

UNIVERSITY OF OKLAHOMA
GRADUATE COLLEGE

SIGNAL PROCESSING TECHNIQUES FOR RADAR OPERATION IN
CONGESTED SPECTRUM

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

By
RYLEE G. MATTINGLY
Norman, Oklahoma
2024

SIGNAL PROCESSING TECHNIQUES FOR RADAR OPERATION IN
CONGESTED SPECTRUM

A DISSERTATION APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Justin Metcalf, Chair

Dr. Cliff Fitzmorris

Dr. Nathan Goodman

Dr. Jay McDaniel

Dr. Mark Yeary

Dr. Sridhar Radhakrishnan

© Copyright by RYLEE G. MATTINGLY 2024
All Rights Reserved.

To Mom and Dad

Acknowledgments

I want to thank Dr. Justin Metcalf and Dr. Jay McDaniel for giving me the opportunity to work at the ARRC and supporting this dissertation work and the numerous side projects that cropped up along the way. Additionally, I thank Dr. Cliff Fitzmorris, Dr. Nathan Goodman, Dr. Mark Yeary, and Dr. Sridhar Radhakrishnan for taking the time to be on my committee and providing thoughtful questions and valuable feedback that significantly improved the quality of my work and this dissertation.

I would also like to thank my Mom, Treva Mattingly, and sister, Cassi Marlow, for always supporting me through challenges and celebrating each success. Additionally, I must acknowledge all the fantastic faculty, staff, and students at the ARRC who make it such an incredible workplace. Specifically, I would like to thank Rachel, Russel, Jon, Shane, Cora, and Cal for their steadfast support, friendship, and the occasional game of pool and Dale for the morning coffee chats that always helped kick start the work day.

I thank the Defense Advanced Research Projects Agency (DARPA) for their support of this research under grant HR0011-20-1-0007. I also thank the Air Force Research Laboratory (AFRL) and DARPA for their support under contract FA2384-23-C-004. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

Table of Contents

List of Tables	ix
List of Figures	x
Abstract	xvii
1 Introduction	1
1.1 Objectives	4
1.2 Outline	5
2 Background	7
2.1 The Pulse-Doppler Radar Model	7
2.1.1 Simulation	11
2.1.2 Experimental Testbed	12
2.2 CFAR Detection	13
2.3 OFDM Communications	16
2.3.1 Capacity	18
2.3.2 Pilots and Structure	20
2.4 Spectrum Sharing	24
3 Intra-CPI Bandwidth Agility	28
3.1 Mismatched Filtering	30

3.1.1	Formulation	32
3.2	Common Matched Responses	33
3.3	Quantitative Results	37
3.4	Experimental Results	43
4	Intra-CPI Frequency Agility	46
4.1	Phase Analysis	49
4.1.1	Phase Correction	52
4.2	Modified Backprojection Algorithm	53
4.2.1	Backprojection	54
4.2.2	Modified Backprojection	55
4.2.3	Range Migration	58
4.3	Why Backprojection?	60
4.4	Experimental Results	62
5	Fully Agile Pulse Processing	65
5.1	Computational Complexity	68
6	Communications Interference Mitigation	70
6.1	Demodulation//Remodulation Estimation: The Ideal Model	72
6.1.1	Initial Performance	76
6.1.2	Blanking Location	80
6.1.3	Cyclic Prefix Stacking	86
6.1.4	Radar Detection Performance	88
6.2	Increasing Model Fidelity	92
6.2.1	Symbol Boundary Estimation	93
6.2.2	Frequency Offset Estimation	98

6.2.3	Channel Estimation	100
6.3	Expanded Demodulation/Remodulation Estimation	102
6.3.1	The Median Filter	108
6.4	Performance Analysis	114
6.5	A Comparison: Adaptive Pulse Compression	124
6.5.1	Ground Truth Performance	127
6.5.2	Alternative Estimates	129
6.6	Summary	135
7	Conclusion and Future Work	138
7.1	Conclusion	140
7.2	Future Work	142
	References	144

List of Tables

2.1	Simulation Parameters	12
2.2	OFDM Parameters	18
3.1	Bandwidth Agility Performance	41
5.1	Full Agility Performance	66
6.1	Probability of detection improvement achieved by full demod/remod interference cancellation processing chain when using a 8 MHz and 25 MHz waveforms.	122

List of Figures

2.1	A reference range-Doppler map containing a target at a range of 40 km traveling at 27 m/s when using a 3 GHz center frequency.	11
2.2	Radar test scene.	14
2.3	The noise distribution and target plus noise distribution split with a detection threshold. The probability of false alarm (left) and probability of missed detection (right) are shown.	15
2.4	Guard and reference cell placement relative to the cell under test in CA-CFAR.	15
2.5	The IFFT used to generate the time domain OFDM signal from sub-carrier.	18
2.6	Symbol constellations for QPSK, 16-QAM, and 64-QAM.	19
2.7	The 7 tap vehicular channel sampled at 30.72 MHz.	21
2.8	The LTE radio frame structure.	23
2.9	Pilots, denoted with a P , are distributed across symbols in time and subcarriers in frequency.	23
2.10	An example of pulses used for processing overlaid on the spectrum collect the band selections were derived from.	26
2.11	The center frequency, $F_H(m)$, and bandwidth, $\beta(m)$ used for the agile cases.	26

3.1	Each pulse at the output of the matched filter [52]. ©IEEE	30
3.2	The result of using traditional range-Doppler processing on a CPI of bandwidth agile pulses [52]. ©IEEE	31
3.3	Normalized PSF resulting from windowed matched filtering and FFT-based Doppler processing [52]. ©IEEE	31
3.4	Normalized PSF generated using the mismatched filter with minimum bandwidth template, (a), and the maximum bandwidth template, (b), both with FFT Doppler processing [52]. ©IEEE	35
3.5	Normalized PSF generated using the mismatched filter with super-resolution template with FFT Doppler processing [52]. ©IEEE	37
3.6	The output of the mismatched filter using the super-resolution template [52]. ©IEEE	38
3.7	The output of the minimum bandwidth matched filter, minimum bandwidth windowed matched filter, and the super-resolution mismatched filter applied to the average bandwidth waveform [52]. ©IEEE	43
3.8	The output of the maximum bandwidth matched filter, maximum bandwidth windowed matched filter, and the super-resolution mismatched filter applied to the average bandwidth waveform [52]. ©IEEE	44
3.9	Experimental data using bandwidth agile CPIs processed using the matched filter (top) and the super resolution mismatched filter (bottom).	45
4.1	Unwrapped agile baseband and agile IF slow time phase history compared to that of a CPI of pulses with no frequency agility with 100 MHz of front end bandwidth [59]. ©IEEE	50
4.2	Unwrapped agile baseband and agile IF slow time phase history compared to that of a CPI of pulses with no frequency agility with 2 GHz of front end bandwidth [59]. ©IEEE	51

4.3	Applying traditional FFT based range-Doppler processing to the baseband matched filtered data with the target's range and Doppler cuts shown [59]. ©IEEE	52
4.4	Result of applying modified backprojection to the baseband match filtered data [59]. ©IEEE	56
4.5	The constant pulse mainlobe boundaries in range and Doppler are shown in (a) and (b), respectively. The agile pulse mainlobe boundaries in range and Doppler are shown by (c) and (d) respectively [52]. ©IEEE	57
4.6	Performance of direct correction FFT processing on range migrating targets and a Doppler cut at the target's true range [52]. ©IEEE	59
4.7	Performance of modified backprojection processing on range migrating targets and a Doppler cut at the target's true range [52]. ©IEEE	60
4.8	Wrapped phase history after direct correction and backprojection correction.	62
4.9	Doppler response of uncorrected, directly corrected, and backprojection corrected phase.	63
4.10	CPI of real-world data with traditional FFT range-Doppler processing (top) and with modified backprojection (bottom) applied at the output of the matched filter [52]. ©IEEE	64
5.1	A real-world CPI of fully agile pulses processed with matched filtering and FFT-based Doppler processing without correction. The range and Doppler cuts at the target's location are shown.	67
5.2	A real-world CPI of fully agile pulses processed with super resolution mismatched filter and modified backprojection. The range and Doppler cuts at the target's location are shown.	67

6.1	A series of OFDM symbols shown with receiver blanking shown in red. [74] ©IEEE	73
6.2	The interference signal with blanked regions highlighted (top left) and the samples that are actually received by the system (top right). The bottom shows the zero inserted vector that is used for demodulation. .	74
6.3	The demodulation/remodulation estimation and extraction technique block diagram [74]. ©IEEE	75
6.4	Average power reduction from demodulation/remodulation estimate and extraction [74]. ©IEEE	77
6.5	The sinc shape changes from the nominal shape as more of the symbol is lost to the system channel.	79
6.6	The sinc shape shows a contribution of every point to each subcarrier shown for a single blank of a quarter and half of a symbol compared to the fully received case.	80
6.7	When the system channel is confined to the cyclic prefix no data is lost and the symbol is perfectly reconstructed. The top shows the region that is blanked and the bottom shows the true signal and the reconstructed signal match perfectly.	81
6.8	When the system channel causes data loss in the middle of the symbol, the BER rate is increased because of the change in the subcarrier orthogonality, however the effect is minimal. The top shows the region that is lost to the system channel and the bottom shows the true signal and the reconstructed signal.	83
6.9	The left plot shows the blank location on the true signal and the right shows the reconstruction performance relative to the true signal. . . .	85

6.10	Range-Doppler map with cyclic prefix interference clearly visible while the other parts of the symbol are extracted [74]. ©IEEE	87
6.11	Range-Doppler map with cyclic prefix stacking used to estimate and extract the cyclic prefix even when the end of the symbol is lost [74]. ©IEEE	88
6.12	Detection performance of detection after interference is canceled as a function of the ISR before cancellation with a 10 kHz[74]. ©IEEE	90
6.13	Range-Doppler map generated without interference cancellation with an ISR of 0 dB [74]. ©IEEE	91
6.14	Range-Doppler map generated after interference cancellation with an precancellation ISR of 0 dB [74]. ©IEEE	91
6.15	An OFDM timeline with two correlation regions shown as they move across the data.	95
6.16	Correlation compared to maximum likelihood estimate of the boundary shown with the true symbol boundary.	96
6.17	This figure shows two test regions for where the first test region's test statistic is lost due to the system channel. Instead of creating a peak at the symbol boundary, the largest value is the rising edge of the boundary peak in test region 2.	97
6.18	Enhanced Processing flow with boundary, frequency, and channel estimation steps included.	103
6.19	Simulated signals shown separately with the interference generated at 0 dB mean power, noise at -25 dB mean power and radar return generated so that the matched filtered, Doppler processed peak is at 0 dB in the range-Doppler map.	104

6.20	Each signal component after matched filtering. The radar return creates a peak at approximately 0 dB and the mean noise and interference power remain the 0 dB and -25 dB respectively.	105
6.21	The interference residue left behind using a demod/remod estimate using time, frequency, and channel estimation.	106
6.22	A closer look at the extra residue that exists around the symbol boundaries.	107
6.23	The range-pulse plot before (top) and after the median filter (bottom). Before the residue caused by ISI is clearly visible and is successfully detected and removed by the median filter.	111
6.24	The first PRI at the output of the median filter. The ISI residual interference has been removed and the near and far range residue from caused by misestimation of the symbol in the presence of the system channel.	112
6.25	Applying the matched filter to the output of the median filter shows a strong peak target response and even further smoothing of the near and far range residue.	113
6.26	The range-Doppler plot shows that after median filtering the target can still be focused successfully at the correct range and Doppler.	114
6.27	The complete processing flow for demod/remod interference estimation and extraction with a complex signal model.	115
6.28	Total power reduction achieved by the demod/remod technique shown with 8 MHz LFM contained completely in the interfere and a 25 MHz LFM centered with the interference.	117
6.29	The peak loss exhibited in the range-Doppler plot when using an 8 MHz waveform (top) and 25 MHz waveform (bottom).	118

6.30	The probability of detection exhibited after interference cancellation when using an 8 MHz waveform (top) and 25 MHz waveform (bottom) relative to the before cancellation probability of detection.	121
6.31	The probability of false alarm exhibited before and after interference cancellation with an 8 MHz waveform (top) and 25 MHz waveform (bottom).	123
6.32	The normalized covariance derived from the ground truth interference data.	128
6.33	APC filter output after 10 iterations using the source truth covariance matrix.	128
6.34	The normalized covariance derived from the demod/remod interference estimate.	130
6.35	APC filter output after 10 iterations using the covariance matrix derived from the demod/remod estimate.	130
6.36	Difference in the normalized ground truth covariance matrix and the demod/remod estimate covariance matrix.	132
6.37	The sample covariance generated using OFDM snapshots with random data, that shares symbol structure and channel effects with the interference signal with a zoomed in look at the channel effects on one of the structures.	134
6.38	The APC filter response after ten iterations using the sample covariance of the interference.	136

Abstract

The electromagnetic spectrum is a finite resource that has become increasingly congested as the use of wireless communications and RF applications have expanded. With increased congestion, dynamic spectrum sensing and spectrum sharing techniques have been of increasing interest. However, dynamically accessing the spectrum on a pulse-by-pulse basis breaks traditional pulse assumptions that are used to inform range-Doppler processing. Additionally, even with a dynamic hopping scheme, more spectrum than is available may be needed to successfully carry out critical radar tasks, creating two unique challenges.

First, this work addresses pulse-by-pulse agility in two parts. This work presents a mismatched filter technique to reduce range-sidelobe modulation when processing a CPI of bandwidth agile pulses. Additionally, a modified backprojection technique enables Doppler processing of frequency agile CPIs. These two techniques work together effectively on CPIs of fully agile pulses, and real-world experimental data verifies their performance.

Second, this work develops a demodulation/remodulation estimation and extraction processing chain that actively cancels interference from orthogonal frequency division multiplexing sources and provides a 19 dB improvement in radar detection performance. When combining active interference cancellation and pulse agile processing, this work provides the tools to ensure radar systems can operate continually in an environment with a congested electromagnetic spectrum.

Chapter 1

Introduction

The electromagnetic spectrum is a finite resource with increasing demand from government, industrial, and commercial interests. The spectrum has long been divided into bands; each allocated to a single user with guard bands and regulations meant to limit cross-band interference. However, as wireless communications proliferate to include more users, more devices, and new data-intensive applications, the network's need for additional spectrum only grows. To provide more spectrum for communications networks, allocations must be taken from other users, including radar.

The competitive and commercial nature of communications provides a very clear utility that justifies additional spectrum allocations for next-generation communications networks. The utility of radar allocations is harder to define, especially in the bits per hertz framework used for communications [1]. Wireless network utilization of spectrum varies heavily with time and geographic location based on the number of active users [2]. Consequently, by leveraging spectrum sharing techniques, a secondary user can identify and leverage unused spectrum, leading to more efficient use of bands that are not actively used by the primary user. Many sharing techniques have been developed to allow secondary users opportunistic and dynamic access to the spectrum. [3, 4, 5, 6, 7, 8].

Various rules-based sharing techniques have been investigated, and regulators have adopted a variety of sharing schemes for various bands [9]. Several techniques have been proposed, including citizens broadband radio service (CBRS), which defines a tiered system of users who can request spectrum access from a centralized database [10, 11]. Additional investigation has been carried out on the coexistence of 5G networks and fixed satellite services at C-band [12].

Non-cooperative techniques have also been adopted. One such non-cooperative technique is the unlicensed national information infrastructure (U-NII), or listen-before-talk, rules that allow wireless network users to share commercial weather radar bands [13]. The listen-before-talk structure requires a secondary user to 'listen' for a primary user of the spectrum before beginning operation and must continue to monitor the environment and cease operation when the primary user is present.

Taking advantage of the pulsed nature and sensitive receivers of radar systems allows for a unique opportunity to share the spectrum by allowing others to take the primary user role. Algorithms like the fast spectrum sensing (FSS) technique take advantage of the quiet time between radar pulses to gather snapshots of the spectrum and apply threshold-based detection to identify available bands of spectrum [14]. A dynamic detection threshold can be calculated from the noise estimate using techniques such as hardware optimized cell averaging estimation (HO-CAE) [15]. Combining FSS and HO-CAE algorithms provides a real-time spectrum sensing technique to identify available spectrum with the sub-millisecond speeds required to match the frequency agility of modern communications networks, allowing for a listen-while-talk type of operation [16]. Pulse-to-pulse agility is important for the radar to operate as a secondary user, thereby avoiding the bands occupied by communications users. Traditional pulsed radar systems typically transmit at high power using directional

antennas [17]. As a consequence, the radar can easily interfere with communication devices in its main beam. Further, even low interference-to-noise ratio (INR) communications signals in the radar band can cause significant interference at the output of constant false alarm rate (CFAR) detectors [18, 19, 20, 21].

Coexisting with primary users that can change frequency allocation on sub-millisecond time scales by changing waveform parameters on a pulse-to-pulse basis also causes processing challenges. Traditional range-Doppler processing utilizes the matched filter to extract range information in fast time, and the fast Fourier transform (FFT) across slow time to measure the Doppler shift, thereby obtaining the target's radial velocity relative to the radar [22]. The matched filter is used with the FFT under the assumption that every pulse in the CPI is identical. By changing the pulses within a CPI, traditional processing techniques produce a degraded or destroyed range-Doppler response.

Non-uniform pulse trains have been used to increase range resolution by transmitting narrowband waveforms in steps across a wide bandwidth [23]. Non-uniform pulse trains cause grating lobes to appear in the Doppler space, requiring intelligent design of the time duration and bandwidth to minimize the periodic phase effects that reduce performance [24]. More recently, work has explored using a non-uniform pulse train to create a Gaussian spectrum shape to reduce the sidelobe level associated with processing frequency hopping pulse trains [25]. While control of the non-uniform pulse train provides many degrees of freedom to mitigate processing effects [26], this work focuses on an arbitrary set of pulses dictated by the spectral environment. Therefore, the degrees of freedom associated with the bandwidth and center frequency are lost.

For environments with strong clutter returns, [27] presents a non-adaptive clutter cancellation technique that mitigates clutter spreading in Doppler using fast convo-

lution. Further, [28] proposed an adaptive technique that recursively optimizes filter weights to adapt in both the range and Doppler dimensions, suppressing unwanted effects from bandwidth and frequency agility. An image processing approach from [29] iteratively applies Richardson-Lucy deconvolution to compensate for the point spread function (PSF) derived from any arbitrary pulse train. Reiterative, fully adaptive techniques are effective methods of correcting for agile pulse trains. However, they are computationally intensive and not feasible for high data rate systems.

Unfortunately, spectrum-sharing schemes will always present the potential for interference when the system fails to detect primary users. Similarly, there are times when an adaptive, mission-critical radar system, like the terminal Doppler weather radar (TDWR) described in [30], may not be able to find any available spectrum, which may require operating in the presence of an interfering user. No matter the reason, when a communications user is present as an in-band interferer, techniques are needed to minimize the impact on radar processing. Ultimately, this work seeks to answer the following two questions: What is the cost of changing the frequency and bandwidth of every pulse to adapt to an open spectrum, and how can a system compensate? How can the interference be mitigated if no spectrum is available and in-band communications users are present?

1.1 Objectives

This work first presents two techniques for mitigating target response degradation caused by intra-CPI bandwidth and frequency agility. First, mismatched filters are introduced as a technique to allow the use of agile bandwidth pulse CPIs. This work identifies the most robust solution to mitigate range-sidelobe modulation by investigating different mismatch filter templates and their trade-offs. Second, a first princi-

ples analysis identifies the cause and impact of frequency agility on the phase history. The identified phase term informs the development of a modified backprojection technique to help focus targets when processing frequency agile pulses at baseband. The two techniques are complementary and, when applied together, successfully mitigate the cost of full pulse agility.

Finally, this work develops a signal processing chain that actively cancels in-band OFDM communications interference via demodulation/remodulation estimation and extraction. Demod/remod estimation reduces the interference duty cycle from 100 percent, even in complex environments with unknown time and frequency synchronization. Median filtering further reduces interference when added to the end of the processing chain while maintaining the performance of the matched filter. This work presents the efficacy of the demod/remod processing chain by examining the radar detection performance and total energy reduction.

1.2 Outline

Following this introduction, Chapter 2 provides a brief background on the pulse-Doppler radar model and traditional processing methods. A light introduction to constant false alarm rate (CFAR) detection is provided to support future discussions on radar performance carried out in Chapter 6. An OFDM signal model and a high level overview of communications structures is given before the chapter concludes with some background on spectrum sensing techniques that were used to inform the ground truth, agile pulse parameters.

Chapter 3 presents an analysis of bandwidth agile pulse processing and introduces the mismatched filter as a technique to mitigate the range sidelobe modulation. Additionally, the chapter provides the performance of different mismatched filter tem-

plates and compares them to a simple windowed matched filter response as a baseline. Chapter 4 presents a phase analysis on a CPI of frequency agile pulses and introduces backprojection as a technique to compensate for the effects of frequency hopping. The imaging function is analyzed, and a metric called relative sidelobe energy growth (RSE) is defined to provide a quantitative performance measure. Chapter 5 brings the mismatched filter and the modified backprojection algorithm together and demonstrates the algorithms' performance in a case where bandwidth and frequency vary simultaneously. Additionally, each pulse agile chapter uses experimental results from data collected with a software-defined radio (SDR) testbed to verify the effectiveness of each algorithm.

Chapter 6 sets aside the idea of agile pulses and examines the radar performance that can be achieved in the presence of a communications emitter. The chapter begins with an analysis of demodulation/remodulation estimation using a simple model, including the effects of radar receiver blanking, its position within the communication signal, and mitigation strategies that improve the robustness of the communications signal estimate. Chapter 6 continues by increasing model fidelity to include more realistic effects, including channel, time, and frequency offsets. The chapter presents techniques that overcome the challenges of the expanded model and carries out a thorough performance analysis. Finally, the chapter compares adaptive pulse compression (APC) as a reiterative method of canceling interference and analyses overall performance and robustness, especially considering the increased computational complexity. Chapter 7 concludes the dissertation by providing a summary of the work carried out and future work that can expand on the tools developed here.

Chapter 2

Background

This chapter serves to provide a background on radar processing, CFAR detection, OFDM interference, and spectrum sensing. A brief introduction to each of these concepts helps to support the discussions carried out in future chapters. This chapter also provides a set of spectrum measurements that are used for the frequency and bandwidth agile parameters to inform simulation and testing for the development of correction algorithms.

2.1 The Pulse-Doppler Radar Model

Pulse-Doppler radar begins with a series of identical pulses, each transmitted by the radar system at some pulse repetition interval (PRI). The pulses reflect off the environment and return to the radar system. The system processes a collection of the received pulses together in a coherent processing interval (CPI). This work will consider the linear frequency modulated (LFM) waveform, which is defined as a complex baseband signal by [22]

$$x(t) = e^{j2\pi\left(\frac{-\beta}{2}t + \frac{\beta}{2T_p}t^2\right)}, \quad (2.1)$$

where β is the bandwidth of the waveform in hertz and T_p is the pulse duration in seconds.

There are many different classes of radar waveforms, each with characteristics that provide performance trade-offs for different operating modes of the radar, and waveform design and optimization is still an active field of research. The LFM was selected as a baseline because it has very well-understood characteristics and is widely used as an introductory waveform [22]. It is worth noting that the basics of pulse-Doppler processing are independent of the waveform used by the system.

The complex baseband waveform is mixed up to the RF frequency, F_0 , and is transmitted by the system. The transmitted pulses reflect off of objects in the scene, and energy is received back at the radar system, where the signal is mixed down from the RF frequency back to the complex baseband. The received baseband radar signal is a scaled, time delayed copy of the transmitted waveform and is given as

$$y_R(t, m) = \alpha(m)x(t - \tau(m))e^{j2\pi F_0\tau(m)} . \quad (2.2)$$

Here, $\alpha(m)$ is the pulse-dependent scaling factor given by the radar range equation defined by [31]

$$\alpha(m) = \frac{G_t G_r \lambda^2 \sigma}{(4\pi)^3 R(m)^4 k T B L} , \quad (2.3)$$

where G_t and G_r is the gain of the transmit and receive antenna, λ is the wavelength of the center frequency, σ is the radar cross section of the target, R is the range to the target, k is Boltzman's constant, T is the system temperature, B is the receiver bandwidth, and L represents the system losses. The time delay $\tau(m)$ is the two way time delay to the target on the m^{th} pulse in the CPI and is given by

$$\tau(m) = \frac{2R(m)}{c} . \quad (2.4)$$

The stop-and-hop assumption assumes an instantaneous range measurement of each target at the pulse transmission time. When using the stop-and-hop assumption, the target is stationary while the pulse interacts with it and moves instantaneously to the next measurement location [22]. Since the target is not moving while the pulse is interacting with it, the pulse will not have any distortions due to the target's motion. The Doppler response, therefore, is only modeled by the phase rotation caused by the continually moving target, and an independent Doppler phase term is not modeled.

At this stage, the pulse return has some duration T_p , which can have a potentially large range extent. The matched filter is used to maximize the signal-to-noise ratio (SNR) of the target return by 'compressing' the long pulse and focusing the energy into a narrow range response [22]. The matched filter for any given waveform is the time-reversed, conjugate copy of the waveform given as

$$z(r, m) = y_R(t, m) * x^*(-t) , \quad (2.5)$$

where $y_R(t, m)$ is the received data for the m^{th} pulse. The matched filter response is the autocorrelation function of the waveform, centered at the target's range. For an LFM, the autocorrelation is approximately sinc shaped. The matched filter cancels all of the waveform phase leaving only the phase caused by the two way time delay to the target. Therefore, the matched filter response can be given as

$$z(r, m) = \alpha(m)\rho_x(r, m)e^{-j2\pi F_0\tau(m)} . \quad (2.6)$$

The waveforms are transmitted as some constant rate, called the pulse repetition frequency (PRF). The PRF and PRI are linked as the PRI is the inverse of the PRF. The duration of the pulse will, at times, be referred to as a percentage of the PRI, this percentage is called the duty cycle. The ADC samples of PRI are collected and

together are called the fast time dimension. When stacking the matched filtered PRIs together, the cross pulse dimension is called the slow time. The residual phase caused by the time delay causes a slow phase rotation in slow time as the target moves from pulse to pulse.

If the target stays in the same range bin for the duration of the CPI, then the discrete Fourier transform (DFT), typically implemented as a fast Fourier transform (FFT), can be applied across the slow time pulses to extract a Doppler frequency from the phase rotation. The relative velocity of a target creates a Doppler frequency shift that depends on the carrier frequency of the radar pulse and is related by

$$v = \frac{cF_D}{2F_0} , \quad (2.7)$$

where F_D is the Doppler frequency caused by the phase rotation and F_0 is the center frequency. The FFT is the matched filter across slow time when using identical pulses, and the Doppler response has a similar sinc like shape as the fast time matched filter response of the LFM.

After applying the matched filter in fast time and then the FFT in slow time, the result is a 2D image that represents range and Doppler. This image is called a range-Doppler map or a range-Doppler plot. Figure 2.1 shows a reference range-Doppler map without noise for a target at 40 kilometers traveling at 27 meters per second. Using traditional range-Doppler processing, the point spread function of the target is a sinc shape in each dimension.

For range-Doppler processing using uniform PRIs, there is the notion of unambiguous range and Doppler. For range, any target whose two-way range delay exceeds the PRI will be ambiguous because the target return occurs after the next pulse is transmitted. In this situation, the system cannot distinguish between the long-range

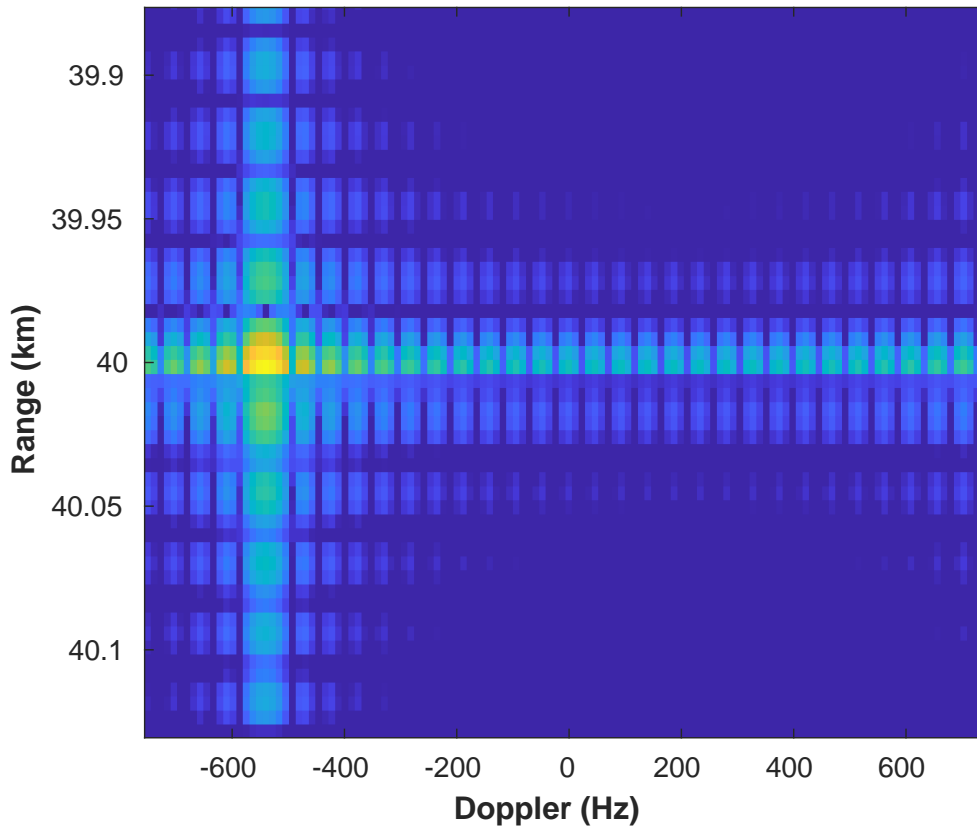


Figure 2.1: A reference range-Doppler map containing a target at a range of 40 km traveling at 27 m/s when using a 3 GHz center frequency.

target return from pulse one and a near-range target return from pulse two. Similarly, If the target's motion creates a Doppler frequency that exceeds $\pm \frac{PRF}{2}$, then the phase rotation aliases, meaning that the phase rotation is completing one or more full rotations around the unit circle between slow time samples, and the true Doppler cannot be determined.

2.1.1 Simulation

Throughout this work a simulation is used to investigate various behaviors and test the proposed techniques. In order to carry out this analysis the point spread function (PSF) is used to identify artifacts of the processing techniques in question. PSFs are

ideal range-Doppler map responses generated for a single target in the absence of noise. Unless noted otherwise, the radar simulation used in this work will utilize the parameters provided in Table 2.1 to simulate a pulse-Doppler radar with a single target. Target locations will be selected so that they do not fall on the sample grid in range or Doppler to provide the most informative performance.

