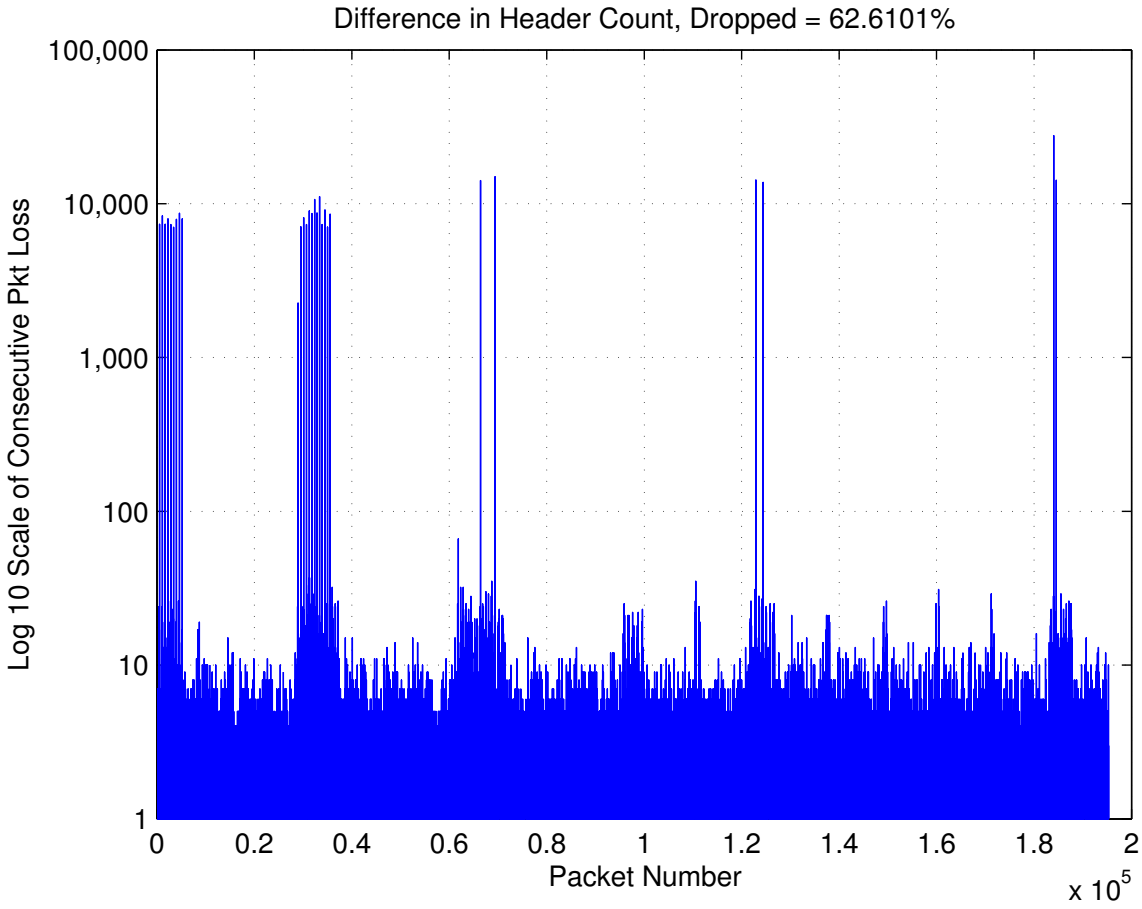

Figure 7.17: Packet Loss for the second second of fast rate data on a log 10 scale. Many instances of bursty packet loss indicate that the measured range and Doppler velocity in Figures 7.13 and 7.15 could be inaccurate.

We now present the equivalent cross-correlations for KZJO in Figures 7.20 and 7.21. As seen with KONG, the first second of data shows little sign of packet loss. The same range is measured, however this plot indicates the true Doppler velocity is around -15 m/s, rather than -17 m/s as seen before. These numbers produce the same packet loss distortion that we saw with KONG: the data experienced about 12% packet loss. Figure 7.21 shows again that the packet loss is so drastic that our close target is lost in the processing side effects.

Reliable detection has been shown for standard data rates, allowing us to see aircraft on four antennas from a single DTV station. The fast rate data transfer mode allows for up to 10 times the aggregate bandwidth, up to 10 DTV stations on four antennas, with the trade-off of packet loss. Care is needed when analyzing this data, and these preliminary results are promising but much more work is left for future students to fully utilize this mode of operation.

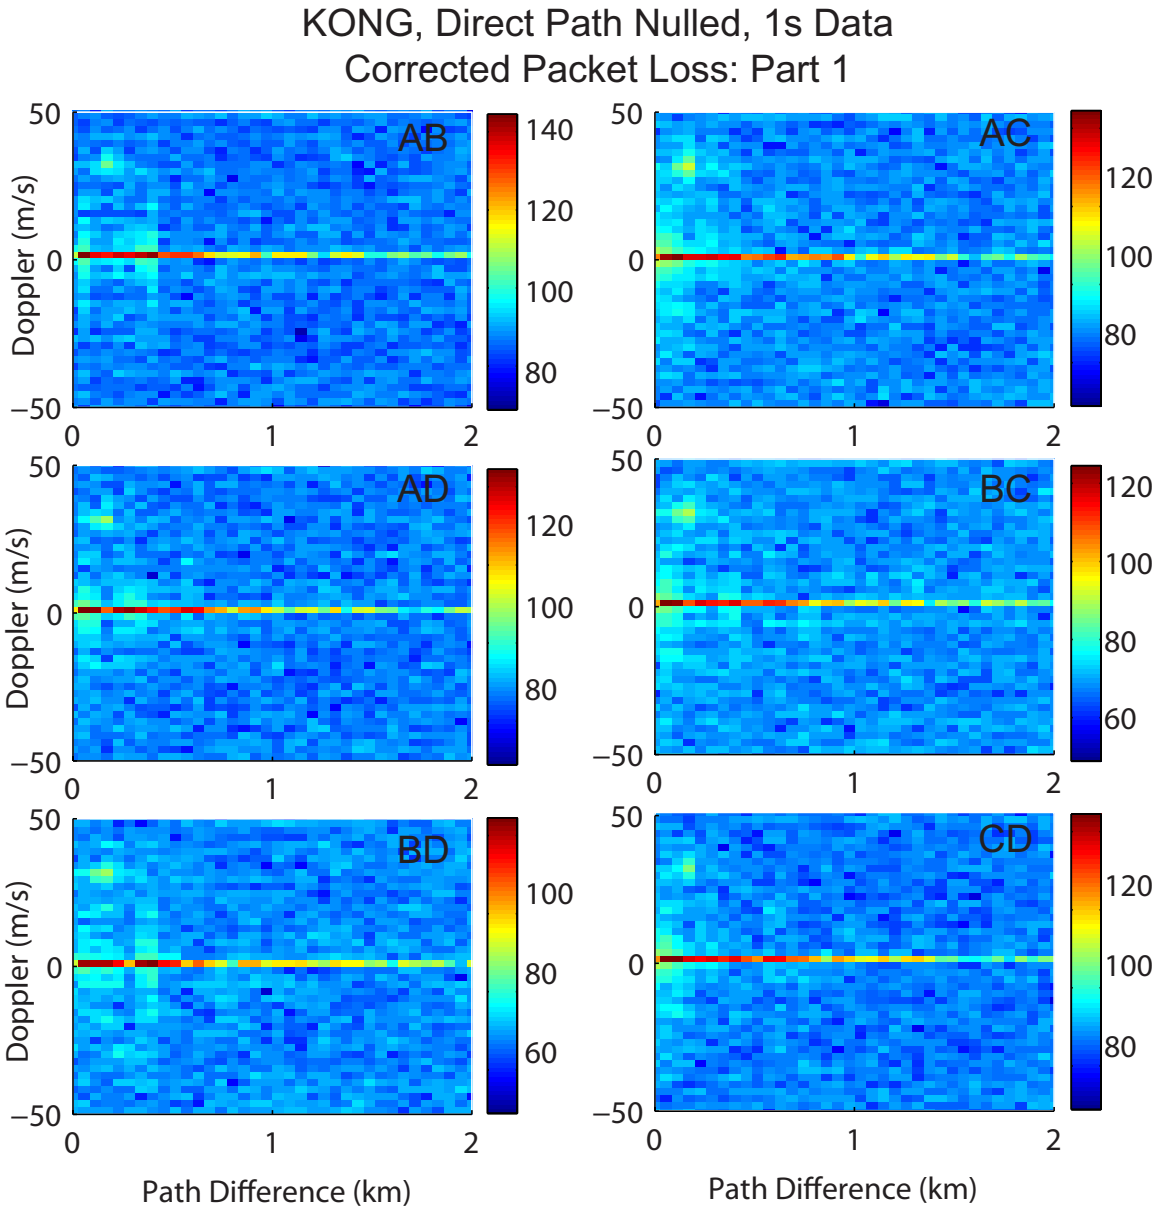

Figure 7.18: First second of KONG data, with packet loss corrected by inserting zeros. Minimal packet loss makes these results very similar to Figure 7.12, showing the true range and Doppler velocity of the target is 170 m at 32 m/s. The difference in the true Doppler velocity and the Doppler velocity measured previously is consistent with 12% packet loss for most of the processing.

