THESIS

DESIGN, DEPLOYMENT, AND COST CONSIDERATIONS FOR DARMA; A LOW-COST AND LIGHTWEIGHT FMCW RADAR

Submitted by

Marshall Bruner

Department of Electrical and Computer Engineering

In partial fulfillment of the requirements

For the Degree of Master of Science

Colorado State University

Fort Collins, Colorado

Summer 2022

Master's Committee:

Advisor: V. Chandrasekar

Margaret Cheney
Ryan Gooch

ABSTRACT

DESIGN, DEPLOYMENT, AND COST CONSIDERATIONS FOR DARMA;
A LOW-COST AND LIGHTWEIGHT FMCW RADAR

The capability of frequency-modulated continuous wave (FMCW) radar to operate in low-power environments has made it a good choice for many mobile systems including automobile radars. While specialized FMCW radars have seen an increase in production recently, there is a lack of general-purpose FMCW radars with the ability to be used in a multitude of applications, especially for volume targets such as precipitation. This thesis presents design considerations for the Dual-polarization phased Array Radar for Measurement of the Atmosphere (DARMA), a low-cost, medium range (km) radar with the versatility to operate mounted on an unmanned aircraft system (UAS) or ground platform. The radar features modular subsystems which allow for easy swapping to support different application requirements as well as upgrades due to rapidly changing technology. Signal processing methods are also introduced, and implemented on COTS systems, to allow for noise mitigation, target detection, and estimation of weather products.

ACKNOWLEDGEMENTS

I would first like to thank my advisor, Dr. Chandra, and my committee members, Dr. Margaret Cheney and Dr. Ryan Gooch. Thank you Dr. Chandra for giving me real engineering work and entrusting so much responsibility in me. Through your guidance, jokes, and advice you have greatly impacted my skill and confidence as an engineer, and for that I am grateful. Thank you Dr. Cheney and Dr. Gooch for giving me the opportunity to defend my thesis.

Thank you to the incredibly talented engineers at the CSU-CHILL radar facility, Jim George, Francesc Junyent, and Pat Kennedy, for letting me act as a sponge while you talked about topics way over my head. Thank you, Jim for providing insightful feedback without sugar-coating it and showing me the importance of choosing a good text editor. Thank you, Francesc for your stories and for taking the time to show and explain your projects. Thank you, Pat for being willing to respond to my curious emails and showing me how fascinating meteorology can be.

Thank you to all my colleagues at Ball Aerospace who were always willing to take time out of their busy schedules to help me in understanding a concept or navigating the maze that is the Westminster office. I would like to especially thank my mentors, Austin Schaller and Nason Newberg. Thank you, Austin for sharing all your projects and inspiring an excitement in me for engineering that made me want to work on side projects before coming into the office and after getting home. Thank you, Nason for being able to swiftly and seamlessly switch between talking about high-complexity RF concepts to providing life advice. I learned so much in one summer and am excited to be back.

Thank you to Steve Narciso and Jon Kraft, two people who have provided invaluable resources and guidance. Thank you, Steve for your constant joyful attitude and introducing me to the analog world; I would not be in this career I love without your guidance. Thank you, Jon for inviting me on to this incredible project and for treating me as a real engineer which I hope to be one day.

Finally, thank you to all my family and friends for providing your time and support over the last few years and beyond. Thank you, Sabrina, my girlfriend and soon-to-be fiancé (assuming

she says yes, of course), for always listening to me talk about my "nerd stuff" and for providing so much encouragement when I need it. Thank you to my mom and dad for always being dedicated to being great parents, even when it is not easy. Thank you to my brother, Hayden, for always having an interesting perspective and asking the questions I would not think to ask.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

1.1 Motivation

Small and short to medium range radar systems are becoming increasingly common for a variety of applications including mobile systems such as cars, planes, and drones. Frequency-modulated continuous wave (FMCW) radar is at the forefront of this innovation due to its ability to operate in low-power environments. This has become possible due to the proliferation of the availability of components driven by the automobile industry. These low-power radars combined with phased array antennas enable the possibility of drone-mounted weather radar capable of navigating around storms and measuring at new points-of-view. In achieving this there are many areas which require careful attention including sizing, power, and polarization, all of which are large topics with significant research. This thesis explores some of these topics.

1.2 Objective

This thesis specifically focuses on low-cost and lightweight FMCW weather radar. In designing such a system there are numerous considerations in both the hardware and software design. The objective of this thesis is to provide a summary of some of the most important design considerations and key concepts for starting development. It is also to develop a generic design for a low-cost, drone-mountable radar system with modular subsystems for easily switching between applications.

1.3 Thesis Organization

This thesis presents the hardware and software considerations taken while designing the Dual-polarization phased Array Radar for Measurement of the Atmosphere (DARMA), a low-cost, portable FMCW weather radar, and is organized as follows:

Chapter 2 provides an introduction to FMCW radar concepts including the design parameters, generic block diagrams, and differences from traditional pulsed radar.

Chapter 3 discusses in-depth the design considerations taken when planning DARMA's structure and choosing components. It begins with a brief overview of the system's requirements and goals then moves on to cover each of the individual subsystems in the system diagram presented.

Chapter 4 presents signal processing methods and algorithms which are standard for many radar systems. They are presented along with considerations of when to use them and how they can be connected to one another in a full system.

Chapter 5 is a comparison between the radar designed in this thesis and the National Oceanic and Atmospheric Association's (NOAA) Snow-level FMCW radar. This radar was used as a reference point throughout the design because of its success and similarity in requirements.

Chapter 6 describes the difficulties presented by the pandemic on sourcing components and designing complex microwave systems and provides suggestions for mitigating these complications.

Chapter 7 provides a summary of the work presented in this thesis and suggestions for future research and new areas to explore.

Appendix A provides schematics of two different ground-based FMCW radars which employ separate beamformers and the fixed gain amplification breakout board described in Chapter 3.

Appendix B provides a detailed cost breakdown for individual components, fabrication, and assembly for each of the three boards described in Chapter 3.

Chapter 2

Principles of FMCW Radar

Frequency-modulated continuous wave (FMCW) radar is becoming increasingly popular for a variety of applications due to its versatility and mobility which is enabled by the the fact that it requires smaller peak transmission power than traditional pulsed radars. FMCW has two main attributes: frequency-modulated (FM) which is how range information is stored and continuous wave (CW) which is what allows for lower peak power transmission. Pulsed radars, for reference, transmit a chirp for a small fraction of the pulse repetition interval (PRI) which means that to transmit the same average power as a continuous signal, a significantly higher peak power must be generated during the chirp. FMCW radar has the advantage of constantly transmitting which means for the same average power, the peak power can be relatively small. For example, a common pulsed radar may transmit for a time, $\tau = 1\mu$s, and have a PRI of $T_{PRI} = 1$ms. Since the fraction of time spent transmitting is $t_{Tx} = \tau/T = 1/1000$, relative to a constant transmission, the output power must be larger by a factor of 1000. Figure 2.1 illustrates the difference between the transmitted waveforms of traditional pulsed radar and FMCW radar.

In contrast to pulsed radar's distance calculation method of transmitting a chirp and measuring the amount of time the signal takes to return, FMCW radars use a clever trick of storing timing information within the transmitted signal. The transmit signal is modulated with a linear waveform, causing the frequency to change with time, as can be seen in Figure 2.2[1]. Once transmitted, the signal travels for some non-zero amount of time and returns to the receive antenna after reflecting off a target. Upon reception of a signal reflected by a target, there are two signals available for use in determining the distance: the received signal and the signal currently being transmitted. Since the signal is constantly ramping in frequency, the received signal (generated earlier at time, $t - t_0$) will have a lower frequency than the currently transmitted signal (assuming the

[1]Figure 2.2 is merely for illustration of the frequency's relationship with time in FMCW radar, therefore the axes do not contain values associated with DARMA, the radar described in this thesis.

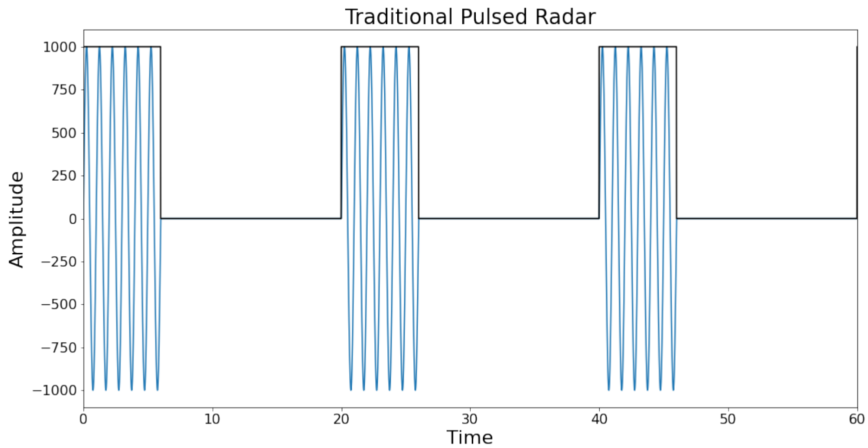

(a) Traditional pulsed radar waveform.

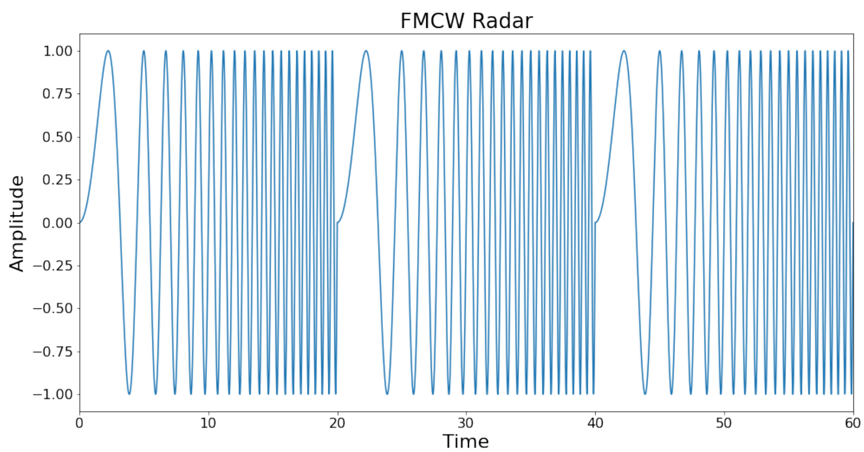

(b) FMCW radar waveform.

Figure 2.1: Transmitted waveforms of traditional pulsed radar and FMCW radar for three periods.

modulation waveform has positive slope as seen in Figure 2.2). Using an RF mixer, the received and transmitted frequencies can be subtracted from one another to obtain the difference as seen in Figure 2.3. This difference is known as the beat frequency, f_{beat}, and will be in the range of $0 \leq f_{beat} \leq B_{RF} - f_{ramp,min}$ where B_{RF} is the RF bandwidth and $f_{ramp,min}$ is the minimum ramping frequency; this is also known as the baseband frequency. Despite the beat frequency's range, in practice the baseband signal will be significantly smaller in frequency than the RF bandwidth

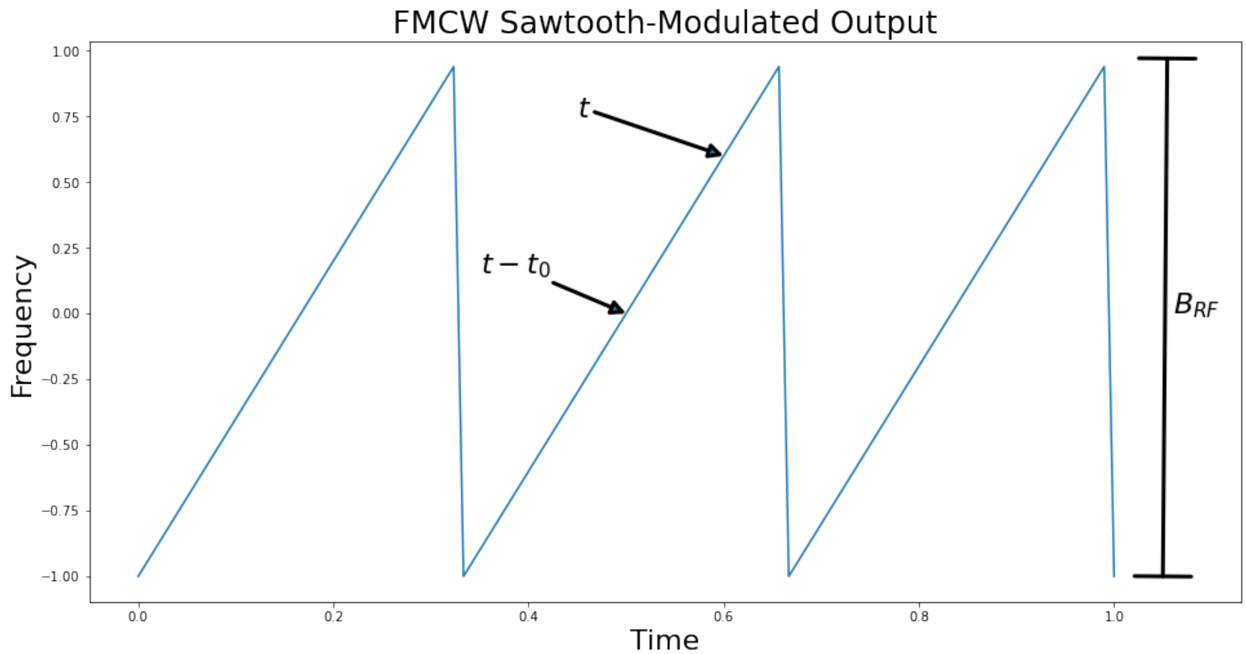

Figure 2.2: Frequency of FMCW sawtooth-modulated signal plotted as a function of time over three periods.

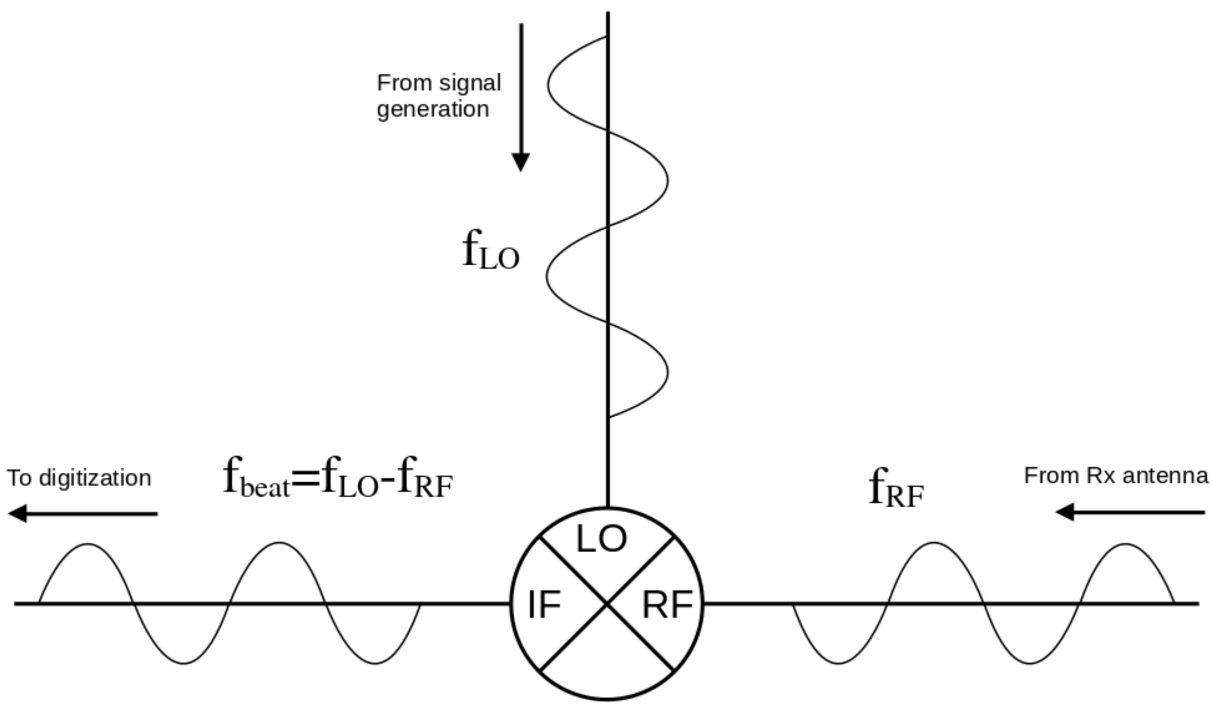

Figure 2.3: RF mixer used to subtract the frequencies from the received and generated signals, f_{RF} and f_{LO}, respectively, to obtain the beat frequency, f_{IF}.

because the ramp time, t_{ramp}, is generally much larger than the signal's round-trip time, or

$$t_{ramp} >> t_0. \tag{2.1}$$

The realization in block diagram form of the signal generation and frequency subtraction will be similar to that in Figure 2.4.

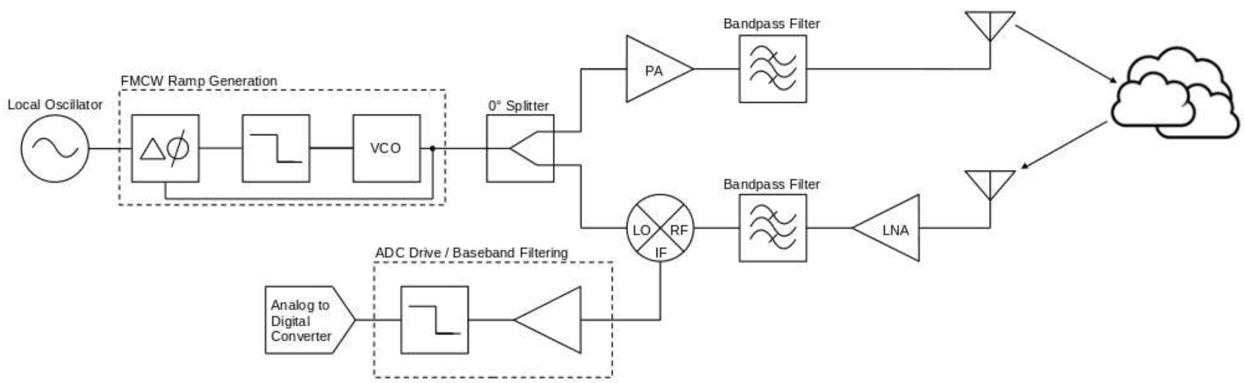

Figure 2.4: Generic block diagram of an FMCW radar.

Two bandwidths exist in FMCW radar: RF and baseband (the baseband frequency can instead be shifted to an IF to allow for filtering of the DC component which can be a large power contributor and can make a signal more difficult to locate). RF bandwidth affects the range resolution and ramp slope while baseband bandwidth is dependant on the maximum range and ramp slope. RF bandwidth is one of the design knobs that can be tuned to achieve finer range resolution, but resolution improvement comes with an increasing bandwidth that degrades noise performance and can be difficult to design an antenna for. This trade-off is discussed in further detail in Section 3.2. RF bandwidth, along with the ramp time, can be used to calculate the linear slope, $S = B_{RF}/t_{ramp}$, which enables the calculation of range from f_{beat},

$$R = \frac{f_{beat} c}{2S} \tag{2.2}$$

where c is the speed of light and the $1/2$ accounts for the signal traveling to and from the target. The calculated range must of course account for the resolution, so it becomes,

$$R \to R \pm \frac{R_{res}}{2}, \tag{2.3a}$$

$$\text{where } R_{res} = \frac{c}{2B_{RF}} \tag{2.3b}$$

is the range resolution.

The range to beat frequency relation (2.2) allows for plotting of received signal amplitude as a function of range rather than frequency. This is known as the range-FFT spectrum and is illustrated in Figure 2.5.

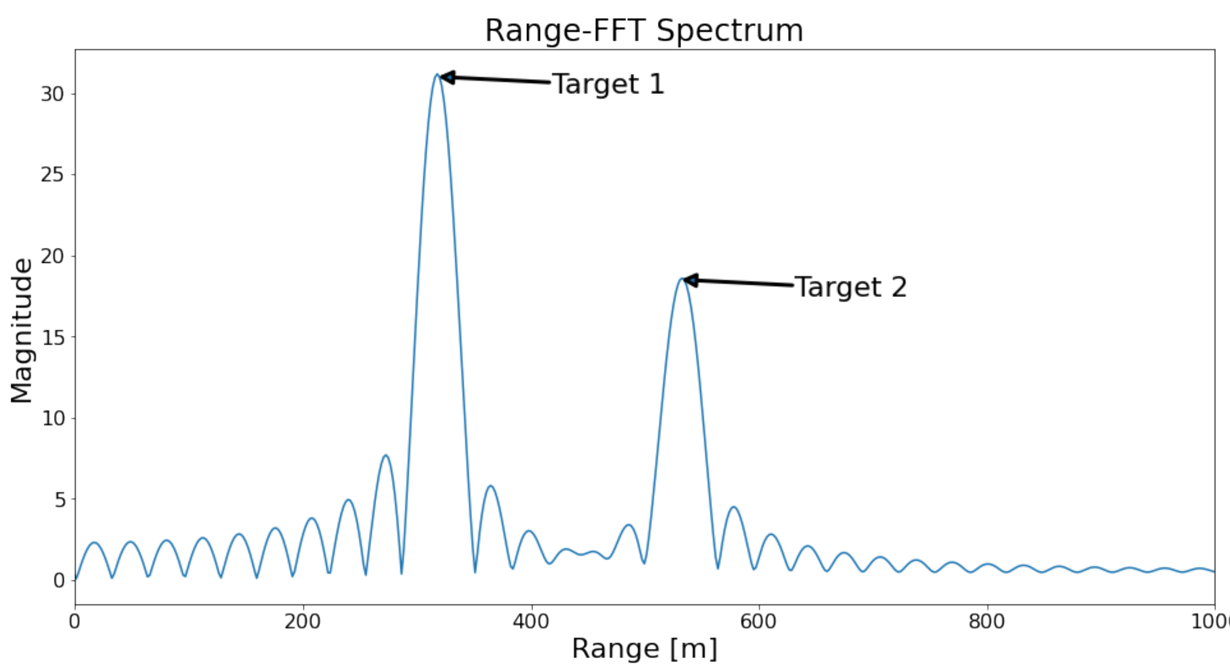

Figure 2.5: Example of a range-FFT spectrum including two targets and noise.

Calculation of distance using the frequency as described above, does not come without limitations, however. Traditional radar systems use the frequency of the received signal to determine the velocity of the target. This is done using the Doppler effect which states that when a wave collides with an object in relative motion, the wave's frequency will be shifted proportionally to the object's relative velocity. The shifted frequency is related to velocity by the equation,

$$v_r = \frac{f_d \lambda}{2} \tag{2.4}$$

where f_d is the Doppler (shifted) frequency, λ is the wavelength of the original signal, and v_r is the velocity relative to the transmitter. Unfortunately, this convenient property cannot be used directly in FMCW systems because it would be ambiguous whether the shifted frequency is due to a target's distance from the transmitter or its relative velocity. This necessitates the use of a different method for velocity calculation. Though the Doppler effect cannot be used when one buffer of samples is present, it is possible to obtain two successive buffers and compare their respective phases. The received signal will have not only a frequency component, but also a phase component and will be of the form, $A \sin(2\pi f t + \phi_0)$, where A is the amplitude and ϕ_0 is the phase. A single phase is not in itself useful but can be compared with another signal's phase to obtain the difference in range (range delta) by the equation,

$$\Delta\phi = \frac{4\pi \Delta R}{\lambda} \tag{2.5}$$

where ΔR is the range delta between the two samples. The range delta can be used along with the time between samples, $T_s = 1/f_s$, to obtain the velocity, $v_r = \Delta R / T_s$. Combining these equations gives the equation for velocity given the phase difference between two successive samples,

$$v_r = \frac{\lambda \Delta\phi}{4\pi T_s} \; . \tag{2.6}$$

This is calculated over all range bins to obtain the range-Doppler spectrum which is a process described in further detail in Section 4.8.3.

Despite the simplicity of the FMCW radar overview provided in this chapter, designing such a system is very involved and has many accompanying design considerations. These considerations are discussed in Chapter 3 in the context of a portable, low-power, X-band radar.

Chapter 3

Design Considerations

3.1 Introduction

FMCW radar, as discussed in Chapter 2, is very versatile. Because of this, it will be useful to test the limits of this versatility by designing a radar system capable of being re-purposed for a variety of use cases, of which the most interesting will be fully mobile systems. As part of this thesis, DARMA was designed with the intent of being used in both ground-based and unmanned aircraft system (UAS) or drone-based environments. The system will have a phased array architecture allowing for both beam-steering and amplitude beamforming when transmitting and receiving. In the interest of being modular, the antenna will not be permanently attached, and processing of the collected data will be offloaded to a computer. The board will have four receive taps for each polarization, horizontal and vertical, which will allow for the detachable antenna to be swapped for different applications, size requirements, and polarizations. The off-board processing is important because stationary applications can take advantage of low latency and high computing power while drone and other portable applications can transmit data over air to be received at a high computing power station but with higher latency. A block diagram of this modular radar system design is shown in Figure 3.1 and Table 3.1 summarizes the radar's design parameters. This section

Table 3.1: Design parameters for the ground- and UAS-based FMCW radar.

Parameter	*Description*	*Value*	*Unit*
R_{max}	Maximum Range	1,000	m
P_t	Transmit Power	30	dBm
f_c	Center Frequency	10,500	MHz
λ	Wavelength	28.57	mm
B_{RF}	RF Bandwidth	3	MHz
t_{ramp}	Ramp Time	10	ms
-	Polarization	Dual-Pol	-

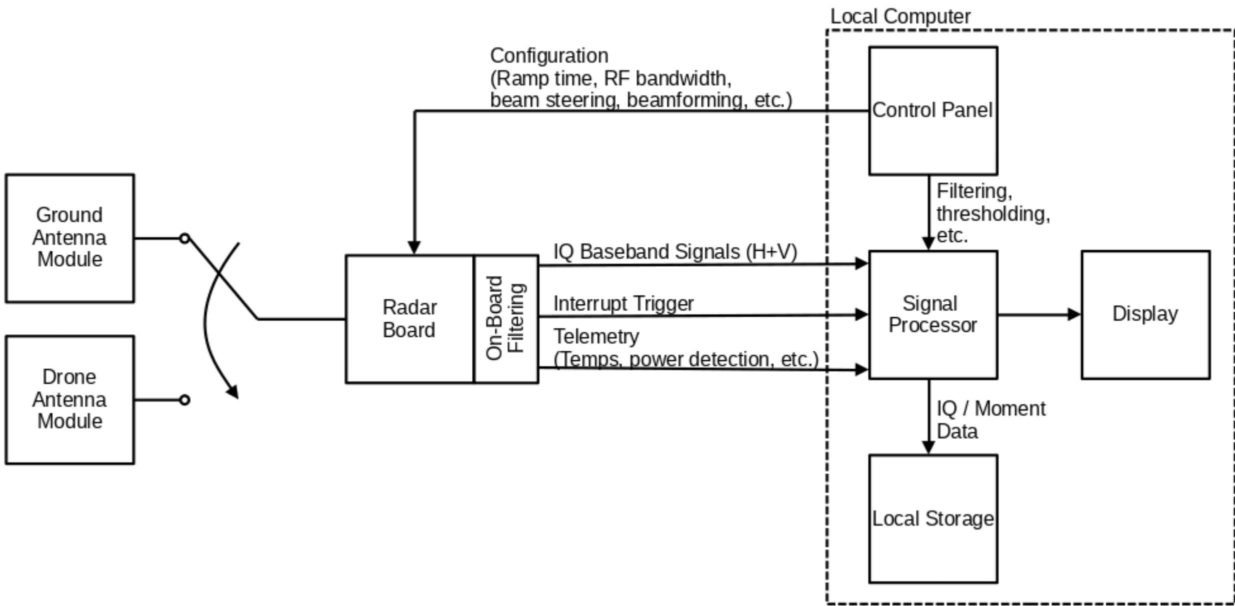

Figure 3.1: Modular FMCW radar system block diagram with support for swapping of the antenna and processing computer.

discusses the necessary trade-offs and available “design knobs” involved in the design of such a system.

3.2 Frequency Planning

3.2.1 Introduction

With the number of wireless devices increasing at such a rapid rate, stringent regulations have been put in place to ensure interference between devices is minimized. In addition to strict regulations it is critical to carefully plan the bands at which the radar system will be transmitting, receiving, and converting because it has a large impact on performance.

3.2.2 Transmission

Selecting a transmission frequency band can be difficult due to the very stringent frequency allocations designated by the International Telecommunication Union (ITU) and Federal Communications Commission (FCC). Their allocations not only outline which frequency bands are intended for public, government, meteorological, and other use but also permissible transmission power ranges. This makes the decision at the same time easier because it narrows the abundant options and more difficult because it decreases the selection of parts available.

Two main allocations were considered for the radar: Industrial, Scientific, and Medical (ISM) and meteorological X-band. The ISM band allows unlicensed use and the primary bands of interest it includes are centered around 2.4 and 5.8GHz. 2.4GHz was quickly removed from contention due to the requirement of the radar to be electrically steerable via phase-shifters and a phased-array antenna. Its wavelength, $\lambda = 12.5$cm, would require an antenna size on the high end of what would be optimal for a system with a small mounting area such as that proposed here. Options for beam-steering rapidly reduce in variety as frequency gets this low as well. The 5.6GHz ISM band is a strong contender because there are more beam-steering options available and the antenna size is manageable.

The meteorological X-band is an enticing option for a number of reasons, some of which are specific to this project and may not apply to others. X-band has a few very solid options for beam-steering including Analog Devices' ADAR1000 and AnokiWave's AWS-0101/0103, both of which include active beamforming and support for receiving in dual-polarization. These options alone were enough to heavily influence the frequency decision. The higher frequency would also allow for smaller antenna arrays.

It is important to note that signal attenuation in free space, or free space path loss, is proportional to the square of the signal's frequency and is determined by the equation,

$$FSPL = \frac{(4\pi R)^2}{G_t G_r \lambda^2} \tag{3.1}$$

where G_t, G_r are the transmitter and receiver gains, $\lambda = c/f$ is the wavelength, and R is the distance traveled. This property is an important consideration in frequency selection. For example, the choice of X-band (10GHz) versus C-band (5GHz) would result in an approximate extra path loss of 6dB which can be significant. There is, however, the upside of the wavelength being smaller in the X-band case allowing for smaller antennas which means more patch array elements can be used in the same surface area. In general, more elements indicates a higher gain which would reduce the negative effects of the higher path loss; fully compensating for the 6dB in loss would most likely require a larger increase in antenna elements than the decrease in wavelength would allow in the same surface area, though.

3.2.3 Reception

While in testing the range of targets may be known, in practice the measurement environment is often not, therefore making the target's corresponding beat frequency, f_{beat}, unknown. However, using the ramp slope, $S = B_{RF}/t_{ramp}$, the speed of light, c, and the maximum range, R_{max}, a range of values within which the beat frequency resides can be calculated, as shown in Chapter 2 as,

$$0 \leq f_{beat} \leq \frac{2R_{max}}{c}S. \tag{3.2}$$

Using the parameters in Table 3.1, the range is calculated as $0\text{Hz} \leq f_{beat} \leq 2\text{kHz}$. Again, this can be shifted to an IF to avoid DC noise, but this system simply uses the baseband signal and this will not mitigate antenna leakage leakage discussed in Section 4.5. With the f_{beat} range known, the receiver system can be designed.

Following the mixing stage will be some filtering, an analog-to-digital converter (ADC) driver, and an ADC, the selection of which depends on f_{beat}. The filter is a type of low-pass filter known as an anti-aliasing filter because of its use in rejecting signals above the Nyquist rate,

$$f_{\text{Nyquist}} = \frac{f_s}{2} \tag{3.3}$$

where f_s is the ADC sampling rate. In the case of this radar system, the beat frequency is far below the sampling frequencies of many ADCs on the market which is a property that can be used to improve noise performance and is discussed in more detail in Section 4.5.3. Before designing an ADC driver, the ADC it drives must be selected. There are numerous factors in the decision but, as far as the frequency planning is concerned, the most important is the sampling frequency must obey the Nyquist theorem,

$$f_s \geq 2 \times f_{beat,max}. \tag{3.4}$$

The driver is another important part of the receiver subsystem because, like high frequency RF lines and components, the ADC has a specific input impedance that, without proper matching, will deliver poor performance. How an ADC is driven greatly affects signal-to-noise ratio (SNR) and further affects detection of targets [1]. The driver specifications can often be calculated with help from manufacturer documentation and the driver itself is often combined with the anti-aliasing filter.

All the above considerations are directly tied to the decisions made for the transmit side which makes the "frequency planning" on the receive side more comprised of filter synthesis and component selection than anything else.

3.2.4 Bandwidth

The two bandwidths mentioned in Chapter 2, RF and baseband, are two very critical design decisions with large effects on the performance of the radar.

RF bandwidth controls one of the most important aspects of a radar system, its range resolution, R_{res}, which in turn directly affects the volumetric resolution, V_{res}. The resolution impacts the ability to differentiate between targets (Section 3.4.1) and is related to RF bandwidth by (2.3b). With only this relation, the obvious choice would be to indefinitely increase the RF bandwidth to achieve a near-perfect range resolution, however in real systems there is the system noise power to account for which is directly related to the RF bandwidth. At every frequency the radar receives there exists some noise, therefore the more frequencies it receives the greater the noise power will be. The noise power is defined by,

$$P_n = kTB_n \tag{3.5}$$

where $k = 1.38 \times 10^{-23}$ is the Boltzmann's constant, T is the ambient temperature in Kelvin, and B_n is the noise bandwidth which can be substituted with B_{RF}. This, along with a desired signal-to-noise ratio, sets an upper limit on the RF bandwidth. The ramp slope, previously defined as $S = B_{RF}/t_{ramp}$, also contributes to the design problem because it determines the beat frequency, which is defined in (2.2). f_{beat} is directly proportional to S which is important because as S decreases, f_{beat} approaches DC which can be highly contaminated with noise, further reducing the noise figure. Some of the noise at DC is avoidable by converting f_{beat} to an intermediate frequency (IF), but others such as transmit to receive antenna leakage will follow the shift to the IF, which is discussed further in Section 4.5.2. The ramp time is one of the "design knobs" with fewer-than-average side-effects which makes it preferable to the RF bandwidth for adjusting the baseband bandwidth.

3.3 Signal Generation

3.3.1 Introduction

Signal generation is an integral component in FMCW radar because it is what allows the extraction of range information from the reflected signal. As discussed in Chapter 2, a linear frequency-modulated signal must be generated. The frequency modulation can take multiple forms, but sawtooth and triangular waves are the most common. Two main methods exist to accomplish this which include using a frequency synthesizer fed into a voltage-controlled oscillator (VCO) and the use of a direct digital synthesizer (DDS) upconverted to the desired RF frequency. Their trade-offs are discussed in this section.

3.3.2 Phase-Locked Loop

The first method, otherwise known as a phase-locked loop (PLL), is very common in a variety of applications including many receiver systems which makes the options abundant (this is less true in the current climate as will be discussed in Chapter 6). PLLs are used to generate a desired frequency from a reference oscillator and lock it in place by taking the output as feedback and applying some correction. Frequency synthesizers are the first component in the PLL. They take in a reference oscillator at a known frequency and feedback from the VCO's output. The reference frequency is multiplied by a value (can be fractional or an integer depending on the synthesizer) to obtain the desired output frequency and compares it with the current frequency which results in some phase difference. The phase difference indicates how much correction is needed and is translated to a voltage to be input to the VCO. VCOs take this voltage and generate a corresponding frequency. The relation between input voltage and output frequency is ideally linear but is not perfect in practice which introduces the need for coherency between the transmit and receive sections (discussed in Section 4.11). The output of the VCO is then fed back to the synthesizer. There is another important component between the PLL and VCO known as the loop filter. The loop filter, which is effectively a low-pass filter, is used to block high frequencies from entering the VCO and interfering with the output. The loop filter bandwidth, or cutoff frequency, can be a difficult design decision because lower bandwidths correspond with lower noise but increase the frequency lock time. In the case of FMCW radar, low lock times are very favorable because for a set ramp time, they allow for finer resolution in frequency.

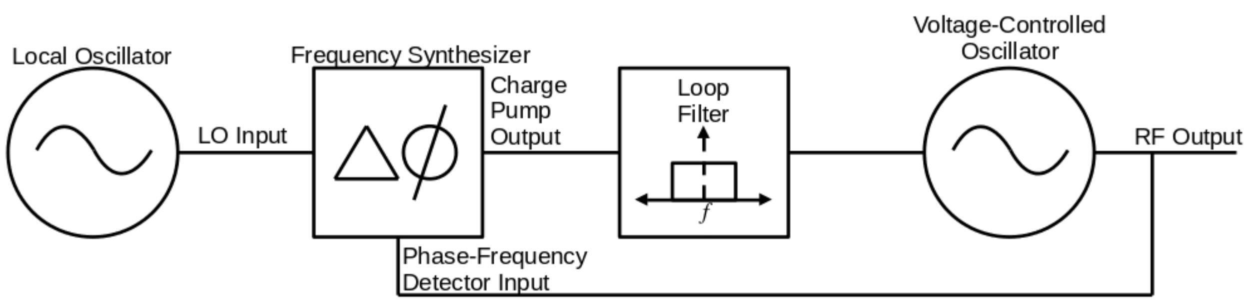

Figure 3.2: Generic phase-locked loop block diagram.

One of the major downfalls of PLLs is their generation of phase noise, or the "cleanliness" of a signal's frequency which is caused by the feedback system. When the first output is created by the VCO, it will be slightly off from the desired frequency so it is sent back to the phase-frequency detector (PFD) to be corrected for. The correction voltage is calculated and sent through the loop filter to the VCO, where it will again be fed back to the PFD. This process is repeated many times within a short period until the error is sufficiently small, at which point the signal is *locked*.

These micro-adjustments cause slight shifts in the frequency surrounding the desired frequency and shows up as a "skirt" around it in the frequency spectrum [2]. Figure 3.3a shows what an ideal signal would resemble and Figure 3.3b shows the same signal with added phase noise.

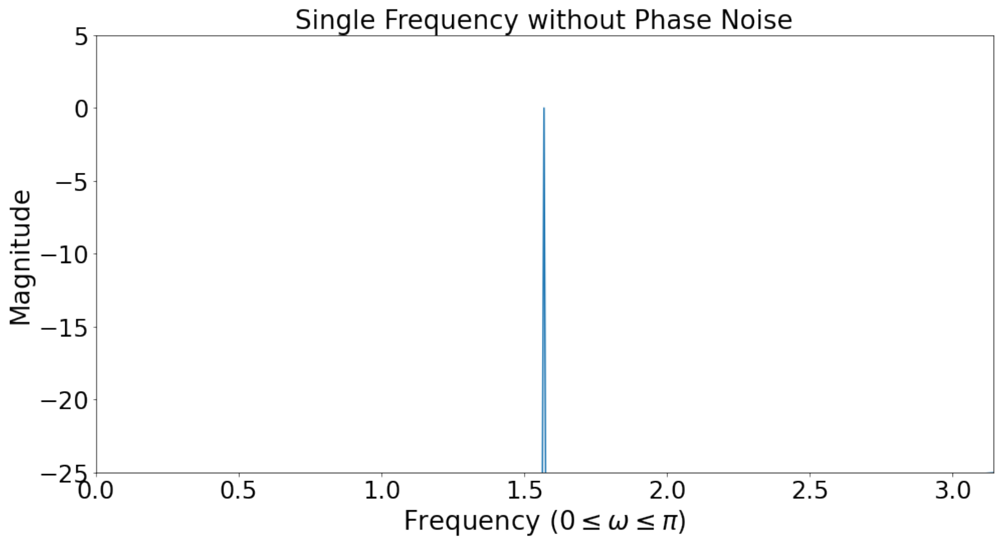

(a) Frequency spectrum of a single tone signal without phase noise.

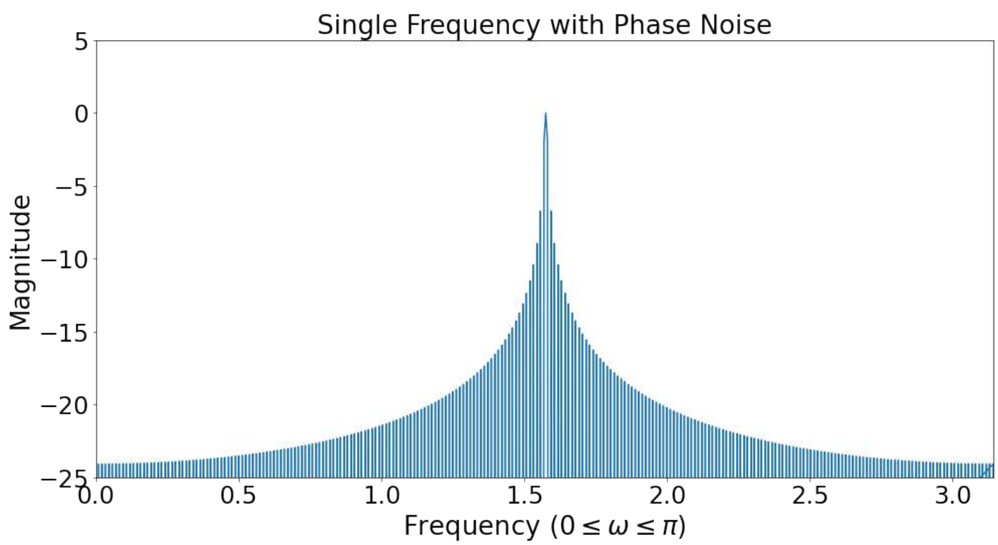

(b) Frequency spectrum of a single tone signal with phase noise.

Figure 3.3: Simulation of two single tone signals, one with and one without phase noise.

VCOs, as briefly mentioned above, also pose some issues in the form on non-linearities and larger-than-ideal frequency resolution. Ideally a VCO's output would be linearly correlated with the input (tuning) voltage, but this is never fully the case in real systems. There is often a section of the bandwidth that can be treated as linear without any real issues, however. It is therefore important to choose a VCO where this ideal bandwidth section encompasses the radar's RF bandwidth. Figure 3.4 gives an example of a common VCO's output frequency and tuning voltage relationship.

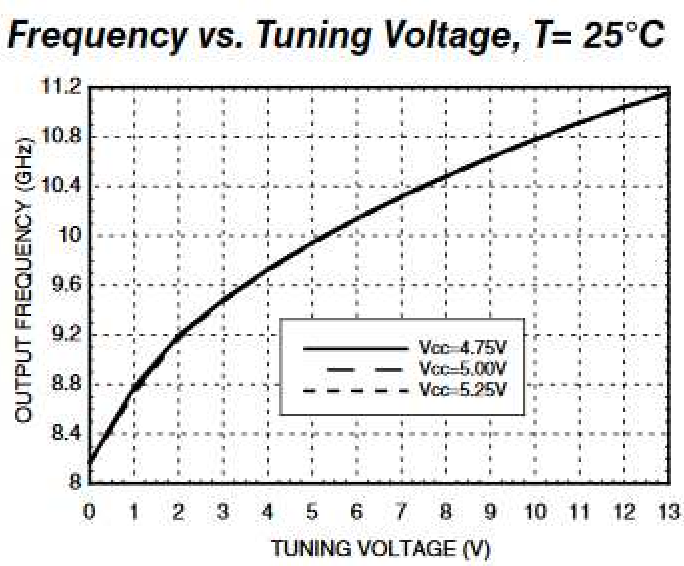

Figure 3.4: Frequency vs. tuning voltage of Analog Device's HMC530 VCO.

The other issue to account for is the frequency resolution which is dependent on the synthesizer's resolution and the VCO's slope. Synthesizers output a current proportional to the phase difference and the smallest detectable difference is the phase resolution. The output current is translated to a voltage via the charge pump which therefore also has a resolution, V_{res}. The frequency resolution is then calculated by $f_{res} = V_{res} \times S$, where S is the VCO's frequency vs. tuning voltage slope in Hz V^{-1}.