Table 2.1: Simulation Parameters

CPI Size, M	256 Pulses
Sample Time, T_s	10 ns
RF Frequency, F_0	5 GHz
PRF	5 kHz
Pulse Duration, T_p	30 μs
Constant Pulse Bandwidth, β	16.211 MHz

2.1.2 Experimental Testbed

The techniques presented in Chapters 3, 4, and 5 are developed from the signal models used for the simulation. When developing from signal models, the actual success of the algorithm depends heavily on the correctness of the model. To verify the techniques, a software-defined radio (SDR) testbed and outdoor test experiment was used to verify the performance shown by the simulated results. An Ettus USRP X310 software-defined radio (SDR) was used as a radar testbed with two Ubiquity airMAX 5 GHz, 30 dBi RocketDish LW Antennas. Two CPIs of pulses were interleaved together: constant pulses without agility to provide a baseline response, and a CPI of the agility under test which is used to verify the processing approach. The radar testbed was operated in simultaneous transmit and receive (STAR) mode with the same parameters described in Table 2.1 except with a CPI of 128 pulses.

Figure 2.2 shows the test scenario setup outside of the Radar Innovations Lab (RIL)

on the campus of the University of Oklahoma. The experimental system was setup at the end of the road and radar captures were taken as a pickup truck drove away from the system. To reiterate, the spectrum was free of interference for the experiment, but for the frequency agile case the radar selected center frequencies according a sample set of real-world experimental data.

2.2 CFAR Detection

This work will use detection performance as a metric to analyze the quality of interference cancellation, so it is important to provide a brief overview of the technique that is used. For detection, two measurement distributions are considered: one distribution represents the null hypothesis from which noise-only measurements are drawn, and the second distribution represents the measurements that contain a target and noise [22]. The goal of detection algorithms is to correctly identify the source distribution of each sample.

When applying radar signal processing, the goal is to focus the target energy and provide separation between the target present and noise-only distributions. With these two distinct distributions, a threshold can be used to classify each measurement as noise or target plus noise. Figure 2.3 shows the noise only and target plus noise distributions separated by a detection threshold. The region under the noise curve to the right of the threshold is the probability of a false alarm (P_{FA}); these are noise-only values greater than the detection threshold. Similarly, the region under the target plus noise, to the left of the threshold is the probability of missed detection.

CFAR detectors, as the name suggests, sets the threshold such that the false alarm rate is always constant. Keeping the P_{FA} consistent helps determine performance metrics for systems upstream of the detector, like trackers. For a CFAR threshold with



Figure 2.2: Radar test scene.

a known Gaussian null distribution, the threshold, T , can be set as [22]

$$T = -\sigma_0^2 \ln(P_{FA}) , \quad (2.8)$$

where σ_0^2 is the noise power. However, the noise distribution is unknown *a priori* and must be estimated. For this work, cell averaging CFAR (CA-CFAR) will be used for detection.

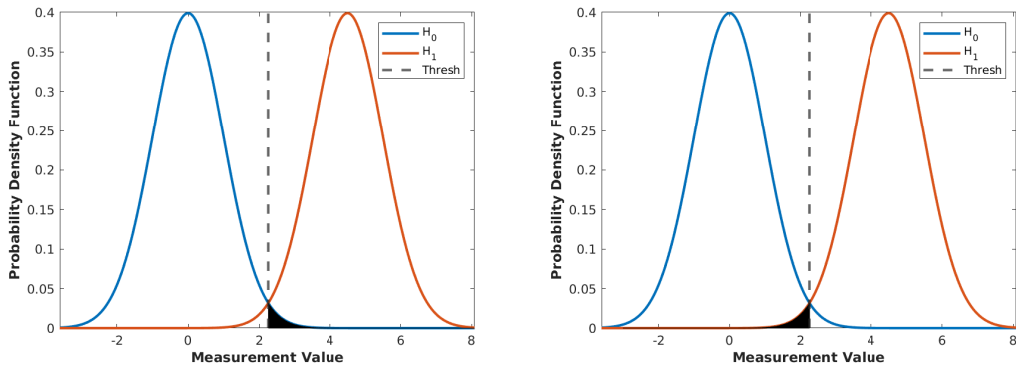


Figure 2.3: The noise distribution and target plus noise distribution split with a detection threshold. The probability of false alarm (left) and probability of missed detection (right) are shown.




Figure 2.4: Guard and reference cell placement relative to the cell under test in CA-CFAR.

It can be shown that the maximum likelihood estimate of σ_0^2 can be found by taking the average of N_R noise only, training samples [22]. CA-CFAR takes advantage of this and averages cells adjacent to the cell under test. However, as discussed in a previous section, targets have sidelobes in the range-Doppler map. To avoid sidelobes contaminating the estimate of the noise distribution, guard cells provide separation from the cell under test and the reference cells. Figure 2.4 shows how the guard and reference cells are placed around the cell under test.

Each noise estimate is only valid for one cell under test and a new threshold value needs to be calculated for each cell. Since the quality of the noise estimate is directly related to the number of cells used, the new threshold will also rely on the number of

reference cells. The CA-CFAR threshold is given as

$$T = N_R \left(P_{FA}^{-1/N_R} - 1 \right) \hat{\sigma}_0^2 \quad (2.9)$$

where $\hat{\sigma}_0^2$ is the noise power estimate. The CA-CFAR threshold equation will be used for all detection performance metrics in this work.

2.3 OFDM Communications

OFDM has become a very popular signal type because of its overall efficiency and flexibility. It is the primary signal type used in modern WiFi, LTE, and 5G networks, and recently, work has shown that it is effective for other applications such as radar [32] and RCS extraction [33]. OFDM and the communication structures will be presented in this section as it is the interference source that is used in chapter 6.

OFDM signals are made of two main pieces: the symbol and the cyclic prefix. A single OFDM symbol can be defined in the continuous time domain using the inverse discrete Fourier Transform (IDFT) [34]

$$x_c(t) = \frac{\text{rect}\left(\frac{t-T_c}{T_c}\right)}{\sqrt{N_{sub}}} \sum_{k=0}^{N_{sub}-1} X_k e^{j2\pi kt/T_c} \quad (2.10)$$

where T_c is the symbol duration in seconds. Each of the N_{sub} subcarriers is modulated with a complex valued symbol, X_k , which is drawn from a constellation of points based on the modulation order, that is the number of bits per subcarrier. Not all subcarriers are modulated with data, LTE and 5G define some number of guard subcarriers which are the outer most subcarriers that do not receive data modulation [35] [36]. The guard subcarriers are defined to provide frequency separation from adjacent users of the spectrum.

The cyclic prefix provides a guard interval in time that prevents the multipath effects of the channel from mixing symbol data together creating inter-symbol interference (ISI) [35]. The cyclic prefix can take on many different lengths based on the length of the free space channel. The cyclic prefix is a copy of the last N_{CP} samples of the symbol appended to the beginning. This cyclic prefix does not carry any data and is stripped away before symbol demodulation.

OFDM symbol generation is highly efficient using the FFT and inverse FFT (IFFT). A vector of complex modulation values represents the subcarrier values in the frequency domain. Taking the IFFT of the subcarrier vector creates the time domain signal. Finally, the last N_{CP} samples from the end of the symbol are prepended to the front of the symbol create the cyclic prefix.

Symbol generation is shown in Figure 2.5 for a symbol with 128 subcarriers, of which 52 are guard subcarriers, generated using a 64 quadrature amplitude modulation (64-QAM) constellation which encodes 6 bits of data per subcarrier. Adding the cyclic prefix after the IFFT is trivial. For the symbol generated in Figure 2.5, the 4G-LTE standard would allow nine samples to create a short cyclic prefix, ten samples to create a normal cyclic prefix, or 32 samples to create an extended cyclic prefix [37].

This work generates OFDM symbols using uniform randomly distributed data bits modulated onto symbols with the parameters shown in Table 2.2. Using random data bits ensures that no particular structure is present in the symbols that skew performance. Additionally, random data bits allow for extension to Monte Carlo simulations without worrying about duplicate data generating identical symbols.




Figure 2.5: The IFFT used to generate the time domain OFDM signal from subcarrier.

Table 2.2: OFDM Parameters

Subcarrier Spacing	15 kHz
Subcarriers	1200
Sample time	32.55 ns
Symbol Duration	66.67 μ s
Cyclic Prefix Duration	4.7 μ s
3 dB bandwidth	18 MHz
Modulation	64-QAM
Samples/Symbol	2048
Samples/Cyclic Prefix	144

2.3.1 Capacity

For digital modulation, various constellations exist, each having a different number of points. With more points, each constellation point can represent more bits. Figure 2.6 shows constellations for quadrature phase shift keying (QPSK), 16-QAM, and 64-QAM constellations. QPSK has four constellation points and a modulation order of 2 and, therefore, supports $\log_2(4)$ bits per symbol. The most significant bit addresses the real part of the complex value, and the least significant bit addresses the imaginary component. Similarly, 16-QAM has a modulation order of 4, with the most significant two bits addressing the real part and the least significant bits addressing the imaginary component. This pattern continues to ever higher modulation orders of 64-QAM (also shown in Figure 2.6) and even 1024-QAM used in 802.11ax [38].




Figure 2.6: Symbol constellations for QPSK, 16-QAM, and 64-QAM.

When noise is added, variance is added to each sample in both time and frequency. As more points are added to the constellation, the points become closer together and the values must vary less, or they will be misidentified. This is directly related to the noise power and the capacity of the channel. The limit of a channel's capacity is described by the Shannon bound, which is given as [39]

$$C = B \log_2 \left(1 + \frac{S}{N_{noise}} \right) , \quad (2.11)$$

where B is the bandwidth, S is the signal power, and N_{noise} is the noise power. From this equation, the network's desire for more spectrum becomes clear. For a fixed SNR,

more spectrum is required to boost the total information capacity of the network. In an ideal world, SNR would be boosted infinitely as capacity increases logarithmically with SNR but only linearly with bandwidth; however, the limits of SNR are often set by physical constraints such as max transmit power and receiver noise figure.

When demodulating, the value on each subcarrier is mapped to the nearest constellation point. As the SNR decreases, the relative variance of each point increases, and if the modulation order is too high, the data values will migrate between constellation points and be decoded incorrectly. So, the use of modulation order is directly limited by the SNR at the receiver as described by (2.11).

2.3.2 Pilots and Structure

Other effects occur when the signal is transmitted through free space. The direct path between the transmitter and receiver is the first signal to arrive, but other multipath copies of the transmitted signal are also received. These multipath effects are caused by the propagation channel and are modeled as a convolution in time. That is, the receive signal becomes

$$y_c(t) = x_c(t) * h(t) , \quad (2.12)$$

where $h(t)$ is the multipath channel response, $x_c(t)$ is the transmitted communications waveform and $y_c(t)$ is the received signal. The channel is represented by some number of taps, each with a delay and magnitude. Figure 2.7 shows the reference vehicular channel response defined by the 3GPP standards consortium, sampled at the rate of the communications signal [40]. Here, the seven peaks correspond to the seven taps.

Because this channel is applied as a time domain convolution, it can be modeled as a multiply in the frequency domain. For OFDM, each subcarrier has a different

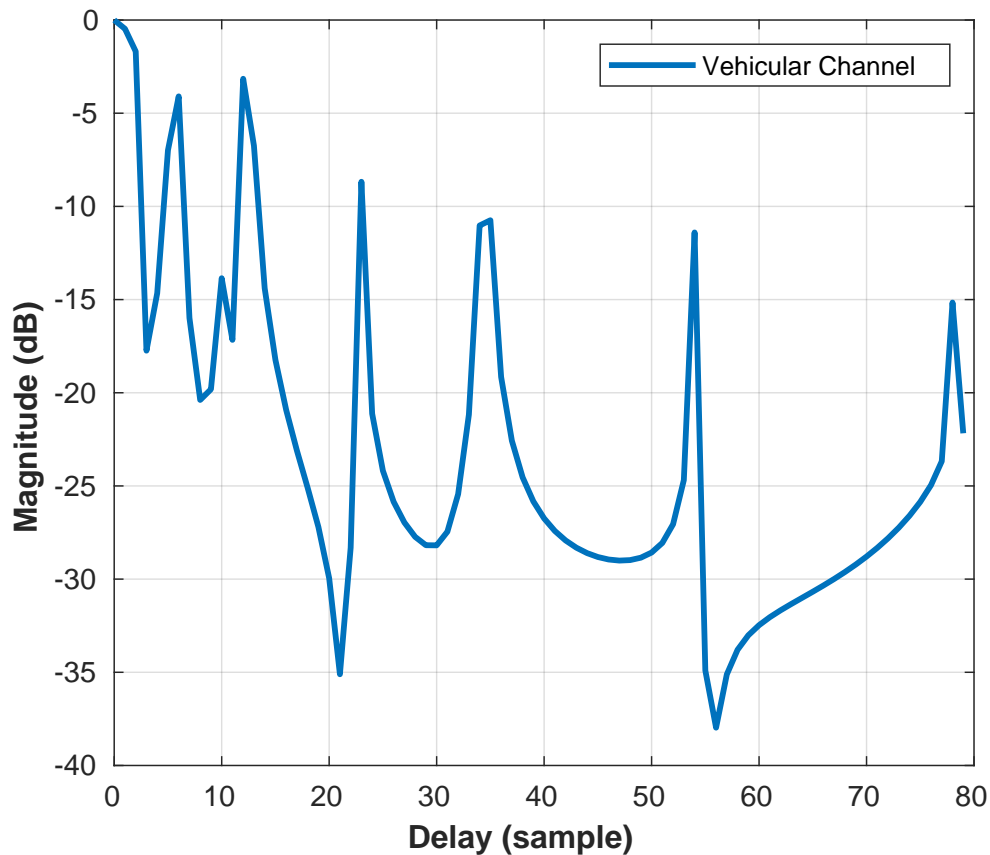


Figure 2.7: The 7 tap vehicular channel sampled at 30.72 MHz.

channel effect that is being multiplied onto it. And if the transmitted subcarrier values were known, an estimate of the channel could be derived. As discussed previously, the cyclic prefix provides a guard interval in time to prevent ISI. However, the channel effects still need to be accounted for to successfully decode data.

To correct these multipath effects, some part of the OFDM signal needs to be known at the receiver. Pilot subcarriers are added to some of the symbols and are known at the transmitter and receiver. The pilots are spaced evenly across the symbol with some subcarrier spacing and some symbol spacing. Since these pilot values are known at the receiver, they can be used to estimate the frequency domain channel values at that subcarrier. Two primary classes exist for the pilot structure: decision-

directed used in WiFi and pilot-assisted used in LTE and 5G [41].

Decision-directed systems send symbols with a pilot on every subcarrier and allows the receiver to estimate the channel across all subcarriers directly. Because OFDM assumes that the channel changes slowly, this symbol containing all pilots are only sent occasionally. Pilot-assisted takes a different approach and spreads a few pilots across the subcarriers of a symbol. From these pilots, the channel can be estimated, and then an interpolation is used to get the channel value of each subcarrier. This work will focus on pilot-assisted OFDM structures.

LTE is defined at a high level by a standard radio frame of 10 ms divided into 10, 1 ms subframes. Each subframe is made up of two 0.5 ms slots. Each slot contains 7 OFDM symbols when using a normal OFDM prefix and 6 OFDM symbols when using an extended prefix [37]. Figure 2.8 shows how a radio frame is divided into symbols.

Because of the slow changing assumption about the channel, channel estimates can be averaged together over the duration of a subframe. In LTE, each slot has two symbols with pilots in them. These two symbols have pilots that are offset in frequency, so both are needed to make a single channel estimate across all of the pilots. For a subframe, the channel estimate can be derived by averaging what is effectively two channel estimates together. Figure 2.9 shows how the pilots are distributed across symbols and subcarriers. For the subframe, two pilots are available for each pilot subcarrier location, one in each slot. The values of these two channel estimates are averaged together before interpolation is used to get a full channel estimate across all subcarriers.

The number of pilots that are used is intentionally minimized. Any subcarrier that is carrying a pilot, is transferring known information and not user data. So, increas-




Figure 2.8: The LTE radio frame structure.

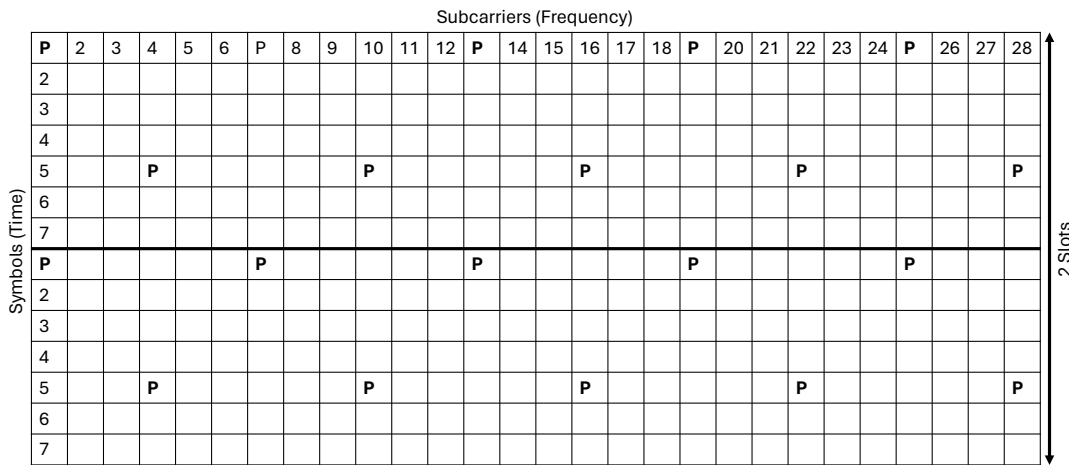


Figure 2.9: Pilots, denoted with a **P**, are distributed across symbols in time and subcarriers in frequency.

ing the number of pilots directly reduces the effective capacity of the network. The values that are assigned to the pilots are pseudorandom sequences called gold codes [42]. These constant modulus codes are useful for their high autocorrelation and low crosscorrelation.

This work will deviate slightly from the LTE pilot configuration described here. The pilot subcarrier spacing will be set at five as in LTE, but the symbol spacing will be set at one, placing pilots on every other symbol. Additionally, the pilots will not be offset in frequency and instead be aligned down the symbol. The pilot structure here

is changed because this work does not assume that the channel is changing slowly. So pilots on the symbols will be used as independent channel estimates that will only be applied to itself and its pilot free neighbors. These changes provide maximum flexibility for processing.

The tools for time synchronization are also buried in the radio frame structure. In LTE, a primary synchronization signal (PSS), a constant modulus Zadoff-Chu sequence, and a secondary synchronization signal (SSS), a constant modulus M sequence, are used to provide precise timing offset [37]. Each of these two signals occur twice in the 10 ms radio frame, in the first and sixth subframes. The PSS and SSS exhibit very high autocorrelation that allows for precise timing synchronization. However, to take advantage of these signals, both instances need to be captured in their full duration. Loss of any part of these signals can limit the user's ability to synchronize. More about this vulnerability will be discussed later in Chapter 6.

2.4 Spectrum Sharing

Spectrum sharing has been of increased interest as a way to use the spectrum more efficiently. As such, many diverse techniques have been proposed that enable dynamic spectrum access. Sharing requires not only sensing but also decision making techniques [43]. Many management strategies have been proposed ranging from simple statistics approaches to intensive machine learning and predictive modeling.

A hierarchical network of agents to control various radar parameters was proposed in [44] that enables decisions to be made at different levels using algorithms and cost functions of various complexity [45]. The cost function designs have been developed that optimize radar operation directly from detection and tracking performance [46]. Predictive machine learning techniques have been developed to predict the type of

emitters that currently use the spectrum [47, 48]. By identifying the type of emitter that is present in the scene, a cognitive radar system can proactively anticipate their use of the spectrum. Predictive schemes are not limited just to machine learning techniques as stochastic approaches have also been presented [49]

Some techniques take a simpler approach, looking directly at the spectrum and making decisions based strictly on power, focusing on optimizing for signal-to-noise ratio (SNR) [14]. With a lower complexity sensing algorithm, an efficient perception action cycle (PAC) can be used to continuously monitor the spectrum and actively make decisions [2, 50]. Previous work expanded on the PAC with hardware optimized, cell averaging estimation (HO-CAE), a greedy dynamic spectrum sensing technique based on CFAR detection [51]. Like CFAR, HO-CAE estimates the noise floor of each spectrum snapshot. Unlike CFAR, the estimate is derived from a set of overlapping averaging windows. An order selection is taken from the ranked values and used as the noise floor estimate to calculate the spectrum occupancy decision threshold. Pairing HO-CAE with the streaming PAC and FSS model allows for real-time dynamic band selection for a radar system.

To provide parameters for a realistic simulation, the HO-CAE/FSS implementation from [16] was deployed to a USRP X310 software-defined radio (SDR) utilizing 100 MHz of front end bandwidth at 2.44 GHz center frequency, collected with a VERT2450 antenna. The spectrum was collected in the Radar Innovations Laboratory at the University of Oklahoma. The spectrum-sharing data was used with an emulated radar system to determine what band selection would be made for each pulse, where HO-CAE/FSS provides the widest contiguous unoccupied band available for the radar to utilize.

Figure 2.10 shows the location of the pulses as they would appear relative to other




Figure 2.10: An example of pulses used for processing overlaid on the spectrum collect the band selections were derived from.




Figure 2.11: The center frequency, $F_H(m)$, and bandwidth, $\beta(m)$ used for the agile cases.

users of the spectrum, where the pulses appear as horizontal bars across frequency, periodically spaced in time (the y axis) while some of the other users are outlined in red. For all simulation results presented, 256 consecutive pulse locations (i.e., center frequency and bandwidth) were selected. These pulse parameters provide an exemplar case for studying the impact of bandwidth and frequency agility. Figure 2.11 shows the center frequency and bandwidth for each pulse that was derived from this real-world data, where the center frequency is referenced relative to complex baseband at 2.44 GHz.

It was assumed that the extreme congestion of 2.44 GHz would require a significant amount of agility to prevent collisions, and Figure 2.11 indicates this assumption was correct. For the frequency agility studies, the center frequencies shown in top plot of Figure 2.11 were used and the bandwidth was held constant at 16.211 MHz. For the bandwidth agility studies the center frequency is held constant to F_0 and the bandwidth per pulse is given in the bottom plot of Figure 2.11. For the fully agile case both the center frequency and the bandwidth are changed according to Figure 2.11. Unless noted otherwise, the remaining simulation parameters are given in Table 2.1. It is important to note that the measured spectrum was only used to generate the waveform parameters. All simulations and measured results are in clear spectrum, to isolate the impact of interference from the signal processing results.

Chapter 3

Intra-CPI Bandwidth Agility

This chapter examines the consequences of changing the waveform bandwidth on a pulse-to-pulse basis. Here, the model from Section 2.1 holds in all but the initial waveform definition. The changing bandwidth can be defined as

$$x(t, m) = e^{j2\pi\left(\frac{-\beta(m)}{2}t + \frac{\beta(m)}{2T_P}t^2\right)} . \quad (3.1)$$

Similar to (2.1), T_P is the pulse duration, but the bandwidth, $\beta(m)$ is defined as a function of the pulse number.

For this chapter only the bandwidth will vary on a pulse-to-pulse basis, while the center frequency will remain constant. Bandwidth agility does not impact the slow time phase at the target's range. As such, the distortion is not directly caused by the slow time phase, which is extracted using FFT-based Doppler processing, but rather the sidelobe magnitudes of the matched filter output.

The bandwidth of a waveform is inversely related to the range resolution as [22]

$$\Delta R(m) = \frac{c}{2\beta(m)} . \quad (3.2)$$

Therefore, the range resolution is directly related to the width of the mainlobe of the

output of the matched filter: the higher the bandwidth, the narrower the mainlobe. The various pulses used for simulation had a minimum bandwidth of 6.54 MHz and a maximum bandwidth of 53.91 MHz. This variation across the associated pulses is shown in the bottom plot of Figure 2.11. The output of matched filtering of these pulses produces many different mainlobe widths, as shown in Figure 3.1.

When waveforms with different bandwidths are used in a single CPI, the peaks of all waveforms are aligned, but the remaining roll-off of the mainlobe will not align. Effectively, there is additional amplitude modulation where the mainlobe roll-off and sidelobes do not align. RSM is created when the FFT is applied across slow time at the range bins with these mismatched sidelobes. Figure 3.2 shows the simulated effect of bandwidth agility without any correction.

When looking at the RSM created by bandwidth agility, the first course of investigation is to apply a window to the matched filter to attempt to mitigate the RSM. By applying a window to the matched filter, sidelobes can be reduced at the cost of a wider mainlobe [53]. By applying a Blackman window to each pulse before the matched filter, the sidelobes of the response can be reduced and therefore the number of range bins that are impacted by RSM is reduced. Figure 3.3 shows the normalized PSF that results from applying a windowed matched filter to each pulse and using FFT Doppler processing.

The windowed matched filter does limit the range spread by pushing down the sidelobes of the matched filter and, therefore, the RSM that is caused by the misalignment of the sidelobes. Although the effective range of the RSM is limited, the mainlobe is broadened by the window and is misshapen which continues to introduce RSM. The misshapen mainlobe is due to the continued misalignment of the range response of the matched filter. Although the Blackman window broadens the mainlobe




Figure 3.1: Each pulse at the output of the matched filter [52]. ©IEEE

of each matched filter response, the broadening is a function of the range resolution so the higher bandwidth waveforms will continue to have narrower mainlobes than the lower bandwidth waveforms. This chapter investigates methods for replacing the match filter because although windowed matched filtering provides an improvement to limit the severity of the RSM, a more robust solution is needed.

3.1 Mismatched Filtering

Optimized mismatched filtering uses a least-mean-square optimization based on some cost function depending on the desired application. Specifically, mismatched

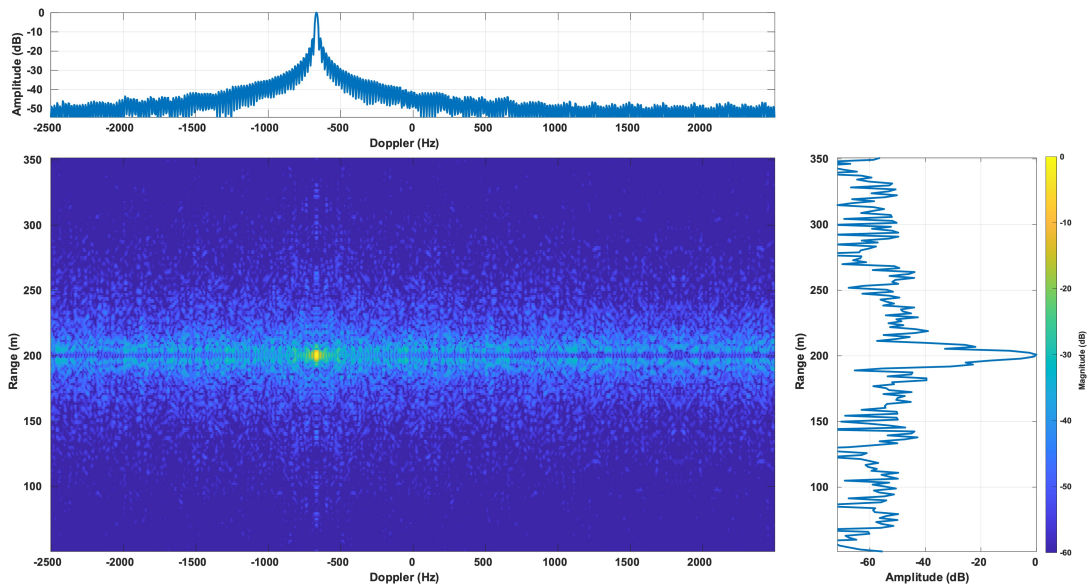


Figure 3.2: The result of using traditional range-Doppler processing on a CPI of bandwidth agile pulses [52]. ©IEEE




Figure 3.3: Normalized PSF resulting from windowed matched filtering and FFT-based Doppler processing [52]. ©IEEE

filtering has been proposed as a technique to enable sidelobe reduction [54]. More recently, the mismatched filtering technique has been proposed to mitigate RSM that

results from using non-identical polyphase-coded FM (PCFM) waveforms within a single CPI [55]. With the mismatched filter technique, a filter is created for each waveform such that the output of the filter has minimized the error relative to some desired output. Here the mismatched filter is used as a way to normalize the response across pulses with varying bandwidths using various templates as the desired output of the mismatched filter.

3.1.1 Formulation

The mismatched filter criterion selected in this paper seeks to minimize the error between the filter output and some template. Minimizing this template error will ensure that the filter outputs will have as common a shape as possible. The cost function to achieve this can be defined as [56]

$$\mathbf{J}_{LS} = \|\mathbf{h}_m \mathbf{A}^H - \mathbf{y}\|^2, \quad (3.3)$$

where \mathbf{y} is the template for the desired output of the filter, and \mathbf{A}_m is constructed as the delay-shift matrix, or convolution matrix, of the m^{th} pulse which takes the form

$$\mathbf{A}_m = \begin{bmatrix} x_1 & 0 & \cdots & 0 \\ \vdots & x_1 & & \vdots \\ x_N & \vdots & \ddots & 0 \\ 0 & x_N & & x_1 \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & 0 & x_N \end{bmatrix} \quad (3.4)$$

\mathbf{A}_m^H is the Hermitian transpose and the dimensions of \mathbf{A} are $((K + 1)N - 1) \times KN$, where K is the over length factor and N is the length of the pulse. Note that for the

purposes of this work while the bandwidth is allowed to change on a pulse-to-pulse basis, the pulse duration and sample time is kept constant.

The mismatched filter, $\hat{\mathbf{h}}_m$, for the m^{th} pulse, can be found via the optimization [55]

$$\hat{\mathbf{h}}_m = (\mathbf{A}_m^H \mathbf{A}_m + \delta \mathbf{I})^{-1} \mathbf{A}_m^H \mathbf{y} . \quad (3.5)$$

Here \mathbf{I} is the identity matrix and δ is a diagonal loading term that prevents the matrix from becoming singular.

To create a filter response with the shape of a given template, the filter needs to be longer than a traditional matched filter. The extra length allow additional degrees of freedom to enable the creation of a filter to match the template shape and reduce sidelobes. The length of the filter is K times the length of the transmitted pulse. It is important to keep K as small as possible, as the size of the matrix A_m increases rapidly and the cost of inverting large matrices is computationally expensive, on the order of $O(n^3)$. Additionally, a longer filter will amplify the noise compared to a shorter filter. With the computation considerations primarily in mind, this work sets $K = 2$, which yields a matrix \mathbf{A}_m of size 8999×6000 for the waveform and system parameters described in Table 2.1.

3.2 Common Matched Responses


When fitting the mismatched filter response to a template, the question naturally arises: what is the best template? The first option considered is the response of the waveform with the smallest bandwidth, corresponding to the worst resolution. Selecting the lowest resolution ensures that all of the mismatched filter's degrees of freedom are used to maximize the filter's match to the template shape.