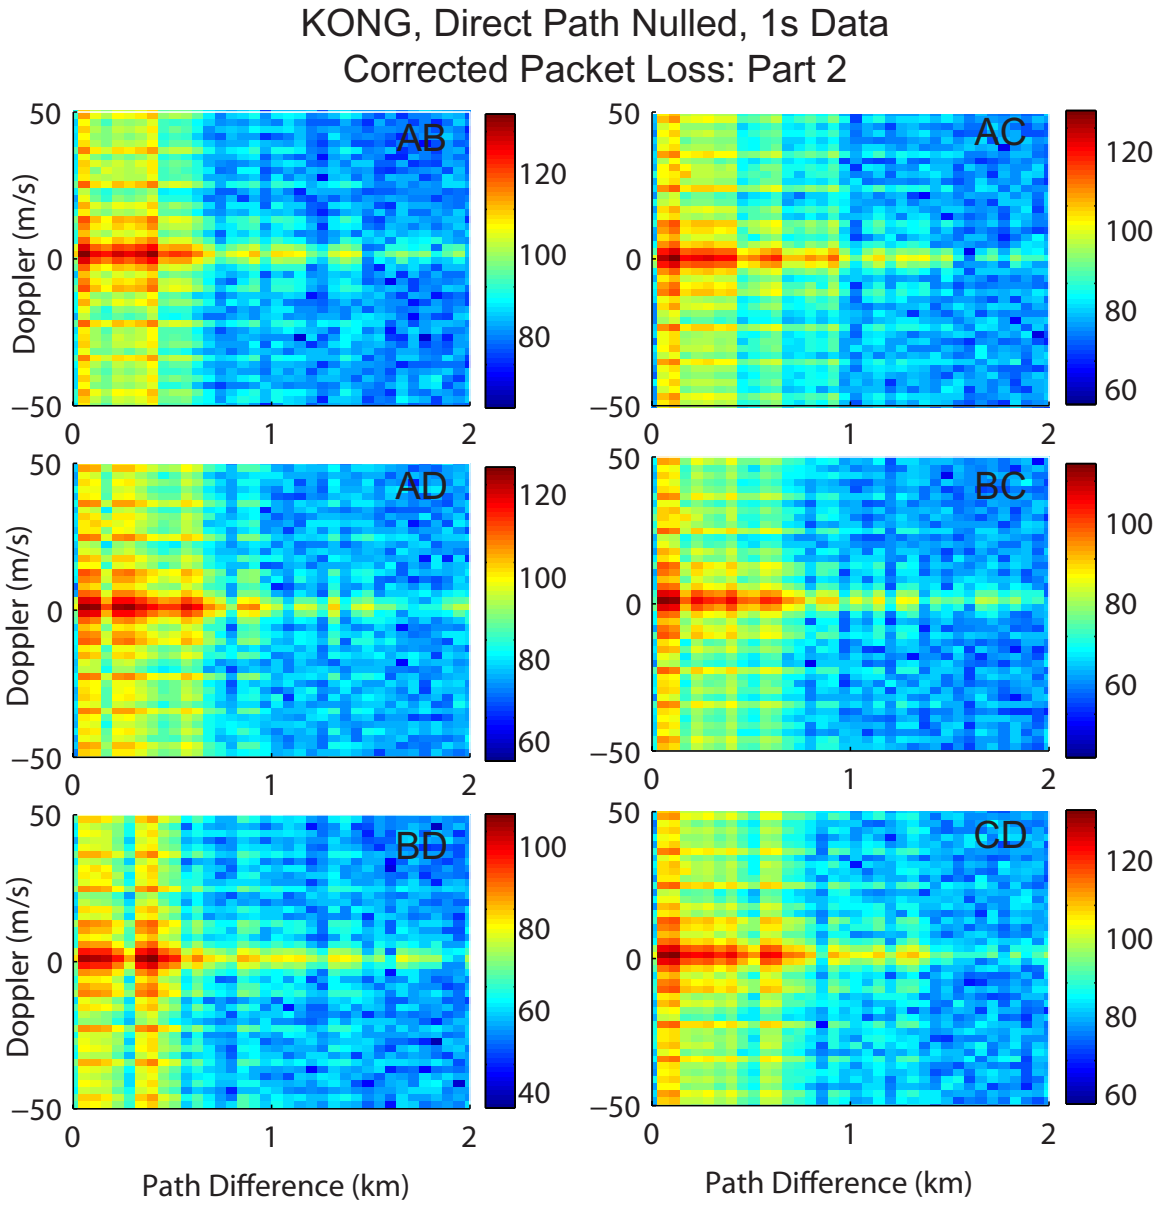

Figure 7.19: Part two of KONG data, with large amounts of packet loss corrected with zeros. The target is completely masked by the effects of the correction scheme.

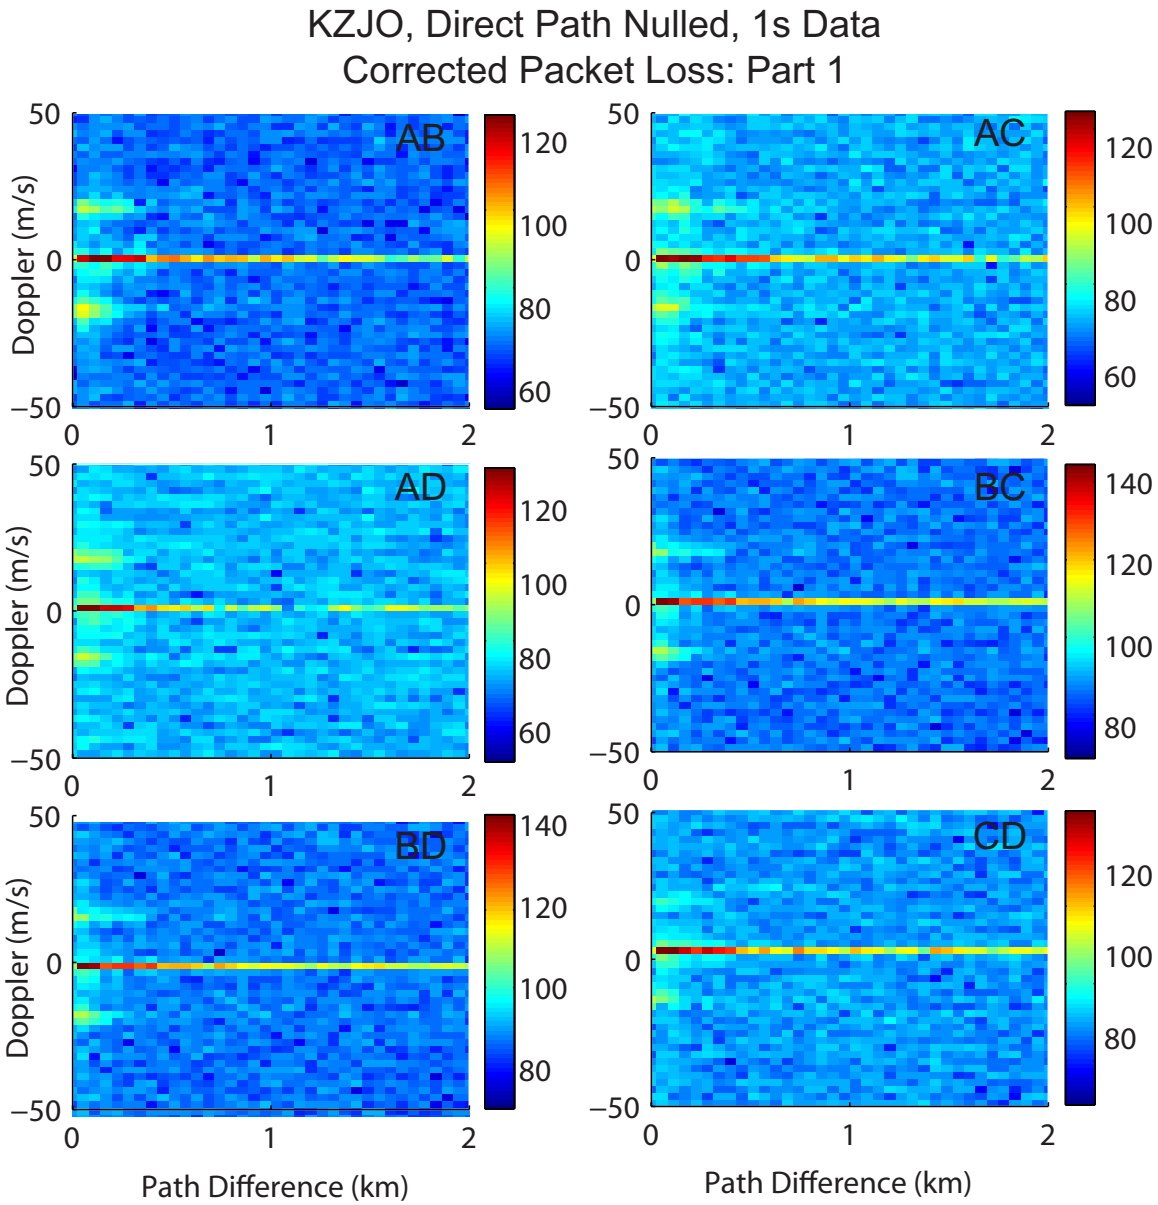

Figure 7.20: First second of KZJO data, with packet loss corrected by inserting zeros. Minimal packet loss makes these results very similar to Figure 7.14, showing the true range and Doppler velocity of the target is ¡75 m at -15 m/s. The difference in the true Doppler velocity and the Doppler velocity measured previously is consistent with 12% packet loss for most of the processing.