3.3.3 Direct Digital Synthesis

The second method uses a DDS to generate the frequency-modulated signal at baseband and is then upconverted to RF using a mixer and RF oscillator. DDS', in contrast to PLLs, are purely digital devices that take in a tuning word to specify the output frequency. To know which frequency to output, the DDS contains a programmable read-only memory (PROM) module with phase values accessed by the tuning word which functions as an address. The tuning word's corresponding phase value is used in conjunction with a local oscillator to generate discrete steps along the sine wave [3]. The length of the tuning word, n, and LO frequency, f_{LO}, determine the frequency resolution allowed by the DDS,

$$f_{res} = \frac{f_{LO}}{2^n}. \tag{3.6}$$

The steps in frequency of size f_{res} can then be accessed using the tuning word, M, to get an output frequency,

$$f_o = \frac{M \times f_{LO}}{2^n} = M \times f_{res} . \tag{3.7}$$

Common tuning word lengths are in the range of 24 to 32 bits or approximately 16 million to 4 billion possible frequencies. For applications like FMCW radar where frequencies must be rapidly modified, the lookup times that accompany this number of possibilities would be entirely impractical. Because of this, the tuning words are often truncated to more manageable lengths which leaves only the first 15 bits or only about 32 thousand possibilities while keeping the same resolution as before. While this makes DDS' more practical for these applications, the noise performance is affected by adding spurs to the output. The magnitude of these spurs is directly correlated to the amount of non-zero bits in the truncated segment of the tuning word [4]. In addition to spurious signals, a DDS will have some phase noise but is generally dominated by the LO input and in most cases is less than its PLL counterpart.

3.3.4 Comparison of Signal Generation Techniques

The DDS and frequency synthesizer considered for this system were the AD5932 and ADF4159, respectively. The AD5932 is a solid contender because of its low power consumption of 6.7mA, frequency resolution of 3Hz, and phase noise at 100kHz of less than -140dBc. In all these important specifications it outperformed the ADF4159 but, in this case, those were not the only deciding factors. The ADF4159 has the advantage of being specifically designed for use in FMCW systems and offers much simpler programming support for generating frequency-modulated signals. In addition to making the programming simpler, the ramping functionality needs very little or no configuration after setup which allows the processor/controller to use its clock cycles for other, more critical applications such as filtering the incoming signal. It is for the reasons listed above that this radar generates the FMCW waveform using a PLL.

As previously discussed, the lock time of a PLL is one of, if not the most, important factors for an FMCW system and is lowered by increasing the loop filter's bandwidth. The designer's job is to find a bandwidth such that the lock time and phase noise are sufficiently fast and low, respectively. Using Analog Devices' PLL design tool, ADISimPLL, the bandwidth was chosen to be 500kHz which resulted in the lock time and noise response shown in Figure 3.5. This configuration results in an approximate lock time of 2μs and phase noise at 100kHz of -95dBc; a worse performance than the DDS configuration but an acceptable trade-off for the ease-of-use. The full signal generation PLL is shown in Figure 3.6.

3.4 Range Selection

3.4.1 Resolution

The range, along with a few other parameters from the antenna and signal generation, dictates the volumetric resolution. This can greatly affect the radar's performance in almost all applications such as in meteorology where it could affect its ability to distinguish between different storms or in aircraft detection where it could affect its ability to differentiate between multiple aircrafts. The volumetric resolution is defined as

$$V_{res} = R_{res} \cdot (R \cdot \theta_1) \cdot (R \cdot \phi_1) \tag{3.8}$$

where R_{res} is the range resolution, R is the distance to the target, and $\theta_1,\ \phi_1$ are the beamwidths in the horizontal and vertical directions, respectively. A cross-section of the volumetric resolution is

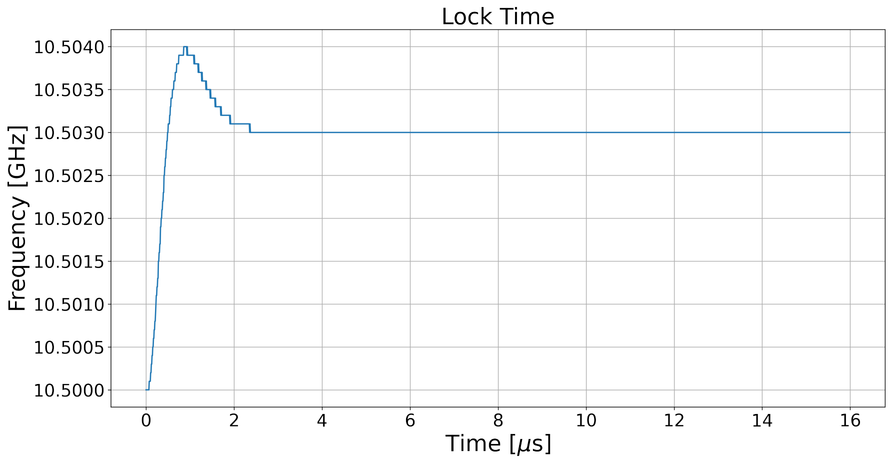

(a) Lock Time.

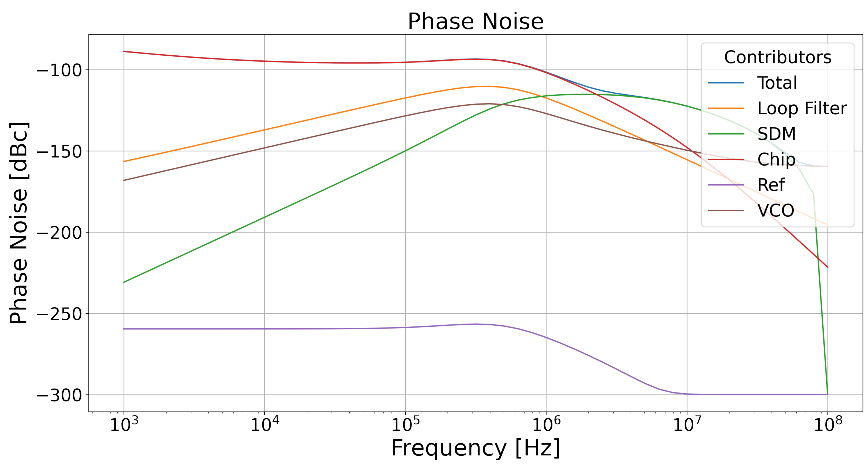

(b) Phase Noise.

Figure 3.5: PLL lock time and phase noise performance with a loop filter bandwidth of 500kHz. The phase noise is shown for the entire system as well as for each of the individual components in the PLL including the loop filter, sigma-delta modulator (SDM), synthesizer IC (Chip), reference clock (Ref), and VCO.

visualized below in Figure 3.7. As can be seen in (3.8) and Figure 3.7, the volumetric resolution increases with distance to the target which means that the further a group of targets is from the radar, the less distinguishable the individual targets will be from one another. The only design knobs available to combat this are the RF bandwidth and antenna beamwidths. Since the goal of this system is to be versatile in different environments (i.e., ground and drone deployment), the choice was made to have the antenna connected via SMA cables rather than be permanently mounted. This allows for the use of small antennas when deployed on a drone and larger when on ground.

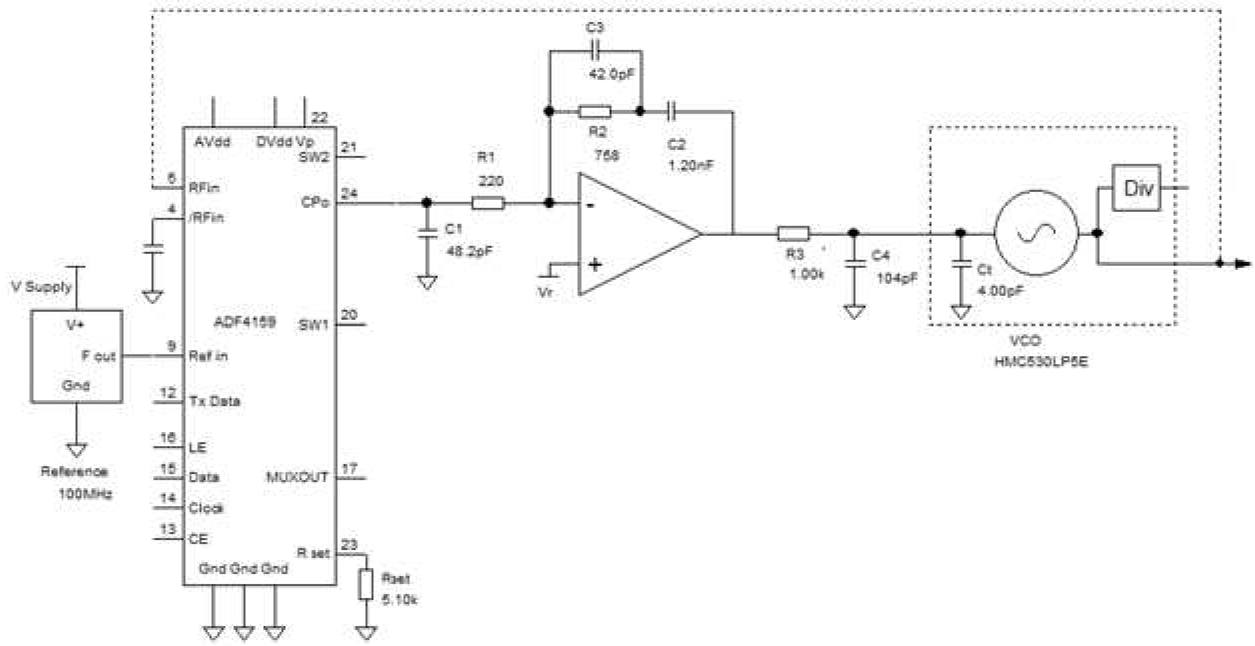

Figure 3.6: PLL schematic including passive component values to achieve the desired loop filter bandwidth of 500kHz, charge pump current of 4.8mA, and phase margin of $45°$.

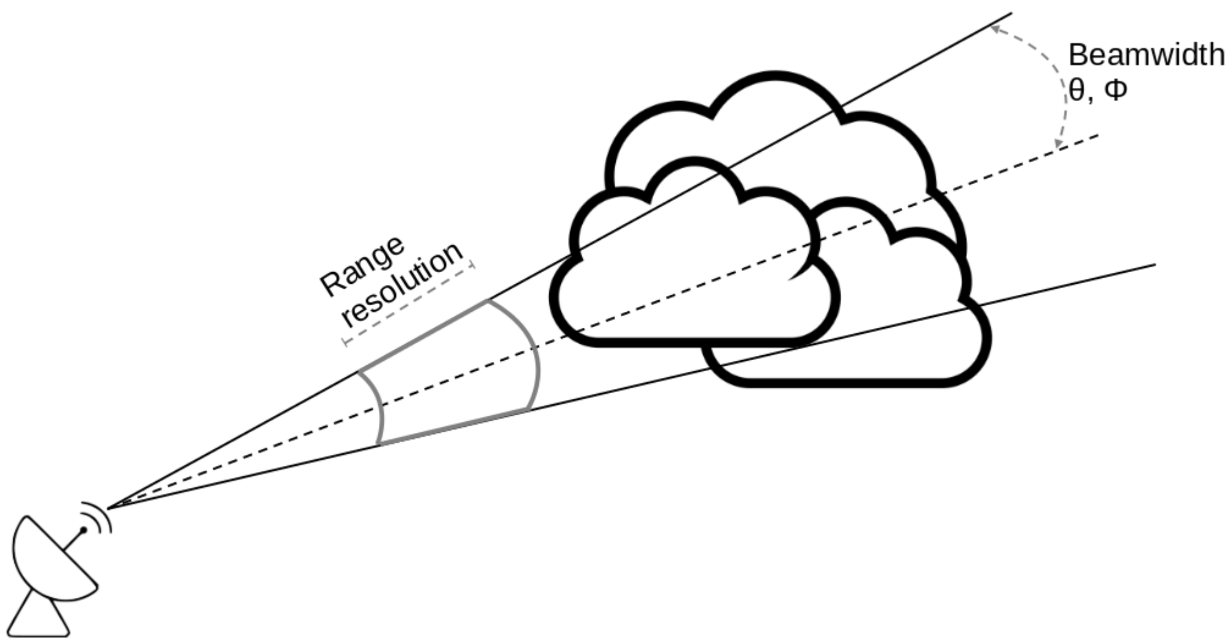

Figure 3.7: Visualization of volumetric resolution, beamwidth, and range resolution.

In addition, antennas with differing beamwidths can be tested for more precise measurement of far away targets. These beamwidths are directly proportional to the number of elements which in this case is four for both transmit and receive. This is limiting but, if needed, the taps could be connected to multiple antennas via a power splitter, allowing for smaller beamwidths but the same beam-steering capability.

3.4.2 Near and Far Field

It is important to understand what the general behavior of the received signal should be, otherwise it can be easy to misinterpret the data. The distance at which targets are being measured can have a large impact on this understanding because propagating waves will have different behaviors at different distances. These behaviors are determined by whether the target resides in the *near field* or *far field*.

RF signals originating from a source in space radiate radially as shown in Figure 3.8a. If such a source were to be placed directly facing a phased array, the signal would not reach each of the elements at the same time; the middle elements would receive the signal first, followed by the surrounding elements. This would not only introduce a phase shift increasing outward from the center element but a difference in received power, effectively applying a gain taper which can produce misleading results. This is the near field effect. Fortunately the near field does not extend indefinitely and eventually starts behaving as a far field shown in Figure 3.8b. The far field is when the radius of the propagating waves to the source become large enough to resemble a flat propagating surface, or plane wave. Receiving a plane wave is ideal because, unlike in a near field environment, all receiving elements will receive at approximately the same time, or at least with a negligible time difference, in the case where the source is located at boresight.

The range boundary between these two regions is not absolute but can be approximated to be

$$R = \frac{2D^2}{\lambda} \tag{3.9}$$

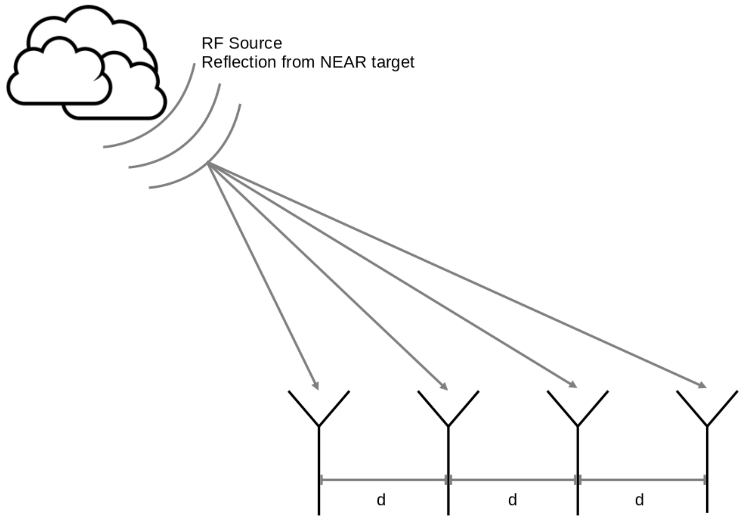

(a) Near field example showing that at short ranges from a target, the radial propagation has a non-negligible effect on the distance to each antenna array element.

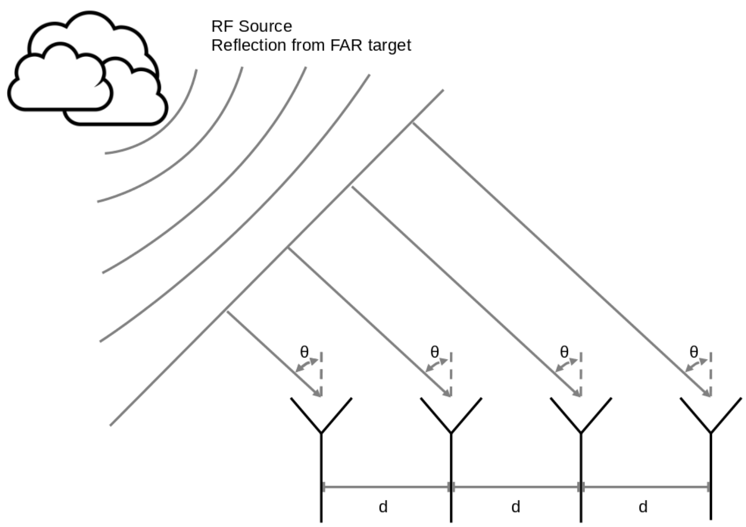

(b) Far field example showing that as range from a target increases, the radius between the wavefront and source becomes large enough to resemble a plane wave.

Figure 3.8: Near field and far field example showing the effect the radius from the wavefront to the source has on the distance to individual antenna array elements. The radiating source is backscatter from a cloud.

where D is the antenna diameter [5]. Phased arrays do not have a diameter, but an *effective* diameter can be approximated for a uniform linear array by

$$D = d(N - 1) \tag{3.10}$$

where d is the element spacing and N is the number of elements. As an example, the Analog Devices' *FMCW Phaser Board* (used in Chapter 4) is taken with the parameters described in Table 3.2. Using (3.10) and (3.9), the effective antenna diameter would be $D = 105$ mm and the field boundary would be $R = 889$ mm. Fortunately this is short relative to the maximum range but, for some applications such as indoor sensing, it could still present issues.

Table 3.2: Parameters for Analog Device *FMCW Phaser Board* relevant to the calculation of the near and far field range boundary.

Parameter	*Description*	*Value*	*Unit*
f_c	Center frequency	10.5	GHz
λ	Wavelength	24.8	mm
N	Number of elements	8	-
d	Element separation	15	mm

The near field effect does not make accurate measurement of targets impossible, but it does require some extra calibration. Depending on the range of a target, the delay at which each element will receive the signal can be calculated and further corrected by applying phase weights to each antenna element [6]. The phase weights effectively add delay to the paths from the antenna element to the analog-to-digital converter. If a delay of equal and opposite magnitude is applied to each respective element, the near field effect can be accounted for. Near field correction is not implemented in this thesis because it was possible to obtain measurements entirely in the far field, but it is an available approach worth noting.

3.4.3 Drone Line-of-Sight

This radar system uses a phased-array antenna that allows for beam-steering. Similar to normal ground-based radar with a mechanical positioner which allows for sweeping in azimuth and elevation, the antenna will be able to scan in one dimension if using a linear array or two if using a two-dimensional array. Because of the ability to scan, the radar will not always be pointing towards nadir and will therefore not always have the specified 1000m of vertical range. The vertical

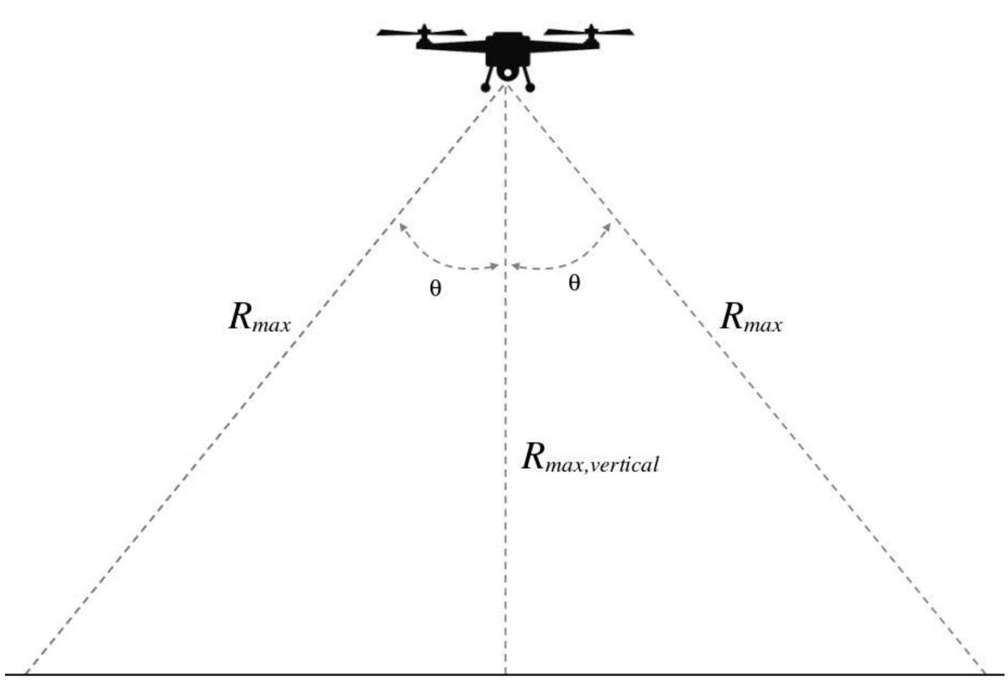

Figure 3.9: Downward-facing radar operation when mounted on a UAS.

range can be found from Figure 3.9 as

$$R_{max,vertical} = R_{max} \cos(\theta) \tag{3.11}$$

where θ is the angle from nadir. This maximum vertical range is important to consider when determining the altitude at which to operate the drone. For example, if the drone is flown to an altitude of its max range of 1000m, it will only be capable of measuring the ground when pointing directly downward which is of course in many cases not useful. Instead, the cross-sectional area of interest should be determined (i.e., an area encapsulating a storm or cloud) and used to calculate the optimal altitude for measurement. With the width and altitude of interest and (3.11), one can determine whether it is possible to measure the entire area by the Pythagorean theorem,

$$R_{max} \geq \sqrt{(\text{Altitude})^2 + (\text{Width})^2}. \tag{3.12}$$

3.5 Beamforming and Beam-steering

Though a phased-array antenna can be used without any beam-steering or beamforming and still achieve better performance than a single patch antenna in terms of directivity and gain, the big advantage of phased-arrays is their ability to be steered electrically and change their antenna pattern. Electronic beam-steering is used as an alternative to or in addition to a mechanical positioner [7], in the case of this radar, an alternative. It works by introducing phase shifts to all or a subset of the elements of the array; this effectively gives the elements' transmitted waveforms a time delay with respect to one another. If the elements are linearly shifted by some angle, $n\psi$, where $n = 0, 1, 2, \ldots, N$ is the element number, the beam will be directed away from boresight at a *scan angle* proportional to the phase difference, ψ. A desired scan angle, shown in Figure 3.10 as θ_{scan} can be used, along with the spacing between array elements, d, to determine the necessary element phase shift by the following equation,

$$\psi = \frac{2\pi d \sin\theta_{\text{scan}}}{\lambda} \tag{3.13}$$

where λ is the wavelength. A common element spacing used in real systems is $d = \lambda/2$ which can be substituted into (3.13) to form the equation,

$$\psi = \pi \sin\theta_{\text{scan}}. \tag{3.14}$$

Using these equations and the beamwidths of the antenna, θ_1, ϕ_1, the entire volume the radar can "see" is calculated. This behavior, of course, becomes more difficult to describe using simple equations when another dimension is added, but the equations above are a good starting point for understanding beam-steering.

It should also be noted that when calculating the phase difference, ψ, between the antenna elements, the elements will most likely have some offset, ψ_0. This offset determines where the antenna is pointing when no beam-steering is applied to the array elements and results in some offset scan angle off of boresight meaning the "boresight" will actually be at some θ_0. Through

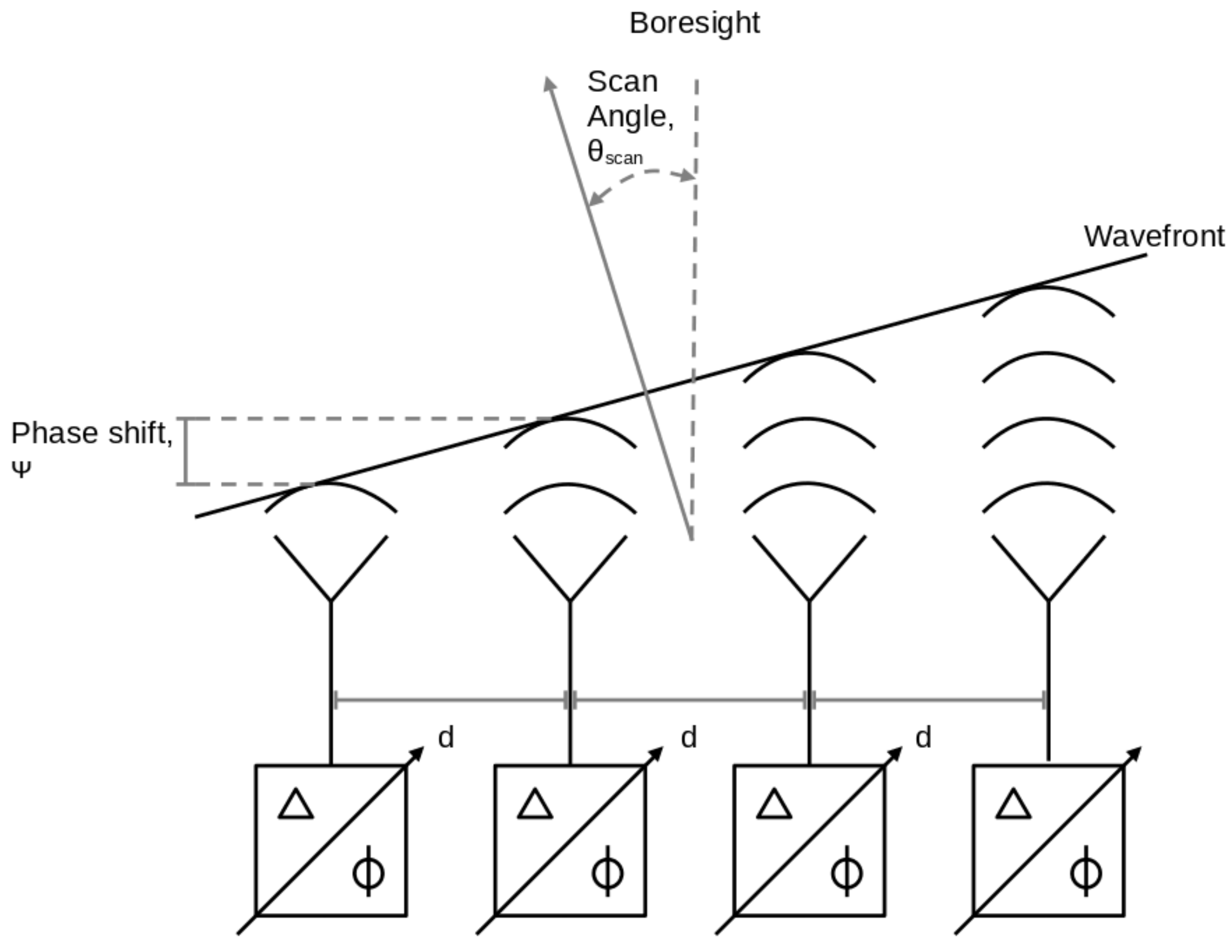

Figure 3.10: Beam-steering of a four-element uniform linear array (ULA).

antenna calibration, ψ_0 can be found and subtracted from the elements for more accurate scanning. When steering an antenna, collected data may seem to be picking up signals at angles away from the scan angle which is undesirable when attempting to detect targets at a specific location. This can be caused by side-lobes of the antenna pattern which can vary in width and magnitude depending on the antenna; an example is shown in Figure 3.11. If the side-lobes are large enough, targets at an angle included in the side-lobe can have very similar gain to those within boresight which can lead to mistaking what is to the side of the radar for what is directly in front of it. To account for this, beamforming techniques can be employed.

Beamforming is another method of beam manipulation that, instead of using phase weights, uses element gains to modify the shape of the beam. If smaller side-lobes are required in the

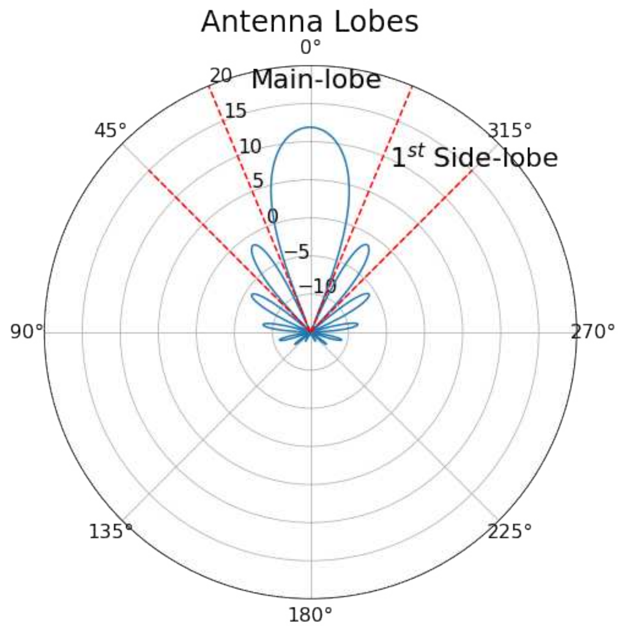

Figure 3.11: Antenna pattern of an ideal 32-element ULA with a Bartlett taper applied to accentuate the main- and first side-lobes which are labeled.

antenna pattern, for example, the engineer can employ gain tapering from the center to the edges. This gives the center element(s) the highest gain and tapers off as the elements get further from the center. This effect is similar to that of windows used in digital signal processing and therefore the general behavior can be visualized using the same techniques. As is a common theme in all engineering disciplines, the *no free lunch* principle applies here in the form of widening the main lobe. The trade-off between no tapering (rectangular window) and two common tapering algorithms, a Taylor taper with $nbar = 4$ and a Kaiser taper with $\beta = 5.65$, can be seen in Figure 3.12 (the visualization uses 64 elements for the purpose of clearly showing the effect). Figure 3.12 clearly shows that with side-lobe reduction comes increasing main lobe width. As a side note, the main lobe width is inversely proportional to the number of array elements meaning, at the cost of space and complexity, it can be reduced.

The beam-steering and beamforming discussed above are realized using two separate components that are many times packaged together, a phase shifter and variable gain amplifier (VGA). The purpose of this radar system is to take advantage of commonly available parts and design a general system for use in a variety of scenarios. Because of this, it would be out of the scope of

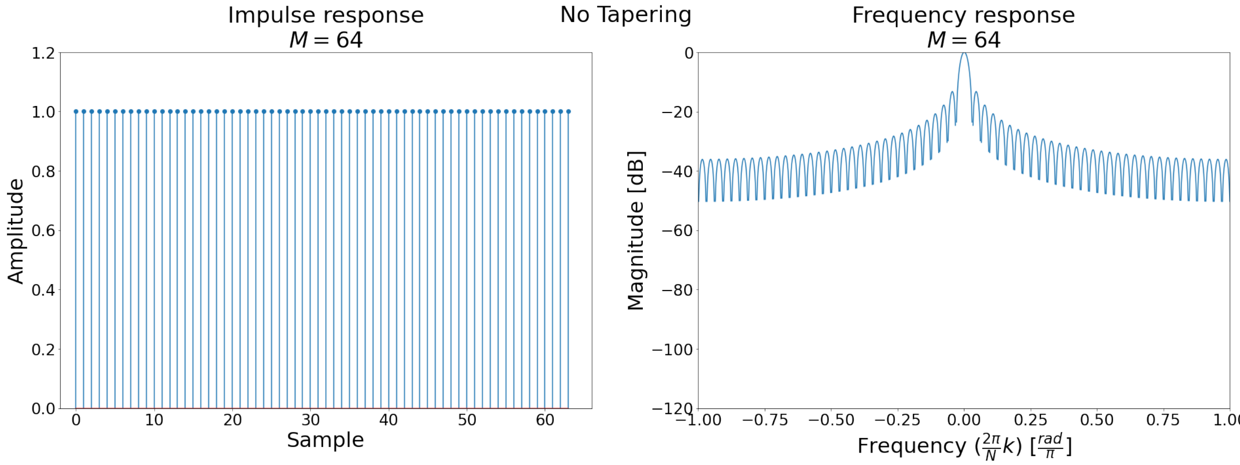

(a) No tapering.

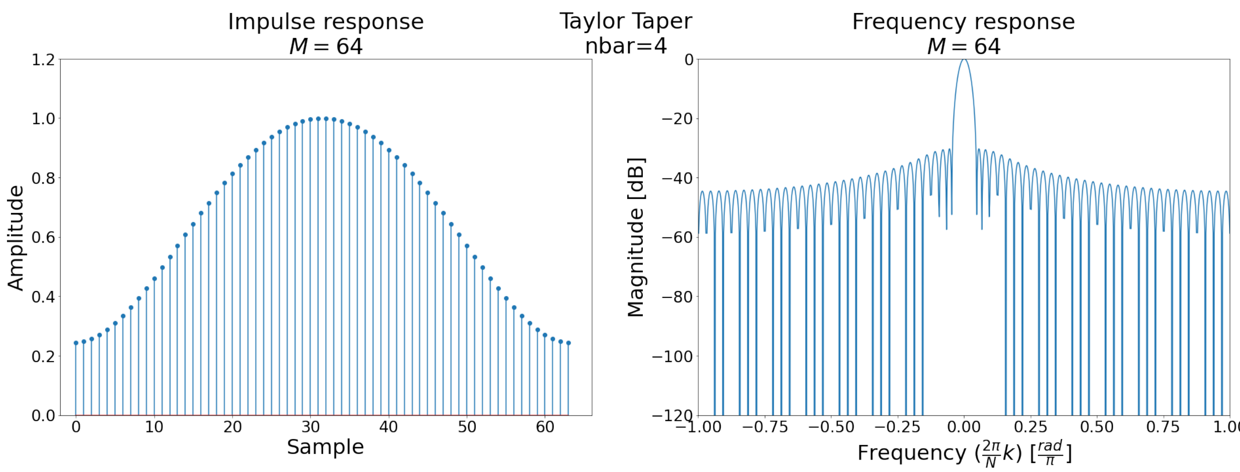

(b) Taylor taper (nbar $= 4$).

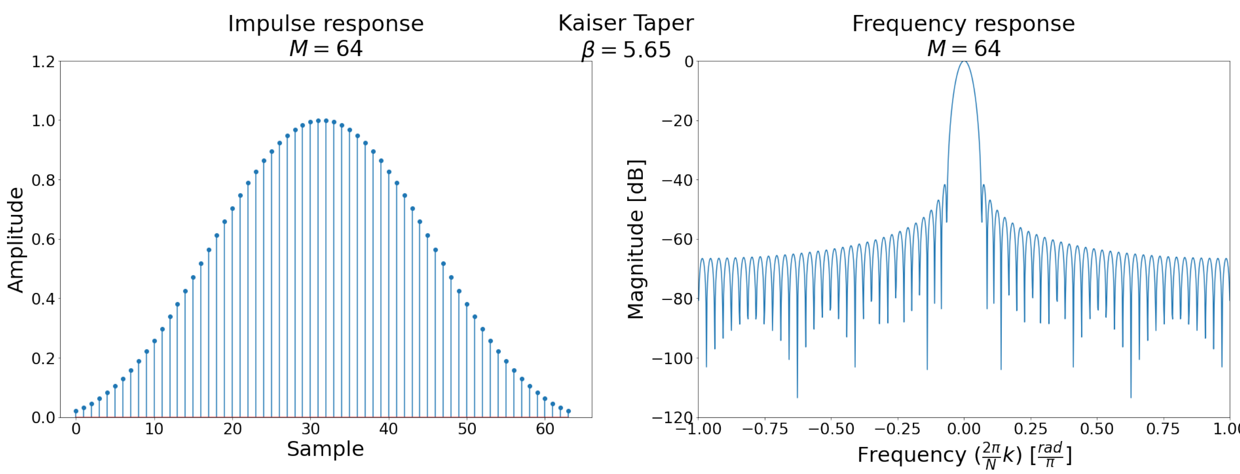

(c) Kaiser taper ($\beta = 5.65$).

Figure 3.12: Effect of three gain tapering functions on a 64-element ULA modeled with temporal sampling.

the project to custom design either of these two components. Unfortunately, however, there are very few options available, and the problem is worsened by the current electronics market climate induced by the worldwide pandemic (Chapter 6). Though the specifications of this radar only require beam-steering, since this is a general concept intended to be used and modified by others, beamforming was also desired. The restrictions outlined above left very few options: AnokiWave's AWS-0101/0103 and Analog Devices' ADAR1000.

Of the two AnokiWave devices, the AWS-0101 was chosen for a few reasons. In a portable system such as this where limited antenna gain is available, it can be difficult to achieve a sufficient SNR which makes receiver noise performance of the utmost importance. The difference in this respect between the two beamformers is massive with the AWS-0101 having a receive noise figure of 4dB while the AWS-0103 has a noise figure of 15dB. The noise figure of the two devices can be misleading, however, because the AWS-0103 is intended for use with an RF front end IC that has significantly better noise performance ($NF = 3.3$dB) and a high coherent gain of 24dB that would mostly negate the second stage noise figure from the beamformer. Though the AWS-0103 has a higher coherent gain than the AWS-0101's 20dB, it does come with two drawbacks in the form of board real estate and an inability to simultaneously transmit and receive. A front-end IC is needed for each of the four taps which, when it is in a 7x7mm, 56-lead package, takes up a significant amount of space. The other issue is that it does not support simultaneous transmit and receive which is a critical part of this FMCW radar. If the AWS-0103 were to be used in this application, it would require at least eight of the RF front-ends for single-polarization or twelve for dual-polarization (assuming a 4-element ULA for both transmit and receive), which is not only a waste of space but also an inefficient use of the component.

Due to the space issues and complexity the AWS-0103 presents, the AWS-0101 was chosen. Though it better suits this application relative to its counterpart, the beamformer introduced a few of its own design challenges. One of these challenges is discussed briefly in the Section 3.7 and is its output power compression point. To achieve a sufficient signal-to-noise ratio at the maximum range of 1000m, it was calculated that the radar would need to be capable of outputting a relatively large output power of 30dBm. Using the beamformer as the final stage in the signal transmission is impossible due to its output power compression point of 13dBm. Accounting for this requires the use of a power amplifier on each element, not just one for the transmit side. Finding a PA to be placed after the beamformer was difficult for a few reasons. First, it needed to be capable of transmitting 30dBm of power, which meant its compression and saturation points needed to be above that which narrowed the selection significantly. In addition, the VCO used prior to the beamformer outputs a relatively high power of 11dBm which becomes 6dBm after going through the power splitter with 2dB of insertion loss. This leaves only a small margin of 7dB below the compression point of the beamformer, so optimally all of it would be available for use in beamforming. At the output of the beamforming stage and input to the PA stage will be $\leq$13dBm meaning that to adhere to transmission standards discussed in Section 3.2, the PA needed to have a gain of near exactly 17dB. All these requirements left a very small subset of options, from which the MAAP-008294 was chosen. The MAAP-008924 is not a perfect option and the design changes that accompanied it are discussed in Section 3.7.

The other promising option was the ADAR1000 by Analog Devices. It boasts a transmit gain of approximately 20dB, output 1dB compression of 10dBm, receive gain of 9dB, input 1dB compression point of -16dBm, and receive NF of 8dB. As is seen in Table 3.3, it is very similar to

Table 3.3: Key specifications for the AWS-0101 and ADAR1000 beamforming ICs.

Parameter	*Description*	AWS-0101	ADAR1000	*Unit*
G_t	Transmit gain	20	20	dB
G_r	Receive coherent gain	15	15	dB
NF	Receive noise figure	4	8	dB
P_{1dB}	Output power at 1dB compression	13	10	dBm
IIP3	Input 3^{rd}-order intercept	-16	-7	dBm

the AWS-0101 but there are a few key differences. One of the most notable differences is the ADAR1000's lack of ability to simultaneously receive in both polarizations and instead offers a programmable output pin, *TR_POL*, which is used to control an RF switch. This is not a huge issue (and will be discussed later to have been a complete non-issue in this application) because, unlike the sampling of in-phase and quadrature (IQ) signals, it is not imperative that they be sampled simultaneously. It is worth noting, however, that this will require an RF switch be added for each receive beam and some extra power and digital routing.

An "issue" mentioned previously was the ADAR1000's lack of simultaneous dual-polarization receiving. This project's requirements made it not an issue, however, because the AWS-0101's simultaneous dual-polarization receiving functionality was decided to not be used. This decision was based on three main factors: the project's objective to compare performance between the ICs, PCB layout difficulties, and code reusability. Designing one board to be simultaneous dual-polarization and the other to not would add complications in achieving a fair comparison between the two beamforming chips because it would add extra differences to be accounted for not directly related to performance. The AWS-0101's footprint also introduced some issues that made it difficult and more time consuming to route the horizontally and vertically polarized RF signals to their respective connectors. In the interest of time, it was also preferable to have as many similarities as possible between the boards so code written to interface with one could more easily be adapted to the other.

3.6 Polarization

3.6.1 Scattering Matrix

Single-polarization radar systems are limited in their measurement of meteorological events because they can only measure phenomena taking place in the same axis they are measuring. Since the goal of weather radar is to measure and classify complex events, it is important to increase the amount of available data. Dual-polarization (dual-pol) radar systems often consist of transmitting in a single polarization, most often horizontal or vertical, and receiving in two. This expands the back-scattered signal, S, from a single dimension to a 2x2 matrix of the form,

$$S = \begin{bmatrix} S_{hh} & S_{hv} \\ S_{vh} & S_{vv} \end{bmatrix} \tag{3.15}$$

where S_{hv} is the signal resulting from transmission in the horizontal polarization and reception in the vertical polarization and the rest of the matrix components can be inferred from this. The resultant electric fields in either polarization from transmission of an electric field is defined by,

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} S_{hh} \cdot E_{ht} + S_{vh} \cdot E_{vt} \\ S_{hv} \cdot E_{ht} + S_{vv} \cdot E_{vt} \end{bmatrix} \tag{3.16}$$

where E_{hr} and E_{ht} are the received and transmit electric fields in the horizontal polarization, respectively. The field received in the same polarization as the transmission is called the *co-polar* signal and the field in the orthogonal polarization is the *cross-polar* signal.

One dual-pol configuration is to transmit in a single polarization and receive in both simultaneously which will give either $E_{vt} = 0$ or $E_{ht} = 0$ and simplify (3.16) to one of the following electric field reception matrices,

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} S_{hh} \cdot E_{ht} + 0 \\ S_{hv} \cdot E_{ht} + 0 \end{bmatrix} \tag{3.17a}$$

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} 0 + S_{vh} \cdot E_{vt} \\ 0 + S_{vv} \cdot E_{vt} \end{bmatrix} \tag{3.17b}$$

for horizontal and vertical transmissions, respectively.

Alternatively, the signal can be transmit and received in both horizontal and vertical polarizations simultaneously which gives the same matrix as in (3.16) and is known as simultaneous transmit simultaneous receive (STSR). Since there is no time duplexing of the polarizations in either transmission or reception, it will be impossible to differentiate between the incoming signals resulting from either polarization. This effectively makes (3.16) become,

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} S_{hh} \cdot E_{ht} + S_{vh} \cdot E_{vt} \\ S_{hv} \cdot E_{ht} + S_{vv} \cdot E_{vt} \end{bmatrix} \rightarrow \begin{bmatrix} S_{hh} + S_{vh} \\ S_{hv} + S_{vv} \end{bmatrix} E_t. \tag{3.18}$$

where E_t could be in either polarization.

Both of these configurations are useful because they provide information about scatterers in both dimensions which is of great use in classifying meteorological events. This section primarily focuses on the hardware implementation of multiple-polarization systems in the context of DARMA, but the applications and products which can be derived from this new information are discussed further in Section 4.9 and in-depth analysis, derivation, and explanation can be found in [8].

3.6.2 Dual-Polarized Receiver

As previously mentioned, to obtain the dual-pol radar products, the radar will need to receive in both the horizontal and vertical polarizations. This creates the need for either simultaneous dual-polarization sampling or sampling one, switching to the other polarization, and sampling again at a fast rate. Of the two beamformers considered, the AWS-0101 and ADAR1000, only the AWS-0101 has multiple outputs for sampling both simultaneously, however, as discussed in Section 3.5, this feature was not used due to the ADAR1000 lacking this option and the desire for uniformity between boards. The choice was made to use the ADRF5019, a high-speed RF switch, to switch between each of the incoming polarizations. The ADRF5019 boasts an isolation between input ports (polarizations) of approximately 30dB at 10GHz which should be sufficiently large to mitigate interference between the two.

For simpler applications (i.e., applications requiring only single-polarization), it is possible to completely ignore the second receive polarization by maintaining the RF switch in the default configuration. In the case of the horizontal transmitter, this will give a single element in the scattering matrix, S_{hh}, which reduces complexity for the signal processor as well as a slight decrease in complexity for the controller.