When creating the template, y , the minimum bandwidth matched filter output is generated by convolving the minimum bandwidth LFM with the time-reversed, complex conjugate of itself. If creating a mismatched filter with $K > 1$, then the matched filter output needs to be padded by $(K - 1)N$ zeros. While the minimum bandwidth template may result in the most uniform output across pulses, it also sacrifices the additional resolution offered by the higher bandwidth pulses. Using this minimum bandwidth matched filter as a template inherently allows the best matching performance, as there are no sidelobe reduction or super resolution constraints. Therefore, all available degrees of freedom for mismatched filter optimization can be used to match the template shape. If this template shape is matched, more traditional window techniques can be used to further reduce sidelobes if so desired.


The second clear option is to use the matched filter response of the maximum bandwidth waveform as the template. In this case, all of the waveforms in the CPI that are not at the maximum bandwidth require filters that super-resolve the mainlobe. The super-resolution requirements of these filters are going to utilize some degrees of freedom. Consequently, when using a filter with the same length as the minimum bandwidth template, the filter output will not fit the template as well. As a result, there is a fundamental tradeoff between resolution and RSM.

Figure 3.4 shows the normalized PSFs that result from a mismatched filter with minimum and maximum bandwidth matched filter templates. Using the lowest resolution template provides a perfect correction to the ideal PSF with the minimum bandwidth resolution. The minimum mismatched filter template wastes the additional bandwidth available in other pulses and, therefore, transmitting the wider bandwidth waveforms would provide no additional resolution.

Looking at the maximum bandwidth result in Figure 3.4 shows that when using



(a)



(b)

Figure 3.4: Normalized PSF generated using the mismatched filter with minimum bandwidth template, (a), and the maximum bandwidth template, (b), both with FFT Doppler processing [52]. ©IEEE

$K = 2$ there are not enough degrees of freedom to enforce the shaping requirements of the template and extract the super-resolution performance from the pulses that have

a narrower bandwidth. As a result, significant RSM is generated from the increased error between the filter output and the template. To improve this performance, a larger value for K could be used to increase the size of the filter and provide the filter with more degrees of freedom. However, as previously discussed, increasing the length of the filter is very computationally expensive due to the matrix inverse that is required. Therefore, it is advantageous to create a template that can enforce super-resolution without significantly increasing the size of the filter.

To improve the resolution of the target response beyond the minimum bandwidth resolution, super-resolution techniques should be explored directly. A super-resolution template can be constructed using a vector of zeros with a one in the central element given as [57]

$$\mathbf{y}_{SR} = \begin{bmatrix} 0 & \cdots & 0 & 1 & 0 & \cdots & 0 \end{bmatrix} . \quad (3.6)$$

Using the simple super-resolution vector allows all degrees of freedom to be used on super-resolving the narrow bandwidth pulses instead of trying to force a particular shape. By setting the desired sidelobes to zero in the super resolution template definition, the RSM will correspondingly be reduced. Figure 3.5 shows the normalized PSF that is generated using the super-resolution vector.

Figure 3.6 shows the output of a mismatched filter using a super-resolution template. Because there is not a desired output shape, the super-resolution vector provides a performance middle ground. Super-resolution provides better resolution performance than the minimum bandwidth template at the cost of some RSM confined to the mainlobe, and better RSM performance relative to the maximum bandwidth template at the cost of some resolution.




Figure 3.5: Normalized PSF generated using the mismatched filter with super-resolution template with FFT Doppler processing [52]. ©IEEE

3.3 Quantitative Results

Other significant parameters need to be considered with the mismatched filter results. The null-to-null mainlobe width, W_{NN} , is important for understanding the achieved resolution with each technique. Additionally, the SNR achieved when noise is present is critical, specifically when comparing against the performance of the windowed matched filter. To measure the SNR, the noise power was compared to the peak of a target return in simulation data. Measuring the noise power at ranges far from the target ensure that the processing of the waveform is not mixed in. Comparing the noise value to the measured target peak effectively provides the *output* SNR of the processing technique on the measured data.

The performance of waveforms can be measured by the integrated sidelobe level (ISL) or the peak sidelobe level (PSL) of their autocorrelation function or ambiguity function [58]. Here the PSF is examined instead of the ambiguity function. Therefore,




Figure 3.6: The output of the mismatched filter using the super-resolution template [52]. ©IEEE

the metric of interest is the ratio of sidelobe energy when using bandwidth agile pulses to the sidelobe energy present when using constant pulses.

First, the total energy in the PSF is defined as

$$E_{tot} = \int_{-\frac{PRF}{2}}^{\frac{PRF}{2}} \int_{R_L}^{R_U} |Z(r, m)|^2 dr dF_d , \quad (3.7)$$

where R_U is the upper bound of range given as

$$R_U = \frac{\left(\frac{1}{PRF} - T_s\right) c}{2} \quad (3.8)$$

and T_s is the sample time. R_L is the minimum range which is less than zero because of the matched filter and defined as

$$R_L = -\frac{T_s \left(\left(\frac{N+N_h}{2} \right) - 1 \right)}{2}, \quad (3.9)$$

where $N = T_p/T_s$ is the sampled waveform length and N_h is the filter length. If a matched filter is used, then $N = N_h$. However, this general definition is adopted here to allow for the use of mismatched filters. The energy in the mainlobe is defined as

$$E_{ML} = \int_{D_1}^{D_2} \int_{\hat{R}_1}^{\hat{R}_2} |Z(r, m)|^2 dr dF_d, \quad (3.10)$$

where D_1 and D_2 are the Doppler bounds of the mainlobe, and \hat{R}_1 and \hat{R}_2 are the range bounds of the mainlobe. These range bounds will vary based on the signal processing used. The relative sidelobe energy growth, RSE , is finally given by

$$RSE = \frac{E_{Tot} - E_{ML}}{E_{Tot}^C - E_{ML}^C}, \quad (3.11)$$

where the superscript C denotes the energy of the PSF derived from the use of constant pulses. Here all quantities are expressed in linear units, but it is useful to define RSE in decibels. As such RSE will be expressed in dB using the $10\log_{10}$ relationship.

So, the sidelobe energy relative to the PSF of a CPI of constant maximum bandwidth pulses and constant minimum bandwidth pulses is used to understand the sidelobe performance and RSM mitigation of each template. As shown in Section 2.4 a minimum bandwidth of 6.54 MHz and a maximum bandwidth of 53.91 MHz was found in the measured spectrum. The desired performance is expected to have more sidelobe energy relative to the constant maximum bandwidth case but less energy than

the minimum bandwidth case. The sidelobe energy will be expressed as relative metrics, RSE_{min} and RSE_{max} , using these constant bandwidth values in (3.11).

Because the sidelobe energy metrics rely on the relative magnitude, RSE_{min} and RSE_{max} values are calculated without normalizing the peak magnitude of the PSFs. Additional considerations need to be made in the case of the super-resolution mismatched filter. The minimum and maximum bandwidth mismatched filter templates have peak values that are greater than 0 dB while the super-resolution vector's peak is 0 dB.

The minimum and maximum resolution matched filter templates had a peak value greater than 1, while the super-resolution template has unity magnitude. In order to provide the fair comparison of sidelobes, the super-resolution vector, y_{SR} , was scaled by the max value of the minimum resolution matched filter templates. This elementary scaling of the template vector causes significant noise amplification through the filter. Because of this noise amplification, the scaled super-resolution is only used for sidelobe comparison purposes, otherwise the regular super-resolution template is used.

Table 3.1 shows the performance metrics for the various mismatched filter template types, traditional matched filtering, and windowed matched filter performance. W_{NN} can be used to compare the range resolution of each technique relative to the maximum bandwidth resolution, given by the maximum bandwidth mismatched filter. The minimum resolution is given by the minimum bandwidth mismatched filter, which perfectly matches the response of the lowest bandwidth waveform. It is also important to note that the Doppler resolution at the target range is unaffected, with only a loss of peak power resulting from the bandwidth agility impacting the Doppler domain.

All of these factors need to be balanced, trading off the SNR and null-to-null main-

Table 3.1: Bandwidth Agility Performance

	SNR (dB)	RSE_{min}	RSE_{max}	W_{NN} (bins)
Matched Filter	58.619 dB	1.18 dB	8.31 dB	16
Windowed Matched Filter	56.385 dB	-4.47 dB	2.66 dB	76
Min. Mismatched	51.500 dB	0.00 dB	7.13 dB	32
Max. Mismatched	32.608 dB	-2.99 dB	4.13 dB	10
Scaled Super Res. Mismatched	56.423 dB	-6.00 dB	1.12 dB	24

lobe width (i.e., range resolution). The minimum bandwidth mismatched filter exactly matches the minimum bandwidth constant case showing a perfect reconstruction of the ideal PSF. But, if the system is going to use the smallest bandwidth template, there is no reason to transmit the higher bandwidth pulses. All of the available spectrum needs to be leveraged, including when additional bandwidth is available. Therefore, a technique that leverages the additional bandwidth is desired such that the additional bandwidth available from other pulses is not wasted.

The maximum bandwidth mismatched filter has worse sidelobe performance than the windowed matched filter, but much better resolution. The scaled, super-resolution technique provides the best sidelobe performance and significantly better mainlobe width than the windowed matched filter and the minimum mismatched filter.

The SNR measure is where the maximum bandwidth mismatched filter fails to perform. Because it cannot enforce the shaping criteria, the noise is significantly amplified through the filter and causes a significant SNR reduction. Ultimately, the super-resolution mismatched filter is selected as the filter to be used to mitigate the

bandwidth agility as it provides the best combination of performance metrics.

The trade-off between these techniques can be seen in the filter's response to various bandwidths. Figure 3.7 shows the super-resolution mismatched filter applied to the average bandwidth pulse compared to the minimum bandwidth matched filter output and minimum bandwidth windowed matched filter output. Here, it is clear that the mismatched filter is able to outperform both windowing and matched filtering when compared against the minimum bandwidth case.

But, the average super-resolution mismatched filter should also be compared to the case where only the maximum bandwidth is transmitted to get a complete understanding of the performance. Figure 3.8 shows the super-resolution mismatched filter applied to the average bandwidth pulse compared to the maximum bandwidth matched filter output and maximum bandwidth windowed matched filter output. As expected from the RSE measurements, the average super-resolution response performs worse than either of the maximum bandwidth matched filters. The trade-off allows for a more consistent shape from the super-resolution filter that reduces RSM at the cost of some range resolution, but that is still better than minimum bandwidth range resolution. Note that all plots are generated with the sample time of 10 ns, corresponding to a Nyquist sampling of 100 MHz bandwidth. Consequently, the sidelobes of Figure 3.8 are not as over sampled as the sidelobes in Figure 3.7.

Ultimately, this represents a very complex trade-space that allows for system dependent decision making. Different filters can be used depending on the performance requirements of the system. If the maximum SNR is needed to resolve one very small target, than matched filtering may be needed to extract the absolute maximum SNR. However, if multiple targets are present, the super resolution mismatched can reduce RSM and reduce the likelihood that a large target will mask a neighboring small target.

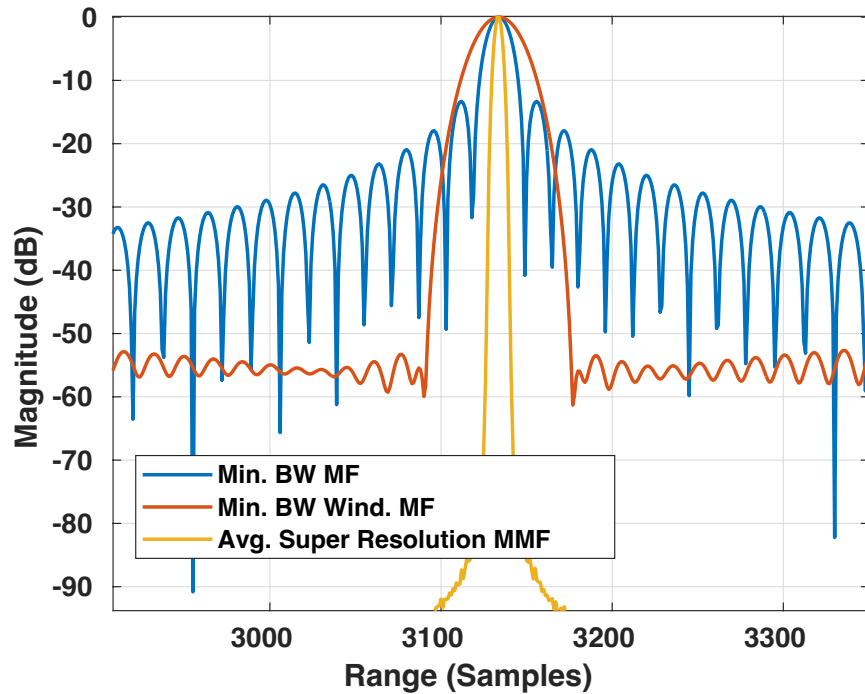


Figure 3.7: The output of the minimum bandwidth matched filter, minimum bandwidth windowed matched filter, and the super-resolution mismatched filter applied to the average bandwidth waveform [52]. ©IEEE

3.4 Experimental Results

Using the experimental testbed described in Section 2.1.2 and a pickup truck as a target, CPIs of bandwidth agile pulses were tested to verify the results on real-world data. Often, optimization-based approaches can become very dependent on the dynamic range of the data. In simulation, the available precision far exceeds what is generally achievable with experimental data. The HO-CAE data set generated a CPI of 128 bandwidth agile pulses. All CPIs were identical, but each had the bandwidth agile data. Figure 3.9 compares data processed with standard matched filtering and the proposed mismatched filtering technique using the super resolution template.

The mismatched filter technique works well on data from the outdoor collection,

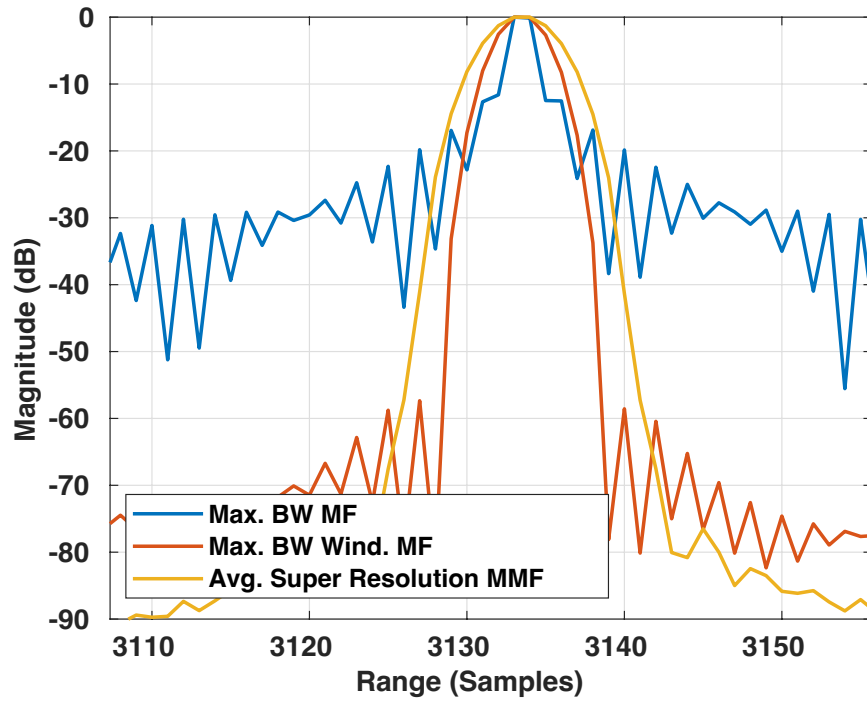


Figure 3.8: The output of the maximum bandwidth matched filter, maximum bandwidth windowed matched filter, and the super-resolution mismatched filter applied to the average bandwidth waveform [52]. ©IEEE

showing that the optimized filters are sufficiently robust to be used with representative data. The experimental data shows the characteristics predicted by the simulation results. RSM is cleaned up across the image, and Doppler sidelobe structures can be identified from the strong clutter returns, indicating the RSM has been suppressed to the noise floor. The cleaner image means smaller targets that might otherwise be masked with standard processing would be more easily distinguished.

This additional performance comes at the cost of 0.5 dB of peak loss measured in this CPI. The total trade-off is well-used. With less than 1 dB of peak loss, the range-Doppler image is significantly cleaner, especially around the strong clutter and direct path returns.

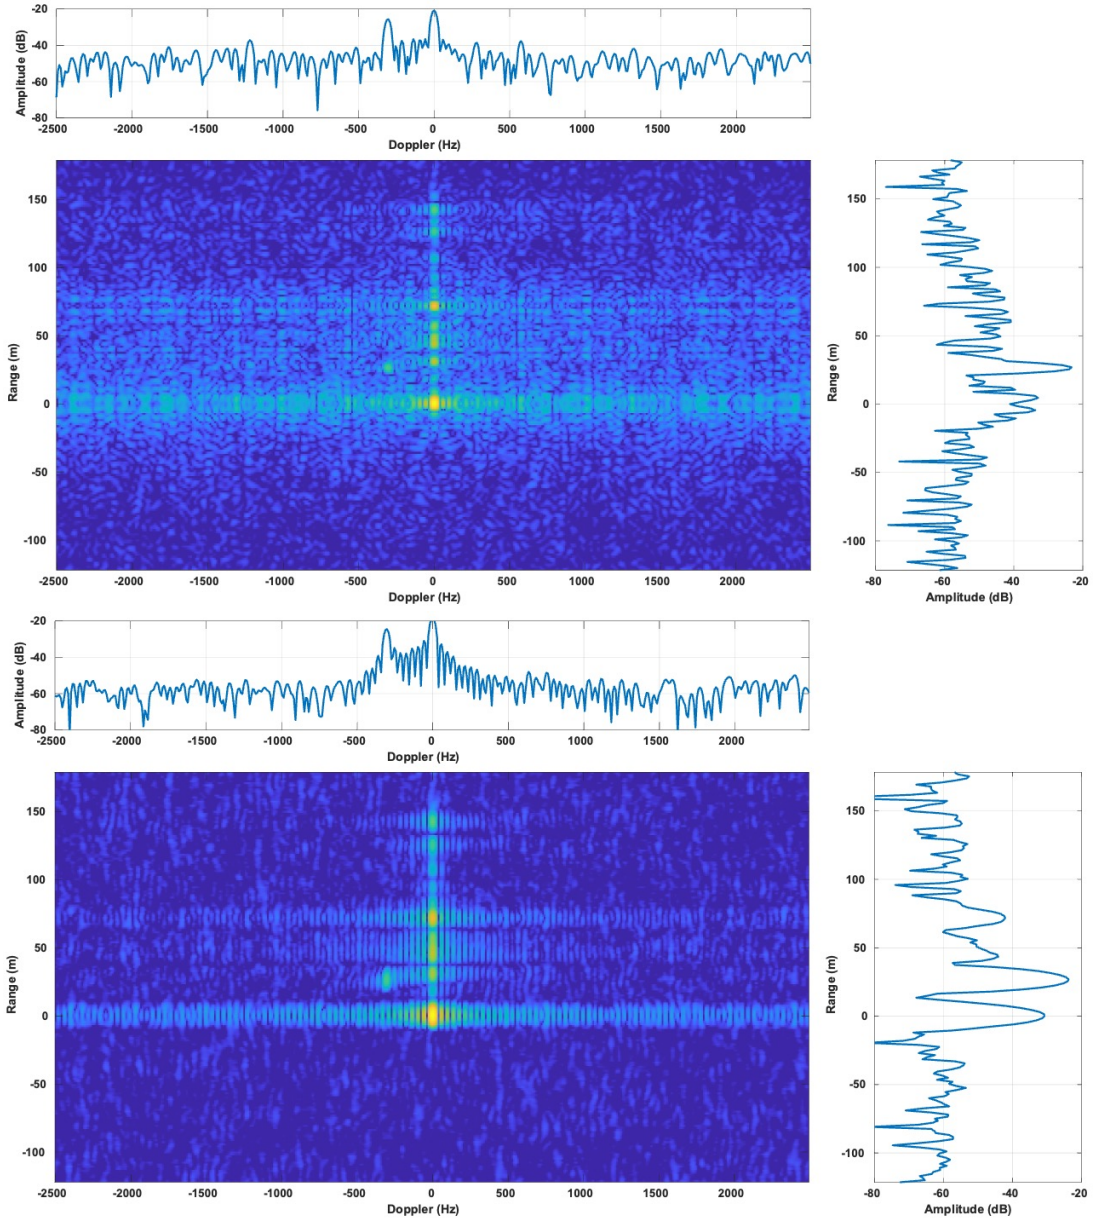


Figure 3.9: Experimental data using bandwidth agile CPIs processed using the matched filter (top) and the super resolution mismatched filter (bottom).

Chapter 4

Intra-CPI Frequency Agility

Now that a technique has been developed to mitigate the bandwidth agility, the focus turns to frequency agile pulses. Unlike bandwidth agile pulse, frequency agile pulses uses the same complex baseband signal for all of the pulses. However, an intermediate frequency step is needed to add the frequency agility.

To generate the frequency agile signal, each pulse from Equation 2.1 is mixed to IF yielding

$$x_{IF}(t, m) = x(t, m)e^{j2\pi F_H(m)t} , \quad (4.1)$$

where $F_H(m)$ is the hopped IF frequency of the m^{th} pulse measured relative to the center frequency. It is important to note that, although it is modeled as mixing up to the IF frequency here, the IF waveform would likely be synthesized directly as it is within the instantaneous bandwidth of the system. In other words, it is assumed that any IF processing is done at digital IF, after sampling. However, for generality the signal model is developed as a continuous time domain signal.

The transmitted signal is achieved by mixing the IF signal to the radio frequency

(RF), F_0 , for transmission as

$$x_{RF}(t, m) = x(t, m)e^{j2\pi F_H(m)t}e^{j2\pi F_0 t} . \quad (4.2)$$

The received signal is a scaled, time-delayed copy of the transmitted signal

$$z_{RF}(t, m) = \alpha(m)x(t - \tau(m), m)e^{j2\pi F_H(m)(t-\tau(m))} . e^{j2\pi F_0(t-\tau(m))} \quad (4.3)$$

where $\alpha(m)$ is the pulse-dependent scaling factor given by the radar range equation give in (2.3) , and the time delay, $\tau(m)$, is the two way time delay to the target on the m^{th} pulse in the CPI.

The receive signal is mixed back down to IF, leaving a residual RF phase term that depends only on target range which yields

$$z_{IF}(t, m) = \alpha(m)x(t - \tau(m), m)e^{j2\pi F_H(m)t} . e^{-j2\pi(F_H(m)+F_0)\tau(m)} . \quad (4.4)$$

Notice, at IF there remains only the IF mixing term and then a term that relies on both the RF carrier and IF term that is related to the target motion at each pulse. The baseband signal can finally be expressed by mixing down the IF frequency resulting in

$$z(t, m) = \alpha(m)x(t - \tau(m), m)e^{-j2\pi(F_H(m)+F_0)\tau(m)} . \quad (4.5)$$

The IF signal defined in (4.4) and baseband signal from (4.5) are both of interest in the investigation of the processing effects and will be examined to understand if the processing should take place at IF or baseband. Pulse compression is applied to the received signals via the matched filter to convert the signal from the fast time domain

to the range domain and is given as function of range, r , and pulse, m , as

$$z_{IF,PC}(r, m) = \alpha(m)\rho_{x_{IF}}(r, m) \cdot e^{-j2\pi(F_0)\tau(m)} \quad (4.6)$$

and

$$z_{PC}(r, m) = \alpha(m)\rho_x(r, m)e^{-j2\pi(F_H(m)+F_0)\tau(m)}, \quad (4.7)$$

where ρ_x is the output of the matched filter of the baseband waveform defined as (2.1) and $\rho_{x_{IF}}$ is the output of the matched filter of the IF waveform defined in (4.1). The Doppler term is expressed by the phase's rate of change due to the target's motion. Here the Doppler term is implicitly represented in $\tau(m)$. To extract the Doppler information the Fourier transform can be applied across slow time at each range bin of the pulsed compressed signal, implemented with the FFT. This process results in the range-Doppler map, or in the case of the simulation the point spread function (PSF), $Z_{IF}(r, F_d)$ and $Z(r, F_d)$, which are each a function of range, r , and Doppler frequency, F_d .

First, an investigation of frequency agility is carried out. From the signal model in the previous section, it is clear that frequency agility will have an impact on the phase of each pulse. Specifically, at the output of the matched filter, there will be an impact on the slow time phase history across pulses that will impact the Doppler response in the range-Doppler image. The phase history will also help to understand the relationship between the Doppler phase term at IF and baseband. To carry out this investigation a radar system was simulated to generate the point spread function using the parameters in Table 2.1 with a simulated target at 200 meters and a radial velocity of 20 meters per second.

4.1 Phase Analysis

The non-agile, slow time phase history should be examined to provide a baseline for analysis. When using a non-agile assumption, the matched filter neutralizes the phase from the waveform and, therefore, using (4.7), the non-agile, slow time phase history from a target can be expressed as

$$\Theta_R = -2\pi F_0 \tau(m). \quad (4.8)$$

Figure 4.1 compares three different phase histories: the fixed (i.e., non-agile) baseband phase history, the IF frequency agile phase history derived from (4.6), and the baseband phase history derived from (4.7). When using the IF matched filter in (4.6), only the ideal phase history term remains. However, because the signal is sampled and not continuous, there exists some phase offset between the sampled time delays and the true delay of the target. The phase offset, dependent on the IF frequency and the true delay of the target can be seen in Figure 4.1 by the slight mismatch between the fixed baseband phase history and the IF phase history.

The baseband phase history deviates from the desired phase history curve more significantly as a function of the IF frequency. Generally, this would mean that the processing should be carried out at IF where the phase history is more closely aligned. There are other practical considerations that make baseband processing a more attractive option. From a practical perspective, by processing at baseband the signal can be further decimated from the IF sample rate down to the bandwidth sample rate. Decimating to the baseband rate can significantly reduce the amount of data that needs to be moved and processed by the system. This can be a key advantage as radar front ends continue to support larger instantaneous IF frequencies and direct RF sampling

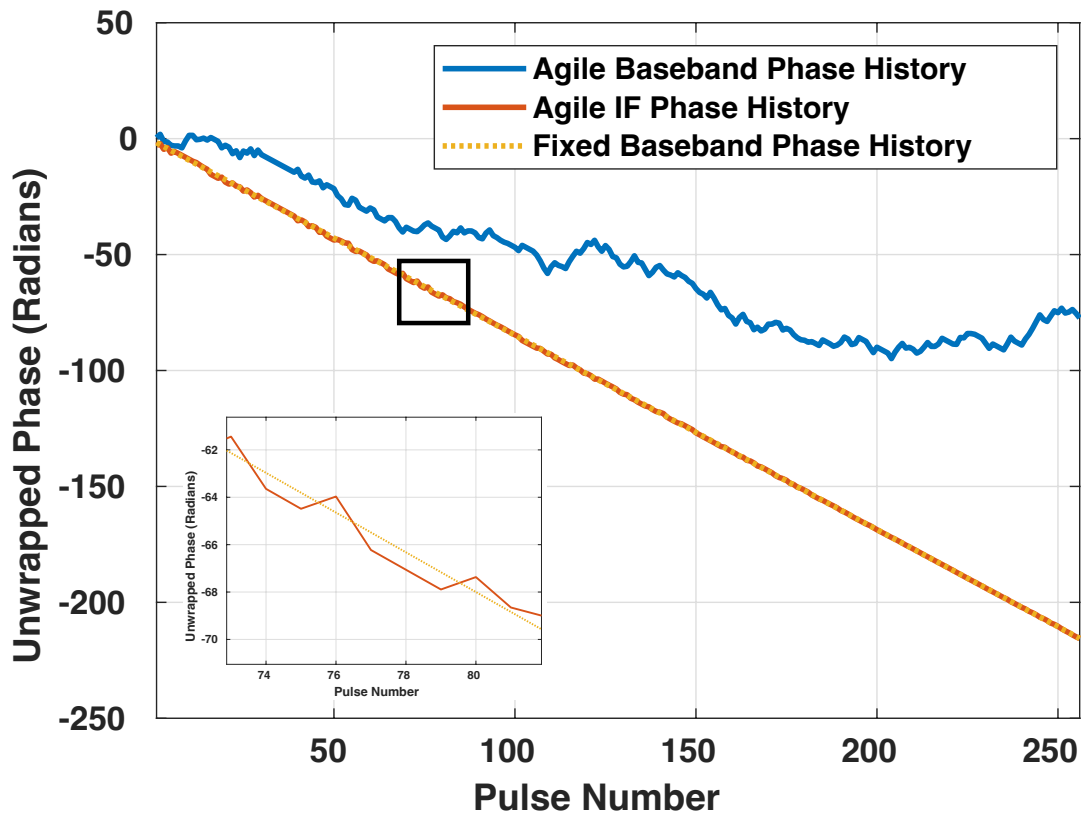


Figure 4.1: Unwrapped agile baseband and agile IF slow time phase history compared to that of a CPI of pulses with no frequency agility with 100 MHz of front end bandwidth [59]. ©IEEE

systems become more popular. This can also have an advantage in digital array radars with many channels.

Similarly, when processing at baseband fewer pulse compression filters need to be stored and calculated than when processing at IF. Baseband processing requires only one filter per bandwidth, while IF processing requires a filter for every bandwidth/IF frequency combination that might be used, which can quickly become impractical for efficient processing. Because of the efficiency benefits for this work the processing is carried out at baseband. This choice increases the importance of correcting the phase offset introduced by frequency agility.




Figure 4.2: Unwrapped agile baseband and agile IF slow time phase history compared to that of a CPI of pulses with no frequency agility with 2 GHz of front end bandwidth [59]. ©IEEE

Generally, the phase performance would indicate that processing and correction should occur at IF. However, because more front end bandwidth would be advantageous to enable additional opportunities for sharing, it is important to investigate how increasing the front end bandwidth relative to the center frequency impacts the behavior of the phase history. As 100 MHz is narrow relative to the carrier frequency of 5 GHz, a 2 GHz front end bandwidth was simulated. The IF frequency of the pulses from the HO-CAE/FSS output were normalized to the new front end bandwidth while the waveform bandwidth was kept the same.

Figure 4.2 shows the resulting phase histories from the 2 GHz bandwidth case.

When the front end bandwidth increases, and the frequency hops increase accordingly, the phase history processed at both IF and baseband diverge significantly from the baseline of constant pulses. Since the behavior of the IF processing changes based on the hopping bandwidth relative to the center frequency, this work will focus on applying the correction when processing at baseband. The importance of correction is increased by choosing to process at baseband but this choice provides more consistent behavior as the available bandwidth for hopping increases.

4.1.1 Phase Correction

The effects of frequency agility on traditional FFT range-Doppler processing at baseband are significant without correction. Figure 4.3 shows the range-Doppler image of the simulated point target at 200 meters with a relative velocity of 20 meters per second. The Doppler and range cuts of the target location are shown at the top




Figure 4.3: Applying traditional FFT based range-Doppler processing to the baseband matched filtered data with the target's range and Doppler cuts shown [59].