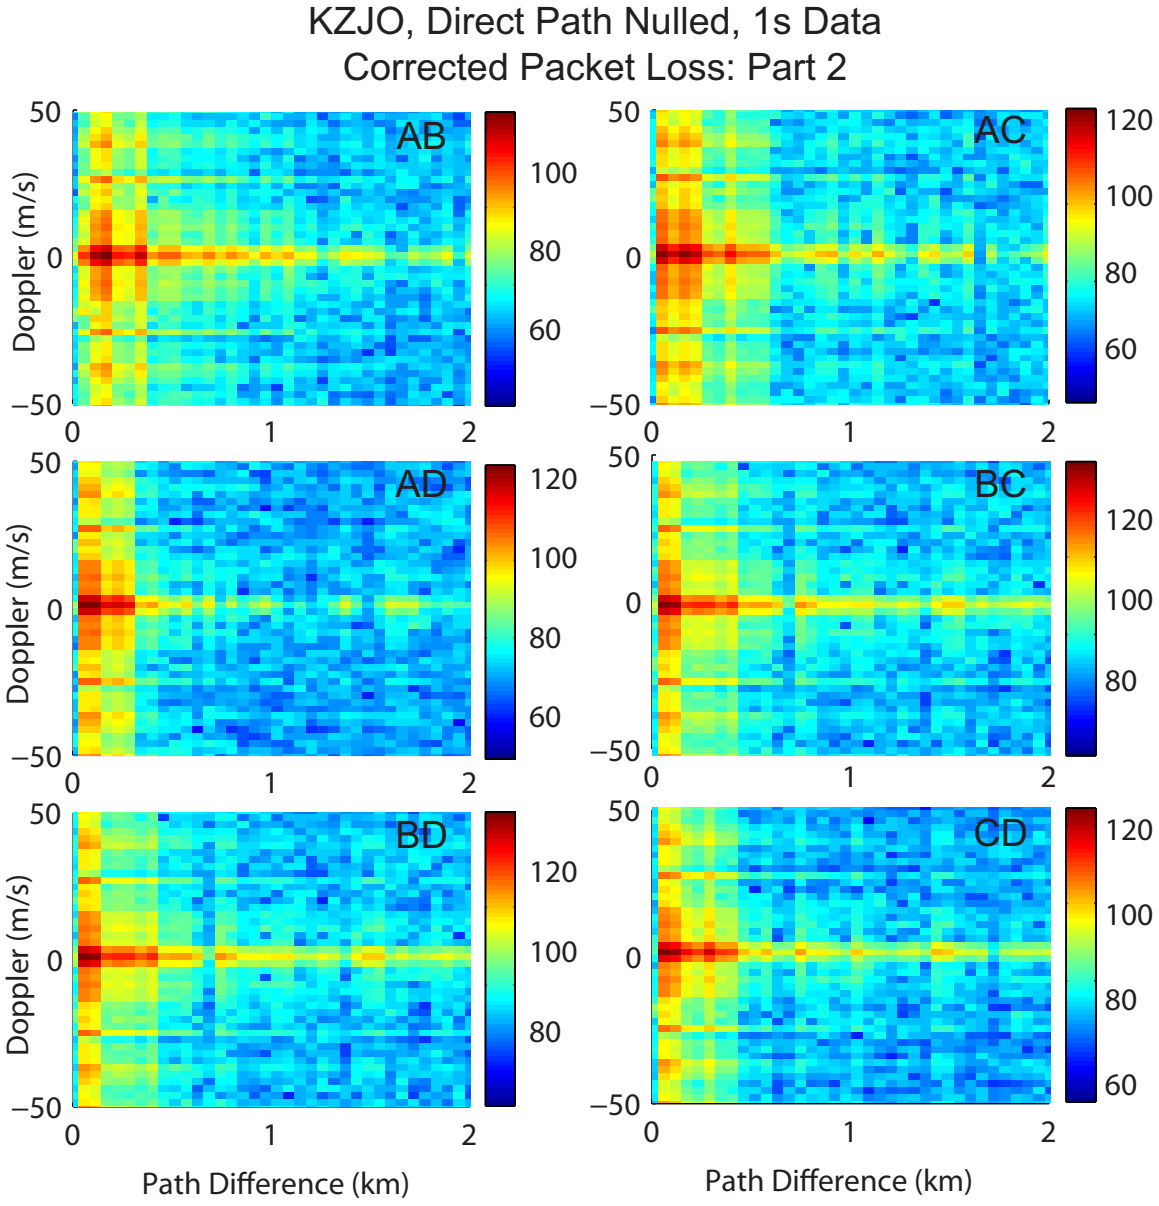

Figure 7.21: Part two of KZJO data, with large amounts of packet loss corrected with zeros. The target is completely masked by the effects of the correction scheme.

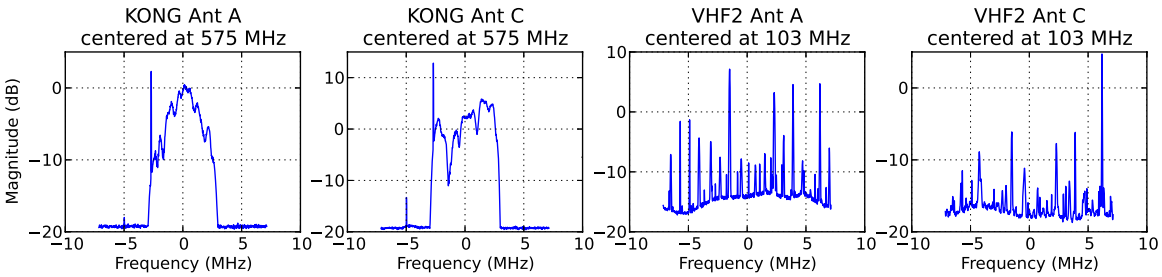

Figure 7.22: FFT of data saved to disk for two antennas, each digitized at 1.6 GSamp/s and split into two streams that isolate 6 MHz of the spectrum. This shows we are observing both FM and DTV from a single antenna input.

7.6 DTV and FM Broadcast on the Same Antenna

We use the FPGA mode for two antennas observing two different 6 MHz portions of the spectrum shown in Figure 5.6. While we have not yet detected aircraft in this mode on both DTV and FM broadcast, we can present the spectrum of the data streams, showing the quality of data we get from the standard rate data transfer. Figure 7.22 shows data that has been digitized from two antennas, where each digitized stream is split on the FPGA to listen to 6 MHz of FM broadcast and 6 MHz of DTV broadcast.