3.6.3 Transmission Polarization

By default, the radar system described in this thesis supports transmission in a single polarization which is assumed to be horizontal for most of the discussion. As mentioned in Section 3.6.1, this only allows for measurement of two elements in the scattering matrix, S_{hh} and S_{hv}, which gives the reception matrix in (3.17a). This can be limiting and it may be preferable to support switching to a second polarization to obtain the other two scattering elements, S_{vh} and S_{vv}, separated by a switching time, T_{PRI}. This would introduce a change to the signal generation approach discussed in Section 3.3 in that instead of transmitting a continuous sawtooth wave, a single sawtooth would be transmit in one polarization, the co-polar (S_{xx}) and cross-polar (S_{xy}) scattering elements measured, then polarizations would be switched and the same measurements taken. This

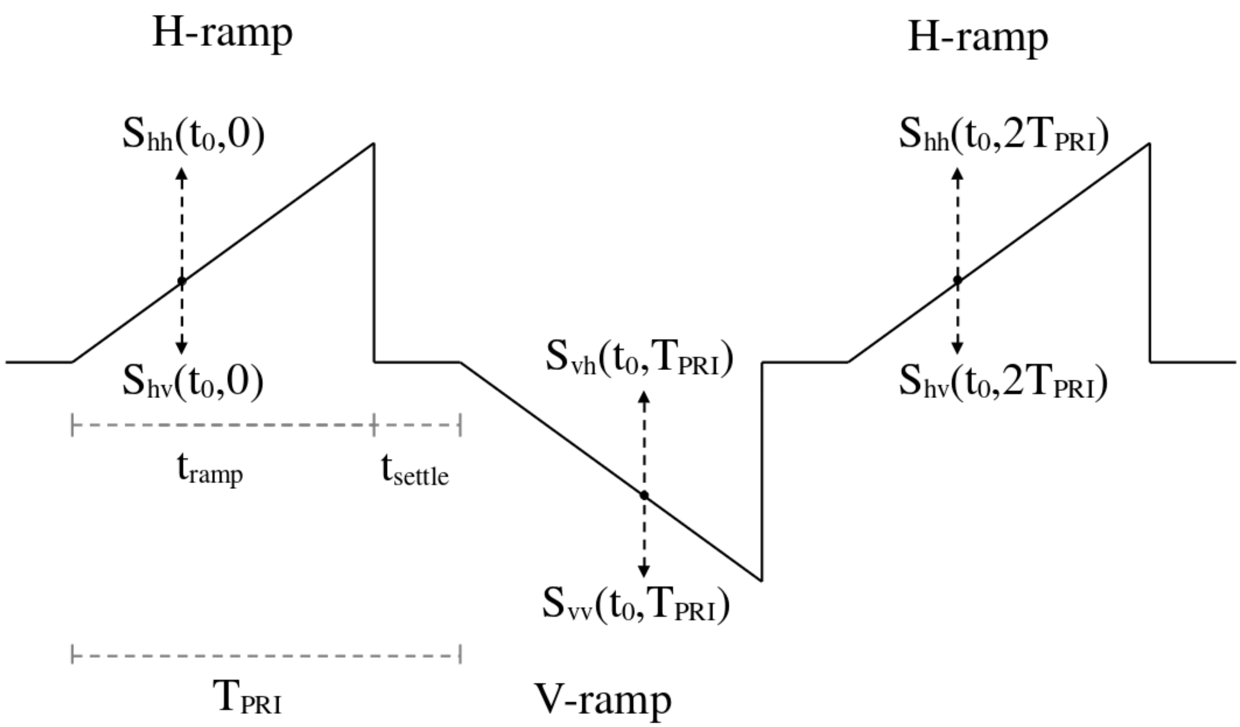

Figure 3.13: Illustration of switching transmission polarizations to obtain both sets of co-polar (S_{hh}, S_{vv}) and cross-polar (S_{hv}, S_{vh}) scattering elements. A single sawtooth wave is transmit and the respective scattering elements are received. The polarization is then switched after some settling time, t_{settle}, and the process repeated. $T_{PRI} = t_{ramp} + t_{settle}$ represents the pulse repetition interval between the horizontal and vertical pulses and t_0 is the round-trip time between the radar and scatterer which corresponds to a beat frequency, f_{beat}.

new process is illustrated in Figure 3.13 and is known as alternate transmit simultaneous receive (ATSR). Some time, t_{settle}, is taken between subsequent ramps for the transmitter to settle which ensures as little power from the "off" polarization is transmit. Leakage from the off polarization would result in ambiguity in the scattering matrix because it would be impossible to differentiate between backscatter from either polarization as in STSR. This is because the received electric field is the sum of that reflected from both transmit polarizations as described by,

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} S_{hh} + S_{vh} \\ S_{hv} + S_{vv} \end{bmatrix} E_t \tag{3.19}$$

where E_t is the transmit electric field which could be in either polarization [9]. For this situation where the transmit polarization is controlled and the leakage into the off polarization can be assumed to be negligible, the elements in (3.19) which result from the off transmission can be assumed to be zero, therefore giving only the co-polar and cross-polar elements for the desired transmission.

Another alternating transmission option is alternate transmit alternate receive (ATAR) which is used for the reception of only co-polar elements. This is done by only receiving in a single polarization when the same respective polarization is being transmit and results in the following scattering and electric field reception matrices,

$$S = \begin{bmatrix} S_{hh} & 0 \\ 0 & S_{vv} \end{bmatrix} \tag{3.20a}$$

$$\begin{bmatrix} E_{hr} \\ E_{vr} \end{bmatrix} = \begin{bmatrix} S_{hh} \cdot E_{ht} \\ S_{vv} \cdot E_{vt} \end{bmatrix}. \tag{3.20b}$$

The transmit waveforms which correspond to horizontal, vertical, and alternating transmission are shown in Figure 3.14 where $\hat{h}_i$ and $\hat{v}_i$ are the horizontal and vertical axes, respectively, and $\hat{k}_i$ is the axis of propagation of the incident wave. A summary of the modulation waveforms used in different polarization configurations including simultaneous, alternate, horizontal, and vertical transmission is shown in Figure 3.15 with the frequency as a function of time.

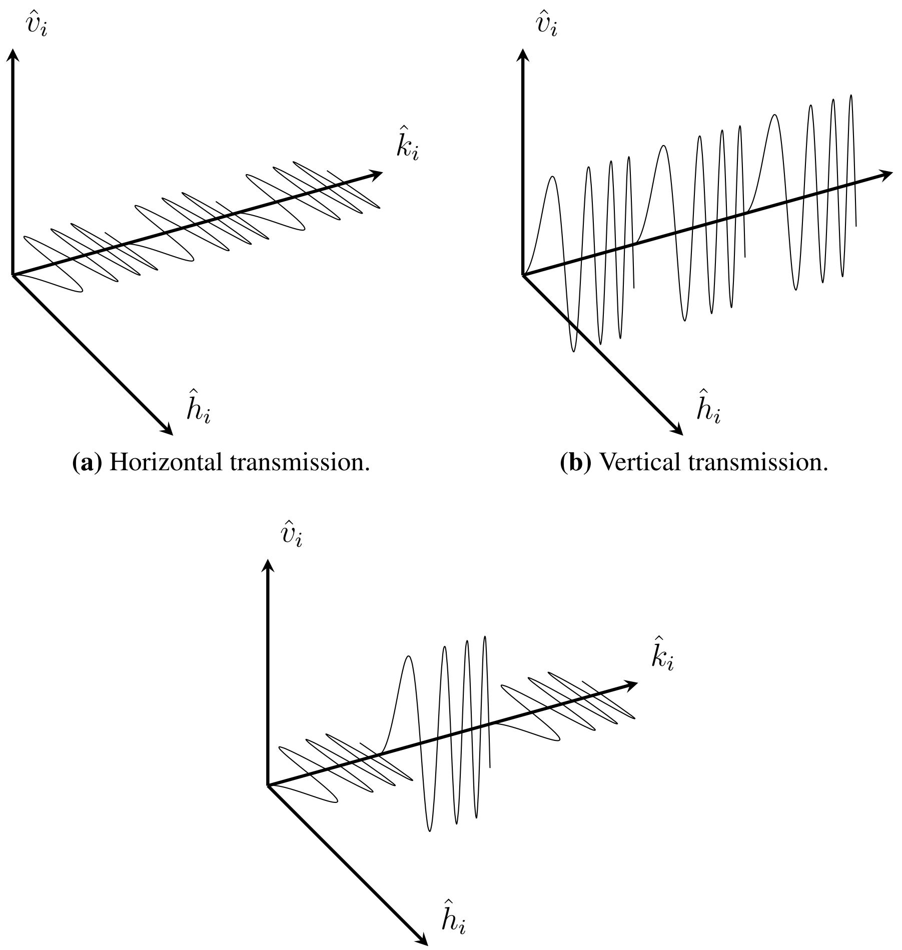

(a) Horizontal transmission.

(b) Vertical transmission.

(c) Alternating transmission.

Figure 3.14: Linear frequency-modulated transmit waveforms in the horizontal and vertical polarizations as well as alternating polarizations (simultaneous horizontal and vertical transmission is not shown). The horizontal and vertical axes are defined for the incident wave as $\hat{h}_i$ and $\hat{v}_i$, respectively, and the axis of propagation is $\hat{k}_i$.

To achieve the switching of transmission polarizations, a similar technique can be used as on the receive side which switches between incoming polarizations. This switch is not built directly into the main radar board due to the added complication and smaller general purpose need for it, but it can be built as a breakout board similar to that in Figure 3.16. Thanks to the modularity of this system, a transmit polarization switch breakout board would be simple to create and connect to the main board via SMA cables. If the design is followed from the receive side switches (Figure A.2), the breakout board would consist of the ADRF5019 (RF switch), control lines, a power supply, and possibly some transmission line length tuning to ensure no additional phase shifts are

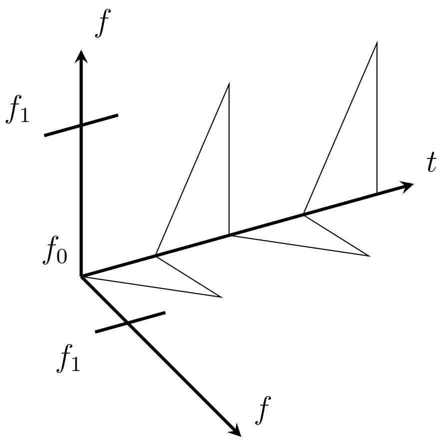

(a) Alternate transmission.

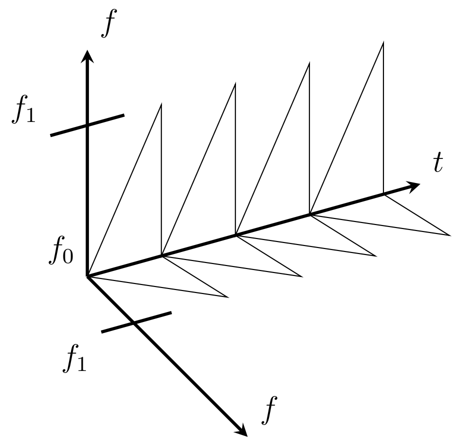

(b) Simultaneous transmission.

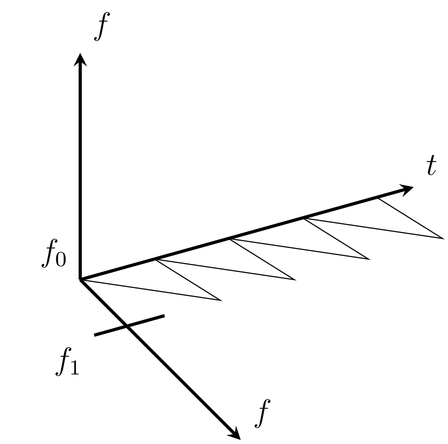

(c) Horizontal transmission.

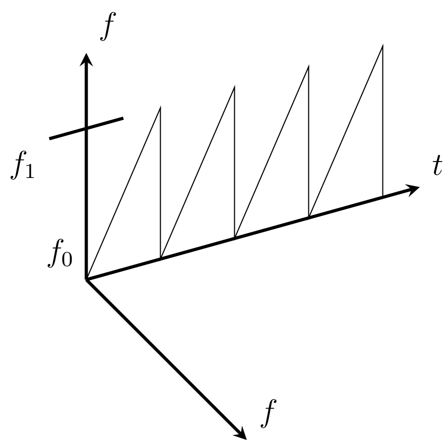

(d) Vertical transmission.

Figure 3.15: Summary of the linear transmit polarization configurations discussed in Section 3.6 including alternate, simultaneous, horizontal, and vertical transmission. This shows the sawtooth frequency modulation waveform as a function of time.

added with respect to the other outputs. There would then be eight outputs which correspond to alternating polarizations (horizontal and vertical) for each of the four signals. These outputs can be connected to any 4-element antenna array with separate horizontal and vertical polarizations. It will also be necessary to have precise timing information about when the transmit polarization is switched; without this information it will be difficult to determine with confidence which transmit polarization the received electric fields result from.

3.7 Lightweight and Portable Operation

Portability is a difficult consideration for a radar of this range (1000m), output power (30dBm), frequency band (X-band), and antenna configuration (phased-array). As previously discussed, since it is meant to drone mountable, the antenna in use must be small which can make achieving high gain, G_t, G_r, difficult. To compensate for relatively low antenna gain, the radar transmits a high output power, $P_t = 30$dBm, relative to other portable radars. Transmission of this much power requires fixed gain, high power amplifiers in addition to the beamforming stage that consume large amounts of power; in the case of the one selected for this application, Macom's MAAP-008924, 1W. This amplifier was chosen due to the reasons outlined in Section 3.5, however an amplifier with as much power consumption as this poses some issues for portability. The ability to be drone-mountable requires the entire system to be relatively small which makes heat dissipation, a critical design feature for high power applications, more difficult. The PAs being on the same board as the rest of the components would mean generating large amounts of heat in close proximity to the other components. For this reason the PAs were moved to a separate breakout board. The addition of another board is not ideal because it does not reduce the heat dissipation of the whole system and it adds more signal loss and financial cost for the interconnects between the two boards. The SMA cables and connectors used as interconnects range in performance which is generally a function of price but can typically be assumed to contribute 1-3dB/ft of loss at X-band. The breakout board does, however, offer the advantage of making the system more modular which allows for swapping of different PA stages or even use without PAs depending on the desired output power. There is also a reduction in number of required power supplies on the main board which makes the power plane routing simpler and reduces the number of possible failure points in the system. A 3-D model of the PA breakout board can be seen in Figure 3.16. It is small relative to the main board with dimensions, 1.5" by 3" which makes it simple to swap out depending on need.

Unfortunately, the AWS-0101 (one of the beamformers used) has an output 1dB compression point, $P_{1dB} = 13$dBm, which is low relative to the 30dBm required for this application. Ideally, the fixed gain power amplifier stage would be placed before the beamformer, but this compres-

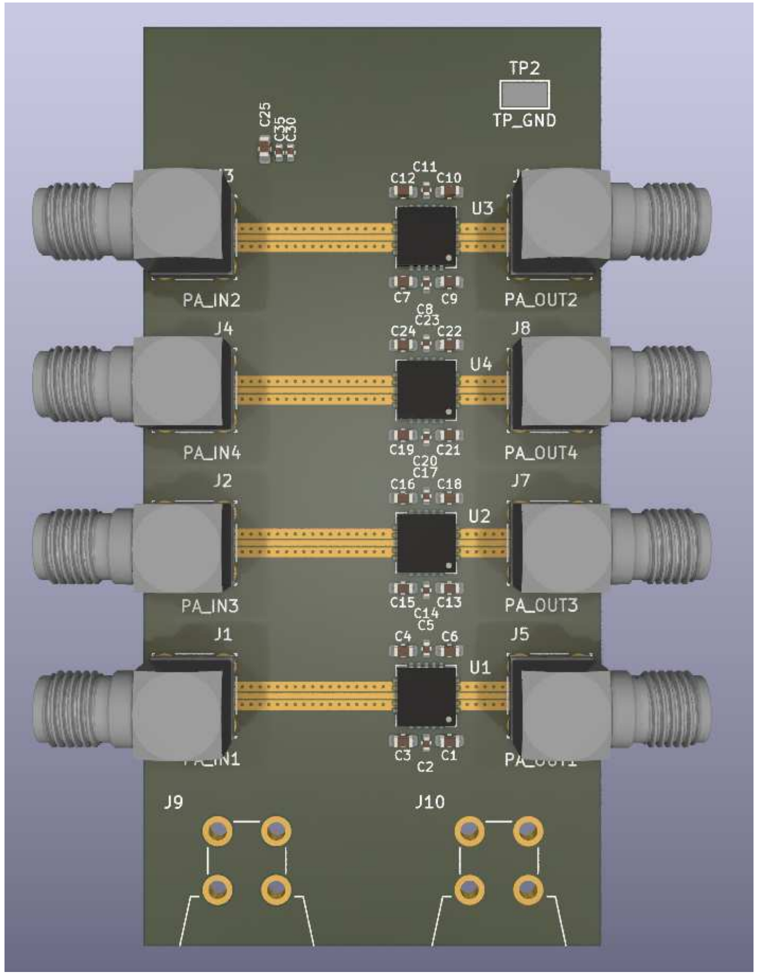

Figure 3.16: 3-D model of the power amplifier breakout board. This is meant as an optional addition to the outputs of the beamforming stage and can be swapped for various other power amplifier stages depending on project requirements.

sion point makes it impossible and must be placed after. The disadvantage of this is that rather than needing one PA, there are now four RF lines needing amplification, so the 1000mA the PA consumes becomes 4000mA; a very large amount for a small, portable radar. This adds multiple concerns that eventually led to the creation of the separate PA breakout board. Finding regulators or voltage converters capable of supplying 4000mA in a surface-mount form factor was near impossible (Section 6.2) and caused the PAs to need to be supplied individually. This complicates board layout by needing more high power fills on the power plane and increases cost. Power dissipation is also a major concern for a board of this size which is made much worse by the addition of this much power consumption. For reference, the main board only consumes approximately 3500mA for both versions, less than half than that of the PAs.

To further the power consumption issue, X-band components generally consume high amounts of power. The major components are listed in Table 3.4 with their power requirements along with board totals.

Table 3.4: Approximate power consumption for some of the most power intensive components and each of the radar boards described by the beamformer IC used.

Type	*Name*	*Power Consumption (mA)*
Components		
Beamformer	AWS-0101	1072
Beamformer	ADAR1000	400
PA	MAAP-008924	1000
VCO	HMC530	390
Mixer	HMC1113	210
Boards		
	AWS-0101	7160
	ADAR1000	5806

In selecting a battery there are many options specifically tailored to drones due to the large market for consumer and commercial drones. They come in different voltage levels in increments of 3.7V, so for this radar where the highest voltage needed was 6V, the choice was made to use a 7.4V battery. If using a low-dropout regulator (LDO), this allows for a dropout voltage up to 1.4V which would be very high for most quality LDOs. Drone batteries, which are typically lithium-ion (Li-Ion), have a wide range of capacities ranging from hundreds of milliamp-hours (mAh) to around ten thousand mAh. The size and weight are generally directly correlated to capacity. Near the upper limit of that range, the radar would last approximately 1-2 hours depending on which radar is in use (Table 3.4).

All these specifications and accompanying complications leave very tight margins for operation of the drone-based radar. For this reason, a ground-based version was built first. Optimally, the two versions would be the same, but for a complex project such as this, it is important to ensure a simpler case works properly before attempting to build a more complex system.

3.8 Composite Design

With much consideration of the topics discussed throughout Chapter 3, all the individual components must come together to form a cohesive system. Figure 3.17 shows the block diagram of the FMCW radar hardware which includes signal generation, transmission, reception, mixing to form an IF or baseband signal, filtering, and data acquisition. Once the data is acquired by the

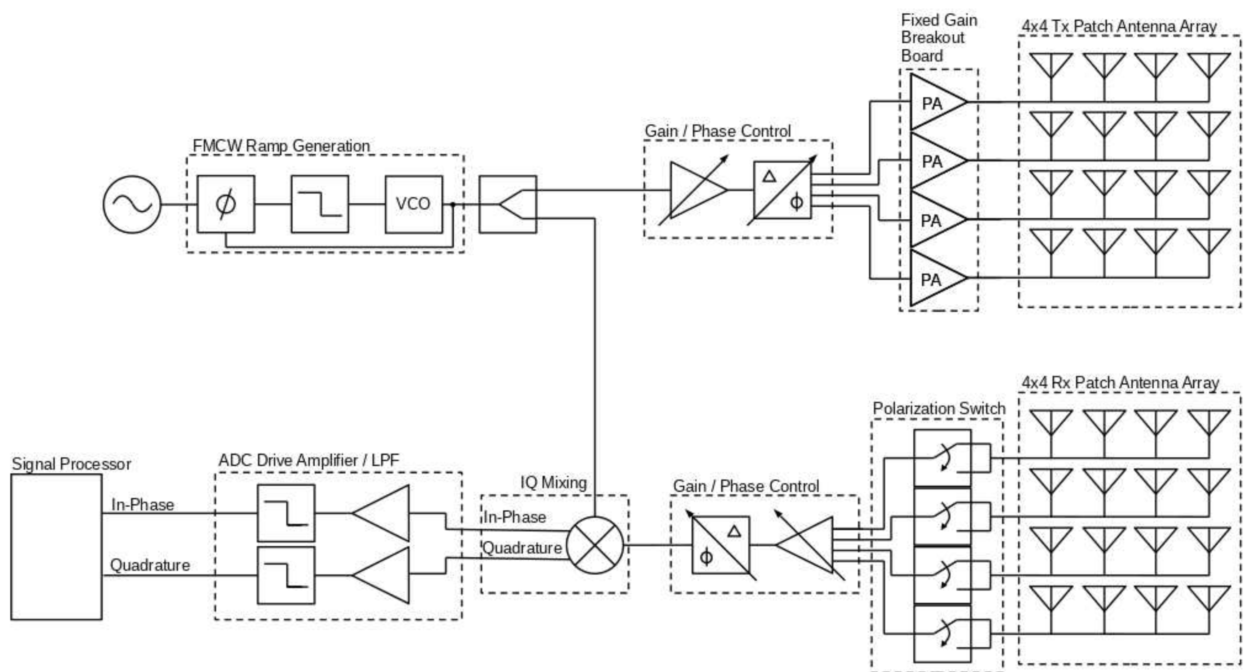

Figure 3.17: Block diagram of the finalized ground-based FMCW radar design including power amplifier breakout board and phased-array antenna.

ADC it is sent to the *signal processor* block which, despite its relatively small size in the block diagram, is one of, if not the most, integral components in the entire system. All the hardware components are centrally controlled by a microcontroller which the signal processor is a part of. Chapter 4 discusses in detail the measures taken to properly translate the voltages seen at the input to human-understandable data. Before beginning signal processing, some important decisions in the design of the board must be made to ensure optimal performance.

Detailed schematic realizations of Figure 3.17 for both radars are shown in Appendix A where they are split into transmit, receive, control, and power sections. The transmit section contains the signal generation, power splitting, transmit beamforming, and antenna subsystems and the receive section contains the polarization switch, receive beamforming, baseband mixing, data acquisition, and antenna subsystems. Digital control and power for these subsystems is shown separately. Component and fabrication costs are also detailed in Appendix B.

3.9 PCB Design

While many of the important hardware design decisions take place in the frequency planning (Section 3.2), range selection (Section 3.4), and component selection, proper physical layout is crucial for optimal radar performance, especially when operating at very high frequencies such as this.

High frequency is a loaded term because it can widely vary between applications. This radar transmits and receives at a center frequency of $f_c = 10.5$GHz and a bandwidth of $B_{RF} = 3$MHz (Table 3.1) which in nearly all applications would be considered high frequency. High frequency signals cause copper traces to diverge from simple circuits fundamentals and start acting as transmission lines. In RF applications, most components, antennas, and transmission lines are matched to an impedance of 50Ω and in the case of any mismatch can cause sharp degradation in performance. Because of this, it is very important to properly design these transmission lines. There are a variety of PCB traces to choose from, of which microstrip, stripline, and coplanar waveguide are some popular options. Microstrip is arguably the most common due to its simplicity and performance and an example is shown in Figure 3.18a. Stripline is very similar but is surrounded by two dielectrics rather than being on an outer layer as shown in Figure 3.18b. Details will be spared in the comparison of the two, but microstrip generally has lower dielectric losses (δ_d), higher propagation constants (ε_r), larger trace widths, and higher radiation losses [10]. Coplanar waveguide (CPW) is a transmission line type that has been growing in popularity due to its performance as frequencies approach millimeter wavelength ($f \approx 30$GHz, $\lambda \approx 10$mm) [11]. The line cross-section shown in Figure 3.18c is a variation of coplanar waveguide called *grounded* coplanar waveguide (GCPW), named for its addition of a reference ground plane on the bottom of the dielectric. This

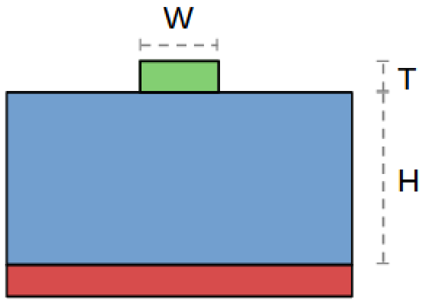

(a) Microstrip transmission line.

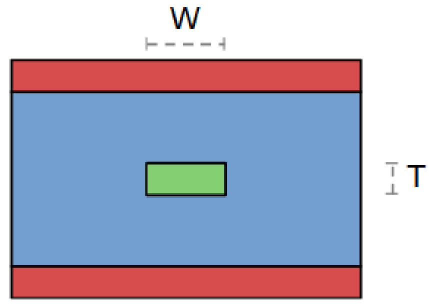

(b) Stripline transmission line.

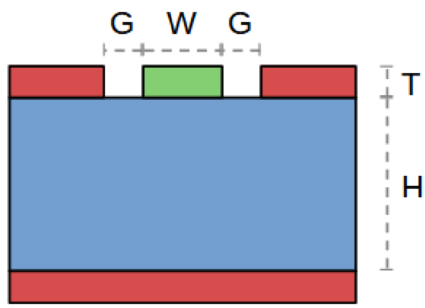

(c) Grounded coplanar waveguide transmission line.

Figure 3.18: Cross-section of three commonly used PCB transmission lines: microstrip, stripline, and grounded coplanar waveguide.

radar is not quite millimeter wave, but it is still beneficial to surround the transmission line with a ground fill separated by a gap. Beyond simulation to ensure optimal performance is met, it is not necessary in the design of a radar system to derive the gap, width, height, and thickness for a desired characteristic impedance and frequency. There exist numerous calculators for finding these parameters, one of which was used to calculate the transmission line parameters in Table 3.5.

The dielectric used is Rogers 4350B, one of the most common for high frequency applications. It has very low dielectric losses, $\delta_d = 0.0037$ at high frequency, $f = 10$GHz, which is very important for maintaining signal integrity while propagating on the PCB. It does, however, come at a significantly higher financial cost relative to dielectrics commonly used at low frequencies such as FR4. There is also limited production and it can be difficult to find fabrication companies with the ability to source it. More common dielectrics, like the aforementioned FR4 which has a

Table 3.5: Calculated grounded coplanar waveguide (GCPW) parameters for the RF traces.

Parameter	*Description*	*Value*	*Unit*
f_c	Center frequency	10.5	GHz
Z_0	Characteristic impedance	50	Ω
G	Gap between signal trace and ground fill	10	mil
W	Width of signal trace	18.5	mil
T	Trace thickness	1.37 (1)	mil (Oz Copper)
H	Dielectric height	10	mil
ε_r	Dielectric relative permitivitty	3.48	-
δ_d	Dielectric dissipation factor	0.0037	-

loss tangent, $\delta_d = 0.016$ at $f = 1$GHz (FR4 performance at $f = 10$GHz is normally not specified), generally do not have acceptable high frequency performance, especially for a system such as this with small margins for error.

Physical layout for mixed signal (digital and RF) boards is difficult because there are many transmission lines, often running at different frequencies and rise/fall times, to be routed in close proximity. RF transmission lines are very sensitive and require careful routing and supporting elements to maintain optimal performance. Since there are four transmission lines routed to each of the beamformer ICs, it would be impractical to keep the lines straight which is sub-optimal due to the performance degradation incurred by line bends. To minimize the drop in performance, there are some general rules and alternate curve types to use. Curved traces can be approximated to have near identical performance to a straight trace if the curve's radius is $r \geq 3W$, where W is the trace width. Following this rule-of-thumb can result in a large area taken up by the trace which is a limited resource. Mitered bends can be used instead if board space is limited and, if properly designed, have as good or sometimes better performance than a bend [12].

An RF line and the ground plane it is referenced to effectively form a waveguide within which the signal propagates. Rather than a traditional waveguide which has air in the center, this waveguide has the dielectric, in this case Rogers 4350, in the center. The dielectric expands to the board edges and it has many other signals running through it which can be very detrimental to electromagnetic interference (EMI). To ensure the RF signal stays within the effective waveguide, a *via fence* is placed surrounding the RF trace. A via fence consists of vias placed closely together

($d \lessapprox \lambda/10$) which, to the RF signal, resembles a continuous surface through which it cannot pass. This is a very common and important design practice for routing RF lines because it is very effective at reducing EMI as much as possible. Figure 3.19 shows the via fence on one of the transmission lines on the ADAR1000 radar board.

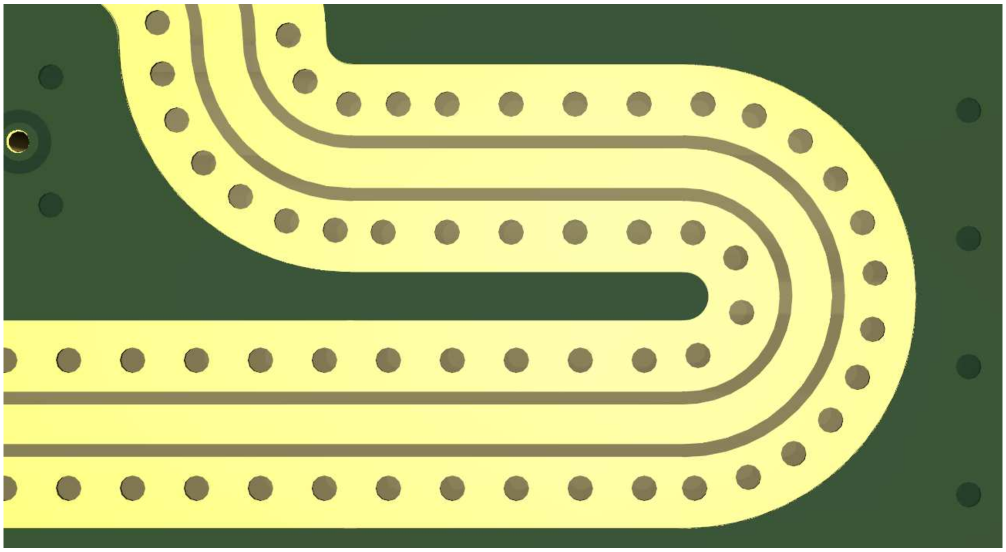

Figure 3.19: Via fence surrounding a curved RF transmission line on the ADAR1000-based radar board. The via fence spacing used is $d = 30$mil.

Taking everything previously mentioned and more into consideration, the two boards for each of the beamforming ICs were designed and are shown in Figure 3.20. They are approximately 5.5"x3.5" in size which most drones should be capable of carrying.

3.10 Summary

The design decisions discussed in this chapter were made to accomplish the specifications in Table 3.1, but there are numerous other applications for which a radar system similar to this would be useful. The application for this radar limits its range because of the need for the antenna to be relatively small. The antenna it uses is an 8x4 microstrip patch array with a gain, $G = 10$dBi at $f = 10.5$GHz, horizontal and vertical beamwidths, $\theta_1 = 10°$ and $\phi_1 = 16°$, respectively, and is seen in Figure 3.21 with its antenna pattern measured and plotted in Figure 3.22. This is also used

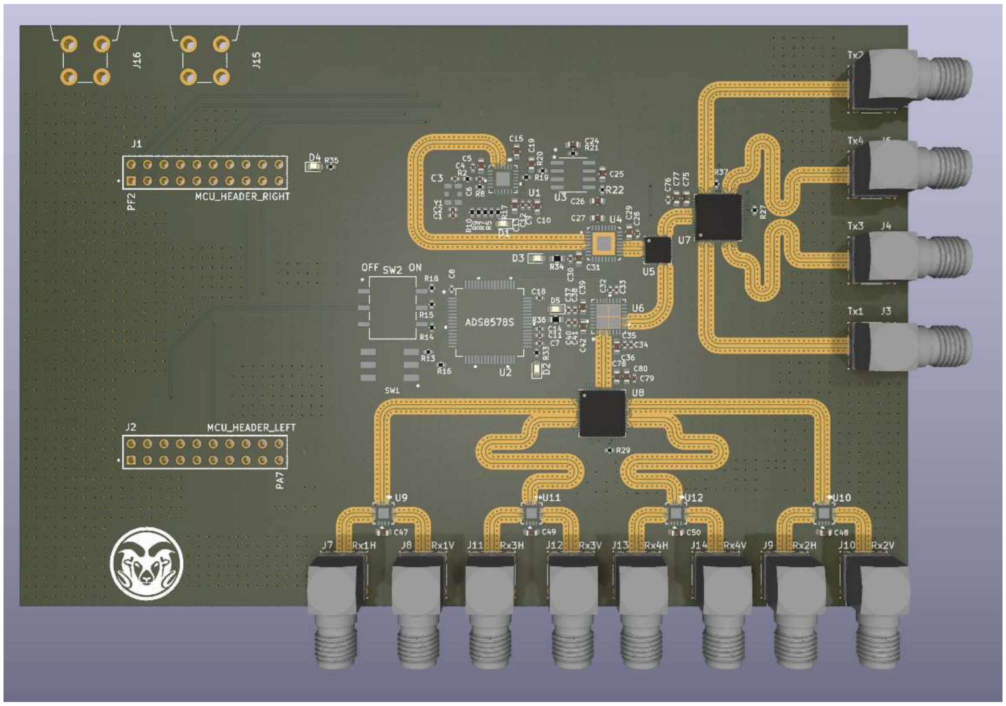

(**a**) 3-D model of the X-band, ground-based radar system PCB which uses the AWS-0101 beamforming IC from AnokiWave.

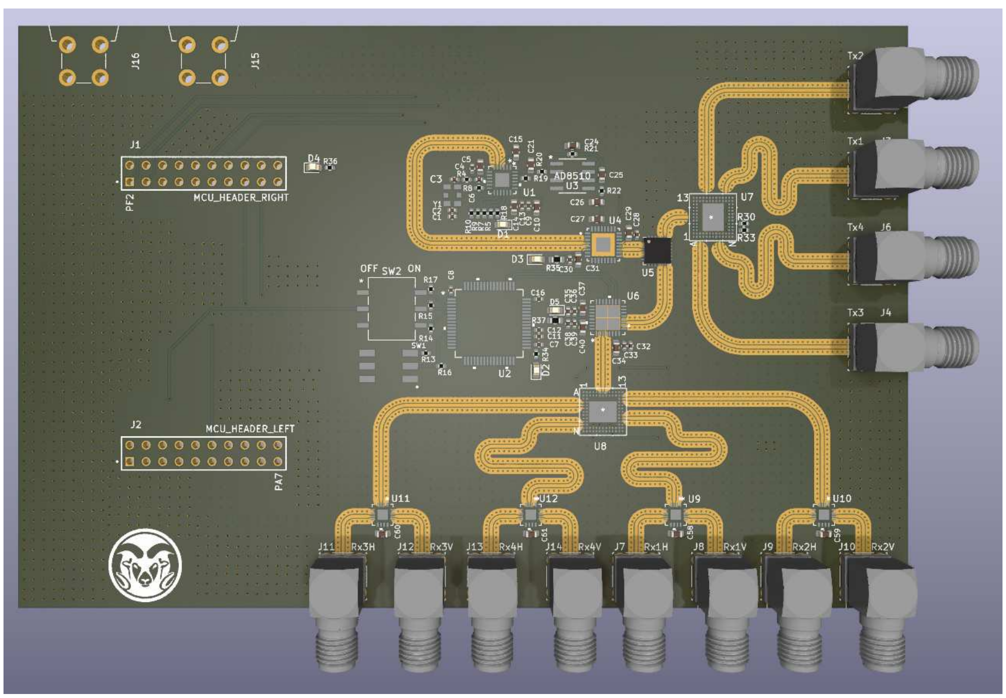

(**b**) 3-D model of the X-band, ground-based radar system PCB which uses the ADAR1000 beamforming IC from Analog Devices.

Figure 3.20: 3-D models of the two radar PCBs designed for ground-based operation using each of the two discussed beamforming ICs, Analog Devices' ADAR1000 and AnokiWave's AWS-0101.

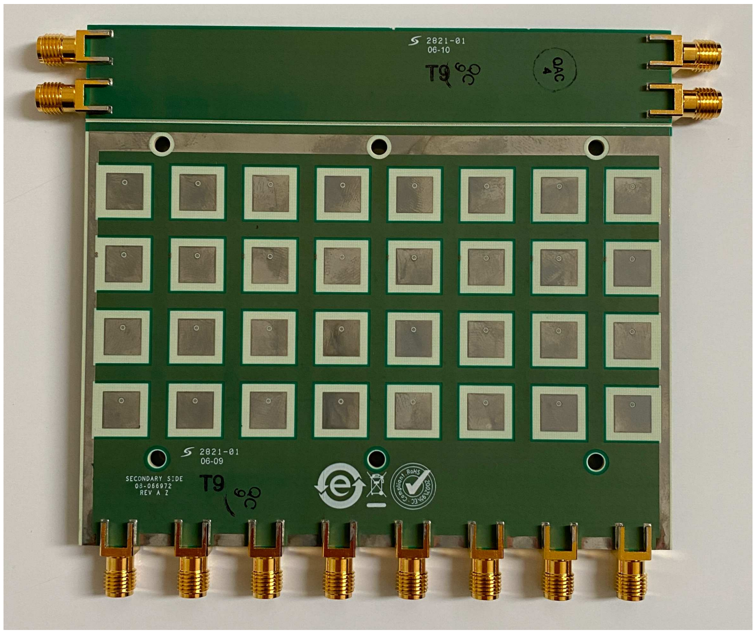

Figure 3.21: 8x4 element microstrip patch array antenna with a gain, $G = 10$dBi at $f = 10.5$GHz and horizontal and vertical beamwidths of $\theta_1 = 10°$ and $\phi_1 = 20°$, respectively.

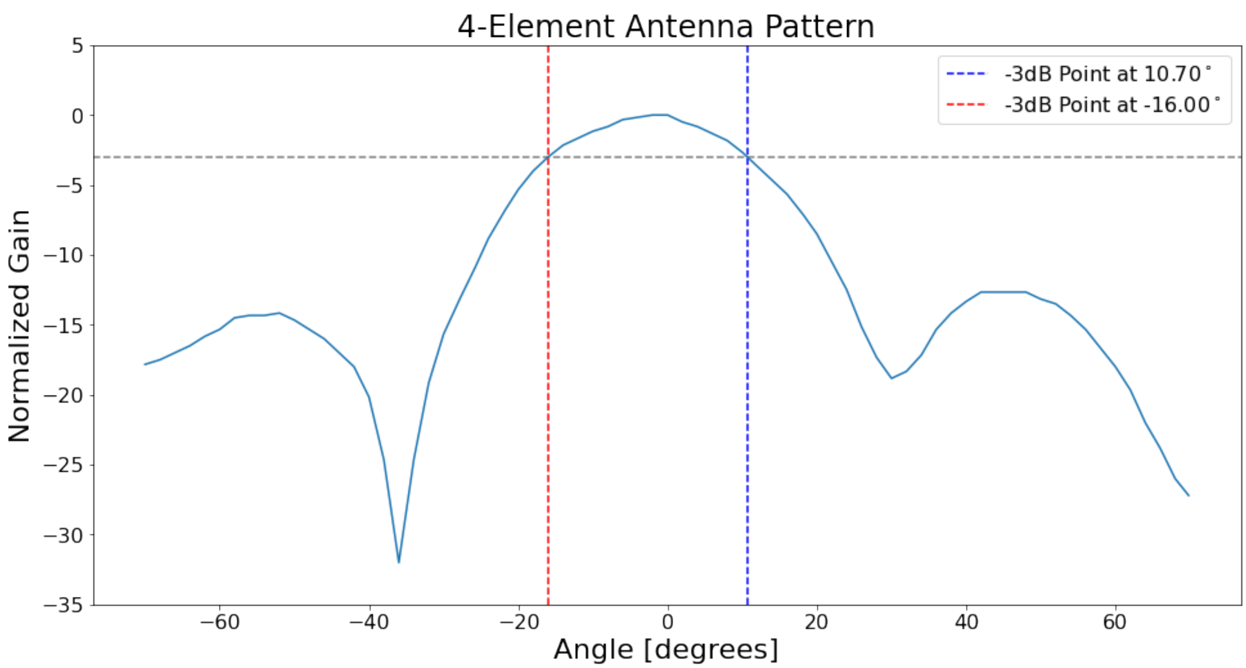

Figure 3.22: Normalized vertical antenna pattern of the 8x4 element patch array using one tap to show performance of the 4-element uniform linear array portion. Not gain tapering is applied.

for measurements obtained by the phaser board in Chapter 4, except the transmissions are sent from a single 1x4 tap, giving a large horizontal beamwidth, and it receives with all eight taps (8x4) for finer beam-steering. The antenna is not capable of dual-polarization which was not an issue for the phaser board, but in the future an antenna upgrade will be needed to achieve DARMA's maximum functionality.

If range is of a higher priority than size, it is possible to replace the small patch array with a larger, higher gain antenna. It is even possible to use only one of the outputs on the transmit and receive sides with a parabolic antenna, completely bypassing the phased array. Parabolic antennas, depending on size, are capable of achieving gains much larger than this patch array. This is one of the key differences discussed in Chapter 5 and is one of the main factors which allows the NOAA snow-level radar to achieve a range larger by a factor of 10, $R_{max} = 10$km. If the radar were used this way, the unused inputs and outputs should have 50Ω SMA loads attached and beamformer gain weights set to the minimum to minimize noise from them.

If switching to a parabolic antenna with higher gain fits the requirements of a certain project more than the phased array presented in this chapter, then it may also be worth considering operating at a lower frequency band as well. In Section 3.2.2 the frequency band was chosen to be X-band ($f_c = 10.5$GHz), mostly because of the desire to have beam-steering via a phased array. S-band was also considered because of its lower path loss (3.1), unrestricted access (ISM band), and simpler design requirements but was removed from contention due to the increasing antenna size requirements. For the situation where size is not a strict requirement, S-band should be considered for lower path loss which increases range. It would also remove the complexity introduced by a phased array architecture.

Conversely, it may be optimal for a radar to have fine resolution with lower range. Though this radar was designed for $R_{res} = 50$m and $R_{max} = 1000$m, it is entirely possible to simply increase the RF bandwidth, B_{RF}, which is directly proportional to resolution (2.3b) and filter the range-FFT spectrum to only include shorter ranges. This was done for the phaser radar in Chapter 4 because most of the testing was done indoors which allowed for a much finer resolution of $R_{res} = 0.3$m at a bandwidth of $B_{RF} = 500$MHz.

In addition to the suggestions mentioned above, there are many other customization options within the radar's software and slight modifications to the hardware can also be made to completely change it for different applications.

Chapter 4

Signal Processing

4.1 Introduction

Hardware design is a very important aspect of radar systems because it sets the limits for the radar's ability to measure, however, the data produced at the output of the ADC are merely voltages waiting to be deciphered. Signal processing is arguably more important than, or at least of comparable importance to, the hardware because it allows the engineer and other users to derive some meaning from the data.

Since in FMCW radar the range information is stored in the baseband frequency (Chapter 2), the processing of incoming signals starts with the discrete Fourier transform (DFT) which transforms the incoming signal from the time domain to the frequency domain. This alone does not enable the range calculation of targets because there exists interference that must not be assumed to be targets. This warrants the use of *detection methods* to separate noise from targets. From here, more meaningful processing can be done to translate the data to the user.

Most of what is discussed and shown in this chapter is in the context of Analog Device's X-band "phaser" board. It uses two ADAR1000 beamforming ICs on the receive side with an 8x4 patch array for eight-element horizontal steering and beamforming with approximate horizontal and vertical beamwidths of $\theta_1 = 9°$ and $\phi_1 = 18°$, respectively (shown in Figure 3.21). The transmitter is a single output connected to a 1x4 patch array allowing for no steering or beamforming with approximate horizontal and vertical beamwidths of $\theta_1 = 72°$ and $\phi_1 = 18°$, respectively. Since the horizontal beamwidth of the test antenna is so large, beamforming and beam-steering will be extremely important for only detecting targets in the desired field-of-view. The radar is highly configurable and will be configured separately for each application and example discussed here. Table 4.1 is a common configuration used for some of the examples and Figure 4.1 shows the radar setup outdoors with minimal clutter at boresight and some clutter at the edges. It uses

Table 4.1: Configured parameters of the Analog Devices FMCW phaser board. This configuration is used for many of the examples and figures discussed in Chapter 4.