©IEEE

and right plots of Figure 4.3, respectively. As shown in the Doppler cut the Doppler information is completely lost due to the effects of the phase residue caused by the baseband frequency hopping. Additionally, the phase distortion is carried throughout the extended range of the matched filter response, smearing across Doppler in every case. Notably, from examination of the range-Doppler plot the distortion presents as a uniform grid, indicating that the nulls in the matched filter output are aligned. Further, it can be seen from the range cut on the plot to the right of Figure 4.3 that the use of a constant bandwidth preserves the pulse compression response of the waveform.

Through examination of the phase term in the baseband radar return after pulse compression from (4.7), a correction term to extract the desired phase history given by (4.8) can be identified as

$$C_m(\tau(m)) = e^{j2\pi F_H(m)\tau(m)} . \quad (4.9)$$

The correction factor relies on τ and although the Doppler effects are implied in this term, it only depends on the range of the target at any given pulse. The output of the matched filter is relative to target range, r . Therefore, this correction factor can be applied directly to each pulse at the output of the matched filter.

4.2 Modified Backprojection Algorithm

The previous section identified a range dependent phase correction that needs to be applied to refocus the target response. Without any knowledge of the target, a technique needs to be developed that will correct for any target in the scene. An algorithm that collects all of the energy for each range bin should be used. A focusing algorithm traditionally used in synthetic aperture radar (SAR), backprojection is an ideal technique for refocusing the target in energy [60].

4.2.1 Backprojection

Although the backprojection algorithm is most commonly found in SAR image processing, the technique has also been proposed as a method for focusing range-Doppler images that contain fast-moving targets [61]. The backprojection algorithm treats every pixel in the range-Doppler map as a unique motion profile with a starting range and velocity corresponding to the location of that pixel in the image. For each pixel in the image, the energy contribution for that motion profile, from each pulse, will be gathered into the final image.

To accomplish this, the algorithm begins by calculating the range to each pixel on the m th pulse. This target range for each pixel at the m th pulse, \hat{R}_m , can be given as

$$\hat{R}_m(r, v) = r + (m - 1)vT_P , \quad (4.10)$$

where r is the range axis and v is the velocity axis of the resulting range-Doppler image.

Once a grid of ranges, \hat{R}_m , has been calculated, then the value from the fast time pulse needs to be interpolated. For this work, interpolation was applied in two steps. First, a high-quality FFT interpolation was used and then a low-quality linear interpolation was used to interpolate to the exact range. After obtaining the interpolated values from the pulse to the grid of range values, a phase correction needs to be applied to focus the energy from a target at that range if it were to exist. The desired phase correction angle, ψ_B , is based on the targets range at that pulse and is given as [61]

$$\psi_B = -2\pi F_0 \frac{2\hat{R}_m(r, v)}{c} . \quad (4.11)$$

The pixel of the range-Doppler image, P , that corresponds to the range, r , and velocity, v , is constructed by interpolating to the hypothetical target's location in each pulse and multiplying by a phase correction that corresponds to that range. Combining the processes for each pixel yields

$$P(r, v) = \sum_{m=1}^M x_{BB}^{Rx} \left(t + \frac{2\hat{R}_m(r, v)}{c}, m \right) e^{j\psi_B} . \quad (4.12)$$

4.2.2 Modified Backprojection

Modifications are made to the phase correction term to enable the processing of the frequency agile pulses [59]. The phase correction in (4.11) takes the same form as the constant frequency slow time phase history given by (4.8), making it an ideal place to implement a correction. Adapting the backprojection factor for the frequency agile pulse case can be extrapolated from the baseband signal given by (4.7) in the signal model.

Adding the hopping frequency to the correction factor yields

$$e^{j\psi_B} e^{j2\pi F_H(m) \frac{2\hat{R}_m(r, v)}{c}} . \quad (4.13)$$

Combining the new terms into the backprojection correction yields the correction phase

$$\psi_{MB} = 2\pi(F_H(m) - F_0) \frac{2\hat{R}_m(r, v)}{c} . \quad (4.14)$$

Using the modified backprojection algorithm by replacing ψ_B with ψ_{MB} in (4.12) successfully sharpens the target and recovers the Doppler information, providing a peak Doppler sidelobe of -14.7 dB while maintaining a strong range response. Figure 4.4 shows the range-Doppler image that results from using the modified backprojec-




Figure 4.4: Result of applying modified backprojection to the baseband match filtered data [59]. ©IEEE

tion method on the frequency agile pulses and the range and Doppler cuts at the target’s true range and Doppler.

There are a few notable features for discussion. The first is the continued presence of RSM. While the correction does focus the target at the correct range, RSM remains. The RSM persists because the phase distortion extends across the length of the matched filter on each side of the target. However, because of the roll-off of the matched filter response the RSM likewise rolls off quickly. Since this correction is applied for every motion profile in the range-Doppler space, regardless of the presence of a target, the sidelobes present in the range bins adjacent to the target bin will decorrelate which accounts for the distortion of the range profile.

Similarly, the mainlobe of the target’s range response is distorted by the decorrelation effect. In the range bins immediately adjacent to the target’s range bin, the phase correction decorrelates with the mainlobe rolloff and sidelobe creating nulls instead of

the smooth mainlobe. The shoulders at -6.2 dB around the peak are a direct result of this decorrelation and technically make the mainlobe narrower. Because of the hopping frequency the pulses have a collectively wider bandwidth, which provides a more narrow mainlobe than would be expected from non-hopped waveforms [23].

Secondly, the Doppler response is affected by this technique. Although the Doppler response is focused, it does not take on the expected shape. Instead of rolling off with a sinc-like shape, as would usually be expected with constant pulses, the Doppler response starts rolling off but flattens below -20 dB, while maintaining a narrow peak at the correct Doppler. The RSE metric from the previous section will be used as a metric to understand the sidelobe performance.

Point spread functions are usually generated at zero range and Doppler, but in this case there is very minimal distortion of a target with zero range and zero velocity




Figure 4.5: The constant pulse mainlobe boundaries in range and Doppler are shown in (a) and (b), respectively. The agile pulse mainlobe boundaries in range and Doppler are shown by (c) and (d) respectively [52]. ©IEEE

because $\tau(m)$ is zero pulse-to-pulse and so it is not impacted by changing center frequency. In order to get a representative metric, an unnormalized version of the PSF that was examined in Figure 4.4 with a target at 200 meters and traveling at 20 meters per second was used. This target location does not fall on the sample grid in either range or Doppler, thus providing an accurate representation of real target performance.

The mainlobe was defined from null-to-null, and the location of the sidelobe boundaries is shown in Figure 4.5. Overall, from (3.11) there was an $RSE = 6.384$ dB increase in sidelobe energy for the agile pulses when compared to a constant bandwidth of 16.211 MHz, as defined in Table 2.1. The increase is primarily caused by the lack of roll-off in the Doppler domain. Additionally, the definition of the mainlobe in the agile case is slightly narrower than the mainlobe in the constant case, expanding the limits of the sidelobe energy calculation. Ultimately, the peak magnitude is unchanged when using the frequency agile pulses, and the PSL is just below the expected value at -13.4 dB.

4.2.3 Range Migration

Backprojection is designed to focus fast moving targets since it works by interpolating to the hypothesized target location. By adapting the backprojection algorithm to compensate for the baseband phase offset fast targets that migrate across range bins can be focused into a single target point.

To compare the FFT and the modified backprojection, the phase correction from (4.9) is applied directly at each range cell at the output of the baseband match filter. An FFT is then taken across each range bin to extract Doppler information. Figure 4.6 shows the result of FFT processing on a target moving at 70 meters per second at 1500 meters. When apply the FFT processing, the phase history is spread across range bins

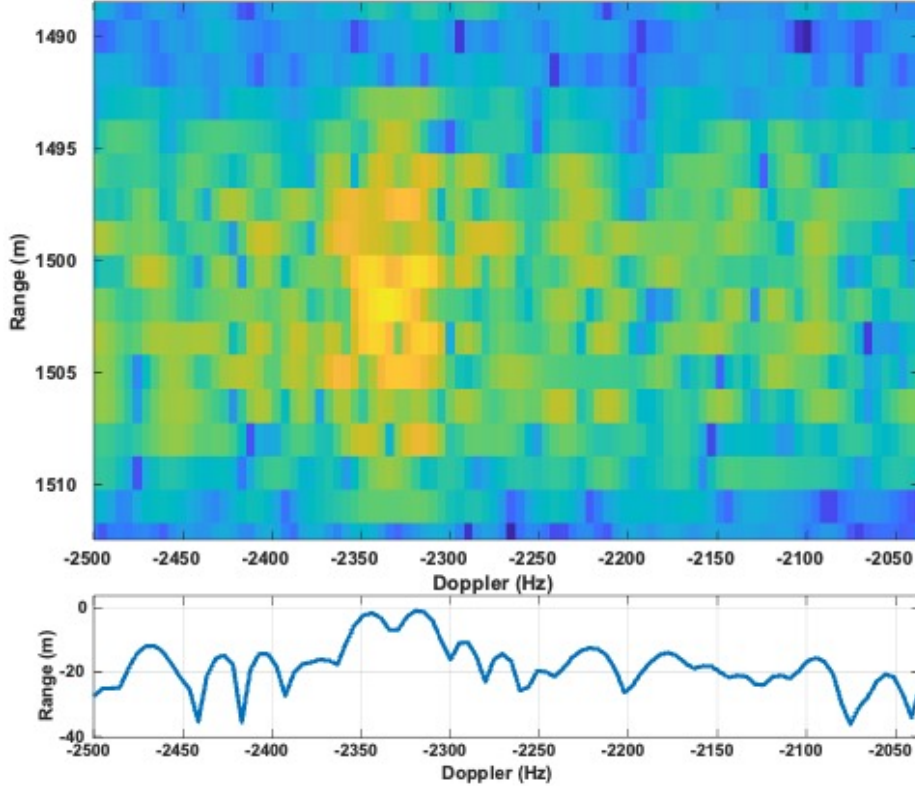


Figure 4.6: Performance of direct correction FFT processing on range migrating targets and a Doppler cut at the target's true range [52]. ©IEEE

which causes target spreading in both range and Doppler.

Figure 4.7 shows the result of applying modified backprojection. With the modified backprojection the target is focused to a single range and Doppler bin and power is recovered compared to the FFT processing. Figure 4.6 and Figure 4.7 are left unnormalized to show the relative power gain from backprojection compared to FFT processing. It is worth noting that the speed of our target is limited by the carrier frequency and PRF, a relatively low PRF of 5 kHz was used with a small number of pulses to keep the simulation comparable to what could be transmitted by our test hardware.

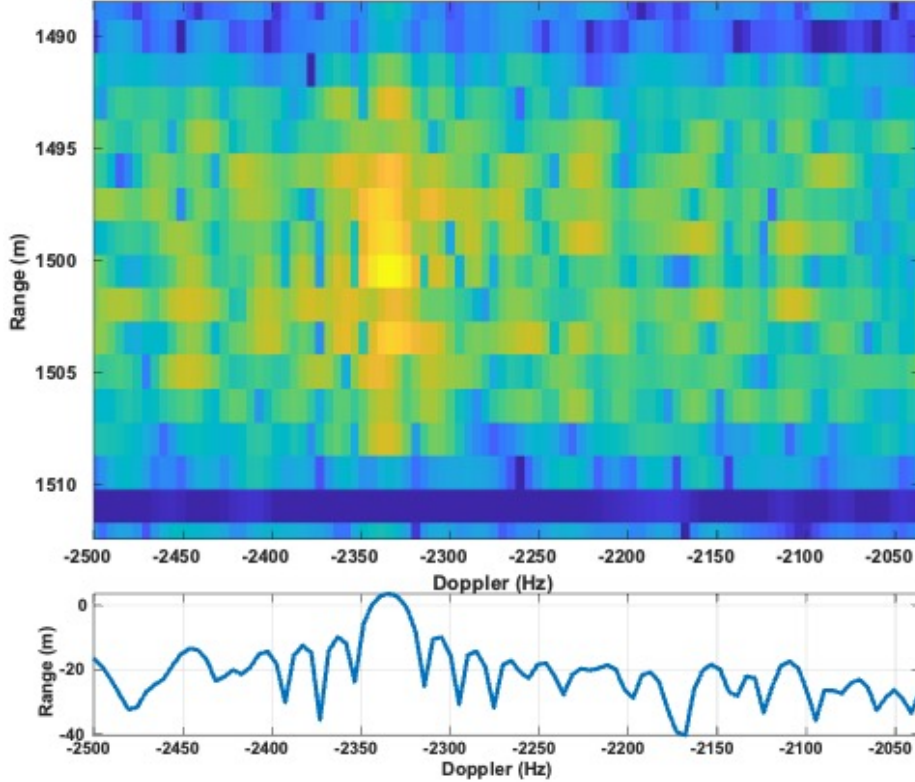


Figure 4.7: Performance of modified backprojection processing on range migrating targets and a Doppler cut at the target's true range [52]. ©IEEE

4.3 Why Backprojection?

As pointed out in the previous section, backprojection is a technique used to compensate for target smearing caused by a fast-moving target that migrates across range cells in a single CPI. The correction factor is only defined in terms of range, so in the nominal case that range migration is not present, backprojection may not seem necessary. But, targets are not stationary for an entire CPI; they are moving, so the correction has a different phase in each pulse based on each motion profile. The error of not accounting for the motion can degrade the Doppler response over time. This section compares backprojection to a direct application of a single correction value,

defined in (4.9), across each range bin. The phase history for a target placed on the range-Doppler grid at 200.86 meters traveling at 20.2008 m/s is examined to show the need for the modified backprojection technique.

Figure 4.8 shows the wrapped phase history of the on-grid target after a correction is applied directly across slow time with a single value. There is a sound agreement between the phase history for the first 20 pulses. However, the direct comparison begins to deviate from the backprojection technique. This target is not moving fast enough to migrate in range, but the error between the correction and target ranges accumulates.

Figure 4.9 shows how the Doppler response changes based on the two correction approaches. First, the reference response is unfocused, as has been shown repeatedly. When applying the direct correction, the target becomes focused, but the sidelobes are significantly elevated and approach the sidelobes of the unfocused target. Finally, for the modified backprojection correction, the Doppler response is perfectly reconstructed for the on-grid target. These curves show a clear improvement of up to 15 dB by using the backprojection over a direct application of the phase correction.

Similarly, the correction relies on the absolute range of the target. The phase error introduced with hopping is more severe with larger hopping frequencies and at longer ranges. The linked dependence between range and the correction mean that an individual augmented DFT matrix could not be used to correct the phase error and would only provide the performance of the direct correction. Again, modified backprojection proves to be necessary to extract the best phase correction performance.

Targets exist in a continuous space that does not fall on a discrete grid, so all simulated results in the previous sections were range and velocity values that fell between the discrete grid. Here, the performance improvement of backprojection over direct

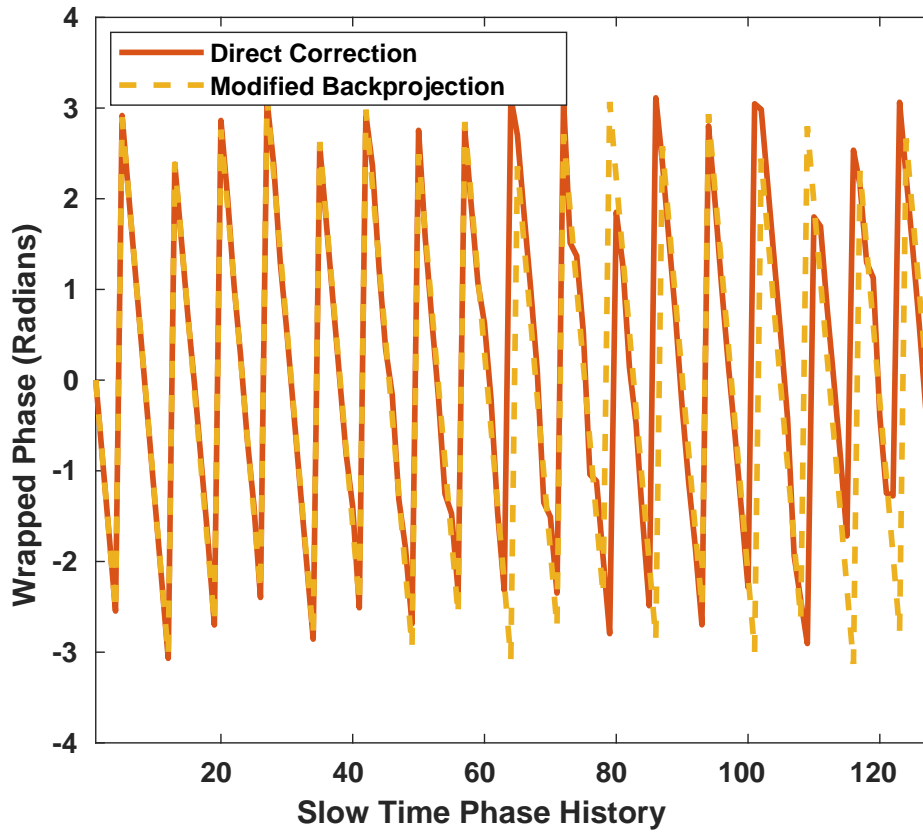


Figure 4.8: Wrapped phase history after direct correction and backprojection correction.

correction was examined using an on-grid target to provide a good reference. As seen before, a target that exists off the grid will not have a perfect Doppler response, but the modified backprojection technique minimizes that error over the other correction method.

4.4 Experimental Results

Using the experimental testbed described in Section 2.1.2, a series of frequency agile pulses were tested to verify the results provided by the simulation. Figure 4.10 shows the real-world data processed with FFT Doppler processing and modified back-

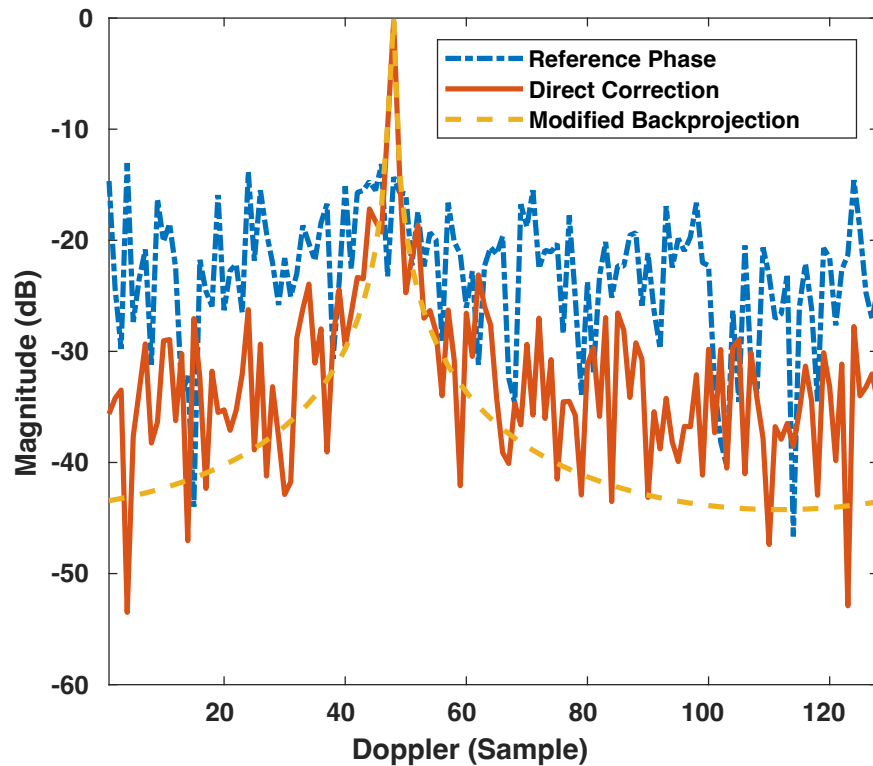


Figure 4.9: Doppler response of uncorrected, directly corrected, and backprojection corrected phase.

projection, both processed with the matched filter. The real-world results reflect an excellent agreement with the simulation and, consequently, the signal model. Before correction the target energy is spread across Doppler with no Doppler peak. Before correction, stationary scatterers are somewhat focused but have significant energy spreading in Doppler.

With correction the vehicle's motion becomes visible in the range-Doppler image moving away from the radar testbed, defined here as negative Doppler. The target is focused in the scene with strong peaks in both range and Doppler. As with simulation, the Doppler sidelobes are flat while the mainlobe in range is affected by decorrelation effects.

Other targets are also focused in the image. The strongest return at zero range and

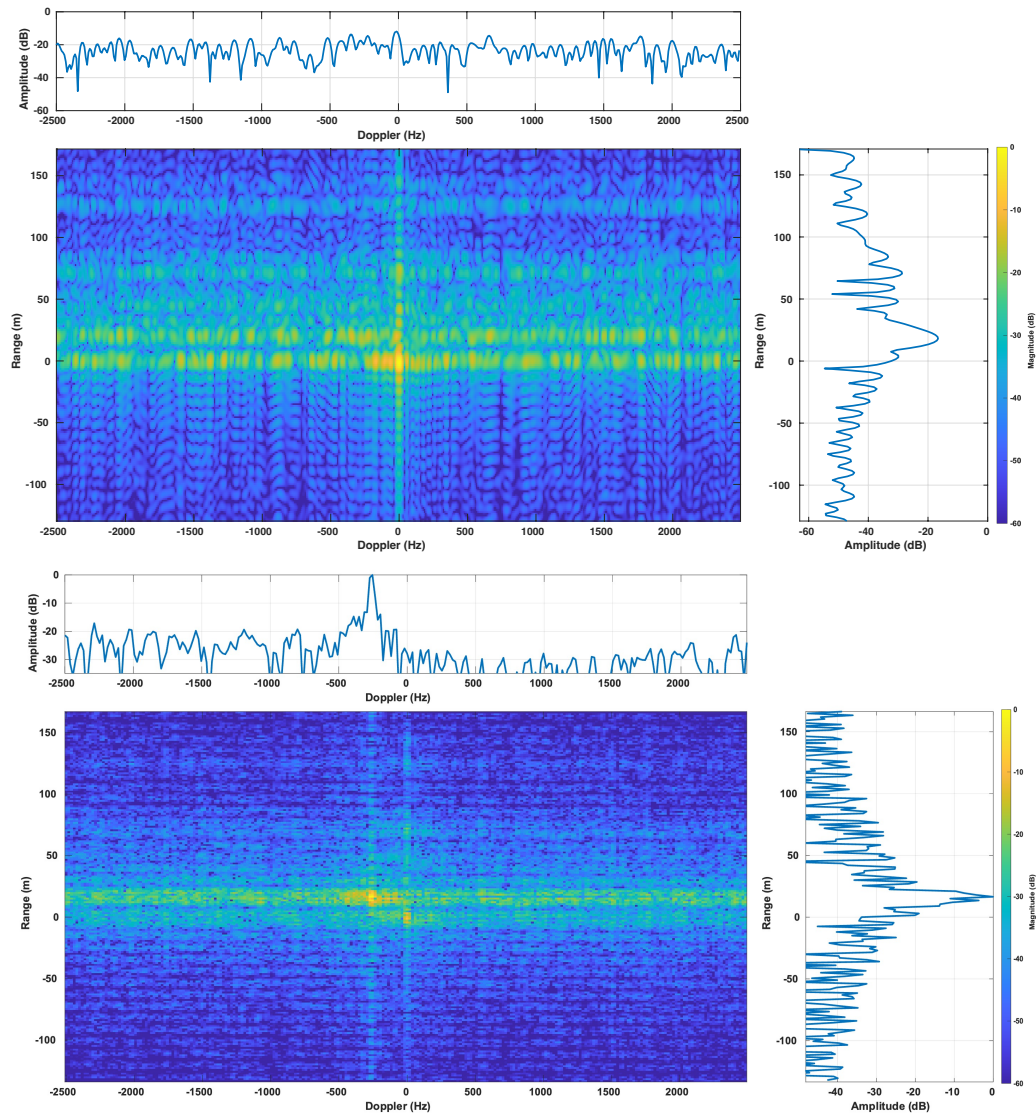


Figure 4.10: CPI of real-world data with traditional FFT range-Doppler processing (top) and with modified backprojection (bottom) applied at the output of the matched filter [52]. ©IEEE

zero Doppler is the direct path antenna coupling return and is the result of the STAR operating mode. As discussed previously, this direct path at zero range and Doppler is mostly focused with and without correction because τ is zero for all pulses. Other stationary scatterers are also present at various ranges at zero Doppler which are also successfully focused using the direct application of the correction.

Chapter 5

Fully Agile Pulse Processing

In the case when the radar is fully agile, that is changing in both frequency and bandwidth from pulse-to-pulse, the phase correction and mismatched filtering techniques need to be combined into a single procedure. First, each received pulse needs to be mixed down to baseband as was discussed in Section 4.1. Then instead of applying a matched filter, a super-resolution mismatched filter is applied as discussed in the previous section. Using the output of the mismatched filters, modified backprojection can be applied to correct the images for the frequency agility and generate the range-Doppler image.

The modified backprojection and mismatched filtering techniques complement each other well. Because the two methods are independent they can be applied sequentially, each mitigating the respective distortion they were designed to compensate for. The super-resolution provides an additional benefit for mitigating the distortion caused by frequency agility. The modified backprojection only focuses the range where the target exists and leaves residual RSM in range. With the super-resolution mismatched filter, the range extent of the RSM caused by frequency agility is reduced just as it was with bandwidth agility.

Table 5.1 shows the RSE and bin width measurements from full agility. As ex-

Table 5.1: Full Agility Performance

	RSE_{min}	RSE_{max}	W_{NN} (bins)
Matched Filter	2.20 dB	9.34 dB	16
Windowed Matched Filter	-1.38 dB	5.74 dB	76
Min. Mismatched	-2.64 dB	4.48 dB	32
Max. Mismatched	-0.30 dB	6.82 dB	10
Scaled Super Res. Mismatched	-4.41 dB	2.72 dB	24

pected, in most cases the addition of the frequency agility increases the RSE as compared to bandwidth agility alone. Because the sidelobes are not focused like the main-lobe of the target, the sidelobes create some smearing at range. Similarly, performance is improved over the frequency agility only case because the window and mismatched filters help to reduce the range spread by reducing the sidelobe level of the waveform.

Figure 5.1 shows a CPI of real data from the scenario discussed previously using fully agile pulses and standard matched filter and FFT Doppler processing with the range and Doppler cuts at the target’s location shown. Without correction only the antenna coupling at zero range and zero Doppler is focused, the target and clutter scatters have their energy spread across Doppler and do not focus to a peak. Figure 5.2 shows the result of processing with the super-resolution mismatched filter and phase correction technique. Using the proposed processing chain, the target and even clutter scatters are focused into sharp responses.

Applying the phase correction at the output of the mismatched filter is trivial and can be done with a lookup table to speed up the process. The mismatched filter is very computationally intensive with a matrix inverse on the order of the number of pulse samples. However, by processing at baseband, only mismatched filters for each bandwidth are required, not each bandwidth/IF pair that is required for IF processing. By processing at baseband, the mismatched filters could reasonably be precomputed and then pulled from storage when a pulse with a certain bandwidth is needed. Recent




Figure 5.1: A real-world CPI of fully agile pulses processed with matched filtering and FFT-based Doppler processing without correction. The range and Doppler cuts at the target's location are shown.




Figure 5.2: A real-world CPI of fully agile pulses processed with super resolution mismatched filter and modified backprojection. The range and Doppler cuts at the target's location are shown.

works have investigated methods for reducing the computational complexity of the mismatched filter, but the reduced computation does have some trade-offs that need to be studied for each application [62].

Other work has investigated methods for handling the effect that non-uniform pulses have on the clutter and clutter statistics, especially with spectrally notched waveforms [63] [64]. Since the present work looks more at the deterministic effects of pulse-to-pulse agility, a more appropriate comparison is the Richardson-Lucy deconvolution technique described in [65] that deconvolves the distorted point spread function from the range-Doppler image. The Richardson-Lucy technique is very computationally efficient as it only uses convolution, which can be efficiently implemented with the FFT, and element-wise multiplication and division.

As with the work in [63], the Richardson-Lucy technique described in [65] processes the pulses at the intermediate frequency. As described in section 4.1, when processing at IF there is a strong target peak that is distorted or presented with ambiguities. The Richardson-Lucy algorithm is ideal for restoring a distorted PSF. Unfortunately, as demonstrated in Figure 4.3, when processing at baseband the target is completely obscured, so there aren't any features in the point spread function that the Richardson-Lucy technique can use to extract the true target return.

5.1 Computational Complexity

The computational complexity of the combined algorithm is of interest to determine if the algorithms could be applied in real-time to the continuous stream of data that a radar system produces. First, the most computationally intensive part of the algorithm is calculating the mismatched filter of each pulse. Calculating the mismatched filter for a given template requires a matrix inverse of the total filter size, which, as

mentioned in section 3.1, is twice the length of the waveform, leaving a complexity of $O((2N)^3)$.

Fortunately, the mismatched filters can be calculated beforehand. If all pulses are mixed down to baseband, then the mismatched filters rely only on bandwidth. For a given system, the band selection algorithm would likely use a maximum and minimum selection bandwidth. Additionally, band selection has a resolution. HO-CAE, for example, had a resolution of 97.656 kHz, which is set by the instantaneous bandwidth of the system and the size of the FFT used to generate the snapshot. In this case, a mismatched filter would be generated for each discrete bandwidth that could be used based on the number of bins that HO-CAE can select.

By precomputing the mismatched filter, its complexity is not considered for real-time operation, leaving only the backprojection method to consider. Backprojection is highly parallelizable, with each range-Doppler cell independent in the calculation. The most intensive part of the calculation for each range-Doppler cell is the coarse oversampling that uses the FFT. The FFT algorithm can be carried out in $O(n \log(n))$, where n is the number of values in the oversampled vector. With sufficient resources for parallel compute, the backprojection algorithm has been demonstrated in real-time in [66].

Because the two algorithms needed to process fully agile pulses operate in different domains, they are complementary for implementation. With the benefit of calculating mismatched filters offline, the ability to operate in real-time becomes limited to the modified backprojection technique, which can be accelerated to operate in real-time due to its highly parallelizable nature. Together, these algorithms provide a highly effective and efficient method for processing pulses that change bandwidth and frequency on a pulse-by-pulse basis.

Chapter 6

Communications Interference Mitigation

Non-cooperative spectrum sharing schemes will always present the potential for interference between users. In fully adaptive schemes of the future, there are times when a mission-critical radar system, like the terminal Doppler weather radar (TDWR) described in [30], may not be able to find any available spectrum, which will require operating in the presence of an interfering user. No matter the reason, when a communications user is present as an in-band interference source, techniques are needed to minimize the impact on radar processing.

Previous work has introduced demodulation and remodulation (demod/remod) as an estimation technique for interference cancellation when the radar is operating in a simultaneous transmit and receive (STAR) mode [67]. Passive radar heavily influenced the demod/remod technique. Instead of transmitting energy, passive radar systems utilize signals of opportunity transmitted by local infrastructure or other emitters present around the scene [68].

Passive radar systems collect two signals of interest. First, a receive antenna is pointed directly at the signal of opportunity source, providing a direct path reference waveform that is used for correlation processing [69]. Second, a receive antenna is pointed in the direction of interest to collect echoes from the scene. The data is pro-

cessed from these two input sources using a bistatic geometry [69].

