

UNIVERSITY OF CAPE TOWN

DEPARTMENT OF ELECTRICAL ENGINEERING

Modelling and Resolving of the Ambiguous Angle of Arrival
Measurements of the SANAE IV SuperDARN Radar



Author:
Tighe BARRIS

Supervisor:
Dr. O'HAGAN

Submitted to the Department of Electrical Engineering at the University of Cape Town
in fulfilment of the academic requirements for a MSc in Radar and Electronic Defence

October 25, 2017

keywords: Angle Of Arrival, Interferometry, SANAE IV Calibration

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Plagiarism Declaration

I know the meaning of plagiarism and declare that all the work in the document, save for that which is properly acknowledged, is my own. This thesis/dissertation has been submitted to the Turnitin module (or equivalent similarity and originality checking software) and I confirm that my supervisor has seen my report and any concerns revealed by such have been resolved with my supervisor.

Tighe Barris - BRRTIG001

October 25, 2017

Abstract

The Super Dual Aural Radar Network (SuperDARN) forms a multinational collaboration consisting of over 30 HF radars used to monitor charged particle convections in the high latitudes of the ionosphere [1]. This research is inspired by the antenna arrangement of one of these HF radars situated at the South African National Antarctic Expedition (SANAE) IV base.

This radar consists of a steered 16-element primary array, and a second 4-element interferometer array. Phase wrapping is introduced into echoed returns due to the two arrays being separated by 100 m, several wavelengths of the carrier frequency. As a result, the radar system is only able to unambiguously determine the Angle of Arrival (AOA) for a given target provided it resides below 33° in elevation when observing along boresight [2].

This dissertation models a subcomponent of this phenomenon and develops an entirely separate pulsed-Doppler radar system on an FPGA platform known as the Red Pitaya. Operating parameters are specifically chosen to be consistent with the SANAE IV radar. An antenna arrangement consisting of two receivers and a transmitter is constructed to replicate the AOA issue experienced by the SANAE radar. Instead of focusing on detecting targets in elevation, as done with the SuperDARN radars, targets are detected in azimuth.

Three algorithms are developed to autonomously measure the AOA and to counteract the effects of 2π phase wrapping. The naïve measurement applies no corrections and is only able to detect targets at angles within $\pm 17.10^\circ$ of boresight. This is the theoretical maximum measurable AOA when operating at 17 MHz. A second algorithm involving a Standard Deviation (STDEV) search process is used to successfully measure angles up to $\pm 70^\circ$ for the same arrangement. It does this by varying the carrier frequency and using numerous received aliased AOAs to correct and expand for a true AOA. Empirical results yield an average Mean Square Error (MSE) of 0.29° between the true and measured target angles. A third algorithm involving the Chinese Remainder Theorem (CRT) was shown to work in theory (albeit only up to $\pm 66^\circ$), but failed to perform in practice.

Results presented in this dissertation indicate that the maximum measurable AOA is expanded by a factor of 4 by using the STDEV algorithm. It is postulated that the same result can be achieved using the actual SuperDARN radar.

Theory relevant to generating the RF Linear Frequency Modulated (LFM) chirp and its use for transmitting, receiving and processing is covered. Topics of pulse compression, range resolution, environment modelling, reconstruction and matched filtering, coherent integration, sidelobe reduction and demodulation are detailed.

Acknowledgements

Foremost, I would like to thank my supervisor, Dr. D. O'Hagan, for his devoted and dedicated interest throughout this project. His consistent presence and insight into this work aided much of the research. I always enjoyed our discussions during meetings.

I would like to thank the South African National Space Agency (SANSa) for making funding, accommodation and access to resources available to conduct this research. Additionally, I would like to extend thanks to Professor M. Kosch and Jonathan Ward for continuously making time available to assist. It was greatly appreciated.

To Skippy and Tom, I honestly don't know how I would have constructed all the antennas alone on that ginormous school field. Thank you for all the help throughout this research and the good laughs. Your countless hours of help were certainly noted.

I extend thanks to my father, Ken, who took the time to read this dissertation and provide insightful input. One day I hope to conquer concord as well as you have.

Finally, to Klee, Indy and Anri. Thank you each for being such wonderful pillars in the Barris family and for always being available for support.

Contents

1	Introduction	1
1.1	Background	1
1.2	Problem Formulation	4
1.3	Applications for Angle of Arrival Estimations	5
1.4	Research Question	6
1.5	Scope & Limitations	6
1.6	Plan of Development	7
2	Supporting Literature	9
2.1	SuperDARN HF Propagation in the Ionosphere	9
2.2	Previous SuperDARN Calibration Attempts	12
2.2.1	SANAE IV - Helicopter Beacon Calibration Attempt	12
2.2.2	ZACUBE-1 Calibration Attempt	15
2.2.3	Buckland Park: Augmented SuperDARN Radar	17
2.3	Angle Measurements Using COTS Equipment	21
2.4	Extended Literature	26
2.5	Chapter Summary	27
3	Design of The Interferometer	28
3.1	Design Requirements	28
3.2	Radar Model	29
3.2.1	Transmitted Waveform	29

3.2.2	Received Signal of Stationary Targets	32
3.2.3	Received Signal of Moving Targets	34
3.2.4	Transceiver Chain, Reconstruction Filter & Range Resolution	37
3.2.5	Matched Filtering	41
3.2.6	Pulse Compression & The Ambiguity Function	45
3.2.7	Coherent Integration	49
3.2.8	Windowing & Sidelobe Reduction	51
3.3	Interferometric Target Detection Algorithm	55
3.3.1	Naïve AOA Algorithm	55
3.3.2	Standard Deviation AOA Algorithm	57
3.3.3	Chinese Remainder Theorem AOA Algorithm	64
3.4	The Interferometer Hardware Design	65
3.4.1	Transceiver Platform: The Red Pitaya	65
3.4.2	Programmable Logic & Processing System Communication	66
3.4.3	Storing Received Data	69
3.4.4	Sinusoidal Unit Tests	71
3.4.5	HF Demodulation Using a Hilbert Transform	72
3.4.6	HF Quadrature Demodulation	77
3.4.7	Transmit DAC Buffer Design	79
3.4.8	Write-Back & Direct-Input Modules: Features of the DAC Buffer	81
3.4.9	Daisy-Chaining The Red Pitaya	83
3.4.10	Switching the Receiver Off During Transmission	84
3.4.11	Discussion on the Hardware Flexibility	85
3.4.12	Front-end Filter Design	86
3.4.13	HF Antenna Design	87
3.5	Chapter Summary	88

4	Experiments and Results	90
4.1	Pre-Field Test Trials	90
4.1.1	General Closed-Loop Testing on the Red Pitaya	90
4.1.2	Characteristics of the 208.42 m coaxial cable	95
4.1.3	Phase delays on the 18 m cables	96
4.1.4	Synchronisation Test using Two Red Pitaya Devices	97
4.2	Field Test Trials	98
4.2.1	Background & Utilised Interferometric Geometry	98
4.2.2	Methodology	99
4.2.3	Default Operating & Test Parameters	99
4.2.4	Estimated Power Budget	101
4.2.5	Empirical Results	103
4.3	Additional Experiment: Moving Aircraft Targets	114
5	Discussion	116
5.1	Comments on the Pre-Field Tests	116
5.2	Comments on the Outdoors Interferometry Experiment	121
5.3	Comments on the Moving Targets Experiment	124
5.4	The Research Question	125
6	Summary & Conclusion	126
6.1	Portability to SANAE IV SuperDARN Radar System	127
6.2	Recommendations	128
A	Appendix	133
A.1	Source Code Access	133
A.2	Proof of Matched Filter Peak SNR Improvement	134
A.3	List of Common Window Functions	135
A.4	Internal Diagram of Zynq IP Core	136
A.5	Red Pitaya Daisy-Chained Phased Array	137

A.6	Vivado FPGA Design	138
A.7	Front-End 12.5-25 MHz Filter	139
A.8	Field Test Experiment Images	140
A.8.1	The Overall Complete FPGA Design & Instruction Set	142

List of Figures

1.1	SuperDARN coverage map [4]	2
1.2	HF wave propagation and electron height density profile in the ionosphere	2
1.3	Example ionospheric convection map over Northern Hemisphere [6]	3
2.1	Computed ray paths for SANA E radar at 11 MHz [13]	10
2.2	Geometrical visualisation of Bragg’s scattering	11
2.3	Geometry of interferometric calculations used to determine an AOA [15] .	13
2.4	SANA E Calibration - Theoretical & Measured Results [15]	14
2.5	Atmospheric effects on LEO signal propagation [19]	16
2.6	ZACUBE-1 (left) and its antenna deployment hardware (right) [17]	17
2.7	Beam patterns of the LPDA and TTFD antenna arrays	18
2.8	Relative gain of the main array with a variously located secondary array .	20
2.9	New proposed layout for Buckland Park SuperDARN radar [2]	20
2.10	Geometry of antennas used to measure the phase difference [23]	22
2.11	Layout of the COTS experiment [23]	23
2.12	Phase difference observed from receiver R and G with integrated view [23]	24
3.1	Instantaneous frequency of LFM and NLFM chirp [35]	30
3.2	PSD of $v_{TX}(t)$, $q_{TX}(t)$ and $p_{TX}(t)$ respectively	31
3.3	$v_{TX}(t)$ alongside its analytical definition, $p_{TX}(t)$, and corresponding PSD	32
3.4	Simulated analytical signal, $q_{RX}(t)$, originating from 3 different targets . .	37
3.5	Radar transmit and receive overview [35]	38

3.6	Spectral illustration of the process used to image an environment [35] . . .	39
3.7	Expected baseband return containing 3 targets, each separated by 0.4λ m	41
3.8	Generic matched filter model [35]	42
3.9	Spectral comparison between the reconstruction and matched filter [35] . .	44
3.10	Delay and Doppler slices of $ X(\tau, f_d) ^2$, where $T = 20\mu\text{s}$, $B=1$ MHz	47
3.11	Ambiguity plot for $p_{TX}(t)$ for various combinations of τ and f_d	48
3.12	Ambiguity plot for $p_{TX_m}(t)$ for various combinations of τ and f_d	49
3.13	Spectral and matched filtering performance of LFM and NLFM pulses [38]	52
3.14	Comparison of LFM, tapered LFM and NLFM responses	53
3.15	Input 13-bit Barker sequence, autocorrelation and accompanying PSD . .	54
3.16	Interpreted Angle of Arrival (AOA) verse true target angle for $D = 2\lambda$. .	56
3.17	Ambiguous AOA chart for 3 carrier frequencies with $D = 30$ m	57
3.18	Planer wave model used to define target ranges relative to each receiver .	58
3.19	Baseband return consisting of 10 evenly spaced targets	58
3.20	True and estimated target locations derived from simulated example . . .	61
3.21	Standard deviation AOA algorithm summary	62
3.22	Error in estimated AOA for varying input SNR levels	63
3.23	Error in estimated AOA for fixed SNR of 5 dB, but with various carriers .	63
3.24	The Red Pitaya: Consumer Based Oscilloscope	65
3.25	Simplified version of the Red Pitaya PS/PL memory model	67
3.26	AXI4 write sequence from master to slave [44]	69
3.27	Two-stage ring buffer used to write recorded data to the SD card	70
3.28	Graphical interpretation of the demodulation process	74
3.29	Filter response used for demodulating the real sampled signal	75
3.30	Receive-only components of the interferometer FPGA design	76
3.31	Low-pass filter frequency response used in the design of the I/Q demodulator	78
3.32	Custom cyclic transmitter buffer	80
3.33	Received signal from loop-back test in channel A (Blue) and B (Red) . . .	82
3.34	Timing sequence of the Red Pitaya initialisation when daisy-chaining . . .	83

3.35	Daisy-chain configuration of the Red Pitaya	84
3.36	External pin can be used to drive a protecting circuit	85
3.37	Circuit diagram of a 12.5-25 MHz anti-aliasing bandpass filter [49]	86
3.38	Simulated & practical measurement of a 12.5-25 MHz anti-aliasing filter [49]	87
3.39	Conical antenna design used for the interferometric trials [50]	87
3.40	Reflection coefficient curve of the conical antenna [50]	88
4.1	Received baseband signal from trial PFT001	94
4.2	Match filtered signal from trial PFT001	94
4.3	Processed and integrated signal from trial PFT001	94
4.4	Power loss measured along 208.42 m coaxial cable for trial PFT020	95
4.5	LFM chirp pulse delay test on a 208.42 m coaxial cable for trial PFT021	96
4.6	Hamming LFM chirp delay test on a 208.42 m coaxial cable for trial PFT022	96
4.7	The split signal travelling along the two 18 m cables for trial PFT023	97
4.8	Received pulse from two independent Red Pitaya devices for trial PFT024	97
4.9	Geometrical layout of the interferometric field test	98
4.10	PSD of background noise power derived from the ITU-R P.372-7 [52]	102
4.11	Trial FM00A with default parameters. Only $f_c = 17$ MHz is shown.	103
4.12	Trial FM00AA with default parameters. Only $f_c = 17$ MHz is shown.	103
4.13	Trial FM00B with default parameters. Only $f_c = 17$ MHz is shown.	104
4.14	Trial FM00BB with default parameters. Only $f_c = 17$ MHz is shown.	104
4.15	Trial FM00C with default parameters. Only $f_c = 17$ MHz is shown.	105
4.16	Trial FM00CC with default parameters. Only $f_c = 17$ MHz is shown.	105
4.17	Trial FM00D with default parameters. Only $f_c = 17$ MHz is shown.	106
4.18	Trial FM00DD with default parameters. Only $f_c = 17$ MHz is shown.	106
4.19	Trial FM00E with default parameters. Only $f_c = 17$ MHz is shown.	107
4.20	Trial FM00EE with default parameters. Only $f_c = 17$ MHz is shown.	107
4.21	Trial FM00F with default parameters. Only $f_c = 17$ MHz is shown.	108
4.22	Trial FM00FF with default parameters. Only $f_c = 17$ MHz is shown.	108

4.23	Trial FM00G with default parameters. Only $f_c = 17$ MHz is shown.	109
4.24	Trial FM00GG with default parameters. Only $f_c = 17$ MHz is shown.	109
4.25	Trial FM00H with default parameters. Only $f_c = 17$ MHz is shown.	110
4.26	Trial FM00HH with default parameters. Only $f_c = 17$ MHz is shown.	110
4.27	Trial FM00I with default parameters. Only $f_c = 17$ MHz is shown.	111
4.28	Trial FM00II with default parameters. Only $f_c = 17$ MHz is shown.	111
4.29	Trial FM00J with default parameters. Only $f_c = 17$ MHz is shown.	112
4.30	Trial FM00JJ with default parameters. Only $f_c = 17$ MHz is shown.	112
4.31	Overview of the construction used to detect moving aircraft	114
4.32	Range-Doppler heat-map generated from moving aircraft experiment	115
4.33	Geographical map of the illuminated area for the moving target experiment	115
A.4.1	Internal configuration of the Zynq IP Core	136
A.5.1	Phased array consisting of multiple synchronised Red Pitayas [49]	137
A.6.1	Final Vivado FPGA design (Only viewable through PDF zoom)	138
A.7.1	Front-end 12.5-25 MHz filter [49]	139
A.8.1	Constructed fishing gut layout used to mimic Figure 4.9	140
A.8.2	Antenna RX_A positioned atop the grid shown in Figure A.8.1	140
A.8.3	Antenna RX_A with antenna RX_B 30 m in the distance past the umbrella.	141
A.8.4	Configuration of the Red Pitaya in the field experiment	141
A.8.5	The transmit antenna, TX_A , located roughly 200 m from RX_A	142
A.8.6	Complete interferometer FPGA design	143

List of Tables

2.1	Maximum unambiguous elevation angles for SuperDARN radars [2]	19
2.2	Summary of results acquired from COTS incident angle experiment [23] .	25
3.1	Expected Doppler shifts for 3 targets, all located 4.5 m from the receiver .	37
3.2	Common tapering function properties relative to a rectangular window [37]	51
3.3	List of all known Barker codes [40]	54
3.4	Simplified example of the CRT algorithm used to expand the AOA	64
4.1	Outline of general pre-field test trials	92
4.2	Measured metrics obtained from lab-environment trials	93
4.3	List of default and varied parameters used at each visited test location . .	101
4.4	Estimated power budget for field experiment	102
4.5	Summary of results depicted in Figures 4.11 to 4.30	113
4.6	Interpreted AOA using the standard deviation algorithm	113
A.3.1	List of common window function definitions	135
A.8.1	Resource utilisation on the 7z010clg400-1 FPGA device	143
A.8.2	The Red Pitaya register set & bit mappings	145

List of Acronyms

AOA	Angle of Arrival
AWG	Arbitrary Waveform Generator
AXI	Advanced eXtensible Interface
COTS	Commercial Off-The-Shelf
DDS	Direct Digital Synthesiser
DSP	Digital Signal Processing
EM	Electromagnetic
FIR	Finite Impulse Response
FOV	Field of View
FPGA	Field Programmable Gate Array
HF	High Frequency
HDL	Hardware Description Language
IP	Intellectual Property
I/Q	In-phase / Quadrature
LFM	Linear Frequency Modulated
LOS	Line of Sight
LPDA	Log-Periodic Dipole Array
NLFM	Non Linear Frequency Modulated
OS	Operating System
PL	Programmable Logic
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Interval
PS	Processing System
PSD	Power Spectral Density
RF	Radio Frequency
SANAE	South African National Antarctic Expedition
SNR	Signal to Noise Ratio
SuperDARN	Super Dual Auroral Radar Network
TTFD	Twin Terminated Folded Dipole

Chapter 1

Introduction

The manipulation of Electromagnetic (EM) radiation is at the forefront of numerous advanced technologies. Applications involving X-ray imaging, early warning systems, telecommunications and synthetic aperture radar are all examples of the versatility of EM propagation. It can be incorporated into many applications including health care, warfare and industry.

This dissertation aims to understand a subcomponent of this wide field, narrowing its scope to a problem of interferometry. The sections to follow provide background, context, objectives and motivation for this study.

1.1 Background

The Super Dual Auroral Radar Network (SuperDARN) is an international collaboration of radar systems dedicated to measuring plasma (or charged particle) convections in the high-latitudes of the ionosphere [1]. Currently, the network consists of more than 34 radars spread between the southern and northern hemispheres. Each radar follows a similar design, and operates within the High Frequency (HF) band from 8 to 20 MHz. Figure 1.1 depicts the location and bearing of each radar station.

The charged particles observed by the network tend to group vertically along the earth's magnetic field lines and demonstrate a specular type of reflection to incident waves [3]. As a result, a perpendicular propagation trajectory from each radar to the plasma densities is required for detection. The left half of Figure 1.2 visually depicts this requirement and shows the reflective nature of the plasma densities.

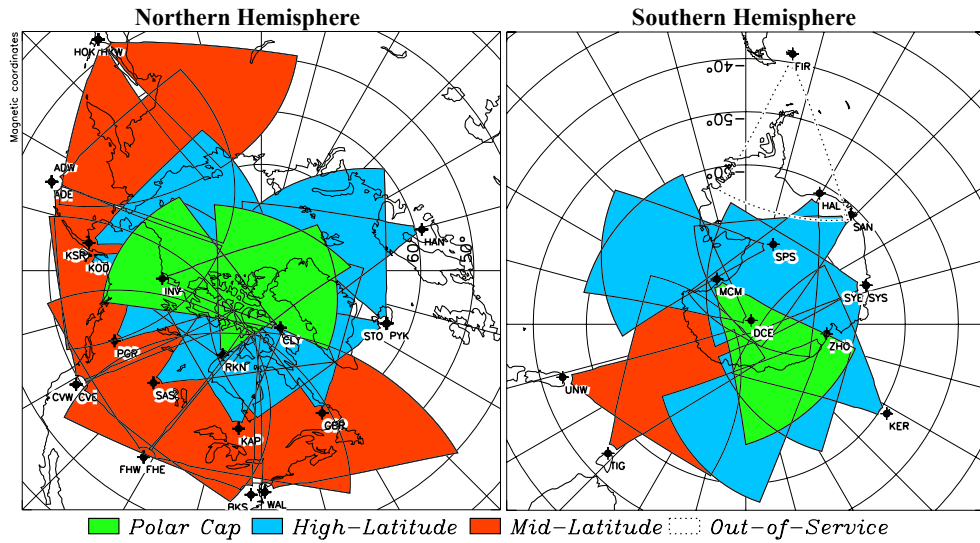


Figure 1.1: SuperDARN coverage map [4]

Observe that frequencies in the HF band tend to curve in the ionosphere. For higher frequencies, the ray-path tends to exit into outer space. The SuperDARN radars are specifically positioned to allow the HF signals to orthogonally illuminate the plasma densities. A detailed explanation of the propagation path is presented in Section 2.1.




Figure 1.2: HF wave propagation and electron height density profile in the ionosphere

Within the ionosphere, there are three definitive regions of interest denoted by D, E and F, as seen in the right of Figure 1.2 [5]. Each region increases in height respectively. Layers F2 and E form a local maximum for a given height profile. These regions are free to change in altitude depending on a multitude of conditions including solar radiation, thermospheric winds, geography and atmospheric composition.

HF wave propagation in the D region tends to be attenuated, while in the E and F regions it can be refracted or reflected. The electron density in each region is separated by orders

of magnitude. The F regions tend to persist, while the D layer subsides at night. In certain cases during the day, the F region can split into two definable regions, F1 and F2. At night, the F2 layer tends to shift and mask the F1 layer. Due to the persistence, altitude and high electron content of the F layer, it is well suited to act as a reflector for long distance radar applications.

Each radar within the SuperDARN is focused on monitoring plasma density irregularities in the E and F regions. By combining the Line of Sight (LOS) Doppler information extracted through paired radar locations (see Figure 1.1), vectored convection maps of the plasma densities can be generated [3]. Paired radars are required to determine absolute velocity from multiple measurements of radial velocity. An example of a vectored convection map for the Northern Hemisphere is shown in Figure 1.3.