Parameter	*Description*	*Value*	*Unit*
f_s	Sample rate	600	kHz
f_c	Center frequency	12.1	GHz
f_{IF}	Intermediate frequency (IF)	100	kHz
B_{RF}	RF bandwidth	500	MHz
N_{buffer}	Buffer size	2^{14}	-
T_{ramp}	Ramp time	1	ms

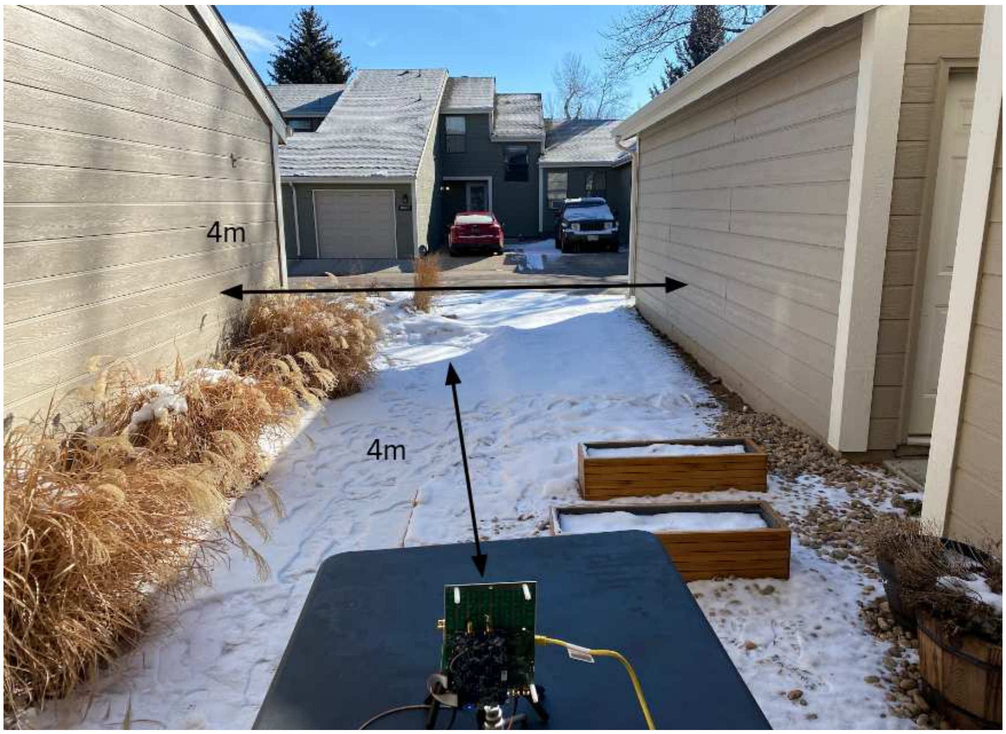

Figure 4.1: FMCW Phaser radar outside setup used to take measurements and generate some of the figures in Chapter 4. At boresight the signal can travel further than the radar configuration allowed, but some clutter can be seen on either side in the form of walls. This is a more ideal setup than the indoor lab setup used for some figures but an open field would be best to avoid this clutter.

the Pluto-SDR which is a software-defined radio (SDR) from Analog Devices for the sampling at a rate in the range, $0.6 \leq f_s \leq 61.44$MHz. The sampling rate is set to the minimum to achieve a fine frequency resolution for a given buffer size which is helpful for signal processing. For the local oscillator (LO) generation, the same frequency synthesizer from Section 3.3 is used, the ADF4159. In addition to these similarities, the phaser board uses the same general block diagram seen in Figure 2.4 which makes it very similar to the radar described in Chapter 3 with a few key differences such as available polarizations. This makes it a good resource for testing signal processing methods and algorithms before designing a custom radar and is why it is used for this chapter.

4.2 Windowing

Windowing is an often overlooked component in digital signal processing despite its large impact on the performance of the ADC. The incoming signal effectively applies a rectangular window without any prompt from the engineer by way of the buffer having finite length. This, along with all other windows, have two main effects, main-lobe width and side-lobe level shown in Figure 4.2. Main lobe width is very important because it affects the frequency resolution and

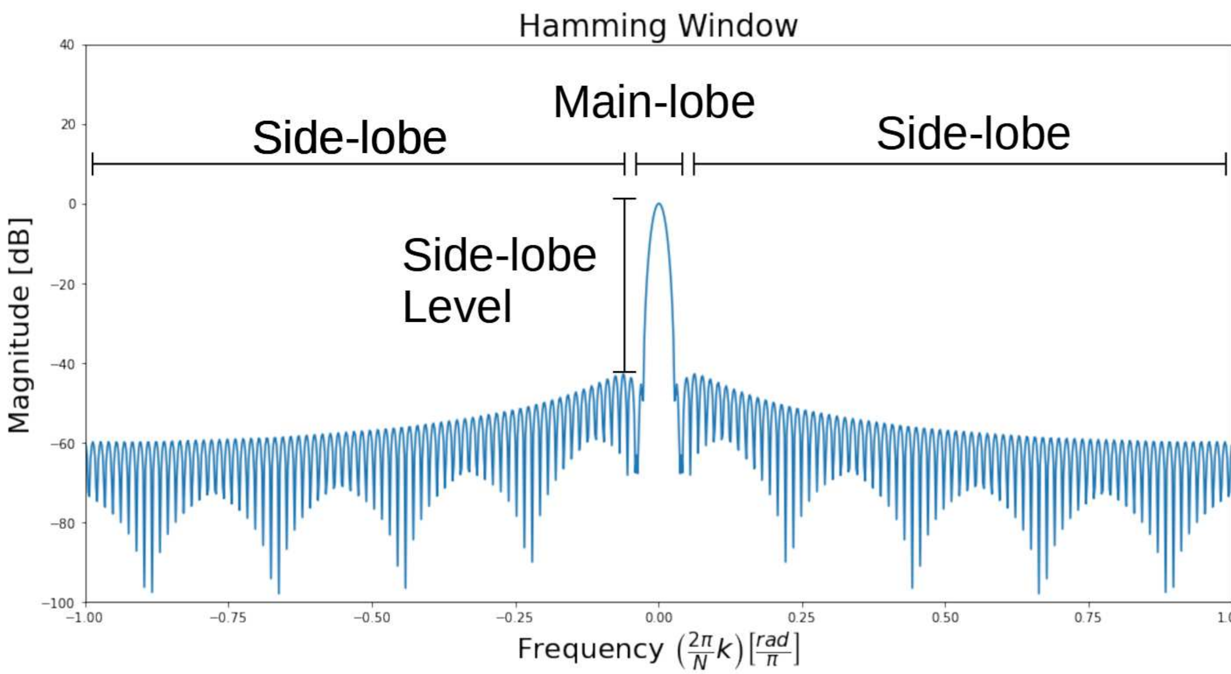

Figure 4.2: Main-lobe and side-lobes labeled on a Hamming window ($M = 145$).

therefore the range resolution. Ideally, it would be infinitesimally small and would not contribute to any beam-widening, but this is only possible with an infinite length window (ADC buffer). The relationship between window length and main-lobe width as well as peak side-lobe level is shown in Table 4.2 and can be researched further in [13]. While a main-lobe width of approximately $\pi/20$

(for length, $M = 145$ as seen in Figure 4.2) may seem inconsequential, in the case of the system in Chapter 3 with a maximum baseband frequency of 2kHz which corresponds to a range of 1km (Table 3.1), a range, R would be expanded to $R \pm 55m$, adding a significant amount of uncertainty as to where in this range the actual target(s) resides.

Table 4.2: Peak side-lobe amplitude and relationship between main-lobe width and window length for common windowing functions.

Type	*Peak side-lobe amplitude [dB]*	*Approximate main-lobe width [rad]*
Rectangular	-13	$\frac{4\pi}{M+1}$
Bartlett	-25	$\frac{8\pi}{M}$
Hanning	-31	$\frac{8\pi}{M}$
Hamming	-41	$\frac{8\pi}{M}$
Blackman	-57	$\frac{12\pi}{M}$

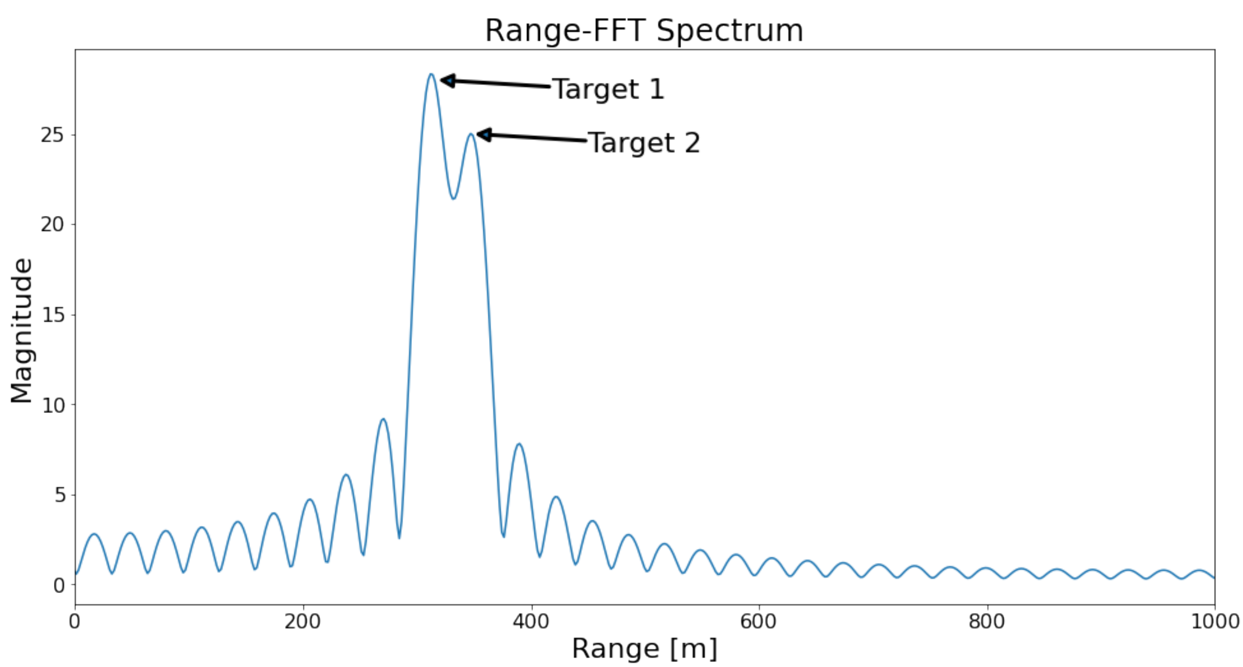

Figure 4.3: Example of a situation where main-lobe width and side-lobe amplitudes contribute to difficult differentiation between multiple targets.

This is made worse by the addition of side-lobes that, when high enough, contribute to the difficulty of differentiating between separate targets. Side-lobes exists next to the main-lobe meaning if they are high enough, they can contribute a non-negligible amount to the main-lobe's width, further increasing the ambiguity region where one or more targets may reside. An example of where this can be an issue is shown in Figure 4.3. Here two targets are very close in range to one another and the use of an imperfect windowing function adds width to the target in the range-FFT spectrum making it difficult to determine whether this is one target or multiple. This can be counteracted by a number of methods, including using the Doppler effect to determine the velocity of the two targets relative to one another (Section 4.10) and the constant false alarm rate (CFAR) method (Section 4.7); the Doppler effect method is only useful if the two targets are moving at unequal velocities.

4.3 Range Binning

Range binning for FMCW and pulsed radars have a fundamental difference in that the range information is stored in the time domain for pulsed radars and the frequency domain for FMCW. For pulsed radars grouping signals into range bins is important because it is time that separates the targets in range within a pulse interval. FMCW radar is different in that the frequency spectrum in one buffer of time contains target information at all possible distances within the maximum range. This changes the range binning process to be more for the purpose of separating parts of the frequency spectrum into range resolution-sized bins.

Often, the frequency corresponding to the range resolution, calculated using B_{RF} in (2.3b), will be larger than the frequency resolution from the buffer size, N_{buffer},

$$f_{res,buffer} = \frac{1}{N_{buffer}T_s}. \tag{4.1}$$

For example, in the case of a radar with parameters in Table 4.1, the range resolution, $R_{res} =$ 3cm corresponds to $f_{res,range} = 100$Hz while the buffer size, $N_{buffer} = 2^{14}$, gives a frequency resolution of $f_{res,buffer} = 36.62$Hz. Approximately 3 frequency points of size $f_{res,buffer}$ fit inside one range bin, $f_{res,range}$, meaning those 3 points cannot be assumed to be accurate individually.

For this case, the mean of the 3 points within each range bin should be taken and a new array of data created as explained by,

$$V = [\underbrace{V_1, V_2, V_3}_{\text{Range bin 1}}, \cdots, \underbrace{V_{N-2}, V_{N-1}, V_N}_{\text{Range bin } M}] \tag{4.2a}$$

$$\rightarrow [V_1, V_2, \cdots, V_{M-1}, V_M] \tag{4.2b}$$

where M is the number of range bins. This also, in many cases, allows for shrinking of the number of points in the frequency spectrum which helps with speed when running a real-time display.

4.4 Range Normalization

Electromagnetic waves suffer high attenuation when propagating through a medium such as air. The attenuation, described in (3.1), is proportional to R^2 which makes more distant targets appear to reflect significantly smaller amounts of power and can further lead to an incorrect inverse correlation between distance and size. It is possible to negate this effect by a process known as *range normalization* which applies normalizing factors, inversely proportional to range, to the received power spectrum. Applying the normalizing factor to a single range bin, $V(R)$, at a distance, R, becomes,

$$V(R) \rightarrow V(R)\left(\frac{R}{RN}\right)^2 \tag{4.3}$$

where RN is the range normalization constant. The R^2 in the numerator counteracts the path loss experienced by the signal at the current range. This makes the signal intensity appear as if it were observed at a distance, RN, from the radar [14]. For demonstration purposes, Figure 4.4 shows a PPI scan with and without normalization applied where RN was chosen to be 1m. There are two targets in focus, a reflector at a range, $R = 1.5$m, 0° off boresight and a human at $R = 2.5$m approximately 20° off boresight. Both figures show good performance in detecting the reflector panel because it reflects a much larger amount of power than the non-metallic human. The scan without normalization (Figure 4.4a) shows a faint bright-band at the specified range and angle off bore-

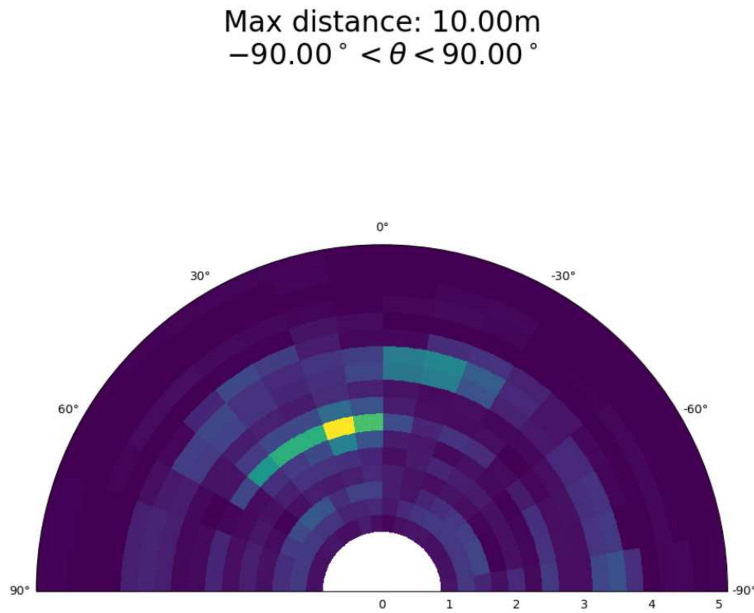

(a) PPI scan without range normalization.

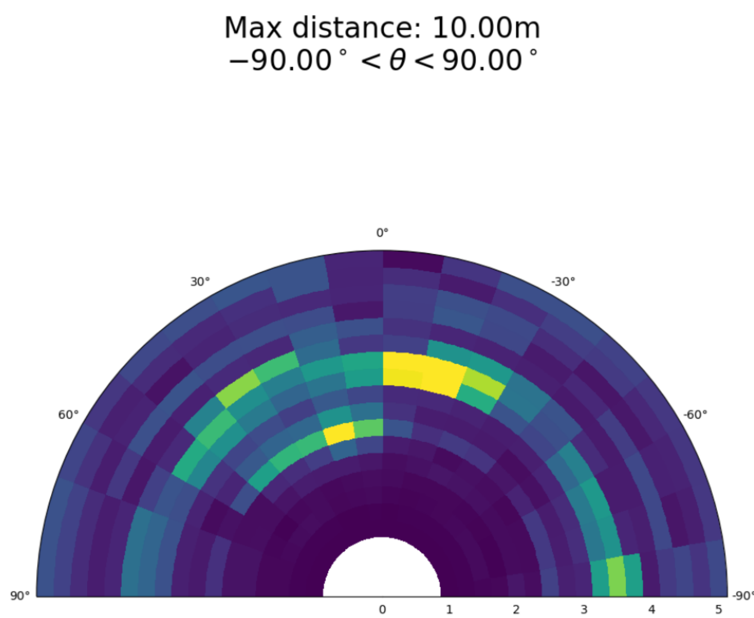

(b) PPI scan with range normalization.

Figure 4.4: PPI scan with and without range normalization applied. Two targets exist, a reflector panel at $R = 1.5$m, 0° off boresight and a human at $R = 2.5$m, 20° off boresight. The normalization range was chosen to be $RN = 1$m and the maximum round-trip range to $R_{max} = 10$m.

sight of the human, but it would be difficult to differentiate between this and noise for the viewer. Figure 4.4b provides a much clearer representation of the human target in the form a very bright band at the expected location. Despite being less reflective than the reflector panel, the human has a higher intensity which is indicative of a larger cross-section. For a user without extensive radar knowledge, this range normalization would prove very useful for proper data interpretation.

While range normalization is effective at bringing distant targets' power levels to match the power levels of near targets, it brings the noise up with it as well. Target detection techniques, described in Section 4.7, may be less effective due to the decreased signal-to-noise ratio at mitigating false alarms thereby increasing the need for proper filtering.

4.5 Filtering

4.5.1 Noise

Noise is very important in all microwave systems and especially in radar because it determines the subset of all targets which are detectable. Large amounts of work, much of which is described in Chapter 3, goes into hardware design for mitigation of noise. There is still a lot to do in the signal processing to ensure it affects the performance as little as possible.

A system's noise level varies significantly based on many factors including system temperature, bandwidth, system components, and external electromagnetic sources such as other radars, cell phones, or even the sun. Because of this it would be a waste of time to define an exact noise level because it varies temporally and spatially. It does, however, generally follow a uniform distribution across the frequency spectrum; this is known as the *noise floor*. Calculation of the noise floor is relatively simple if a frequency spectrum can be obtained and the noise can be assumed to have zero mean. It is defined by the equation,

$$N_{floor} = \frac{\sigma^2}{2} \tag{4.4}$$

where σ^2 is the variance or standard deviation squared [15]. This is a very simplistic model for the system's noise because objects, especially non-target objects, in the radar's view will contribute a non-uniform amount of clutter across the frequency spectrum. Other contributors include transmit to receive antenna leakage, phase noise from the frequency ramp generator, and much more.

4.5.2 Antenna Leakage and Low-Frequency Noise

Noise exists at all frequencies at varying magnitudes, but it generally averages to some *noise floor* which is discussed further in Section 4.5.1. One of the noisiest sections of the frequency spectrum is that surrounding 0Hz, or DC. Low frequency noise can be caused by a number of factors, one of which is transmitter to receiver leakage. In a small form factor such as that on a drone, it is more difficult to mitigate leakage from the transmitting antenna to the receiving antenna because they must be placed in close proximity. The leakage has effectively no time delay or, $f_{Rx} \approx f_{Tx}$ which when mixed with together corresponds to near zero frequencies. This can present an issue for FMCW radars because the target's beat frequency, f_{beat}, is directly correlated to its range which means, for collision avoidance systems such as one of the possible applications for the drone radar described in this thesis, closer targets will be more obstructed by noise and more difficult to detect. This is partially self-mitigated, however, because RF signal power when propagating is relative to R^{-2} which means closer targets will be attenuated significantly less and have far higher return powers. An example of this is shown in Figure 4.5a where the radar is pointed directly at a reflecting panel approximately a meter away. Here the leakage is overshadowed by the high reflected power. In addition, the signal is shifted to an IF of $f_{IF} = 100\text{kHz}$ avoiding the DC interference.

The frequency spectrums obtained in Figure 4.5a and Figure 4.5b use sampling rates of $f_s = 600\text{kHz}$ and $f_s = 2400\text{kHz}$, respectively with a large buffer size of $N = 2^{14} = 16384$. All else equal, the sampling rate in Figure 4.5a allows for a resolution four times smaller than that in Figure 4.5b as is described by

$$\Omega_{res} = \frac{2\pi}{NT_s} \tag{4.5}$$

where Ω_{res} is the periodic frequency resolution in radians and T_s is the sampling period. This finer resolution allows for a greater distinction between close peaks.

If the radar in use transmits low power, has very poor transmit-to-receive antenna isolation, or is attempting to measure very small or non-reflective targets, DC and leakage noise can overshadow real targets when a detection algorithm, discussed in Section 4.7, is applied. It is there-

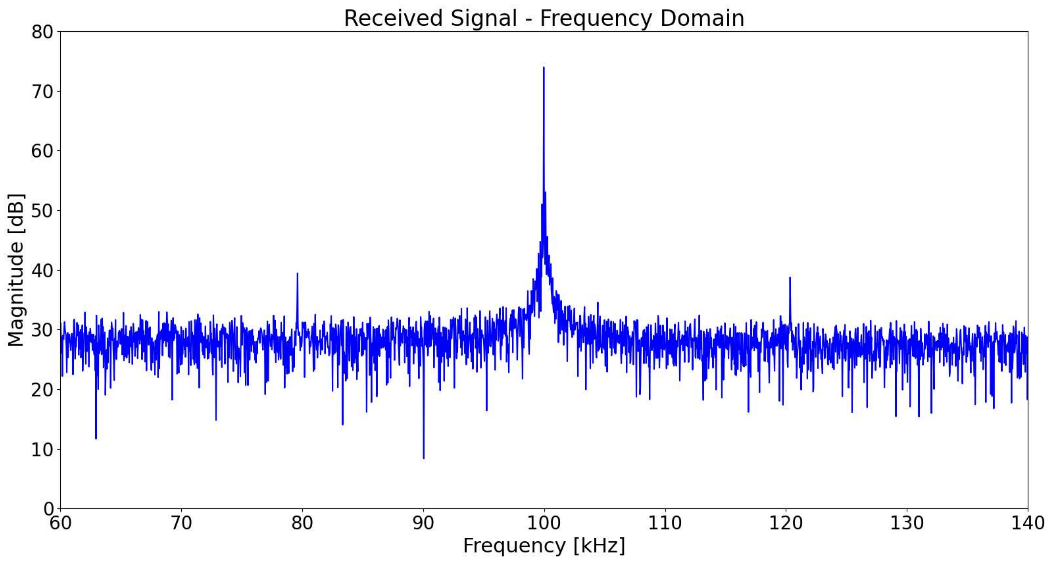

(a) Fine resolution frequency spectrum obtained using $f_s = 600\text{kHz}$ and $N = 2^{14}$ samples.

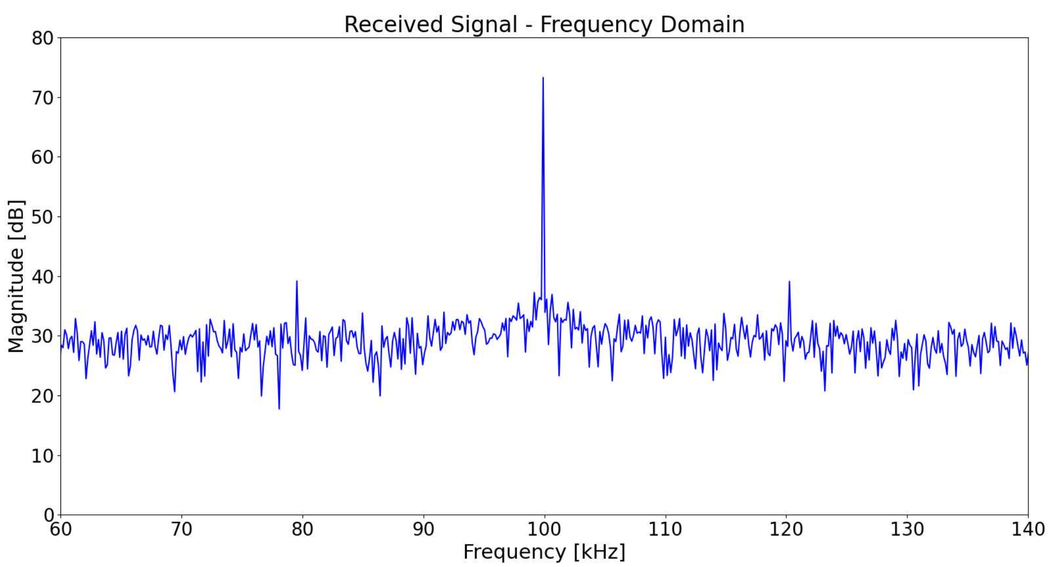

(b) Rough resolution frequency spectrum obtained using $f_s = 2400\text{kHz}$ and $N = 2^{14}$ samples.

Figure 4.5: Frequency spectrum output from the Analog Devices phaser board showing the proximity of DC and leakage noise to radar targets with varying sample rates, f_s. Antenna is facing a wall approximately three meters away.

fore necessary to filter these low frequencies out which is a more difficult task due to the close proximity of desired targets. Since no filter is capable of a zero-length transition band (shown in Figure 4.6), there will be some low frequencies filtered out corresponding to short ranges. In

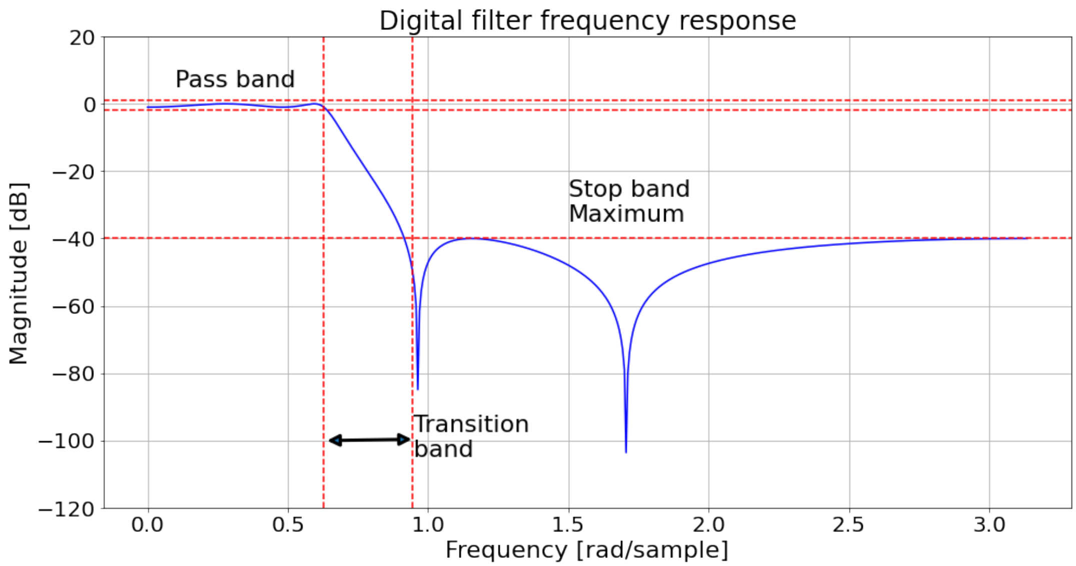

Figure 4.6: Frequency spectrum of a Chebyshev Type 1 digital low-pass filter with the pass, transition, and stop bands labeled.

an attempt to emulate the low power, low isolation radar incapable of shifting to an IF described previously, transmission was disabled, only allowing for leakage to be transmitted. This still gave interpretable results because frequency ramping was active. The setup is summarized in Table 4.3.

Table 4.3: Configured parameters of the Analog Devices FMCW phaser board for emulation of a radar transmitting very low power, detecting small targets, and the beat frequency at baseband.

Parameter	*Description*	*Value*	*Unit*
f_s	Sample rate	600	kHz
f_c	Center frequency	12.1	GHz
f_{IF}	Intermediate frequency (IF)	0	kHz
B_{RF}	RF bandwidth	500	MHz
N_{buffer}	Buffer size	2^{14}	-
T_{ramp}	Ramp time	1	ms

From the RF bandwidth, B_{RF}, and ramp time, T_{ramp}, in Table 4.3, the transmitted signal slope, S, can be found to be 500GHz/s. This, along with (2.2), gives the necessary tools to calculate how much noise filtering will affect the radar's ability to detect targets at short ranges. As shown in Figure 4.7, the DC interference occupies a bandwidth of approximately 300Hz which would correspond to a distance of 9cm from the radar and would, in most applications, be a negligible amount. $B_{RF} = 500$MHz is a relatively large bandwidth that some radar systems may be incapable

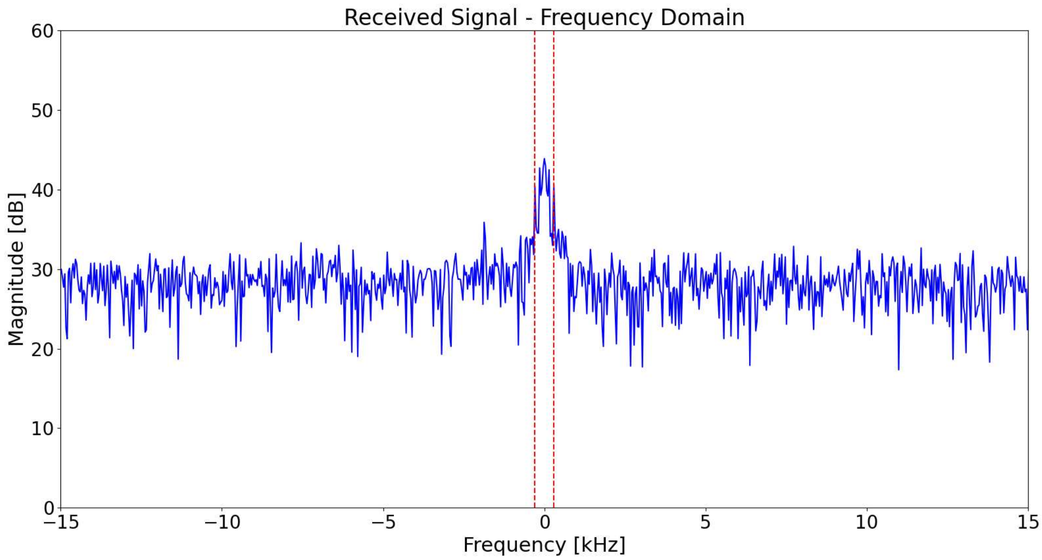

Figure 4.7: Narrowed frequency spectrum with the extent of the DC interference enclosed.

of generating or may degrade the antenna pattern which is a trade-off discussed in Section 3.2.4. The RF bandwidth for the radar proposed in Chapter 3 is a much more modest 3MHz. Lowering the bandwidth this much greatly affects the filtered low ranges which in this case, for the same DC interference bandwidth of 300Hz, would be 15m which is more problematic. The decreased bandwidth from 500MHz to 3MHz does greatly decrease the noise power at the receiver input by (3.5), however, which makes target detection easier and possibly reduces the amount the DC noise interferes with the signal. Unfortunately though, this will not help with the separate issue of antenna leakage. The signal processing will clearly be more difficult at these lower ranges with narrower signal-to-noise ratios but not impossible.

The DC interference bandwidth and power will not be constant, especially when the antenna is scanned across non-uniform environments which is discussed in more detail in Section 4.7. Because of this, when performing real-time signal processing it is favorable to use simpler filters

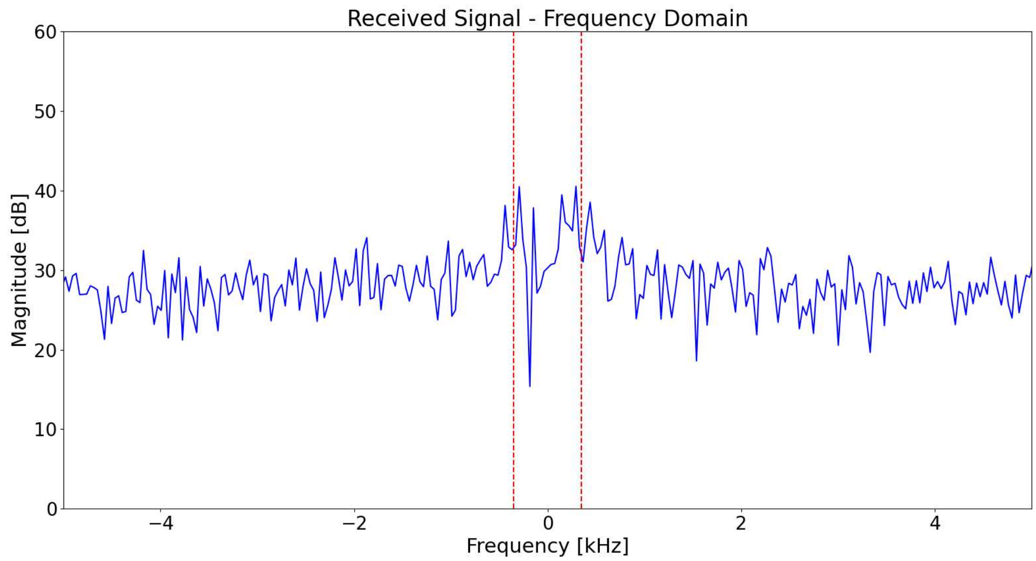

(a) Filter order = 2.

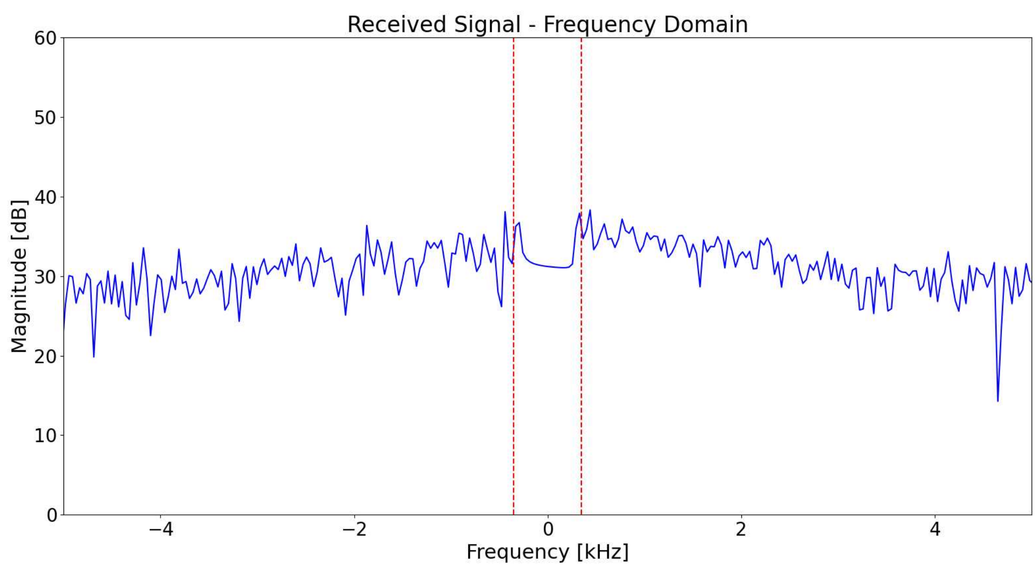

(b) Filter order = 10.

Figure 4.8: Comparison of a Type 1 Chebyshev high pass filter for removing leakage noise interference with $f_{\text{cutoff}} = 300\text{Hz}$, maximum permissible pass-band ripple of 5dB, and orders 2 and 10.

that take less time to compute. This worsens the previously mentioned issue of non-zero width transition bands and leads to either more of the desired signal or less of the interference being filtered. Figure 4.8 shows the trade-off between complexity, in this case filter order, and transition band width.

In designing a filter the *pass band ripple* is one of the parameters that can be seen in Figure 4.6 and it specifies how much the signal's magnitude deviates from unity gain. This is important because a poorly designed or simple filter with high ripple can reduce the SNR of targets which may already bordering the noise floor depending on their range, size, and material. Fortunately, many digital filter design tools exist which allow direct specification of pass band ripple. Another, arguably more important parameter is the *stop band suppression*, of which the maximum value is shown in Figure 4.6. This can have a very similar effect to a window's side-lobe level discussed in Section 4.2 in that if it is too high, parts of the stop band will interfere with the pass band causing false alarms. Again, the same tools can be used to directly specify the suppression, but it may come at the cost of computation time. This is of little or no importance for the majority of modern computers and even low-powered microcontrollers, especially if the filter weights are computed before real-time sampling and processing starts.

4.5.3 Maximum Range

With the availability of high-performance ADCs it is many times practical to use one with higher-than-needed sampling rate. This allows for a process known as oversampling which can be used to increase the SNR by sampling at a rate, $f_s >> f_{\text{Nyquist}}$, above the Nyquist rate and down-sampling to leave only the necessary frequencies. Depending on the radar's configuration (output power, bandwidth, etc.), there may still exist some high frequencies that will not realistically be used due to the high attenuation signals from targets at such a distance would experience. While processing may be very quick on modern computers, data visualization in real-time applications can be difficult because, to attain a quality resolution and frame-rate, the screen must update large amounts of data very quickly. It may, therefore, be useful to truncate this unused part of the high end of the frequency spectrum. This is less of a frequency filter because it is simply removing part of the spectrum that is not needed for display.

4.6 Coherent Change Detection

Coherent change detection (CCD) is useful if the only targets of interest are moving, as is the case in weather radars. CCD is similar to CBS described in Section 4.11 except instead of taking one baseline measurement before scanning for targets to obtain clutter presence, it subtracts every measurement's magnitude and phase from the subsequent one. This removes all stationary targets, effectively removing clutter and leaving only targets which have moved since the last measurement.

Since CCD subtracts magnitude and phase, it is important for maximum performance to operate in a coherent radar system, otherwise the phases of the chirp signal, $\phi_{\text{chirp},1}$ and $\phi_{\text{chirp},2}$, will not cancel as in (4.19) and will produce unexpected results. However, like CBS, CCD can be used in non-coherent systems by only subtracting the magnitude but is far less effective. Including phase information allows for detection of very minute changes in range which is very helpful when the sampling period is short relative to target velocity.

4.7 Detection Methods

4.7.1 Introduction

After the reflected signal is received, digitized, filtered, and transformed to the frequency domain, the processor still cannot differentiate between targets and noise. Even after filtering, the signal includes a large amount of noise from various sources that, depending on the signal-to-noise ratio, can be very difficult to distinguish from the targets of interest. Because of this, detection algorithms must be employed to decipher which range bins are desired.

4.7.2 Static Threshold

There are a wide variety of these algorithms that range in intricacy and effectiveness. At the lower end of the intricacy scale is a static threshold, shown in Figure 4.9, which is a constant value everything below which is considered noise and everything above a target. This static threshold method benefits from its simplicity by needing a small amount of "engineering hours" to implement as well as little computing costs which, in the case of many radars that use low power,

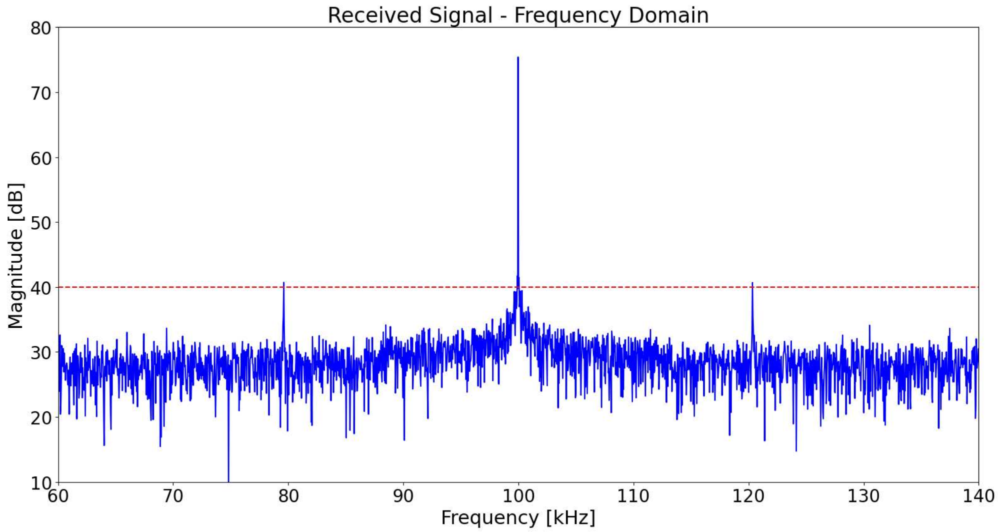

(a) Static threshold chosen for a single buffer in a known environment.

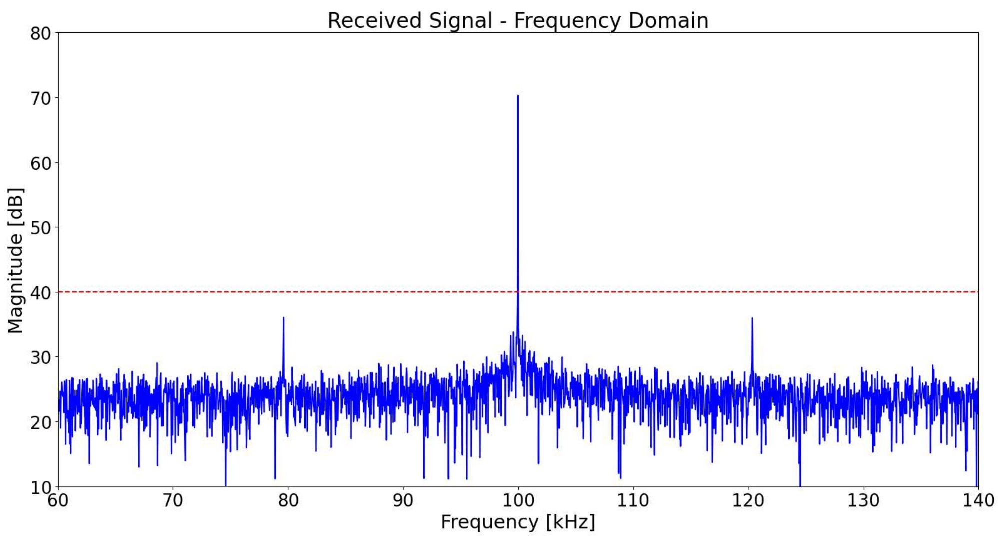

(b) Same static threshold applied after scanning to a location with less clutter and further, more attenuated targets.