Over the last two decades, passive radar systems have been developed and investigated to operate on many signals of interest such as FM radio, DAB radio, and DVB-T [70], and the feasibility of many other emitters of opportunity has been examined [71]. As wireless communication devices proliferate and telecommunications networks rapidly expand, recent work has shown that 5G parameters can be extracted so that 5G networks can be used as a source for passive radar [72].

While passive and STAR systems are able to receive continuously without interruption, many high-power radar systems must blank their receiver when transmitting, which results in the loss of some reference signal information [31]. This chapter presents an active interference mitigation technique using demodulation and remodulation to estimate the reference signal for extraction. The main focus of this chapter assumes that the radar system only has a single antenna or that the interference is in the main beam of a steered multi-element array.

The OFDM communications model presented in Chapter 2.3 is used to simulate an OFDM time signal, which is received by a pulse-Doppler radar system that blanks its receiver on transmit. When combining the radar and interference signals, the interference power is set to some *output* interference-to-signal ratio (ISR), which describes the ratio of the average interference power to the target peak in the output range-Doppler map. Complex white Gaussian noise is generated with an *output* interference-to-noise ratio (INR) of -25 dB. Each of the three vectors, the radar return vector, the interference vector, and the noise vector, are added together after the radar timeline has been applied to remove data that would be lost from the receiver blanking but before any processing has occurred.

This chapter will begin by exploring the ideal performance of the demod/remod

technique. Initially, the communications signal is generated without a channel or Doppler offset. The initial performance evaluation provides insight into the effects of receiver blanking. The next section of this chapter expands the OFDM model to include a channel and a large Doppler offset. The performance will be evaluated by examining the technique's raw fast-time power reduction, probability of detection, false alarm statistics, and target peak loss after interference cancellation and traditional range-Doppler processing.

The chapter will conclude by investigating and comparing the demod/remod process with a statistics-based approach called adaptive pulse compression (APC) [73]. APC's sensitivity to model mismatch and the notional covariance will be examined to understand the future feasibility of such a technique.

6.1 Demodulation//Remodulation Estimation: The Ideal Model

A radar system sends a pulse and then collects a PRI of samples. The pulse-and-listen cycle is carried out continuously and some number of PRIs are collected into a CPI and processed together. When the radar system is transmitting the receiver is blanked. When the blanking occurs, information about the continuous interference is lost. This blanking is referred to as the system channel. The time duration of the waveform, and the data loss from the disabled receiver, can be relatively short compared to the total duration of an OFDM symbol. Figure 6.1 shows receiver blanking from a uniform pulse train relative to a series of OFDM symbols.

Pulses are likely to fall within the boundaries of a symbol and are unlikely to be time aligned. If the scope of OFDM demodulation were confined to a PRI, more data would be lost than necessary because the rest of the symbol duration beyond the PRI boundary would need to be zero-padded. Zooming out and demodulating relative




Figure 6.1: A series of OFDM symbols shown with receiver blanking shown in red.
[74] ©IEEE

to a CPI means that only the samples blanked by the receiver and symbols at the beginning and end of the CPI need to be zero padded. It is, therefore, advantageous to demodulate an entire CPI of symbols at once instead of one PRI at a time.

Figure 6.2 shows the blanked region of a full interfering communications signal and the resulting sample vector that is received by the radar system. Interference from an OFDM communication network source is constantly present at the antenna, as shown in the top left of Figure 6.2. The top right of the figure shows the samples that are collected by the radar system. Notably, the sampled full CPI vector does not accurately reflect the true time as the blanked times are “missing.” Consequently, zeros must be inserted to make the vector represent a true time axis.

Estimating the communications signal begins by examining a full CPI worth of data, and zeros are inserted into the vector to represent the samples lost by the receiver blanking. Additionally, zeros are appended to the end of the CPI so that there is enough samples for an integer number of OFDM symbols. The CPI vector is divided on the symbol boundaries, and each symbol is demodulated. Standard demodulation is used as described in section 2.3. First the cyclic prefix is removed and then the FFT is applied to the remaining samples. Finally, the complex symbol values are removed




Figure 6.2: The interference signal with blanked regions highlighted (top left) and the samples that are actually received by the system (top right). The bottom shows the zero inserted vector that is used for demodulation.

from each occupied subcarrier. Using the modulation order constellation, the symbol values can be translated to bits. Each of the complex values are snapped to the nearest constellation point to decode data.

Once bits have been extracted, the appropriate modulation constellation is used to modulate the data back onto an OFDM symbol using the same modulation parameters that were used for demodulation. By decoding down to the bits, a robust estimate of the symbol can be constructed where error results captured by the bit error rate (BER).

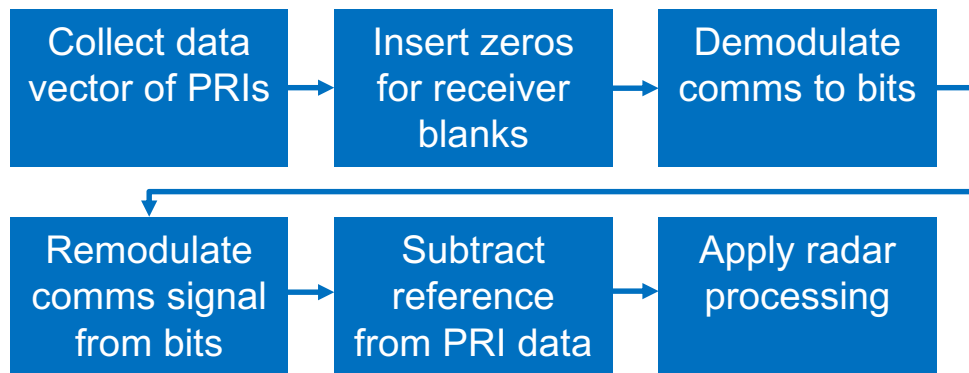


Figure 6.3: The demodulation/remodulation estimation and extraction technique block diagram [74]. ©IEEE

At higher levels of the communications model, forward error correction is utilized in the data packets that can account for low BERs [75]. This work will stay confined to the lowest level of the communication model and assume that no information is available about the data packet level structure.

For sufficiently high INRs, BER would be low enough to effectively extract all of the bits and provide a perfect reconstruction of the time domain signal. This would continue to hold if the entire interfering signal were captured; however, some data is lost due to receiver blanking. The data loss results in a degraded estimate of the transmitted signal. After remodulation, the reference signal vector is subtracted from the original data record. With a perfect estimate of the interference signal, subtracting the reference signal from the data record would result in only noise and radar returns. However, because the estimate is imperfect, a residue of the communication interference is left behind.

Once the estimated signal has been removed from the radar returns, standard matched filtering and Doppler processing can be applied. Figure 6.3 shows a block diagram of the demodulation/remodulation (demod/remod) estimation and extraction technique relative to the radar processing flow.

6.1.1 Initial Performance

To understand the demod/remod performance with the system channel included, a Monte Carlo simulation is carried out for several combinations of PRF and pulse duty cycles. Here the duty cycle defines the pulse duration as a percentage of the PRI. To begin, uniformly distributed random data bits are generated and modulated to generate a series of OFDM symbols whose total duration is at least the length of a CPI. The OFDM symbols are blanked, and zeros are added according to the radar timeline for the given PRF and duty cycle under test.

Once the blanked symbols are generated, a blanked noise vector is added to create an interference-to-noise ratio of -25 dB. The set of noisy, blanked symbols are then demodulated and remodulated to create a reference signal and the reference signal is subtracted from the initial receive vector of blanked symbols. The result is the reduction in interference and noise power. If the symbols are perfectly estimated, the total reduction would be 25 dB since the total performance is bounded by the noise power. For each test case, 5000 CPI iterations were carried out and the results are reported in Figure 6.4.

The results show that total interference reduction performance is maximized with a lower PRF and shorter pulses, which is in direct contention with traditional radar performance metrics. While a radar system optimizes parameters to maximize energy on target, here longer pulses and more energy on target increase the symbol data loss and degrades the estimate of the interfering signal. Similarly, as the PRF increases, although the pulses duration is shorter when using the same duty cycle, having more discrete blanks in the symbol has a greater impact than a larger single blank.

It is intuitive that interference cancellation performance is better when data loss

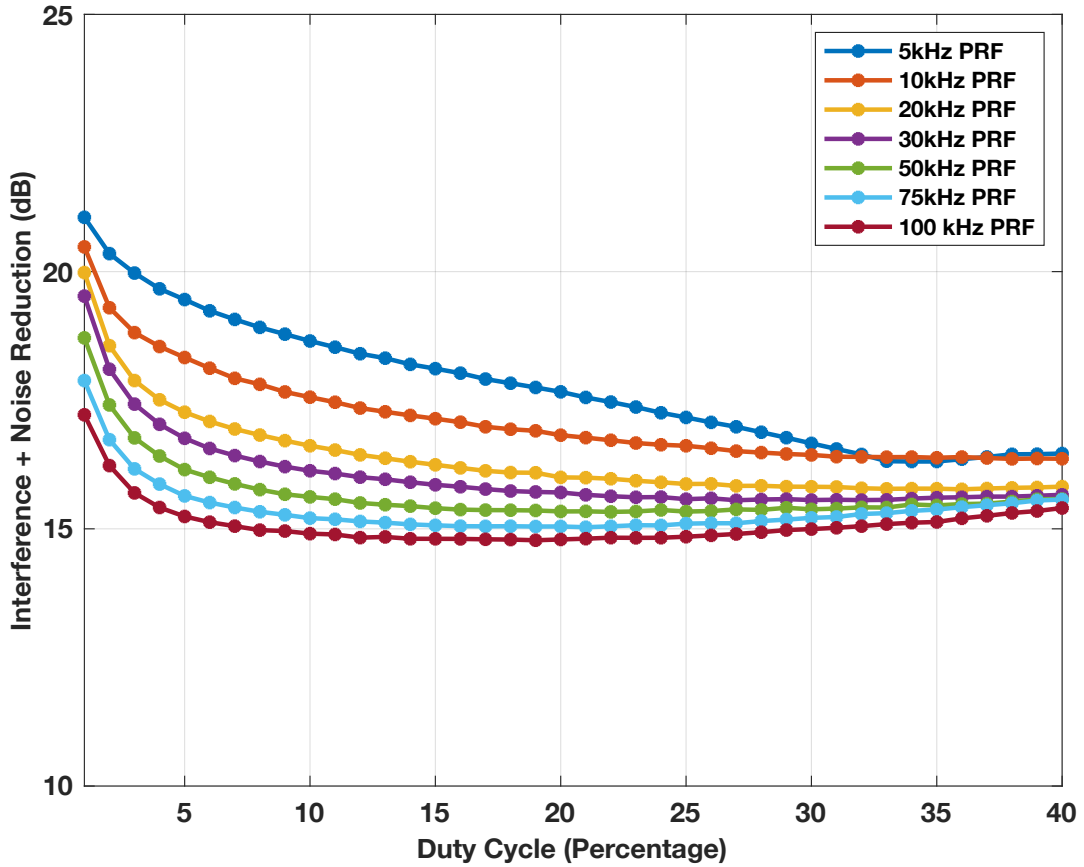


Figure 6.4: Average power reduction from demodulation/remodulation estimate and extraction [74]. ©IEEE

is reduced. Looking closer at the OFDM interference signal model and the receiver blanking provides a better understanding of the underlying behavior. Recall that in section 2.3 OFDM communications symbols are defined as

$$x_i(t) = \frac{\text{rect}\left(\frac{t-T_c}{T_c}\right)}{\sqrt{N_{sub}}} \sum_{k=0}^{N_{sub}-1} X_k e^{j2\pi kt/T_c} . \quad (6.1)$$

Looking at just the OFDM definition, which is just a sum of sinusoids, each with a phase and amplitude term, one might expect that the power of each sinusoidal component is simply reduced. However there is more happening when blanking is intro-

duced. The system channel can be defined as

$$\rho(t) = u(t) + \sum_{i=0}^{M-1} -u(\Upsilon_i) + u(\Upsilon_i + T_p) , \quad (6.2)$$

where $u(t)$ is the unit step function, Υ_i is the start time of the i^{th} pulse and T_p is the pulse duration. Unlike the free space channel described in section 2.3, the system channel is multiplicative. The receive signal is therefore defined as

$$x_{iRX}(t) = x_i(t) \cdot \rho(t) . \quad (6.3)$$

Multiplying the communication signal in the time domain is equivalent to convolving with a sinc function in the frequency domain. When expanded, the rect function that defines the time duration of the OFDM symbol is what implicitly changes. Therefore, the system channel simply replaces the rect function to yield

$$x_{iRX}(t) = \frac{\rho(t)}{\sqrt{N_{sub}}} \sum_{k=0}^{N_{sub}-1} X_k e^{j2\pi kt/T_c} . \quad (6.4)$$

Changing the shape of the rect function changes the orthogonality of the subcarriers. Without the system channel each subcarrier has a perfect sinc like shape where all other subcarriers reside in a null of the each other's sinc function. Figure 6.5 shows how a single subcarrier's oversampled sinc function changes when the system channel is used in place of the rect function. Looking at the oversampled sinc functions, the sidelobes' shape changes such that adjacent subcarriers no longer land in the sinc nulls and the sidelobe structure energy is increased.

When the rect function is replaced by the system channel, the shape, while still sinc like, changes and some interaction between the symbols can occur. When the orthogonality of the subcarriers changes, an increase in the BER rate should be expected.




Figure 6.5: The sinc shape changes from the nominal shape as more of the symbol is lost to the system channel.

This sinc shape is further degraded when the pulse duration is increased which leads to worse performance.

Figure 6.6 shows how a subcarrier contributes to all other subcarriers when the middle of the OFDM symbol is blanked in time. When the full signal is captured, each point is independent and the subcarrier only provides energy to itself, as every other point is orthogonal and falls into a null in the sidelobes. When some of the signal is blanked, then the points are no longer orthogonal and each subcarrier contributes energy to every other subcarrier.




Figure 6.6: The sinc shape shows a contribution of every point to each subcarrier shown for a single blank of a quarter and half of a symbol compared to the fully received case.

6.1.2 Blanking Location

In addition to the sinc effects discussed above, the location of the data loss also plays an important roll in defining reconstruction performance. Three distinct performance categories can be defined based on the regions that are lost to the system channel.

First, if the blank is confined exclusively to the cyclic prefix, the symbol data can still be demodulated without loss, because the cyclic prefix is stripped off before symbol demodulation and so none of the data is changed. Additionally, since there is no

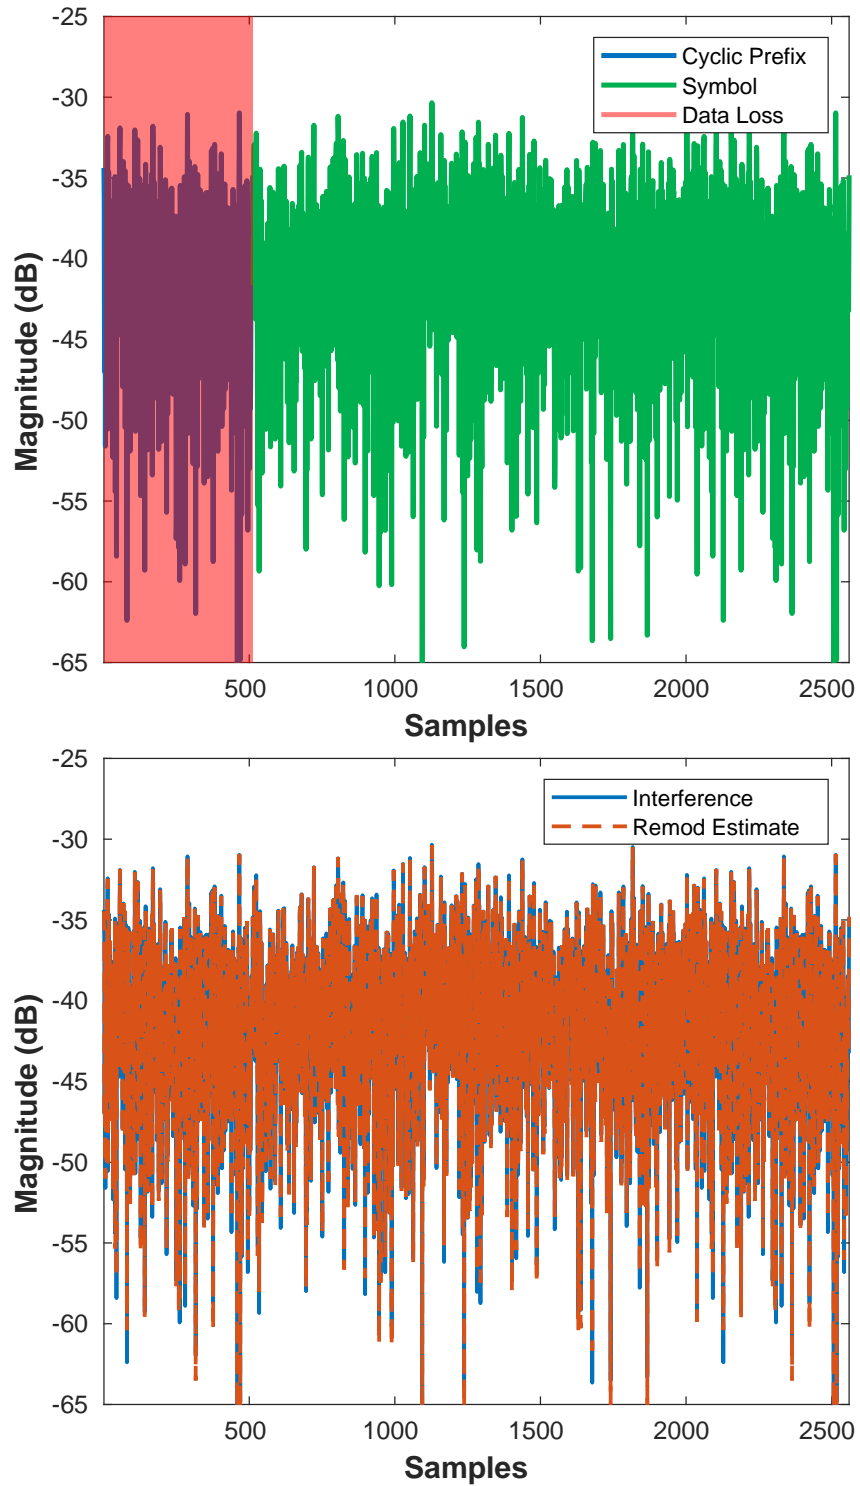


Figure 6.7: When the system channel is confined to the cyclic prefix no data is lost and the symbol is perfectly reconstructed. The top shows the region that is blanked and the bottom shows the true signal and the reconstructed signal match perfectly.

data loss, the orthogonality of the subcarriers is not changed. Without any demodulation loss, the estimate of the signal allows for complete extraction of the interference. Figure 6.7 shows the result of blanking only the cyclic prefix.

Although total reconstruction of the signal is a desired result, confining a system's pulses to the cyclic prefix is not feasible. The LTE and 5G standards define different cyclic prefix durations with some having non-uniform cyclic prefix duration from symbol-to-symbol over a slot [36], [37]. Within a single 10 millisecond frame, some configurations use a longer cyclic prefix on the first symbol and a short cyclic prefix on the remaining six symbols. This would apply a significant constraint on the PRF and pulse duration that may not be compatible with the mission of the radar system.

Additionally, if multiple communication channels are present in-band of the radar system, their respective cyclic prefixes may not be aligned, especially if they are from two different networks. All considerations combined, it is impracticable to confine the radar pulse to the cyclic prefix.

Second, when the receiver blanking occurs during the symbol, the estimated signal has a reduced power compared to the actual signal. Since all subcarriers are present at all time, when some of the body of the symbol is blanked then energy is lost in all symbols. Here is where the change in the sinc shape begins to impact the symbol demodulation. The energy loss combined with the increase in BER from the cross subcarrier interaction mean that a perfect reconstruction is no longer possible.

When remodulating with data loss in the middle of the symbol, the reduced power and subcarrier interaction create a signal that varies from the true reference, resulting in some interference residue left over in the radar return. The residual energy can cause a loss of stationarity in fast time. For example if a PRI captures multiple complete symbols and a partial symbol, full symbols yield better estimates than partial symbols.

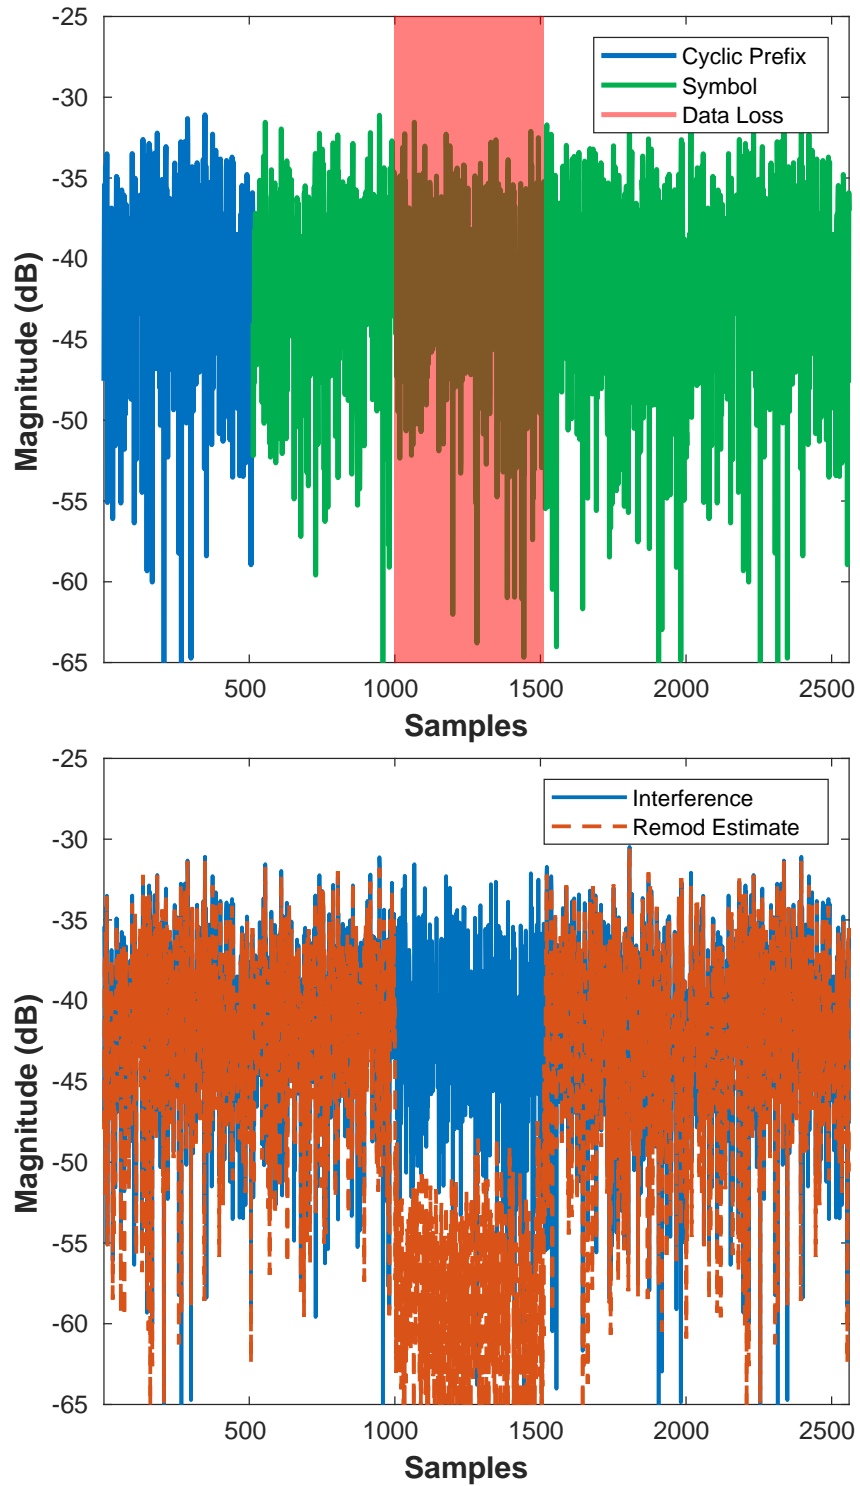


Figure 6.8: When the system channel causes data loss in the middle of the symbol, the BER rate is increased because of the change in the subcarrier orthogonality, however the effect is minimal. The top shows the region that is lost to the system channel and the bottom shows the true signal and the reconstructed signal.

From a CFAR detector's perspective, this is a change in noise power with respect to range. Meaning that the noise and interference statistics are not stationary over fast time which will affect the probability of false alarm over the transition period between the symbols.

Figure 6.8 shows the result of blanking in the middle of the symbol. The top plot shows the true symbol and the region that is lost to the system channel, while the bottom figure shows the true signal and the reconstructed symbol. It is notable that the reconstructed signal is estimated with much lower power in the region where the receiver was blanked. This low power region does not impact the interference cancellation performance because there are not any radar returns of interest in that region as the data vector actually doesn't have data in that region at all.

Additionally, the loss in power can be seen in the peaks of the true and reconstructed signal. The reconstructed signal does not meet the peak heights of the true signal with some peaks being flattened all together. It is these peak losses that cause residual power to be left behind during subtraction.

Finally, the worst case occurs when the system channel causes a loss in the last N_{CP} samples of the symbol. When considering just the symbol, the reconstruction performance is the same. However, the cyclic prefix reconstruction is significantly degraded. As before, the blanked samples are reconstructed at much lower power, in this case, the last samples.

Because the last samples are used to generate the cyclic prefix, if those samples are lost, then the samples that result from remodulation will be significantly lower power. In the previous two cases, it doesn't matter how good the lost samples are reconstructed because they aren't needed to extract the interference as no data was actually collected in that time. However, the loss in reconstruction at the end of the

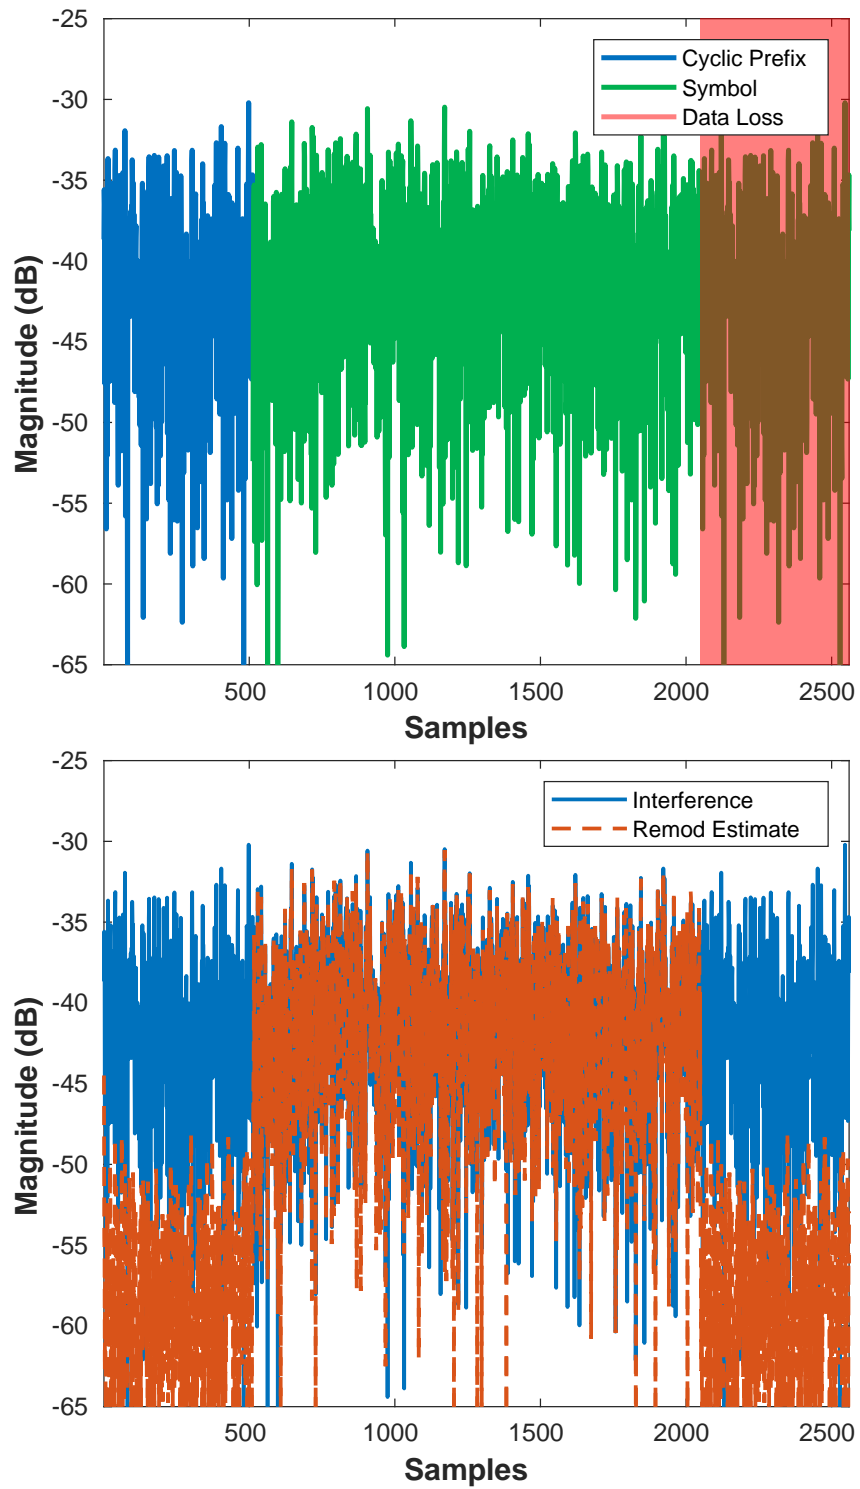


Figure 6.9: The left plot shows the blank location on the true signal and the right shows the reconstruction performance relative to the true signal.

symbol still have a significant impact.

Figure 6.9 shows the reconstructed symbol when the last 512 samples are lost to the system channel. Note the top shows the region of the symbol that is lost while the bottom shows that the end of the symbol and the cyclic prefix fail to be reconstructed correctly.

The consequences of losing these last samples can be seen in the range-Doppler map that is generated after interference mitigation. Figure 6.10 shows a range-Doppler plot, without any targets present, where one of the pulses falls within the last N_{CP} samples of one of the symbols. The cyclic prefix interference can be seen clearly from 5 to 6 km. This type of interference cancellation is especially concerning when detection is considered. The cyclic prefix for this simulation includes 144 samples, but depending on the number of cyclic prefix samples and the size of the guard window used with the CFAR detector, the entire cyclic prefix region may be detected as targets. For a solid region like the one shown here, the false alarm rate would be significantly impacted by such a large number of false detections.