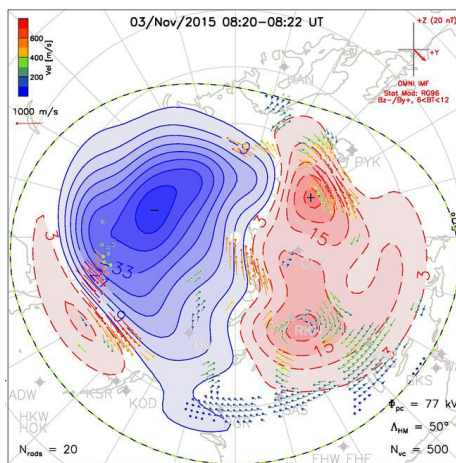


Figure 1.3: Example ionospheric convection map over Northern Hemisphere [6]

The SuperDARN data has major relevance to the scientific community and provides valuable input to other extended fields of research involving, but not limited to gravity waves, mesospheric winds and magnetohydrodynamics [3].

While complex mathematical and numerical models exist to characterise and explain the behaviour of the ionosphere, it is not made a focus of this research. Instead, attention is focused towards a subcomponent of a SuperDARN radar.

More specifically, in addition to Doppler feedback, the Angle of Arrival (AOA) from received signals can be used for added insight in determining the propagation modes of transmitted HF signals [7]. Furthermore, it can be used to decipher at which altitudes high plasma densities reside. The mechanism employed to achieve an AOA forms the subject of this dissertation and will be elaborated below. The present introductory section provides foundation and context to the project.

1.2 Problem Formulation

The motivation for this dissertation is based on the angular ambiguities of the interferometric measurements of the SuperDARN radar located at the South African National Antarctic Expedition (SANAE) IV base station. This radar was initially constructed in 1997, upgraded in 2013, and consists of two phased antenna arrays [8].

The two arrays are separated by 100 m and are constructed from horizontally polarized Twin Terminated Folded Dipoles (TTFDs). The primary array consists of 16 elements that can transmit and receive. The secondary array has 4 elements and is only capable of receiving.

This SuperDARN radar operates within 9 to 18 MHz and makes use of beamforming to span across its 52° azimuthal Field of View (FOV). The total FOV is segmented into 3.25° increments, amounting to 16 different beam positions. As its operating frequency increases, its azimuth beamwidth reduces. At 10 MHz and 14 MHz, its 3 dB beamwidth is 4° and 3° respectively. Its vertical beamwidth is designed to be around 50° depending on the frequency. The radar is capable of adapting its transmitted signal and can achieve range resolutions between 15 and 45 km. The flexibility of the radar's parameters is designed to cater for the variability in altitude, range and refractive properties of the E and F regions in the ionosphere.

As with the other SuperDARN radars, the second smaller array is intended to be used as an interferometer array. This allows for the time at which an arriving echo hits each array to be independently captured. Using basic geometry, the difference in arrival times can be converted into a phase shift and used to infer an angle at which a signal is received.

While this is simple in theory, it is challenging to realise as the characteristics of the array's beam pattern is not fully known and calibrated [7]. In other words, the effective transfer function of the radar system itself is not fully understood. This directly influences measured results. Furthermore, the calculation of the AOA is dependent on continuously changing variables. Such effects include gradual changes in antenna properties, electrical hardware, thermal conditions and environment. Thus, fully understanding and confidently measuring an AOA is considered a delicate and difficult task.

Depending on the architecture and layout of the SuperDARN antenna arrangement, only angles below roughly 33° can be measured unambiguously when operating in the HF band [2]. For mid-latitude radars pointed towards the earth's poles, this issue is assumed to be less impacting as ionospheric backscatter returns are predicted to be within the range of around $0-40^\circ$ [7]. For SuperDARN radars located at higher latitudes, this problem becomes more severe as echoed returns are expected to alias into the principle measuring range. These issues are elaborated in Section 2.2.3.

This study extends from this problem, and aims to investigate and build a localised scaled version of the SANAE IV interferometer. Using similar parameters to the SuperDARN

radar, a smaller and confined experiment is designed in which the AOA calibration is investigated. In this experiment, a known beacon signal is established and transmitted. Correspondingly, a system consisting of two antennas is used to receive the beacon signal and determine its location. This is done in azimuth as opposed to elevation as done by the SuperDARN radars. The experiment aims to establish what kinds of ambiguities are related to the problem.

Local experimentation circumvents several problems associated with using the actual SuperDARN radar. Foremost, its location in Antarctica makes it impractical to work with. Its far-field range extends several kilometres and makes testing in those ranges impractical. As an alternative, testing in the near-field can be transformed to equivalent far-field measurements through the application of a Fourier transform [8]. However, this has the disadvantage of introducing unnecessary complexities. This study therefore makes use of local equipment and resources to allow for more flexibility and freedom when researching this problem.

The theory from modulating the beacon source signal to sampling and digitally processing it will be covered.

1.3 Applications for Angle of Arrival Estimations

While determining the angle of arriving signals in this context will allow for improved interpretation of received echoes, its principle has relevance in many different applications, three of which are mentioned to motivate the design of this project.

- **Enhanced 911:** In 1996 the Federal Communications Commission in the USA began mandating an effort to improve emergency services for cellular callers [9]. In the interests of safety, legislation was passed that would enable the location of an emergency caller to be identified from their mobile device. One of the many mechanisms used to achieve this involves using principles of interferometry. In this scenario, the angle between a cellular device and its surrounding base stations can be computed using the phase delays of transmitted telecommunication signals. The position of the emergency caller can be identified by combining the measured angles. This location can then be passed on to appropriate emergency services such as the police or fire department.
- **RFI Direction Finding System:** Gowans identifies that various radio astronomy sites require locations with little spectral interference in order to detect weak signals from space [10]. In his study he develops a customised interferometer which is able to locate both impulsive and narrowband sources of interference. He develops an algorithm which makes use of both phase interferometry and time difference of arrival to do this. This application is intended to identify sources of interference within a 360° field of view such that they can be remediated as to not interfere

with larger and more sensitive radar installations.

- **Robotic Guidance:** Kim makes use of Radio Frequency Identification (RFID) tags to navigate a robot through space [11]. By placing RFID tags on neighbouring objects, the robot is able to identify and navigate its way towards them. It does this by measuring the relative signal strength of received echoes as it manoeuvres. This is somewhat different to the other two applications as it does not rely on phase information. Although the study makes large assumptions about the robot's environment, it shows promising results.

The theory adopted in all these neighbouring studies make use of similar principles, which justifies further research.

1.4 Research Question

To create a prominent centre and focus for this study, a research question is proposed:

To what degree of accuracy can an HF beacon signal be located in azimuth using the phase difference between two spatially separated receiving antennas? Furthermore, what ambiguities exist and can they be circumvented?

The research question has been agreed with SANSA and the answer will be addressed in subsequent chapters.

1.5 Scope & Limitations

A subset of the SANAE radar capabilities is modelled in order to feasibly conduct this work. Namely, only angles in azimuth are measured as opposed to elevation. Physically orientating antennas on ground-level over an incremental span of 180° is far simpler than doing the equivalent in elevation. On the contrary, this does not model the vertical trajectories and returns which the SANAE IV radar would observe. Effects of multipath distortion, refraction and polarisation rotation are not modelled.

A total of 3 antenna elements are used, as opposed to 20 in the SANAE radar. The SuperDARN beamforming capabilities are not modelled as part of this work in the interests of simplifying the problem financially, practically and theoretically. Focus is drawn on the AOA, and not necessarily steering abilities.

Two antennas are used to receive. One represents the main array and the other the secondary array of the SuperDARN radar. Alternative arrangements consisting of more receive elements are not implemented as the problem solution requires as few hardware modifications.

The third transmitting antenna is also modelled as a point source target to simplify the problem. Developed algorithms do not accommodate for extended targets which is measured by the actual SuperDARN radar.

Consumer based Field Programmable Gate Arrays (FPGA) hardware is used to capture received signals and execute basic Digital Signal Processing (DSP) operations. An input sampling rate of 125 MHz and primary carrier frequency of 17 MHz is used in both local trials and on the SANAE radar. A full explanation of these parameters is provided.

Conical antennas are built and locally tested. External filters used on the front-end are similarly constructed. All Verilog code is independently written for the FPGA and low-level C code designed to extrapolate raw binary data. More intricate algorithm designs are post-processed in Matlab. Issues of synchronisation between internal clocks, soft-time Operating Systems (OSs) and SD write speeds when demodulating the received signals are addressed. All theory related to the transmitted waveform, received signal, amplification, conditioning, filtering and processing is covered.

It is worth noting that as conical antennas are used, they exhibit an omnidirectional beam pattern. This is different to the Twin Terminated Folded Dipoles (TTFDs) used by the SANAE IV radar which have a roughly flat gain around 5-60° in elevation.

Results used to prove correct functionality of the developed hardware and software is included as evidence of expected behaviour. A marriage between theory and practice is attempted to show a proof of concept in which a low unambiguous azimuthal angle is autonomously expanded using carrier frequency agility.

1.6 Plan of Development

This work is separated into five remaining chapters. Each chapter is written with the research question in mind, where the overall goal is to create a cohesive discussion on interferometry.

Chapter 2 continues with a literature review, in which three different studies are identified on the calibration of various different SuperDARN radars. A further case study is also presented in which an AOA is measured using Commercial Off-The-Shelf (COTS) equipment. The relative success of each study is assessed. In addition, a background on the ionosphere and its refractive properties is included.

Chapter 3 focuses on the design of the interferometer used for local experimentation. All theory related to capturing the received signals, developed hardware, processing algorithms and antenna structures is presented. Expected gains, frequency shifts and received angles are tested, evaluated and documented in this section.

Chapter 4 defines the experiment conducted to evaluate the success of the interferometer. Various parameters are interchanged along with their recorded results.

Chapter 5 critically analyses the results presented in Chapter 4. Explanations for all deviations are theorised. The research question is evaluated, and an assessment is provided.

Chapter 6 concludes the dissertation, and makes recommendations for further research.

Chapter 2

Supporting Literature

The mechanisms adopted to compute the angle of an impinging signal have been much discussed. This chapter will collate, review and draw on the results from various studies. It begins with an overview of HF propagation modes in the ionosphere to further contextualise the purpose of this study. Following this, a discussion in which three different attempts to calibrate the SuperDARN radars are examined. Further, the difficulties associated with the required interferometric calculations are addressed. In addition, the chapter ends with a discussion on using Commercial-Off-The-Shelf (COTS) equipment to measure an inbound angle without coherent synchronization between multiple receivers.

2.1 SuperDARN HF Propagation in the Ionosphere

Before addressing the interferometric problem, it is important to understand how the SuperDARN HF signals propagate within the ionosphere. Although this study will not make extensive use of this subsection as it relies on a much more confined scenario, it is provided to develop context.

Extending from Section 1.1, the ionosphere, located some 60 km above the earth's surface, consists of a plethora of molecules. These molecules, neutral in nature, become ionized due to highly inbound energised solar radiation [5]. This causes electrons in the molecular make-up of the atmosphere to separate from their respective atoms. Effectively, this results in an ionized layer of particles.

This phenomenon becomes so well-defined that it is capable of refracting radio waves. In literature, it is common to estimate and define the *plasma frequency* as follows [12]:

$$f_p \approx 9 \times 10^{-6} \sqrt{N_e} \quad (2.1)$$

where N_e is the electron density per cubic metre. A vertically transmitted signal towards the ionosphere will be reflected provided that the carrier frequency, f_c , is less than the plasma frequency. If f_c exceeds f_p , the signal will penetrate the ionosphere and exit into outer space. In order for reflection to occur at elevated transmission angles, α , the following condition must hold [12]:

$$f_c \leq \frac{f_p}{\sin(\alpha)}. \quad (2.2)$$

These relationships are better visualised as shown in Figure 2.1. The grey lines extending outward from 0 km range are separated by 1° and depict a series of ray paths for different transmit elevations. Note that for larger elevations, the ray paths tend to penetrate the ionosphere. For smaller angles, the ray path bends back towards earth. This is found to be congruent with Equation 2.2.

The nearly straight vertical lines spaced roughly 250 km apart represent the earth's magnetic field lines. Recall that in order for a detection to occur, orthogonal angles between the magnetic field line and ray path are required. The black painted 'S' region indicates where the ray path falls within a degree of being orthogonal to the magnetic field line. The white curved vertical lines represent equivalent contours of range. Each successive contour is 225 km apart, with the first contour denoting a distance of 180 km from the radar. The contours take on distorted shapes due to the refractive properties of the ionosphere on HF signals.




Figure 2.1: Computed ray paths for SANAE radar at 11 MHz [13]

It is worth briefly mentioning that the computed ray paths in Figure 2.1 correspond to the 7th fixed beam position of the SANAE SuperDARN radar. Its location and field of view is shown in Figure 1.1 and indicated by the acronym 'SAN'.

The region from 0 to 800 km is denoted the *skip* zone as no ground returns can be achieved within this range [12]. The skip zone can be reduced by lowering the transmit

frequency. The illuminated stretch from 800 to 1900 km is known as the *footprint* of the transmitted signal. It extends several hundred kilometres and requires ground clutter be mitigated through Doppler filtering in order to detect the plasma densities.

The necessity of determining the AOA becomes apparent from Figure 2.1. Without knowledge of the AOA, one cannot determine the altitude of returns. Less obvious, a signal arriving at its first projected footprint or *hop* may reflect again further onwards and upwards into the ionosphere. This can result in a return from several hops away. The AOA thus forms crucial for determining the true path and characteristics of a received signal.

The complexity of this task cannot be overstated. Accurate knowledge of the electron content in the ionosphere is required to predict ray paths and allows one to map angles of arrival to target ranges and altitudes. Furthermore, transmitted signals are subjected to a multitude of effects including attenuation in the D layer and refraction or possibly reflection in the E layers of the ionosphere (Section 1.1). This can lead to multipath propagation issues and ambiguities in measured returns. These factors need to be taken into account.

A second condition for detection requires that the plasma irregularities be separated by half a wavelength to achieve the strongest returns with in-phase constructive interference [1]. This is consistent with the principles of Bragg's scattering [14].

Figure 2.2 demonstrates that constructive interference between two imminent wave-fronts is achieved when:

$$2d \sin \theta = N\lambda \quad N = \pm 1, \pm 2, \pm 3 \dots \quad (2.3)$$




Figure 2.2: Geometrical visualisation of Bragg's scattering

The symbol λ designates the wavelength of the carrier. Assuming the orthogonal condition previously mentioned is satisfied, maximum optimal constructive interference is achieved when $d = \frac{N\lambda}{2}$. As a result, the strength of the echoed signals are dependent on the spacing of the irregularities and the transmit frequency.

In order to calibrate the SuperDARN radars for optimal returns, prior knowledge of the ionosphere is needed [12]. The International Reference Ionosphere (IRI) forms an

empirical best-fit data model that can be used for computing ray paths and predicting conditions [8]. Coincidentally, Figure 2.1 is partially computed from the IRI. Alternatively, the Chapman layer model can be used and is instead founded on a theoretical basis.

In general, the use of extended long distance Over the Horizon (OTH) transmission via the ionosphere allows for a multitude of applications including military surveillance, aircraft tracking and remote sensing. In each case, the agility of the radar's operating parameters need to account for continuous fluctuations in the ionosphere's D, E and F regions as well as for the target under consideration.

This concludes the discussion on HF propagation. Several topics were addressed including the formation of the ionosphere, its refractive qualities on HF propagation and the conditions required for plasma irregularities to be detected.

2.2 Previous SuperDARN Calibration Attempts

There have been several attempts to calibrate the SuperDARN radar without overwhelming success. This section elaborates three such instances and evaluates the approaches taken to achieve accurate interferometric data. The final case study differs as it proposes a change to the hardware layout, whereas the first two discuss methods to characterise the radar.

2.2.1 SANAE IV - Helicopter Beacon Calibration Attempt

In January 2007, Cilliers conducted a field test in which the phase response of the SANAE radar was investigated [15]. He hypothesised that the radar applied internal phase delays through its cables, receivers and electronics. Further analysis revealed that the radar's attenuators introduced a constant phase shift to received signals independent of the amount of suppression, warranting the hypothesis.

The experiment aimed to develop a calibration curve such that recorded data could be corrected against phase distortions within the radar system. In other words, a transfer function describing the radar was to be determined. It was hoped that one could invert the effects of the radar system. In order to develop data for such a calibration curve, a set of known input signals were applied to the radar at different elevation angles.

A continuous wave beacon transmitter at 13.165 MHz was placed on a helicopter and flown at boresite to the radar's field of view. Measurements were taken at 0, 2, 5, 10 and 20 km from the radar at altitudes of 1 and 3 km respectively. The radar was then used to receive the signal and measure the corresponding phase difference between its main and interferometer array.

Figure 2.3 was used to derive expected results, where the dotted and solid vertical lines represent the main and interferometer arrays respectively. The main array is 10 m tall, while the secondary array is 13 m high. The distance between the two arrays, d_1 , is 100 m. Two different models were proposed to compute the angle of arrival.




Figure 2.3: Geometry of interferometric calculations used to determine an AOA [15]

The first made use of an assumption that all the elevation angles were equal and that $h_1 = h_2$. Under these conditions, the theoretical phase shift is given by:

$$\Psi_1 = \frac{d_1 \cos(e_n)}{\lambda} \times 360^\circ \quad \text{where } e_n = e_1 = e_2 = e_3 \quad (2.4)$$

The second adopted model accounted for the antenna height discrepancies and for subtle differences in the arrival angles. This model aligns with what was implemented in practice, and derives its phase shift as follows:

$$t_1 = \frac{\sqrt{(x_3)^2 + (h_3 - h_2)^2}}{c}, \quad t_2 = \frac{\sqrt{(x_1)^2 + (h_3 - h_1)^2}}{c} \quad (2.5)$$

$$\Psi_2 = \frac{t_2 - t_1}{T} \times 360^\circ \quad (2.6)$$

The symbol T denotes the period of the transmitted signal (76 ns) and c the speed of light. The graphical interpretations of these models along with the measured results are provided in Figure 2.4.

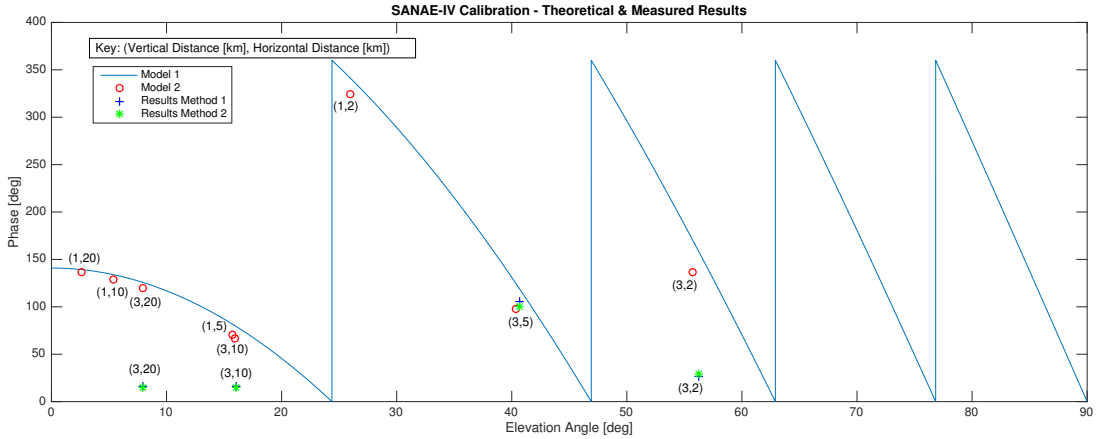


Figure 2.4: SANA-E Calibration - Theoretical & Measured Results [15]

It is interesting to observe the aliasing in phase exhibited by the first model, where the maximum unambiguous elevation angle measurable is 24.3° . This limit can be increased by either reducing the transmit frequency or narrowing the distance between the arrays. Aliasing occurs provided the separation length of the two arrays exceed half a wavelength. A similar problem of ambiguity exists for the second model. The elevation angle is shown to wrap at several locations:

$$e_\Delta = \cos^{-1} \left(\frac{\lambda n}{d_1} \right) , \quad n = 0, 1, 2 \dots \text{floor} \left(\frac{d_1}{\lambda} \right) , \quad n\lambda \leq d_1 \quad (2.7)$$

In order to measure the time difference between the arrays ($t_2 - t_1$), the signal arriving at the main array was cross-correlated with the received signal at the interferometer array. The time delay at which maximum correlation occurs was used to quantify the phase difference.

An alternative method was also used in which the signals arriving at each array were correlated with a modelled sine wave of 159 Hz. The net phase difference was calculated by subtracting the individually observed phase shifts at each element. The original transmitted signal contained a 159 Hz harmonic. The results captured from these two methods are indicated by the blue and green markers respectively in Figure 2.4

The remainder of the experiment was concerned with fitting the measured results to the theoretical data using the Mean Square Error (MSE). This was subsequently used to compute the ‘phase correction’ calibration curve encompassing the radar’s implicit phase delays. Despite this, according to Cilliers, the study was found to be inconclusive [15]. From Figure 2.4 it is clear that the practical results do not appear to follow a consistent or predictable pattern.

It is important to note that not all of the planned experiments were fully recorded for logistical reasons. This resulted in a smaller data set than anticipated. While in theory

this experiment was well planned, there are a few concerns regarding the accuracy of the measured results.

Foremost, roughly half of the data was captured using photographed oscilloscope monitors [15]. These displays were digitised and used for analytical results. Inherently this introduces inaccuracies. Similarly, imprecision due to the helicopter's on-board GPS (used to infer the beacon's true position) was not taken into account.

The experiment also did not consider alternative frequency responses. No strong evidence was provided to suggest that the radar maintained a constant phase delay. For instance, a signal transmitted at 19 MHz could have produced a different calibration curve to one at 11 MHz. It also would have been interesting to observe the response at non-boresite angles. Milan provides an interesting perspective in which his interpretation of the CUTLASS HF radar takes this into account [16].