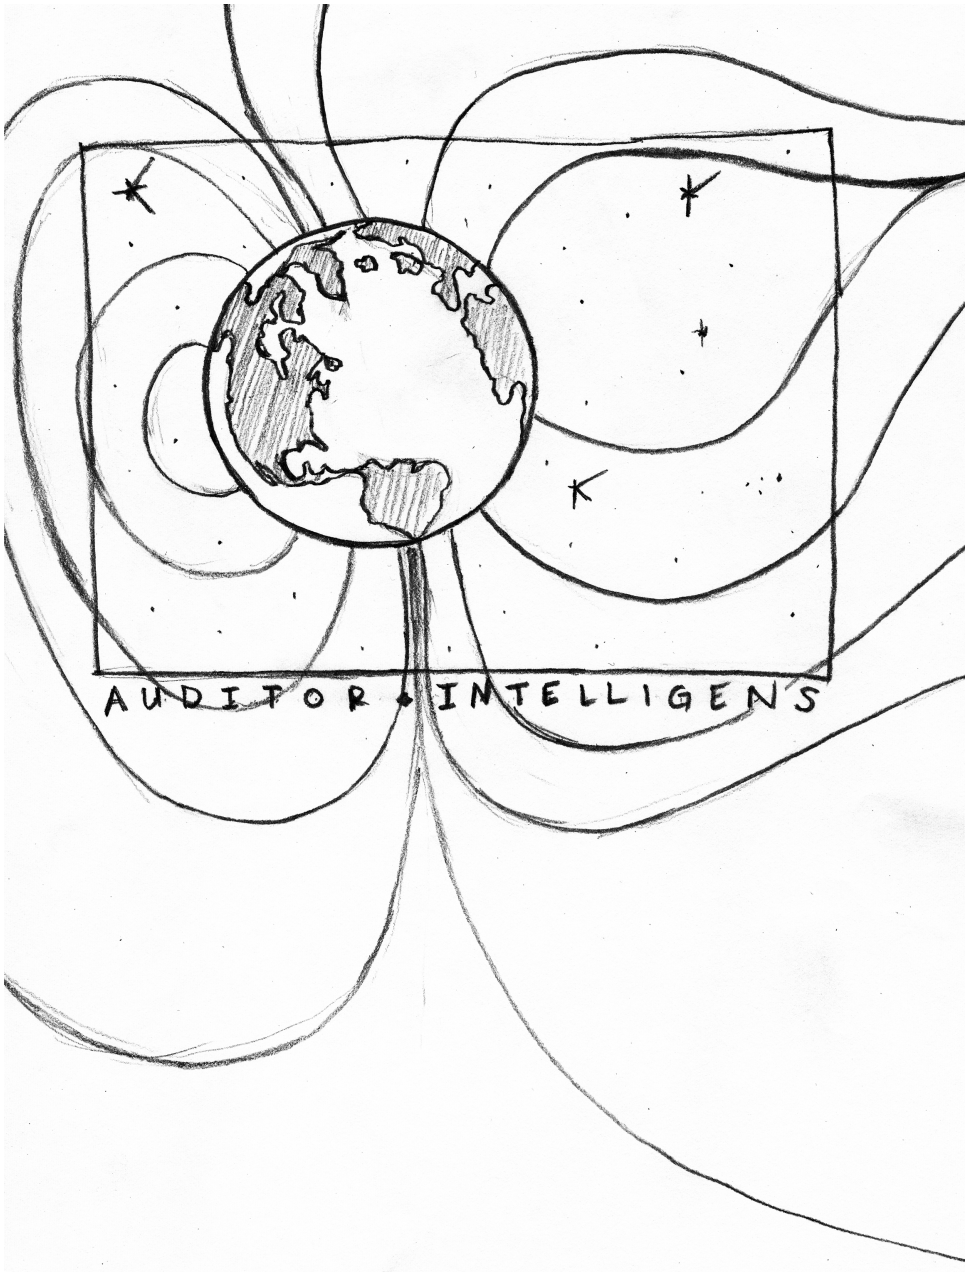
AUDITOR • INTELLIGENS

Chapter 8

FUTURE EXPERIMENTATION WITH THIS SYSTEM

In this brief chapter we explore ideas for upcoming research with this new system.

8.1 Hybrid Transmitter Detection of Aircraft

Multiband or hybrid systems have been developed ([52]) and proposed, [31], noting the use of FM an initial track-and-detect, then moving on to DAB and DVB-T for finer resolution processing. While bandwidth of DTV in the United States and DVB-T are comparable (with different modulation shemes), it does not suffer from the single-frequency transmission structure of DVB-T operation in particular geographies. The proposed system in [31] also intends to have independent antennas and front end components. This will allow them to take advantage of antenna gain patterns with small sidelobes. Our system would rely on a single broadband antenna to detect both FM and DTV, offering a technology exploration that can be compared when both systems are complete. We have thus far demonstrated the ability to observe both FM and DTV on a single antenna for 2 antennas, but have yet to achieve the necessary processing to make detections, combine the independent measurements across all transmitters, and track targets. The first step would be detecting an aircraft simultaneously with DTV and FM transmission.

8.2 Testing New Coherent Processing Strategies

With single-site passive radar, the power from the direct path and the indirect path from a scatterer are received simultaneously. Not only that, but with 100% duty cycle radars, the targets are constantly reflecting power. This means that detectability of a particular target of interest is competing against all scattered power in the system. While the target reflection will integrate coherently at the correct range and Doppler transformation of the matched filter, all other power will add incoherently. The better we can subtract this unwanted

clutter the further we can drive the cross-ambiguity floor down, thus making weaker targets more detectable. Developing further Space Time Adaptive Filtering (STAP) techniques, rather than the very simplistic transmitter nulling algorithm currently used, will allow us to better remove ground clutter.

Another challenge with this type of radar is the imperfect matched filter. Better understanding of the DTV modulation scheme [8] will allow us to decode the signal and recreate a clean matched filter.

8.3 Auroral Observations

DTV-illuminated aircraft detection was an excellent problem to solve that allowed us to apply our direct sampling strategy to new areas. Auroral turbulence has historically been the RRSL's main research interest. Space weather is linked to solar activity, which falls naturally into an 11-year cycle. In fact, upon beginning the PhD program and planning my course load and research track, things were lined up nicely for me to take classes during the solar minimum and then ramp up the research and graduate in 2013 in the middle of solar maximum. Nature had different plans. In 2008 and 2009, NASA Headline News reported on the deepest solar minimum in almost a century ([44], [42]). 2008 was a record-breaking year in terms of minimal solar activity. It is being referred to ([43]) as the "blankest year of the space age." Not only was it weak, but the solar minimum lasted an extra 15 months longer than expected, leading scientists to refer to this as an extended solar minimum, and are looking to explain it by variations in the plasma flow from the sun's equator to the poles [24]. The National Center for Atmospheric Research (NCAR) predicted in 2006 that the next solar maximum would begin a year late and that it would be the strongest in half a century [12]. While their particular model was correct about the delay, scientists have discovered this current maximum, peaking now in 2013, is turning out to be the smallest sunspot cycle in over a century, as mentioned in a report by Dr. David H. Hathaway [25].

It was a convenient cycle for our lab to develop new techniques and advance our own instrumentation. For the last few years of this solar maximum, we can develop a mode that simply isolates the entire FM band, allowing us to probe the ionosphere on four antennas with every transmitting FM station. Previously with the ISIS radar, we could observe 4

FM stations on 6 antennas. While the new system is more limited in number of antennas, we have redundancy in the spectrum. To create a mode to directly compare with past operation, we may need to build an identical system and place it at our remote observatory. The potential exists, however, to explore single-site detection of auroral events.

Having now developed a mobile device, the opportunity exists to bring our equipment over the border to Canada and take data at a location better suited for auroral turbulence observations.

With a future mode that reliably processes both DTV and FM data, we will be able to probe the ionosphere with a wide frequency range of transmitters, giving us a brand new type of measurement to analyze.

8.4 Considerations for Wi-Fi or GPS transmitter applications

While this system was built with the initial goal of aircraft detection with DTV, the fast sampler and high aggregate bandwidth allow us to explore other applications that arise from isolating different transmitters.

Passive Wi-Fi systems have been developed for tracking people in buildings. To adapt our system for Wi-Fi (centered at 2.4 GHz) will take some care. If we sample at 3 GHz, we will alias Wi-Fi down to 600 MHz, where we know DTV stations may exist. If the spectrum happens to be empty there, it may be possible to operate without the use of a band-pass filter other than the roll-off of the antennas. Sampling at 1.6 GHz results in aliasing at 800 MHz, where there will be complications with cellphone transmission. An 800 MHz sampling rate will actually not work at all for a Wi-Fi channel centered on 2.4 GHz as it will alias onto itself at baseband (2.4 GHz is a perfect multiple of 800 MHz). Experimentation with sampling rate can be done in the lab with a signal generator.

Undersampling is really our only option for this high carrier frequency. A front end band-pass filter that isolates the Wi-Fi band will help eliminate aliasing onto other transmissions.

The system could be altered similarly to observe GPS signals. Examining the GPS L1 P-code waveform, we see a similar bandwidth to Wi-Fi, but a center frequency of 1.575 GHz. This frequency is very close to Nyquist folds of each sampling frequency we have examined. A tight band-pass filter would be necessary, or of course a front end mixer. Again we can

explore different sampling frequencies with a signal generator, and buy a dedicated oscillator if we would like to take the system mobile.

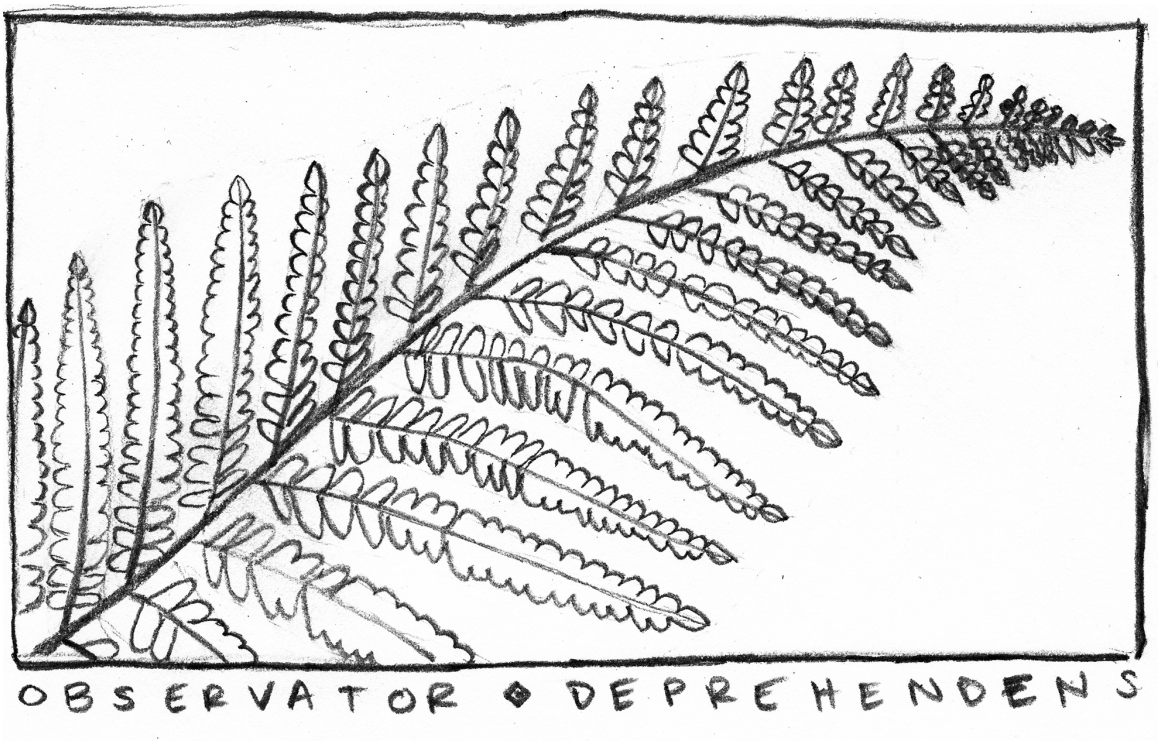
OBSERVATOR ◆ DEPREHENDENS

Chapter 9

CONCLUSIONS

Passive radar is currently an important research topic. It offers advantages over active radar in terms of cost, sensitivity, and although not directly applicable to our research project, stealth. Many systems designed today are concerned with aircraft detection from a mobile platform. We have created the next generation of the Manastash Ridge Radar (MRR), designed to look at a broad spectrum of transmitters (with the current analog-to-digital converters (ADCs) the Nyquist rate is as high as 1.5 GHz) and is designed for light weight, low cost, low power consumption, and mobile operation. This new system significantly increases the versatility of the MRR, increases the range resolution, and decreases the cost. Our key contribution to the field of passive radar is to demonstrate the efficacy of direct sampling through UHF without the use of analog down-conversion.