6.1.3 Cyclic Prefix Stacking

The cyclic prefix which is causing such a problem can also serve as the solution. If the end of the symbol is lost and the cyclic prefix is not lost, then the cyclic prefix samples can fill the lost samples at the end of the symbols. They are effectively stacking the cyclic prefix onto the back of the symbol when the back of the symbol is lost. The cyclic prefix filling allows for a better estimate of the symbol, especially the end of the symbol, which means that the cyclic prefix can be successfully canceled.

In the case that the cyclic prefix and the end of the symbol are lost, cyclic prefix stacking is not possible, but it is also not necessary. If the cyclic prefix is lost, then




Figure 6.10: Range-Doppler map with cyclic prefix interference clearly visible while the other parts of the symbol are extracted [74]. ©IEEE

no residual interference, like that seen in 6.10, will be present. Figure 6.11 shows a range-Doppler plot using the same pulse locations and symbol train as in Figure 6.10; however, cyclic prefix stacking successfully removes the high energy region that was caused by the loss of the end of the symbol.

There are some issues that can occur with cyclic prefix stacking that will need consideration later in this chapter. The cyclic prefix serves as a guard interval between symbols in time. When the symbol is transmitted it is convolved with the environment which can smear residue from the symbol in time. The cyclic prefix prevents the smearing from the preceding symbol from contaminating the symbol it was attached




Figure 6.11: Range-Doppler map with cyclic prefix stacking used to estimate and extract the cyclic prefix even when the end of the symbol is lost [74]. ©IEEE

to. Without the cyclic prefix, the two symbols would smear together to create intersymbol interference (ISI). Accordingly, by stacking the cyclic prefix into the back of the symbol, some of the previous symbol information comes with it.

6.1.4 Radar Detection Performance

The previous sections have shown that the technique provides an interference power reduction and cyclic prefix stacking has been introduced to remove outlier high-power regions that may result from the technique. This section examines the performance from a radar's perspective. Once the interference power has been re-

duced matched filtering and FFT Doppler processing is applied to the data.

The improvement in radar performance is examined in terms of detection performance and false alarm performance. To measure this, three targets were generated at random ranges with a uniform distribution. The only constraint on target range is that the radar return must be fully captured to provide a full match with the matched filter. Similarly, the velocity of each target is selected randomly selected from a uniform random distribution that allows for a target velocity to be up to 90 percent of the total unambiguous velocity based on the PRF.

For this simulation, the output SIR is set using the precancellation ISR, which is the relationship between the fast time interference power and the peak of the target in the range-Doppler map. For example, a 0 dB precancellation ISR means the interference signal's average power is set to the same level as the target peak in the range-Doppler map. From here, noise is generated to provide 25 dB of INR.

The simulation was run with 5000 iterations for each combination of precancellation ISR and duty cycle with a 10 kHz PRF. A CFAR detector is used with a guard interval of 32 bins and a training window of 16 bins, and the false alarm rate was set to 1×10^{-6} . Each detection registered by the CFAR detector is classified as a detection or false alarm. If any of the bins adjacent to the true target bin contained a detection, then the target was considered detected. The immediate bins around the true target bin were considered to account for any range straddling effects. Because the targets were randomly generated, they did not fall exactly on the grid points. Similarly, after all of the true target locations were checked, all of the remaining detections in the range-Doppler map were counted as false alarms. This means that sidelobes of targets that are well above the noise and interference would have any sidelobe detections count as a false alarm.




Figure 6.12: Detection performance of detection after interference is canceled as a function of the ISR before cancellation with a 10 kHz[74]. ©IEEE

Figure 6.12 shows the detection performance as a function of the precancellation ISR. The baseline detection performance, that is, the performance without any cancellation, is shown with the dashed lines. As expected, the baseline detection performance is the same for all duty cycles where the target peaks need to be well above the interference power to be detected.

The detection performance mirrors the results shown in Figure 6.4. By using the demodulation/remodulation estimation and extraction method, the targets can be detected even if they are 15 dB below the interference when using short pulses relative to the PRI. The demod/Remod technique provides up to a 20 dB performance improvement over the baseline performance.

Figure 6.13 shows range-Doppler processing applied to a CPI of data with 0 dB




Figure 6.13: Range-Doppler map generated without interference cancellation with an ISR of 0 dB [74]. ©IEEE




Figure 6.14: Range-Doppler map generated after interference cancellation with a precancellation ISR of 0 dB [74]. ©IEEE

precancellation ISR. Here, the targets are completely obscured and cannot be detected by a CFAR detector. After applying the interference cancellation, the range-Doppler map in Figure 6.14 is produced, and the two targets at 4.06 km and 4.13 km are easily identifiable and, importantly, are detected by CFAR.

The interference mitigation does have a negative impact on the detector's false alarm performance. The false alarm rate for the CFAR detector was set aggressively to 1×10^{-6} ; it is expected that the total number of false alarms divided by the total number of bins tested should approach this value. For the baseline case when interference cancellation was not used, the false alarm rate was approximately 1×10^{-6} as desired. After interference cancellation, the average false alarm rate increased to 5×10^{-5} , which is an order of magnitude increase.

6.2 Increasing Model Fidelity

The previous model showed that using demodulation/remodulation estimation and extraction is feasible even when the receiver is blanked, and some of the symbols are lost. The section provided an analysis of performance and the source of estimation errors in the simple case when there is no channel or Doppler offset and the receiver is assumed to be time aligned with the symbol boundaries.

This section will expand the model to include a complex channel, large Doppler offset, and random time offsets relative to the receiver. Each of these parameters must be estimated, removed from the symbol, and reapplied to the reconstructed estimate of the communication signal before the estimate can be subtracted from the received data vector.

With each new piece of the model this section will introduce a technique for estimating each of the above parameters. Once all parameters are estimated the section

will conclude section with a technique to mitigate errors that result from misestimation. The only parameters that are assumed to be known are the subcarrier spacing, cyclic prefix length, and modulation order.

Work has been done that estimates the subcarrier spacing using order statistics [76], various clustering techniques [77], and likelihood ratio tests [78]. However, this work assumes knowledge of subcarrier spacing and the modulation. Previous work in passive radar has shown that it is possible to extract basic parameters from the network with a subscribed side channel device [72].

6.2.1 Symbol Boundary Estimation

The first step to demodulating an incoming OFDM signal is to identify symbol boundaries. Detecting the first symbol boundary helps to locate the first full symbol and provides time alignment to the symbol data stream. As discussed in Section 2.3, traditional LTE and 5G networks utilize primary and secondary synchronization sequences placed throughout a frame to provide coarse and fine timing synchronization, respectively.

Using the PSS and SSS signal has significant drawbacks when considering it for use as the timing synchronization mechanism for demod/remod. First, these signals only occur twice in each ten millisecond OFDM frame. Both instances of the synchronization sequence must be captured to be effective. This constraint means that if the transmission of either pair of sequences arrives at the radar system while the receiver is blanked, then time alignment would not be possible.

Second, if the PSS and SSS were used then the minimum CPI duration would need to 15 milliseconds or longer to ensure that two PSS and SSS sequences are captured. This work declines to use any mechanisms or techniques that would enforce minimum

or maximum constraints on the radar system. Since the synchronization signals are not going to be used another structure of OFDM needs to be exploited.

The cyclic prefix provides a repeating structure in every symbol. The cyclic prefix is a copy of the last samples in the symbol appended to the front. Therefore, an auto-correlation technique can be used [79]. By correlating a pair of vectors with length N_{CP} that are spaced a symbol apart, a peak will be formed when the two vectors are aligned with the symbol.

Figure 6.15 shows the two regions that are correlated as they are dragged across the received data. At each time step, the two regions are correlated to produce a test statistic for that sample. The correlation operation can be defined as [79]

$$\gamma[n] = \sum_{i=n}^{n+N_{CP}} x_{iRX}[i]x_{iRX}^*[i + N_{SC}] , \quad (6.5)$$

where N_{CP} is the number of samples in the cyclic prefix and N_{SC} is the number of subcarriers that make up the symbol. When applying (6.5) to the zero inserted receive data, a peak will be produced when the two regions align with the symbol boundary at index n .

The correlation term alone does not express the maximum likelihood (ML) estimate of the time offset. In order to calculate the maximum likelihood estimate an additional energy term is needed on that takes the form of [79]

$$\Phi[n] = \frac{1}{2} \sum_{i=n}^{n+N_{CP}} |x_{iRX}[i]|^2 + |x_{iRX}[i + N_{SC}]|^2 . \quad (6.6)$$

The ML estimate is then expressed as [79]

$$\gamma_{ML}[n] = |\gamma[n]| - \rho\Phi[n] , \quad (6.7)$$




Figure 6.15: An OFDM timeline with two correlation regions shown as they move across the data.

where

$$\rho = \frac{\text{SNR}}{\text{SNR} + 1} . \quad (6.8)$$

Figure 6.16 shows the γ and γ_{ML} values relative to the true symbol boundaries. Both the correlation-only technique and the ML estimate produce a peak at the boundary of each of the symbols. However, the ML estimate provides a larger dynamic range by pushing down the off-boundary estimates.

Multiple symbols will be present in each CPI vector of data. Even if k symbols can fit within a CPI and k symbols are expected, the maximum k samples will not yield the appropriate offsets as one single rounded boundary peak may contain many of the largest values of γ_{ML} values. To find the appropriate peak values, test regions are created that are length $N_{SC} + N_{CP}$. The regions chosen here are smaller than the test regions used in [79]. Here, the start of the CPI is used as the beginning point of the first region, and the symbol offset must exist within one symbol length of the start of the CPI. Therefore, the offset found in each test region, relative to the test region starting point, provides the global boundary offset.

An estimate for each symbol boundary can be found by simply taking the maximum value in each of these test regions. However, because of the system channel




Figure 6.16: Correlation compared to maximum likelihood estimate of the boundary shown with the true symbol boundary.

effects, taking an average across all offset values is insufficient for finding the true symbol offset. If the radar pulse blanks any of the cyclic prefixes or last samples of the symbol, the correlation peak will be reduced. Figure 6.17 shows two test regions where the first test region fails to make a peak because some of the cyclic prefix is lost to a pulse. In this case, the value chosen in a test region is selected from the rising slope of the peak caused by a boundary peak that lands in the proceeding test region.

To avoid erroneous values from poorly behaved test regions, the result values from the test regions are reduced. The test region starts with a number of values equal to the number of symbols that fit in the CPI. Each pulse has the potential to corrupt a test region. So, the smallest N samples are removed from the possible values, one for each pulse. An additional two of the smallest values are removed from the set to account




Figure 6.17: This figure shows two test regions for where the first test region’s test statistic is lost due to the system channel. Instead of creating a peak at the symbol boundary, the largest value is the rising edge of the boundary peak in test region 2.

for effects caused by the beginning and end of the CPI.

The remaining values are representative of the true symbol offset and can be averaged together to be used as the true offset. The culling carried out on these values has the effect of removing outliers from the data, where the pulses occur less often than the number of symbols. In testing, using median absolute deviation tests for outlier removal provided the same results.

The performance of this boundary estimation technique was carried out by simulating a PRI of symbols that each have a random time offset such that a symbol does not begin at the start of the data vector. For each of the 1,000 trials, OFDM symbols were generated, a channel was added with a Doppler offset, and the noise was added

25 dB below the mean OFDM signal power. Using the ML estimate of the boundary provided an average error of 4.105 samples. Other techniques have been proposed that utilize the pilot subcarriers present in OFDM [80] [81]. However, the described timing estimator is used since it can be utilized directly from the zero-padded vector without pilot information.

6.2.2 Frequency Offset Estimation

The frequency offset between the transmitter and receiver usually just describes the difference in system oscillators. However, in OFDM, this can also include the Doppler frequency offset caused by the transmitter and receiver moving relative to one another. Traditionally, OFDM communications networks assume that the motion of the user, and therefore the Doppler offset, is small relative to the subcarrier spacing. A maximum Doppler of 300 hertz is considered for a multi-tap channel and 1500 hertz for the high-speed train single-tap channel [40].

As described in previous sections, a Doppler shift or frequency offset can be expressed as a progressive phase shift in the time domain. The phase shift for a frequency or Doppler offset of f_D hertz takes the form $e^{j2\pi F_D t}$.

The cyclic prefix can again be exploited as a way to measure the center frequency offset. Since the cyclic prefix and the last samples are the same at transmit, then they should only have a phase difference related to the Doppler frequency. In the previous subsection, an ML estimator for symbol boundary estimation was presented.

The peak values of γ_{ML} that were identified as valid boundary estimates also contain phase information that can be used for the relative center frequency offset using [79]

$$\hat{\epsilon} = -\frac{1}{2\pi} \angle \gamma(\hat{n}_{ML}) \quad (6.9)$$

where \angle is the argument of a complex number and \hat{n}_{ML} is the index of the ML boundary estimate. The estimate $\hat{\epsilon}$ yields a value between -0.5 and 0.5, which is the frequency offset relative to the subcarrier spacing. Because of the relationship between the symbol's time duration and subcarrier spacing, the resolution of the frequency offset estimate is $-\frac{\Delta f}{2} \leq f_d < \frac{\Delta f}{2}$. If the frequency offset is outside of the resolvable range, the phase rotation across the symbol aliases, making one or more trips around the unit circle.

For large Doppler shifts that approach the $\frac{\Delta f}{2}$, even small errors can cause a sign change in the estimated phase, which results in large incorrect estimates. Again, the use of multiple estimates from the culled set used in boundary estimation can be used. The sign of the phase term of each valid estimate is considered, and a majority vote is used to decide which sign should be used. After the majority vote provides the appropriate sign of the Doppler offset, the absolute value of all of the valid estimates are averaged together to collect the final estimate. As with the boundary estimate, 1000 trials were carried out, each with a stream of symbols that had a random offset and a Doppler shift of 1500 hertz. The trials resulted in a mean error of 10.16 hertz.

Additional timing and frequency offset estimators have been proposed for various modifications to the OFDM symbol structure. Constant modulus OFDM is particularly sensitive to timing and frequency offset, and therefore, a more robust estimator that takes advantage of a training symbol [82]. If control over the network is possible, additional robustness in the estimates can be achieved by adding null subcarriers into the symbols [83].

Many options for timing and frequency offset estimation are available in the literature under different assumptions. The presented techniques were selected because of their low complexity and the minimal number of assumptions. Additionally, there

is room for future expansion with this technique using many parallel correlation operations with different boundary widths and region spacing to help deduce the cyclic prefix and symbol duration based on the largest correlation result.

6.2.3 Channel Estimation

Once the time and frequency offset information are gathered, the Doppler component is removed by applying a phase shift in the opposite direction of the phase shift caused by the frequency offset. A single Doppler estimate is used because, as with radar processing, the velocity of targets in the scene or the platform have a constant velocity over time.

If the source is a cell user moving relative to a stationary radar, the CPI is short enough to extend the assumption that their motion is constant over the CPI. Additionally, if the interference originates from a stationary tower and the platform is moving, the CPI will be short enough to assume that the platform's motion relative to the interference source is constant.

After the Doppler component is removed, the vector is broken apart on the symbol boundaries and each symbol is processed separately. First, each the cyclic prefix is removed from each symbol and an FFT is applied to the remaining points. The resulting grid, X_{dRX} , is $k \times N_{SC}$ with the number of columns equal to the number of subcarriers and the number of rows is equal to the number of symbols.

Some symbols in the communications stream have pilots that are known at the receiver. These pilots are inserted at uniform intervals throughout the symbol. Each of the pilot subcarriers can be used to estimate the free space channel effects on the symbols. In the time domain, the free space propagation channel is represented as a set of complex-valued multipath taps. This creates a flat fading response for each of

the subcarriers and is slowly changing across subcarriers [84]. The channel is then modeled in the frequency domain as

$$\mathbf{X}_{dRX}(f) = \mathbf{X}_{dTX}(f) \odot \mathbf{H}(f) , \quad (6.10)$$

where \odot is the Hadamard product. The dTX notation indicates the transmitted subcarrier grid and dRX is the received subcarrier grid. For pilot subcarriers, the transmitted symbol is known by the receiver and therefore a technique called zero-forcing can be used to extract the channel value at each of the pilot locations [41]. For the subcarriers at the pilot values, the channel can be extracted by simply reformulating (6.10) as

$$\mathbf{H}(f_p) = \mathbf{X}_{dRX}(f_p) \oslash \mathbf{X}_{dTX}(f_p) , \quad (6.11)$$

where f_p are the pilot containing subcarriers and \oslash is an element wise divide. The extracted values are evenly spaced across the occupied subcarriers, and interpolation can be used to estimate the channel values for all other points. For this work, spline interpolation was used to estimate the channel effects for the data-carrying subcarriers. Using the full channel estimate, $\hat{H}_i(F)$, the estimate fo transmitted symbols can be derived as

$$\hat{\mathbf{X}}_{TX}(f) = \mathbf{X}_{RX}(f) \oslash \hat{\mathbf{H}}_i(f) . \quad (6.12)$$

Not every symbol has pilots that can be used to extract the channel response. LTE and 5G communications networks assume that the channel is constant over long time periods of up to 10 milliseconds [36]. However, with the presence of the system channel vector, the channel has to be estimated more locally.

For this work, symbols that contain pilots use their own channel estimate only. Symbols that do not contain pilots will average the channel estimates from the ad-

adjacent symbols except when one of the symbols occurs in a symbol corrupted by a pulse. Channel estimates derived from a contaminated symbol can only be used on that symbol to ensure that the system channel does not reduce the effectiveness of channel estimation.

6.3 Expanded Demodulation/Remodulation Estimation

Now that techniques for estimating symbol time, frequency offset, and free space channel have been described, an updated processing flow can be established. Figure 6.18 shows the enhanced processing chain that enables estimation of the new unknowns that were added in the previous section. Many of the blocks in the new processing flow are familiar. Collect data, insert zeros, demod/remod, and estimate subtraction are all described in the initial investigation of section 6.1. Additionally, the cyclic prefix stacking described in section 6.1.3 has been added to the total processing flow after the time and frequency alignment steps and before the free space channel estimation that were described in the previous subsections.

An additional step has been added to the processing flow that was not previously discussed. Before subtracting the reconstructed signal from the radar data, the free space channel estimate and frequency offset estimate are reapplied to the remodulated signal. Although demod/remod provides an estimate of the exact packet that was transmitted by the interfering emitter, what actually needs to be subtracted from the radar data is the interference as it appeared when it arrived at the aperture. Therefore, the free space channel and Doppler frequency offset need to be added back to the remodulated signal so that an accurate representation of the received signal is subtracted from the radar data.

Similar to the work carried out in section 6.1, a simulation was developed to test

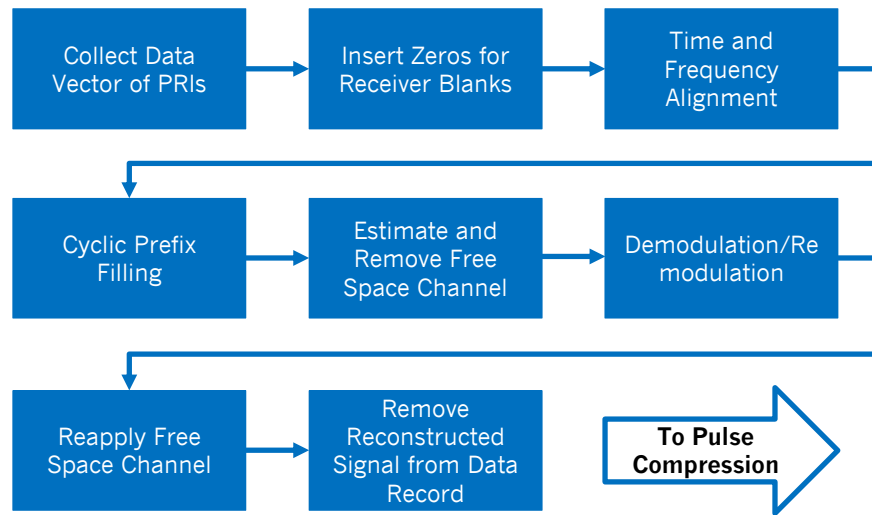


Figure 6.18: Enhanced Processing flow with boundary, frequency, and channel estimation steps included.

the effectiveness of the enhanced processing chain and the robustness of parameter estimation. Initial performance analysis is carried out on a single PRI of fast time data to see the total fast time performance. For fast time analysis, a 1.5 kHz PRF was used to generate data with a 49 microsecond LFM pulse generated such that the return has a matched filter peak response of 0 dB. The communications interference is generated as a continuous stream of OFDM symbols with the parameters described in Table 2.2. Noise is generated 25 dB below the average power of the OFDM interference signal.

The nine tap vehicular validation free space channel model is applied to the stream of OFDM symbols with a random time offset relative to the PRI [40]. Instead of the 70 Hz Doppler that is generally used with the vehicular channel model, a 1500 Hz Doppler shift was used to provide a more difficult interference scenario and robustly test the estimation techniques. At a 5 GHz center frequency, 1500 Hz represents a target motion of 45 meters per second. Figure 6.19 shows the signal components with a 0 dB INR with the interference generated with a mean power of 0 dB.

Each of the three signals are blanked according to the appropriate radar timeline




Figure 6.19: Simulated signals shown separately with the interference generated at 0 dB mean power, noise at -25 dB mean power and radar return generated so that the matched filtered, Doppler processed peak is at 0 dB in the range-Doppler map.

and then combined together to form the total received vector with zeros already appropriately inserted. When each of the simulated signal components are matched filtered with the transmitted LFM waveform the noise and interference maintain approximately the same power level with the radar return formed into a peak. Figure 6.20 shows the result of each of the individual components being matched filtered, which results in the radar peak still falling below the interference but above the noise component.

It should be noted that although the target response peak is approximately 0 dB, it does not reach 0 dB because of range straddling effects. When the target delay is not




Figure 6.20: Each signal component after matched filtering. The radar return creates a peak at approximately 0 dB and the mean noise and interference power remain the 0 dB and -25 dB respectively.

an integer multiple of the sample rate, the energy is split between the two samples that the target return straddles. These targets are considered to be off of the range grid.

The combined signals are processed through the more robust processing chain and the reconstructed signal is subtracted from the receive data vector. Figure 6.21 shows a signal that is left over after the subtraction with the true symbol boundaries marked. In addition to the regions of higher power at near and far ranges that were observed in Section 6.1, there are regions of higher power near each of the symbol boundaries.

Figure 6.22 provides a closer look at one of the symbol boundaries where significant amounts of interference residue is left behind. On closer inspection, the large




Figure 6.21: The interference residue left behind using a demod/remod estimate using time, frequency, and channel estimation.

interference residue strictly follows the symbol boundary, which indicates that the residue exists within the cyclic prefix. The residue is caused by the phenomenon that the cyclic prefix is designed to capture. When the stream of OFDM symbols is convolved with the free space channel, some of the leading symbol is smeared into the symbol that trails, creating inter symbol interference (ISI). The cyclic prefix captures the ISI which protects the data symbol. However, the radar system captures the cyclic prefix interference and so it must also be estimated to effectively cancel it.

The channel is estimated using zero-forcing to provide a estimate for each subcarrier in the frequency domain. Subsequently, the channel is added back to the symbol in the frequency domain before the subcarrier grid is modulated. Consequently, the




Figure 6.22: A closer look at the extra residue that exists around the symbol boundaries.

cyclic prefix has the channel applied but does not contain any ISI that is present in the receive data.

To accurately capture ISI within the cyclic prefix, a time domain estimate of the channel would be needed. When attempting to convert a frequency domain channel estimate to the time domain, determining the length of the channel is critical. When convolving the symbol stream with the time domain signal, which is required to get ISI, the length of the time domain determines how deep the ISI penetrates into the cyclic prefix.

Techniques have been developed that estimate the length of the free space channel [85]. However, as seen in Figure 6.22, the estimate subtraction is very sensitive to

the length of the ISI. Even with the ISI interference residue left behind, the problem has been made significantly easier because the duty cycle of the interference has been significantly reduced.

6.3.1 The Median Filter

Reducing the effective duty cycle from 100% provides the opportunity to examine and apply an existing interference subtraction technique. A class of interference identification and cancellation techniques that have been explored extensively is the median filter [86]. The median filter identifies inference by using the median value across slow time to create a threshold for each range bin [87].

These detectors are energy detectors that require some information about the null hypothesis to get information for a threshold requires some data points that contain no interference. The power distribution of the null hypothesis yields a Rayleigh distribution [88]. Traditional techniques consider each range bin independently. However, the quality of the null distribution estimate can be improved by increasing the number of values used for the estimate.

In this case, the estimate is derived from the median value across slow time at each range bin. The ratio of the cell power and the slow time median provides the baseline test statistic for each cell. Averaging together power-to-median ratios from adjacent range bins improves the estimate, allowing the same probability of false alarm to be achieved with a lower threshold [89]. This work uses the two-dimensional median filtering technique presented in [89], taking advantage of the increased detection performance.

Just like the previous steps in the demod/remod processing chain, median filtering will be applied before any pulse compression or radar processing takes place. First, the

slow time median is calculated for each range bin. Next, the power-to-median ratio is calculated for each bin in the range-pulse plot. Finally, for each value, the mean across range cells adjacent to the cell under test is taken, and the value is compared to the threshold to determine if it is an outlier.

Each of the cells that result in an outlier is flagged as containing interference and needs to be corrected. For this work, a simple interpolation will be carried out in a short time to replace the values that are flagged as containing interference. However, the interpolation in this case needs to consider both magnitude and phase. First, an interpolation in magnitude is carried out using a linear interpolation. Next, the unwrapped phase is linearly interpolated across slow time, and then the wrapped phase is applied to the corresponding magnitude value, resulting in a replacement value for the flagged interference bin.

A higher probability of false alarm of 10^{-4} is used for this application. Because the work considers interferers that are on the order of the target's range-Doppler peak response, it will be much greater than the target return before pulse compression, which makes it unlikely to rise even above the noise as shown in Figure 6.20. Because the pulse return will be so low, it is unlikely that a false alarm will act negatively on the pulse return. Additionally, because the pulses are long in range before pulse compression, they will align across slow time, even in instances of range migration, and contribute to all values, not just values above the mean, allowing for the use of higher false alarm probability settings.

The median filter's performance, when applied to the data, can be seen in the range-pulse plots in Figure 6.23. The top range-pulse plot shows the data after the demod/remod estimate has been subtracted from the received data, and the ISI residue is clearly visible, appearing in every pulse. The median filter cleans up the ISI residue

with the added benefit of softening and shortening the residue in the near and far ranges.

The impact on fast time is of the most interest to provide a comparison to the output of the subtraction algorithm shown in Figure 6.22. Figure 6.24 shows the first PRI at the output of the median filter. As shown in the range-pulse plots, the first PRI shows the successful removal of the ISI residue that was present after symbol subtraction. Additionally, at near and far ranges, the residual interference is caused by system channel-induced symbol misestimation.

Next, Figure 6.25 shows the result of applying the matched filter to the first PRI after the median filter. Applying the matched filter reveals the target peak at the correct location. Additionally, the matched filter also helps smooth the transition from higher residue in the near and far ranges to the rest of the range extent.

Although the median filter successfully removes the impulsive interference, it is important to verify that the target will still focus in the range-Doppler map after the interpolation in amplitude and phase. Basic FFT-Doppler processing was applied across 32 PRIs after demod/remod extraction, and median filtering was applied. The target was generated such that the peak in the range-Doppler plot was the same magnitude as the mean input interference power.

Figure 6.26 shows the resulting range-Doppler plot showing successful interference mitigation and a strongly focused target in the correct range and Doppler bin. Ultimately, the median filter was shown to perform well at removing the interference left over by ISI residue while still allowing the target to focus in both range and Doppler. Therefore, the median filter is added to the final processing chain after the remodulation estimate is removed from the received signal but before any radar processing is applied.

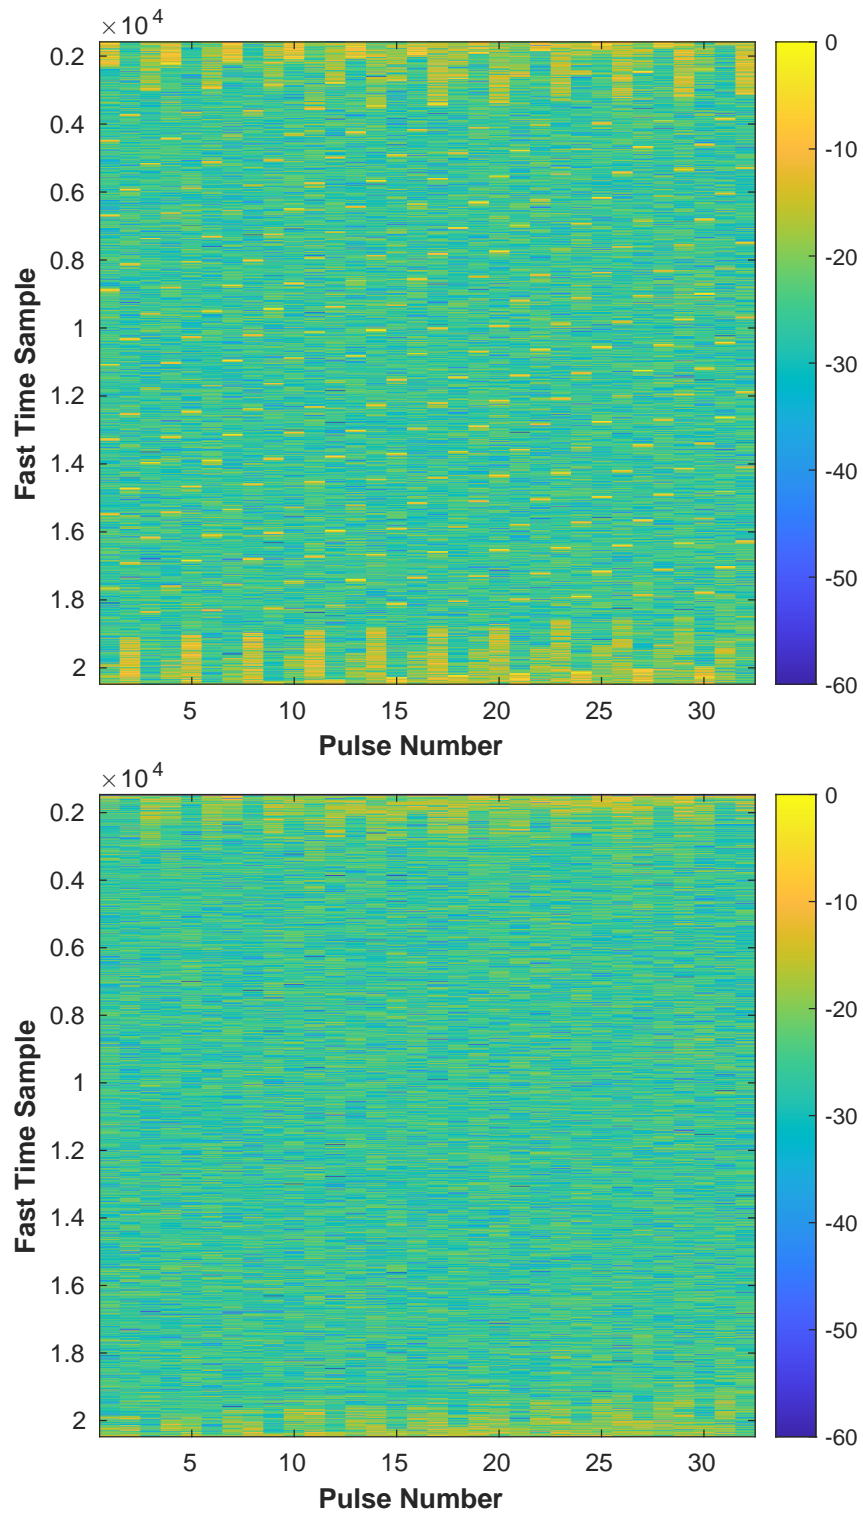


Figure 6.23: The range-pulse plot before (top) and after the median filter (bottom). Before the residue caused by ISI is clearly visible and is successfully detected and removed by the median filter.