Nevertheless, it is important to recognise the models and methods adopted to measure the phase difference between the received signals as they are directly related to this study. Despite the experiment being considered inconclusive, the lessons learnt and approaches explored provide valuable insight for future work.

2.2.2 ZACUBE-1 Calibration Attempt

The ZACUBE-1 encompasses South Africa's first 10 cm³ nano-satellite [17]. The CubeSat weighs 1.3 kg and is designed to operate on 3 W. Its primary payload consists of a single 14.099 MHz transmitter and low resolution camera [18].

The satellite was launched in 2013, with the primary goal of enabling the beam patterns of the SANAE radar to be characterised. Using similar theory to the previous case study, the CubeSat was designed to illuminate a signal from its position in Low Earth Orbit (LEO) from roughly 600 km altitude towards the ground. As it passes into the radar's field of view, its signal becomes a reference for interferometric calibrations. By knowing the position of the satellite, experiments can be conducted to qualify the measurement of the elevated AOA.

This approach has several advantages over the helicopter case study [17]. Foremost, its recurring availability and predictability allow for repeatable tests to be conducted, whereas attempting to fly a helicopter in the Antarctic region tends to face logistical and safety complications. The experiment administrated by Cilliers concluded with a single trial. Secondly, the distance and manoeuvrability of the beacon satellite allows for a more dynamic range of results to be captured in the far-field of the radar. In the former case, the flight path was restricted and tests were conducted in the near-field.

On the other hand, in the ZACUBE-1 project, additional complications are introduced. As the beacon signal travels through the ionosphere towards earth, it is subject to a multitude of distortions, including refraction, attenuation, and polarisation shifting. The

relative influence of these factors is largely dependent on the trajectory of propagation and gradient of the electron density in the ionosphere. Figure 2.5 presents a summary of potential distortions.




Figure 2.5: Atmospheric effects on LEO signal propagation [19]

Consider in the context of non-geosynchronous LEO satellites, the effect of a Doppler shift on a transmitted signal. According to Davies, this can be approximated by [19]:

$$f_d = \mu \frac{v_l}{\lambda} + \frac{40.3}{c f_c} \frac{dN_\tau}{dt} \quad (2.8)$$

where v_l , μ and N_τ is the line of sight velocity, refractive index at the satellite and Total Electron Content (TEC) respectively. The first term relates to the motion of the satellite with respect to the receiver, where the second term accounts for the varying electron content. For high carrier frequencies, the Doppler shift can approach several KHz and needs to be catered for in the receiver [20]. Furthermore, the observed Doppler shift will fluctuate in magnitude non-linearly depending on the geometry of the satellite's pass. This Doppler shift inherently alters the instantaneous phase which would be measured on earth.

Incorporating these factors proves crucial as a naïve measurement of the elevation angle will produce incorrect results. One cannot simply reproduce the true position of the beacon using the same methods developed in the previous case study. To alleviate the problem, the IRI and Chapman models (Section 2.1) can be used to predict and correct for distortions in amplitude and phase of the beacon signal. Data provided from Antarctic Global Navigation Satellite Systems (GNSS) can also be used to improve models [17]. In general, this problem extends to all satellites.

Prior to the launch of the device, much discussion was held about where to place the CubeSat in orbit. Primary concerns were around the visibility of the beacon, and whether

its footprint would coincide with SANAE. Furthermore, the launch vehicle used to deliver the satellite into orbit had to be taken into account.

Several vehicles were considered including the Cosmos, Epsilon and Dnepr rockets. Each potential choice would affect the cube's final orbital path [18]. Rigorous ray and orbital path simulations were conducted in which the Dnepr rocket was confirmed to meet the project requirements [17].

At the time writing, the ZACUBE-1 is confirmed operational and maintains an orbital period of 98 minutes, and inclination of 97.8° [8]. An inclination near 90° is required in order for the satellite to be visible to radars positioned in the upper latitudes. It has a perigee and apogee altitude of 600 and 720 km respectively. The satellite orbits 14 times a day, and maintains sporadic visibility from SANAE.

However, despite successful communication with the CubeSat, and the retrieval of images from its onboard camera, its objective to enable the HF beacon failed. More specifically, its 10 m long antenna coiled within the cube failed to deploy. The antenna was originally intended to eject in a straight line while the cube spun at a nominal rate of 2° [17].




Figure 2.6: ZACUBE-1 (left) and its antenna deployment hardware (right) [17]

The inability to deploy the antenna meant that the HF beacon signal could not be transmitted. While this approach had interesting and potentially promising results, its implementation unfortunately failed. Despite this, the collaborative research involved in realising Africa's first CubeSat proved largely beneficial. Current research is ongoing and has extended towards developing a second larger CubeSat [21].

2.2.3 Buckland Park: Augmented SuperDARN Radar

Research by McDonald *et al.* at La Trobe University offers insightful information on this subject [2]. McDonald proposes an alteration to the SuperDARN HF antenna arrangement in order to improve the AOA calculations.

Where this study focuses on the SANA E HF radar, McDonald’s work encompasses the history and development of the generic SuperDARN radar. He draws attention to the evolution of the antennas. Prior to 2005, it was common for SuperDARN radars to make use of Log-Periodic Dipole Array (LPDA) elements [2]. However, the use of Twin Terminated Folded Dipole (TTFD) antennas has become common in recent years due to their lower cost, desirable beam pattern, and steering capabilities.

Figure 2.7 depicts the beam patterns of two 16-element arrays, where one consists of LPDA antennas and the other TTFDs. Notice that the ratio between the forward pointing and reverse facing gain is significantly better with regard to the TTFD antennas. Ideally, one hopes to achieve minimal gain in the back lobe as this mitigates unwanted echoes originating from behind the antenna structure. Note as well that the sidelobes for the TTFD arrangement are notably reduced in comparison. This is advantageous as the returns from undesired azimuth are reduced.

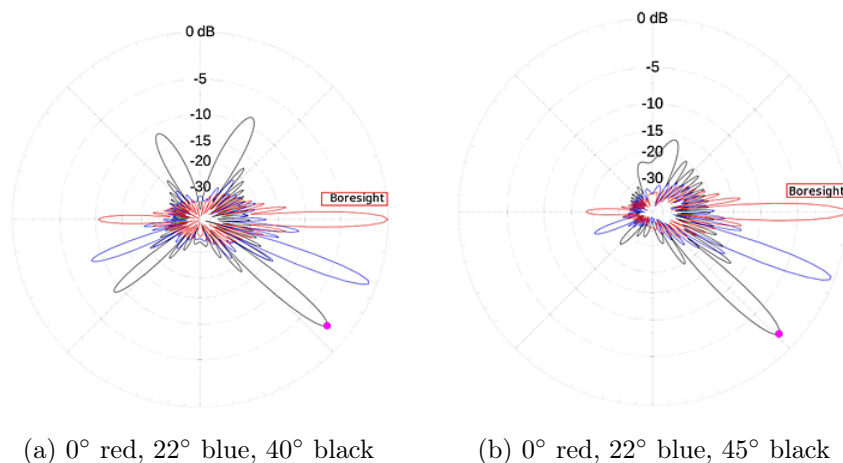


Figure 2.7: Simulated azimuthal beam patterns for a 16-element LPDA (left) and TTFD (right) antenna array. Transmission is at 14 MHz at an elevation of 20°. The legends of each field plot denote at which angle the beams are phased from boresight [2]

Interestingly, the SANA E radar was initially constructed from LPDA antennas. However, a storm in 2008 caused them to collapse [18]. The elements were shortly replaced with TTFDs because of their low wind profile and resilience in polar storms.

In order for the sidelobes to be optimally compressed, the spacing between the individual antennas must be kept within $\frac{\lambda}{2}$. This translates to a minimum gap of 7.5 m, provided the radar does not exceed 20 MHz. Traditionally, the LPDA antennas have been separated by 15.24 m. The size of the elements constrain this dimension and therefore it cannot be made smaller. As a result, beamforming efficiency tends to degrade significantly for operating frequencies above 10 MHz. For phased angles larger than around 30°, the power in the sidelobes becomes proportional to its main lobe [2]. This effect is clearly seen in

the left of Figure 2.7, where the sidelobe gain increases substantially as the azimuthal angle increases. The gain from the sidelobes introduces unwanted returns.

The use of TTFD antennas circumvents this problem, as the adjacent elements can be spaced closer due to their slimmer profile. This allows for steerable angles much larger than 30° to be achieved with sidelobes of at least 10 dB below its main beam. On the contrary, reducing the adjacent spacing increases antenna coupling which degrades performance. A more mathematical description regarding the beamforming process is provided by Greenwald [22].

While the spacing between adjacent elements has an effect on the formed beam pattern, the distance between the main and interferometer array influences the overall gain and quantifiable arrival angle. McDonald draws attention to this, and provides Table 2.1 for insight. As the distance between the two arrays decrease, the maximum measurable AOA increases. The discussion in Celliers’s case study arrives at a similar conclusion [15].

Table 2.1: Maximum unambiguous elevation angles for SuperDARN radars [2]

Separation Distance (m)	Δ_{max} 8 MHz	Δ_{max} 18 MHz
180	37.64°	24.84°
150	41.39°	27.26°
120	46.55°	30.55°
100	51.30°	33.55°
80	57.89°	37.64°
60	67.95°	43.75°

Naturally, one would think to place the two arrays within a half a wavelength from each other to achieve a Δ_{max} of at least 90° . However, due to effects of shadowing, this is avoided. To improve the gain of the TTFD arrays, corner reflectors are used to redirect radiated energy in the forward direction (or conversely, focus received energy). These corner reflectors tower some 17 m above ground, where the TTFD elements are no taller than 10 m [2]. Therefore, a nominal separation distance of 100 m is used to prevent occlusion of received signals at the cost of reducing Δ_{max} .

McDonald provides multiple simulations that expand on the shadowing effect. Figure 2.8 depicts the impact of placing the interferometer array in front of the 16-element main array at different distances. The dependent variable measures the relative gain of the main array as if it has no obstructions. Note that the use of more antennas in the secondary array increases the relative gain. For low elevation angles, the interferometer array shadows the main array and reduces its received power. This effect is reduced for angles roughly above 30° as signals can pass over the peak of the obstructing reflectors. Observe that displacements of 100 m produce significantly better results.




Figure 2.8: Relative gain of 16-element main array for increasing elevation angles with differentially positioned interferometer array. The left graph depicts results taken at 14 MHz [2]

Having discussed the geometry of the arrays and their respective elements, McDonald suggests an alternative antenna scheme shown in Figure 2.9. He prefers the use of two auxiliary arrays. One would consist of a single antenna and the other 3-elements, where each is to sandwich the main array. Modifications should not be made to the 16-element array in order to encourage lower costs for implementation.




Figure 2.9: New proposed layout for Buckland Park SuperDARN radar [2]

The respective dimensions for D_1 , D_2 and D_3 are 67 m, 80 m and 14 m. McDonald notes that the SNR in an N -element array tends to scale proportionally to \sqrt{N} . The gain of the single front antenna is likely to be some 2.3 dB below the rear array. For applications that require improved SNR in the auxiliary arrays, more elements can be appended at the expense of shadowing and coupling effects (as evident in Figure 2.8).

The use of two auxiliary arrays allow for the ambiguous elevation angle to be resolved. It can be shown that the phase shifts at the front and rear interferometer arrays, Ψ_1 and Ψ_2 respectively, can be calculated as follows [2]:

$$\frac{\Psi_1}{2\pi} + m = a(\cos^2 \phi_0 - \sin^2 \Delta)^{\frac{1}{2}} = aC \quad (2.9)$$

$$\frac{-\Psi_2}{2\pi} + n = b(\cos^2 \phi_0 - \sin^2 \Delta)^{\frac{1}{2}} = bC \quad (2.10)$$

$$a = \frac{D_1}{\lambda}, \quad b = \frac{D_2}{\lambda}, \quad \Psi_1, \Psi_2 \in [-\pi, \pi), \quad n, m \in \mathbb{Z} \quad (2.11)$$

$$n = \begin{cases} m & \Psi_2 \geq \Psi_1 \\ m + 1 & \Psi_2 < \Psi_1 \end{cases} \quad (2.12)$$

The symbol Δ denotes the elevation angle, while m and n represent unknown integers due to the aliasing effect. The azimuthal direction in which the phased array is pointing at zero elevation is given by ϕ_0 . Equation 2.12 formalises the relationship between the phases shifts measured at the two independent auxiliary arrays. Thus one can solve for Δ by expressing n in terms of m and substituting 2.10 into 2.9. This solution caters for the fact that the measured elevation angle is coupled to the azimuthal pointing direction of the phased array. In terms of realising a measurable phase shift, the methods adopted from Cilliers’s case study can be used.

McDonald makes for further analysis on impedance matching in the TTFD and LPDA antennas. This, however, is not elaborated in favour of maintaining scope. In general, his work shows delicate consideration for efficient beamforming and antenna design. Unlike the previous two case studies, this method aims to resolve the ambiguous arrival by making appropriate changes to the SuperDARN array design.

2.3 Angle Measurements Using COTS Equipment

Chen *et al.* conducts a study in which an incident angle created between two receivers and a target is measured [23]. Each receiver consists of three Universal Software Radio Peripherals (USRPs) operating at 900 MHz, three rubber ducky omnidirectional antennas, a single 10 MHz clock and a standard PC. The clock is used to synchronise each directly connected Software Defined Radio (SDR), while the PC is used to communicate with the SDRs over a gigabit Ethernet hub. Each receiver is stored in a weatherproof and shockproof case and mounted on a tower appropriately 6 m above ground. Although each receiver contains 3 SDRs, only 2 are used.

Figure 2.10 illustrates a simplified layout of the antennas within their protective containers. The phase measured at the two antennas, Φ_1 and Φ_2 , can be translated into an incident angle:

$$\Phi_1 - \Phi_2 = \frac{2\pi H}{\lambda} = \frac{2\pi D \sin \theta}{\lambda}. \quad (2.13)$$

It is worth noting that this definition is equivalent to the one provided in Equation 2.4, except that the complementary angle has been used here. The phase difference is largely a function of the additional propagation distance, H . In order for this definition to work, coherency needs to be maintained while sampling the antenna feeds.

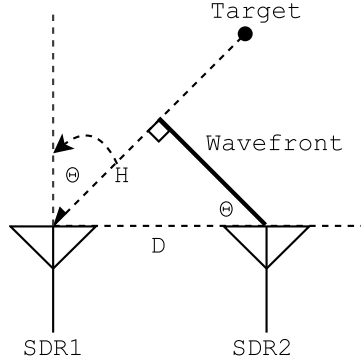


Figure 2.10: Geometry of antennas used to measure the phase difference [23]

The use of individually operated SDRs makes this study particularly interesting, as each SDR operates with its own oscillator. While an external master 10 MHz clock is used to set a common oscillator rate, their individual starting phases are unknown. In the Super-DARN radar, all the components are coordinated by a common clock. This proposes an issue as coherency is not naturally maintained across all the components. To circumvent this, an additional reference transmitter is used to perform Phase Difference of Arrival (PDOA) corrections.

The theory behind PDOA is relatively simple. In order to illustrate this method, Equation 2.13 is modified to represent the phase difference between two SDRs with unknown initial phases γ_1 and γ_2 , given a target signal A :

$$\Phi'_{A1} - \Phi'_{A2} = (\Phi_{A1} + \gamma_1) - (\Phi_{A2} + \gamma_2). \quad (2.14)$$

If one artificially introduces a reference transmitter at a known location and repeats the same experiment, one expects to observe a similar result. The difference of the two independent readings can be used to extract the true phase difference:

$$\Phi'_{B1} - \Phi'_{B2} = (\Phi_{B1} + \gamma_1) - (\Phi_{B2} + \gamma_2) \quad (2.15)$$

$$\Phi_{A1} - \Phi_{A2} = (\Phi'_{A1} - \Phi'_{A2}) - (\Phi'_{B1} - \Phi'_{B2}) + (\Phi_{B1} - \Phi_{B2}). \quad (2.16)$$

The true phase difference $(\Phi_{B1} - \Phi_{B2})$ can be derived from Equation 2.13 based on the location of the reference transmitter and receiver. The initial unknown parameters γ_1 and γ_2 are thus eliminated.

The study draws attention to three notable points [23]. First, the PDOA method is independent on the number of elements in the array. Each angle computation is reliant on a pair of SDRs. Second, advanced or specialised waveforms are not required to maintain

synchronization. The reference signal is simply provided with sufficient energy to eradicate the unknown initial phase offset. Third, this method makes use of entirely common variables to both the reference and target signals. Therefore, unaccounted for influences when computing the true phase difference are likely to be mitigated.

This technique is desirable as the initial offsets due to each oscillator do not have to be actively measured before experimentation. A single reference transmitter can be used to correct the measurements of numerous paired receivers.

Despite these advantages, the initial phase offsets are known to drift over time and to invalidate results. As a preventative measure, the PDOA computation can be periodically repeated. Alternatively, the drift rate can be estimated and used to compensate measurements over time. The initial phase offset is observed to be stable for at least one minute in practice, long enough to record meaningful measurements [23]

Figure 2.11 illustrates the layout of the field test conducted. Labels R and G represent the two receivers used, each of which are mounted on top of a 6 m tower in their protective containers. The reference transmitter used to mitigate the initial unknown phase is positioned at B , 7.5 m away from receiver R . A mobile target transmitter is located at A and operates at 916 MHz. The location and azimuth of A is to be measured in the field test. Both transmitters A and B are positioned approximately 35 cm above the ground.

Throughout the experiment, B , R and G remain fixed in their positions. Ten recordings are conducted, where transmitter A is rotated in a fixed arc at 5° increments between trials. Both A and B operate on the same frequency and take turns communicating with the receivers. Between each run, the receivers are recalibrated to reduce the phase offset. The layout is used to provide a ground truth against which to compare the measured position of A .




Figure 2.11: Layout of the COTS experiment [23]

The position of the target is not directly measured from Equation 2.13. Instead, a two-

dimensional grid is applied to the layout of the field test. For every location, there is an expected phase difference. By comparing the measured or estimated phase difference to the expected phase difference, the likelihood of a target being present at a given location can be evaluated. The best estimate is used to infer the target's location. The likelihood function used to define the position of a target is given by [23]:

$$\mathbf{L}(x) = p_x(\hat{\Phi}_{A1} - \hat{\Phi}_{A2}) \quad (2.17)$$

where $p_x(\Delta)$ is the probability of observing a phase difference, Δ , at the receiver provided a target is at position x . The phase measurement is assumed to be conducted in Gaussian distributed noise. Therefore, $p_x(\Delta)$ is modelled by:

$$p_x(\Delta) \sim \mathcal{N}(\Phi_{x1} - \Phi_{x2}, \sigma^2). \quad (2.18)$$

In other words, the probability density function of an observed phase difference is normally distributed with mean or expected phase difference ($\Phi_{x1} - \Phi_{x2}$) and variance σ^2 . The variance of the target's position is empirically chosen by calculating the deviation of the measured phase difference ($\hat{\Phi}_{A1} - \hat{\Phi}_{A2}$) over a period of time. As there are two receivers, the joint likelihood is used to estimate the position of the target:

$$\mathbf{L}_{(G,R)} = \mathbf{L}_G(x)\mathbf{L}_R(x). \quad (2.19)$$

The first two heat maps in Figure 2.12 depict the phase difference likelihoods as observed by receivers R and G respectively. The black symbol is used to represent the target which is situated at 5° along the arc as shown in Figure 2.11. The yellow dot represents the calibrating transmitter, while the squares represent the SDRs. The final panel shows the joint likelihoods as described by Equation 2.19.

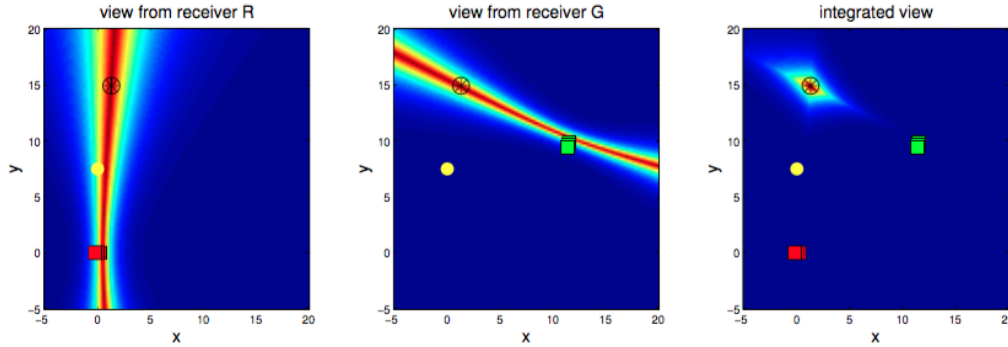


Figure 2.12: Phase difference observed from receiver R and G with integrated view [23]

Lines of constant phase are indicated by the same colour. The lines curve and distort outwards from the receivers due to the 3D geometry of the configuration. The model used

to represent this simulation incorporates the relationship between the elevation angle and the azimuthal angle. As a result, curves are introduced into the projection. Should all the receivers and transmitters be on the same plane, the phase difference lines would be straight.

It is also evident that the individual SDRs within each receiver box are sufficiently spaced as no aliasing occurs in the phase contours. In order for aliasing in the incident angle to occur, the antenna elements would need to be spaced by more than $\lambda/2 = 0.16$ m apart. Recall that the separation distance between the main and interferometer array in the SANAE radar results in a low unambiguous elevation angle of approximately 33° .

The results from the 10 recordings are listed in Table 2.2. Trial 3 performs the best with a localisation error of 0.17 m. This error measures the Euclidean distance between ground truth and the target’s estimated position. Trials 8 and 9 are considered outliers. The poor accuracy is hypothesised to be a result of the uneven ground plane on which the experiment was conducted. The field test relied simply on the geometry of the SDRs and did not account for environmental or hardware-specific corrections.