The system offers many advantages over current technology:

- The system provides new capabilities for ionospheric studies; there is no other radar like it for that application
- The system provides an important technology study which allows the user to combine data from several transmitters spread far in frequency space (such as FM at 100 MHz and DTV at 540 MHz) without replication of the RF path, including the digitizer
- We will be able to run experiments to contribute to the fresh field of mobile passive radar receivers, especially in terms of ground clutter studies
- The system is low-cost, $24k
- The system is open source. There are (I'm sure) many classified systems and commercial systems that protect their IP. We have a chance to offer our technological developments to the open scientific community.

- The system uses the latest technologies, which are new developments for radar and will only become less expensive and more efficient, specifically: FPGA, 10 GbE data link, and graphical processing units (GPUs).

- The system provides high aggregate bandwidth, up to 6.4 Gbps, which we can allocate as desired. A current choice is to produce 8-bit IQ samples on 4 antennas, each observing 60 MHz of the spectrum. While data loss occur at these high rates, we demonstrate the ability to detect aircraft and retrieve accurate range and Doppler velocity measurements.

In this document we presented in-depth details of performance analysis as well as design choices, strategies and challenges.

Additionally, we present real data taken with this system to detect aircraft using DTV as an illuminator of opportunity. We present aircraft detection for both mobile and stable operation. We also present the stationary detection of several different kinds of aircraft, including a helicopter as noted by the blade modulation. Preliminary results of the road-vehicle mobile campaign, in addition to aircraft, show Doppler-spread ground clutter. Additionally we include early results of aircraft detection using fast data transfer and discuss the effects of packet loss on this detection.

We demonstrate dynamic range of 60 dB in stationary mode, and 55 dB in mobile operation, close to the theoretical expectations of our performance analysis.

The results we present demonstrate the capabilities and versatility of this new system.

BIBLIOGRAPHY

[1] CJ Baker, HD Griffiths, and I. Papoutsis. Passive coherent location radar systems. part 2: Waveform properties. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 152, pages 160–168. IET, 2005.

[2] J. Baniak, G. Baker, A.M. Cunningham, and L. Martin. Silent sentry passive surveillance. *Aviation week and space technology*, 7:134–139, 1999.

[3] P. Bezousek and V. Kubecek. A 3d passive surveillance system vera accuracy analysis. In *Microwaves, Radar and Wireless Communications. 2000. MIKON-2000. 13th International Conference on*, volume 1, pages 25–28. IEEE, 2000.

[4] S. Carson, D. Kilfoyle, M. Potter, and J. Vance. A passive, multi-static radar system. In *Radar Systems, 2007 IET International Conference on*, pages 1–4. IET, 2007.

[5] Hong Chen and Mark Wagner. Casper. In *External Adjustment for Atmel/e2v iADC*, August 2010.

[6] M. Cherniakov, D. Nezlin, and K. Kubik. Air target detection via bistatic radar based on leos communication signals. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 149(1), pages 33–38. IET, 2002.

[7] Mikhail Cherniakov. *Bistatic radars: Emerging technology.* Wiley, 2008.

[8] Gerald W Collins. *Fundamentals of digital television transmission.* Wiley, 2001.

[9] S. Cross, R. Crida, T. Bennett, and M. Welz. Guidelines for communicating with devices, 2009.

[10] Ivo De Lotto and Gian Enrico Paglia. Dithering improves a/d converter linearity. *Instrumentation and Measurement, IEEE Transactions on*, 1001(2):170–177, 1986.

[11] A. Di Lallo, A. Farina, R. Fulcoli, P. Genovesi, R. Lalli, and R. Mancinelli. Design, development and test on real data of an fm based prototypical passive radar. In *Radar Conference, 2008. RADAR'08. IEEE*, pages 1–6. IEEE, 2008.

[12] Mausumi Dikpati and Peter Gilman. Scientists issue unprecedented forecast of next sunspot cycle. *NCAR (National Center for Atmospheric Research) e UCAR (University Corporation for Atmospheric Research) Office Programs*, 6, 2006.

[13] J. Dupont. Strange tale of french passive radar project. *Air & Cosmos, Aviation Magazine International*, 1984:22–23, 2005.

[14] ADC Enhancing. resolution by oversampling, avr121 application note, rev. Technical report, 8003A-AVR-09/05, www. atmel. com, 2005.

[15] P. Falcone, F. Colone, C. Bongioanni, and P. Lombardo. Experimental results for ofdm wifi-based passive bistatic radar. In *Radar Conference, 2010 IEEE*, pages 516–521. IEEE, 2010.

[16] Z. Gao, R. Tao, Y. Ma, and T. Shao. Dvb-t signal cross-ambiguity functions improvement for passive radar. In *Radar, 2006. CIE'06. International Conference on*, pages 1–4. IEEE, 2006.

[17] M. Glende, J. Heckenbach, H. Kuschel, S. Müller, J. Schell, and C. Schumacher. Experimental passive radar systems using digital illuminators (dab/dvb-t). In *Proc. of the International Radar Symposium (IRS 2007), Cologne, Germany*, pages 5–7, 2007.

[18] E.F. Greneker and J.L. Geisheimer. The use of passive radar for mapping lightning channels in a thunderstorm. In *Radar Conference, 2003. Proceedings of the 2003 IEEE*, pages 28–33. IEEE, 2003.

[19] H.D. Griffiths and C.J. Baker. Measurement and analysis of ambiguity functions of passive radar transmissions. In *Radar Conference, 2005 IEEE International*, pages 321–325. IEEE, 2005.

[20] HD Griffiths and CJ Baker. Passive coherent location radar systems. part 1: Performance prediction. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 152, pages 153–159. IET, 2005.

[21] HD Griffiths, AJ Garnett, CJ Baker, and S. Keaveney. Bistatic radar using satellite-borne illuminators of opportunity. In *Radar 92. International Conference*, pages 276–279. IET, 1992.

[22] HD Griffiths and NRW Long. Television-based bistatic radar. *Communications, Radar and Signal Processing, IEE Proceedings F*, 133(7):649–657, 1986.

[23] F.J. Harris, C. Dick, and M. Rice. Digital receivers and transmitters using polyphase filter banks for wireless communications. *Microwave Theory and Techniques, IEEE Transactions on*, 51(4):1395–1412, 2003.

[24] David H. Hathaway and Lisa Rightmire. Variations in the suns meridional flow over a solar cycle. *Science*, 327(5971):1350–1352, 2010.

[25] Dr. David H. Hathaway. Solar cycle prediction, June 2013.

[26] X. He, T. Zeng, and M. Cherniakov. Signal detectability in ss-bsar with gnss non-cooperative transmitter. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 152(3), pages 124–132. IET, 2005.

[27] Texas Instrument. Graychip 4016 data sheet rev1. 0, 2001.

[28] Akira Ishimaru. *Electromagnetic wave propagation, radiation, and scattering*. Prentice Hall Englewood Cliffs, NJ, 1991.

[29] Elliott D Kaplan and Christopher J Hegarty. *Understanding GPS: principles and applications*. Artech House Publishers, 2006.

[30] Specification sheet: Adc08d1520 low power, 8-bit, dual 1.5 gsps or single 3.0 gsps a/d converter, March 2013.

[31] H Kuschel, J Heckenbach, D O'Hagan, and M Ummenhofer. A hybrid multi-frequency passive radar concept for medium range air surveillance. In *Microwaves, Radar and Remote Sensing Symposium (MRRS), 2011*, pages 275–279. IEEE, 2011.

[32] H. Kuschel and D. O'Hagan. Passive radar from history to future. In *Radar Symposium (IRS), 2010 11th International*, pages 1–4. IEEE, 2010.

[33] F.D. Lind, A.J. Coster, P.J. Erickson, J.C. Foster, and J.D. Sahr. Intercepted signals for ionospheric science. In *General Assembly and Scientific Symposium, 2011 XXXth URSI*, pages 1–1. IEEE, 2011.

[34] F.D. Lind, J.D. Sahr, and D.M. Gidner. First passive radar observations of auroral e-region irregularities. *Geophysical research letters*, 26(14):2155–2158, 1999.

[35] KA Lukin. Noise radar technology. *Telecommunications and Radio Engineering*, 55(12), 2001.