Figure 6.24: The first PRI at the output of the median filter. The ISI residual interference has been removed and the near and far range residue from caused by misestimation of the symbol in the presence of the system channel.

Finally, Figure 6.27 shows the final processing chain that is used to cancel the continuous stream of OFDM symbols. Starting with a full CPI of received data in a vector, zeros are padded to get an accurate time vector. The symbol boundaries are estimated, and the Doppler offset is removed based on the estimated value. Any residual Doppler can be removed by the channel estimation.

The time align vector is divided into a matrix, with each row containing one symbol, and cyclic prefix filling is used to repair the end of any malformed symbols. Each symbol has its cyclic prefix removed, and an FFT transforms the symbol into the modulation domain. A channel estimate is taken using the zero-forcing technique. The




Figure 6.25: Applying the matched filter to the output of the median filter shows a strong peak target response and even further smoothing of the near and far range residue.

estimated channel is removed, and each subcarrier is demodulated down to bits.

The demodulated stream of bits is then used to remodulate each subcarrier, and the channel estimate is applied before the IFFT is used to return the time domain and cyclic prefixes are reapplied. The symbol matrix is vectorized, and the time offset and Doppler frequency are reapplied. Finally, the vector is subtracted from the received data vector, and the median filter is applied to the remaining signal, which contains radar returns, noise, and interference residue.




Figure 6.26: The range-Doppler plot shows that after median filtering the target can still be focused successfully at the correct range and Doppler.

6.4 Performance Analysis

Now that a complete demod/remod processing chain has been established to consider the more complex OFDM signal model, the performance can be tested and analyzed. First the fast time power reduction is examined. Additionally, radar parameters such as target peak loss, probability of detection, and probability of false alarm are provided as the key metrics.

A Monte Carlo simulation was set up to capture a robust set of performance metrics that inform a detailed performance analysis. First, the radar PRF was set to 1.5 kHz utilizing a CPI of 32 pulses. Again the pulse duration was varied as a function of duty

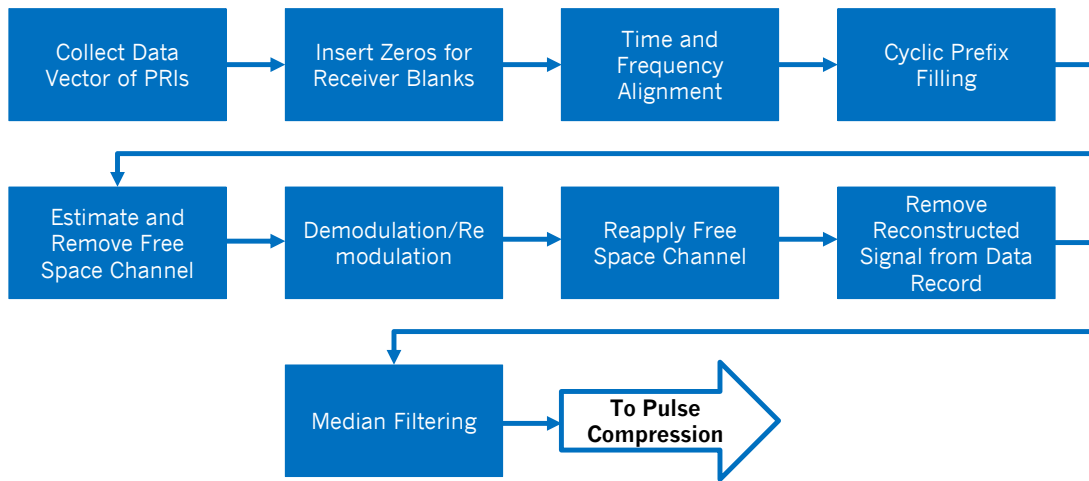


Figure 6.27: The complete processing flow for demod/remod interference estimation and extraction with a complex signal model.

cycle examining 1, 5, and 10 percent duty cycles. For each CPI one target was placed at a random range selected from a uniform distribution of all fully matched ranges and a random velocity selected from uniform distribution of all unambiguous ranges.

Each of the radar returns is generated such that the target peak in each range-Doppler map is 0 dB, and then the interference is generated to create a mean power to achieve some ISR. Additionally, noise is generated below the interference signal, so as the ISR decreases, the SNR increases. For each generated stream of interference, a uniform random sequence of bits is modulated onto each symbol using the parameters in Table 2.2. Once the interference stream has been generated, a random time offset is applied relative to the start of the CPI. Finally, the vehicular free space channel is added to the OFDM symbol through convolution, and a 1500 Hz Doppler shift is added.

The system channel is applied to each of the three signal components, and they are added together to create the total received signal. One thousand CPIs of data are generated for each duty cycle and ISR. For each trial, information is gathered about the

probability of detection and probability of false alarm before and after demod/remod interference estimation, and subtraction is applied using a false alarm rate of 10^{-5} .

Two different LFM waveforms are used for experimentation. The first is an 8 MHz waveform that is contained completely within the 18 MHz of the active subcarriers in the interference waveform. The second is 25 MHz, which provides 7 MHz of bandwidth on each side of the interference that only contains roll-off from the interference.

First, the total energy reduction is examined to understand how the interference and noise reduction performs in the fast time. The average received energy over the duration of the CPI is measured and compared with the mean energy of the remaining residue signal after the median filter. Comparing these two values captures the total energy reduction in the signal, providing a starting point for the performance analysis.

Figure 6.28 shows the total power reduction achieved by the demod/remod interference cancellation technique. It is notable how close the performance is when the interference is the dominant signal. This shows that the channel estimation has little impact on the waveform as long as the waveform is a very small component.

When the radar signal becomes a very large component, the 25 MHz waveform has better performance because it hinders the channel estimate uniformly as a constant offset. The 8 MHz waveform, on the other hand, causes the free space channel to be misestimated in just the middle of the waveform leading to slightly lower performance when the radar signal is dominant over the interference.

Since we can see that the channel estimate is being affected by the waveform, it is important to examine the target's peak loss to determine if any of the radar returns are being suppressed by the interference cancellation. To truly consider peak loss, the range and Doppler straddling must be accounted for to get an accurate measure. To account for this, on each CPI, a reference range-Doppler map is created without




Figure 6.28: Total power reduction achieved by the demod/remod technique shown with 8 MHz LFM contained completely in the interfere and a 25 MHz LFM centered with the interference.

interference or noise, and the peak magnitude in the appropriate range bin is measured as the reference. The difference is then taken between this reference peak and the peak achieved in the range-Doppler map generated after interference removal.

Figure 6.29 shows the peak loss, with the top plot providing the loss from an 8 MHz waveform and the bottom plot showing the 25 MHz waveform loss performance. Three distinct regions emerge when considering peak loss performance.

First, when the radar signal is well below the interference, from 13 dB to 25 dB of ISR, the 8 MHz and 25 MHz cases behave similarly. In this region, the signal is barely detectable even after interference cancellation. Values that show a peak gain





Figure 6.29: The peak loss exhibited in the range-Doppler plot when using an 8 MHz waveform (top) and 25 MHz waveform (bottom).

show an incidentally detected target or a target peak that has significant energy from noise variance. In the high ISR case, not all traces in the plot are complete, as only CPIs that registered a detection in the correct range bin were counted. At the extreme end, there were ISRs with zero detection.

The second region is from -13 dB and 13 dB ISR where the signal is generally undetectable without interference cancellation and detectable after. Here we see a similar peak loss around 5 dB from both waveforms with a slight increase in the 8 MHz case. Finally, in the low ISR case where the radar signal is a dominant component of the signal.

The peak loss at the upper end of the ISR remains flat for the 25 MHz case where the waveform has an even effect on the channel estimate. However, the 8 MHz waveform has a significant rise in peak loss as it begins to dominate the signal. The 8 MHz signal impacts less than half of the used subcarriers in the OFDM signal and becomes a dominant part of the channel estimate and is able to be partially canceled.

With a high peak loss for the radar dominate case it is worth considering the need for interference cancellation. When the target is easily visible above interference, it may not be necessary to apply demod/remod subtraction. However if the interference is detectable above the noise floor, there is always a possibility that it is masking a target that would otherwise be visible.

Looking at the peak loss and fast time reduction together shows that there is still a performance increase even when the radar signal is dominant. The interference is still being pulled down an additional 10 dB further below the target peak for the 8 MHz cases and up to 20 dB when the radar signal is well below the communications signal. This result shows that strong targets will be impacted more by the interference subtraction while the weak signals will not. This means that when both types of targets

are present, strong and weak, they will both be resolved as long as both could be independently uncovered from the interference.

The most important and representative metric for the performance of this technique is the radar detection improvement. Increasing the detectable signal level provides a direct performance that can be measured in magnitude. Figure 6.30 shows the probability of detection curves using both 8 MHz (top) and 25 MHz (bottom) waveforms.

With respect to the probability of detection, the results perform remarkably similar. This should not be a surprise given that the transition from detection to not detected occurs over the region that shared similar peak loss and fast time energy reduction. Both waveforms are slightly slower, rising between 90 and 100 percent compared to the nominal curve. This slower transition at the high end of the probability curve is caused by the higher residue regions in the very close and far range. The increased residue makes the targets harder to detect and means that it takes longer to detect all of the targets as they are uniformly distributed over the range.

Looking at the total amount that the probability curves are pushed to the right gives the most meaningful measure of performance. In this case, the change in detection performance was measured where the curves crossed 50 percent probability of detection. For the 8 MHz case the baseline reference curve crossed the 50 percent at -11.02 dB precancellation ISR and provides a 19.38 dB increase in performance in the best case. Similarly, the baseline for 50 percent detection rate for the 25 MHz case is -9.45 dB providing a maximum improvement of 18.98 dB.

Table 6.1 shows the calculated results at a 50 percent probability of detection. Each of the values measuring performance coincides very closely across all duty cycles and bandwidths. Together, the total performance of the technique can be squarely placed at **19 dB** improvement, allowing the radar to uncover targets that fall nearly two orders




Figure 6.30: The probability of detection exhibited after interference cancellation when using an 8 MHz waveform (top) and 25 MHz waveform (bottom) relative to the before cancellation probability of detection.

of magnitude below the input interference level.

Table 6.1: Probability of detection improvement achieved by full demod/remod interference cancellation processing chain when using a 8 MHz and 25 MHz waveforms.

	Baseline	1% Improvement	5%	10%
8 MHz	-11.02 dB	19.38 dB	19.18 dB	18.97 dB
25 MHz	9.45 dB	18.89 dB	18.98 dB	18.84 dB

Again, the probability of detection cannot be considered alone. The probability of false alarm should also be considered, especially given the increase that was observed in section 6.1.4. Figure 6.31 shows the probability of false alarm before and after interference cancellation has been carried out.

Again, the performance between the 25 MHz case and 8 MHz case are very similar until the radar return dominates the received signal. From 0 dB to 25 dB ISR, the probability of false alarm in both the 8 MHz and 25 MHz cases are flat and fall right around the 10^{-5} rate that was set for the CFAR algorithm.

ISRs below 0 dB start to have an increase in the probability of false alarm. But the rise in false alarm is not because of the processing directly but because of the way that false alarms are counted. The target response in the range-Doppler plot are not perfect thumbtacks and exhibit sidelobes in both range and Doppler. Below 0 dB ISR, the cancellation is performance is good enough that even sidelobes are being uncovered from the interference.

A detection is marked when it lands in the correct range-Doppler bin and those on either side in each dimension. Therefore, a detection on a sidelobe near the true target on the same range or Doppler axis are registered as false alarm even though they are caused by the target. So, the increased false alarm rate exhibited in the simple case is not present where the complete channel model and full processing chain are used.





Figure 6.31: The probability of false alarm exhibited before and after interference cancellation with an 8 MHz waveform (top) and 25 MHz waveform (bottom).

The median filter contributes heavily to resolving the issues with the probability of false alarm. In addition to removing the large residues that are caused by the ISI, the median filter also removes any sample points that far exceed the median. This removal smooths the distribution of the output vector back to what is expected by the CFAR assumptions. While, the range distribution is still not stationary, it tapers more slowly so that it is quasi-stationary within the training interval of CFAR.

Ultimately, this chapter has provided a technique for active interference cancellation that uses demodulation/remodulation estimation and extraction, which provide **19 dB** increased detection performance. The chapter opened with a simple model to test the feasibility of the technique for non-STAR radar systems. The second part of the chapter expanded the technique to consider a more robust model and developed a suite of tools to overcome the challenge of estimating OFDM with real-world considerations. A robust look at performance and the operation of each step in the processing chain was presented, providing an end-to-end approach to OFDM communications interference mitigation.

6.5 A Comparison: Adaptive Pulse Compression

Now that a direct estimation approach has been developed in this chapter, a look at other adaptive approaches for canceling continuous interference is warranted. One adaptive radar processing technique is adaptive pulse compression (APC), which utilizes a reiterative minimum mean square estimate (RMMSE) to reduce the sidelobes of the target response [90]. The RMMSE techniques have expanded rapidly into other domains, including beamforming [91] and RCS extraction [33].

First, APC starts by defining the matched filter of the received data as

$$z[l] = \mathbf{A}^H y_r[l] + v[n] , \quad (6.13)$$

where \mathbf{A} is the delay shift matrix of the transmitted waveform $x(t)$ as defined in (3.4) and v is the noise vector. From this formulation, a cost function can be defined as [92]

$$J_{MMSE} = E \{ |y_I[l] - \mathbf{u}[l]y[l]|^2 \} , \quad (6.14)$$

where $y_I[l]$ is the impulse response of the radar channel and $E\{\cdot\}$ is the expected value operation. Here, the difference between the impulse response and the filter response of the channel is going to be sidelobes and noise. Minimizing this quantity results in reduced sidelobes that approach the noise floor.

A gradient decent algorithm can be used to find the optimal solution [93]. In addition to the optimized result, the gain of the optimized filter should be constrained so that the target response maintains the appropriate magnitude. The gain constraint can be achieved by using a minimum variance distortionless response (MVDR) framework, common in the beamforming literature [94], to create a filter solution which can be defined as [95]

$$\mathbf{u}[l] = \frac{\mathbf{Q}[l]^{-1} s_r}{s_r^H \mathbf{Q}^{-1} s_r} , \quad (6.15)$$

where \mathbf{Q} is defined as

$$\mathbf{Q}[l] = (\mathbf{S}_r \mathbf{P}[l] \mathbf{S}_r^H) \odot \mathbf{I} + \mathbf{R}_n , \quad (6.16)$$

\mathbf{S}_r is the steering vector matrix of delay shift transmit waveforms, s_r is the delayed vector of the transmit waveform for a target return in the l^{th} range bin, \mathbf{P} is the covariance matrix from the matched filtered received data, \mathbf{I} is the identity matrix, and \mathbf{R}_n is the noise covariance. The covariance matrix of the received matched filtered data can

be calculated as

$$\mathbf{P} = z(t)z^H(t). \quad (6.17)$$

On each iteration this matrix P is updated with the covariance matrix derived from the output of the previous iteration's filter bank.

Here, L is defined as the number of fast time samples and K is the number of samples in the matched filtered range extent. Accordingly, \mathbf{S}_r is an $L \times K$ matrix, \mathbf{P} is a $K \times K$ matrix, \mathbf{I} is an $L \times L$ identity matrix, and \mathbf{R}_n is an $L \times L$ matrix.

Recently, APC has been demonstrated as a technique for suppressing interference that is present in the sidelobe of the receive beam [73]. To cancel interference, the covariance matrix of the interference is added to the noise covariance to yield [73]

$$\mathbf{Q}[l] = (\mathbf{S}_r \mathbf{P}[l] \mathbf{S}_r^H) \odot \mathbf{I} + \mathbf{R}_{\text{canc}} + \mathbf{R}_n, \quad (6.18)$$

where \mathbf{R}_{canc} is the interference covariance matrix. For a direct estimate of the interference, the interference covariance can be calculated similar to (6.17) as

$$\mathbf{R}_{\text{canc}} = \hat{y}_c(t) \hat{y}_c^H(t), \quad (6.19)$$

where $\hat{y}_c(t)$ is the estimate of the communication interference received by the system.

Using a phased array radar, when the interference is emitted from a different direction than the radar is scanning, the interference can be separated from the radar return by beamforming in the radar and interference direction individually. Beamforming acts as a spatial filter that allows the interference to be isolated from the radar returns. When the interference estimate is derived from a spatial filter using a STAR radar system, the estimate is free of any system effects. If the emitter is well spatially separated from the radar direction, the APC robustly cancels the interference to isolate the target

as shown in [73]. However, this work considers interference that occurs on the main beam of the radar, that is, in the same direction as the radar is scanning. In this case, beamforming cannot be used to isolate the communications interference, and other methods must be examined.

There are many implementation options for APC that help combat the high computational complexity. Traditionally, as described above, APC requires an inversion of the matrix Q at full rank. However, work has been carried out to reduce the rank of the matrix that must be inverted [96] [97]. This work will utilize a block APC approach that requires just one large matrix inversion per iteration, similar to the spatial RMMSE algorithm called reiterative super-resolution (RISR) [73] [91].

Using block APC, the filter equation becomes

$$u_k = \frac{\mathbf{Q}^{-1} s_{r,k}}{s_{r,k}^H \mathbf{Q}^{-1} s_{r,k}}, \quad (6.20)$$

where each of the k^{th} columns of the filter matrix uses the same full rank inverted Q matrix and the corresponding steering vector $s_{r,k}$. By reusing the same Q matrix, the matrix inverse only needs to be calculated once for each iteration.

The full rank matrix is used to capture the long range covariance structure that may exist in the interference. Without interference, the maximum range effects of the covariance are from the length of the waveform, which allows for blocking regions to be used. However, here no assumptions are made about the interference covariance structure.

6.5.1 Ground Truth Performance

The investigation of APC's performance begins by verifying the behavior when using the ground truth interference data to create the interference covariance matrix.




Figure 6.32: The normalized covariance derived from the ground truth interference data.

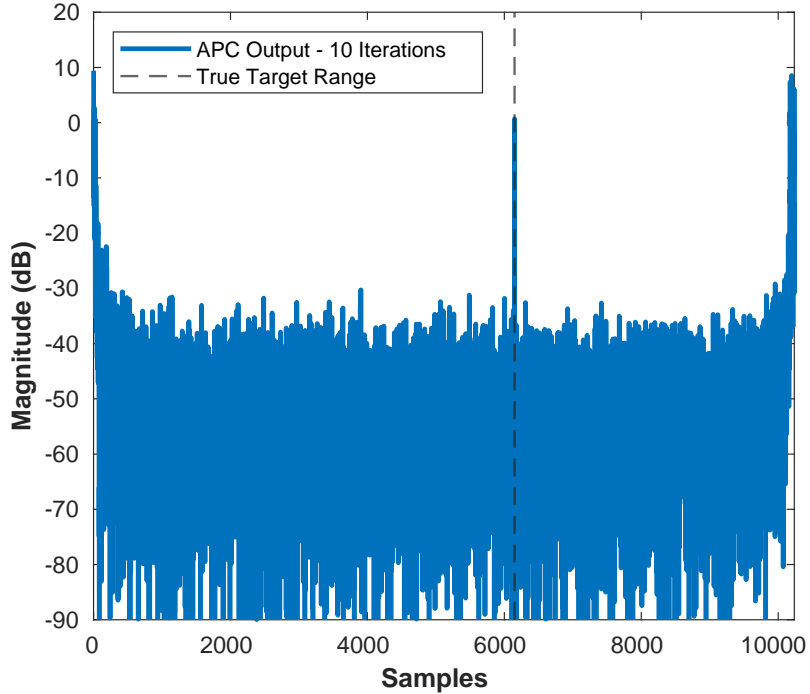


Figure 6.33: APC filter output after 10 iterations using the source truth covariance matrix.

The same simulation from the previous sections is used to generate radar and interference data with a 0 dB precancellation ISR. The ground truth interference data is used to create a covariance matrix according to (6.19).

Figure 6.32 shows the source interference covariance that has been normalized using the Frobenius norm. As expected with a single snapshot, the covariance matrix does not have any structure, and there are very long-range covariance values. However, because the covariance matrix is defined from the truth data, APC performs very well.

Figure 6.33 shows the result of applying the APC filter to the received data after ten APC iterations are run. Here, the target is easily uncovered from the interference with an incredibly smooth and flat range profile. The only real shape seen here are the convolutional tails on each end of the range profile, which are expected and do not impact performance.

Using a phased array, sidelobe interference can be well canceled because the interference can be spatially separated and provides ground truth like data as shown in [73]. Because sources are well separated, beamforming has very little impact on the shape of the collected data. Therefore, for a STAR system, the reference signal derived from a sidelobe beam would be very close to the truth data of what was actually received.

6.5.2 Alternative Estimates

This work considers interference that exists in the mainlobe of the radar and cannot be spatially separated from the radar returns. Without a domain to separate the radar and interference, a pure interference covariance cannot be generated. Therefore, another avenue of estimating the interference covariance needs to be established. First, the demod/remod estimate of the signal is considered. When crafting the demod/remod estimate, the goal is to recreate the interference signal that was received at the




Figure 6.34: The normalized covariance derived from the demod/remod interference estimate.

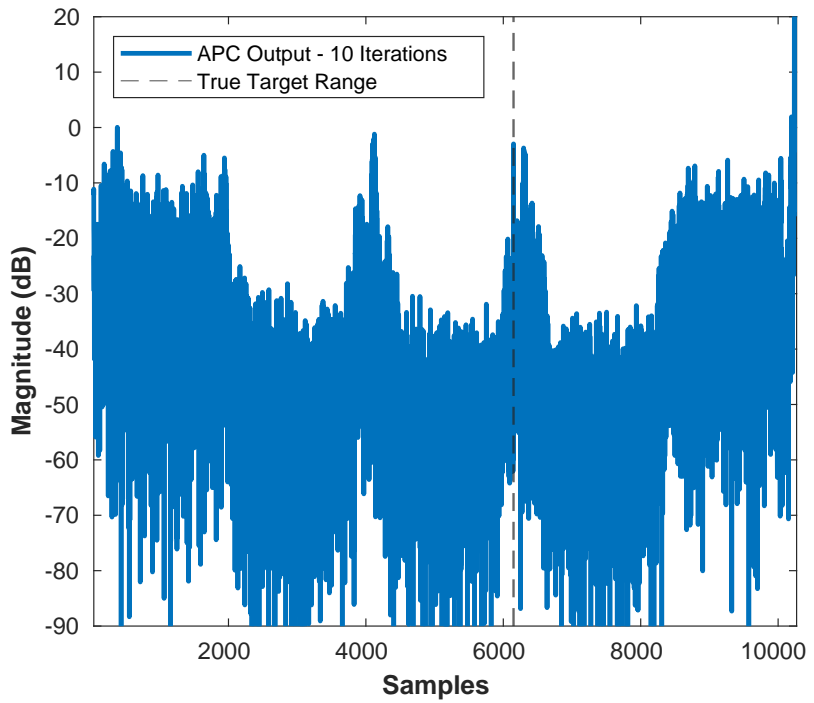


Figure 6.35: APC filter output after 10 iterations using the covariance matrix derived from the demod/remod estimate.

radar system separate from the radar return. This goal means that the demod/remod estimate should provide a reasonable starting point for constructing a useful interference covariance matrix.

Figure 6.34 shows the Frobenius normalized covariance matrix calculated from the robust demod/remod estimate described in the previous section. Again, this covariance is derived from a single snapshot of data; therefore, there is not any visible structure, and there are long-range covariance values populated. The demod/remod covariance is an estimate that shares some similarity with the truth covariance in scale and structure.

Figure 6.35 shows the output of the filter after ten iterations of APC. Several familiar characteristics are represented in the filter output. First, the target response is successfully focused in the correct location and approximate magnitude. Next, increased residual interference is present in the near and far ranges. This residue is similar to the interference residue that was left after demod/remod subtraction in section 6.2.

Additionally, the residue is gathered on the symbol boundaries near samples 4,000 and 6,500. Some of the residue present maintains much of the original interference power and is not reduced at all. Residue at the symbol boundary was also seen in section 6.2, causing ISI in the cyclic prefix that was not captured by the demod/remod estimate. Here, the model mismatches that were present in the previous section, much of which was alleviated by the addition of the median filter, are now impacting the performance of the APC algorithm in very significant ways.

The difference in these demod/remod covariance estimates relative to the ground truth covariance can be captured by looking at the difference in the normalized covariance. Figure 6.36 shows the difference in the covariance matrix derived from ground truth compared to the demod/remod estimate.




Figure 6.36: Difference in the normalized ground truth covariance matrix and the demod/remod estimate covariance matrix.

Here, the major differences at the near and far ranges exhibit the strongest difference, with smaller differences existing at the symbol boundaries. From this image the square sub-regions can be seen that correspond to the size of the symbols, bounded by small rows with a larger difference. The covariance behavior exhibits exactly the behavior that would be expected given the characteristics seen in the filter output.

The regions of residual interference left behind by APC are larger in both range and magnitude compared to the direct subtraction technique. Although the residue is present in many of the same locations as the demod/remod technique, the APC residue takes a wildly different shape. Ultimately, APC proves to be far more sensitive to misestimation of the interference source than the direct subtraction technique.

Finally, the notional sample covariance matrix can be used as a possible alternative to supply APC with an interference covariance matrix. When creating a sample co-

variance matrix, the goal is to capture the actual structure of the data as opposed to the single sample variances that have no structure, as seen above in Figures 6.32 and 6.34. In order to capture a sample covariance that minimizes loss to within 3 dB, $2N - 3$ snapshots are needed for the covariance matrix, where N is the number of elements in the set [98].

Accordingly, for a fast time range extent of 8,735 samples, 17,470 snapshots of interference data need to be generated. So each snapshot was generated with uniform random data, a timing offset, frequency offset, and free space channel estimates derived from the demod/remod process were applied to each, providing the same structural changes to each OFDM as was received. Figure 6.37 shows the sample covariance matrix created using 17,470 snapshots and a closer look at the channel effects on the covariance matrix.

There are many interesting features present in the sample covariance matrix. First, the strong diagonal through the matrix represents a Gaussian component that comes from the fact that OFDM symbol bodies appear Gaussian to any unsynchronized receiver over a short enough interval. Each symbol can be easily distinguished when looking down the matrix diagonal.

On each symbol around the diagonal, the cyclic prefix is easily visible, caused by the relationship between that prefix and the last samples of the symbols. Additionally, every other symbol has pilots attached. The pilots in the symbol result in a set of super diagonals above and below the main diagonal. The number of super diagonals on each side of the main diagonal is directly related to the spacing between pilots. For this work, the pilot spacing is five, resulting in five super diagonals above and below the main diagonal. The sample covariance also shows different levels of correlation with one another, represented by the different baseline energy levels in the off-diagonal

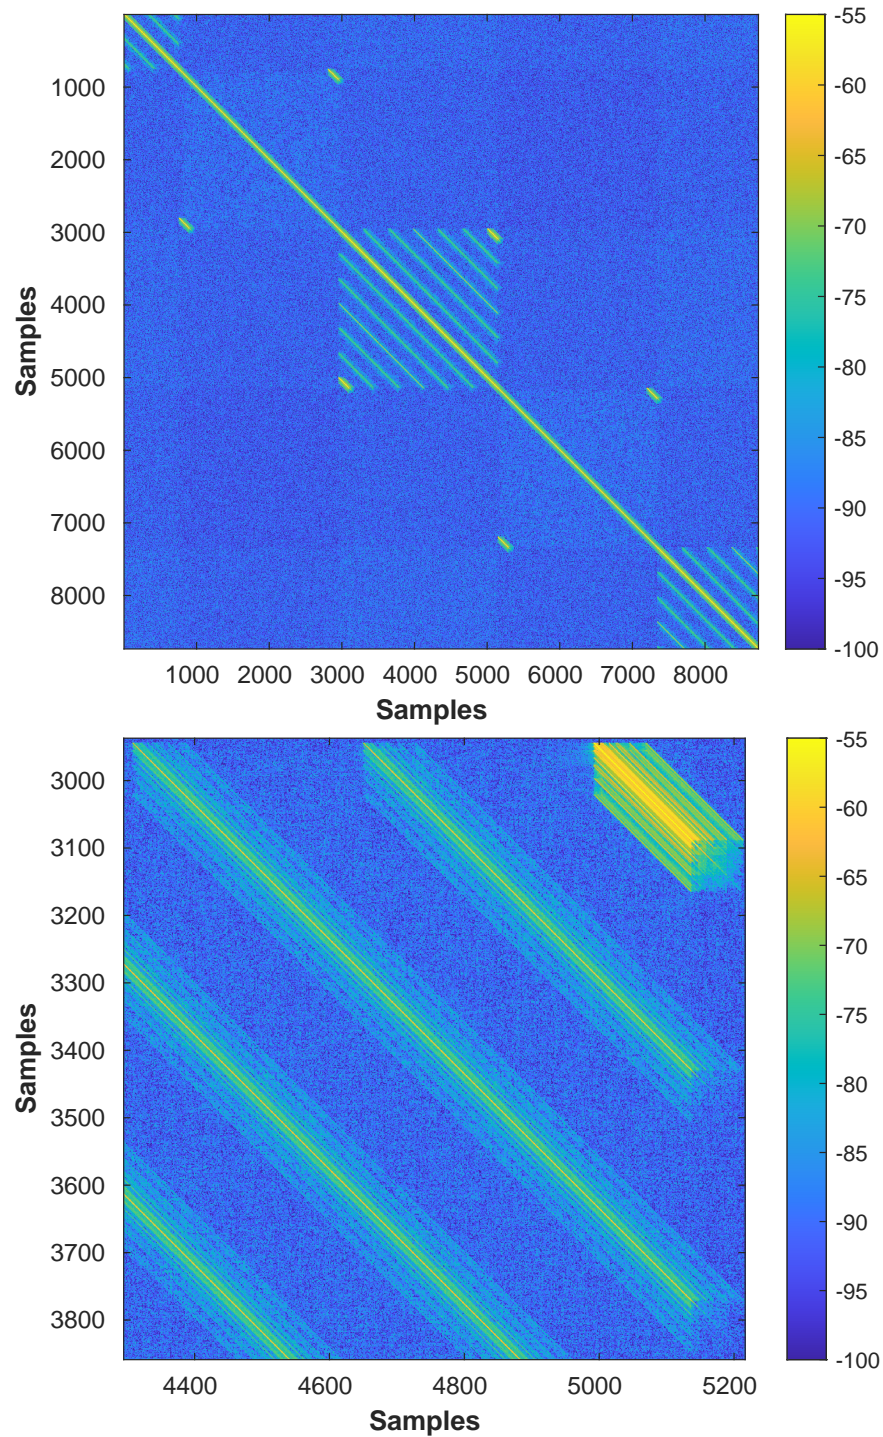


Figure 6.37: The sample covariance generated using OFDM snapshots with random data, that shares symbol structure and channel effects with the interference signal with a zoomed in look at the channel effects on one of the structures.