Figure 4.9: Target detection example using a static threshold. This shows the inefficacy of static thresholds in dynamic environments.

embedded processors, is a limited resource. It does however come with the drawback of being a relatively ineffective tool at classification. This may not be obvious when viewing Figure 4.9a because it would seem that this threshold does a good job at removing everything below the signals of interest. This is, however, a special case in that it is a single buffer in a known environment which makes choosing a proper threshold simple. In real radar systems, and especially this one that is designed to be flown on a drone in unfamiliar environments, noise levels vary greatly depending on the distance from clutter (stationary) objects and the presence of other RF emitters nearby. Not only can the noise level be different between deployments but, in scanning radars, between angles of detection. This means that the static threshold chosen can be a good choice for one environment or scan angle but can then be completely invalidated when the radar is pointed towards some other objects as shown in Figure 4.9b. Figure 4.10 gives an example where the radar will receive a significant amount more power reflected from clutter at one scan angle than another which would result in differing noise floors and different threshold requirements. This issue justifies the need

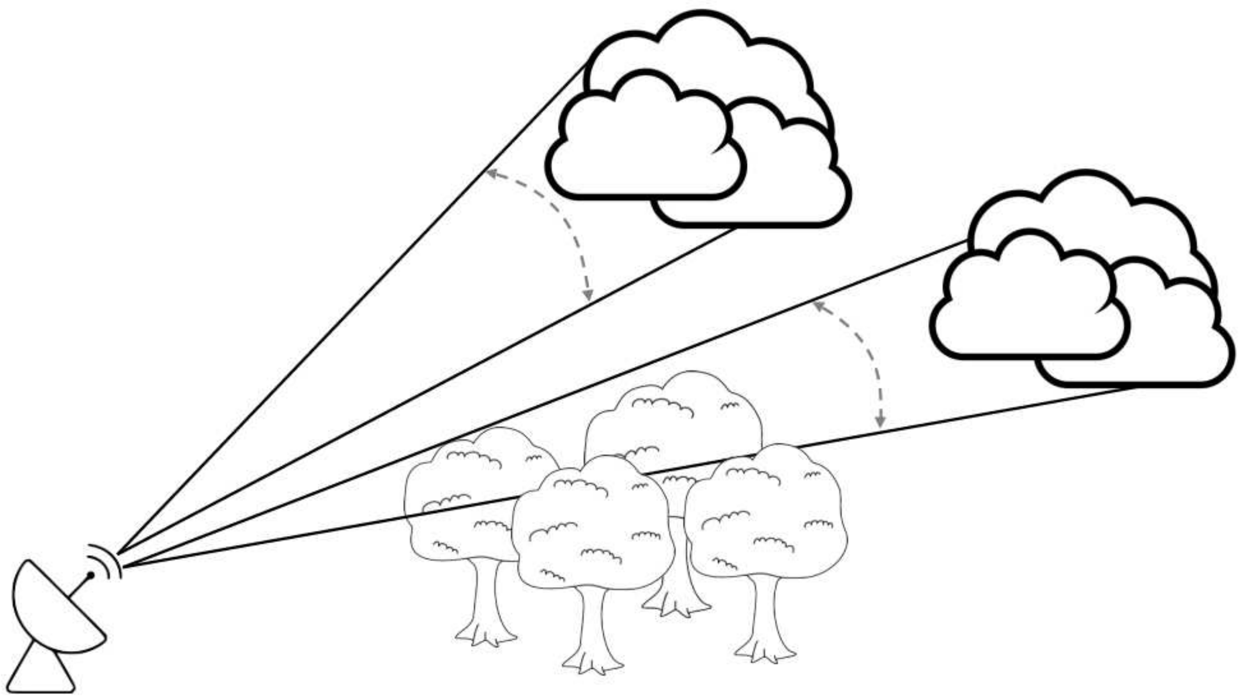

Figure 4.10: Illustration showing how changing the scan angle can greatly affect the amount of noise produced by ground clutter which is detected by the radar.

for a constantly varying threshold that is a function of the noise surrounding the range or volume being observed.

4.7.3 Constant False Alarm Rate

Constant false alarm rate (CFAR) is a step up in complexity and performance relative to the static threshold method which it achieves by addressing the issue of temporal and spatial noise variation. CFAR calculates a threshold at each point in time and for each cell in the range-FFT spectrum by taking the mean of the surrounding cells' powers and calculating a value at a point above the average power. Cells with a power above this calculated value, like in the static threshold method, are classified as targets.

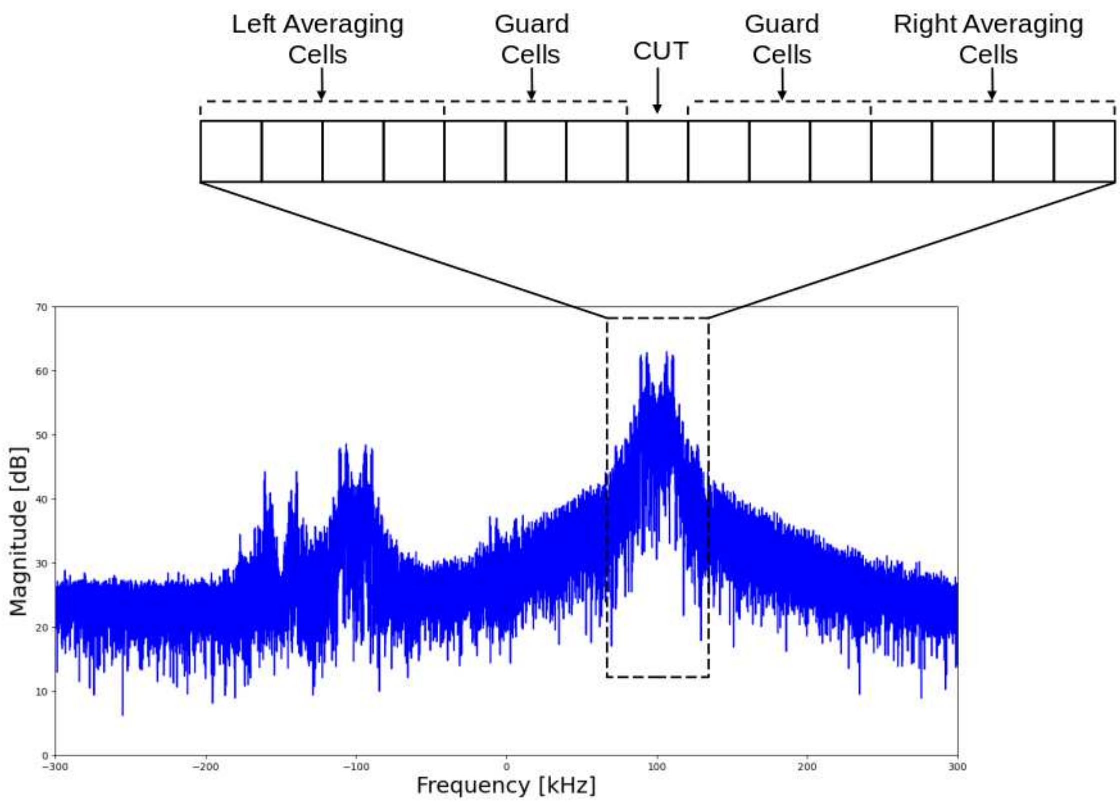

Figure 4.11: CFAR cells of interest visualized with an example of a received signal.

As mentioned above, the method is an iterative process that goes through the entire frequency spectrum and assigns a threshold to each of the cells. Since the process is repeated for every cell, CFAR can be described for one cell-under-test (CUT) and its amplitude A_i. To avoid interference from the CUT, a number of *guard cells*, N_{guard}, surrounding it are chosen to be omitted from the mean amplitude calculation. This serves the purpose of stopping leakage from the CUT into nearby cells from affecting the mean calculation which could artificially raise the mean power and block classification of the target. Outside the guard cells, the mean amplitude of some number of *reference cells*, N_{ref}, to the left and right of the CUT is calculated, giving the two quantities $A_{i,left}$ and $A_{i,right}$, respectively. Once the means are calculated, there are a variety of ways to determine which, if not both, of them will be used to classify targets [15]. The simplest of these methods are the aptly named *average*, *greatest*, and *smallest* methods which respectively take the average, greatest, and smallest of the two mean amplitudes. The threshold is then multiplied by a *bias*, C, which determines how far above the noise level a target must be to be classified as such. This effectively sets the minimum signal-to-noise ratio a cell must have to be considered a target. As a summary, these three detection thresholds can be described as

$$A_i > C \times \frac{1}{2}\left(A_{i,left} + A_{i,right}\right), \tag{4.6a}$$

$$A_i > C \times \max\left([A_{i,left}], [A_{i,right}]\right), \tag{4.6b}$$

$$A_i > C \times \min\left([A_{i,left}], [A_{i,right}]\right). \tag{4.6c}$$

These methods applied to the same frequency spectrum as in Figure 4.11 are shown below in Figure 4.12 with the design parameters in Table 4.4. The choice between the various CFAR meth-

Table 4.4: Design parameters used in CFAR examples. The parameters were chosen to fit a relatively large buffer, $N_{buffer} = 2^{14}$, taken in a close-proximity, clutter-contaminated environment.

Parameter	*Description*	*Value*	*Unit*
N_{guard}	Cells surrounding the CUT not included in mean.	10	Cells
N_{ref}	Cells surrounding the CUT included in mean.	30	Cells
C	Bias factor multiplied by mean.	3	-

ods depends on the environment and expected signal-to-noise ratio. For example, if the system's SNR is expected to be relatively low, the smallest method (Figure 4.12c) may be optimal because the threshold will be smaller than the other two, allowing for a finer line between noise and target. This does, however, increase the likelihood of misclassifications, or false alarms, which can lead to doing more processing than what is necessary as well as the obvious downside of misinterpreting data and giving misleading results. Figure 4.13 shows the same frequency spectrums as Figure 4.12

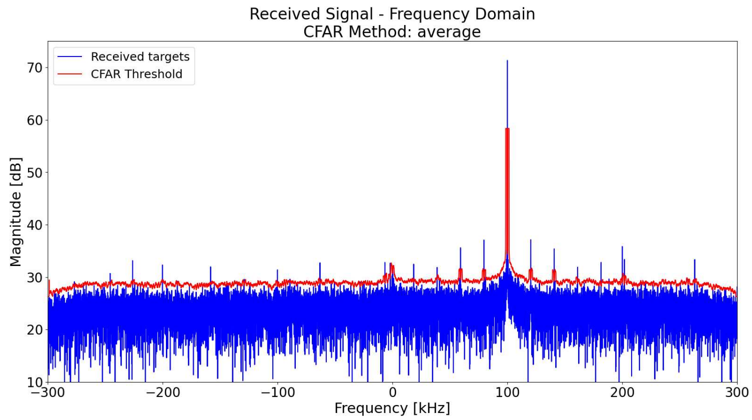

(a) Average method.

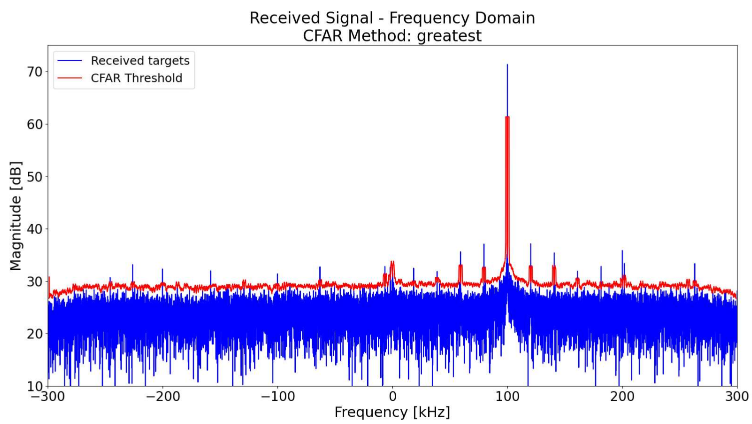

(b) Greatest method.

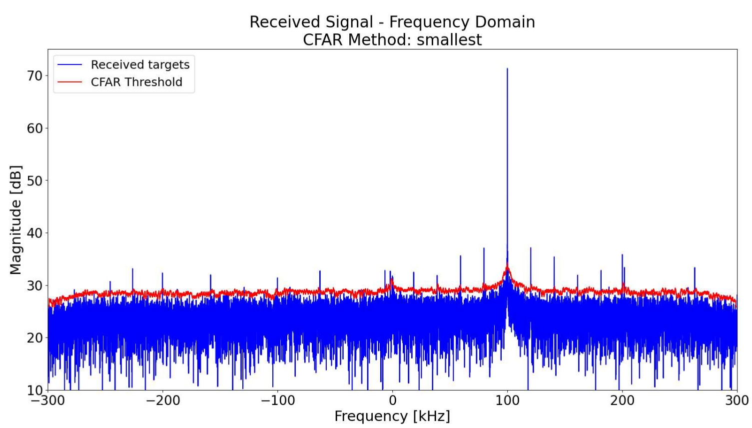

(c) Smallest method.

Figure 4.12: Comparison of three CFAR methods (average, greatest, and smallest) on the frequency spectrum of a received signal.

but with the cells below the CFAR threshold masked, leaving only the cells classified as targets. With the cells below the threshold masked, the difference between the techniques becomes more

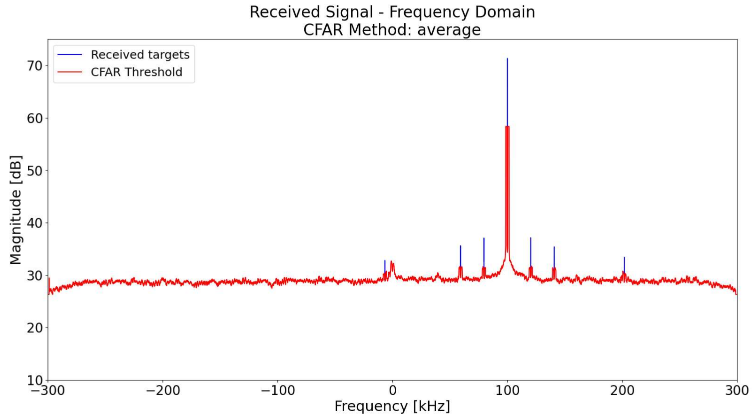

(a) Average method.

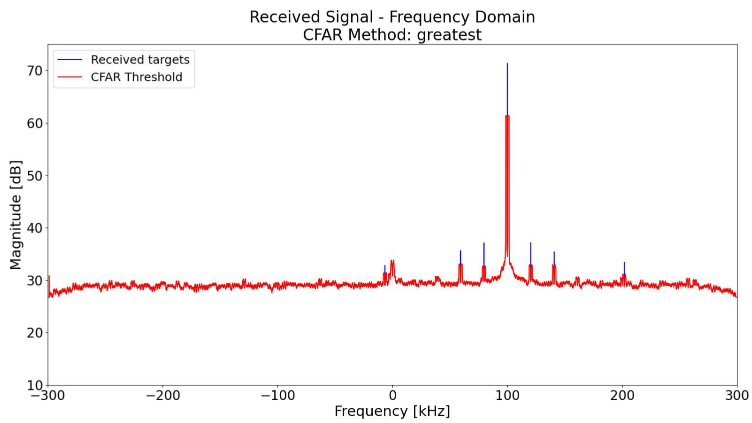

(b) Greatest method.

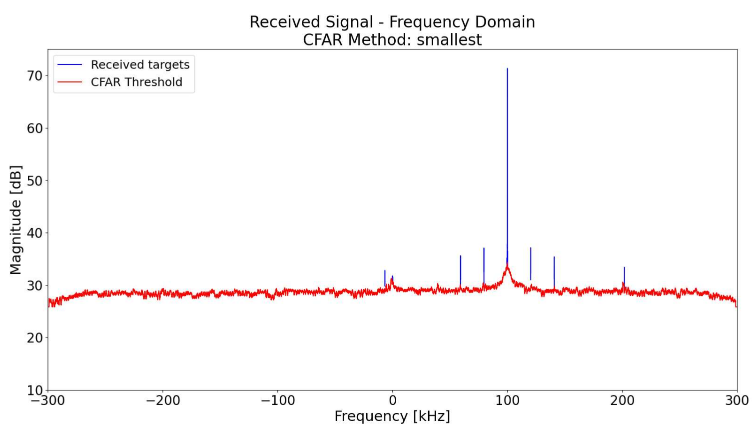

(c) Smallest method.

Figure 4.13: Comparison of three CFAR methods (average, greatest, and smallest) on the frequency spectrum of a received signal with all values below the threshold removed.

apparent. For this example signal, the greatest method (Figure 4.13b) outperforms the other two methods in terms of misclassifications. The CFAR method which will work best for different applications completely depends on the environment and equipment used. It will also greatly benefit the accuracy of classification if filtering is done prior to CFAR. To prove this, Table 4.5 shows results of measuring a reflective panel approximately 1 meter away from the antenna in a small room. This non-ideal environment will generate large amounts of reflections and stationary clutter. The filtering applied (Figure 4.15) was a relatively small bandwidth, $B = 20$kHz, band-pass

Table 4.5: Comparison of the effectiveness of three CFAR methods (average, greatest, smallest) with and without prior filtering measured by number of bins classified as targets. Antenna is pointed at a reflective panel approximately 1 meter away in a small room. The CFAR parameters used are those listed in Table 4.4.

Method	*With Filtering*	*Without Filtering*	$\lvert Actual - Classified \rvert$
Average	413	92	0
Greatest	375	114	22
Smallest	518	195	103

filter surrounding the intermediate frequency, $f_{IF} = 100$kHz, and a high-pass filter with a cutoff frequency, $f_{cutoff} = 100$kHz, both of which were 10^{th} order Chebyshev filters. The number of classifications only includes those with corresponding positive frequencies and does not include the image about the intermediate frequency. With a posteriori knowledge of the data it is possible to apply a filter custom-fit to the beat frequency received from the reflector panel giving very narrow and accurate results. Doing this obtains the optimal solution (92 cells) for what the CFAR algorithm would classify and is compared to the methods in this predictive filtering example in Table 4.5. The average method emerged as the best solution for this specific example with an error of 0%, with the greatest and smallest having respective errors of 23.9% and 112%. If run multiple times these could change due to dynamic environments and noise levels, but it is worth testing to get a baseline before aimlessly applying a detection method to a system.

One interesting behavior of the CFAR detection method can be seen surrounding the signal in Figure 4.12. The threshold gets significantly higher than surrounding areas because the target is captured in the reference cells, raising the mean. This helps in refining the classification by blocking out nearby cells that may have been influenced by the large amount of power reflected from the target. It also effectively decreases the resolution influenced by the RF bandwidth (2.3b)

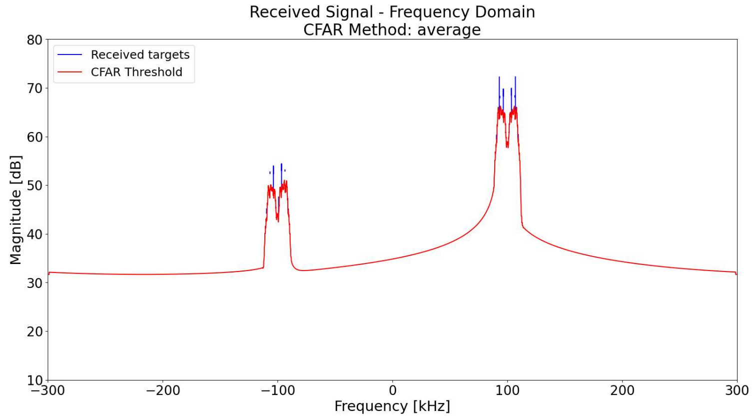

(a) Average method.

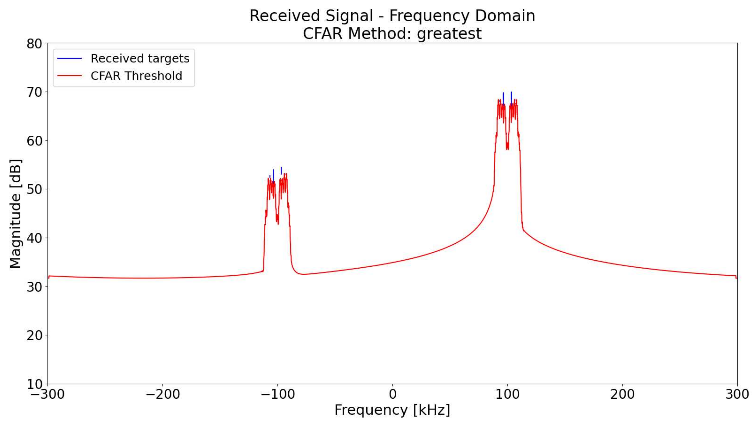

(b) Greatest method.

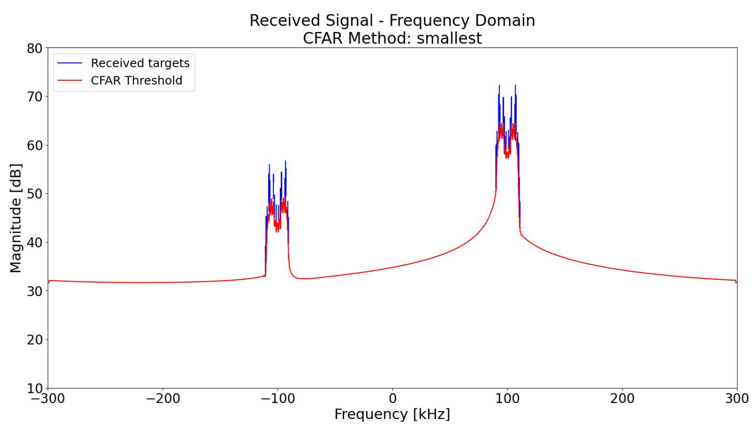

(c) Smallest method.

Figure 4.14: Comparison of three CFAR methods (average, greatest, and smallest) on the filtered frequency spectrum of a received signal with all values below the threshold removed. The filters used are a 10^{th} order band-pass ($90\text{kHz} \leq B_{\text{pass}} \leq 110\text{kHz}$) and high-pass ($f_{\text{cutoff}} = 100\text{kHz}$).

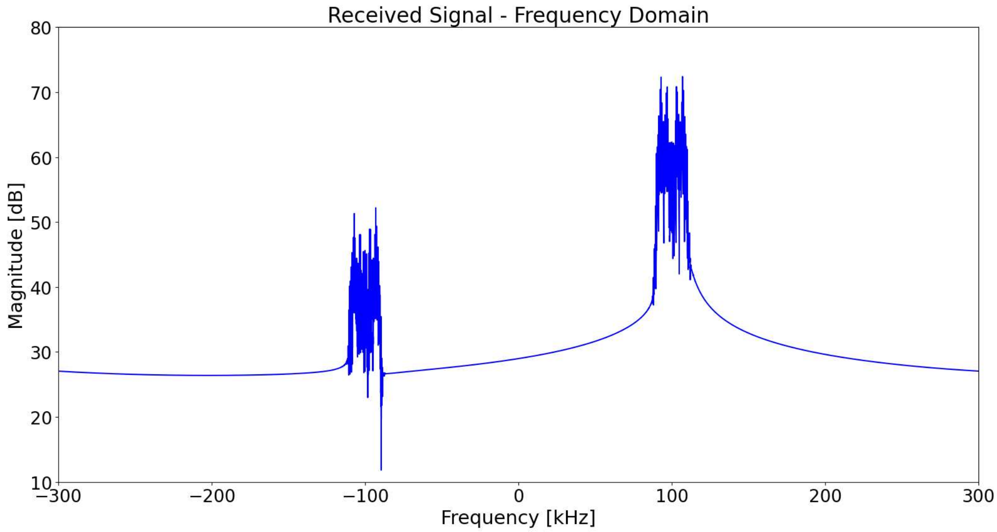

Figure 4.15: Frequency spectrum from the phaser board measuring a reflective panel approximately 1 meter away in a small room with a 10^{th} order band-pass (90kHz $\leq B_{\text{pass}} \leq$ 110kHz) and high-pass ($f_{\text{cutoff}} =$ 100kHz) filter applied.

and window used. To more strictly block the interference, the number of guard cells, N_{guard}, can be reduced. This will reduce the distance of the threshold peaks from the signal peaks, but it also has the effect of lowering the isolation from leakage from the CUT to nearby cells which can give artificially higher peaks.

Since the CFAR algorithm computes a separate threshold for each range bin, its speed is inversely proportional to the size of the buffer, N_{buffer}. For each bin, it must also access and average $2N_{ref}$ surrounding cells which is a detriment to speed. This effect on computation speed is an important consideration for real-time applications and is discussed further in Section 4.12.

Range normalization, discussed in Section 4.4, can make the process of target detection more difficult because the noise power is increased with range where the signal-to-noise ratio is already the lowest due to high path loss (3.1). Large bias values such as $C = 3$ used in Table 4.4 can completely mask distant targets. More measurements were taken with the same setup used to capture Figure 4.4b, this time as a comparison of the efficacy of varying biases, $C = 1.2, 1.5, 2, 3$, and is shown in Figure 4.16. The bias used previously, $C = 3$, which produced satisfying results clearly is not suitable for use after range normalization because it completely masks the reflector target. Though the metallic reflector is the only visible target before normalization, it is the increased noise

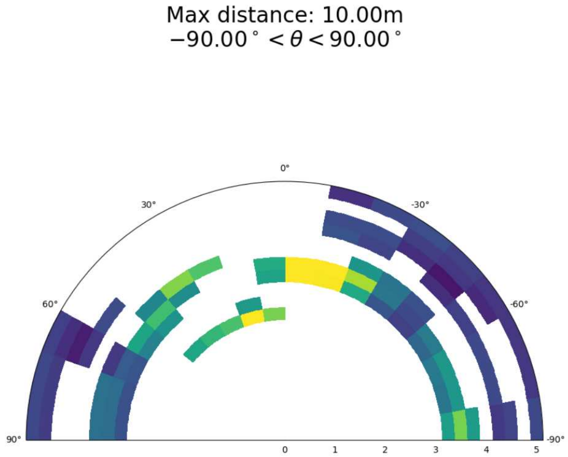

(a) $C = 1.2$.

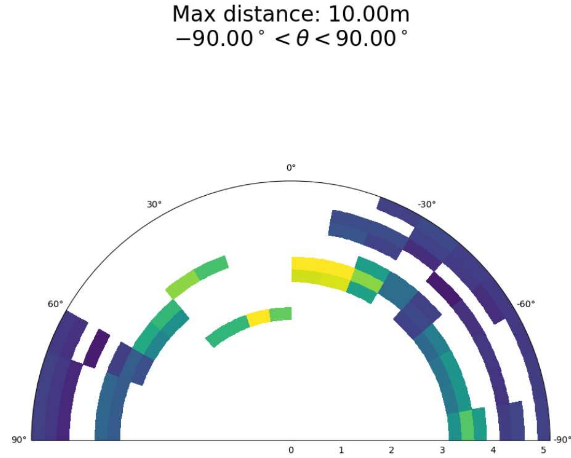

(b) $C = 1.5$.

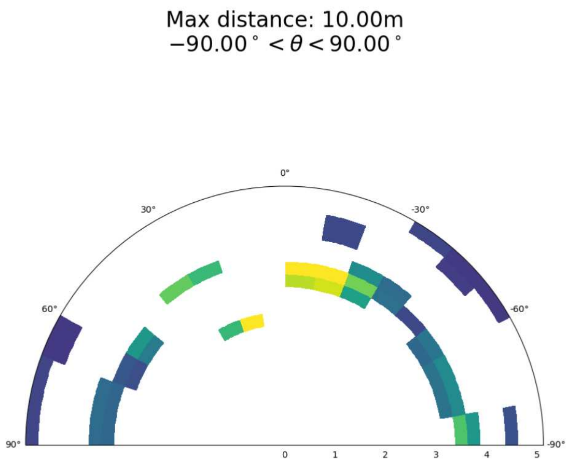

(c) $C = 2$.

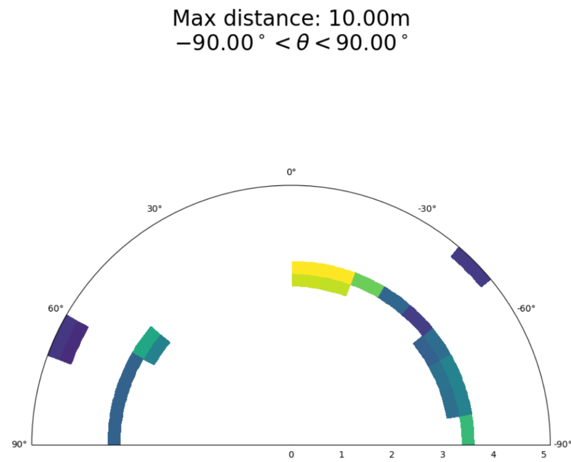

(d) $C = 3$.

Figure 4.16: PPI scans with range normalization ($RN = 1$m) and CFAR "greatest" method applied with varying biases ($C = 1.2, 1.5, 2, 3$). Range binning is also applied ($R_{res} = 3$cm) which necessitates a reduction in N_{guard} and N_{ref} to 0 and 3 bins, respectively. Two targets exist, a reflector panel at $R = 1.5$m, 0° off boresight and a human at $R = 2.5$m, 20° off boresight. The maximum round-trip range is $R_{max} = 10$m.

which brings the dynamic threshold higher than $C = 3$ times below its power, therefore causing it to be classified as noise. It should be noted that, assuming an equal noise power on either side of the target in range, the threshold past the target will have a disproportionately higher intensity than the noise before due to the normalization (4.3). Lowering to a bias, $C = 2$, allows the reflector target to emerge from below the threshold. Further lowering to the smallest biases, $C = 1.2, 1.5$, also shows both targets but allows an increasing amount of false alarms. They do, however, allow for more range depth to be shown for the targets. The choice is subjective to the application and expected targets, but a bias which completely disregards targets such as that in Figure 4.16d should be removed from consideration.

To avoid the need for testing multiple biases and methods, there is a method which allows for specification of a false alarm rate, $0 \leq P_{FA} \leq 1$, to obtain the threshold [6]. Implementation of this method first requires calculating the *background noise variance*,

$$\sigma_n^2 = \sum_{i=1}^{N_{ref}} \frac{A_i^2}{N} \tag{4.7}$$

where A_i are all N_{ref} reference cells defined previously. Since the variance is simply the square of the background noise's standard deviation, the standard deviation, σ_n, can be obtained and further used to calculate the target detection threshold,

$$T = \sigma_n \sqrt{-2 \ln P_{FA}}. \tag{4.8}$$

Identically to the previous methods, the cell is then classified as a target if it satisfies the condition, $A_i > T$. If a cell is a target, its signal-to-noise ratio can be calculated by,

$$SNR = \frac{\sigma_s^2}{\sigma_n^2} = \frac{A_i^2}{\sigma_n^2}. \tag{4.9}$$

This method may be the most desirable because of its minimization of changes needed when switching environments.

4.8 Single-Polarization Radar Products

4.8.1 Introduction

Previous sections of this chapter have primarily focused on the "cleansing" of the signal with the purpose of making it usable for further processing and interpretation. This section provides a discussion of meteorological radar products, some of the most common and useful interpretations of the received data. Meteorological radar products have been invaluable resources for classifying weather events since the mid-twentieth century. They help understand and predict storm behaviour which is beneficial for convenience and safety. Even with just a single polarization, very useful products can be derived from a scattered signal, of which two will be discussed in this section, reflectivity and Doppler velocity.

4.8.2 Reflectivity

Reflectivity is the most common of the radar products and is what most think of when they imagine a radar display. It is used to show the radar cross section (RCS or σ_b) of all the objects within a range bin which *scatter* the signal, otherwise known as scatterers. In this way it is very similar to reflected power which is directly proportional to the RCS. This section will provide a brief explanation of reflectivity, but a more in-depth analysis can be found in [8].

Reflectivity is, by default, range normalized, a process described in Section 4.4. It is described by the equation,

$$Z_e = \frac{\lambda^4}{\pi \left|K\right|^2}\eta \tag{4.10a}$$

$$\text{where } |K|^2 = \left|\frac{\varepsilon_2 - 1}{\varepsilon + 2}\right|^2 \tag{4.10b}$$

is the dielectric factor which often uses the permitivitty for water, $\varepsilon_r \approx 80$, since the weather radars are mostly interested in precipitation. η is the RCS per-unit-volume which in a volume, ΔV, can be described by,

$$\eta(R, \theta, \phi)\Delta V = \sum_{k}^{N} \langle \sigma_b \rangle \tag{4.11}$$

where $\langle \sigma_b \rangle$ is the expected value of the RCS for the k^{th} scatterer in the volume. It is intuitive that the RCS is a function of the scatterer's size in the direction of the wave's polarization, but less obviously it also depends on the scatterer's dielectric factor and can be summarized by,

$$\sigma_b = 4\pi|S|^2 \tag{4.12a}$$

$$= \frac{4\pi}{k_0^2}\,|K|\,(k_0 a)^6 \tag{4.12b}$$

$$= \frac{\pi^5}{\lambda^4}|K|^2 D^6 \tag{4.12c}$$

where S is the scattering matrix, $k_0 = 2\pi/\lambda$ is the wave number, a is the radius of a spherical scatterer, and $D = k_0 a$ is the diameter of the sphere. This is useful for simulation to show the magnitude of the expected reflectivity for a given environment with N scatterers, but since the distribution of scatterers in a range bin may be unknown, another method must be used to obtain the reflectivity from a received signal.

For pulsed radars an alternate method for this computation uses the fact that mean received power can be related to the reflectivity by,

$$\bar{P}_r = \underbrace{\left(\frac{cT_0}{2}\right)}_{\substack{\text{Range}\\\text{resolution}}} \underbrace{\left[\frac{\lambda^2 P_t G_0^2}{(4\pi)^3}\right]}_{\substack{\text{Power and}\\\text{gain}}} \underbrace{\left[\frac{\pi\theta_1\phi_1}{8\ln(2)}\right]}_{\substack{\text{Narrow beam}\\\text{approximation}}} \underbrace{\left(\frac{\pi^5|K|^2}{\lambda^4 R^2}\right)}_{\text{Path loss}} Z_e(R) \tag{4.13}$$

where R is the range and T_0 is the pulse repetition interval for a pulsed radar. G_{rx}, the total gain of the receiver and signal processing path, can be applied to (4.13) to achieve the received power at the output of the signal processing path, P_o. This equation bypasses a number of steps because they are explained in detail in [8] and the derivation is unnecessary to describe in this thesis. The *narrow beam approximation* is an approximation for the volume enclosed in the beamwidths, θ_1 and ϕ_1, proposed in [16] which is useful for determining which scatterers reside within the radar's scan. Before (4.13) can be used for this radar, some of the parameters must be modified to suit an FMCW radar. The range resolution in (4.13) cannot be used directly for FMCW radar and instead should be substituted with (2.3b). The power and gain portion must also be slightly modified because this radar uses separate antennas for transmit and receive giving separate gains, G_t and G_r, respectively. After revision for FMCW radar, the equation for the power at the output of the signal processing path becomes,

$$P_o = G_{rx} \underbrace{\left(\frac{c}{2B_{RF}}\right)}_{\substack{\text{Range}\\ \text{resolution}}} \underbrace{\left[\frac{\lambda^2 P_t G_t G_r}{(4\pi)^3}\right]}_{\substack{\text{Power and}\\ \text{gain}}} \underbrace{\left[\frac{\pi\theta_1\phi_1}{8\ln(2)}\right]}_{\substack{\text{Narrow beam}\\ \text{approximation}}} \underbrace{\left(\frac{\pi^5|K|^2}{\lambda^4 R^2}\right)}_{\text{Path loss}} Z_e(R). \tag{4.14}$$

Many of these terms are set at the beginning of operation and remain constant throughout which enables them to be grouped into a constant variable,

$$C = \frac{G_{rx}P_t G_t G_r \pi^3 |K|^2 c\theta_1\phi_1}{\lambda^2 1024 B_{RF} \ln 2}. \tag{4.15}$$

Rearranging (4.14) and substituting in the radar constant (4.15) gives an equation for reflectivity all in terms of variables attainable without knowledge of the scatters within the scan volume,

$$Z_e(R) = \frac{P_o R^2}{C}. \tag{4.16}$$

From (4.16) the reflectivity from an amount of precipitation can be simulated as a function of received power which is in turn a function of range.

Volumetric resolution is proportional to range by (3.8), therefore, as distance from the radar increases, the number of scatterers able to fit inside the range bin will increase. It is useful then to determine how a radar's reflectivity will vary with range. This relationship for DARMA is shown in Figure 4.17a. A reflectivity of 100dBZ is relatively high (many radar displays use limits similar to $-20 \leq Z_e \leq 70$) which can mostly be attributed to the beamwidths, $\theta_1 = 10°$ and $\phi_1 = 20°$. Most long-distance radars use high-performance antennas with fine resolutions (e.g., the CHILL radar in Figure 4.17b which uses a parabolic reflector antenna with beamwidths of $\theta_1 = \phi_1 = 1.1°$ [17]) which makes the amount of precipitation in each range bin smaller. The range resolution can also be decreased but $R_{res} = 50$m is fairly standard.

It should be noted that the units of reflectivity are mm^{6}m^{-3} therefore, if base units (m, W, Hz, etc.) are used in (4.16), the calculation will be off by a factor of 10^{18}. This can simply be multiplied

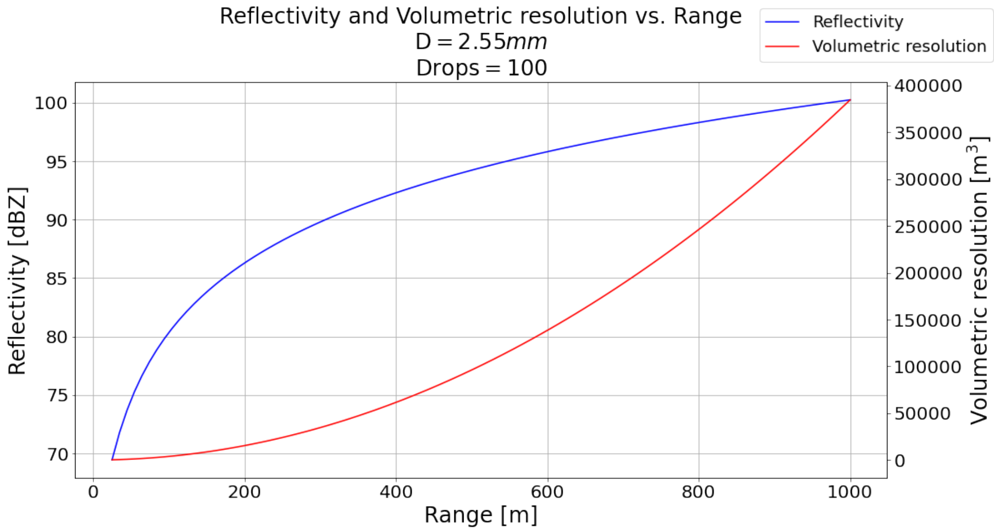

(a) Simulation using the DARMA radar's parameters over the range $0.025 \leq R \leq 1$km since $R_{max} = 1$km. This radar has a range resolution of $R_{res} = 50$m and horizontal and vertical beamwidths of $\theta_1 = 10°$ and $\phi_1 = 20°$, respectively.

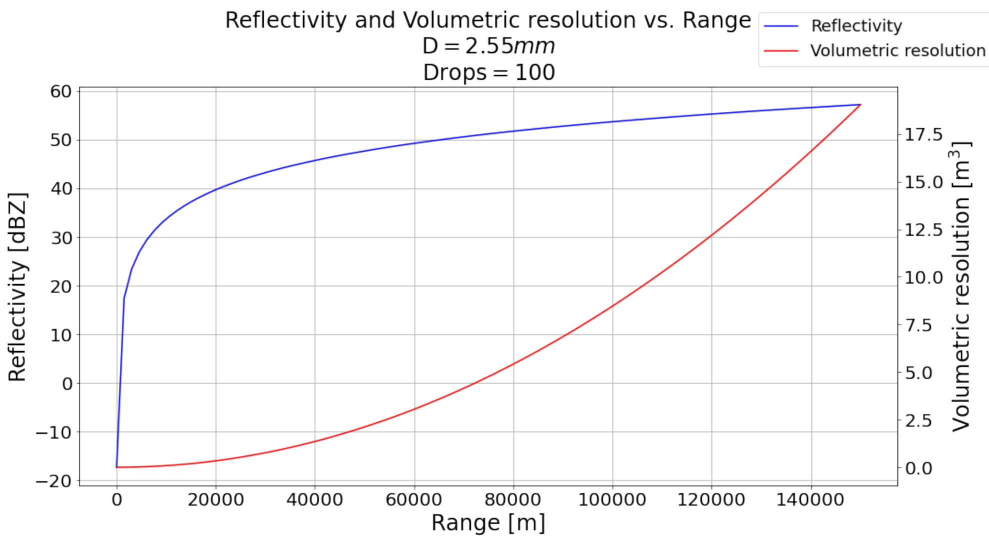

(b) Simulation using the CHILL radar's parameters over the range $0.025 \leq R \leq 150$km since $R_{max} = 150$km. CHILL has a range resolution of $R_{res} = 45$m and identical beamwidths in the horizontal and vertical directions of $\theta_1 = \phi_1 = 1.1°$.

Figure 4.17: Simulation of reflectivity and volumetric resolution as a function of range with the radars (DARMA and CHILL) measuring 100 rain drops of diameter, $D = 2.55$mm. Both simulations are plotted to the respective radars' maximum ranges, 1km and 150km.

by $Z_e(R)$ to obtain the expected value,

$$Z_e(R) \to 10^{18} Z_e(R). \tag{4.17}$$

4.8.3 Doppler Velocity

Doppler velocity is another very important radar product because it allows the viewer to understand the behaviour of the measured targets; for weather radar this can help predict the path and intensity of a storm which is very useful information. As mentioned in Chapter 2, unlike in pulsed radars, the Doppler shifted frequency, f_d, cannot be directly obtained from the frequency spectrum because it is ambiguous whether the shift results from non-zero range to the target or velocity. An alternative for finding velocity in FMCW radar is to collect two subsequent buffers between which the phase difference can be calculated. This phase difference can show minute changes in range between samples which in turn can be used with the time between samples, T_{diff}, to calculate the range delta over time, or velocity.

Phase difference calculation is allowed by *coherent* FMCW radar [18] which adds the ability to transmit every pulse not only at a controlled chirp frequency, f_{chirp}, but at a constant phase, ϕ_{chirp}. The phase difference is therefore,

$$\Delta\phi = (\phi_1 + \phi_{\text{chirp}}) - (\phi_0 + \phi_{\text{chirp}}) = \phi_1 - \phi_0 \tag{4.18}$$

where ϕ_0 and ϕ_1 are the phases at $t = 0, 1$, respectively. Without the constant chirp phase, the phase difference would be

$$\Delta\phi = (\phi_1 + \phi_{\text{chirp},1}) - (\phi_0 + \phi_{\text{chirp},0}) \tag{4.19}$$

where $\phi_{\text{chirp},1} \neq \phi_{\text{chirp},0}$, leaving two unknown phases and making the phase difference incalculable. For this reason, it should be noted that the following methods will not work predictably in a non-coherent system and all results should be assumed to have been obtained from a coherent system.

For a meaningful velocity calculation the radar's antenna was pointed directly facing a reflector panel moving with constant velocity of approximately $v_{\text{reflector}} = 0.3\text{m s}^{-1}$ toward the antenna. A large buffer ($N = 2^{14}$ samples) was taken at a sample rate of $T_s = 600\text{kHz}$ and split into two subsequent buffers, $N_1 = N_2 = N/2$, with a time difference between the buffers of $T_{\text{diff}} =$

13.65ms. The two subsequent buffers were filtered to only include the range at which the reflector resides and are shown overlaid in Figure 4.18. The two signals have very similar temporal behavior

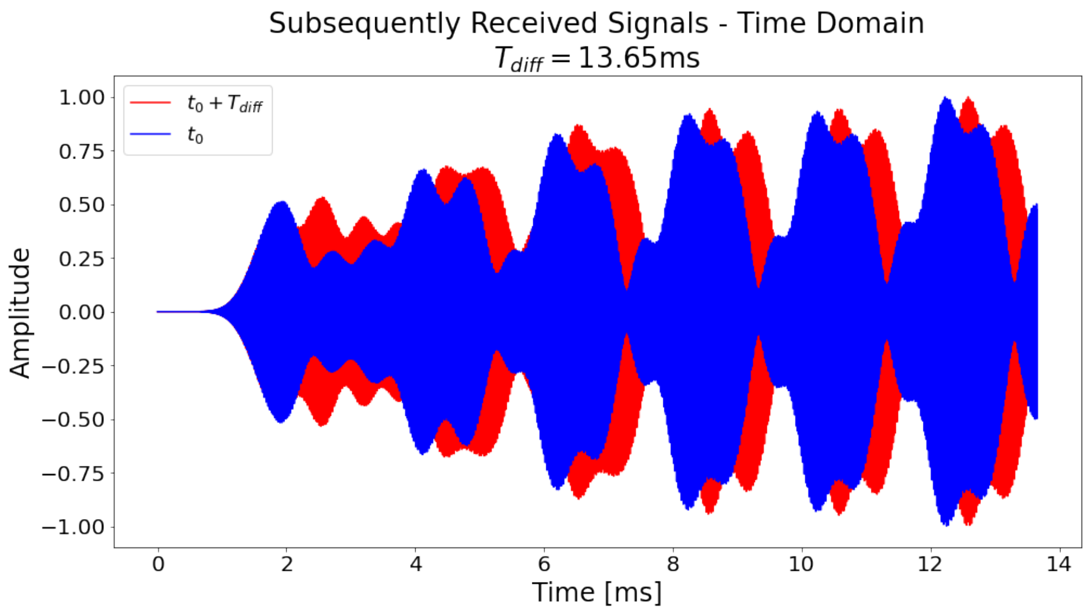

Figure 4.18: Two subsequently received signals with a time difference between buffers of $T_{\text{diff}} = 13.65$ms and amplitudes normalized. Antenna is pointed at a reflector panel moving with constant velocity of approximately $v_{\text{reflector}} = 0.3\text{m s}^{-1}$ toward the antenna. The received signal is filtered to only contain one range bin.

except for the buffer at $t = t_0 + T_{\text{diff}}$ having a slightly larger amplitude and a positive phase shift. The amplitude difference can be attributed to the fact that the reflector is moving towards the antenna therefore making the signal travel a shorter distance which results in lower attenuation (3.1). To account for errors in the velocity calculation this may incur, both signals are normalized. The other component, the phase difference, needs to be obtained for calculation of the velocity relative to the antenna. The cross-correlation is used for this calculation.