[36] M. Malanowski, K. Kulpa, and J. Misiurewicz. Parade-passive radar demonstrator family development at warsaw university of technology. In *Microwaves, Radar and Remote Sensing Symposium, 2008. MRRS 2008*, pages 75–78. IEEE, 2008.

[37] N. Millet and M. Klein. Passive radar air surveillance: last results with multi-receiver systems. In *Radar Symposium (IRS), 2011 Proceedings International*, pages 281–285. IEEE, 2011.

[38] A.N. Morabito, J.D. Sahr, Z.M.P. Berkowitz, and L.E. Vertatschitsch. Phase noise analysis in distributed coherent noise radar systems. In *Radar Symposium, 2008 International*, pages 1–4. IEEE, 2008.

[39] Karl Erik Olsen and Karl Woodbridge. Performance of a multiband passive bistatic radar processing schemepart i. *Aerospace and Electronic Systems Magazine, IEEE*, 27(10):16–25, 2012.

[40] A.V. Oppenheim, R.W. Schafer, J.R. Buck, et al. *Discrete-time signal processing*, volume 2. Prentice hall Englewood Cliffs, NJ:, 1989.

[41] Donald B Percival and Andrew T Walden. *Spectral analysis for physical applications*. Cambridge University Press, 1993.

[42] Tony Phillips. Solar wind loses power, hits 50-year low. *NASA Headline News*, 2008.

[43] Tony Phillips. Spotless sun: Blankest year of the space age. *NASA Headline News*, 2008.

[44] Tony Phillips. Deep solar minimum. *NASA Headline News*, 2009.

[45] M.A. Ringer and G.J. Frazer. Waveform analysis of transmissions of opportunity for passive radar. In *Signal Processing and Its Applications, 1999. ISSPA'99. Proceedings of the Fifth International Symposium on*, volume 2, pages 511–514. IEEE, 1999.

[46] George T Ruck, Donald E Barrick, William D Stuart, and Clarence K Krichbaum. *Radar cross section handbook*, volume 1. Plenum press New York, 1970.

[47] J.D. Sahr. An investigation of a distributed passive radar system. *University of Washington Department of Electrical Engineering Technical Report*, 2001. https://www.ee.washington.edu/techsite/papers/refer/UWEETR-2002-0008.html or by request of author.

[48] J.D. Sahr and F.D. Lind. The manastash ridge radar: A passive bistatic radar for upper atmospheric radio science. *Radio science*, 32(6):2345–2358, 1997.

[49] R. Saini and M. Cherniakov. Dtv signal ambiguity function analysis for radar application. In *Radar, Sonar and Navigation, IEE Proceedings-*, volume 152, pages 133–142. IET, 2005.

[50] R. Saini, M. Cherniakov, and V. Lenive. Direct path interference suppression in bistatic system: Dtv based radar. In *Radar Conference, 2003. Proceedings of the International*, pages 309–314. IEEE, 2003.

[51] J. Scheer and J.L. Kurtz. *Coherent radar performance estimation*. Artech House, 1993.

[52] Alexander Schroeder and Michael Edrich. Cassidian multiband mobile passive radar system. In *Radar Symposium (IRS), 2011 Proceedings International*, pages 286–291. IEEE, 2011.

[53] I. Skolnik Merrill. Radar handbook, 1990.

[54] D. Trizna and J. Gordon. Results of a bistatic hf radar surface wave sea scatter experiment. In *Geoscience and Remote Sensing Symposium, 2002. IGARSS'02. 2002 IEEE International*, volume 3, pages 1902–1904. IEEE, 2002.

[55] unknown. Celldar - cellphone radar system, November 2012. The original website has been taken down, this reference is hosted by a web archive and is missing any original graphics.

[56] LE Vertatschitsch, JD Sahr, P Colestock, and S Close. Meteoroid head echo polarization features studied by numerical electromagnetics modeling. *Radio Science*, 46(6), 2011.

[57] Mathias Weiß. Multi-sensor systems. In *Lecture Series Set-157on Multisensor Fusion: Advanced Methodologies and Applications*. NATO RTO SET-157, 2012.

[58] N.J. Willis. *Bistatic radar*. SciTech Publishing, 2005.

[59] VP Yakubov, VB Antipov, DN Losev, and IA Yuriev. Passive radar detection of radioactive pollution. In *Application of the Conversion Research Results for International Cooperation, 1999. SIBCONVERS'99. The Third International Symposium*, volume 2, pages 397–399. IEEE, 1999.

[60] C.L. Zoeller, J. MC Budge, and M.J. Moody. Passive coherent location radar demonstration. In *System Theory, 2002. Proceedings of the Thirty-Fourth Southeastern Symposium on*, pages 358–362. IEEE, 2002.

Appendix A

POLYPHASE DECIMATION

A detailed description of polyphase decimation can be readily found; I like [40]. The mathematical proof relies heavily on the Z transform and proof of some correlaries. It is done nicely in many resources, however here I opt for a time domain motivation.

Decimation is the act of reducing the rate of a signal by first using an anti-alias filter and then downsampling as shown in Figure A.1.

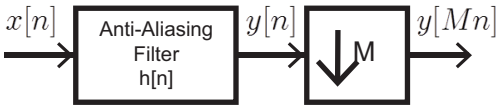

Figure A.1: Block diagram of decimation.

Sequence $y[n]$ is computed at high rate, and after all of this computation, only every Mth sample is saved to reduce the data rate. This process is wasteful, as the filter calculates every term of output $y[n]$ and then we simply toss out most of the samples. With a finite impulse response (FIR) filter of order N, it takes N multiplies and N-1 adds to compute each output sample in $y[n]$. Only after M samples are computed in $y[n]$ is one saved to produce the signal $y[Mn]$. Therefore it takes $M(2N-1)$ computations, $\mathcal{O}(MN)$, to compute each sample in final decimated sequence. $(M-1)(2N-1)$ computations are ultimately wasted.

Polyphase decimation is a multi-rate filter bank that will only compute the values of $y[n]$ intended to keep. The rate change happens first, decimating the input signal $x[n]$ into streams of parallel data running at $1/M$ the input rate. These streams are created as:

$$x_k[n] = x[Mn + k], k = 0..M-1 \tag{A.1}$$

In interleaved mode, the KATADC will feed the FPGA $M = 8$ parallel streams of data.

If our initial sample rate is 1.6 GHz, we can design a high-rate filter $h[n]$ that will select the band we want. To implement this as a polyphase, we can break the long filter into subbands, each with identical magnitude by different phase.

$$h_k[n] = h[Mn + k], k = 0..M - 1 \tag{A.2}$$

The streaming data is matched to the filters as follows:

$$\begin{aligned}
y_0[n] &= x_0[n] * h_0[n] && \text{(A.3)}\\
y_1[n] &= x_1[n-1] * h_7[n] && \text{(A.4)}\\
y_2[n] &= x_2[n-1] * h_6[n] && \text{(A.5)}\\
y_3[n] &= x_3[n-1] * h_5[n] && \text{(A.6)}\\
y_4[n] &= x_4[n-1] * h_4[n] && \text{(A.7)}\\
y_5[n] &= x_5[n-1] * h_3[n] && \text{(A.8)}\\
y_6[n] &= x_6[n-1] * h_2[n] && \text{(A.9)}\\
y_7[n] &= x_7[n-1] * h_1[n] && \text{(A.10)}\\
& && \text{(A.11)}
\end{aligned}$$

Finally we can create the output sequence $z[n]$ (mathematically equal to $y[8n]$) by:

$$z[n] = \sum_{k=0}^{k=7} y_k[n] \tag{A.12}$$

Figure A.2 shows an example of the Simulink layout for the polyphase filter. The filter coefficients are loaded from a file for testing, and are hard-coded into the file when compiled for the FPGA.

Each filter output sample $y_k[n]$ requires N/M adds and N/M+1 multiplies. There are M of these streams, creating a total of 2N-1 calculations. Finally each $y_k[n]$ sample must be added together (M-1 adds) for a total of 2N-1, exactly reduced by 8, calculations per output sample $z[n]$.

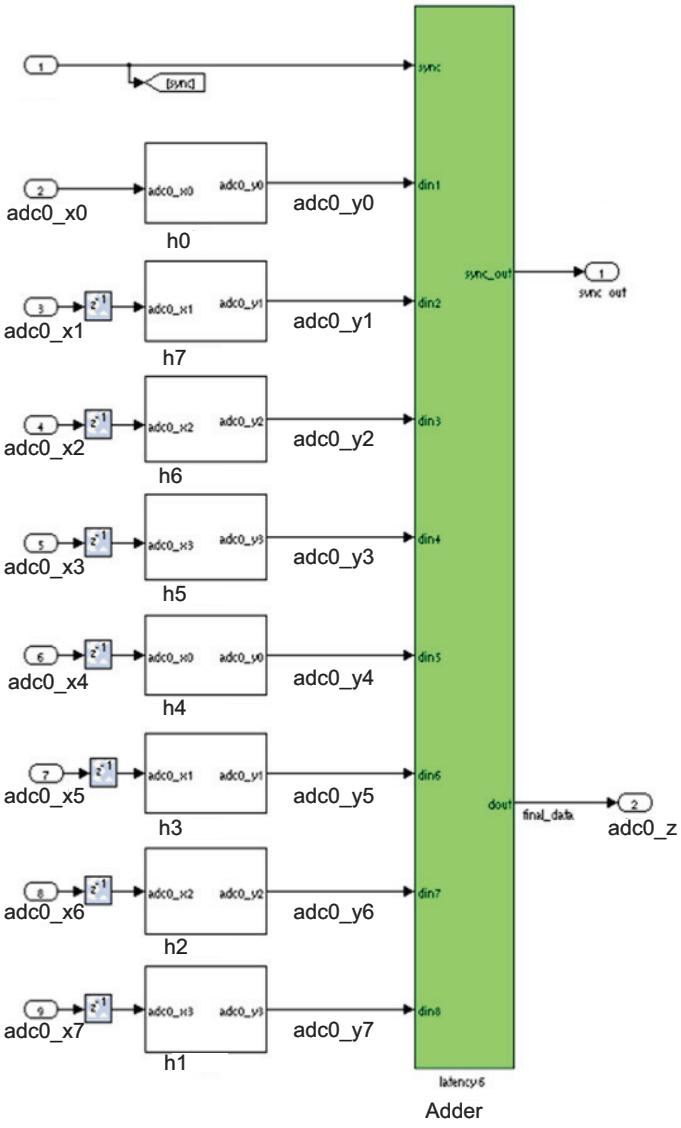

Figure A.2: Example of polyphase decimation as seen in Simulink.