Table 2.2: Summary of results acquired from COTS incident angle experiment [23]

Run	1	2	3	4	5	6	7	8	9	10
Angle ($^\circ$)	0	0	5	10	15	20	25	30	35	40
Target to R distance (m)	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15
Target to G distance (m)	14.02	14.02	12.95	11.88	10.83	9.81	8.84	7.94	7.16	6.54
Localisation Error (m)	2.03	0.82	0.17	2.43	2.00	1.77	1.23	4.52	3.39	1.63
R Phase Error (%)	8.36	3.24	0.03	10.58	0.62	4.83	6.13	18.54	12.57	0.10
G Phase Error (%)	5.52	0.63	0.49	6.98	13.18	10.14	8.65	5.32	4.22	1.65
R measured phase variance ($^\circ$)	2.37	4.01	4.02	2.65	5.43	4.64	25.41	7.40	5.28	17.91
G measured phase variance ($^\circ$)	2.66	2.38	3.17	2.66	5.58	3.09	2.56	2.26	2.29	2.22

The majority of the results achieve a localisation error below 3 m. There is clear evidence that the proposed method works, especially with the applied PDOA technique. However, the results do not appear to be overly consistent or accurate. For example, Trials 1 and 2 are conducted from the same location, yet have a difference in localisation error of 1.21 m. The phase variance due to receiver R is also consistently larger than G . There appears to be a strong correlation between the target distance and the accuracy, as R is positioned further from the target than G . The use of additional receivers is likely to improve the results substantially.

In the context of the SANAE radar, the PDOA technique may prove too inaccurate. Given that this radar operates in the HF band and is intended to detect targets several hundred kilometres away, the algorithm’s accuracy may degrade substantially and translate into implausible measurements.

Despite these concerns, Chen highlights three advantages of using the modularised COTS equipment over the dedicated beamforming SANAE radar. Foremost the COTS solution allows for flexible design. Each SDR can operate over a wide bandwidth, be easily

separated and redistributed into new geometries. In Section 2.2.3, the positioning of the elements were shown to have a substantial impact on the performance of the radar. With more conservative solutions, the elements are fixed into a structure and tuned to operate over a narrow band. Secondly, the use of COTS equipment results in lower expenses. The price of each receiver box fell below \$5000, while the SuperDARN radars are known to cost several orders of magnitude more. Finally, the COTS equipment is better suited to ad hoc deployment.

The study concludes by evaluating the use of the Received Signal Strength (RSS) as an alternative for target localisation. It is shown that the signal strengths measured by receivers R and G decay slower than $1/R^2$, different to that predicted by the free space propagation model. Inaccuracies due to environmental effects and SNR fluctuations cause this method to be over-sensitive and reliant on fine tuning. A localisation error of up to 8 m is predicted by repeating the same experiment with RSS measurements. This is substantially worse by approximately 5 m compared to the PDOA algorithm.

Chen draws attention to numerous works that make use of similar principles, not far from that presented in Section 1.3. Blanco *et al.* conduct a study on the effectiveness of beamforming to improve SNR gain and spatial reuse of indoor WAN receivers [24]. The study evaluates directionality and localisation capabilities of an 8-element phased array in a multipath rich environment. Kaspar and Sayrafian-Pour conduct research on direction finding techniques using beamforming with a transmitter-receiver pair in closed environments [25]. All these approaches make use of active phased arrays similar to the SuperDARN radars, while Chen's study makes use of passive receivers.

Maróti *et al.* investigate a similar study to Chen's, where they use two transmitters to identify and locate a target [26]. Each transmitter operates simultaneously at different frequencies to create an embedded low frequency component in the received signal. Collectively, the specialised waveforms are used to achieve localisation errors as narrow as 3 cm up to the range of 160 m.

Overall, the use of a calibrating transmitter to circumvent the initial phase problem is shown to be successful. The inexpensive COTS solution is considered adaptable, reliable and suitable for determining incident angles in ad hoc deployments. The ease of demonstrating the PDOA algorithm using COTS equipment clearly substantiates its use. The ideas and methods discussed in this case study have strong influence over the interferometer designed in Chapter 3.

2.4 Extended Literature

There are a number of other related and useful texts that are not explicitly incorporated into this dissertation. Kono *et al.* discuss the use of multi-frequency Very Long Baseline Interferometry (VLBI) to precisely position the location of spacecraft with a means to

resolve cycle ambiguities [27]. Rose holds a patent which determines the angular position of frequency agile transmitters from non-monopulse interferometric measurements [28]. Lioio *et al.* similarly detail a patent which resolves the ambiguous angle made between two receiving antennas using Time Difference of Arrival (TDOA) techniques [29], while Goodwin describes methods which make use of three and four element interferometers for Direction Finding (DF) purposes [30].

Similar to McDonald, Shepherd discusses the elevation angle discriminations of various layouts and designs of the SuperDARN radar [31] [2]. Burrell *et al.* discuss and assess phase delays introduced by the difference in length of the electrical path between the main and interferometer array in the SuperDARN radar. They make use of existing observations to improve the elevation angle uncertainties from tens of degrees to less than one [32].

Greenwald *et al.* discuss the importance of determining the elevation angle for HF investigations on the ionosphere [33]. Finally, similar to ZACUBE, Bernhardt *et al.* develop the Precision Expandable Radar Calibration Sphere (PERCS), a radar target positioned in space to be identified by ground based HF radars [21] [34].

2.5 Chapter Summary

The chapter began with a discussion on basic HF propagation modes in the ionosphere. Context for the SuperDARN project was developed. The need to accurately measure an elevation angle in order to determine the height of a given plasma density was elaborated. Conditions for atmospheric refraction, particle detection and over the horizon transmission was also addressed.

Three case studies concerning the calibrations of two different SuperDARN radars were addressed. The experiment conducted by Cilliers to determine a phase calibration curve proved to be unsuccessful [15]. A similar non-rewarding conclusion was reached for the ZACUBE-1 project. The proposal to augment the currently existing Buckland Park SuperDARN radar proved to be more promising [2]. Its newer antenna arrangement was shown to have strong beneficial influence on the achievable gains and measurable unambiguous angles in comparison to the older generation SuperDARN radars.

Finally, the chapter concluded on the use of COTS equipment to localise a target through azimuth. An initial phase problem was circumvented using PDOA to achieve a Euclidean range error of less than 3 m. Additional works pertaining to the same field were also highlighted.

Each of the applications discussed in this chapter provide guidance on the design of the interferometer. In Chapter 3, a more mathematical and congruent understanding to the problem at hand is developed.

Chapter 3

Design of The Interferometer

This chapter is separated into four main sections. The first elaborates on the interferometer design requirements. Specifications regarding the characteristics of the planned small-scale experiment are described. The bandwidth, carrier frequency and transmit power are bounded to appropriate and logical constraints.

Section 3.2 expands on necessary background theory. Models regarding the transmitted waveform, environment, received signal, phase extraction and matched filtering are developed. This section forms the foundation to the physical design and implementation of the interferometer.

In Section 3.3, three algorithms are discussed and used to interpret an AOA. Simulations of expected performance are illustrated and examined.

Section 3.4 concludes with the construction of the interferometer. The development of the antennas, FPGA hardware and software are documented. A discussion on the operating system, filter designs and data acquisition methods is included.

Careful consideration has been applied to the consistency of the function names and symbols used within this chapter. Figure 3.5 acts as a global reference of all the signals as they are mentioned.

3.1 Design Requirements

In Chapter 2, the generic SuperDARN radar was identified as being composed of a primary and secondary array. The main consists of 16 elements, and the auxiliary 4 elements. As the two arrays are separated by 100 m, an ambiguity is introduced when determining the AOA. This is primarily due to the distance between the two arrays spanning several wavelengths at HF. This was illustrated in Figure 2.4.

To reconstruct a similar ambiguous AOA phenomenon, three HF antennas are constructed. While the SuperDARN radar is composed of a phased array with beamforming capabilities, the design for this project uses individual isotropic elements. Two of the antennas receive, while the third transmits.

The carrier frequency is required to be set within the HF band as to maintain its resemblance to the SuperDARN radars. It is considered satisfactory to draw conclusive results by expanding an azimuthal angle of arrival as opposed to an elevation angle as measured by the SANA radar. This allows for a more practical field test to be conducted.

The interferometer must effectively expand the ambiguous AOA, whilst at the same time require as few hardware modifications as possible. This is to mitigate the number of required physical adjustments, should any of the results from this dissertation be applied to the larger SANA radar. Where the study in Buckland Park (Section 2.2.3) proposes a modification to the SuperDARN antenna arrangement, this research attempts to achieve the same result by adapting the carrier frequency [2].

3.2 Radar Model

To develop and construct the interferometer, a basic mathematical radar model is required. This is useful as it allows one to anticipate results, predict returns and derive conclusions. Many of the mathematical descriptions provided in this section extend from the work of Wilkinson [35]. Topics involving pulse compression, range resolution, blind range and Doppler are elaborated.

3.2.1 Transmitted Waveform

A Linear Frequency Modulated (LFM) chirp function is chosen as the transmitted signal. Its Radio Frequency (RF) definition is given by [35]:

$$\text{rect}\left(\frac{t}{T}\right) = \begin{cases} 1 & |t| \leq \frac{T}{2} \\ 0 & \text{elsewhere} \end{cases} \quad (3.1)$$

$$v_{TX}(t) = A \text{rect}\left(\frac{t}{T}\right) \cos\left(2\pi\left[f_c t + \frac{1}{2} k t^2\right]\right) \Big|_{A=\frac{1}{\sqrt{T}}} \quad (3.2)$$

where k defines the chirp rate in Hz/s, A the amplitude and T the duration of the pulse. Note that this expression has been defined to have an energy of $0.5 J$. The instantaneous frequency can be determined by examining the derivative of the phase [35]:

$$\phi(t) = 2\pi\left[f_c t + \frac{1}{2} k t^2\right] \quad (3.3)$$

$$f_i(t) = \frac{1}{2\pi} \frac{d\phi(t)}{dt} = f_c + kt. \quad (3.4)$$

Recognising the structure of the instantaneous frequency and how it relates to the definition of the chirp carries importance as it is used in the hardware implementation to dynamically generate the transmitted signal at run-time. Figure 3.1 illustrates this relationship, where the total bandwidth of the LFM chirp is given by $B = f_{\text{end}} - f_{\text{start}}$. The chirp rate is defined in terms of the bandwidth and signal duration: $k = \frac{B}{T}$.

A non-linear (NLFM) chirp can alternatively be utilised to improve the sidelobe levels of a received signal following pulse compression. A more detailed discussion comparing the advantages of NLFM chirps to tapered LFM pulses follows in Section 3.2.8.

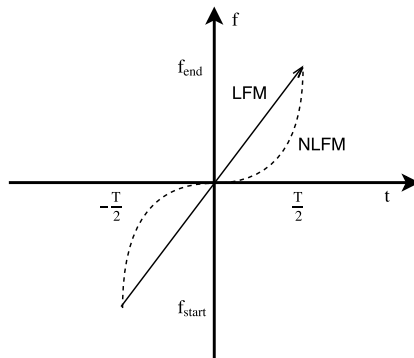


Figure 3.1: Instantaneous frequency of LFM and NLFM chirp [35]

For any given RF signal, its corresponding analytical definition can be derived by considering the positive-only frequency components [35]:

$$Q_{TX}(f) = \begin{cases} 2V_{TX}^+(f) & f > 0 \\ V_{TX}(f) & f = 0 \\ 0 & f < 0 \end{cases} \quad (3.5)$$

The superscript ‘+’ denotes the positive only spectrum of $V_{TX}(f)$, and is used consistently. When converting $v_{TX}(t)$ to its analytical form, Euler’s equations are used to extrapolate the positive-only frequency components to yield [35]:

$$q_{TX}(t) = A \text{rect}\left(\frac{t}{T}\right) e^{j2\pi(f_c t + \frac{1}{2}kt^2)} \Big|_{A=\frac{1}{\sqrt{T}}} \quad (3.6)$$

Inherently the total energy contained within $q_{TX}(t)$ is observed to have doubled to unity in comparison to its RF definition. This is purposefully defined to unify the output of the matched filter used to compress the chirp (Section 3.2.6). All analytical definitions

are shown to contain twice as much energy as their RF counterparts, using Parseval's theorem [35]:

$$\begin{aligned} E_v &= \int_{-\infty}^{\infty} v_{TX}^2(t) dt = \int_{-\infty}^{\infty} |V_{TX}(f)|^2 df = 2 \int_0^{\infty} |V_{TX}(f)|^2 df \\ E_q &= \int_{-\infty}^{\infty} q_{TX}^2(t) dt = \int_{-\infty}^{\infty} |2V_{TX}(f)|^2 df = 4 \int_0^{\infty} |V_{TX}(f)|^2 df \end{aligned} \quad (3.7)$$

Clearly, $E_q = 2E_v$. Careful consideration of this fact is taken into account when performing expected power calculations.

A further generalisation allows one to obtain $v_{TX}(t)$ from its analytical form by simply extracting the real components of $q_{TX}(t)$. Analytical expressions are used for the remainder of this dissertation for the sake of mathematical convenience. The baseband LFM chirp can be obtained by down-converting [35]:

$$p_{TX}(t) = \text{Arect}\left(\frac{t}{T}\right) e^{j2\pi(f_c t + \frac{1}{2}kt^2)} e^{-j2\pi f_c t} = \text{Arect}\left(\frac{t}{T}\right) e^{j\pi kt^2} \Big|_{A=\frac{1}{\sqrt{T}}} \quad (3.8)$$

Figure 3.2 shows the Power Spectral Density (PSD) of an LFM chirp of bandwidth 10 MHz and carrier 15 MHz. Its duration is set to 5 μ s. The top-panel indicates its RF definition, while the lowest panel its baseband equivalent.




Figure 3.2: PSD of $v_{TX}(t)$, $q_{TX}(t)$ and $p_{TX}(t)$ respectively

This illustration is provided to explicitly show the spectral relationship between the various forms, where it is evident that [35]:

$$P_{TX}(f) = 2V_{TX}^+(f + f_c). \quad (3.9)$$

The top-most panel of Figure 3.3 depicts the time domain of $v_{TX}(t)$. The instantaneous frequency is observed to linearly increase with time. The centre panel depicts $p_{TX}(t)$,

where its real and imaginary components are collectively plotted. The lowest panel depicts the PSD of $p_{TX}(t)$, where B has been fixed at 10 MHz. Three pulses of duration 1, 5 and 20 μs are shown.

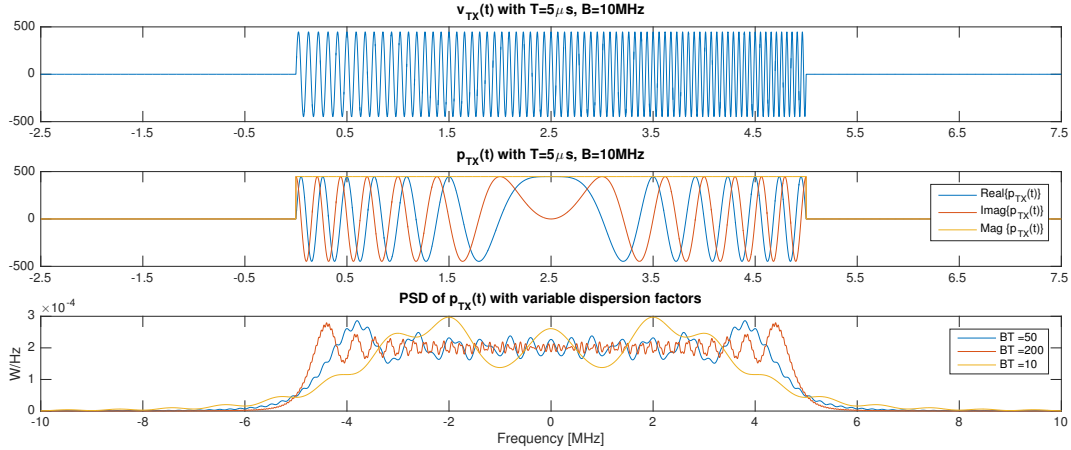


Figure 3.3: $v_{TX}(t)$ alongside its analytical definition, $p_{TX}(t)$, and corresponding PSD

The dispersion factor, $D_f = kT^2 = BT$, is a common measurement used to define the quality of the LFM chirp. The larger the factor, the more box-like the spectrum of the chirp becomes. The symbol B refers to the bandwidth traversed by the LFM chirp and not the width of the chirp's Fourier spectrum. The dispersion factor or pulse compression factor has further implications for the achievable gains after match filtering the received signal (discussed in Section 3.2.6).

In summary, the developed signals $p_{TX}(t)$, $q_{TX}(t)$ and $v_{TX}(t)$ form the foundation to this dissertation, and are used to design and construct many of the FPGA signal processing functions.

3.2.2 Received Signal of Stationary Targets

As with any typical radar system, a signal is transmitted from an antenna, after which it is received following some delay. Usually this signal would have undergone some transformation in amplitude, phase and possibly frequency - depending on whether a return originated from a moving target.

To mathematically model this, the environment is assumed to consist of infinitely spaced point targets. These targets can be thought to cater for direct line-of-sight objects, multi-path returns, specular echoes or even interference. Collectively, each target is modelled as a series of Dirac-delta functions [35]:

$$\delta(t) = \begin{cases} +\infty & t = 0 \\ 0 & \text{else} \end{cases}, \quad \int_{-\infty}^{\infty} \delta(t) dt = 1 \quad (3.10)$$

$$\zeta(t) = \sum_{i=0}^{\infty} \alpha_i \delta(t - \tau_i) \quad (3.11)$$

where α_i and τ_i represent a constant scaling and delay factor due to the i 'th target. The amplitude of the received signal can be correlated to the radar range equation, while τ_i the range of a target [35]:

$$\alpha_i = \sqrt{\frac{P_t G_t G_r \lambda^2 \sigma_i}{(4\pi)^3 K T_s B R_i^4 L_s}} \propto \frac{\sqrt{\sigma_i}}{R_i^2}, \quad \tau_i = \frac{2R_i}{c}. \quad (3.12)$$

The symbols P_t , σ_i , G_t and G_r represent the transmitted power, Radar Cross Section (RCS) of the illuminated target, and the antenna transmit and receive gain. The other symbols K , T_s , B and L_s denote Boltzmann constant, system temperature, signal bandwidth, and system loss respectively. A square root has been applied to the radar range equation as the voltage of the received signal is considered, and not its power. A more detailed radar budget is supplied in Chapter 4, where the units of the respective symbols and values are detailed. The delay factor, τ_i , is halved for one-way propagation.

The transmitted signal, $p_{TX}(t)$, is up-converted and then radiated into the environment or system to acquire the received signal after some time τ_i [35]:

$$q_{RX}(t) = q_{TX}(t) * \zeta(t) = \sum_{i=0}^{\infty} \alpha_i q_{TX}(t - \tau_i) \quad (3.13)$$

which can be further expanded to achieve:

$$q_{RX}(t) = \sum_{i=0}^{\infty} \alpha_i e^{-j2\pi f_c \tau_i} p_{TX}(t - \tau_i) e^{j2\pi f_c t}. \quad (3.14)$$

The value of using the analytical representation becomes apparent when considering down-converting $q_{RX}(t)$ to represent the received baseband signal [35]:

$$p_{RX}(t) = \sum_{i=0}^{\infty} \left[\alpha_i e^{-j2\pi f_c \tau_i} \right] p_{TX}(t - \tau_i). \quad (3.15)$$

The term surrounded in large brackets indicates a complex scalar which is largely dependent on the range of the target. This factor is arguably the most important term as it

contains and defines all the required phasor information needed to perform the interferometric calculations.