The cross-correlation is a function used to determine how similar two signals are, not only at the current time but also at various time delays, called *lags*. The cross-correlation between two signals, a and b, is defined by the equation,

$$R_{ab}[\tau] = \sum_{n}^{N} a[n]b^*[n-\tau] \tag{4.20}$$

where τ is the time lag and "*" denotes the conjugate. Figure 4.19 gives an example of a sine wave at lags, $\tau = 0, 1$. Computing the cross correlation is effectively multiplying the two signals,

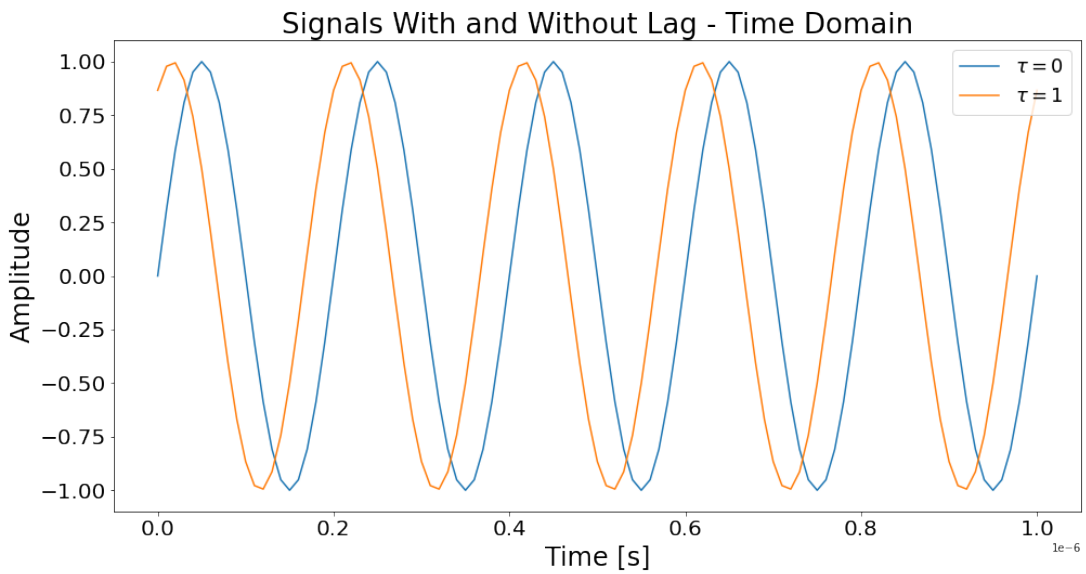

Figure 4.19: Time-domain representation of one signal at lags, $\tau = 0, 1$.

a and b, together with different shifts in time which results in varying amounts of constructive or destructive interference. If a and b are the same signal phase shifted by $\phi = \phi_0$, at some lag, τ_0, the two signals, $a[n]$ and $b[n - \tau_0]$, will be identical and the cross-correlation maximized. This lag, τ_0, is the time corresponding to the phase difference, ϕ_0, between the two signals. Since the same amount of interference will occur at lags corresponding to phase shifts, $\phi = \phi_1 + 2k\pi$ for $k = 0, 1, 2, \ldots, \infty$, the cross-correlation need only be calculated over the range, $-\pi \leq \phi \leq \pi$. From (4.20) it can be seen that the output will be a function of the lag and have length $2N-1$ which maps to a time axis in the range of $-NT_s \leq t \leq NT_s$. Figure 4.20 shows the cross-correlation of two signals with a respective phase shift of $\phi = \pi/3$. Without the known phase shift it still

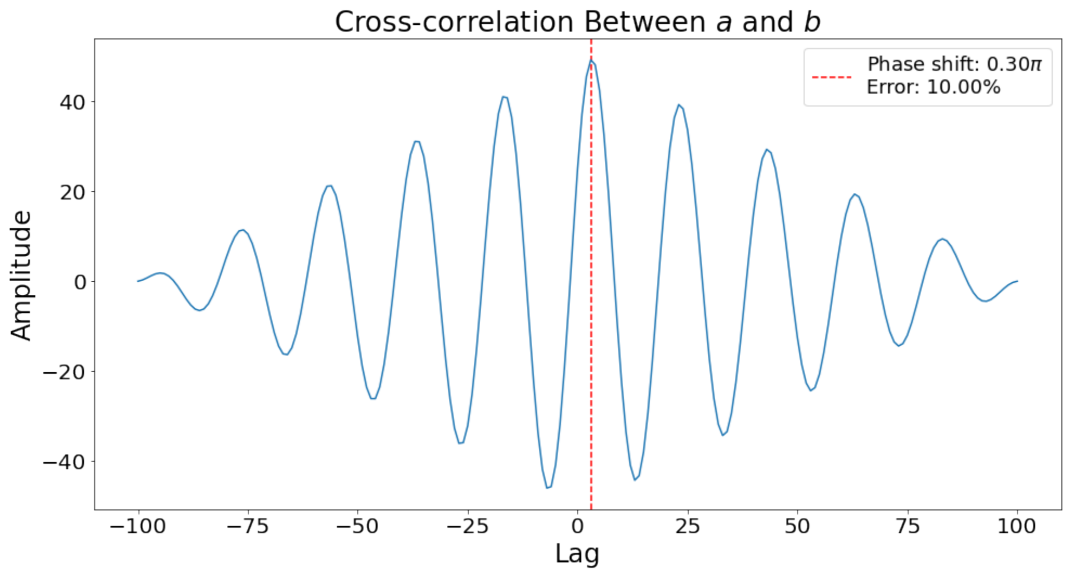

Figure 4.20: Cross-correlation of two identical sine waves with a phase shift of $\phi = \pi/3$. The maximum correlation is labeled with the corresponding calculated phase shift and error relative to the actual phase shift.

can be seen that the maximum correlation occurs at some non-zero lag which indicates the signals have some phase shift with respect to one another. Since the time range is known from the buffer size and sampling rate, the lag can be converted to a time shift and further used to obtain a phase shift of 0.3π giving an error of 9.09%. This should be accurate enough to distinguish targets from stationary clutter using the phase shift to Doppler velocity relation in (2.6).

Applying the cross-correlation to the two subsequent buffers taken in Figure 4.18 yields Figure 4.21. The maximum correlation occurs at lag, $\tau_0 = -232$, or a phase shift, $\phi_0 = 0.27\pi$. Following (2.6) where $\Delta\phi = \phi_0$, the Doppler velocity is found to be $v_r = 0.12\text{m s}^{-1}$ which, relative to the approximated reflector velocity, $v_{\text{reflector}} = 0.3\text{m s}^{-1}$, gives 60% error. There are multiple levels of uncertainty in this calculation.

With the process for calculating the velocity of one range bin defined, the same process can be used for the rest of the range bins to obtain the range-velocity spectrum. This process, however, can be made inaccurate when phase errors are introduced. Mitigating phase errors, especially in

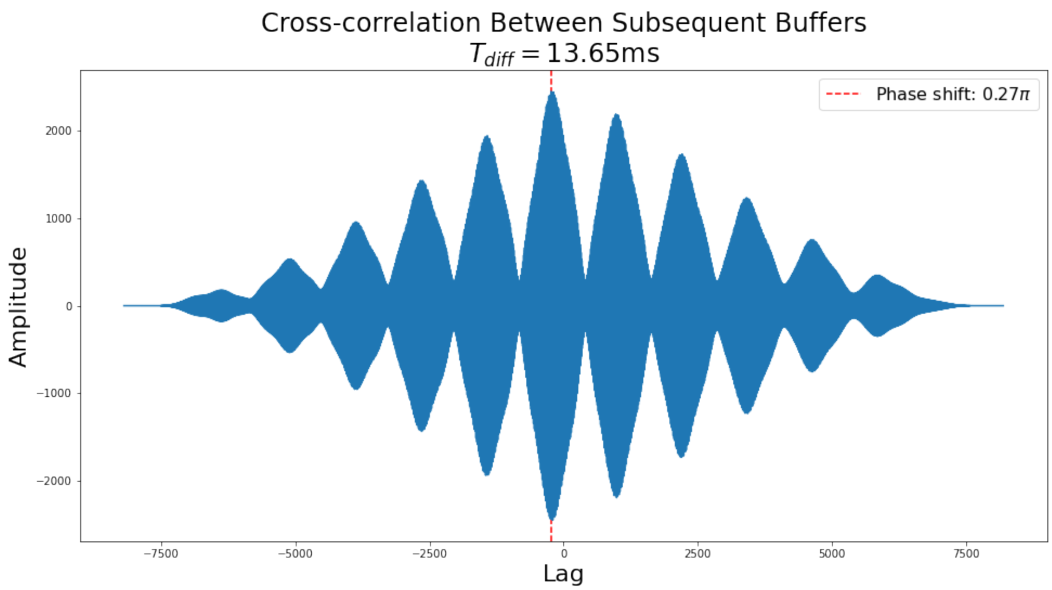

Figure 4.21: Cross-correlation of two subsequently received signals with a time difference between buffers of $T_{\text{diff}} = 13.65$ms. Antenna is pointed at a reflector panel moving with constant velocity of approximately $v_{\text{reflector}} = 0.3$m s^{-1} toward the antenna.

a low-cost system such as that in Chapter 3 and the phaser radar, can be very difficult. A useful alternative is introduced in Section 4.10.

4.9 Dual-Polarization Radar Products

4.9.1 Motivation

Moving from reception in a single polarization, most commonly horizontal, to two polarizations greatly increases the amount of information that can be obtained from a volume which is very useful for classifying weather events and predicting their future behavior [8]. This extra information comes from the introduction of new radar products, of which some of the most common are differential reflectivity (Z_{dr}), differential phase (Φ_{DP}), and cross correlation ratio (ρ_{HV}). These base products can then be used to derive others such as specific differential phase (K_{dp}). There are multiple others but, as these are the most commonly used, the focus will be on them.

4.9.2 Differential Reflectivity

Differential reflectivity is a very useful product for determining the distribution of scatterers' cross-sections (σ_b) in the horizontal and vertical planes. It achieves this by calculating the reflectivity received by each polarization separately and computing the ratio in logarithm form of horizontal to vertical, or

$$Z_{dr} = 10\log_{10}\left(\frac{Z_H}{Z_V}\right). \tag{4.21}$$

This means if a scatterer has a circular cross-section, $Z_{dr} = 0$dB and if its cross-sectional width is larger or smaller than the height, Z_{dr} will be positive or negative, respectively, which is shown in Figure 4.22. Knowing the ratio of horizontal to vertical reflectivity can be used to determine

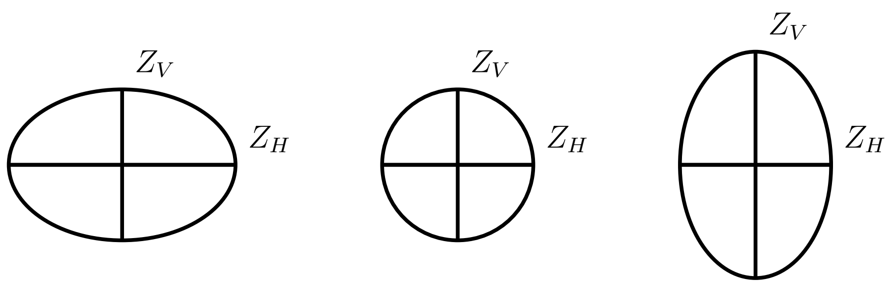

$Z_{dr} > 0$dB $\quad$ $Z_{dr} = 0$dB $\quad$ $Z_{dr} < 0$dB

Figure 4.22: Visualization of the impact of scatterer cross-section dimensions on differential reflectivity in logarithm form.

raindrop size and shape [19] and detect updrafts. When plotting Z_{dr} the limits are commonly set to $-1 \leq Z_{dr} \leq 8$ since the horizontally-polarized signal is often larger. This is due to the effect where as a scatterer is increasingly composed of liquid, air resistance on the bottom of the scatterer will cause its base to widen, further increasing the reflectivity measured in the horizontal polarization. Typical values for Z_{dr} and the effect "wetness" has on the measurement is shown in Figure 4.23 [20]. Given this, not only can Z_{dr} be used to extrapolate information from a dual-polarized signal about the shape of scatterers, but also their composition.

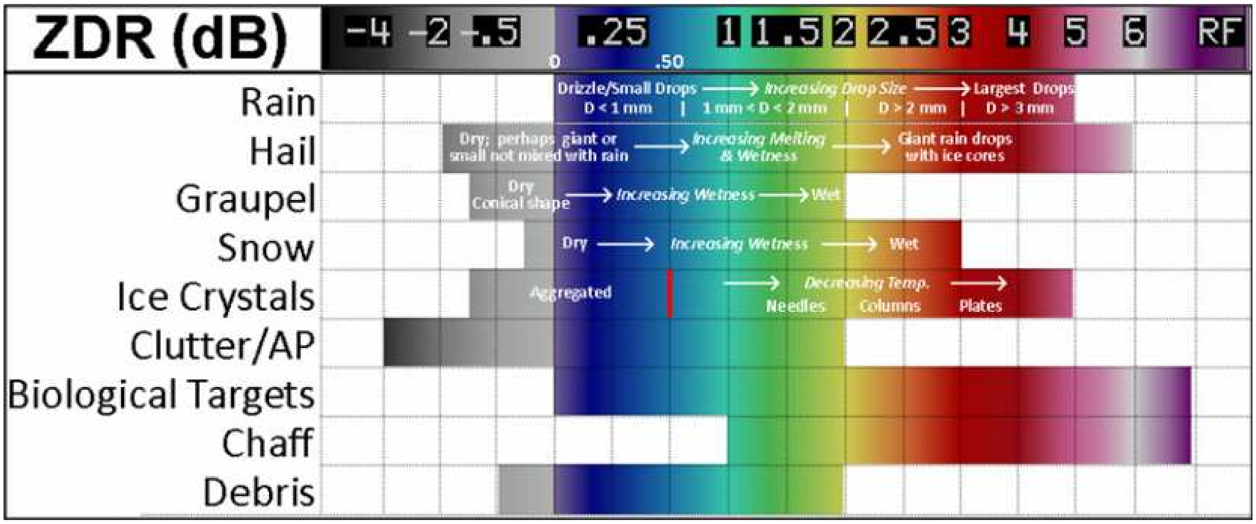

Figure 4.23: Typical values for differential reflectivity and the precipitations events with which they correspond.

4.9.3 Differential Phase and Specific Differential Phase

Electromagnetic waves are slowed when passing through a non-vacuum medium ($\varepsilon_r > 1$) which effectively adds a phase shift to the signal which can be measured by the receiver. Subtracting the phase difference seen by the horizontal (Φ_H) and vertical (Φ_V) receivers results in the differential phase, Φ_{DP}. Φ_{DP} is related to the horizontal and vertical axes of the scatterer's cross-section in the same way Z_{dr} is. There is, however, the addition of sensitivity to the distribution of scatterers within a volume since the signal is phase shifted with every medium it passes through. Differential phase is limited because it is an accumulation of phase shifts incurred by every scatterer within the signal's path. There is therefore a need for a radar product capable of relaying phase information about scatterers within a single range bin. This derived product is known as specific differential phase, K_{dp}, which takes the range derivative of Φ_{DP} giving the phase change over range instead of a phase accumulation. This makes the measurements specific to each range bin and gives units of degrees per kilometer (deg/km). Typically K_{dp} is shown in the range of $-2° \leq K_{dp} \leq 5°$ and Figure 4.24 shows what these common values correspond to for various meteorological events [20]. Non-meteorological events contribute large amounts of noise and make the K_{dp} measurement very difficult to read which is why they are not included.

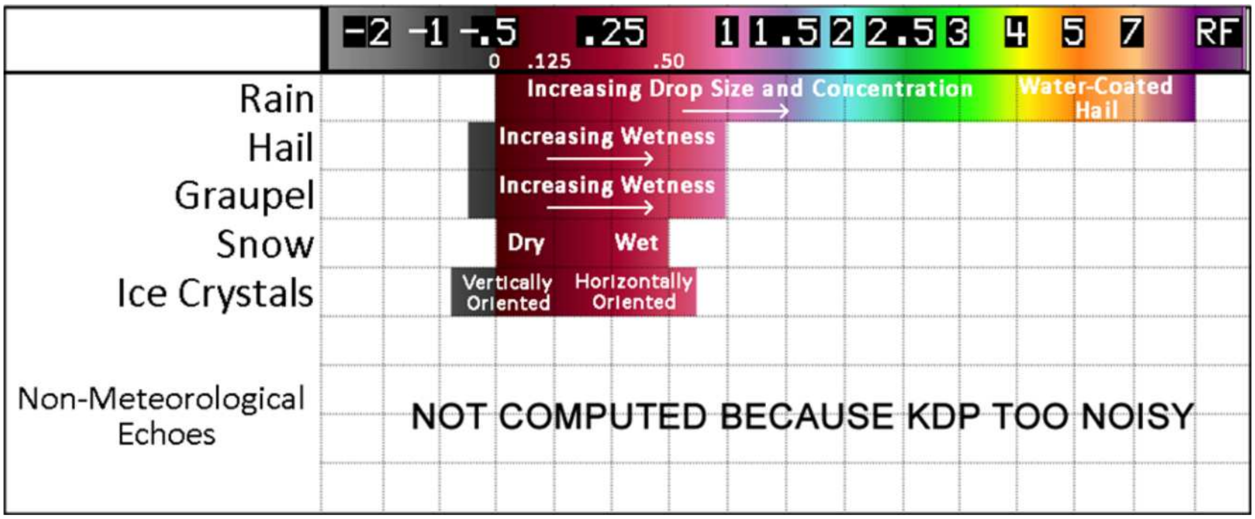

Figure 4.24: Typical values for specific differential phase and the precipitations events with which they correspond. Non-meteorological events are not described due to their large contribution of noise.

4.9.4 Cross-Correlation Ratio

Cross-correlation ratio, ρ_{HV}, is the correlation between the horizontal and vertical polarizations from pulse-to-pulse. This is used to determine the degree of uniformity of targets within a volume which is useful information because it can show whether there exists non-meteorological scatterers (low correlation, $\rho_{HV} < 0.8$) or if there are multiple types of precipitation (medium correlation, $0.8 < \rho_{HV} < 0.97$). Discrimination between meteorological ($\rho_{HV} > 0.8$) and non-meteorological ($\rho_{HV} < 0.8$) scatterers is useful in the minimization of interference from birds, planes, etc. in classifying a storm. Multiple types of precipitation being present in one volume could be an indicator of snow or ice melting to form rain at a certain altitude, otherwise known as the melting layer. It can also be used to show where different types of precipitation such as snow and rain are forming. Many of these phenomena are determined by the use of cross-correlation ratio in conjunction with reflectivity or another radar product which gives a more complete set of information for classification. It can be used by itself, however, for fitting a volume into one of the three regions defined in Table 4.6.

Table 4.6: Regions of cross-correlation values and commonly measured scatterers which fit into them. This includes non-meteorological and meteorological scatterers.

Cross-Correlation	*Description*	*Examples*
$0.00 < \rho_{HV} < 0.80$	Erratic movement relative to precipitation and complex scattering from pulse-to-pulse.	Birds, insects, planes, etc.
$0.80 < \rho_{HV} < 0.97$	More uniform scattering but not completely and possibly multiple types of precipitation.	Hail, melting snow, etc.
$0.97 < \rho_{HV} < 1.00$	Uniform scattering from pulse-to-pulse.	Rain, snow, etc.

4.10 The Range-Doppler Spectrum

Phase can be difficult to control even in a coherent system which makes calculation of target velocity by the method in Section 4.8.3 difficult as well. As previously mentioned, since range information is stored in the frequency of a signal, it is ambiguous whether the signal's frequency is a result of the range or Doppler velocity but it is possible to quantify this by the use of the *ambiguity function*.

Attaining the ambiguity function is done by combining the linear frequency modulation (LFM) waveform from FMCW radar with pulse compression which effectively introduces a time delay between the end of the LFM ramp and the start of the next ramp. The radar transmits the ramp for T_{ramp} then waits for some amount of time, T_{off}, within each pulse repetition interval (PRI), T_{PRI}. This process is repeated and the PRIs are divided into M buffers of size N giving two time axes, *fast time* and *slow time*. The fast time axis corresponds to the times between individual points in a buffer and is separated by the ADC's sampling period, T_s while slow time corresponds to the time intervals between buffers, T_{PRI}. This two-dimensional representation of M buffers can be seen in Figure 4.25.

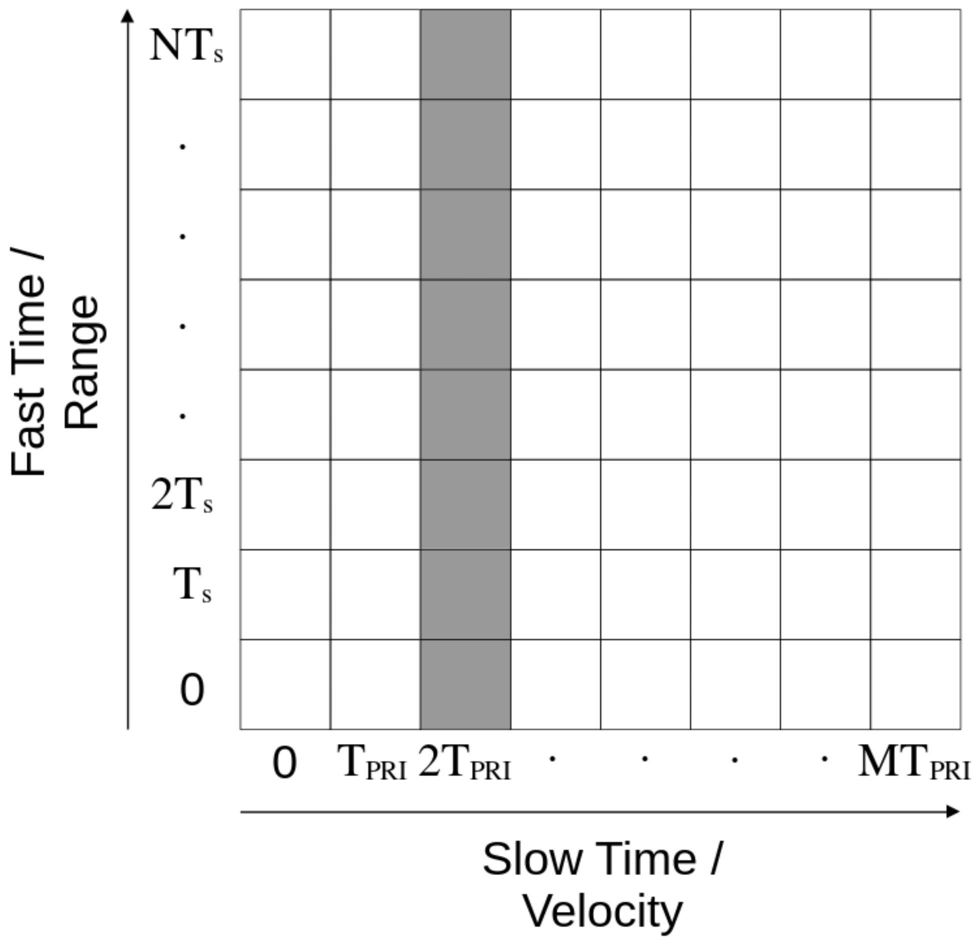

Figure 4.25: Diagram showing the relationship between fast and slow time with one buffer of size N shaded.

With the received signal now in two dimensions, the beat signal's time domain equation can be related to fast and slow time by,

$$b[l,m] = \underbrace{a \exp\left[j2\pi \left(\underbrace{\frac{2f_c v}{c}}_{\text{Doppler shift}} + \underbrace{\frac{2B_{RF}R}{t_{ramp}c}}_{f_{beat}} \right) lT_s \right]}_{\text{Fast time}} \underbrace{\exp\left[j2\pi \frac{2f_c v}{c} mT_{PRI} \right]}_{\text{Slow time}} \tag{4.22}$$

where l and m are the indices of fast and slow time, respectively, v is the object's velocity, f_c is the frequency of the transmitted signal, and a is the amplitude. Applying the two-dimensional discrete Fourier transform (DFT) results in the following equation for the corresponding frequency spectrum,

$$B[p,k] = \frac{1}{\sqrt{NM}} \sum_{l=0}^{N-1} \sum_{m=0}^{M-1} b[l,m] \exp\left[-j2\pi\left(\frac{lp}{N} + \frac{mk}{M}\right)\right] \tag{4.23}$$

where N and M are the number of fast and slow time indices, respectively and p and k are the indices of the two-dimensional DFT spectrum. Each point is a function of the range and velocity, so a target with range, R, and velocity, v, can be found at the point, (p, k) where

$$p = \left(\frac{2f_c v}{c} + \frac{2B_{RF}R}{T_{ramp}c}\right) T_{ramp} \tag{4.24a}$$

$$k = \frac{2f_c v}{c} M T_{PRI}. \tag{4.24b}$$

p can be seen to be a function of both the range and velocity but k is solely a function of velocity. When plotted, this shows this ambiguity of range and velocity for all targets. This is shown as the range-Doppler spectrum in Figure 4.26 for the same target at three different velocities, $v = 0, \pm 0.3$m s^{-1}, and roughly the same range, $R = 2$m. For Figure 4.26c and Figure 4.26b where the targets have some non-zero velocity there is still a significant amount of power around $v = 0$m s^{-1} which can be attributed to the stationary clutter in the indoor environment. This shows the range-Doppler spectrum's effectiveness at classifying clutter when all targets are in motion and how it can suffer when some targets are stationary as in Figure 4.26a. In fact, without a priori knowledge of the stationary target's existence at $R = 2$m, it would be difficult to discern it from clutter.

To understand how the two-dimensional DFT outputs information about the range and velocity of the measured targets, it is important to understand how it is computed. First, the DFT is computed for the first dimension, i.e., for each of the individual fast time buffers separated in time by T_s. This results in M range-FFT spectrums as in Figure 2.5 which means each point in the fast time axis now corresponds to a range. Next, the DFT is computed for the second dimension, i.e., for each of the individual ranges just calculated separated in time by T_{PRI}. This effectively gives a spectrum showing the rate at which the power at individual ranges change, or the velocities of targets in that range. The range resolution (2.3b), repeated in (4.25a), can be modified to this

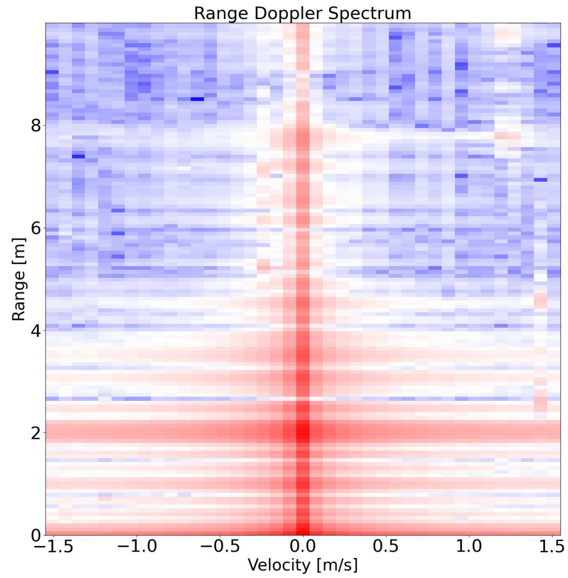

(a) Range-Doppler spectrum with a stationary target at $R = 2\text{m}$.

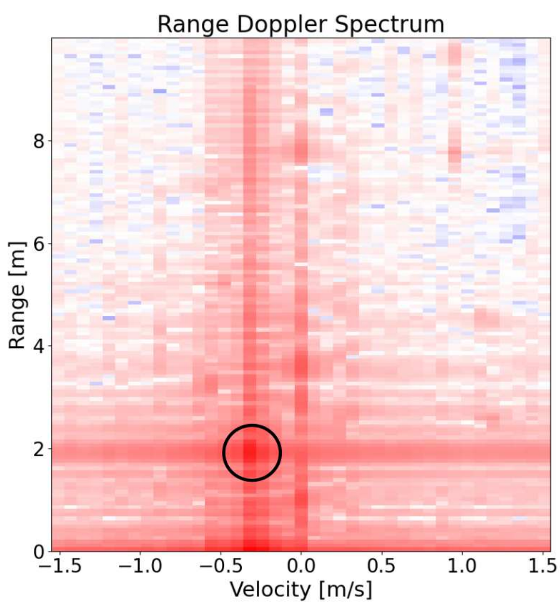

(b) Range-Doppler spectrum with a target at $R = 2\text{m}$ traveling at $v = -0.3\text{m s}^{-1}$.

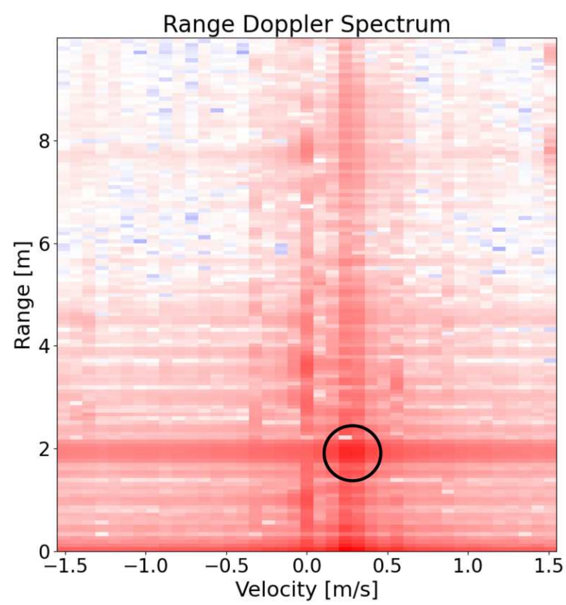

(c) Range-Doppler spectrum with a target at $R = 2\text{m}$ traveling at $v = 0.3\text{m s}^{-1}$.

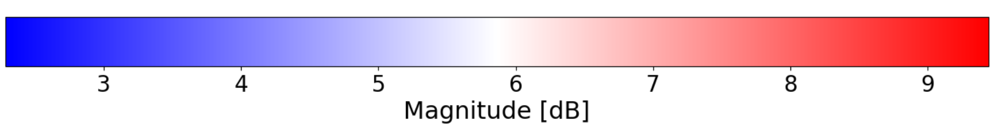

(d) Colorbar for the above range-Doppler spectrums.

Figure 4.26: Range-Doppler spectrums of a stationary target and two moving targets with velocities, $v = \pm 0.3\text{m s}^{-1}$, all at a range of roughly $R = 2\text{m}$. Targets are circled for clarity.

second dimension with sampling period, T_{PRI}, and size, M to get the velocity resolution,

$$R_{res} = \frac{c}{2B_{RF}} \tag{4.25a}$$

$$v_{res} = \frac{\lambda}{2MT_{PRI}} = \frac{\lambda}{2T_{CPI}} \tag{4.25b}$$

where $T_{CPI} = MT_{PRI}$ is the time for the entire coherent processing interval. This means if two targets are very close in range, as in Figure 4.3, they can be distinguished from one another in the range-Doppler spectrum if their velocities relative to one another are at least v_{res} apart, or $|v_1 - v_2| \geq v_{res}$.

The bounds for the range-Doppler spectrum are defined by the sampling rates, f_s for range and the pulse repetition frequency (PRF $= 1/T_{PRI}$) for velocity and are,

$$-\frac{f_s c}{4S} \leq R \leq \frac{f_s c}{4S} \tag{4.26a}$$

$$-\frac{\lambda}{4T_{PRI}} \leq v \leq \frac{\lambda}{4T_{PRI}} \tag{4.26b}$$

where $S = B_{RF}/t_{ramp}$ is the ramp slope. The PRI can also be expressed as a function of the ramp and settle times from Figure 3.13 as $T_{PRI} = t_{ramp} + t_{settle}$ which makes the range directly proportional to t_{ramp} and the velocity and its resolution inversely proportional to t_{ramp}.

$$-\frac{f_s t_{rampc}}{4B_{RF}} \leq R \leq \frac{f_s t_{ramp} c}{4B_{RF}} \tag{4.27a}$$

$$-\frac{\lambda}{4(t_{ramp} + t_{settle})} \leq v \leq \frac{\lambda}{4(t_{ramp} + t_{settle})} \tag{4.27b}$$

$$v_{res} = \frac{\lambda}{2M(t_{ramp} + t_{settle})} \tag{4.27c}$$

Given this relation, a trade-off must be made between maximum unambiguous range and velocity, where t_{ramp} is main the *design knob*. This choice will depend on the expected targets and their behaviour. The environment used to obtain the range-Doppler spectrum in Figure 4.26 had high-clutter and a maximum expected target velocity of $v \approx 0.5$m s^{-1}. This made it preferable to have a low range of velocities and fine velocity resolution to make differentiation between clutter and targets easier. For this reason, a high ramp slope, $t_{ramp} = 10$ms, was chosen. This relationship can be seen plotted over the range 1μs $\leq t_{ramp} \leq 1$ms in Figure 4.27.

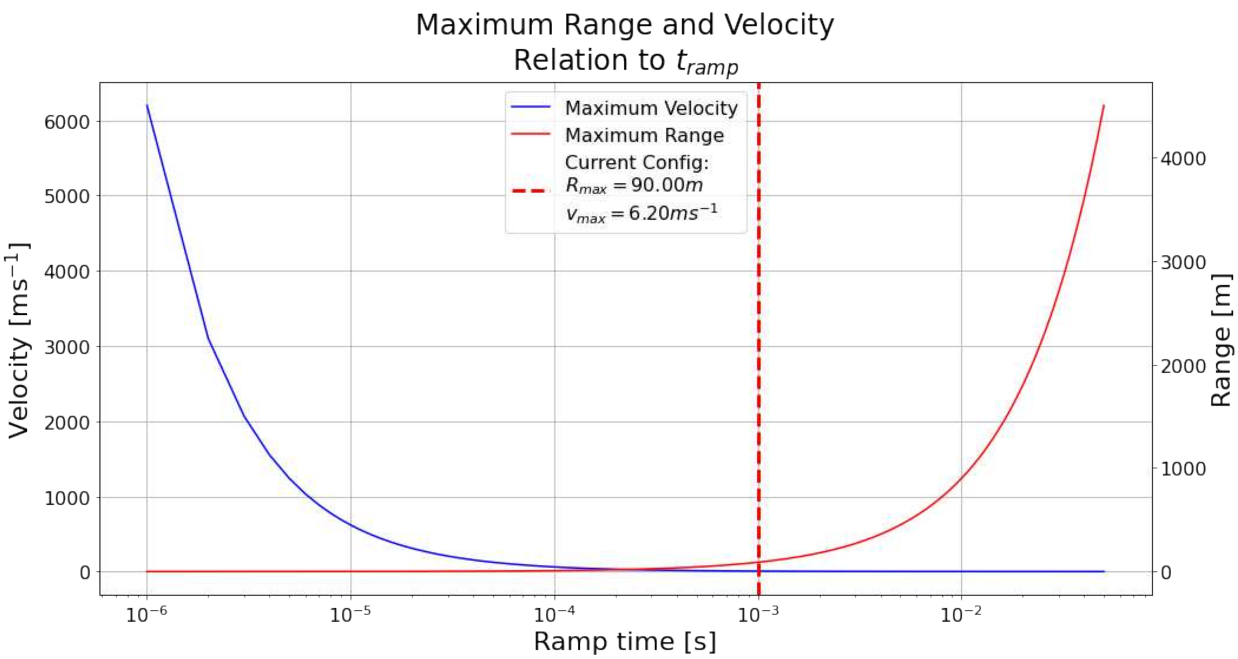

Figure 4.27: Maximum range and velocity as a function of the ramp time, t_{ramp} for the phaser radar with parameters in Table 3.2. The ramp time is swept from 1μs to 50ms and the sampling rate is assumed to be set to $f_s = 2f_{IF} + 1$MHz, effectively giving a baseband sampling rate of $f_s = 1$MHz.

For DARMA, the maximum unambiguous velocity, $v = 6.2$m s^{-1} will most likely need to be increased for proper measurement of meteorological events. As previously mentioned, a lower ramp time could achieve this but at the cost of range which, for the configuration in Table 3.2, is already well below the desired range, therefore other parameters will need to be modified. According to (4.27a), the range is also inversely proportional to the RF bandwidth, which for DARMA's base configuration is 3MHz as opposed to the phaser's configured bandwidth of 500MHz. This difference brings the maximum range significantly higher and allows for more variation of the ramp time to increase the maximum unambiguous velocity. Modifying DARMA's configuration to use a lower ramp time, $t_{ramp} = 1$ms would achieve a maximum unambiguous range of 5km, assuming the use of an ADC with $f_s = 200$kHz, and a maximum unambiguous velocity of 7.14m s^{-1}. The maximum unambiguous range of 5km indicates that the ramp time could be decreased by another factor of 5 while maintaining the desired range of 1km, effectively increasing the velocity by a factor of 5 as well (for $t_{settle} = 0$s). This relation will not be as direct in systems with $t_{settle} \neq 0$ or with multiple ramps within a single PRI, however.

4.11 Clutter

Clutter, sometimes called *ground clutter*, is the reflected signal from generally stationary objects outside of what is considered a target. Included in the term clutter are buildings, walls, trees, and other objects that are, in many cases, larger than the target and further contribute larger amounts of power to the range-FFT spectrum. An example can be seen in Figure 4.28 where trees

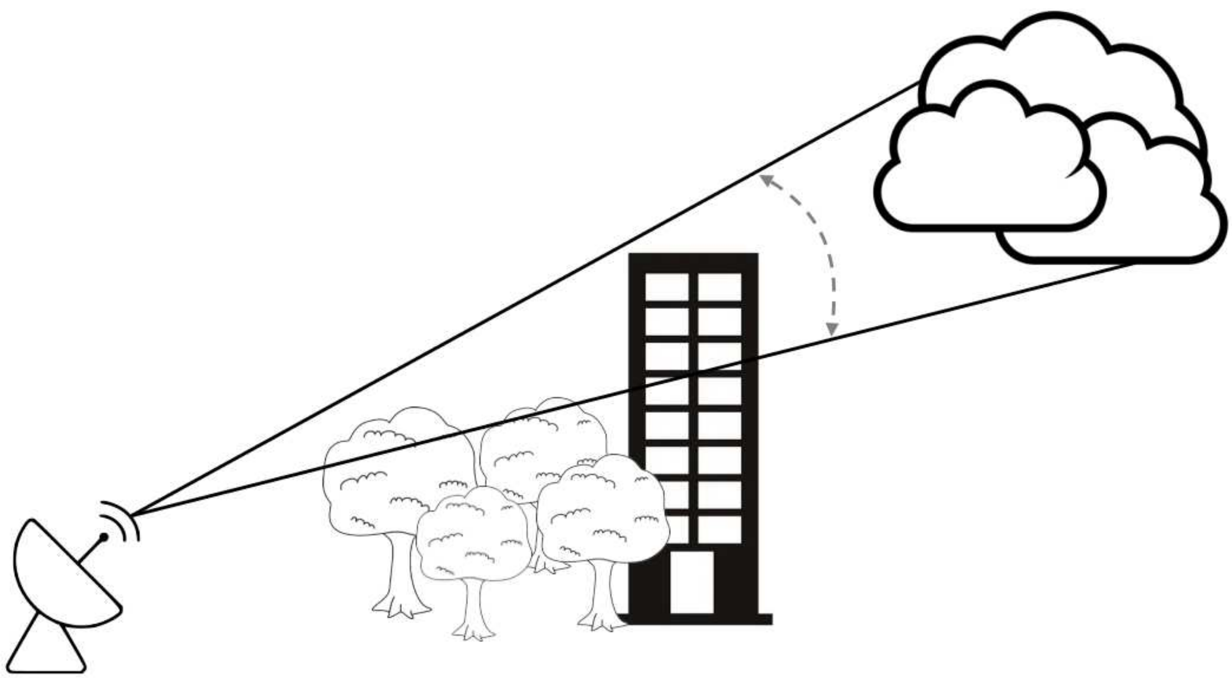

Figure 4.28: Example of ground clutter interfering with the measurement of a precipitation volume.

and a building will reflect large amounts of power back to the antenna. This is detrimental to target detection because it can overshadow targets of interest or be misclassified as one. Clutter therefore requires its own form of filtering and is generally more involved relative to the leakage and maximum range filters discussed in Section 4.5.2 and Section 4.5.3, respectively.

In differentiating between high reflected powers from clutter and targets, an object's velocity, introduced in Chapter 2, becomes very useful because clutter is generally stationary while a target, such as a human or drone, is generally in motion. Even in the case where the target is not walking or flying in some direction, there will most likely be some small movements (swaying, heartbeat, etc.) in the human case or wind interference and a control loop accounting for it in the case of a drone. Such minuscule changes can often be detected by obtaining the phase shift, ϕ_0, between two time samples then using (2.6) to obtain the velocity. The range-Doppler spectrum is another method of differentiating between clutter and targets which expands the received signal into two dimensions by collecting multiple buffers and calculating the Fourier transform of individual ranges between the buffer times. This expands the range-FFT spectrum to another dimension of velocity. These processes are described in detail in Section 4.8.3 and Section 4.10, respectively. Once the velocity is calculated, there are numerous methods for classifying certain parts of the signal as clutter which range from setting a maximum velocity threshold, similar to the static threshold target detection technique in Section 4.7.2, to more complex methods like the popular gaussian model adaptive processing (GMAP) method described in [21].

To demonstrate the static velocity threshold clutter filtering technique, the range-Doppler spectrum must first be estimated via the method described in Section 4.10 and can be seen in Figure 4.29a. In the range-Doppler spectrum there is a high-intensity line spanning most of the ranges centered at $v = 0\text{m s}^{-1}$ which can be considered clutter. The measurement was taken in an indoor room with significant clutter, hence the high intensity. There is another, less intense, region of interest at approximately $R = 2\text{m}$ and $v = 0.3\text{m s}^{-1}$ which is a moving reflector panel at boresight. Looking at one of the range-FFT spectrums used to obtain this range-Doppler spectrum (Figure 4.30a), it is difficult or impossible to differentiate between the noise and target at $R \approx 2\text{m}$. Since there is a known target moving at $v = 0.3\text{m s}^{-1}$, a static threshold can be applied to the velocity spectrum to mitigate the effect clutter has on target detection. The velocity threshold was

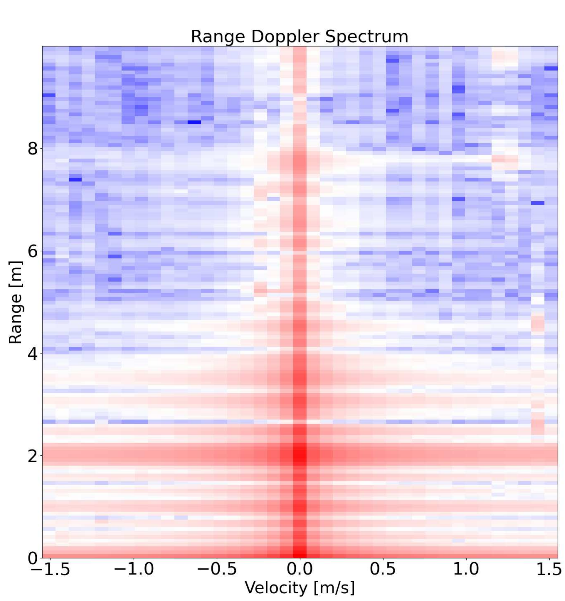

(a) No clutter filtering applied.