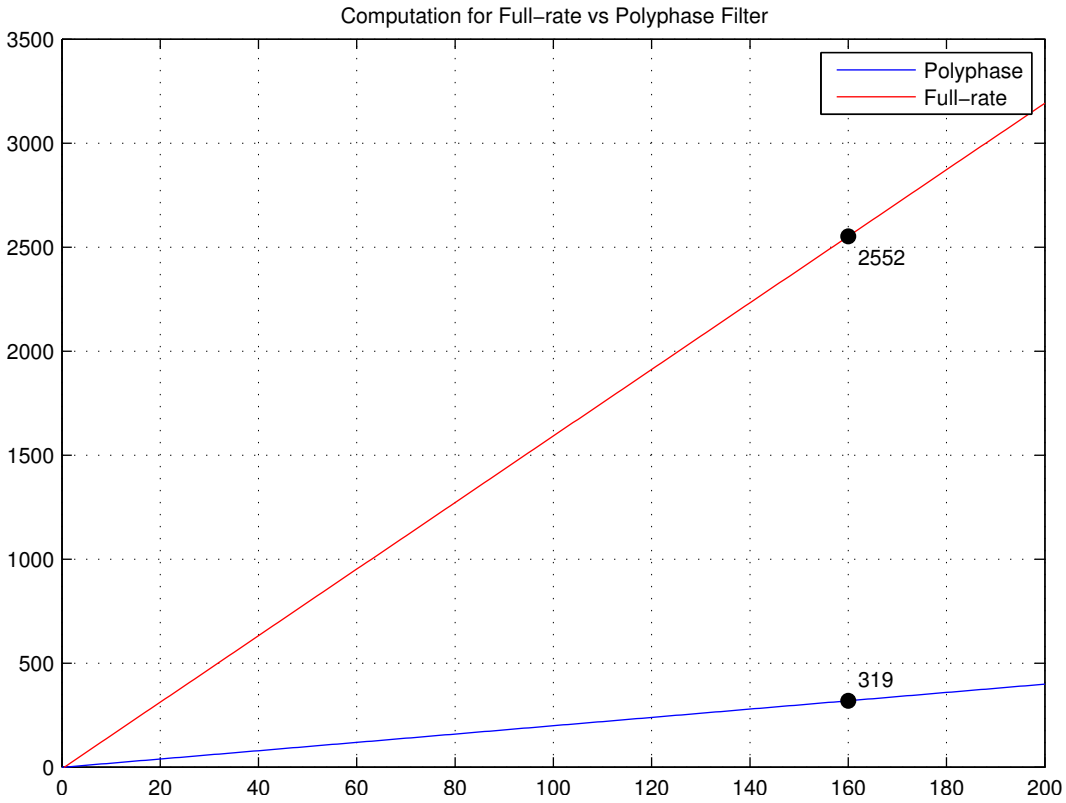

Figure A.3: Polyphase computation comparison to high rate decimation. For 160 taps, polyphase filtering only takes 319 computations per clock whereas straight-forward decimation would take over 2500.

The number of computations for $M = 8$ and variable tap length is plotted in Figure A.3. One of our initial filters uses an FIR equiripple filter with length 160 taps. To complete this with traditional decimation would take 2552 operations per clock cycle, whereas the polyphase filter implementation takes only 319 computations per clock.

Appendix B

ANIMATION OF DETECTIONS

Animations of the stationary detection of a helicopter can be found in the supplemental files.

Animated GIF of helicopter detection using DTV transmitter KONG: `KONGadc01_null.gif`

Animated GIF of helicopter detection using DTV transmitter KZJO: `KZJOadc01_null.gif`

Animated GIF of bistatic range plotted on map of local Seattle area for helicopter detection from both KZJO and KONG transmitters: `mapKZJOKONG.gif`

Appendix C

SIGNAL MODELLING AND CROSS-AMBIGUITY CALCULATIONS

We begin by modeling the signal $x_m(t)$ received by antenna m as the sum of a direct path and an indirect path, delayed by τ and doppler shifted by f_0. The (perfect, ideal) transmitted signal is $x(t)$.

$$x_m(t) = e^{j\theta_m}x(t) + \alpha e^{j\phi_m}x(t-\tau_0)e^{j2\pi f_0(t-\tau_0)} \quad \text{(C.1)}$$

$$\theta_m = k\vec{r_m}\cdot\widehat{tr}, \text{ antenna } m \text{ phase associated with the direct path} \quad \text{(C.2)}$$

$$\phi_m = k\vec{r_m}\cdot\widehat{sr}, \text{ antenna } m \text{ phase associated with the scatter} \quad \text{(C.3)}$$

There are a variety of cross ambiguities that can be calculated. If we assume that the antenna spacing is very small compared to the correlation length of the transmitter signal (a very reasonable assumption for our system), then

$$\chi(\tau,\nu) = \langle x(t)x^*(t-\tau)e^{-j2\pi\nu t}\rangle \text{ the perfect ambiguity} \quad \text{(C.4)}$$

$$\chi_{pq}(\tau,\nu) = \langle x_p(t)x_q^*(t-\tau)e^{-j2\pi\nu t}\rangle \quad \text{(C.5)}$$

where the $\langle\rangle$ notation represents the statistical expected value operator. To perform this computation we begin by stating $x_p(t)$ and $x_q^*(t-\tau)$.

$$x_p(t) = e^{j\theta_p}x(t) + \alpha e^{j\phi_p}e^{j2\pi f_0(t-\tau_0)}x(t-\tau_0) \quad \text{(C.6)}$$

$$x_q^*(t-\tau) = e^{-j\theta_q}x^*(t-\tau) + \alpha^* e^{-j\phi_q}e^{-j2\pi f_0(t-\tau-\tau_0)}x^*(t-\tau-\tau_0) \quad \text{(C.7)}$$

$\chi_{pq}(\tau,\nu)$ can be broken into the 4 individual cross terms where the expected value operation

distributes across addition. We will solve for each term individually.

$$\langle e^{j\theta_p} x(t) e^{-j\theta_q} x^*(t-\tau) e^{-j2\pi\nu t} \rangle = e^{j(\theta_p-\theta_q)} \langle x(t) x^*(t-\tau) e^{-j2\pi\nu t} \rangle \quad \text{(C.8)}$$

$$= e^{j(\theta_p-\theta_q)} \chi(\tau,\nu) \quad \text{(C.9)}$$

$$\langle e^{j\theta_p} x(t) \alpha^* e^{-j\phi_q} e^{-j2\pi f_0(t-\tau-\tau_0)} x^*(t-\tau-\tau_0) e^{-j2\pi\nu t} \rangle = \quad \text{(C.10)}$$

$$= \alpha^* e^{j(\theta_p-\phi_q)} \langle x(t) x^*(t-(\tau+\tau_0)) e^{-j2\pi f_0(t-(\tau+\tau_0))} e^{-j2\pi\nu t} \rangle \quad \text{(C.11)}$$

$$= \alpha^* e^{j(\theta_p-\phi_q)} e^{-j2\pi\nu(\tau+\tau_0)} \langle x(t) x^*(t-(\tau+\tau_0)) e^{-j2\pi(\nu+f_0)(t-(\tau+\tau_0))} \rangle \quad \text{(C.12)}$$

$$= \alpha^* e^{j(\theta_p-\phi_q)} e^{-j2\pi\nu(\tau+\tau_0)} \chi(\tau+\tau_0, \nu+f_0) \quad \text{(C.13)}$$

$$\langle \alpha e^{j\phi_p} e^{j2\pi f_0(t-\tau_0)} x(t-\tau_0) e^{-j\theta_q} x^*(t-\tau) e^{-j2\pi\nu t} \rangle = \quad \text{(C.14)}$$

$$= \alpha e^{j(\phi_p-\theta_q)} \langle x(t-\tau_0) x^*(t-\tau) e^{j2\pi f_0(t-\tau_0)} e^{-j2\pi\nu t} \rangle \quad \text{(C.15)}$$

$$= \alpha e^{j(\phi_p-\theta_q)} \langle x(t) x^*(t-(\tau-\tau_0)) e^{j2\pi f_0 t} e^{-j2\pi\nu(t+\tau_0)} \rangle \quad \text{(C.16)}$$