3.2.3 Received Signal of Moving Targets

In addition to the stationary interferometric trials, an experiment with a moving target is performed. It is therefore necessary to expand on Section 3.2.2 and develop a model for moving targets within the environment, $\zeta(t)$. To do this, the radar's pulse and target's instantaneous position is modelled as a function of time respectively [35]:

$$r(t) = ct \quad (3.16)$$

$$R(t) = R_0 + vt \quad (3.17)$$

where c , R_0 and v represent the speed of light, initial target position and target radial velocity. The time at which the pulse and the target intersect can be computed by equating $r(t)$ and $R(t)$, and solving for t [35]:

$$t_{int} = \frac{R_0}{c - v}. \quad (3.18)$$

The time taken for a zero-length pulse to return to the radar is given by $t_d = 2t_{int}$. For a non-zero duration pulse, the received signal is modelled slightly differently. Consider the echo of a single finite length pulse given by [35]:

$$q_{RX}(t) = \alpha \times q_{TX}(t - t_d) \quad (3.19)$$

It is worth observing the similarities to Equation 3.13. The leading edge of a return may be disproportionate to its lagging edge depending on the presence of a Doppler shift. The amplitude of the signal is assumed to remain constant and proportionate to the scaling factor, α . An expression for t_d is obtained by considering [35]:

$$t_d = \frac{2R(t - t_d)}{c - v} = \frac{2(R_0 + v[t - t_d])}{c - v} \quad (3.20)$$

and solving for t_d to acquire:

$$t_d = \frac{2(R_0 + vt)}{c + v} = \frac{2R(t)}{c + v}. \quad (3.21)$$

If one substitutes this result into Equation 3.19, one finds [35]:

$$q_{RX}(t) = \alpha \times q_{TX}\left(t - \frac{2[R_0 + vt]}{c + v}\right) \quad (3.22)$$

which can be alternatively written as:

$$q_{RX}(t) = \alpha \times q_{TX}\left(\frac{c - v}{c + v}\left[t - \frac{2R_0}{c - v}\right]\right) \quad (3.23)$$

and significantly simplified as follows:

$$\gamma = \frac{c - v}{c + v} = \frac{1 - v/c}{1 + v/c} \approx (1 - v/c)^2 \approx 1 - 2\frac{v}{c} \quad (3.24)$$

$$\frac{1}{\gamma} = \frac{c + v}{c - v} = \frac{1 + v/c}{1 - v/c} \approx (1 + v/c)^2 \approx 1 + 2\frac{v}{c} \quad (3.25)$$

$$q_{RX}(t) = \alpha \times q_{TX}(\gamma[t - \tau]) \quad (3.26)$$

where $\tau = \frac{2R_0}{c - v}$ in this context. This expression can be used for non-moving targets provided one sets $v = 0$ m/s. The usefulness of this equation becomes apparent when considering its Fourier transform. Before doing so, the scaling constant γ is substituted by performing an exchange of dummy variables and computing the total received energy [35]:

$$E_{RX} = \int_{-\infty}^{\infty} |q_{RX}(t)|^2 dt = \int_{-\infty}^{\infty} \alpha^2 |q_{TX}(\gamma[t - \tau])|^2 dt \quad (3.27)$$

$$E_{RX} = \frac{\alpha^2}{\gamma} \int_{-\infty}^{\infty} |q_{TX}(t - \tau)|^2 dt = \frac{\alpha^2}{\gamma} E_{TX}. \quad (3.28)$$

As the energy in the received pulse remains constant, whether it originated from a moving or stationary target, it can be assumed that $\alpha = \alpha_s \sqrt{\gamma}$, since [35]:

$$\frac{\alpha^2}{\gamma} E_{TX} = \alpha_s^2 E_{TX} \Big|_{\gamma=1} \quad (3.29)$$

The subscript 's' indicates a stationary target. The received signal, $q_{RX}(t)$, is thus scaled vertically by a factor proportional to $\sqrt{\gamma}$. The Fourier transform of the received pulse is therefore shown to be [35]:

$$Q_{RX}(f) = \alpha_s \sqrt{\gamma} \times \mathcal{F}\left\{q_{TX}(\gamma[t - \tau])\right\} \quad (3.30)$$

$$Q_{RX}(f) = \frac{\alpha_s}{\sqrt{\gamma}} Q_{TX}\left(\frac{f}{\gamma}\right) e^{-j2\pi f\tau}. \quad (3.31)$$

The Doppler shift is obtained by introducing the approximation defined in Equation 3.25 [35]:

$$Q_{RX}(f) \approx \frac{\alpha_s}{\sqrt{\gamma}} Q_{TX}\left(f + f\frac{2v}{c}\right) e^{-j2\pi f\tau} \approx \frac{\alpha_s}{\sqrt{\gamma}} Q_{TX}(f - f_d) e^{-j2\pi f\tau} \quad (3.32)$$

$$f_d = \frac{-2v}{\lambda} \quad (3.33)$$

Several observations can be made:

- For targets nearing the radar, f_d and γ are positive in sign. As a consequence, the received pulse is compressed in the time-domain and appears shorter in length. Its vertical amplitude increases to conserve an equivalent amount of energy. The spectral content of the pulse is shifted to higher frequencies.
- For targets travelling away from the radar, the opposite occurs. The pulse is lengthened in time and diminished in amplitude. Similarly, its spectral content is shifted nearer to DC.
- For targets travelling tangentially to the radar's Line of Sight (LOS), no Doppler shift is yielded. Only the radial velocity component of a target influences the Doppler shift.
- The Doppler shift is clearly dependent on the wavelength. For narrow-band radar systems, one can assume λ to be fairly consistent and for the spectral content around the centre frequency to shift by a similar amount. The larger the carrier frequency, the bigger the observed Doppler shift.
- The Pulse Repetition Frequency (PRF) of a coherently pulsed radar system has implication for the maximum measurable Doppler shift. If the system uses synchronous I/Q demodulation, the minimum required PRF should be set to twice the largest two-way anticipated Doppler frequency.
- A matched filter can be used to significantly improve the SNR of a received pulse; however, its performance is degraded by the presence of a Doppler shift. This is discussed in Section 3.2.6 with the use of an ambiguity function and illustrated in Figure 3.11.

Figure 3.4 is used as an example to illustrate the expected returns of 3 different targets, each travelling at speeds outlined in Table 3.1. The speed of propagation was fixed at 340 m/s to mimic an acoustic sonar transducer. This was done to exaggerate the effects of a Doppler shift. For each simulation, an identical transmitted pulse with carrier of 10 kHz was used.




Figure 3.4: Simulated analytical signal, $q_{RX}(t)$, originating from 3 different targets

The relative amplitude of an echoed return from an inbound target is comparatively larger than a stationary object. Furthermore, its spectral content is wider and up-shifted. Its perceived position is falsely interpreted nearer to the receiver. Equations 3.6, 3.26 and 3.31 were used to derive these results, all of which are found to be congruent with anticipated observations.

Table 3.1: Expected Doppler shifts for 3 targets, all located 4.5 m from the receiver

Direction of Target	Target Speed [m/s]	f_d [kHz]	$f_d + f_c$ [kHz]
Away from transceiver	20	-1.176	8.824
-	0	0	10
Toward the transceiver	-35	2.0588	12.058

This concludes the theory on modelling the received pulse for moving targets.

3.2.4 Transceiver Chain, Reconstruction Filter & Range Resolution

Figure 3.5 depicts the relationship between the overall described mathematics and the generalised radar model. The transmitted signal, $p_{TX}(t)$ is observed to be injected into the radar system, after which it is up-converted and radiated through the antenna. It is then convolved with the environment, $\zeta(t)$, down-converted and digitally processed.

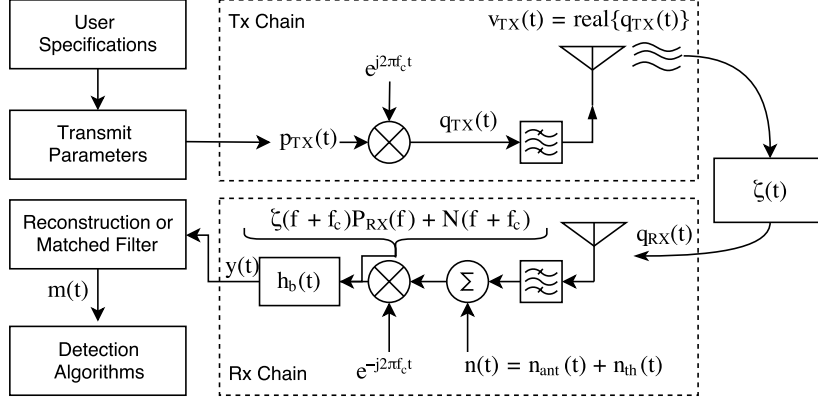


Figure 3.5: Radar transmit and receive overview [35]

Additive noise, $n(t)$, is introduced via the antennas and thermal environment. One could choose to cater for additional sources of multiplicative noise, but this is not done here. The block labelled $h_b(t)$ models the baseband transfer function of the receiver hardware. The overall received signal, $y(t)$, can thus be written as [35]:

$$y(t) = [q_{TX}(t) * \zeta(t) + n(t)]e^{-j2\pi f_c t} * h_b(t) \quad (3.34)$$

$$s(t) = [q_{TX}(t) * \zeta(t)]e^{-j2\pi f_c t} = \left[\int_{-\infty}^{\infty} p_{TX}(\tau) e^{j2\pi f_c \tau} \zeta(t - \tau) d\tau \right] e^{-j2\pi f_c t} \quad (3.35)$$

$$s(t) = \int_{-\infty}^{\infty} p_{TX}(\tau) e^{-j2\pi f_c (t-\tau)} \zeta(t - \tau) d\tau = p_{TX}(t) * \zeta(t) e^{-j2\pi f_c t} \quad (3.36)$$

$$y(t) = \zeta(t) e^{-j2\pi f_c t} * p_{TX}(t) * h_b(t) + n(t) e^{-j2\pi f_c t} * h_b(t). \quad (3.37)$$

This derivation is useful for two different applications, each of which have implication on the interferometer. The first is for imaging and the second for pulse compression. By considering the Fourier transform of $y(t)$:

$$Y(f) = [\zeta(f + f_c) P_{TX}(f) + N(f + f_c)] H_b(f) \quad (3.38)$$

it becomes clear that one can only observe the portion of the environment's spectrum relative to the bandwidth of $P_{TX}(f)$ and $H_b(f)$. For instance, if $P_{TX}(f)$ is modelled as a boxcar function, then one can only effectively observe the spectrum of the environment within $f_c \pm B/2$. Note that the LFM chirp is similarly shown to approximate a rectangular spectrum in Figure 3.3. Targets which are to be detected need to 'resonate' or reflect energy within this bandwidth.

Figure 3.6 illustrates this principle. The left half depicts the spectrum of the environment, transmitted pulse and the received baseband signal. The right half shows the equivalent operations as described by Equation 3.38 which is partially illustrated in Figure 3.5.

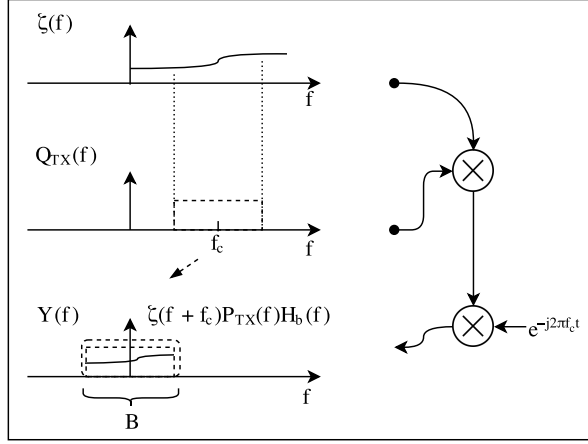


Figure 3.6: Spectral illustration of the process used to image an environment [35]

In imaging applications, the effects of the radar system and transmitted pulse are hoped to be mitigated. For this reason, a reconstruction filter, $H_c(f)$, is developed to reverse the linear effects of the receiver. The filter is applied to the output $Y(f)$ to obtain the impulse response of the scene (with additive noise):

$$Y(f)H_c(f) = Y(f) \frac{1}{P_{TX}(f)H_b(f)} = \zeta(f + f_c) + \frac{N(f + f_c)}{P_{TX}(f)}. \quad (3.39)$$

Unfortunately, in practice this cannot be achieved as the bandwidth of $P_{TX}(f)$ is limited. An infinite amount of bandwidth would be required to fully extract $\zeta(f)$. A more realistic approach would be to redefine $H_c(f)$ as [35]:

$$H_c(f) = \begin{cases} \frac{1}{P_{TX}(f)H_b(f)} & |f| \leq \frac{B}{2} \\ 0 & \text{else} \end{cases} \quad (3.40)$$

in which case a practical solution can be realised. If one momentarily assumes $H_b(f) = 1$, ignores noise and defines $P_{TX}(f)H_b(f)$ as a boxcar function of bandwidth B (see Figure 3.6), one can anticipate the output for a series of point targets to be [35]:

$$G(f) = \zeta(f + f_c) \text{rect}\left(\frac{f}{B}\right) \quad (3.41)$$

$$g(t) = \zeta(t) e^{-j2\pi f_c t} * \text{BSa}(\pi t B) \quad (3.42)$$

where:

$$\text{Sa}(x) = \frac{\sin(x)}{x} \quad (3.43)$$

and:

$$g(t) = \sum_{i=0}^{\infty} \alpha_i \delta(t - \tau_i) e^{-j2\pi f_c t} * \text{BSa}(\pi t B) \quad (3.44)$$

$$g(t) = \sum_{i=0}^{\infty} \left[\alpha_i B e^{-j2\pi f_c \tau_i} \right] \text{Sa}(\pi(t - \tau_i) B) \quad (3.45)$$

The complex scalar wrapped in large brackets of Equation 3.45 contains all the phasor information [35]:

$$\Psi = \arg \left\{ g(t) \right\} = \left(\left(-2\pi f_c \frac{2R_i}{c} \right) \right)_{2\pi} = \left(\left(-\frac{4\pi R_i}{\lambda} \right) \right)_{2\pi} \quad (3.46)$$

where the double bracket notation is used to indicate modulo 2π and enforce that $\Psi \in [-\pi, \pi]$. The phase appears periodic for ranges separated by integer multiples of λ . It is important to note that Equation 3.46 assumes two-way propagation. For one-way transmission, the measured phase is simply halved.

Figure 3.7 illustrates the expected baseband return derived from three targets located at 7500, 7508 and 7516 m respectively. The transmitted signal has a carrier wavelength of 20 m and bandwidth of 3 MHz. The top most panel indicates the real components of the received signal, while the centre panel shows the received magnitude in dB. The lowest plot depicts the argument of the returned signal as governed by Equation 3.46.

It is important to notice that the phase of the received signal does not change over the width of the mainlobe for each return. This is a critical requirement for interferometry as the instantaneous phase between two spatially separated receiving antennas needs to be consistent in order to compute an AOA.

The sudden toggle of π radians occurring at each null in the phase plot is due to the polarity conversion in the real signal as $e^{j\pi} = -1$. Each null occurs on integer multiples of $\frac{c}{2B} = 50$ m.

The -3 dB half power width of the mainlobe is approximately $\delta_r = 0.89 \frac{c}{2B} = 44.5$ m. This measurement is important as it is used to approximate the range resolution of the received signal. Any two targets which are within δ_r metres are superimposed and considered amalgamated and difficult to distinguish.




Figure 3.7: Expected baseband return containing 3 targets, each separated by 0.4λ m

3.2.5 Matched Filtering

The use of a spectral reconstruction filter, $H_c(f)$, was discussed in Section 3.2.4 to recover the impulse response of the scene, $\zeta(t)$. Instead, the application of a matched filter to optimise the peak instantaneous SNR is now elaborated. Shortly thereafter, its relation to the reconstruction filter is drawn. Finally, the use of windowing to improve the sidelobe response in the received signal concludes in Section 3.2.8.

Matched filtering is a technique used to optimise the instantaneous peak signal to average noise power ratio at a given time instance, t_d [35]:

$$\frac{|m(t_d)|^2}{|n_o^2(t)|} = \frac{\text{Peak signal at time } t_d}{\text{Average output noise power}} \quad (3.47)$$

where $n_o(t)$ and $m(t)$ represent the output noise and signal functions. The theoretical derivation of the matched filter is provided in Appendix A.2. The generic structure, as proved by Skolnik, is quoted [12]:

$$h(t) = \Gamma \times z^*(-t + t_d) \quad (3.48)$$

$$H(f) = \Gamma \times Z^*(f)e^{-j2\pi ft_d} \quad (3.49)$$

where the symbol Γ represents a constant scaling factor (often set to 1), and $z(t)$ the signal to be matched. This definition is used to derive the output SNR for a given input

noise signal, $n_b(t)$, and its corresponding potentially non-white PSD, $S_{n_b}(f)$.




Figure 3.8: Generic matched filter model [35]

Figure 3.8 is used to illustrate the assumed model of the matched filter, and represents an expansion of the filtering block shown in Figure 3.5. The matched filter, $H_m(f)$, is momentarily decomposed into $H_1(f)$ and $H_2(f)$. In addition, the definition of $y(t)$ (Equation 3.37) is split into a pure baseband signal and noise component:

$$y(t) = y_s(t) + n_b(t) \quad (3.50)$$

The block labelled $H_1(f)$ is denoted as the pre-whitening filter. It is defined [35]:

$$H_1(f) = \frac{1}{\sqrt{S_{n_b}(f)}} \quad (3.51)$$

such that one can show the output PSD after the first filter stage to be unity:

$$S'_{n_b}(f) = \lim_{T \rightarrow \infty} \frac{|N_b(f)H_1(f)|^2}{T} = S_{n_b}(f)|H_1(f)|^2 = 1. \quad (3.52)$$

The second filter follows from the definition of the matched filter (Equation 3.49) and is designed to optimise the peak instantaneous SNR [35]:

$$H_2(f) = [H_1(f)Y_s(f)]^* e^{-j2\pi ft_d} = \frac{1}{\sqrt{S_{n_b}^*(f)}} Y_s^*(f) e^{-j2\pi ft_d} \quad (3.53)$$

The combined matched filter definition thus becomes:

$$H_m(f) = H_1(f)H_2(f) = \frac{Y_s^*(f)}{S_{n_b}(f)} e^{-j2\pi ft_d}. \quad (3.54)$$

The improvement in the SNR can be calculated by first considering the total noise power by integrating the noise PSD at the output of the matched filter [35]:

$$\overline{|n_o^2(t)|} = \int_{-\infty}^{\infty} |H_m(f)|^2 S_{n_b}(f) df = \int_{-\infty}^{\infty} \frac{|Y_s(f)|^2}{S_{n_b}(f)} df \quad (3.55)$$

and the peak signal value at time t_d via an inverse Fourier transform:

$$m(t_d) = \int_{-\infty}^{\infty} Y_s(f) H_m(f) e^{j2\pi f t_d} df = \int_{-\infty}^{\infty} \frac{|Y_s(f)|^2}{S_{n_b}(f)} df \quad (3.56)$$

followed by combining the two results in a ratio as stated by Equation 3.47:

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} = \int_{-\infty}^{\infty} \frac{|Y_s(f)|^2}{S_{n_b}(f)} df. \quad (3.57)$$

If one assumes $S_{n_b}(f)$ to be modelled by thermal white noise, one can show [35]:

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} = \frac{2E_y}{\eta} \Big|_{S_{n_b}(f)=\frac{\eta}{2}} \quad (3.58)$$

where $\eta = KT_s$ is the input noise power (in $\text{W}\cdot\text{Hz}^{-1}$) and E_y the total energy in $y_s(t)$. In this context K represents Boltzmann's constant and T_s the receiver system noise temperature in Kelvin. Normally η is defined for positive frequencies only and hence its value has been halved to cater for the negative frequencies in the integral of Equation 3.57.

It is interesting to note that the SNR improvement is independent of the signal's bandwidth, but instead on the total energy. This is important to recognise, as the frequency modulation in the transmitted signal, $p_{TX}(t)$, has no effect. Intuitively, transmitting longer pulses or of greater amplitude will improve the peak SNR.

The usefulness of these equations become apparent when one expresses the matched filter in terms of $p_{TX}(t)$. To do this, the return of a single point target irrespective of noise is considered (derived from a simplified version of Equation 3.15 and 3.37) [35]:

$$y_s(t) = \zeta(t) e^{-j2\pi f_c t} * p_{TX}(t) * h_b(t) \quad (3.59)$$

$$y_s(t) = \left[\alpha_0 e^{-j2\pi f_c \tau_0} \right] p_{TX}(t - \tau_0) * h_b(t) \quad (3.60)$$

$$Y_s(f) = \left[\alpha_0 e^{-j2\pi f_c \tau_0} \right] e^{-j2\pi f \tau_0} P_{TX}(f) H_b(f). \quad (3.61)$$

From $Y_s(f)$, the form of $H_m(f)$ becomes evident. It is additionally defined to reverse the effects of the receiver transfer function as well as to peak at time $t_d = \tau_0$ [35]:

$$H_m(f) = \frac{1}{H_b(f)} \left\{ e^{-j2\pi f\tau_0} P_{TX}(f) \right\}^* e^{-j2\pi f t_d} \Big|_{t_d=\tau_0} \quad (3.62)$$

$$H_m(f) = \frac{P_{TX}^*(f)}{H_b(f)}. \quad (3.63)$$

The design of this filter is non-causal. In order for causality, one requires $t_d \geq \tau_0 + T$, where T is the pulse length. Since all the interferometric computations are performed post-capture, this is not strictly required. The received phase and amplitude of $Y_s(f)$ does not feature in the design of $H_m(f)$.

The reconstruction filter, $H_c(f)$, is related to the matched filter as [35]:

$$H_c(f) = \begin{cases} \frac{1}{|P_{TX}(f)|^2} H_m(f) & |f| \leq \frac{B}{2} \\ 0 & \text{else} \end{cases} \quad (3.64)$$

The two filters are shown to have an identical phase response over the passband:

$$\arg \{H_c(f)\} = \arg \{H_m(f)\} = \left(-\arg \{P_{TX}(f)\} \right)_{2\pi} \quad (3.65)$$

and to differ by a scaled magnitude of $|P_{TX}(f)|^{-2}$. If one approximates both the transfer functions of the receiver hardware and transmit pulses to be $\text{rect}(\frac{f}{B})$, then one finds $H_m(f) = H_c(f)$. This simplification becomes more accurate as the chirp time-bandwidth product increases provided $H_b(f)$ remains ‘flat-topped’ over the bandwidth of interest.

The discussion on range resolution in Section 3.2.4 applies to the matched filter as well. Figure 3.9 is provided to show the conceptual difference between the two filters. The matched filter emphasises the spectral shape of $P_{TX}(f)$ to improve the SNR, while the reconstruction filter inverts it to obtain the impulse response of the scene, $\zeta(f)$.




Figure 3.9: Spectral comparison between the reconstruction and matched filter [35]

3.2.6 Pulse Compression & The Ambiguity Function

The decision to make use of an LFM chirp was intentional due to its desirable properties exhibited through a process called ‘pulse compression’. This technique is used to improve the range resolution in the received echo, increase the output SNR and to allow for sufficient energy to be transmitted. In reality, the design of $H_m(f)$ has already taken this into account. It is simply formalised in this section.

To contrast the benefits of the LFM chirp, the use of a single fixed frequency monochromatic pulse is temporarily considered. Its baseband and RF definition is respectively given by [35]:

$$p_{TX_m}(t) = \text{Arect}\left(\frac{t}{T}\right) \quad , \quad v_{TX_m}(t) = \text{Arect}\left(\frac{t}{T}\right) \cos(2\pi f_c t) \Big|_{A=\frac{1}{\sqrt{T}}} \quad (3.66)$$

The range resolution and peak SNR at the output of a matched filter for this particular signal is expected to be [35]:

$$\delta R \approx \frac{cT}{2} \quad , \quad SNR = \frac{2E}{\eta} \Big|_{E=A^2T} \quad (3.67)$$

As evident, one is required to transmit a shorter pulse in order to improve the resolution. On the contrary, this leads to a reduced peak SNR at the filter output and subsequent worsened probability of detection. The properties of the LFM pulse and its matched filter allow for it to act as a suitable solution to this predicament.