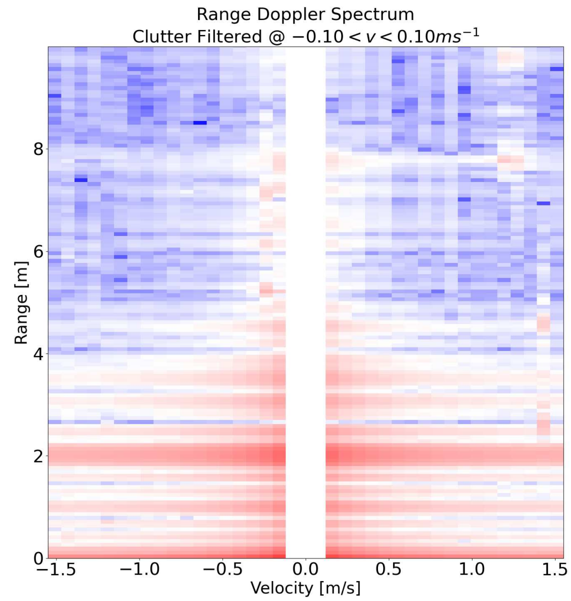

(b) Static velocity threshold ($-0.1 \leq v \leq 0.1\text{m s}^{-1}$) applied as a method of clutter filtering

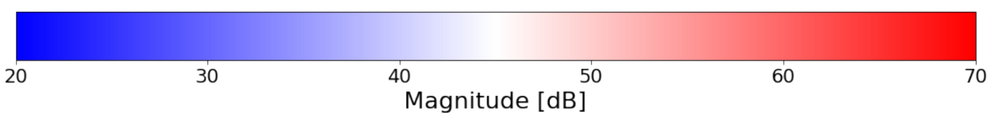

(c) Colorbar for the above range-Doppler spectrums.

Figure 4.29: Range-Doppler spectrum estimated as a baseline for a clutter filtering demonstration. The measurement was taken with buffer size, $N = 40000$, and $M = 40$ range-FFT spectrums. The first range-FFT spectrum is ignored to allow for settling of the ramp generation circuitry resulting in $M = 39$.

chosen to filter all targets within the range, $-0.1 \leq v \leq 0.1\text{m s}^{-1}$. In this static threshold case, these filtered targets are simply set to zero magnitude, but more advanced filtering can be done to not completely remove targets with low velocities [21].

The range-Doppler spectrum with clutter filtering applied can be seen in Figure 4.29b. The inverse, two-dimensional DFT is then applied to the range-Doppler spectrum to obtain M clutter filtered range-FFT spectrums. The spectrum with slow time index, $m = 6$, is shown in Figure 4.30b. Comparing Figure 4.30a and Figure 4.30b shows the large impact clutter filtering has on the visibility of moving targets. The reflector panel can now be differentiated from its environment with ease whereas before it was completely masked by noise. There is the downside, however that some desired targets may reside within the low Doppler velocities, but this can be remedied

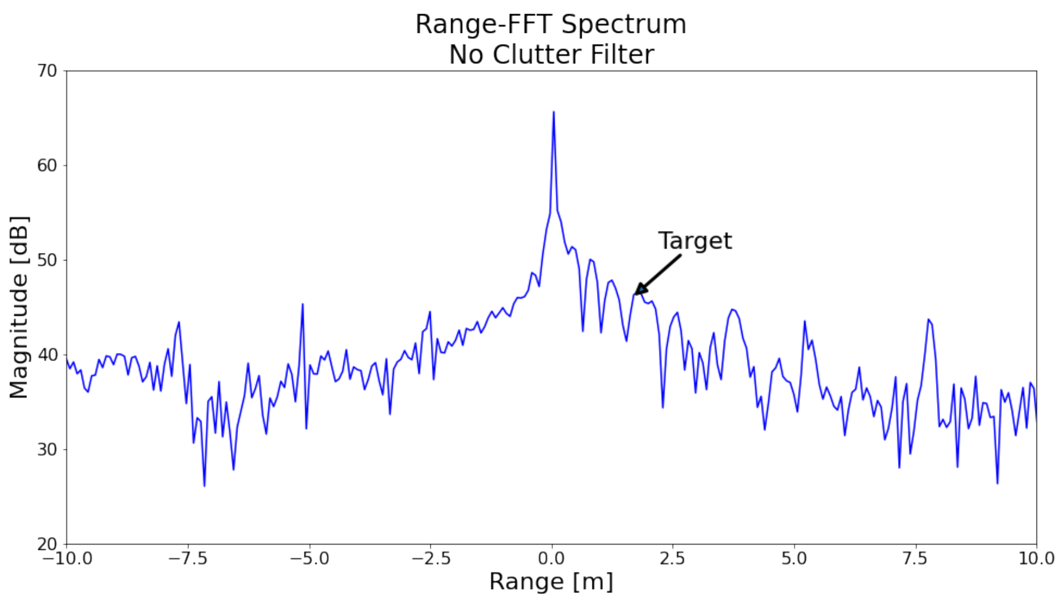

(a) No clutter filtering applied.

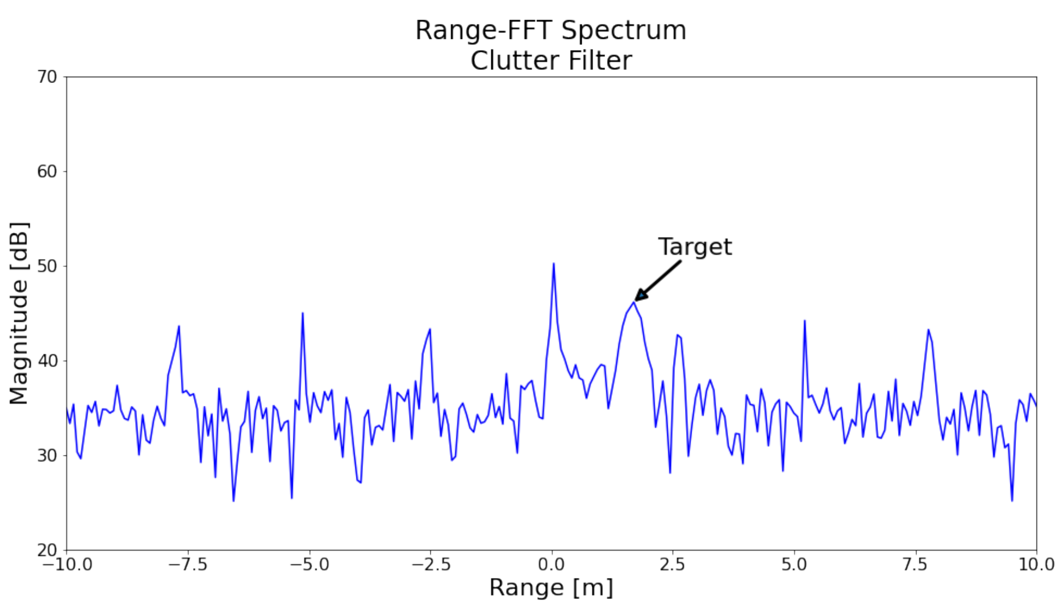

(b) Static velocity threshold ($-0.1 \leq v \leq 0.1$m s^{-1}) applied as a method of clutter filtering.

Figure 4.30: Range-FFT spectrum at slow time index, $m = 6$, of the range-Doppler spectrum shown in Figure 4.29a. A moving reflector panel is at $R \approx 2$m with a velocity, $v = 0.3$m s^{-1}.

by using a less strict filter. The aforementioned GMAP technique can be employed in this case in an attempt to recover targets within this range.

If a constant environment is being measured, a more simple process known as coherent background subtraction (CBS) can be used to mitigate clutter. It is effectively a calibration technique which takes a scan of the static environment before measurement of targets to get a baseline for the clutter present in the environment. The magnitude and phase obtained in this baseline can then be subtracted from the subsequent measurements to mitigate the effect of clutter. It is also possible to use this technique in a non-coherent system by only subtracting the magnitude of the baseline, but it is far less effective. This method has largely become obsolete, however, due to the ease with which the above range-Doppler spectrum method can be employed and its effectiveness.

4.12 Display

The signal processing explained in previous sections is integral for understanding the radar's readings but it is not the final step. Often radar systems are meant to be viewed by customers with minimal or no radar knowledge which makes the displaying of the data very important. There are a multitude of methods for displaying radar data, some of which are plan-position indicator plots (PPI), range-height indicator plots (RHI), and waterfall plots. PPI plots are the "classic" example most are familiar with; they sweep in azimuth and plot range bins radially from the center. This is useful if all the targets one is expecting are in the horizontal plane, but it is not necessarily limited to one elevation angle as often times the radar can be positioned at different elevations before sweeping in azimuth. RHI plots are different in that they sweep in elevation, but they still plot range radially from the radar. This is useful for the situation where one expects targets to be in the vertical plane. An example of this which is commonly used is taking a cross section of a storm in the vertical direction. This is helpful in determining the height of different precipitation layers within the storm. Examples of both PPI and RHI scans taken with the CSU-CHIVO radar [22] can be seen in Figure 4.31.

Both PPI and RHI scans can be very computationally intensive to display because they require setting and updating large amounts of pixels at very short intervals. Failure to be *sufficiently fast* can mean missing quickly-moving targets or frustrating the user. Sufficiently fast is a term relative to the current application and can vary widely depending on the expected velocity of targets and numerous other factors. Meeting the speed standard for a specific application's display can be

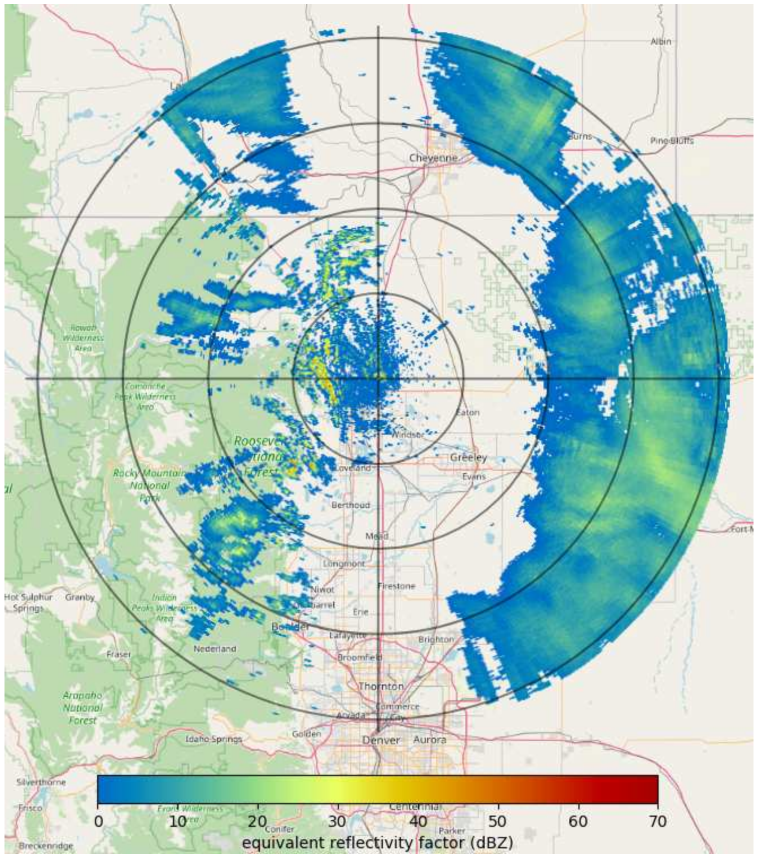

(a) Plan-position indicator scan.

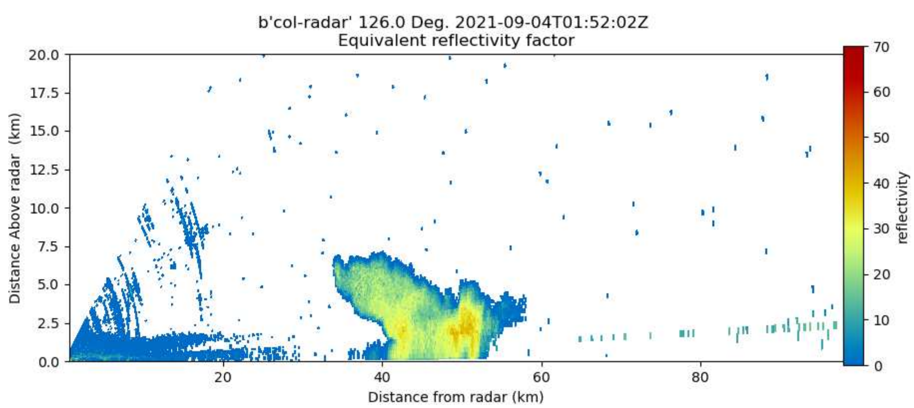

(b) Range-height indicator scan.

Figure 4.31: Example of plan-position indicator (PPI) and range-height indicator (RHI) reflectivity scans taken with the CSU-CHIVO radar.

difficult because of available computing power and might require making trade-offs in terms of the signal processing. Target detection algorithms such as CFAR that calculate different thresholds for each range bin are examples of processing that can greatly degrade the real-time performance. It is therefore necessary to determine where on the spectrum of speed to precision needs a certain project resides.

4.13 Summary

The signal processing used in a certain application varies widely depending on the environment, desired targets, and many other factors. This section describes some of the basic blocks many radar systems employ before doing some more in-depth processing. Figure 4.32 provides a summary of these common blocks and the general order in which they are used. Some of these blocks may not

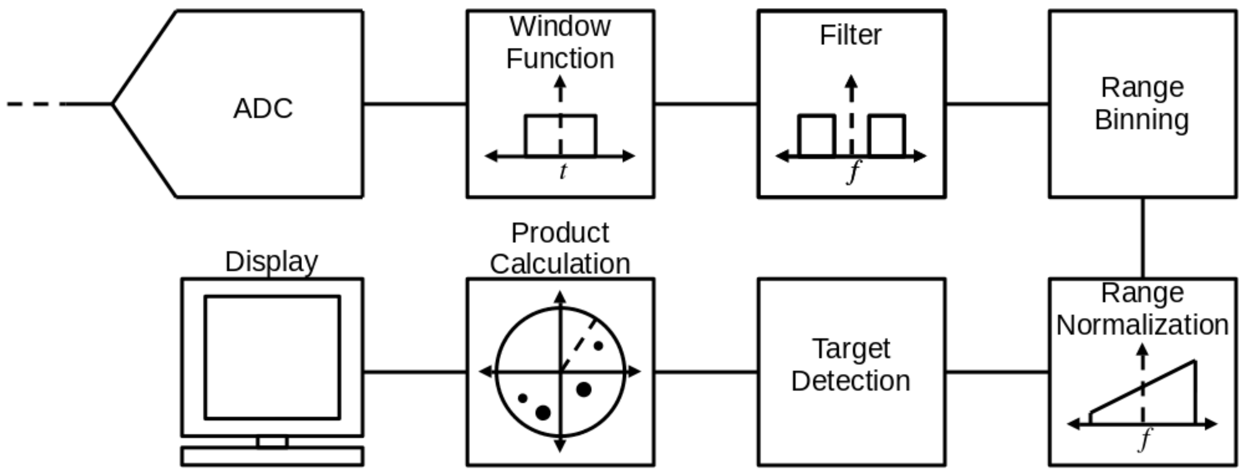

Figure 4.32: Block diagram showing the signal processing modules discussed in Chapter 4 connected together to show the process from sampling to displaying of radar data.

be necessary in all applications such as the product calculation and display in a collision avoidance system which may not need to relay detailed information to the user but simply be used to course correct for a drone. Available computing power, timing requirements, and hardware may also make an application deviate greatly from that described in this section.

Most of the material in this section is oriented towards real-time processing which is only a small component in many radar systems. Post-processing offers superior computing power and time which allows it to be used for immensely complex radar algorithms which are far beyond the scope of this thesis but may be implemented in the future which is a possibility discussed in further detail in Section 7.2.

Chapter 5

NOAA FMCW Snow-Level Radar Comparison

5.1 Introduction

The original goal of this project was to propose an upgrade to an FMCW radar system designed by the National Oceanic and Atmospheric Association (NOAA) shown in Figure 5.1 [23]. This

Figure 5.1: Picture of the NOAA snow-level radar taken by Clark King, NOAA.

goal changed over the course of the project but the NOAA radar was still used as a reference throughout because of its success in measuring meteorological events. Currently the NOAA radar operates in one polarization and points directly upward with the purpose of measuring snowfall. Though the use case is different than the that of DARMA, it is a useful comparison for determining the efficacy of this radar. This chapter discusses the differences and between the two radars and the motivation behind these design choices.

5.2 Key Differences

It is heavily emphasized that both these radars were designed to serve separate purposes, but it is useful to compare current work to successful projects in the past with similar characteristics. A high-level comparison of some of the key parameters from the ground-based FMCW radar described in this thesis and the NOAA FMCW Snow-Level Radar can be seen in Table 5.1. It

Table 5.1: Comparison of some of the key parameters for the ground-based FMCW radar presented in Chapter 3 and NOAA's *FMCW Snow-Level Radar*.

Parameter	*Description*	CSU	NOAA	*Unit*
R_{max}	Maximum Range	1,000	10,000	m
P_t	Transmit Power	30	30.04	dBm
λ	Wavelength	28.57	106	mm
f_c	Center Frequency	10,500	2,830	MHz
B_{RF}	RF Bandwidth	3	3.75	MHz
$f_{beat,max}$	Baseband Bandwidth	0.2	22.11	kHz
t_{ramp}	Ramp Time	10	1.13	ms
R_{res}	Range Resolution	50	40	m
-	Polarization	Dual-Pol	Single-Pol	-
G	Antenna Gain	10	27.6	dBi
θ_1	Horizontal Beamwidth	10	6	degrees
ϕ_1	Vertical Beamwidth	10	6	degrees

should be noted that DARMA's antenna gain is specific to the 8x4 phased array antenna shown in Figure 3.21 but it is capable of swapping antennas which allows for various gains to be achieved depending on requirements. One of the most notable differences between the two is the intended application which affects many components of the design. NOAA designed their radar intending for it to be stationary and pointed directly upward for measurement of snow. This reduces some of the complexity because there is no need for beamforming and beam-steering which is a large topic of concern for this radar and is discussed in depth in Section 3.5. The stationary nature of the radar also allows its size to be less strict which was another big concern for this radar (Section 3.7). Despite those relative simplifications, the snow-level radar is possibly a more difficult system to deploy due its desired target, snow. This increases the need for a robust design able to withstand extreme temperatures, wet environments, precipitation build-up, and more. One of the effects of this is the need for custom antenna covers which are slanted to allow for the snow to slide off, reducing the interference from build-up.

The maximum range, R_{max}, the snow-level radar is able to achieve is greater than DARMA's by a factor of 10. This ability can be attributed to a few factors, most notably the antenna gain, G_t, G_r, and center frequency, f_c. This would also include the transmit power, P_t, but both radars are near identical in this area. As previously mentioned, the snow-level radar has less strict size requirements and can therefore use larger, higher gain antennas as seen in Figure 5.1 surrounded by a metal tube and topped with an electrically invisible tarp. They are 1.2m parabolic reflectors illuminated by a linearly-polarized microstrip patch antenna which results in transmit and receive gains of $G = 27.6$dBi, roughly 17dB above what the small microstrip patch arrays used in this thesis are capable of. Another benefit for long-distance propagation is its wavelength, $\lambda = 106$mm, roughly 4 times as large as this radar ($\lambda = 28.57$mm). According to (3.1), which states that the free-space path loss is inversely proportional to λ^2, the snow-level radar's signal will suffer approximately 16 times less attenuation at a given range. Since the main feature of DARMA is its modularity and ability to operate in multiple environments, it is entirely possible to connect a higher gain antenna to either or both sides for improved range as suggested in Section 3.10. This radar does not, however, possess the ability to operate at S-band like the snow-level radar does which will still be a limiting factor for range, but it was necessary for the reasons outlined in Section 3.2.2.

Besides the maximum range, the baseband bandwidth is one of the most obvious discrepancies between the two radars with DARMA having $f_{beat,max} = 0.2$kHz and the snow-level radar having $f_{beat,max} = 22.11$kHz. This difference can be explained by (3.2) which states that beat frequency is directly proportional to range and RF bandwidth and inversely proportional to ramp time. In this case, the RF bandwidths are relatively close, $B_{RF} = 3,\ 3.75$MHz, which means $f_{beat,max}$ is mostly affected by the maximum range and ramp time which are 10 times larger and smaller, respectively, for the snow-level radar.

Dual-polarization can be very helpful in characterizing weather because it introduces a few radar products: differential reflectivity, Z_{dr}, differential phase, Φ_{DP}, and cross-correlation ratio, ρ_{HV} [24] which can then be used to obtain additional derived radar products, of which the most common is specific differential phase, K_{dp}. Differential reflectivity is the ratio in logarithm form between the reflectivity measurements obtained by the horizontally- and vertically-polarized receivers, where $Z_{dr} > 0$ indicates greater reflectivity in the horizontal polarization. It can be used to determine raindrop size, detect updrafts, and other phenomena. Electromagnetic waves are slowed when passing through a non-vacuum medium ($\varepsilon_r > 1$) which results in a phase shift. For a non-uniform object, phase shifts will be unequal in the horizontal and vertical polarizations which is measured as differential phase. This is not exceptionally useful because phase shift is accumulated as the wave passes through subsequent objects, so a new product, specific differential phase, is derived. K_{dp} shows the change in Φ_{DP} with range and is a good indication of size and concentration of precipitation within a range bin. Cross-correlation ratio is a measure of the variability in the shapes and sizes of targets within the volume, where higher values indicate less variability. These are the base dual-polarized products which are explained in more depth in Section 4.9, but more exist for more specific classification. They are all very useful in meteorological measurement and the NOAA radar could most likely benefit greatly from the addition. DARMA has four dedicated connections for a second polarization which is one of its advantages.

Despite the comparison of the two as if they were alike, the NOAA radar was designed for a completely separate application and it is stressed that the two are not in competition. The comparison is more a useful tool for exploring other options and methodologies for radar design.

Chapter 6

Component Sourcing

6.1 Introduction

Due to circumstances outside the scope of this thesis, the electronics market has grown sparse with extremely long lead-times. This electronics market climate is a large issue for radar systems because they employ complex and specific components which are often not of a very high priority to manufacturers. Throughout this thesis there have been multiple references to this section when discussing the design trade-offs of certain components because sourcing has become a much larger concern than in the past. The number of times a long lead-time for a product affected the component choice or even the entire system structure far exceeds those discussed in Chapter 3, but it was often omitted because reading would become frustrating if it were mentioned every time.

6.2 Components

Many of the components used in a radar system such as this are complex and have stringent requirements for them to operate properly with the other components. This complexity and specificity leaves few options even before introducing the market's situation. This has led RF components to be heavily sought-after which can make procurement very difficult or sometimes completely impractical for a given timeline. Some of the most affected components are discussed in this section, including beamformers, frequency synthesizers, etc.

Beamformers and phase-shifters were among the most difficult to find because they are very intricate components and are needed for phased array systems which led them to being purchased quickly, extending their lead-times. This issue almost made the final creation of the PCB impossible in the short time frame. After the two designs with their respective beamformers were decided upon, both went out of stock or had other procurement complications leaving the radar system inoperable. Fortunately, Analog Devices and AnokiWave worked to find other suppliers or sent samples to get the prototypes built. This would not be sufficient if the goal were to put radars into production, however.

Frequency synthesizers are arguably the most important component in an FMCW radar because they generate the frequency-modulated output signal. As discussed in Section 3.3, the ADF4159 was chosen due to its large suite of options specifically tailored to FMCW signal generation which was going to be a large time-saving choice. Like the beamformers, however, this component is one of the best options on the market for signal generation which has made it a difficult item to procure. Other RF components such as voltage-controlled oscillators (VCO) and RF amplifiers also caused issues, but to a lesser degree because they are easier to swap for one another.

To avoid these complications, it is imperative when using complex components such as these, or most electronics in this market, to begin procurement immediately when components are finalized. This will not completely remedy the problem, but it will help in knowing which components need changing further from a project's deadline.

6.3 Fabrication and Assembly

Since most of the discussion surrounding the availability and lead-time is about components and specifically ICs, fabrication and assembly of the PCB was mostly overlooked until it was needed which, in hindsight, was a mistake. As discussed in Section 3.9, the high frequency of the radar made it impractical to use a more common substrate such as FR4; a more performance and less common substrate, Rogers 4350, was needed. Rogers material is already significantly more expensive than FR4 but the limited availability worsened this to a large degree. In addition to the price increase, labor is another factor to consider when working on a specialized electronics product.

The amount of available labor has taken a massive hit as well which is another contributor to the expense of PCB fabrication. Because of this, it is not only the components with high lead-times but fabrication turn-times are long as well. For example, DARMA, with important ordering parameters in Table 6.1, was quoted a minimum fabrication turn-time of 15 days or 20 days at a more reasonable rate. This much lower rate was still approximately 6 times what was expected when beginning the project and the assembly was a few thousand dollars in addition. It should

Table 6.1: Important parameters for fabrication of the PCB. These were some of the largest contributors toward the price and turn-time of both fabrication and assembly.

Parameter	*Value*
X Dimension	5.406"
Y Dimension	3.536"
Layers	4
Dielectric material	Rogers 4350B
Finish	Immersion Gold
Impedance Control?	Yes
Blind Vias?	Yes
Micro Vias?	Yes

be noted that this PCB is a very complex board with many components, vias, micro-vias, and impedance control which makes it more difficult to fabricate. That being said, the price and turn-time was still significantly higher than expected, most likely due to the market issues.

These issues were not specific to a single fabrication company but extended across the field. Many companies did not have availability for high frequency material or the capability to fabricate a board with these requirements which narrowed the options even further. The important take-away for the current electronics climate is to begin sourcing components and coordinating with fabrication companies as early as possible to minimize delays. For exact cost details, the reader is referred to Appendix B.

Chapter 7

Summary and Future Work

7.1 Summary

The objective of this thesis is to present DARMA, a low-cost, modular FMCW radar capable of being deployed on an unmanned aircraft system (UAS) and ground platform and to provide a discussion of the important design, deployment, and cost considerations taken when designing such a system. DARMA followed the general FMCW radar structure, seen in Figure 7.1, with some modifications such as adding support for dual-polarization receiving. The considerations were provided

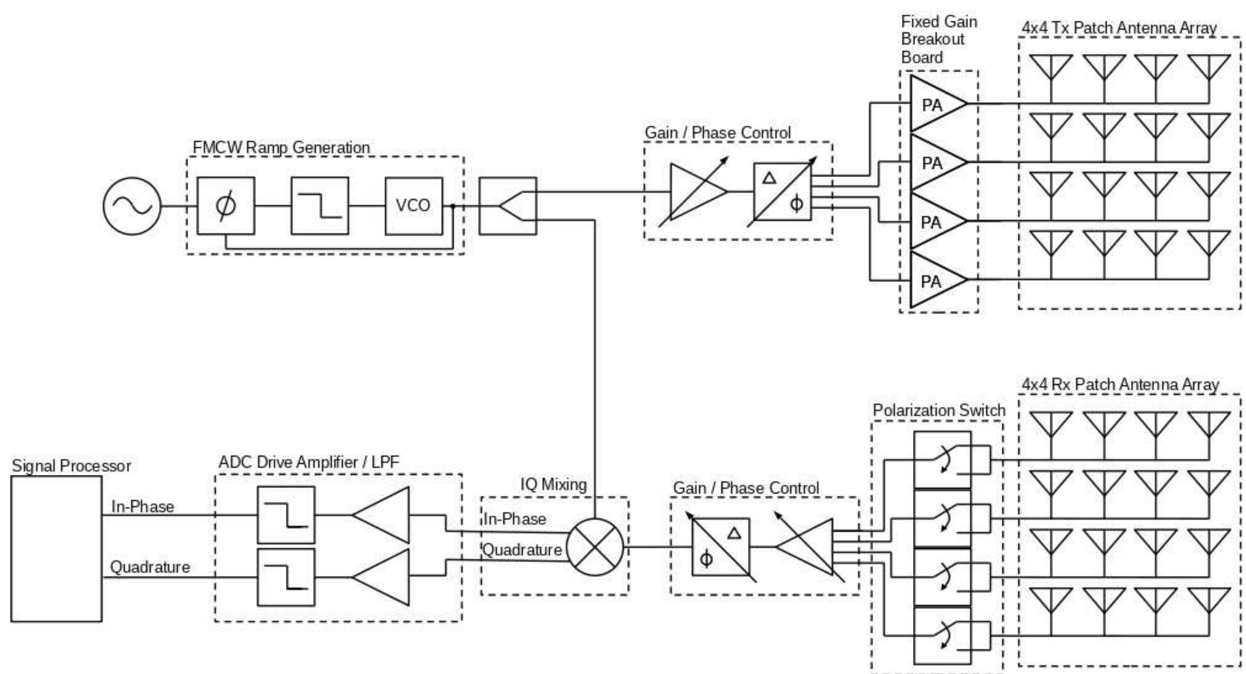

Figure 7.1: Block diagram for FMCW radar modified to support the requirements outlined in Section 3.1 such as the addition of dual-polarization receiving and support for modular subsystems.

with accompanying suggestions for optimal performance as well as options to consider for application requirements different than those in this thesis. Due to this design's inclusion of modular subsystems, available modifications were presented to work within these different application requirements. Common signal processing methods including noise mitigation, target detection, and estimation of weather radar products were also introduced and implemented on an FMCW radar. This was done with the intent of providing a basis for obtaining a clean signal from which further processing and interpretation can be done. It was also to describe the differences in processing methods for traditional pulsed and FMCW radars. These implemented and discussed methods are shown in block diagram form in Figure 7.2.

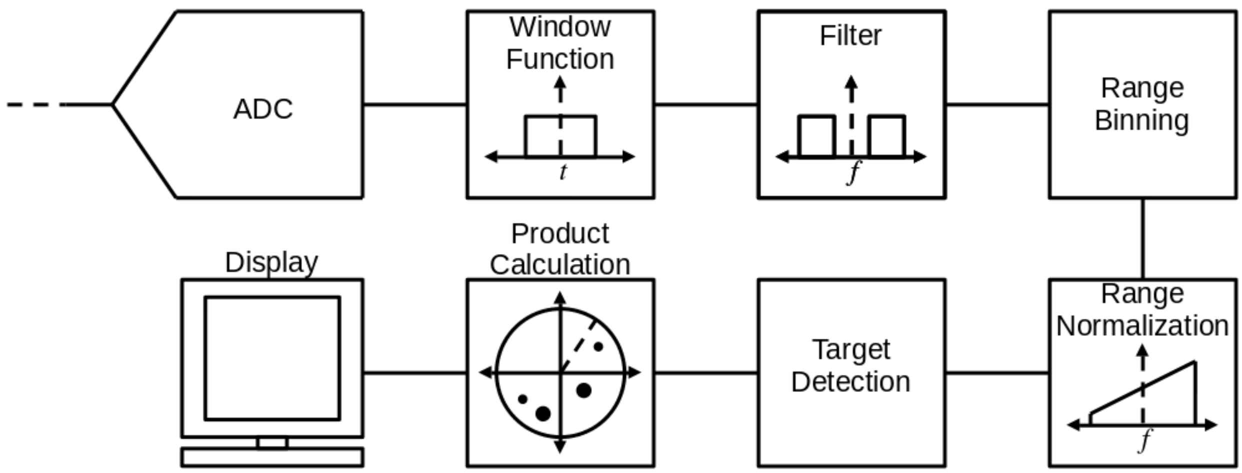

Figure 7.2: Block diagram showing the flow of data through the signal processing methods and algorithms described in Chapter 4.

The radar presented in this thesis, DARMA, was designed to operate at X-band ($f_c = 10.5$GHz) and be able to measure targets at distances up to $R_{max} = 1$km with dual-polarization receiving in the base configuration. It also supports four element beam-steering and beamforming on transmit and receive which is accomplished by the use of two beamforming ICs, Analog Device's ADAR1000 and AnokiWave's AWS-0101. The two are used separately on two separate radars for the purpose of comparing their respective performances. The base configuration achieves this maximum range by transmitting $P_t = 30$dBm with an antenna gain of $G_t = G_r = 10$dBi. The transmitted bandwidth is $B_{RF} = 3$MHz at a ramp slope of $t_{ramp} = 10$ms which, along with the maximum range, produces a baseband bandwidth of 200Hz. These design parameters are summarized in Table 7.1 and a 3D model of both the ADAR1000- and AWS-0101-based radars can be seen in Figure 7.3a and Figure 7.3b, respectively.

Table 7.1: Base configuration for DARMA presented in Chapter 3. These parameters can be heavily modified to suit different application requirements.

Parameter	*Description*	*Value*	*Unit*
R_{max}	Maximum Range	1,000	m
P_t	Transmit Power	30	dBm
f_c	Center Frequency	10,500	MHz
λ	Wavelength	28.57	mm
B_{RF}	RF Bandwidth	3	MHz
t_{ramp}	Ramp Time	10	ms
G_t	Transmit Gain	10	dBi
G_r	Receive Gain	10	dBi
-	Polarization	Dual-Pol	-

One of the major advantages of this radar are its modular subsystems and high software/firmware configurability which can be used to vastly modify the radar's performance and tailor it for use in very specific environments. Figure 7.4 shows some of the high-level swappable subsystems. The antenna is entirely detachable and can be swapped for higher gain antennas which is very useful given the low gain of the small antenna used in this thesis (Figure 7.5). It also possesses the ability to completely bypass the phased array architecture and use a high-gain parabolic reflector antenna with only one of the inputs and outputs. This would be useful for achieving further maximum ranges or higher signal-to-noise ratios at the same range. The RF bandwidth in the base configuration allows for a range resolution, $R_{res} = 50$m which would be too large for shorter range sensing applications. Due to the use of a frequency synthesizer specifically designed with FMCW radar in mind, it is very simple to modify this RF bandwidth to achieve much smaller range resolutions. As an example of the possible configuration, the phaser radar used for signal processing uses the same synthesizer configured to transmit a bandwidth, $B_{RF} = 500$MHz which corresponds to $R_{res} = 0.3$m.

The fixed gain amplification subsystem (shown as a 3D model in Figure 7.3c) is another which allows for easy swapping depending on application requirements. It consists of four inputs and

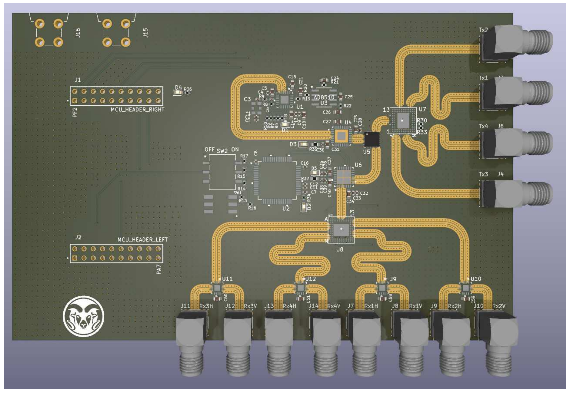

(a) ADAR1000-based FMCW radar PCB.

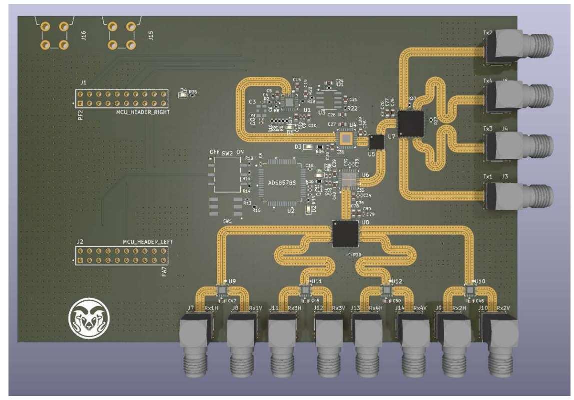

(b) AWS-0101-based FMCW radar PCB.

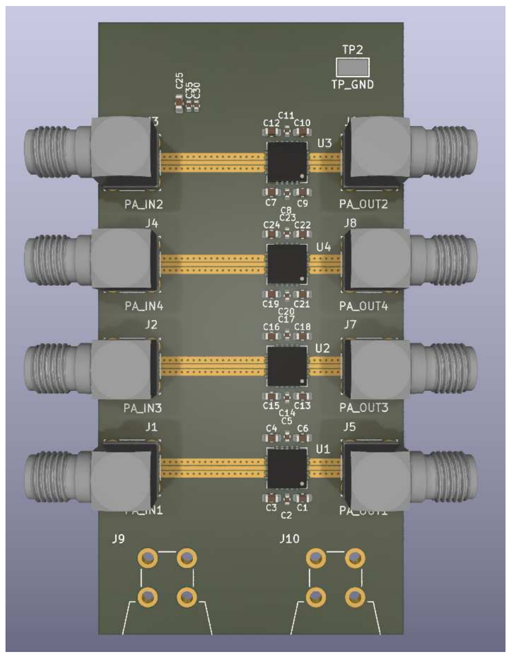

(c) Fixed gain power amplifier breakout board used to increase output power to 30dBm. This simple PCB can be redesigned and swapped out for different power requirements with ease.

Figure 7.3: 3D models of the two FMCW radars using the ADAR1000 and AWS-0101 beamforming ICs and the optional fixed gain breakout board used to increase output power to 30dBm.

outputs and each path includes a high power amplifier with supporting power supplies. Replacing the breakout board in this thesis with a custom one will mostly affect the range and signal-to-noise ratio but also has added complications of system heat dissipation and cost which should be considered if using even higher power outputs.

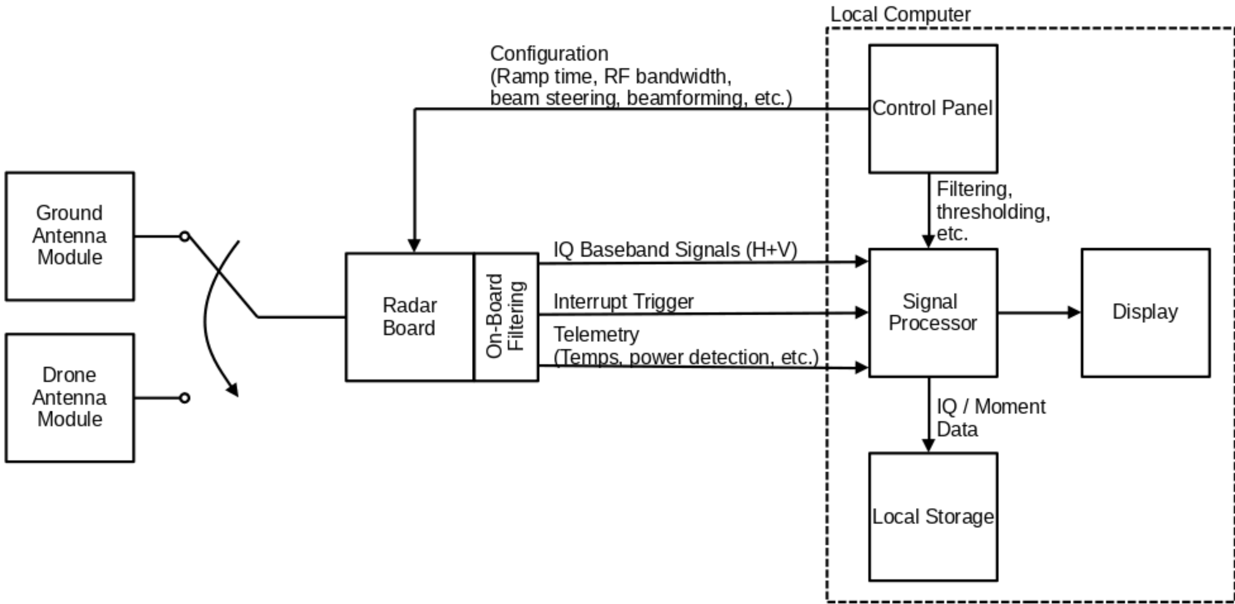

Figure 7.4: System diagram of the FMCW radar presented in Chapter 3. This shows the high-level subsystems (antenna, local computer, etc.) capable of being swapped for different applications and the basic control and information flow of the radar.

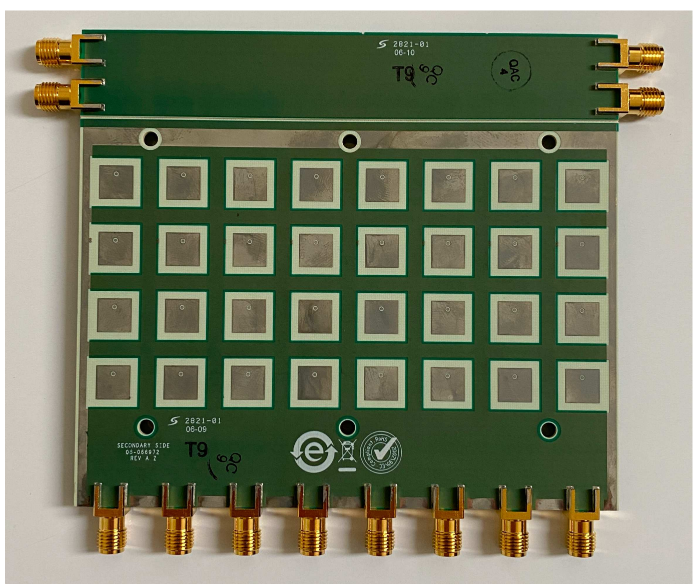

Figure 7.5: Phased array antenna used for both transmit and receive for the radar in Chapter 4 and for the simulations of the radar in Chapter 3.

The work in this thesis provides a comprehensive discussion of the critical design, cost, and deployment considerations taken when designing FMCW radars in general but is tailored towards multi-purpose FMCW radar with the ability to operate in mobile systems. Two designs are presented for a radar with the purpose of being deployed on a UAS as well as a ground platform. The two designs use separate beamforming options to compare available ICs and one, the ADAR1000 version, was fabricated and assembled. Signal processing techniques useful for all FMCW radars were also discussed and implemented. This thesis can serve as a reference for the design and implementation of multi-purpose FMCW radars.

7.2 Future Work

Future work for DARMA is suggested in this section with the intent of allowing the radar to operate at its full capacity of characterizing storms within a fully mobile system. Once realized, this will be a very useful tool for understanding meteorological events.

This thesis provides the first realization of the low-power, portable FMCW radar system and additions and improvements are envisioned for the future. Modifications should be relatively simple since the system follows a common FMCW radar block diagram and some subsystems are completely swappable. This customizability is one of the radar's main features and should be maintained through later iterations.

Necessary signal processing methods and algorithms were presented in this thesis but there is more to be done for DARMA to be deployed and to function as a weather radar mounted on a mobile system. All processing was implemented on a similar FMCW radar but, due to time constraints, was not implemented on the radar designed in Chapter 3 which will be a necessity in the future. Filtering will need to be employed on the microcontroller following the digitization of the measurement, after which it will be sent to a computer for further, more intensive processing. All the algorithms described in Chapter 4 will be applicable to DARMA and in some cases the code realization will be able to be recycled with some slight tweaking to support the different parameters such as RF bandwidth, ramp time, etc. To ensure optimal operation of the radar, calibration for determining phase offset of the antenna as well as other components will be an integral addition.

Significant improvement to DARMA's performance can be seen with upgrades to the system's antenna. This thesis describes general concepts for alternatives such as bypassing the phased array architecture for a parabolic antenna which could achieve high gains. It would therefore be useful to provide a comparison between using this configuration and a parabolic antenna with a mechanical positioner and detail the important trade-offs between the two.

Since the main purpose of this radar is the measurement of precipitation events, it will be necessary to design or obtain a radome to protect the radar from the precipitation. While adding a radome solves the issue of damaging electronics, it introduces the issue of attenuation caused by build-up of water on the radome through which the radar's signal must pass. There are multiple options for mitigation of precipitation build-up (one of which is described and employed in [23]) which should be considered, tested, and reported.

Implementing single and dual-polarization products will be critical to this radar's success as a weather radar. As discussed in Section 4.8 and Section 4.9, this information is invaluable for the classification of precipitation events. Since the radar will be in motion, additional care will be needed to ensure measurements are not misinterpreted giving incorrect classifications. One possible solution which should be tested is attaching the radar to the drone via a stabilization gimbal, a tool commonly used for cameras. Post-processing will most likely also be necessary to correct for the drone's movement. For this it will be helpful to include other sensors such as accelerometers, gyroscopes, etc. on the drone which can be used for calibration. Since there is little work on a fully-mobile, drone-based radar, extensive testing will be important for understanding the additional complications this new environment introduces.