$$= \alpha e^{j(\phi_p-\theta_q)} e^{-j2\pi\nu\tau} e^{j2\pi f_0(\tau-\tau_0)} \langle x(t) x^*(t-(\tau-\tau_0)) e^{-j2\pi(\nu-f_0)(t-(\tau-\tau_0)} \rangle \quad \text{(C.17)}$$

$$= \alpha e^{j(\phi_p-\theta_q)} e^{-j2\pi\nu\tau} e^{j2\pi f_0(\tau-\tau_0)} \chi(\tau-\tau_0, \nu-f_0) \quad \text{(C.18)}$$

$$\langle \alpha e^{j\phi_p} e^{j2\pi f_0(t-\tau_0)} x(t-\tau_0) \alpha^* e^{-j\phi_q} e^{-j2\pi f_0(t-\tau-\tau_0)} x^*(t-\tau-\tau_0) e^{-j2\pi\nu t} \rangle \quad \text{(C.19)}$$

$$= |\alpha|^2 e^{j(\phi_p-\phi_q)} e^{j2\pi f_0\tau} \langle x(t-\tau_0) x^*(t-\tau-\tau_0) e^{-j2\pi\nu t} \rangle \quad \text{(C.20)}$$

$$= |\alpha|^2 e^{j(\phi_p-\phi_q)} e^{j2\pi f_0\tau} e^{-j2\pi\nu\tau_0} \langle x(t) x^*(t-\tau) e^{-j2\pi\nu t} \rangle \quad \text{(C.21)}$$

$$= |\alpha|^2 e^{j(\phi_p-\phi_q)} e^{j2\pi f_0\tau} e^{-j2\pi\nu\tau_0} \chi(\tau,\nu) \quad \text{(C.22)}$$

If we can assume that $\chi(\tau,\nu) \approx \delta(\tau,\nu)$, then

$$\langle e^{j\theta_p} x(t) e^{-j\theta_q} x^*(t-\tau) e^{-j2\pi\nu t} \rangle = e^{j(\theta_p-\theta_q)} \langle x(t) x^*(t-\tau) e^{-j2\pi\nu t} \rangle \quad \text{(C.23)}$$

$$\approx e^{j(\theta_p-\theta_q)} \delta(\tau,\nu) \quad \text{(C.24)}$$

$$\langle e^{j\theta_p} x(t) \alpha^* e^{-j\phi_q} e^{-j2\pi f_0(t-\tau-\tau_0)} x^*(t-\tau-\tau_0) e^{-j2\pi\nu t} \rangle = \quad \text{(C.25)}$$

$$\approx \alpha^* e^{j(\theta_p-\phi_q)} \delta(\tau+\tau_0, \nu+f_0) \quad \text{(C.26)}$$

$$\langle \alpha e^{j\phi_p} e^{j2\pi f_0(t-\tau_0)} x(t-\tau_0) e^{-j\theta_q} x^*(t-\tau) e^{-j2\pi\nu t} \rangle = \quad \text{(C.27)}$$

$$\approx \alpha e^{j(\phi_p-\theta_q)} e^{-j2\pi f_0\tau_0} \delta(\tau-\tau_0, \nu-f_0) \quad \text{(C.28)}$$

$$\langle \alpha e^{j\phi_p} e^{j2\pi f_0(t-\tau_0)} x(t-\tau_0) \alpha^* e^{-j\phi_q} e^{-j2\pi f_0(t-\tau-\tau_0)} x^*(t-\tau-\tau_0) e^{-j2\pi\nu t} \rangle \quad \text{(C.29)}$$

$$\approx |\alpha|^2 e^{j(\phi_p-\phi_q)} \delta(\tau,\nu) \quad \text{(C.30)}$$

And we can finally approximate $\chi_{pq}(\tau,\nu)$ as:

$$\chi_{pq}(\tau \geq 0, \nu) \approx \left(e^{j(\theta_p-\theta_q)} + |\alpha|^2 e^{j(\phi_p-\phi_q)} \right) \delta(\tau,\nu) + \alpha e^{j(\phi_p-\theta_q)} e^{j2\pi f_0\tau_0} \delta(\tau-\tau_0, \nu-f_0) \quad \text{(C.31)}$$

The resulting Equation C.31 shows us that the peak of this cross-ambiguity at τ_0 and f_0 contains information about the angle of the incoming scattered signal relative to the direct transmission. It is also important to note that this peak scales with α, the voltage level attenuation of the scattered signal compared to the direct path. This proves that this

simple cross-ambiguity calculation is still in unites of voltage, and to obtain power we need to compute a full four-term cross-ambiguity, $|\chi(\tau,\nu)|^2$. If our assumptions hold and α is sufficiently small ($|\alpha|^2 \ll 1$), then:

$$|\chi(\tau \geq 0,\nu)|^2 \approx \delta(\tau,\nu) + |\alpha|^2\delta(\tau-\tau_0,\nu-f_0) \tag{C.32}$$

VITA

Laura Vertatschitsch earned a Bachelor of Arts degree from Whitman College for a combined major of Mathematics-Physics in 2006. She is a PhD Candidate at the University of Washington Department of Electrical Engineering, with expected graduation in 2013. A move across the country awaits this Washingtonian as she embarks on a new adventure as a post-doctoral fellow at the Harvard-Smithsonian Center for Astrophics in Cambridge, MA. There she will continue her research in high bandwidth instrumentation design by contributing to the Event Horizon Telescope (EHT) project.

Publications

L. Vertatschitsch, J. Sahr, P. Colestock, S. Close. Meteoroid head echo polarization features studied by numerical electromagnetics modeling. *Radio Science*, 46(6):RS6016, 2011.

S. Close, M. Kelley, L. Vertatschitsch, P. Colestock, M. Oppenheim, and J. Yee. Polarization and scattering of a long-duration meteor trail. *Journal of Geophysical Research*, 116(A1):A01309, 2011.

A.N. Morabito, D.B. Percival, J.D. Sahr, Z.M.P. Berkowitz, and L.E. Vertatschitsch. Ricean parameter estimation using phase information in low snr environments. *Communications Letters, IEEE*, 12(4):244246, 2008.

A.N. Morabito, J.D. Sahr, Z.M.P. Berkowitz, and L.E. Vertatschitsch. Phase noise analysis in distributed coherent noise radar systems. In *Radar Symposium, 2008 International*, pages 14. IEEE, 2008.

A.N. Morabito, J.D. Sahr, Z.M.P. Berkowitz, and L.E. Vertatschitsch. Post-processing phase noise mitigation performance comparison in a coherent distributed passive radar system. In *Radar Symposium, 2008 International*, pages 14. IEEE, 2008.

AH Tewfik, AB Tchagang, and L. Vertatschitsch. Parallel biclustering of genes with

coherent evolutions: Algorithm and biological significance of the biclusters. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on,* volume 2, pages IIII. IEEE, 2006. 433

A.H. Tewfik, A.B. Tchagang, and L. Vertatschitsch. Parallel identification of gene biclusters with coherent evolutions. *Signal Processing, IEEE Transactions on,* 54(6):24082417, 2006.

L. Vertatschitsch, A.H. Tewfik, and A.B. Tchagang. Biological significance of a novel biclustering technique on genetic expression data. In *Signal Processing and Information Technology, 2005. Proceedings of the Fifth IEEE International Symposium on,* pages 3539. IEEE, 2005.

Conference Talks

L. Vertatschitsch, J. Sahr, "A New Mobile Passive Radar Designed For Transmitters Up to 1.5 GHz For Detection of Aerospace and Geoscience Targets," Jan 2013 UNSC-USRI National Radio Science Meeting, Boulder, CO.

L. Vertatschitsch, J. Sahr, W. Sun, "First Passive Radar Receiver Implementation with ROACH," Aug 2012 CASPER 2012 Workshop, Greenbank, WV.

L. Vertatschitsch, J. Sahr, P. Colestock, S.Close, "Polarization of Meteor Head Echoes Numerical Modeling ," Jan 2011 UNSC-USRI National Radio Science Meeting, Boulder, CO.

L. Vertatschitsch, J. Sahr, P. Colestock, S.Close, "Polarization of Meteor Head Echoes Numerical Modeling ," May 2010 Meteoroids Meeting, Breckenridge, CO.

L. Vertatschitsch, S. Close, P. Colestock, J. Sahr, "Simultaneous VHF/UHF Detection and Analysis of Polarization Properties of Head Echoes," Jan 2010 UNSC-USRI National Radio Science Meeting, Boulder, CO.

L. Vertatschitsch, S. Close, P. Colestock, "Polarization Properties of Meteor Head Echoes," June 2009 CEDAR Workshop, Santa Fe, NM.

L. Vertatschitsch, S. Close, P. Colestock, J. Sahr, "Polarization Properties of Head Echoes Collected at ALTAIR," Jan 2009 UNSC-USRI National Radio Science Meeting, Boulder, CO.

L. Vertatschitsch, J. Sahr, "A Model of the MRR Distributed Passive Radar System,"

Jan 2008 UNSC-USRI National Radio Science Meeting, Boulder, CO.

Poster Sessions

L. Vertatschitsch, P. Colestock, S. Close, J. Sahr, "Using FDTD Simulations to Understand Observed Polarization Properties of Meteor Head Echoes," Dec 2009 AGU Fall Meeting, San Francisco, CA.