The ambiguity function is used to assess the matching performance of a given filter. Its definition for the LFM chirp is derived by considering the transmitted signal:

$$p_{TX}(t) e^{j2\pi f_d t} \quad (3.68)$$

and evaluating the output of the matched filter for a given Doppler shift f_d and time delay τ [36]:

$$|X(\tau, f_d)|^2 = \left| \int_{-\infty}^{\infty} h(\tau - t) p_{TX}(t) e^{j2\pi f_d t} dt \right|^2 \quad (3.69)$$

$$|X(\tau, f_d)|^2 = \left| \int_{-\infty}^{\infty} p_{TX}^*(t - \tau) p_{TX}(t) e^{j2\pi f_d t} dt \right|^2 \quad (3.70)$$

$$|X(\tau, f_d)|^2 = \left| e^{-j2\pi k\tau^2} \int_{-\infty}^{\infty} \text{rect}\left(\frac{t - \tau}{T}\right) \text{rect}\left(\frac{t}{T}\right) e^{j2\pi t(k\tau + f_d)} dt \right|^2 \quad (3.71)$$

One can continue the integration, to resolve the closed-loop solution [36]:

$$|X(\tau, f_d)|^2 = \begin{cases} \left| \left(1 - \frac{|\tau|}{T}\right) \frac{\sin(\pi T(k\tau + f_d)(1 - \frac{|\tau|}{T}))}{\pi T(k\tau + f_d)(1 - \frac{|\tau|}{T})} \right|^2 & \tau \leq T \\ 0 & \textit{otherwise} \end{cases} \quad (3.72)$$

Two relevant properties are derived from this result [36]. The first states that the function is greatest in magnitude at the origin $|X(0, 0)|^2 = E^2$. Since $p_{TX}(t)$ (Equation 3.8) was defined as having unity energy, the matched filter is expected to peak with a value of 1. The maximum gain is exclusively dependent on the total energy in the matched signal. The second property stipulates that the total volume under the surface remains constant [36]:

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |X(\tau, f_d)|^2 d\tau df_d = E^2 \quad (3.73)$$

This is important to consider when designing the transmit signal, since it inherently states that any compensation to reshape the matched filter response will have adverse counteracting effects in order to conserve the total energy.

The monochromatic pulse can be derived from the LFM chirp by simply defining the chirp rate as $k = 0$ Hz/s. Therefore, the ambiguity function of the monopulse is given by a simplification of Equation 3.72:

$$|X_m(\tau, f_d)|^2 = \begin{cases} \left| \left(1 - \frac{|\tau|}{T}\right) \frac{\sin(\pi f_d(T - |\tau|))}{\pi f_d(T - |\tau|)} \right|^2 & \tau \leq T \\ 0 & \textit{otherwise} \end{cases} \quad (3.74)$$

Figure 3.10 illustrates the radar ambiguity diagram for the two signals, where the axes have been scaled in terms of T . The top half depicts the Doppler cut, where only the delay τ is taken into consideration. The lower half shows the delay slice by only considering f_d as the independent variable.

The 3 dB (or ‘half-power’) time resolution of the monopulse is approximately equivalent to the signal duration, $\delta t \approx T$ (as inferred by Equation 3.67). Recall that the bandwidth of this signal is $B \approx \frac{1}{T}$. The narrower the pulse becomes, the more bandwidth the transceiver needs to handle.

The second signal shows the output of the LFM chirp for a dispersion factor of 20. Its 3 dB resolution is compressed to approximately $\delta t = \frac{0.89}{B}$ seconds. The equivalent observation was made when discussing the use of the reconstruction filter (Section 3.2.4).

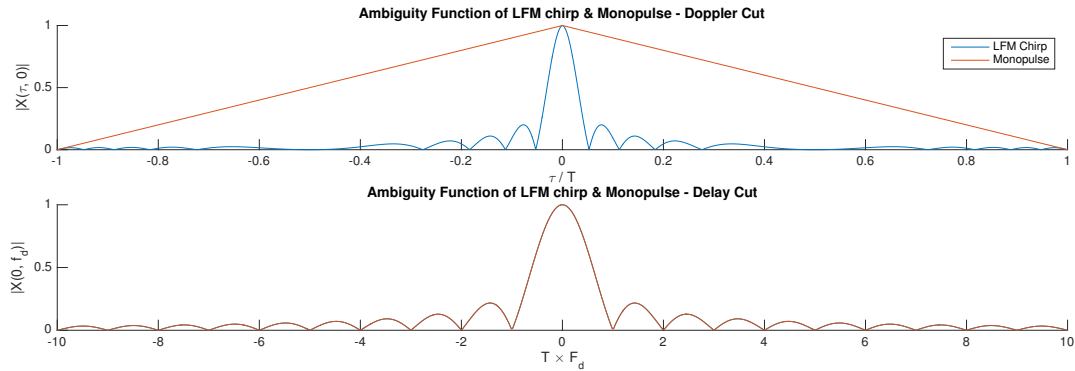


Figure 3.10: Delay and Doppler slices of $|X(\tau, f_d)|^2$, where $T = 20\mu\text{s}$, $B=1\text{ MHz}$

The lower panel of Figure 3.10 shows the adverse effects of a Doppler shift on the gain of the matched filter. Nulls occur periodically on intervals $f_d = \frac{N}{T}$, where $N \in \mathbb{Z}$. The dispersion factor has no influence on this response. Objects travelling at velocities which map to a null will not be detected by the interferometer.

Performing pulse compression with an LFM chirp proves to be advantageous as it allows one to adjust the time (or range) resolution and pulse energy independently of each other. The former property is reliant entirely on the bandwidth B , while the latter is governed by the amplitude A and duration T . As mentioned previously, this is not the case for the monochromatic pulse.

The ambiguity function can be better illustrated by extracting slices for fixed τ and f_d . Figure 3.11 shows this, where the gain of the filter is observed to steadily decay as the Doppler shift increases. The shape of the response however remains consistent.

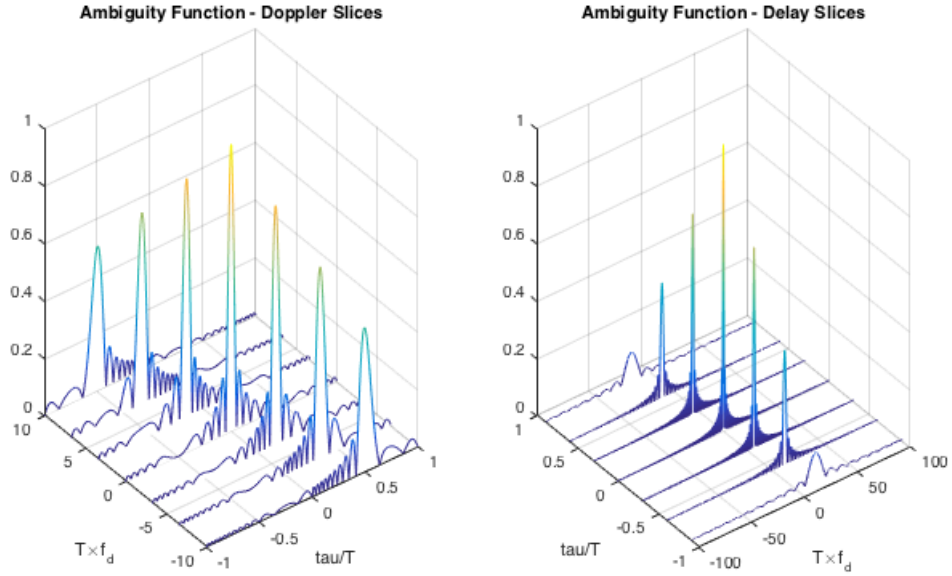


Figure 3.11: Ambiguity plot for $p_{TX}(t)$ for various combinations of τ and f_d

The peak output location drifts proportionately to f_d . This introduced range error, ϵ_r , is given by [36]:

$$\epsilon_t = \frac{-T}{B} f_d \quad , \quad \epsilon_r = \frac{c}{2} \epsilon_t \quad (3.75)$$

This offset can be diminished by either increasing the bandwidth or reducing the pulse length. If, for example, one transmits a 3 MHz chirp, at 18 MHz and of duration 20 μ s, one expects to observe a range error in the order of 3 cm for cruising aircraft. This is considered fairly inconsequential, but becomes relevant to applications using Sonar or for signals which utilise bandwidths in the kHz region. This is formally known as range-Doppler coupling and is an undesired phenomenon.

In contrast, Figure 3.12 is presented to show the equivalent ambiguity plots for the monopulse. Its output response does not exhibit range-Doppler coupling. It does, however, degrade in shape for more extreme Doppler shifts. This worsens target detection performance, as a single object may appear as multiple targets.

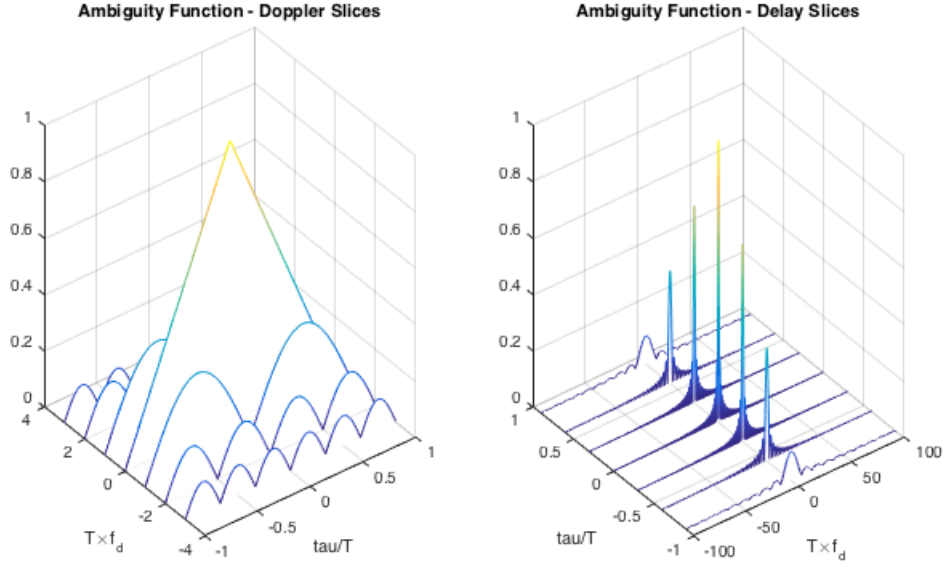


Figure 3.12: Ambiguity plot for $p_{TX_m}(t)$ for various combinations of τ and f_d

The improvement in the SNR, which is partially derived from Equation 3.58, after the matched filter for the LFM chirp is shown to be [35]:

$$SNR_{gain} = \frac{SNR_{output}}{SNR_{input}} = \frac{A^2 T / (\eta/2)}{A^2 / ((\eta/2)B)} = BT = D_f \quad (3.76)$$

The monochromatic pulse does not experience this improvement, since $B \approx \frac{1}{T}$ and hence, $SNR_{gain} = 1$ by using the equivalent derivation.

While the chirp compression appears as an attractive DSP technique, it has a downfall for certain monostatic radars. If the radar is unable to transmit and receive concurrently, it will be blind for the duration of transmission, T . Improving the resolution and SNR gain worsens the blind range. Therefore, this additional trade-off needs to be considered. For the interferometric trials, this is not a concern.

In summary, both the range resolution and SNR after the matched filter for the LFM chirp improve as D_f is increased. Consideration for the range-Doppler coupling property needs to be taken into account when interpreting results.

3.2.7 Coherent Integration

In order to further improve the SNR after the matched filter, coherent integration is performed. To explain the benefits, one considers a scenario in which N equal PRI

recordings are measured:

$$m_0[n], m_1[n], m_2[n] \dots m_{N-1}[n]. \quad (3.77)$$

The discretised definition of the received and processed analogue signal $m(t)$ is used (see Figure 3.5). If one integrates the PRIs together, and considers a sample at any particular index, one finds [37]:

$$u[i] = \sum_{n=0}^{N-1} (Be^{j\theta} + w_n[n]) \quad (3.78)$$

$$u[i] = NB e^{j\theta} + \sum_{n=0}^{N-1} w_n[n] \quad (3.79)$$

where B and θ model the amplitude and phase of the target signal at a given index i . Each coherently integrated signal component adds in phase and magnitude, provided the given target produces an identical echo within each PRI. This is a valid assumption for the trials conducted in Chapter 4. Since the noise, $w_n[n]$, is assumed to be independent and identically Gaussian distributed between PRIs, it has an average power and variance of σ_n^2 [37]. The total power present in the noise component after integration is [37]:

$$E\left\{\left|\sum_{n=0}^{N-1} w_n[n]\right|^2\right\} = E\left\{\left(\sum_{n=0}^{N-1} w_n[n]\right)\left(\sum_{l=0}^{N-1} w_n^*[l]\right)\right\} = \sum_{n=0}^{N-1} \sum_{l=0}^{N-1} E\{w_n[n]w_n^*[l]\} = N\sigma_n^2 \quad (3.80)$$

where E represents the expected value operator:

$$E\{x\} = \frac{1}{N} \sum_{i=0}^{N-1} x. \quad (3.81)$$

This operator is used to compute the power by averaging the sum of the absolute squared integrated signal. The output SNR after this integration step is [37]:

$$SNR_{output} = \frac{N^2 B^2}{N \sigma_n^2} = N \frac{B^2}{\sigma_n^2}. \quad (3.82)$$

Since the SNR before integration is given by B^2/σ_n^2 , one can show the SNR to improve linearly with the number of coherently integrated intervals:

$$SNR_{gain} = \frac{SNR_{output}}{SNR_{input}} = N \quad (3.83)$$

3.2.8 Windowing & Sidelobe Reduction

The first sidelobe of the filter response shown in Figure 3.7 is found to reside 13.2 dB below its mainlobe. This is considered undesirable as the sidelobes have the potential to be interpreted as false targets when applying detection thresholds. Intuitively by reducing them, probabilities of detection can be improved for targets in rich cluttered environments.

In response to this problem, it is common practice to taper the received signal, $m(t)$, as it lowers the sidelobes. As a consequence, it broadens the mainlobe, lowers the achievable SNR and worsens the range resolution. Table 3.2 lists and identifies the losses for various windowing functions, where each value has been computed for a discretised window length of $N = 32$.

Table 3.2: Common tapering function properties relative to a rectangular window [37]

Window Type	3 dB Width	Peak Gain [dB]	Peak Sidelobe [dB]	SNR Loss [dB]
Rectangular	1.00	0.0	-13.2	0.00
Hann	1.68	-6.3	-31.5	-1.90
Hamming	1.50	-5.6	-41.7	-1.44
Kaiser, $\beta=6.0$	1.63	-6.3	-44.1	-1.80
Dolph-Chebyshev (50 dB equiripple)	1.54	-5.6	-50.0	-1.54

The discrete definition of each window function, $w[n]$, is included in Appendix A.3. The SNR loss for each can be computed by first normalising it to ensure $\max\{w[n]\} = 1$ and calculating [37]:

$$SNR_{loss} = \frac{\left| \sum_{n=0}^{N-1} w[n] \right|^2}{N \sum_{n=0}^{N-1} |w[n]|^2} \quad (3.84)$$

The reduction in the SNR can be avoided by tapering and modelling the frequency response of the transmitted signal as a square-root window function instead [35]:

$$P_{TX}(f) = \sqrt{W(f)} \quad (3.85)$$

such that when it is passed through the matched filter, its output models the same response of the chosen window function [35]:

$$P_{TX}^*(f) \times P_{TX}(f) = W(f). \quad (3.86)$$

This method is advantageous as lower sidelobes can be achieved without necessarily reducing the SNR. In the interferometric experiment, the Hamming window is applied to $p_{TX}(t)$, where its influence on the results is discussed in Chapter 5.

Non-linearly frequency modulating the chirp can achieve a similar result. Leśnik *et al.* conduct a conclusive analysis on the design of NLFM signals [38]. Figure 3.13 illustrates the outcome of their research, where they depict the spectral difference between an LFM and NLFM chirp with a bandwidth of 50 MHz each.

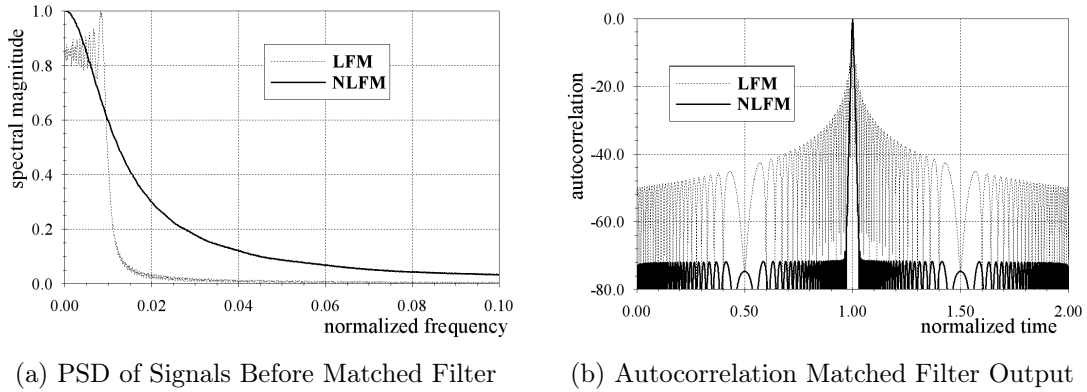


Figure 3.13: Spectral and matched filtering performance of LFM and NLFM pulses [38]

The PSD shows that the rectangular spectrum of the LFM chirp has been softened to envelope a window function shape which results in substantially reduced sidelobes. Evidently, this has a cost in range resolution.

Wang makes similar observations, where trade-offs between different transmit signals used for Chirp Spread Spectrum (CSS) applications are compared. These include 3rd power NLFM chirps and sinusoidally frequency modulated functions [39]. He aims to optimise the performance of CSS based Binary Orthogonal Keying (BOK) communication protocols, where the available RF bandwidth is considered an expensive commodity. Similar conclusions are derived.

Figure 3.14 is included to show the disparity between the non-windowed and tapered LFM pulses. The NLFM chirp has an improved range resolution in comparison to the LFM chirp, however suffers in sidelobe performance. In this instance, a cosine based NLFM chirp has been used, where its baseband definition is given by [39]:

$$p_{NL}(t) = \text{rect}\left(\frac{t}{T}\right) e^{-jBT \cos\left(\pi \frac{t}{T}\right)}. \quad (3.87)$$

Depending on the application, different NLFM definitions can be used and tailored for various optimisations.

The NLFM pulse can be considered more desirable for physical implementation as the amplitude of the signal can be held constant. This is not the case for the tapered chirp. This is useful for applications which involve non-linear power amplifiers which perform better with constant envelope amplitudes.

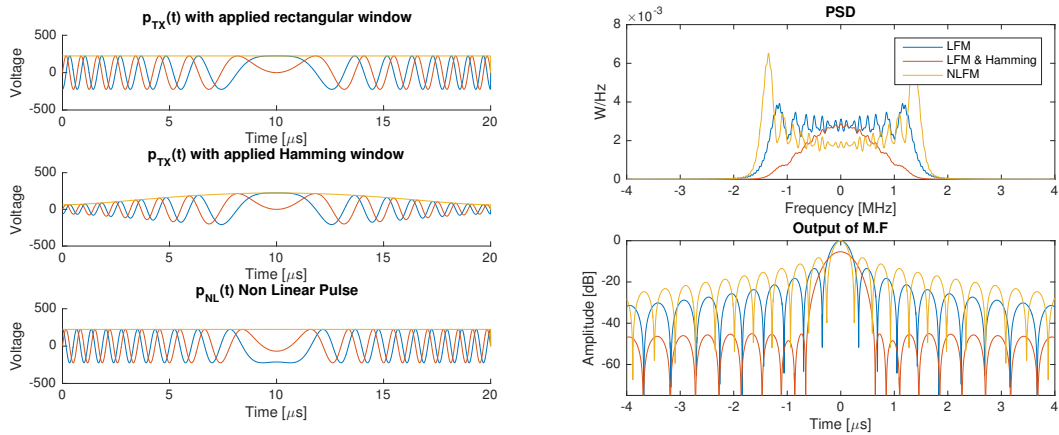


Figure 3.14: Comparison of LFM, tapered LFM and NLFM responses

Concepts of pulse compression apply to Barker codes, which are used to achieve similar improvements with regards to sidelobe reduction. This DSP method segments the monochromatic pulse duration, T , into N smaller interval ‘chip lengths’, denoted T_c . The phase of the signal is then altered by π radians in order to acquire a particular matched filter output with a SNR gain of N and sidelobe reduction factor of $20 \log(1/N)$. Mathematically, the baseband signal is defined as [35]:

$$p_b(t) = \sum_{n=0}^{N-1} \text{rect}\left(\frac{t - nT_c}{T_c}\right) e^{j\phi_n} \quad (3.88)$$

where $\phi_n \in \{0, \pi\}$ represents the particular phase shift for the n^{th} chip. In other words, a new signal is created from a composition of monochromatic pulses with structured 180° phase shifts. The bandwidth in the signal increases to $B \approx N \times 1/T$, as the signal length has been effectively reduced from T to $T_c = T/N$ seconds. As a result, the resolution is improved by N .

Table 3.3 lists the entire set of known Barker codes. The positive sign indicates a 0° phase shift, while a negative sign represents a 180° shift. Figure 3.15 depicts the matched filter output for the longest Barker sequence of length 13.

Table 3.3: List of all known Barker codes [40]

Barker Length	Barker Code	Sidelobe Reduction [dB]
2	+1 +1	-6.02
3	+1 +1 -1	-9.54
4	+1 +1 +1 -1	-12.04
5	+1 +1 +1 -1 +1	-13.98
7	+1 +1 +1 -1 -1 +1 -1	-16.90
11	+1 +1 +1 -1 -1 -1 +1 -1 -1 +1 -1	-20.83
13	+1 +1 +1 +1 +1 -1 -1 +1 +1 -1 +1 -1 +1	-22.27
N	...	$-20 \log(1/N)$

As with all the sequences, the sidelobes appear uniform in amplitude, N times below the peak. An equivalent duration monochromatic signal is included in the figure for reference. There is a clear and obvious improvement in resolution and sidelobe level.