Figure 6.38: The APC filter response after ten iterations using the sample covariance of the interference.

grids representing the correlation between symbols.

For the biggest chance of success, the channel estimate was built into the interference snapshots that were used to make the sample covariance. The effect of the channel on the covariance is clearly visible when looking closely at one of the cyclic prefix structures. Just like the channel starts with large multipath taps and fades lower, the covariance matrix demonstrates a correlation between adjacent symbols and tapers off, creating a smeared grid of each structure.

Figure 6.38 shows the APC filter output that is generated after ten iterations using the sample covariance matrix. Here, the target is resolved at the correct location and approximate magnitude, with less interference residue demonstrated across the range extent. However, there is impulsive residue left behind with a narrow target-like re-

sponse in range; a detector would register the peaks as target detections, drastically increasing false alarms. The use of the sample covariance provides a middle ground between the performance of the ground truth covariance and the demod/remod covariance. Although it is a significant improvement over the demod/remod performance, it still performs worse than the robust demod/remod estimation and extraction technique provided above.

Ultimately, the APC technique requires a very good estimate in order to properly cancel the interference and is very sensitive to the quality of the estimate. This sensitivity, held in tandem with the computational complexity of the APC, makes it less desirable than the demod/remod estimate approach.

6.6 Summary

This chapter provided a wide-ranging exploration of demod/remod interference estimation both as a method of direct subtraction and as a method for informing a more advanced adaptive interference cancellation technique. This chapter used a simple model to verify that the demod/remod estimation and extraction approach would be effective when data is lost due to receiver blanking. After verifying the efficacy of the technique, the model was expanded to include a more robust real-world model.

Several techniques were introduced from the literature to help estimate and counteract the effects of time and frequency offsets and the free space channel. Each of these techniques were presented and added to a total design flow. Finally, median filtering was introduced as the last step of the processing chain to help clean up the residual interference left behind due to misestimation.

The performance of the full processing chain was examined using absolute interference reduction, target peak loss, and radar detection metrics. Altogether, these met-

rics properly present the trade space of the technique's performance across a range of duty cycles and ISR values. The results show a final net **19 dB** performance improvement in target detection, uncovering targets that are buried well below the interference with minimal increase in false alarm rate.

Finally, the chapter concluded by examining the performance of an RMMSE based technique call APC. The interference can be perfectly canceled when APC is provided with a perfect estimate of the interference. However, the technique struggles when applying an estimate or sample covariance in place of truth data. Ultimately, APC fails to provide results comparable to the demod/remod estimation and extraction approach developed earlier in the chapter.

Chapter 7

Conclusion and Future Work

This work has looked at a pair of problems across two separate domains in order to provide a solution to a singular problem: How can a radar system ensure continual operation as a secondary user in congested spectrum?

First, Chapter 1 presents an introduction to the congested spectrum problem and existing cognitive radar and spectrum sharing techniques were presented. The HO-CAE algorithm was presented to motivate the discussion of dynamic spectrum access and frequency and bandwidth agility. Additionally, Chapter 2 provides an overview of traditional range-Doppler processing, CFAR detection, and OFDM communications for LTE, 5G, and WiFi were presented to set the baseline of discussion.

As a first step, this work considers the case when spectrum is available, but the radar system must hop around, changing both frequency and bandwidth on a pulse-to-pulse basis in order get access to the spectrum. The hopping radar problem is examined in two parts. First, chapter 3 presents the case when the center frequency of the radar system stays fixed over the course of a CPI but the bandwidth used by the radar changes on each pulse.

Sidelobe misalignment of the matched filter response was identified as the main contribution to the range-sidelobe modulation that was spread across range and Doppler.

As a result, section 3.1 presented the mismatched filter as a method for creating a common pulse compression response across all pulses in the CPI. Further, section 3.2 investigates the different common responses that can be generated by the mismatched filter, with section 3.3 identifying the super-resolution response as an appropriate trade-off between resolution and RSM.

After the mismatched filter was investigated, Chapter 4, the bandwidth is held constant but the center frequency of the pulse changes on each pulse. When changing the center frequency, the slow time phase history is distorted. Section 4.1 analyses the phase history and identifies a range dependent phase correction term to re-align the phase history. However, because the correction term is range dependent and the target changes range over the course of the CPI, the correction cannot be applied directly as even minor changes can result in an increased phase error over the CPI.

Since every range-Doppler cell represents a target with a different velocity, resulting in a different change in position over the CPI, each cell needs its own set of correction. Section 4.2 presented a modified backprojection algorithm that can focus each range-Doppler cell independently, each with its own set of corrections. Section 4.4 verifies the performance of the modified backprojection algorithm using real data from an outdoor radar experiment.

Because of the complimentary nature of the mismatched filter and the modified backprojection algorithm, both can be applied together. Chapter 5 looks at the performance of combining the techniques to process a CPI where each pulse changes in both bandwidth and center frequency. The full agile processing chain is demonstrated on a real-world, outdoor data.

With a technique developed to process frequency and bandwidth agile pulses, the focus then turned to the second mode of operation. When spectrum is not available to

support a mission critical radar task, the radar will need to accept interference from other users of the spectrum. Chapter 6 presented the demodulation/remodulation estimation and extraction technique that can be used to actively cancel OFDM interference.

Starting in section 6.1 a simple model was used to measure initial performance before the model was expanded in section 6.2. Techniques were presented from the literature to overcome new complexity introduced by the more robust model. Median filtering was presented in 6.3.1 as a method for removing the low duty cycle residue that was left behind due to misestimation of the received OFDM. The full performance of the technique was analyzed in section 6.4 and identified a 19 dB performance improvement in CFAR detection.

Section 6.5 investigated an alternative reiterative technique called adaptive pulse compression. APC uses the interference covariance to reject the interference. Section 6.5.2 showed that APC is incredibly sensitive to misestimation of the covariance matrix including that the demod/remod estimate and a sample covariance of a stream of OFDM symbols were insufficient to fully cancel the interference. Ultimately, the chapter concluded by identify the demod/remod estimation and extraction technique as a more performant and computationally efficient technique.

7.1 Conclusion

This work provides a set of signal processing techniques that enable a radar system to operate in congested spectrum. When changing bandwidth and center frequency on a pulse-to-pulse basis cause targets to smear in both range and Doppler, usually completely scattering the target energy in Doppler so that there is no longer a target peak. Separately, if there is not any available bandwidth and interference is present in

the received signal, the targets may be completely masked by interference.

This work presented the mismatched filter to correct range sidelobe modulation caused by changing the bandwidth on a pulse to pulse basis. Mismatched filtering had traditionally been applied to unify the response of phase attached waveforms. Here the mismatched filter is used to coerce a common pulse compression response to align the mainlobe and sidelobes across pulse and reduce RSM.

Second, a modified backprojection technique was developed that corrected the range dependent phase term that contaminates the slow time phase history when the center frequency changes on a pulse-to-pulse basis. The modified backprojection technique was able to correctly focus targets when applied to both simulated and real-world experimental data.

Because mismatched filtering and the modified backprojection algorithm operate across different domains, they fit well together to provide a robust solution to CPIs of data where the center frequency and bandwidth change on every pulse. The combined processing chain was applied and verified on real data captured on a COTS SDR.

When there is not enough spectrum to hop to an available band, interference is accepted at the radio. When interference is present, targets may be completely masked. This work developed a robust OFDM interference cancellation technique that estimates and demodulates the OFDM signals down to an estimate of the bits, while also estimating time offset, Doppler frequency, and free space channel that is present as part of the interference. Using all of these parameters an estimate of the interference signal is generated by remodulating the data onto OFDM symbols and reapplying all of the estimated channel effects.

The remod signal is subtracted from the real data. After subtraction some interference residue persisted at the symbol boundaries because of ISI. However, the effective

duty cycle of the interference was reduced down from 100%, enabling a median filter to be applied to the data.

Using the full processing chain developed in this work, interference can be effectively reduced to unmask targets. Using the demod/remod estimation and extraction technique yields a **19 dB** increase in CFAR detection performance. Additionally, the demod/remod processing chain provided better interference cancellation performance compared to an the computationally intensive APC.

APC is very successful at canceling interference when the interference covariance is well known. However, the technique is very sensitive to covariance mismatch and failed to outperform demod/remod when the using an estimate covariance or sample covariance.

In conclusion, this work developed a wide range of tools to allow radar systems to operate in congested spectrum. When utilizing real time spectrum sensing, modified backprojection technique and mismatched filter can be utilized to focus targets from a set of non-coherent pulses that are changing bandwidth and frequency over a CPI.

Alternatively, when there is not sufficient spectrum to enable hopping, the demodulation/remodulation estimation and extraction can be used to successfully remove OFDM interference and uncover targets. Both of these techniques together will ensure that a mission critical radar system can always produce useful data products, even in the presence of a congested spectrum.

7.2 Future Work

With the development of techniques to enable frequency agile radar brings additional challenges that have not been considered in this work. As the center frequency changes significantly, the returned energy of the target can change dramatically as

the primary scattering surface of a target can change relative to the transmitted wavelength. Future work should consider these effects in tandem with the results presented here to inform a cost function to make more intelligent decisions about dynamic spectrum sensing.

Further, sidelobes of the pulse compression response carry the phase history of the target peak. When applying the modified backprojection technique, these sidelobes create a phase error when collecting energy for the adjacent range bins, so the sidelobe energy is distributed across Doppler instead of focused to form an x-shape as would traditionally be effected. Future work could be carried out to create a phase sensitive mismatch filter that could carry the appropriate phase shift out along the sidelobes to so that even the sidelobes will be focused.

When accepting communication interference, more work can be carried out to identify shortcomings in the sample covariance matrix to provide more robust cancellation from APC. When considering the sample covariance matrix, it would be worthwhile to investigate structures that could be added to the OFDM signal model that would enable easier cancellation. Further, can radar waveforms be developed that provide additional separation between that APC can leverage to enable better sample covariance performance.

Overall, many new and exciting directions of investigation are available that would further enrich the tools that were developed here.

References

- [1] J. G. Metcalf, "An examination of the spectral utility of radar," in *2022 IEEE Radar Conference (RadarConf22)*, 2022, pp. 1–6.
- [2] B. H. Kirk, R. M. Narayanan, K. A. Gallagher, A. F. Martone, and K. D. Sherbondy, "Avoidance of time-varying radio frequency interference with software-defined cognitive radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, pp. 1090–1107, 6 2019.
- [3] S. Bhattarai, J. M. Park, B. Gao, K. Bian, and W. Lehr, "An overview of dynamic spectrum sharing: Ongoing initiatives, challenges, and a roadmap for future research," *IEEE Transactions on Cognitive Communications and Networking*, vol. 2, pp. 110–128, 6 2016.
- [4] J. H. Reed, A. W. Clegg, A. V. Padaki, T. Yang, R. Nealy, C. Dietrich, C. R. Anderson, and D. M. Mearns, "On the co-existence of td-lte and radar over 3.5 ghz band: An experimental study," *IEEE Wireless Communications Letters*, vol. 5, pp. 368–371, 8 2016.
- [5] P. Stinco, M. Greco, F. Gini, and B. Himed, "Cognitive radars in spectrally dense environments," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 10, pp. 20–27, 2016.
- [6] P. Stinco, M. S. Greco, and F. Gini, "Spectrum sensing and sharing for cognitive radars," *IET Radar, Sonar & Navigation*, vol. 10, no. 3, pp. 595–602, 2016.
- [7] M. S. Greco, F. Gini, P. Stinco, and K. Bell, "Cognitive radars: On the road to reality: Progress thus far and possibilities for the future," *IEEE Signal Processing Magazine*, vol. 35, no. 4, pp. 112–125, 2018.
- [8] A. F. Martone *et al.*, "Closing the loop on cognitive radar for spectrum sharing," *IEEE Aerospace and Electronic Systems Magazine*, vol. 36, pp. 44–55, 9 2021.
- [9] J. M. Peha, "Sharing spectrum through spectrum policy reform and cognitive radio," *Proceedings of the IEEE*, vol. 97, pp. 708–719, 2009.

- [10] L. Kulacz, P. Kryszkiewicz, A. Kliks, H. Bogucka, J. Ojaniemi, J. Paavola, J. Kalliovaara, and H. Kokkinen, “Coordinated spectrum allocation and coexistence management in cbrs-sas wireless networks,” *IEEE Access*, vol. 7, pp. 139 294–139 316, 2019.
- [11] A. W. Clegg, S. A. Seguin, R. J. Marks, and C. Baylis, “Radar sharing in the u.s. 3 ghz band,” in *2022 IEEE Radar Conference (RadarConf22)*, 2022, pp. 1–5.
- [12] E. Lagunas, C. G. Tsinos, S. K. Sharma, and S. Chatzinotas, “5g cellular and fixed satellite service spectrum coexistence in c-band,” *IEEE Access*, vol. 8, pp. 72 078–72 094, 2020.
- [13] J. Jeon, H. Niu, Q. Li, A. Papathanassiou, and G. Wu, “Lte with listen-before-talk in unlicensed spectrum,” in *2015 IEEE International Conference on Communication Workshop (ICCW)*, 6 2015, pp. 2320–2324, corrected.
- [14] A. F. Martone, K. I. Ranney, K. Sherbondy, K. A. Gallagher, and S. D. Blunt, “Spectrum allocation for noncooperative radar coexistence,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, pp. 90–105, 2 2018, corrected.
- [15] R. G. Mattingly and J. G. Metcalf, “Hardware optimized cell averaging estimation (ho-cae) for threshold based signal detection,” in *2022 IEEE Radar Conference (RadarConf22)*. IEEE, 3 2022, pp. 1–6.
- [16] R. G. Mattingly and J. G. Metcalf, “Fast adaptive spectrum sensing using hardware optimized, cell averaging estimation for cognitive radio and radar applications,” *IEEE Transactions on Aerospace and Electronic Systems*, 2023.
- [17] H. Griffiths, L. Cohen, S. Watts, E. Mokole, C. Baker, M. Wicks, and S. Blunt, “Radar spectrum engineering and management: Technical and regulatory issues,” *Proceedings of the IEEE*, vol. 103, no. 1, pp. 85–102, Jan. 2015.
- [18] F. Sanders, R. Sole, B. Bedford, D. Franc, and T. Pawlowitz, “Effects of rf interference on radar receivers, ntia report tr-06-444,” *US Department of Commerce, Tech. Rep*, 2006.
- [19] F. H. Sanders, R. L. Sole, J. E. Carroll, G. S. Secrest, and T. L. Allmon, *Analysis and resolution of RF interference to radars operating in the band 2700-2900 MHz from broadband communication transmitters*. US Department of Commerce, National Telecommunications and Information . . . , 2012.
- [20] D. P. Zilz and M. R. Bell, “Statistical modeling of wireless communications interference and its effects on adaptive-threshold radar detection,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 2, pp. 890–911, 2018.

- [21] J. G. Metcalf and S. Flandermeyer, "On spectrum sharing for pulse-doppler radar and ofdm communications," in *2020 IEEE Radar Conference (RadarConf20)*, 2020, pp. 1–6.
- [22] M. A. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill Education, 2014.
- [23] D. J. Rabideau, "Nonlinear synthetic wideband waveforms," *IEEE National Radar Conference - Proceedings*, pp. 212–219, 2002.
- [24] N. Levanon and E. Mozeson, "Nullifying acf grating lobes in stepped-frequency train of lfm pulses," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, pp. 694–703, 4 2003.
- [25] N. Neuberger and R. Vehmas, "Range sidelobe level reduction with a train of diverse lfm pulses," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, pp. 1480–1486, 4 2022.
- [26] R. Calderbank, S. D. Howard, and B. Moran, "Waveform diversity in radar signal processing: A focus on the use and control of degrees of freedom," *IEEE Signal Processing Magazine*, vol. 26, pp. 32–41, 2009.
- [27] T. Higgins, K. Gerlach, A. K. Shackelford, and S. D. Blunt, "Aspects of non-identical multiple pulse compression," *IEEE National Radar Conference - Proceedings*, pp. 895–900, 2011.
- [28] T. Higgins, S. D. Blunt, and A. K. Shackelford, "Time-range adaptive processing for pulse agile radar," *2010 International Waveform Diversity and Design Conference, WDD 2010*, pp. 115–120, 2010.
- [29] B. H. Kirk, A. F. Martone, K. A. Gallagher, R. M. Narayanan, and K. D. Sherrbondy, "Mitigation of clutter modulation in cognitive radar for spectrum sharing applications," *IEEE Transactions on Radar Systems*, pp. 1–1, 6 2023.
- [30] J. E. Carroll, "Case study: Investigation of interference into 5 ghz weather radars from unlicensed national information infrastructure devices, part 3," Institute for Telecommunication Sciences, Tech. Rep., 2012.
- [31] M. I. Skolnik, "Introduction to radar systems," *New York*, 1980.
- [32] C. Munnell, R. Mattingly, S. Flandermeyer, and J. G. Metcalf, "On the practical use and experimentation of lte signals for radar-communications." *IEEE*, 3 2022, pp. 1–6.

- [33] R. E. Jarvis, J. G. Metcalf, and J. W. McDaniel, "Application of adaptive pulse compression in cluttered radar cross section measurements," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–8, 2022.
- [34] J. G. Proakis, *Digital signal processing: principles, algorithms, and applications, 4/E*. Pearson Education India, 2007.
- [35] T. F. Collins, R. Getz, D. Pu, and A. M. Wyglinski, *Software-Defined Radio for Engineers*. Artech House, 2018.
- [36] A. Mukherjee, *5G New Radio: Beyond Mobile Broadband*. Artech House, 2019.
- [37] T. Innovations, "Lte in a nutshell," *White paper*, 2010.
- [38] E. Khorov, A. Kiryanov, A. Lyakhov, and G. Bianchi, "A tutorial on ieee 802.11ax high efficiency w lans," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 197–216, Firstquarter 2019.
- [39] P. Mogensen, W. Na, I. Z. Kovacs, F. Frederiksen, A. Pokhariyal, K. I. Pedersen, T. Kolding, K. Hugel, and M. Kuusela, "Lte capacity compared to the shannon bound," in *2007 IEEE 65th Vehicular Technology Conference - VTC2007-Spring*, April 2007, pp. 1234–1238.
- [40] 3GPP, "Base station (bs) radio transmission and reception," 3rd Generation Partnership Project; Evolved Universal Terrestrial Radio Access (E-UTRA), Tech. Rep. TR 36.104 V8.2.0 (2008-05), 2006.
- [41] Y. Liu, Z. Tan, H. Hu, L. J. Cimini, and G. Y. Li, "Channel estimation for ofdm," *IEEE Communications Surveys and Tutorials*, vol. 16, pp. 1891–1908, 4 2014.
- [42] 3GPP, "Multiplexing and channel coding," 3rd Generation Partnership Project; Evolved Universal Terrestrial Radio Access (E-UTRA), Tech. Rep. TS 36.212, 2006.
- [43] S. M. Dudley *et al.*, "Practical issues for spectrum management with cognitive radios," *Proceedings of the IEEE*, vol. 102, pp. 242–264, 2014.
- [44] A. E. Mitchell, G. E. Smith, K. L. Bell, A. J. Duly, and M. Rangaswamy, "Hierarchical fully adaptive radar," *IET Radar, Sonar & Navigation*, vol. 12, pp. 1371–1379, 12 2018.
- [45] A. E. Mitchell, G. E. Smith, K. L. Bell, A. Duly, and M. Rangaswamy, "Fully adaptive radar cost function design," *2018 IEEE Radar Conference, RadarConf 2018*, pp. 1301–1306, 6 2018.

- [46] K. L. Bell, C. J. Baker, G. E. Smith, J. T. Johnson, and M. Rangaswamy, “Cognitive radar framework for target detection and tracking,” *IEEE Journal on Selected Topics in Signal Processing*, vol. 9, pp. 1427–1439, 12 2015.
- [47] J. A. Kovarskiy, B. H. Kirk, A. F. Martone, R. M. Narayanan, and K. D. Sherbondy, “Evaluation of real-time predictive spectrum sharing for cognitive radar,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 1, pp. 690–705, Feb 2021.
- [48] J. A. Devault, R. M. Narayanan, and A. F. Martone, “Tdd mode identification for spectrum sharing applications,” *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–11, 2024.
- [49] J. A. Kovarskiy, R. M. Narayanan, A. F. Martone, and K. D. Sherbondy, “A stochastic model for prediction and avoidance of rf interference to cognitive radars,” *2019 IEEE Radar Conference, RadarConf 2019*, 4 2019.
- [50] B. H. Kirk, M. A. Kozy, K. A. Gallagher, R. M. Narayanan, R. M. Buehrer, A. F. Martone, and K. D. Sherbondy, “Cognitive software-defined radar: Evaluation of target detection with rfi avoidance,” 4 2019, pp. 1–6, corrected.
- [51] R. G. Mattingly, “Implementation and analysis of adaptive spectrum sensing,” Master’s thesis, University of Oklahoma, 2021.
- [52] R. G. Mattingly, A. F. Martone, and J. G. Metcalf, “Techniques for mitigating the impact of intra-cpi waveform agility,” *IEEE Transactions on Radar Systems*, vol. 2, pp. 24–40, 2024.
- [53] F. Harris, “On the use of windows for harmonic analysis with the discrete fourier transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978.
- [54] M. H. Ackroyd and F. Ghani, “Optimum mismatched filters for sidelobe suppression,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-9, pp. 214–218, 1973.
- [55] C. Sahin, J. G. Metcalf, and S. D. Blunt, “Filter design to address range sidelobe modulation in transmit-encoded radar-embedded communications,” *2017 IEEE Radar Conference, RadarConf 2017*, pp. 1509–1514, 6 2017.
- [56] L. A. Harnett and S. D. Blunt, “Least-squares optimal mismatched doppler processing,” *2019 IEEE Radar Conference, RadarConf 2019*, 4 2019.
- [57] D. Henke, P. McCormick, S. D. Blunt, and T. Higgins, “Practical aspects of optimal mismatch filtering and adaptive pulse compression for fm waveforms,” *2015 IEEE Radar Conference (RadarCon)*, pp. 1149–1155, 5 2015.

- [58] S. D. Blunt and E. L. Mokole, "Overview of radar waveform diversity," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 11, pp. 2–42, 2016.
- [59] R. G. Mattingly, A. F. Martone, and J. G. Metcalf, "Enabling intra-cpi frequency agility via backprojection based range-doppler processing," *2023 International Radar Conference (RADAR)*, pp. 1–6, 2023.
- [60] R. P. Perry, R. C. DiPietro, and R. L. Fante, "Sar imaging of moving targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, pp. 188–200, 1999.
- [61] F. Uysal and N. Goodman, "The effect of moving target on range-doppler map and backprojection algorithm for focusing," *2016 IEEE Radar Conference, RadarConf 2016*, 6 2016.
- [62] M. B. Heintzelman, J. W. Owen, S. D. Blunt, B. Maio, and E. D. Steinbach, "Practical considerations for optimal mismatched filtering of nonrepeating waveforms," in *2023 IEEE Radar Conference (RadarConf)*, 2023, pp. 1–6.
- [63] J. W. Owen, B. Ravenscroft, and S. D. Blunt, "Devoid clutter capture and filling (deccaf) to compensate for intra-cpi spectral notch variation," *2019 International Radar Conference, RADAR 2019*, 9 2019.
- [64] B. Ravenscroft, J. W. Owen, B. H. Kirk, S. D. Blunt, A. F. Martone, K. D. Sherbondy, and R. M. Narayanan, "Experimental assessment of joint range-doppler processing to address clutter modulation from dynamic radar spectrum sharing," *2020 IEEE International Radar Conference, RADAR 2020*, pp. 448–453, 4 2020.
- [65] B. H. Kirk, A. F. Martone, K. D. Sherbondy, and R. M. Narayanan, "Mitigation of target distortion in pulse-agile sensors via richardson-lucy deconvolution," *Electronics Letters*, vol. 55, pp. 1249–1252, 11 2019.
- [66] Y. Cao, S. Guo, S. Jiang, X. Zhou, X. Wang, Y. Luo, Z. Yu, Z. Zhang, and Y. Deng, "Parallel optimisation and implementation of a real-time back projection (bp) algorithm for sar based on fpga," *Sensors*, vol. 22, no. 6, 2022.
- [67] J. Price, "The intersection of radar and communications: A study on spectrum management for addressing rf interference," Master's thesis, The University of Oklahoma, May 2023.
- [68] H. Kuschel, D. Cristallini, and K. E. Olsen, "Tutorial: Passive radar tutorial," *IEEE Aerospace and Electronic Systems Magazine*, vol. 34, pp. 2–19, 2 2019.
- [69] M. Malanowski, *Signal processing for passive bistatic radar*. Artech House, 2019.

- [70] M. Edrich, S. Lutz, and F. Hoffmann, "Passive radar at hensoldt: A review to the last decade," *Proceedings International Radar Symposium*, vol. 2019-June, 6 2019.
- [71] W. Klembowski, A. Kawalec, and W. Wizner, "Critical views on present passive radars performance as compared with that of active radars," in *2013 14th International Radar Symposium (IRS)*, vol. 1, 2013, pp. 131–135.
- [72] A. Ksiezyk, M. Plotka, K. Abratkiewicz, R. Maksymiuk, J. Wszolek, P. Samczynski, and T. P. Zielinski, "Opportunities and limitations in radar sensing based on 5g broadband cellular networks," *IEEE Aerospace and Electronic Systems Magazine*, 2023.
- [73] L. A. Harnett, "Mismatched processing for radar interference cancellation," Ph.D. dissertation, University of Kansas, August 2022.
- [74] R. Mattingly, N. Goodman, and J. Metcalf, "Performance analysis of ofdm interference mitigation via demodulation/remodulation estimation and extraction," in *2024 IEEE Radar Conference (RadarConf24)*, 2024, pp. 1–6.
- [75] S.-J. Lee, M. Goel, Y. Zhu, J.-F. Ren, and Y. Sun, "Forward error correction decoding for wimax and 3gpp lte modems," in *2008 42nd Asilomar Conference on Signals, Systems and Computers*, Oct 2008, pp. 1143–1147.
- [76] L. Han, F. Gao, Z. Li, and O. A. Dobre, "Low complexity automatic modulation classification based on order-statistics," *IEEE Transactions on Wireless Communications*, vol. 16, no. 1, pp. 400–411, Jan 2017.
- [77] J. P. Mouton, M. Ferreira, and A. S. Helberg, "A comparison of clustering algorithms for automatic modulation classification," *Expert Systems with Applications*, vol. 151, p. 113317, 2020.
- [78] F. Hameed, O. A. Dobre, and D. C. Popescu, "On the likelihood-based approach to modulation classification," *IEEE Transactions on Wireless Communications*, vol. 8, no. 12, pp. 5884–5892, December 2009.
- [79] J. J. Van De Beek, M. Sandell, and P. O. Börjesson, "Ml estimation of time and frequency offset in ofdm systems," *IEEE Transactions on Signal Processing*, vol. 45, pp. 1800–1805, 1997.
- [80] A. Baghaki and B. Champagne, "Joint frequency offset, time offset, and channel estimation for ofdm/oqam systems," *Eurasip Journal on Advances in Signal Processing*, vol. 2018, pp. 1–19, 12 2018.

- [81] A. Filippi and S. Serbetli, "Ofdm symbol synchronization using frequency domain pilots in time domain," *IEEE Transactions on Wireless Communications*, vol. 8, pp. 3240–3248, 6 2009.
- [82] Z. Zhang, K. Long, M. Zhao, and Y. Liu, "Joint frame synchronization and frequency offset estimation in ofdm systems," *IEEE Transactions on Broadcasting*, vol. 51, pp. 389–394, 9 2005.
- [83] Q. Huang, M. Ghogho, J. Wei, and P. Ciblat, "Practical timing and frequency synchronization for ofdm-based cooperative systems," *IEEE Transactions on Signal Processing*, vol. 58, pp. 3706–3716, 7 2010.
- [84] M. K. Ozdemir and H. Arslan, "Channel estimation for wireless ofdm systems," *IEEE Communications Surveys and Tutorials*, vol. 9, pp. 18–48, 6 2007.
- [85] N. Van Duc, H. P. Kuchenbecker, and M. Pätzold, "Estimation of the channel impulse response length and the noise variance for ofdm systems," *IEEE Vehicular Technology Conference*, vol. 61, pp. 429–433, 2005.
- [86] M. Juhola, J. Katajainen, and T. Raita, "Comparison of algorithms for standard median filtering," *IEEE Transactions on Signal Processing*, vol. 39, no. 1, pp. 204–208, 1991.
- [87] J. L. Lake, M. Yeary, and C. D. Curtis, "Adaptive radio frequency interference mitigation techniques at the national weather radar testbed: First results," in *2014 IEEE Radar Conference*, 2014, pp. 0840–0845.
- [88] H. Urkowitz, "Energy detection of unknown deterministic signals," *Proceedings of the IEEE*, vol. 55, no. 4, pp. 523–531, April 1967.
- [89] J. Y. Cho, "A new radio frequency interference filter for weather radars," *Journal of Atmospheric and Oceanic Technology*, vol. 34, pp. 1393–1406, 7 2017.
- [90] S. Blunt and K. Gerlach, "Adaptive pulse compression via mmse estimation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 2, pp. 572–584, April 2006.
- [91] S. D. Blunt, T. Chan, and K. Gerlach, "Robust doa estimation: The reiterative super-resolution (risr) algorithm," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 1, pp. 332–346, January 2011.
- [92] S. Blunt and K. Gerlach, "Adaptive pulse compression," in *Proceedings of the 2004 IEEE Radar Conference (IEEE Cat. No.04CH37509)*, April 2004, pp. 271–276.
- [93] S. S. Haykin, *Adaptive filter theory*. Pearson Education India, 2002.

- [94] H. Van Trees, *Optimum array processing: Part IV of detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [95] T. Higgins, S. D. Blunt, and K. Gerlach, "Gain-constrained adaptive pulse compression via an mvdr framework," in *2009 IEEE Radar Conference*, May 2009, pp. 1–6.
- [96] S. D. Blunt and T. Higgins, "Dimensionality reduction techniques for efficient adaptive pulse compression," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 1, pp. 349–362, Jan 2010.
- [97] S. D. Blunt and T. Higgins, "Achieving real-time efficiency for adaptive radar pulse compression," in *2007 IEEE Radar Conference*, April 2007, pp. 116–121.
- [98] I. Reed, J. Mallett, and L. Brennan, "Rapid convergence rate in adaptive arrays," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-10, no. 6, pp. 853–863, Nov 1974.