Finally, a GUI should be created to interact with DARMA. It will provide controls with preset scan profiles as well as more detailed control of each subsystem. It will also serve as a display for the measurements taken. This will make DARMA more accessible to individuals with little knowledge of radar hardware and programming experience.

Bibliography

[1] Manuel Chavez. Antialiasing filter circuit design for single-ended adc input using fixed cutoff frequency. Technical report, Analog Devices Inc., 2019.

[2] Ian Collins. Phase-locked loop (pll) fundamentals. Technical report, Analog Devices Inc., 2018.

[3] Ken Gentile and Rick Cushing. A technical tutorial on digital signal synthesis. Technical report, Analog Devices Inc., 1999.

[4] Jordi Figueras i Ventura and Herman Russchenberg. Improvement of the performance of fm-cw radar systems by using direct digital synthesizers: Comparison with voltage controlled oscillators. In *2006 International Radar Symposium*, pages 1–4, 2006.

[5] Peter Delos, Bob Broughton, and Jon Kraft. Phased array antenna patterns—part 1: Linear array beam characteristics and array factor. *Analog Dialogue*, 54:1–8, 2020.

[6] Gregory L. Charvat. *Small and Short-Range Radar Systems*, volume 1. CRC Press, 2017.

[7] Alexander Morin. System design and development of front-x: An x-band dual-polarization phased array weather radar. Master's thesis, Colorado State University, 2019.

[8] V.N. Bringi and V. Chandrasekar. *Polarimetric Doppler Weather Radar*, volume 1. Cambridge University Press, 2001.

[9] Oleg A. Krasnov, Leo P. Ligthart, Zhijian Li, Pidio Lys, and Fred van der Zwan. The parsax - full polarimetric fmcw radar with dual-orthogonal signals. In *2008 European Radar Conference*, pages 84–87, 2008.

[10] Zachariah Peterson. Stripline vs microstrip pcb routing differences and guidelines. Webpage, September 2017.

[11] John Coonrod. Comparing microstrip and grounded coplanar waveguides. https://rogerscorp.com/blog/2020/comparing-microstrip-and-grounded-coplanar-waveguides, July 2020.

[12] R.J.P. Douville and D.S. James. Experimental study of symmetric microstrip bends and their compensation. *IEEE Transactions on Microwave Theory and Techniques*, 26(3):175–182, 1978.

[13] Alan V. Oppenheim and Ronald W. Schafer. *Discrete-Time Signal Processing*, volume 1. Prentice Hall, 1999.

[14] Edwin Kessler. Note on range normalization of digitized radar data. *Journal of Applied Meteorology (1962-1982)*, 2(1):178–180, 1963.

[15] Andy Harrison. *Introduction to Radar Using Python and MATLAB*, volume 1. Artech House, 2020.

[16] J. R. Probert-Jones. The radar equation in meteorology. *Quarterly Journal of the Royal Meteorological Society*, 88(378):485–495, 1962.

[17] David Brunkow, V. N. Bringi, Patrick C. Kennedy, Steven A. Rutledge, V. Chandrasekar, E. A. Mueller, and Robert K. Bowie. A description of the csu–chill national radar facility. *Journal of Atmospheric and Oceanic Technology*, 17(12):1596 – 1608, 2000.

[18] Guochao Wang, José-María Muñoz-Ferreras, Roberto Gómez-García, and Changzhi Li. Clutter interference reduction in coherent fmcw radar for weak physiological signal detection. In *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*, pages 1–4, 2014.

[19] Kenneth V. Beard and Catherine Chuang. A new model for the equilibrium shape of raindrops. *Journal of Atmospheric Sciences*, 44(11):1509 – 1524, 1987.

[20] National Oceanic and Atmospheric Association (NOAA). Dual-polarization radar training for nws partners, 2013.

[21] Alan D. Siggia and Richard E. Passarelli. Gaussian model adaptive processing (gmap) for improved ground clutter cancellation and moment calculation. 2004.

[22] Ivan Arias, V. Chandrasekar, and Shashank S. Joshil. Cross-validation of csu-chivo radar and gpm during relampago. In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 7586–7589, 2019.

[23] Paul E. Johnston, James R. Jordan, Allen B. White, David A. Carter, David M. Costa, and Thomas E. Ayers. The noaa fm-cw snow-level radar. *Journal of Atmospheric and Oceanic Technology*, pages 249–267, 2017.

[24] V. Chandrasekar, Arthur Hou, Eric Smith, V. N. Bringi, S. A. Rutledge, E. Gorgucci, W. A. Petersen, and Gail Skofronick Jackson. Potential role of dual-polarization radar in the validation of satellite precipitation measurements: Rationale and opportunities. *Bulletin of the American Meteorological Society*, 89(8):1127 – 1146, 2008.

Appendix A

Schematic and Subsystem Overview

Chapter 3 provides a detailed description of the considerations taken when designing the subsystems of the FMCW radar. A brief subsystem overview is presented in Table A.1. For reproducibility, schematics of all subsystems for both the ADAR1000 and AWS-0101-based radar boards are presented in Section A.1 and Section A.2, respectively. The boards are split into four main sections: transmit, receive, control, and power. The transmit section contains the signal generation, power splitting, transmit beamforming, and antenna subsystems and the receive section contains the polarization switch, receive beamforming, baseband mixing, data acquisition, and antenna subsystems. Digital control and power for these subsystems are shown separately. The fixed gain amplification breakout board schematics are shown in Section A.3 split into the RF and power subsystems.

Table A.1: Description of the subsystems for an FMCW radar introduced in Chapter 3.

Subsystem	*Description*
Signal Generation	Responsible for generating the linear frequency-modulated (LFM) signal by use of a PLL. Frequency modulation can take multiple waveforms; this frequency synthesizer (ADF4159) is capable of generating sawtooth and triangular modulation patterns.
Transmit Beamforming	Splits signal from the signal generation subsystem into four separate signals and individually applies phase weights for beam-steering and gain weights for beamforming. Beamforming is used primarily in this radar for tapering which reduces side-lobes in the antenna pattern and makes steering to outer angles possible without interference from the side-lobes.
Receive Beamforming	Phase weighting and variable gain amplification is applied to each of the four signals received from the transmit antenna which allows for beam-steering and beamforming. The weighted signals are then combined and output as one.
Baseband Mixing	Takes inputs from the signal generation and receive beamforming subsystems and mixes them together, effectively subtracting the signals' respective frequencies and generating a baseband signal. A 90° phase shift is then applied to half the signal's power to create two outputs: in-phase and quadrature (IQ).
Fixed Gain Amplification	Following the transmit beamforming, additional fixed gain is applied to each of the signals to achieve the desired output power, $P_t = 30$dBm. This is realized as a breakout board.
Data Acquisition	IQ signals from the baseband mixer have an anti-aliasing filter applied and are impedance-matched to the ADC which digitizes them for further processing.
Control	A swappable microcontroller header allows for control of each of the subsystems and collecting of the digitized data from the ADC. Available control includes: setting the LFM ramp profile, beamforming phase and gain weights, oversampling the ADC, and controlling the polarization switch.
Polarization Switch	Switches between the two available receive polarizations, each of which have four antenna inputs.
Antenna	Separate antennas are used for transmit and receive. Transmit uses a 4x4 microstrip patch array with beam-steering and beamforming available in one dimension. Receive is the same except with the availability for use in two polarizations.
Signal Processor	Digitized IQ signals from the ADC have simple filtering applied via the microcontroller then are sent to a local computer for further, more complex processing. This includes, but is not limited to: range binning, range normalization, clutter filtering, target detection, and radar product calculation.

A.1 ADAR1000-based Radar

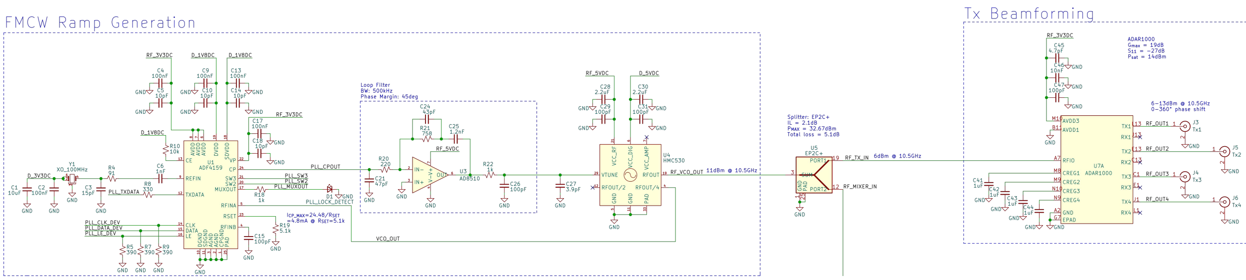

Figure A.1: Transmit side schematic for the ADAR1000-based radar including the signal generation, transmit beamforming, antenna, and power splitting subsystems. Pins for the ADAR1000 beamforming subsystem are not complete; some are connected in the control subsystem shown in Figure A.3. This is identical to the AWS-0101-based radar except for the use of a different beamforming IC.

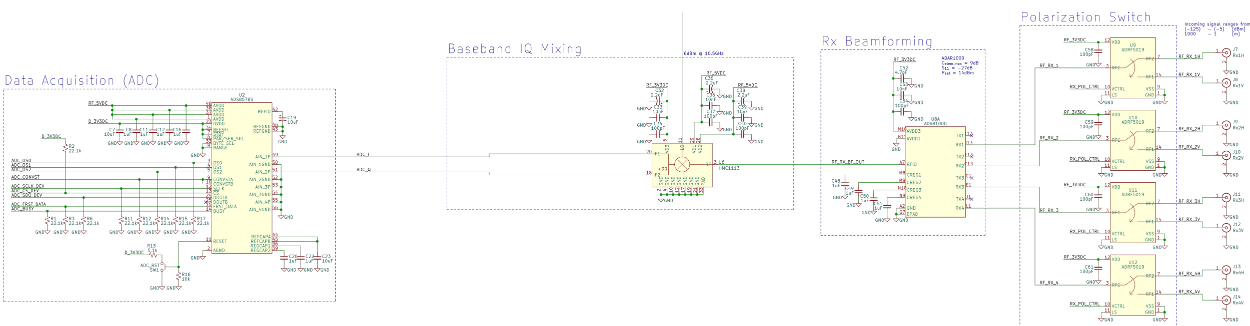

Figure A.2: Receive side schematic for the ADAR1000-based radar including the polarization switch, receive beamforming, baseband mixing, data acquisition, and antenna subsystems. Pins for the ADAR1000 beamforming subsystem are not complete; some are connected in the control subsystem shown in Figure A.3. The baseband mixer's outputs IF1 and IF2 correspond to the in-phase and quadrature baseband signals, respectively. The ADC (ADS8578) includes internal anti-aliasing filtering and impedance matching which explains the lack of external circuitry. This is identical to the AWS-0101-based radar except for the use of a different beamforming IC.

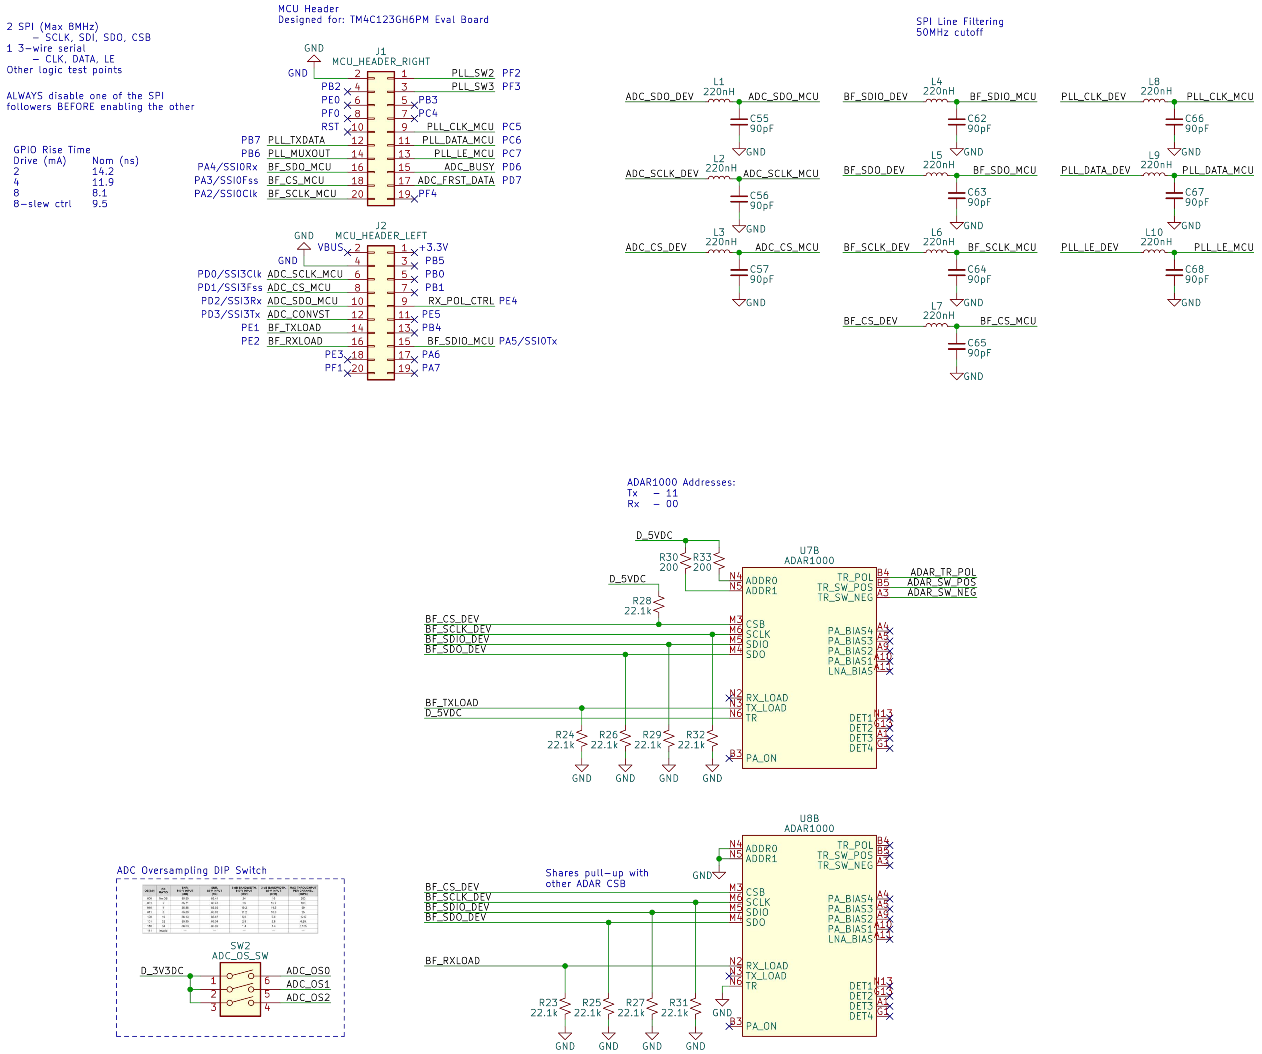

Figure A.3: Control subsystem schematic for the ground-based FMCW radar board. This subsystem is specific to the ADAR1000-based radar but the general functionality remains the same as in the AWS-0101-based radar.

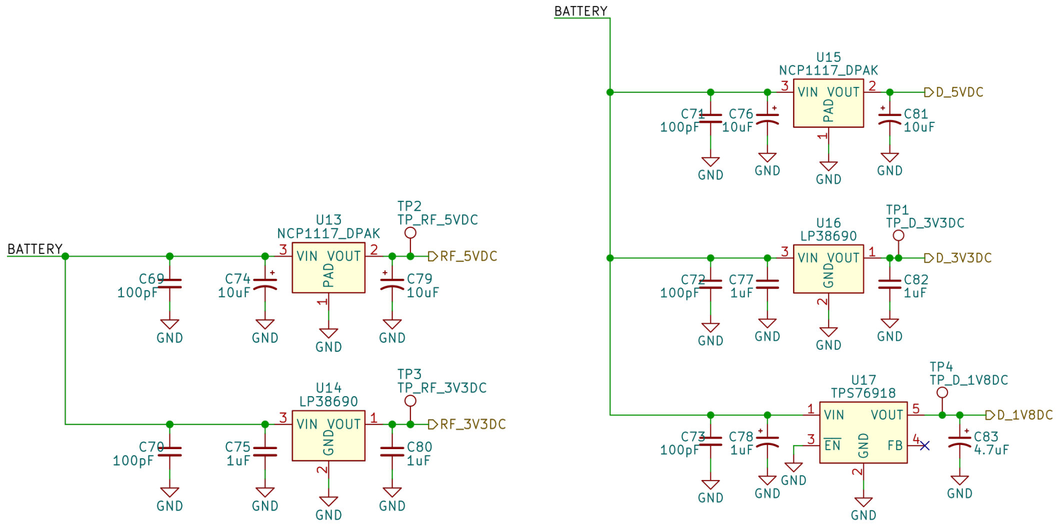

(a) RF power supplies at $V_{out} = 3.3,\ 5$V used to supply the frequency synthesizer, loop filter amplifier, VCO, beamformer, mixer, polarization switch, and ADC.

(b) Digital power supplies at $V_{out} = 1.8,\ 3.3,\ 5$V used to supply the frequency synthesizer, VCO, beamformer, and ADC.

Figure A.4: Power subsystem schematic for the ground-based FMCW radar board. This subsystem contains power supplies at voltages specific to the ADAR1000-based radar but the general functionality remains the same as in the AWS-0101-based radar.

A.2 AWS-0101-based Radar

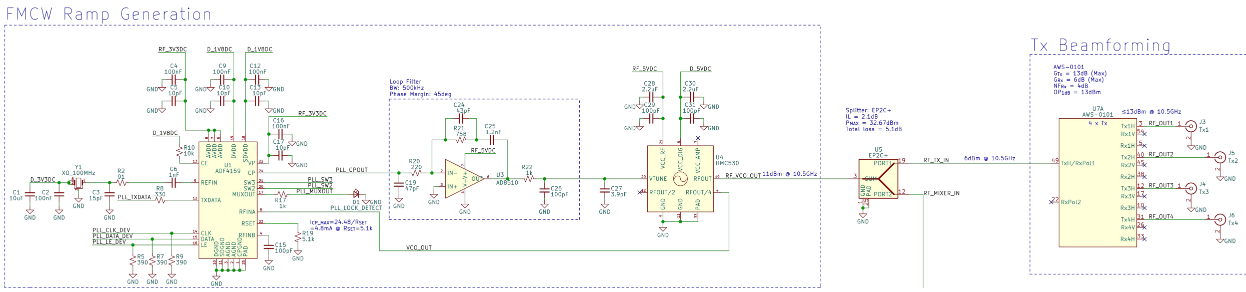

Figure A.5: Transmit side schematic for the AWS-0101-based radar including the signal generation, transmit beamforming, antenna, and power splitting subsystems. Pins for the AWS-0101 beamforming subsystem are not complete; some are connected in the control subsystem shown in Figure A.7 and the power subsystem shown in Figure A.8c. This is identical to the ADAR1000-based radar except for the use of a different beamforming IC.

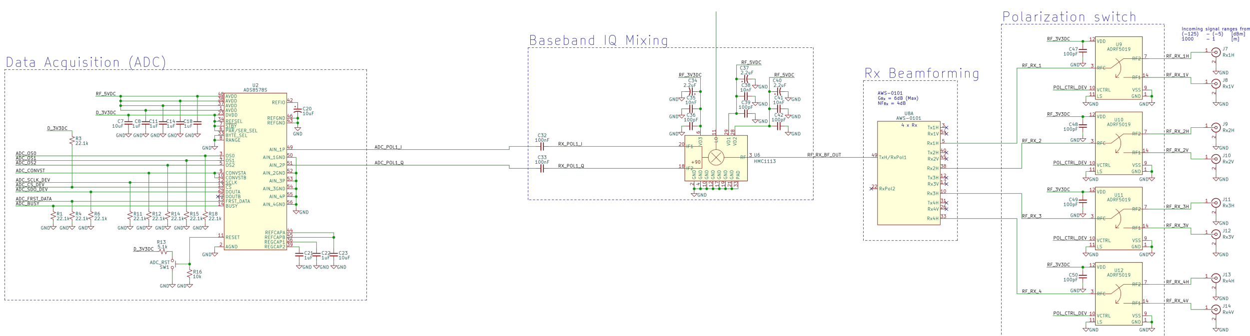

Figure A.6: Receive side schematic for the AWS-0101-based radar including the polarization switch, receive beamforming, baseband mixing, data acquisition, and antenna subsystems. Pins for the AWS-0101 beamforming subsystem are not complete; some are connected in the control subsystem shown in Figure A.7 and the power subsystem shown in Figure A.8c. The baseband mixer's outputs IF1 and IF2 correspond to the in-phase and quadrature baseband signals, respectively. The ADC (ADS8578) includes internal anti-aliasing filtering and impedance matching which explains the lack of external circuitry. This is identical to the ADAR1000-based radar except for the use of a different beamforming IC.

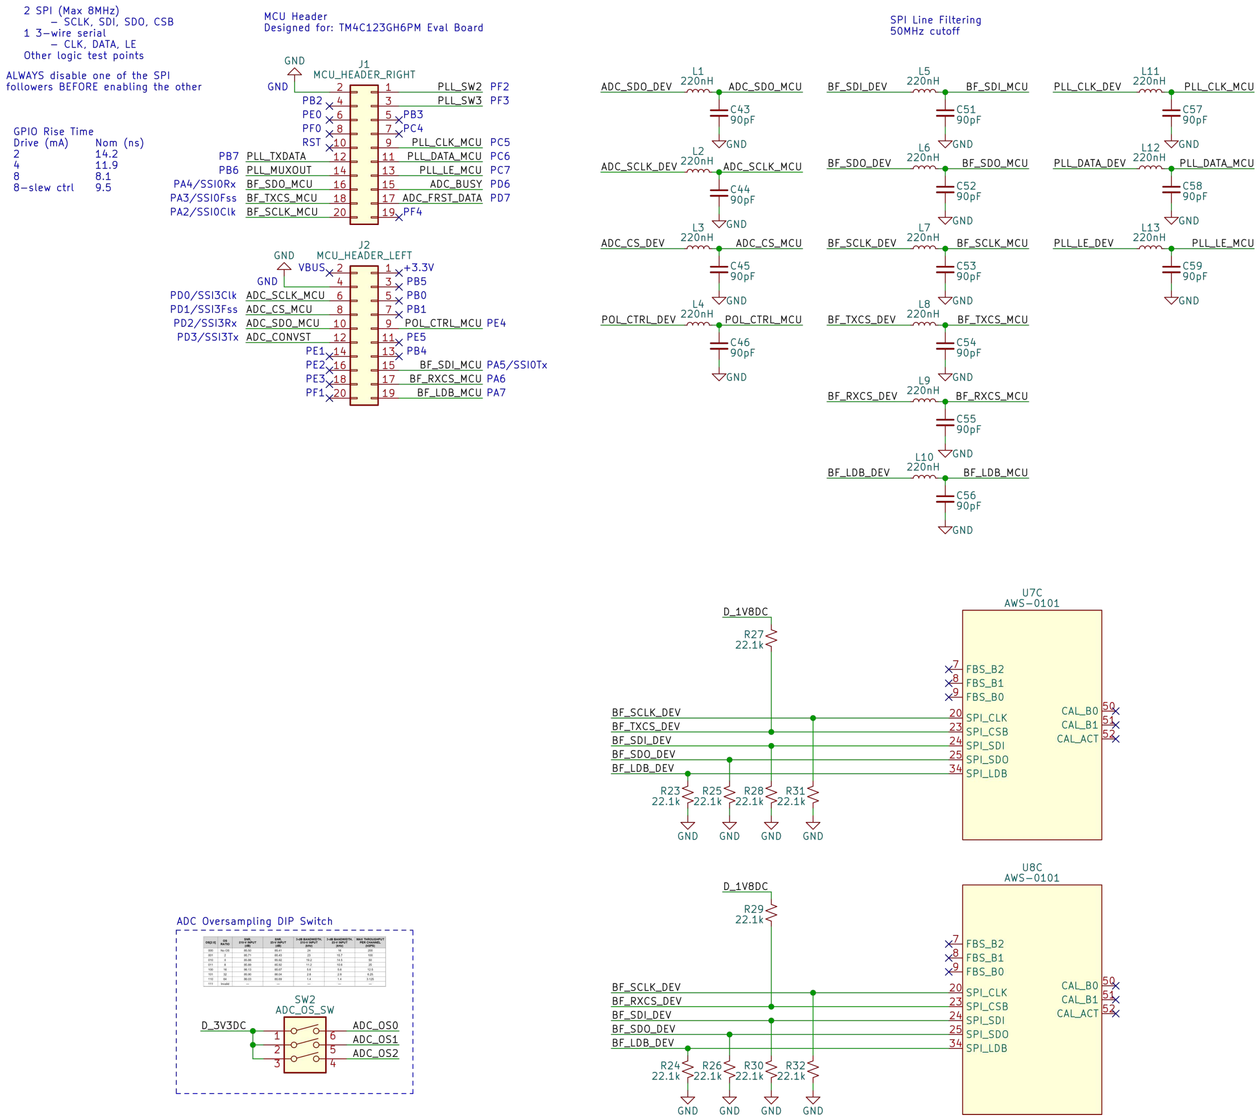

Figure A.7: Control subsystem schematic for the ground-based FMCW radar board. This subsystem is specific to the AWS-0101-based radar but the general functionality remains the same as in the AWS-0101-based radar.

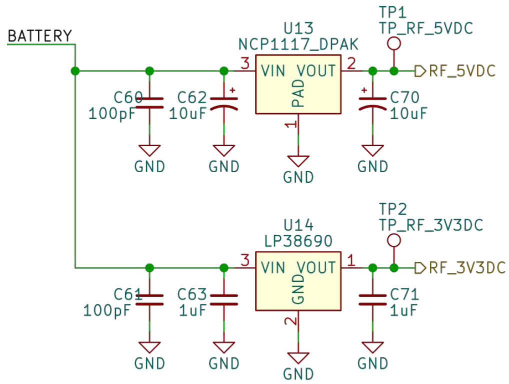

(a) RF power supplies at $V_{out} = 3.3,\ 5$V used to supply the frequency synthesizer, loop filter amplifier, VCO, mixer, polarization switch, and ADC.

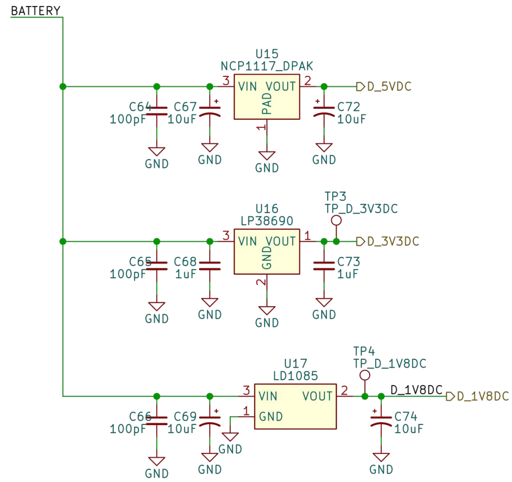

(b) Digital power supplies at $V_{out} = 1.8,\ 3.3,\ 5$V used to supply the frequency synthesizer, VCO, beamformer, and ADC.

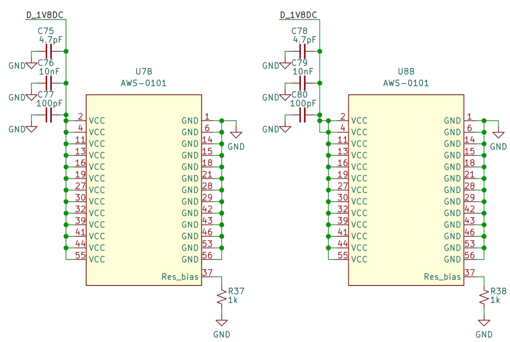

(c) Power supply connections to the AWS-0101 beamformer ICs.

Figure A.8: Power subsystem schematic for the ground-based FMCW radar board. This subsystem contains power supplies at voltages specific to the AWS-0101-based radar but the general functionality remains the same as in the ADAR1000-based radar.

A.3 Fixed Gain Amplification Breakout Board

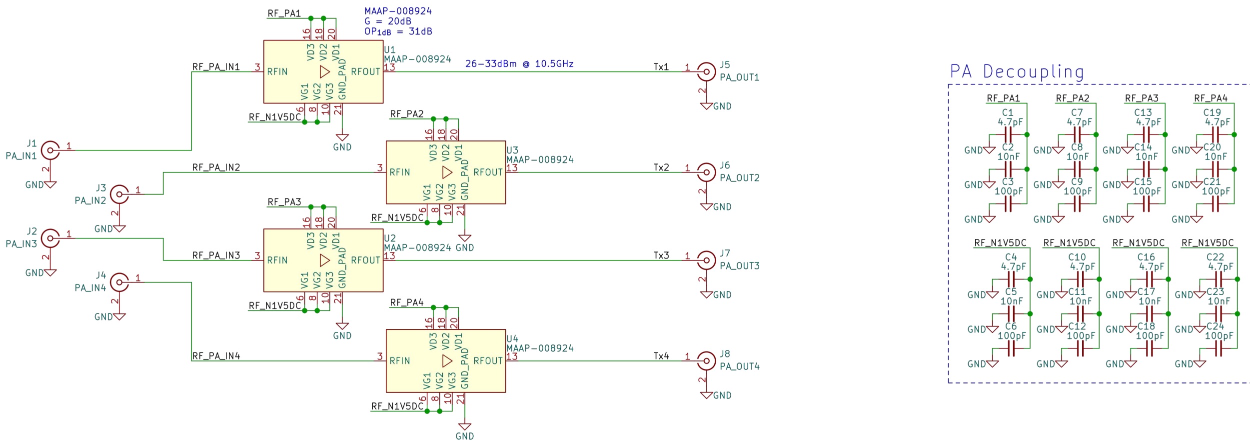

Figure A.9: Fixed gain amplification breakout board schematic including RF components, exterior connectors, and power supply decoupling.

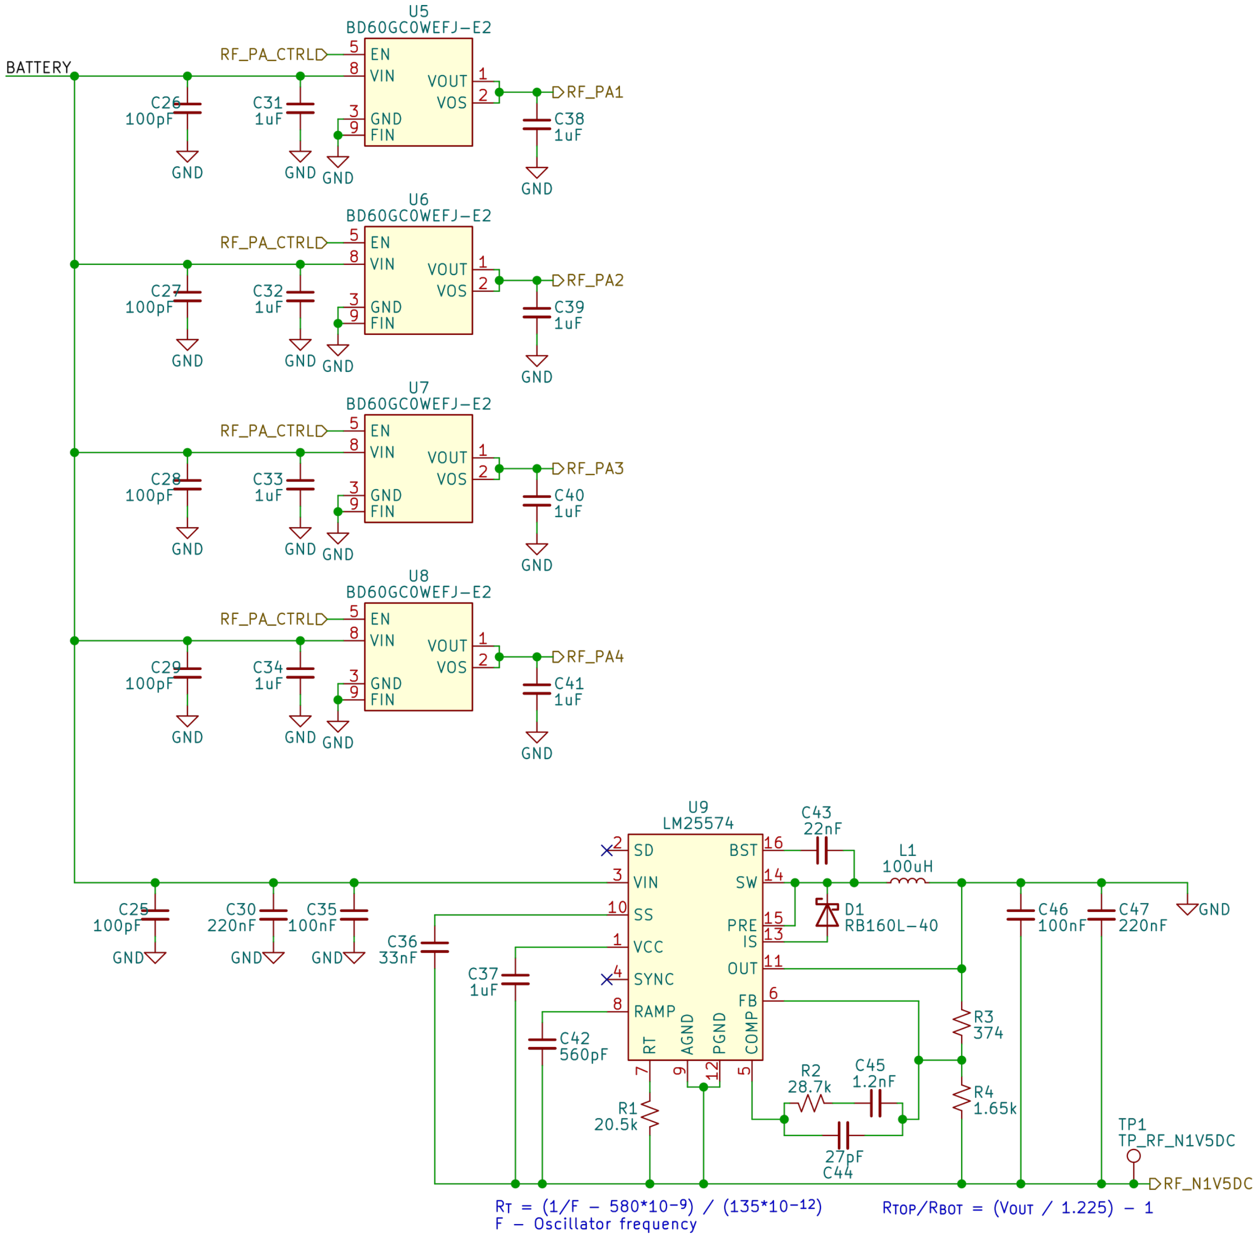

Figure A.10: Power subsystem schematic for the fixed gain amplification breakout board. This contains RF power supplies at $V_{out} = 6, -1.5$V used to supply the high power amplifiers.

Appendix B

Cost Breakdown

This appendix contains a detailed breakdown of the cost for each of the three boards including components, fabrication, and assembly as well as totals depending on fabrication and assembly time. Chapter 3 discusses offloading computation to an external computer but does not specify any requirements for the computer because the type needed can be application-specific. All signal processing and display for the phaser radar in Chapter 4 is capable of being done on a Raspberry Pi which is a cheap embedded Linux computer. Plotting speeds suffer slightly due to the low-power of the Pi so most figures and real-time displays were generated on more powerful Windows and Linux desktops and laptops. Both the desktop and laptop were fully capable of generating real-time plots while simultaneously collecting data and processing it in Python (their specifications are listed in Table B.1).

Though two FMCW radar boards and a fixed gain amplification breakout board was designed, only the ADAR1000-based radar was fabricated and assembled due to the reasons outlined in Chapter 6. For a prototyping quantity of 2 boards and fabrication and assembly times of 20 and 5 days, respectively, the total cost was approximately \$3,705.52 per board (Table B.7), or \$7,411.04. It should be noted, however, that prices for fabrication and assembly decrease significantly with increasing order volume and component costs decrease slightly. For example, in a production environment with order quantities greater than 150, the fabrication cost decreases by a factor of 10, or to ~\$200 for a 20-day turn-time. Some of the more specialized and expensive ICs ($\gtrapprox$ \$50) also decrease in cost significantly which, in production, will greatly affect the total price.

There is a large cost not considered in this section, the non-recurring engineering (NRE) costs which include research, design, development, and testing. NRE costs are highly dependent on numerous factors such as engineer experience, available tools, team size, and much more. Since these factors vary widely between organization, it would be irresponsible to present an NRE price estimate.

B.1 Components

Table B.1: Important components and specifications for the two computers used for signal processing.

Specification	Desktop	Laptop	*Unit*
CPU	Intel i5-6600k	Intel i7-7500u	-
CPU Clock Speed	3.5	2.7	GHz
GPU	Nvidia GTX960	Nvidia 940MX	-
Memory	8	16	GB
Memory Speed	2133	2133	MHz

Table B.2: Individual component costs for the ADAR1000-based radar.

Type	*MPN*[2]	*Manufacturer*	*Cost* [$]	*Quantity*
Oscillator	AOCJY1-100.000MHZ	Abracon	62.25	1
PLL	ADF4159	Analog Devices	16.67	1
Loop Filter Amplifier	AD8510	Analog Devices	6.66	1
VCO	HMC530	Analog Devices	38.26	1
Power Splitter	EP2C+	Mini-Circuits	8.95	1
Beamformer	ADAR1000	Analog Devices	360.00	2
Baseband Mixer	HMC1113	Analog Devices	60.18	1
Polarization Switch	ADRF5019	Analog Devices	13.51	4
ADC	ADS8578S	Texas Instruments	24.88	1
DIP Switch	219-3MST	CTS Electrocomponents	0.74	1
MCU Header	TSW-110-23-S-D	Samtec	2.33	2
SMA Connector	132136-12	Amphenol RF	17.07	12
RF 3.3V Supply	LP38690DTX-3.3	Texas Instruments	2.03	1
RF 5V Supply	NCP1117IDT50T4G	onSemi	0.67	1
Digital 5V Supply	NCP1117IDT50T4G	onSemi	0.67	1
Digital 3.3V Supply	LP38690DTX-3.3	Texas Instruments	2.03	1
Digital 1.8V Supply	TPS76918	Texas Instruments	1.32	1
Passive Components (Resistors, Capacitors, Inductors, LEDs)			66.47	-
		Total	**1275.32**	

[2]Manufacturer part number.

Table B.3: Individual component costs for the AWS-0101-based radar.

Type	*MPN*	*Manufacturer*	*Cost* [$]	*Quantity*
Oscillator	AOCJY1-100.000MHZ	Abracon	62.25	1
PLL	ADF4159	Analog Devices	16.67	1
Loop Filter Amplifier	AD8510	Analog Devices	6.66	1
VCO	HMC530	Analog Devices	38.26	1
Power Splitter	EP2C+	Mini-Circuits	8.95	1
Beamformer	AWS-0101	AnokiWave	186.42	2
Baseband Mixer	HMC1113	Analog Devices	60.18	1
Polarization Switch	ADRF5019	Analog Devices	13.51	4
ADC	ADS8578S	Texas Instruments	24.88	1
DIP Switch	219-3MST	CTS Electrocomponents	0.74	1
MCU Header	TSW-110-23-S-D	Samtec	2.33	2
SMA Connector	132136-12	Amphenol RF	17.07	12
RF 3.3V Supply	LP38690DTX-3.3	Texas Instruments	2.03	1
RF 5V Supply	NCP1117IDT50T4G	onSemi	0.67	1
Digital 5V Supply	NCP1117IDT50T4G	onSemi	0.67	1
Digital 3.3V Supply	LP38690DTX-3.3	Texas Instruments	2.03	1
Digital 1.8V Supply	LD1085	ST Microwave	1.89	1
Passive Components (Resistors, Capacitors, Inductors, LEDs)			66.47	-
		Total	**866.48**	

Table B.4: Individual component cost for the fixed gain amplification breakout board.

Type	*MPN*	*Manufacturer*	*Cost* [$]	*Quantity*
Power Amplifier	MAAP-00896	Macom	39.71	4
SMA Connector	132136-12	Amphenol RF	17.07	8
RF 6V Supply	BD60GC0WEFJ-E2	Rohm Semiconductor	0.81	4
RF -1.5V Supply	LM25574MT	Texas Instruments	4.40	1
Passive Components (Resistors, Capacitors, Inductors, LEDs)			24.09	-
		Total	**327.13**	

B.2 Fabrication and Assembly

Since the radar was built as a prototype, only two boards were ordered which is a large contributing factor in the price. Beyond two boards, fabrication and assembly became much cheaper per unit.

Table B.5: Costs and turn-times for fabrication of the two radar boards and the fixed gain amplification breakout board. Price was quoted for a quantity of two boards which was optimal for prototyping; price per board decreases significantly as quantity increases.

	Fabrication Cost **[$]**		
Turn-time [Days]	*ADAR1000*	*AWS-0101*	*Fixed Gain Breakout*
10	3,067.79	3,067.79	2,972.93
15	2,801.02	2,801.02	2,714.41
20	2,667.64	2,667.64	2,585.16

Table B.6: Costs and turn-times for assembly of the two radar boards and the fixed gain amplification breakout board. Price was quoted for a quantity of two boards which was optimal for prototyping; price per board decreases significantly as quantity increases.

	Assembly Cost **[$]**		
Turn-time [Days]	*ADAR1000*	*AWS-0101*	*Fixed Gain Breakout*
1	-	-	947.10
2	1,574.71	1,408.91	744.15
3	1,385.74	1,239.84	654.85
4	1,216.82	1,088.70	575.02
5	1,037.88	928.60	490.46
7	930.51	832.53	439.72
10	758.72	678.84	358.54
15	658.51	589.18	311.19
20	629.88	563.56	297.66

B.3 Total Cost

Table B.7: Total cost [$] per board for components, fabrication, and assembly for the ADAR1000-based FMCW radar board as a function of fabrication and assembly time [days]. Price was quoted for a quantity of two boards which was optimal for prototyping; price per board decreases significantly as quantity increases.

		Fabrication		
	Days	10	15	20
Assembly	1	-	-	-
	2	5,917.82	5,651.05	5,517.67
	3	5,728.85	5,462.08	5,328.70
	4	5,559.93	5,293.16	5,159.78
	5	5,380.99	5,114.22	4,980.84
	7	5,273.62	5,006.85	4,873.47
	10	5,101.83	4,835.06	4,701.68
	15	5,001.62	4,734.85	4,601.47
	20	4,972.99	4,706.22	4,572.84

Table B.8: Total cost [$] per board for components, fabrication, and assembly for the AWS-0101-based FMCW radar board as a function of fabrication and assembly time [days]. Price was quoted for a quantity of two boards which was optimal for prototyping; price per board decreases significantly as quantity increases.

		Fabrication		
	Days	10	15	20
Assembly	1	-	-	4,481.22
	2	5,343.18	5,076.41	4,278.27
	3	5,174.11	4,907.34	4,188.97
	4	5,022.97	4,756.20	4,109.14
	5	4,862.87	4,596.10	4,024.58
	7	4,766.80	4,500.03	3,973.84
	10	4,613.11	4,346.34	3,892.66
	15	4,523.45	4,256.68	3,845.31
	20	4,497.83	4,231.06	3,831.78

Table B.9: Total cost [$] per board for components, fabrication, and assembly for the fixed gain amplification breakout board as a function of fabrication and assembly time [days]. Price was quoted for a quantity of two boards which was optimal for prototyping; price per board decreases significantly as quantity increases.

		Fabrication		
	Days	10	15	20
Assembly	1	4,342.02	4,075.25	3,941.87
	2	4,139.07	3,872.30	3,738.92
	3	4,049.77	3,783.00	3,649.62
	4	3,969.94	3,703.17	3,569.79
	5	3,885.38	3,618.61	3,485.23
	7	3,834.64	3,567.87	3,434.49
	10	3,753.46	3,486.69	3,353.31
	15	3,706.11	3,439.34	3,305.96
	20	3,692.58	3,425.81	3,292.43