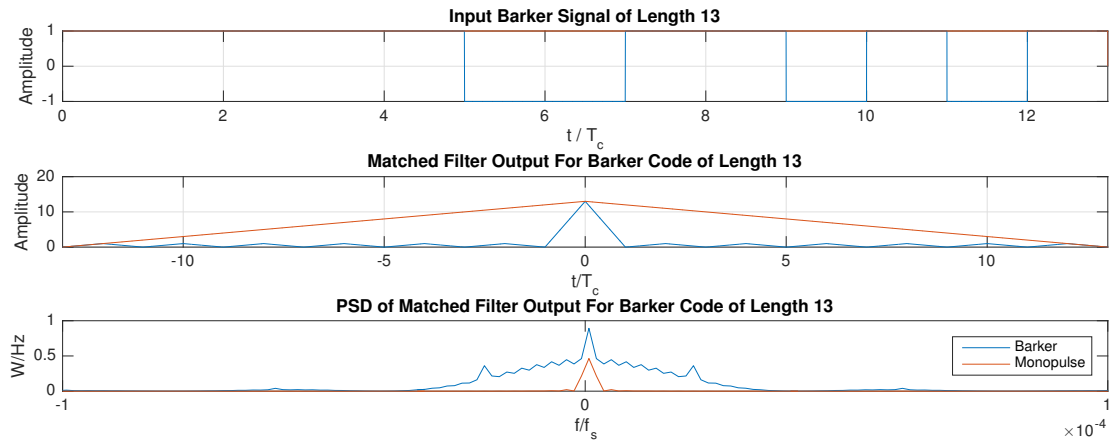


Figure 3.15: Input 13-bit Barker sequence, autocorrelation and accompanying PSD

The Barker sequences are not used in this study, as $P_b(f)$ is known to have poor spectral band-limiting properties in comparison to the LFM chirp. Various other poly-phase coded options exists but are not explored.

This concludes the discussion on the transmitted signal, received signal, environment model, scaling factors, bandwidth, resolution, matched filters, Doppler influences and sidelobe reduction techniques.

3.3 Interferometric Target Detection Algorithm

In this section, three algorithms used to perform the AOA measurements based on the configuration shown in Figure 2.10 are discussed. The geometry consists of two receiving antennas and a single transmitter. The latter forms the beacon target to be identified which generates the measurable one-way EM wavefront.

The two intercepting antennas are connected to a hardware platform consisting of two receiving channels. Each produces a series of complex baseband samples:

$$m_A[n] = [m_{A_0}, m_{A_1} \dots m_{A_{N-1}}] \quad , \quad m_B[n] = [m_{B_0}, m_{B_1} \dots m_{B_{N-1}}] \quad (3.89)$$

$$0 \leq n \leq N - 1 \quad (3.90)$$

where N represents the number of samples in the captured signal. The signal $m_{A/B}[n]$ is assumed to have been conditioned through integration and matched filtering and to be the input to the ‘Detection Algorithms’ block as illustrated in Figure 3.5.

This section adopts a mix of discretised mathematics as the majority of the DSP is performed in Matlab. The influence of sampling the continuous signal, $y(t)$ and $m(t)$, is better characterised in Section 3.4.

3.3.1 Naïve AOA Algorithm

Recall that the phase difference produced from the assumed antenna geometry (Equation 2.13) is given by [23]:

$$\Phi_1 - \Phi_2 = \Psi = \frac{2\pi D \sin(\theta)}{\lambda} + 2\pi K \quad (3.91)$$

where $K \in \mathbb{Z}$ has been additionally defined to account for 2π wrap-around. One can rearrange the equation to define the measured angle in terms of the phase difference:

$$\theta = \arcsin \left[\frac{\lambda}{D} \left(\frac{\Psi - 2\pi K}{2\pi} \right) \right]. \quad (3.92)$$

Note that the argument of the arcsine function must be within $[-1, 1]$:

$$-1 \leq \frac{\lambda}{D} \left(\frac{\Psi - 2\pi K}{2\pi} \right) \leq 1 \quad (3.93)$$

and hence:

$$\frac{-D}{\lambda} + \frac{\Psi}{2\pi} \leq K \leq \frac{D}{\lambda} + \frac{\Psi}{2\pi} \quad (3.94)$$

The measured phase difference, Ψ , always wraps to within the principle range $[-\pi, \pi]$. Therefore, two main boundary conditions are identified by setting $\Psi = \pi$ and $\Psi = -\pi$ respectively:

$$\frac{-D}{\lambda} + \frac{1}{2} \leq K \leq \frac{D}{\lambda} + \frac{1}{2} \quad , \quad \frac{-D}{\lambda} - \frac{1}{2} \leq K \leq \frac{D}{\lambda} - \frac{1}{2}. \quad (3.95)$$

If one assumes $D = \frac{\lambda}{2}$, then $K \subset \{-1, 0, 1\}$. Interpreting this result indicates no ambiguity in the measured angle. This is true for any distances smaller than half a wavelength. If however, one sets $D = 2\lambda$ then one finds $K \subset \{2, 1, 0, -1, -2\}$ which translates to a maximum unambiguous angle of 14.48° . As D increases, the maximum unambiguous angle diminishes. This is congruent with the data provided and discussed in Table 2.1 of the literature review.

Figure 3.16 illustrates this result for all measured phase differences and valid K . The grey zone identifies the unambiguous region. Targets which are situated within 14.48° of boresight are interpreted correctly. However, if they reside at larger angles, say 63.93° , they are perceived to be residing at -5.84° . The receiver cannot determine if this same target is also perhaps at 23.47° or -37.00° .

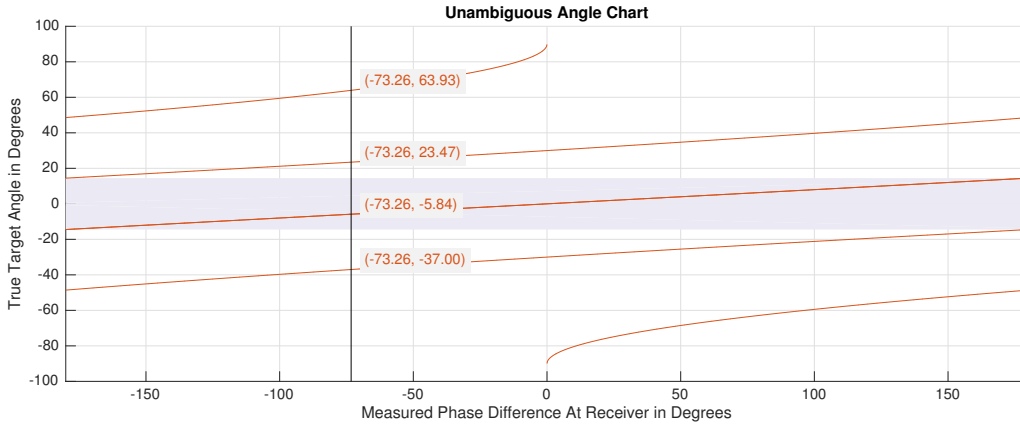


Figure 3.16: Interpreted Angle of Arrival (AOA) verse true target angle for $D = 2\lambda$

The naïve algorithm uses Equation 3.92 with $K = 0$ to determine the AOA. It always assumes the target to be within the first unambiguous zone. This algorithm offers no solution to the problem, and is used to emphasise the ambiguous AOA issue.

3.3.2 Standard Deviation AOA Algorithm

The Standard Deviation (STDEV) based AOA algorithm extends from the naïve method by varying the carrier frequency of the transmitted signal. In order for ease of explanation, the separation distance is fixed to be $D = 30$ m. This is the identical distance used in the interferometric trials of Chapter 4.

The same simulation conducted in Figure 3.16 is repeated for a carrier frequency of 17, 18 and 19 MHz, all of which are superimposed. The unwrap patterns differ between the carriers as the measurement is dependent on λ .

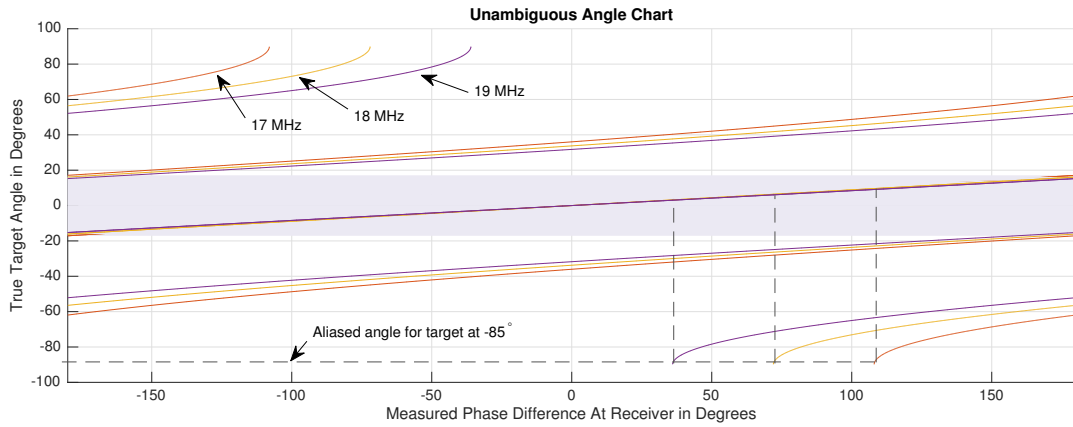


Figure 3.17: Ambiguous AOA chart for 3 carrier frequencies with $D = 30$ m

Naturally, two questions present themselves: how can one detect a target and compute its phase difference in order to calculate the AOA? Furthermore, how can one simplify this process, make it autonomous and reduce the overall required computation? The answers are explained by means of a logically worked example.

An environment of 10 targets is simulated, with each separated evenly between 0.2 and 1.8 km. In addition, each is equally and respectively distributed from -85° to 85° from boresite (see Figure 3.20). The ground-truth locations of the targets are denoted in polar coordinates:

$$R[n] = [R_0, R_1 \dots R_8, R_9] \quad , \quad \theta[n] = [\theta_0, \theta_1 \dots \theta_8, \theta_9]. \quad (3.96)$$

The wavefront distance from a given target to each receiver is slightly different, and is denoted R_A and R_B . Figure 3.18 shows the Cartesian model used to determine these respective ranges. In this illustration, the signal arrives at antenna B before A. The time delay due to each target, τ_i (Equation 3.12), is individually computed for each antenna and channel.

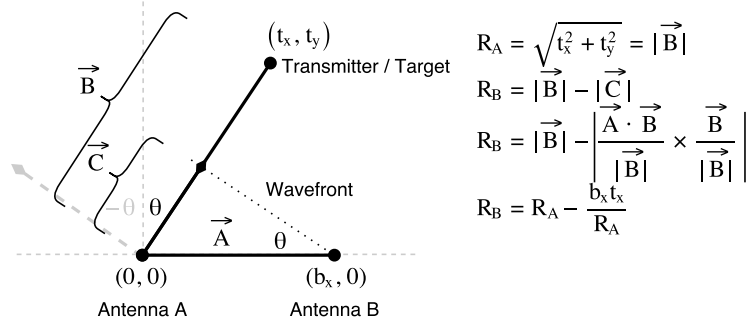


Figure 3.18: Planer wave model used to define target ranges relative to each receiver

If the target is located in the 2nd quadrant (or for negative θ), the wavefront range for the second receiver is redefined to be: $R_B = |\vec{B}| + |\vec{C}|$. All the geometry is done with reference to antenna A.

Three trials are simulated at 17, 18 and 19 MHz as done in Figure 3.17. Equations 3.15 and 3.37 define the returned echo for each target. Influences related to noise and the range scaling factor, α_i , are temporarily ignored.

Figure 3.19 illustrates what one observes after matched filtering. An LFM chirp of bandwidth 3 MHz, carrier 17 MHz and duration $1 \mu\text{s}$ has been used. The range resolution is $\delta_r = \frac{0.89c}{2B} = 44.5 \text{ m}$. A similar plot exists for the other two trials conducted at 18 and 19 MHz, but is not shown.




Figure 3.19: Baseband return consisting of 10 evenly spaced targets

Targets located near right angles, such as $\pm 85^\circ$ produce a visible lead or lag in the received signal of each channel. For instance, at 1800 m, channel B clearly leads A, and hence the target is located somewhere between $[0, 90]$ degrees. For the target at 200 m,

the opposite occurs suggesting the transmitter resides within $[-90, 0]$ degrees.

Each peak is marked with an ‘X’. These locations are automatically determined from Matlab’s `findpeaks()` function which specialises in locating local maxima based on a given set of parameters. The use of a Constant False Alarm Rate (CFAR) detector is suggested as an alternative, but not implemented nor discussed in favour of maintaining scope.

A single frequency is chosen as the primary carrier. In this example, 17 MHz has been designated as the primary, with 18 and 19 MHz being auxiliary frequencies. Target detection is done solely using the primary frequency and on channel A. No target detection is performed on channel B. Ten targets are thus identified, each located at a particular index within $m_A[n]$ (see yellow crosses in Figure 3.19):

$$t_{hit}[n] = [t_0, t_1, t_2, t_3, \dots, t_8, t_9]. \quad (3.97)$$

Further questions arise. Should 18 or 19 MHz rather be used as the primary? Should a target only be declared valid if both channels detect it within several μs of each other? Perhaps target detection should be done on channel B instead? These are all valid concerns. The method detailed and chosen is one of many possibilities, however, and is considered suitable as it is shown to work in Chapter 4.

The phase difference is subsequently computed for each utilised carrier frequency by subtracting the argument in each channel with applied phase wrapping:

$$\Psi[i, n] = \left(\left(\arg \{m_{Ai}[n]\} - \arg \{m_{Bi}[n]\} \right) \right)_{2\pi}. \quad (3.98)$$

In this example, $\Psi[i, n]$ represents a $3 \times N$ matrix. The variable $i \in \{0, 1, 2\}$ maps to the three chosen carriers respectively. Thus, the lower panel of Figure 3.19 illustrates $\Psi[0, n]$. The phase difference at each target index location, $t_{hit}[n]$ is extracted and converted to a naïve angle:

$$\hat{\theta}[i, n] = \arcsin \left[\frac{\lambda[i]}{D} \frac{\Psi[i, t_{hit}[n]]}{2\pi} \right]. \quad (3.99)$$

Note that $\hat{\theta}[i, n]$ is represented by a 3×10 matrix. The function $\lambda[i]$ maps to the respective wavelengths of the 3 chosen carriers. In other words:

$$\lambda[0] = \frac{c}{17 \times 10^6} = 17.65, \quad \lambda[1] = \frac{c}{18 \times 10^6} = 16.66, \quad \lambda[2] = \frac{c}{19 \times 10^6} = 15.78 \quad (3.100)$$

To simplify the explanation, only the angle of the first target positioned at 200 m is explicitly computed. The first column of $\hat{\theta}[i, n]$ is therefore extracted, with the carrier

frequencies being denoted in the subscript:

$$\hat{\theta}[i, 0] = \begin{bmatrix} \hat{\theta}_{17} \\ \hat{\theta}_{18} \\ \hat{\theta}_{19} \end{bmatrix} = \begin{bmatrix} 10.4253 \\ 6.6357 \\ 3.2701 \end{bmatrix} \quad (3.101)$$

Clearly these angles are wrong as the first target resides at -85° . Nevertheless, this aliased result is expected if one identifies the dashed lines drawn in Figure 3.17. The largest carrier of 19 MHz translates to the smallest measured angle of 3.2701° .

If instead, one expands each angle by using Equation 3.92 for all K within the boundaries of Equation 3.95, one expects to find:

$$\hat{\beta} = \begin{bmatrix} \hat{\theta}_{17} & \hat{\theta}_{17}' & \hat{\theta}_{17}'' & \hat{\theta}_{17}''' \\ \hat{\theta}_{18} & \hat{\theta}_{18}' & \hat{\theta}_{18}'' & \hat{\theta}_{18}''' \\ \hat{\theta}_{19} & \hat{\theta}_{19}' & \hat{\theta}_{19}'' & \hat{\theta}_{19}''' \end{bmatrix} = \begin{bmatrix} -84.57 & -24.03 & 10.43 & 50.28 \\ -84.60 & -26.10 & 6.64 & 42.15 \\ -84.62 & -27.99 & 3.27 & 35.68 \end{bmatrix} \quad (3.102)$$

The rows of this matrix contain all the possible aliased angles at which the first target can reside for a given carrier frequency. This can also be graphically interpreted from Figure 3.17. Clearly, the correct measured angles reside within the first column. In order to programmatically determine this solution, a vectored cross-product like subspace is created:

$$\hat{\rho} = \begin{bmatrix} \hat{\theta}_{19} & \hat{\theta}_{18} & \hat{\theta}_{17} \\ \hat{\theta}_{19} & \hat{\theta}_{18} & \hat{\theta}_{17}' \\ \hat{\theta}_{19} & \hat{\theta}_{18} & \hat{\theta}_{17}'' \\ \hat{\theta}_{19} & \hat{\theta}_{18} & \hat{\theta}_{17}''' \\ \hat{\theta}_{19} & \hat{\theta}_{18}' & \hat{\theta}_{17} \\ \hat{\theta}_{19} & \hat{\theta}_{18}' & \hat{\theta}_{17}' \\ \hat{\theta}_{19} & \hat{\theta}_{18}' & \hat{\theta}_{17}'' \\ \hat{\theta}_{19} & \hat{\theta}_{18}' & \hat{\theta}_{17}''' \\ \vdots & \vdots & \vdots \\ \hat{\theta}_{19}''' & \hat{\theta}_{18}''' & \hat{\theta}_{17} \\ \hat{\theta}_{19}''' & \hat{\theta}_{18}''' & \hat{\theta}_{17}' \\ \hat{\theta}_{19}''' & \hat{\theta}_{18}''' & \hat{\theta}_{17}'' \\ \hat{\theta}_{19}''' & \hat{\theta}_{18}''' & \hat{\theta}_{17}''' \end{bmatrix} \quad (3.103)$$

which follows the pattern of a binary counting system. Every possible combination is considered. The number of rows in this resulting matrix is given by:

$$P = \prod_{i=0}^{L-1} 2 \times \text{round}\left(\frac{D}{\lambda[i]}\right) + 1 \quad (3.104)$$

The symbol $L = 3$ represents the number of carrier frequencies used. Note that for every potential target, P row vectors are generated. In this example, $P = 125$. As there are 10 identified targets, a total of 1250 combinations are evaluated.

The matrix $\hat{\rho}$ is then parsed row by row. The row of angles with the lowest standard deviation is selected as the autonomously expanded angle. In this case:

$$\text{Search Operation} \left\{ \hat{\rho} \right\} = [\hat{\theta}_{19} \quad \hat{\theta}_{18} \quad \hat{\theta}_{17}] = [-84.62 \quad -84.60 \quad -84.57] \quad (3.105)$$

The mean of this row is then quoted as the final expanded angle: -84.59° . In this exercise, the AOA has successfully been interpreted for a target which resides at -85° to within 0.41 of a degree. Without using the STDEV procedure to determine the AOA, the theoretical maximum unambiguous measurable angle is 17.10° at 17 MHz.

It is important to recognise that the matrix $\hat{\rho}$ may contain complex angles. This occurs when Equation 3.92 is used to unwrap a given phase difference with a value of K which makes the arcsine undefined. These entries are simply removed from the standard deviation calculation. In this example, only 64 of the 125 rows of matrix $\hat{\rho}$ actually need to be searched, thus reducing computational overhead.

Vector 3.106 represent the true target angles, while 3.107 shows the estimated.

$$\theta[n] = [-85.00, -66.11, -47.22, -28.33, -9.44, 9.44, 28.33, 47.22, 66.11, 85.00] \quad (3.106)$$

$$\hat{\theta}[n] = [-84.59, -66.02, -47.18, -28.31, -9.43, 9.43, 28.31, 47.18, 66.02, 84.59] \quad (3.107)$$

It is interesting to observe that the algorithm produces a better estimate for angles closer to boresight. Figure 3.20 illustrates the discussed example on an artificial Plan Position Indicator (PPI) display. The true targets are located with a '+' symbol. The red 'X' shows the autonomously determined target positions using the standard deviation algorithm.




Figure 3.20: True and estimated target locations derived from simulated example

The diamond symbols indicate the aliased target positions for each of the carrier frequencies. The colours correspond to the same frequencies used in Figure 3.17. The pink circles show the anticipated target positions using the Chinese Remainder Theorem (CRT) algorithm as discussed in Section 3.3.3. The straight lines extending outward from the origin illustrate the theoretical maximum measurable angle when using the naïve algorithm. This is an equivalent boundary to the grey zone depicted in Figure 3.17.

There is a strong correlation between the anticipated and actual target position when using the standard deviation algorithm. Figure 3.21 illustrates a summary of this procedure.

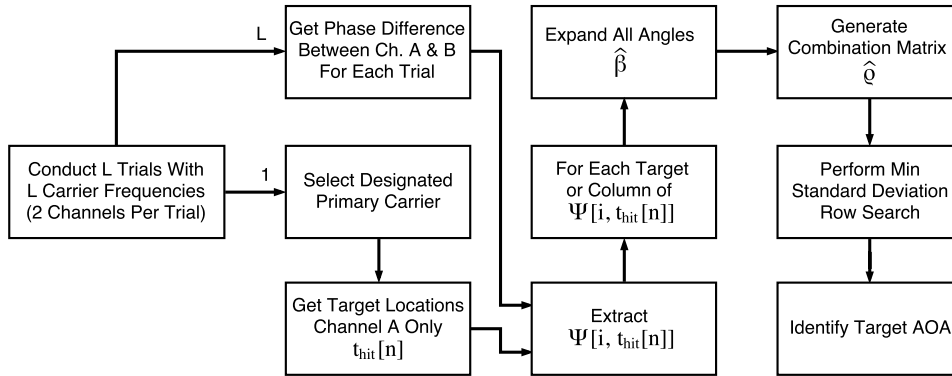


Figure 3.21: Standard deviation AOA algorithm summary

An extensive Monte Carlo simulation is conducted to evaluate the performance of this algorithm. To do this, an identical experiment is run. In this case, however, a single target is migrated from -90° to 90° in $180/500 = 0.36^\circ$ increments along an arc of 300 m. For each angle, a mean square error is calculated from 5 repeated tests:

$$\epsilon_\theta = \frac{1}{5} \sum_{i=1}^5 (\hat{\theta} - \theta)^2 \quad (3.108)$$

If the estimated AOA matches the true target angle θ , then an error of 0 is produced. Figure 3.22 illustrates the outcome of this Monte Carlo simulation for various SNR levels at the output of the matched filter. The SNR is artificially degraded by superimposing additive white Gaussian noise into $m_{A/B}[n]$.

It is evident that as the target drifts from boresight, the performance of the algorithm exponentially declines. The error also increases substantially as the SNR becomes lower, which is to be expected. Provided the SNR remains above 5 dB, this algorithm appears fairly stable. Targets in excess of $\pm 80^\circ$ are expected to produce an estimation error greater than 0.1° .




Figure 3.22: Error in estimated AOA for varying input SNR levels

Figure 3.23 illustrates the use of multiple spaced carriers with a maintained SNR of 5 dB. It is evident that by reducing the gap between the carrier frequencies, the performance worsens. This is expected as the translational mapping space reduces when doing this. One can consider an extreme case where 3 carrier frequencies of 1 Hz apart are utilised. It roughly models a single frequency, in which no new information can be derived to unwrap the AOA.

In addition, this figure depicts an improvement in performance when more carrier frequencies are used. The test case involving 17, 17.125 and 17.25 MHz is significantly improved with the addition of the frequencies 17.375 and 17.5 MHz. The same is noted for the case which additionally includes 18.5 and 19 MHz.

In general, it is also found that using a larger carrier frequency degrades performance. This is expected since the maximum unambiguous angle is proportional to λ .




Figure 3.23: Error in estimated AOA for fixed SNR of 5 dB, but with various carriers

In conclusion, the best performance is realised by maximising the margin between each utilised carrier frequency and acquiring the best SNR at the output of the matched filter. The error is also reduced by using more carriers, L . This comes at the expense of increasing the dimension of the search matrix, $\hat{\rho}$ (Equation 3.104) and computational search time.

A recursive programmatic version of this algorithm is written in Matlab and stored in an online repository which is detailed in Appendix A.1.

3.3.3 Chinese Remainder Theorem AOA Algorithm

The Chinese Remainder Theorem (CRT) is commonly used to circumvent the trade-off between the maximum measurable Doppler velocity and target range of a monostatic radar system. In this context, it is adapted to expand the unambiguous AOA.

Table 3.4 is presented to act as an over simplification of the algorithm. The first row represents a discretised set of angles which effectively range from index 0 to 11. Each index represents an incremented amount of $180/12 = 15^\circ$. This is excessively wide, however it allows for easier explanation.

Three example carrier wavelengths are chosen: λ_0, λ_1 and λ_2 . Each of these is shown to start wrapping after a different number of indices. Similarly to the STDEV algorithm, a much larger AOA can be determined by combing the ‘code’ of the three ambiguous measurements.

Table 3.4: Simplified example of the CRT algorithm used to expand the AOA

AOA Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
λ_0	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2
λ_1	0	1	2	0	1	2	0	1	2	0	1	2	0	1	2
λ_2	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0

If a target resides at an angle represented by index 9, it will be measured to be positioned at index 1, 0 and 1 when operating at λ_0, λ_1 and λ_2 respectively. This unique mapping can be used to correlate the measured results to an unfolded index.

Note that index 12 and above cannot be effectively represented as the codes begin to repeat thereafter. All indices from 0 to 11 can be uniquely mapped. This is clearly an improvement as λ_0 , the best performing wavelength, can only measure up to index 3.

Unfortunately this algorithm degrades fairly substantially in accuracy as the true target angles begin to wrap non-linearly for larger AOAs (see Figure 3.17). Linearity is required for the CRT. Nevertheless, it works for small angles beyond the maximum unambiguous angle. Its performance is contrasted to the standard deviation AOA procedure in Figure 3.20 (represented by the pink circles). For targets at 66° and greater, the algorithm does not find a solution within the span $[-90^\circ, 90^\circ]$ and hence they are not shown. In this study, each index is made to represent 0.5° as opposed to 15° . It is interesting to note that the STDEV algorithm is at its core the CRT algorithm with applied linearisation and searching.

3.4 The Interferometer Hardware Design

In the context of the mathematical radar model, the physical implementation of the interferometer is realised. This process was conducted in multiple stages, where its design naturally developed from a process of redundant and cyclic testing. The hardware, register sets, programmable logic and software was developed in isolation before everything was amalgamated to form the interferometer.

This section details the internal FPGA design of the utilised transceiver, external filters and antennas. Focus and priority is given to the design of the DSP hardware. The order of explanation follows that implemented in reality.

See Appendix A.8.1 for an expanded discussion on the internal design of the transceiver.

3.4.1 Transceiver Platform: The Red Pitaya

The Red Pitaya, developed by STEM^{Lab}, is sold as a portable consumer based LCR meter, spectrum analyser, and oscilloscope. The product, illustrated in Figure 3.24, is roughly the size of a credit card. Its out-of-the-box functionality is stripped clean, leaving its physical hardware to be used to construct the DSP heart of the interferometer. All FPGA development is done in Vivado Studio v.2016.2.




Figure 3.24: The Red Pitaya: Consumer Based Oscilloscope

The device specifications comprise [41]:

- A single Xilinx Zynq 7010 System On Chip (SOC), preconfigured to support a dual core ARM[®] Cortex[®]-A9 CPU with spare FPGA capacity consisting of 17.6k LUTs, 28k logic cells and 80 DSP slices.
- Up to 32 GB of micro-SD storage. A lightweight soft-time Ubuntu/Debian OS flavour is run from this non-volatile memory.

- 1×512 MB of DDR3 RAM.
- 1×USB 2.0, serial console and gigabit Ethernet port.
- 2×14-bit ADCs with a refresh rate of 125 MHz. They are connected via high impedance SMA RF inputs. The ADCs can be configured to either handle $2 V_{pp}$ or $46 V_{pp}$ maximum. The former limit is used for this study as it provides a better voltage resolution per bit. The least significant bit represents $122.07 \mu V$ and $2.81 mV$ for each option respectively.
- 2×14-bit DACs which operate at 125 MHz which are connected via two 50Ω SMA RF outputs. The DACs produce a maximum of $2 V_{pp}$.
- 2×Low speed 12-bit ADCs operating at 100 kHz
- 2×Low speed 12-bit DACs operating at 100 kHz
- 2×SATA connectors (used for daisy-chaining)
- 8×General purpose LEDs
- 16×GPIO pins

The Red Pitaya is powered from a single $5 V$, $2 A$ regulated micro USB port. It can alternatively be powered via its external E2 and ground pins (which has no voltage protection). For the study, the former option is used.

3.4.2 Programmable Logic & Processing System Communication

The complete FPGA schematic encompasses a fairly complex system. To articulate its design, it is important to elaborate on three main components: the Processing System (PS), Operating System (OS) and Programmable Logic (PL). The PS runs the OS and can be considered as fixed hardware, where as the PL is flexible and can be redefined in Verilog. In order for the PS to communicate with the PL, two main Zynq system interfaces are used.

The High Performance (HP) 32/64-bit wide Advanced eXtensible Interface (AXI) is used to direct data from the PL to the PS in a simplex configuration. This interface is designed to operate at high speeds, and can support a phenomenal streaming rate of over $32 \times 125 \times 10^6 = 3.73$ Gbps. The OS, which is dependent on the PS and SD card write-speed, cannot effectively receive data at this rate. This predicament is dealt with below.

The General Purpose (GP) 32-bit AXI interface on the other hand is used to support half-duplex communication from the PS to the PL. Its purpose is to simply configure and read from various control and status registers. It does not handle any data-streams and is designed to operate at significantly slower rates as much of the communication is performed over a shared AXI interconnect.

The 512 MB RAM chip is used as a shared memory space between the PL and PS. It acts as means for the two to exchange data. Each memory address is 32-bits long, and references an 8-bit wide memory location. A minimum of 4 bytes is required to be written to any address (the data spans 4 memory locations). The complexity of physically transferring data between the PL, RAM and PS is encapsulated in the design of the Zynq Intellectual Property (IP) core.

An IP core is analogous to a software library. It consists of pre-written Hardware Description Language (HDL) code and contains a set of adjustable parameters which tailor its intended behaviour. Various common IP cores exist for things like adders, multipliers, FIR filters, buffers and combinatorial logic. One can combine and link various IP cores to each other to form an intricate design of some sort. A given core may well consist of various cores in itself, giving rise to a hierarchical recursive design pattern. In this study, a combination of pre-existing Xilinx IP and custom cores are used to build the FPGA design. An illustration of the internal Zynq IP core is included in Appendix A.4

Figure 3.25 illustrates a simplified version of the memory model used to communicate data between the PL and PS. Note that a 125 MHz clock source is used to trigger the PL which is either provided from the on-board crystal clock or the SATA connector. In this example, the value 0xDDCCBBAA is written to the address 0x40000000. The memory is organised in little-endian format [42]. Bilateral arrows indicate two-way communication, while the single ended arrows illustrate simplex communication. The Synchronisation Logic block represents the circuitry required to daisy-chain multiple Red Pitayas (elaborated in Section 3.4.9).

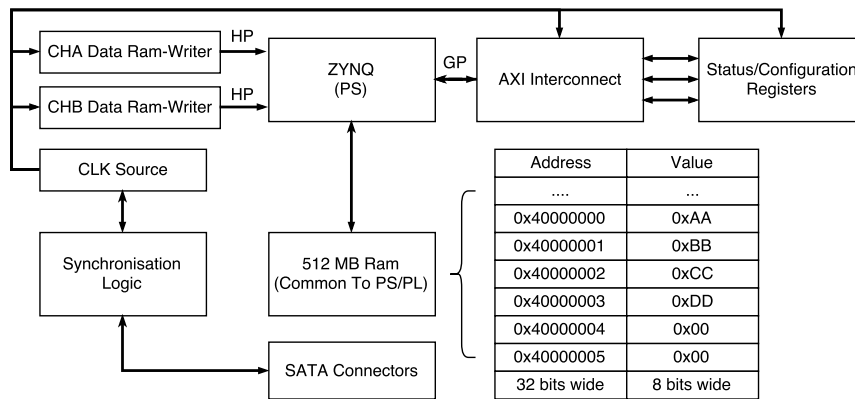


Figure 3.25: Simplified version of the Red Pitaya PS/PL memory model

A given memory address can be written to (or read from) in either one of two ways over the HP/GP interface. The first is to use the PS. Here, the OS writes a value to a physical memory address. This is done using pointers in a C program, where their respective memory locations are mapped using the Linux system command: `mmap()`. Programmatically, the ‘pseudo-code’ may look very similar to as follows:

Listing 3.1: Pseudo code used to communicate memory transfers between the PS and PL

```

#include <stdio.h>
#include <stdlib.h>
#include <stdint.h>
#include <sys/mman.h>
#include <fcntl.h>
#include <unistd.h>

int main(int argc, char *argv[]) {
    //Open file mapped to physical memory
    int fd = open("/dev/mem", O_RDWR | O_SYNC); //Should check returns 0 or errno

    //Define pointers & variables
    uint32_t value = 0;
    void *add_ptr;

    //Map add_ptr to 16KB of memory, starting at address 0x40000000
    add_ptr = mmap(NULL, 4<<10, PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0x40000000);

    //write 32 bits to address 0x40000000
    *((uint32_t *) add_ptr) = 0xDDCCBBAA;

    //Value holds 0x0xDDCCBBAA
    value = *((uint32_t *)add_ptr);

    //Clean up
    munmap(add_ptr, 4 << 10); //Should check returns 0 or errno
    close(fd); //Should check returns 0 or errno

    exit(EXIT_SUCCESS);
}

```

The second alternative is to make use of the AXI communication protocol via the PL. This is the standard used over the HP and GP interfaces because it allows for a reliable transfer of data between different IP cores. Therefore, since the Zynq IP core supports AXI, one can instruct it to write or read data from a particular memory address in RAM from the PL.

Various versions of the protocol exist, including: AXI3, AXI4 and AXI4-Lite [43]. Each requires handshaking, transmit-when-ready flags and buffering queues to avoid packet loss or corruption. These three versions all support memory mapped designs, and require a target address and memory location for data to be transferred.

Similarly, the AXI4-Stream protocol is designed to cater for data-flow paradigms, where there is no concept of a memory address. It is simply used to transfer data directly from one unit to another reliably. The internal design of the interferometer uses a mix of these protocols to communicate between various IP cores.

Figure 3.26 shows an example timing diagram of an AXI4 data transfer used to transfer data from one IP core to another. The letter ‘M’ designates the master, and ‘S’ the slave. In the protocol, the master device will always make a request to a slave - which it is obligated to fulfil. The prefixes ‘AW’, ‘W’ and ‘B’ denote the Address Write, Write and Write Response signals respectively.

Write Burst

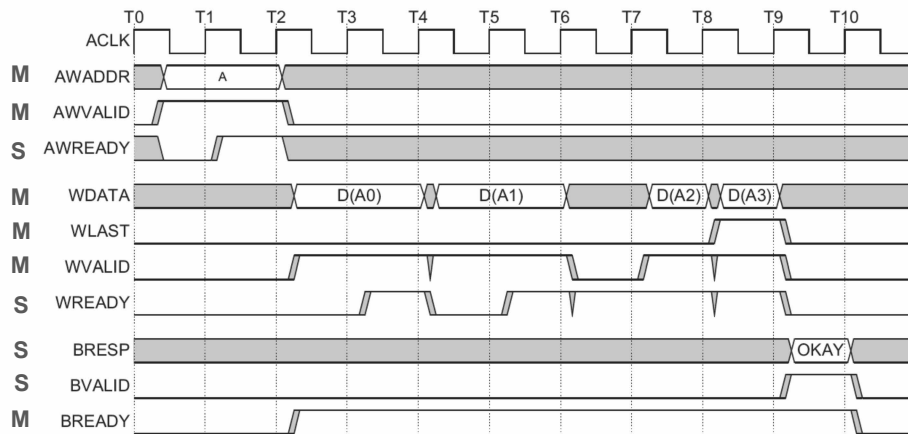


Figure 3.26: AXI4 write sequence from master to slave [44]

In the timing sequence, the master asserts an address and designates it as valid. The slave indicates it is ready to receive the appropriate address information. Thereafter, the transfer of data from the master to the slave begins. Note that the slave’s ‘WREADY’ signal is low whenever the master is not transmitting valid data. This line may be held low for other reasons, such as a full buffer on the slave’s receive line. The master will continuously attempt to transfer the same bit until the slave is ready to receive it. On completion of the transfer, the slave responds with an ‘OKAY’ signal. State engines governing the relationship and behaviour between these signals are well defined by the AXI4 standards [43].

It is important to note that the address space to which the PL writes to RAM is separate to the space used by the OS. Therefore, any clashes or accidental overwrites between the PS, OS and PL cannot occur.

3.4.3 Storing Received Data

In order for received samples to be recorded permanently, data must be transferred from the RAM chip to the SD card. As the OS runs non-deterministically, in comparison to the PL which is routinely clocked at 125 MHz, one is required to synchronise transfers between the two without any loss. This is formally done using a two-stage circular buffer (see Figure 3.27 for explanation).

Within the FPGA design resides an IP core called the ‘Ram-Writer’. It is constructed to accept a 32-bit value and simply record it to RAM. It interfaces directly with the Zynq IP core over the AXI protocol (see Figure 3.25) and is designed to contiguously record chronologically fed data into a 4 MB ring buffer, starting from the address 0x1E000000. For example, a 32 MB input would result in the buffer being overwritten in a cyclic fashion 8 times. An external status register is continuously updated with the current memory location being written to by the Ram-Writer and is stored at 0x40001000. The

data being written to RAM is passed over the HP interface, while the status register is read from the GP interface by the OS.

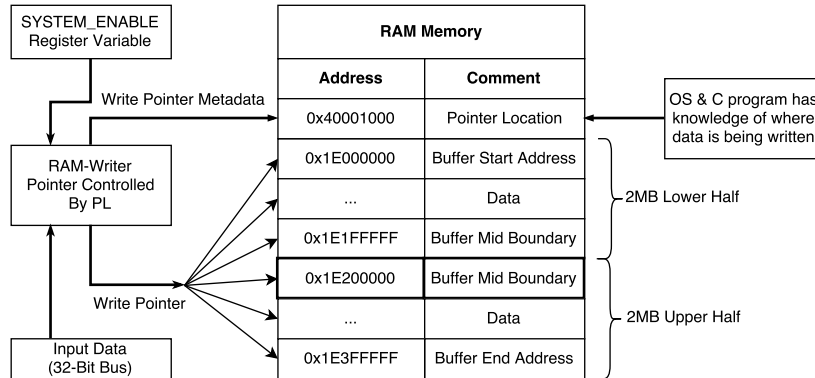


Figure 3.27: Two-stage ring buffer used to write recorded data to the SD card

A C program is concurrently executed on the OS, which copies 2 MB of data at a time from either half of the buffer to file on the SD card. The program is designed to interrogate the position of the Ram-Writer, and to only copy valid data from the half of the 4 MB buffer which is not actively being written to. In effect, the data stored within the circular buffer is transferred and ‘unfolded’ onto the SD card before it is overwritten. The **SYSTEM_ENABLE** register variable is used to reset the Ram-Writer and to instruct it when to start recording data (see Table A.8.2).

A secondary 4 MB buffer, starting at address 0x1E400000 is used for the additional receive channel on the Red Pitaya. The C program is multi-threaded to allow for data to be concurrently exchanged from both buffers via the OS to the SD card.

This extended transfer algorithm is used to circumvent a potential data-race condition. Typically, the time taken for the C program to copy data from RAM to the SD card is several orders of magnitude slower than the HP interface. In other words, data from the PL is overwritten multiple times before the C program can retrieve it via the RAM. A Direct Memory Access (DMA) solution could not be realised as one can only write to the SD card from the OS - it is not accessible from the PL.

In order to validate the algorithm, a fixed 32-bit counter was fed into the Ram-Writer, where each count value was recorded to file. This file was then interrogated to ensure no values in the linear series had been lost. After conclusive experimentation, it was found that the PS could keep up with the PL provided a decimation rate of 5 was utilised. The effective and usable maximum recording rate is experimentally determined to be 25 MHz. This was exhaustively proved to work for recordings in excess of 500 MB. No samples were dropped at this rate.

It is worth noting that an UHS class I micro SD card with a maximum 10 MB/s sequential write speed was used to derive these results. This is a considerably limiting factor.

Should one require lower decimation rates, one is encouraged to use more progressive SD technology which can support up to 90 MB/s write speeds. The low performance SD card allowed for a sufficient Nyquist sampling rate, and hence there was no need to purchase a more expensive card for the interferometric trials.

Alternative solutions such as streaming the recorded data out the 1 Gbps Ethernet port, instead of writing the results to the SD card is also considered. In reality, data was lost when networking at 25 MHz (or equivalently $25 \times 10^6 \times 32 / 2^{30} = 0.745$ Gbps), realising corrupted results and diminished performance. The OS simply cannot prepare TCP packets consistently at this rate. Lightweight network protocols, including UDP and COAP, were considered - but could not compare to the performance obtained by simply writing to the local file system on the Red Pitaya.

3.4.4 Sinusoidal Unit Tests

As the development on the Red Pitaya unfolded, it was evident that a means to consistently and continuously test the interferometer was required. For this reason, two Xilinx Direct Digital Synthesiser (DDS) IP cores were appended to the FPGA design and fed into the system via a set of multiplexers for each receive channel.

The DDS cores are constructed on the FPGA and designed to accept two arguments of which each would adjust the phase and frequency of the generated fixed sinusoidal signal. The 64-bit GEN register, located at 0x40003000, is used to control these cores at runtime. Each argument is specified using $64/4 = 16$ -bits (see Table A.8.2).

The particular output frequency of the DDS core is selected by adjusting the phase increment argument, $\Delta\phi \in \mathbb{N}$ [45]:

$$f_{out} = \frac{f_{clk}\Delta\phi}{2^{N_b}} \Big|_{f_{clk}=125 \times 10^6, N_b=16} \quad (3.109)$$

The phase width in bits, N_b , is set at 16 and f_{clk} to 125 MHz to match the PL clock. The lowest generatable frequency is given by $\Delta f = 1907.35$ Hz (when $\Delta\phi = 1$). This represents the frequency step size. Any desired frequency needs to be a multiple of this amount $f_{out} = N \times \Delta f$, $N \in \mathbb{N}$ provided $0 \leq f_{out} < f_{clk}$. This step size may be considered too coarse, and may be improved by increasing the number of bits N_b . Since these generators are simply used for testing, the step size of 1907.35 Hz is considered adequate.

Being able to inject an artificially generated signal into channels A and B of the Red Pitaya proves to be useful for a number of reasons. Foremost, it allows for the demodulation and down-conversion FPGA designs to be tested. This can be done without the use of the ADCs and RF front-end. No noise is introduced into the inserted signal - it is purely mathematically derived. Secondly, it allows tests to be written after each new

developed feature - where expected outputs can be modelled and evaluated. Therefore, if a new feature breaks a subcomponent in the system or previous test, it can be detected. This encourages future development and stability in the design should anyone extend the project.

In summary, the two DDS cores were implemented purely to provide test tones to be utilised for stable and consistent development.

3.4.5 HF Demodulation Using a Hilbert Transform

Since all the interferometric processing was tailored for baseband signals, a means to translate the received HF signal to DC was required. Two modes of demodulation were implemented on the Red Pitaya. The first made use of simply sampling the input signal, and down-converting it using a Hilbert transform. The second adopted a fully fledged quadrature I/Q demodulator. Differences exist between the two, which are addressed. The former option is now discussed.

Much of the theory on the radar model developed in Section 3.2 is utilised. Where previously all the mathematics made use of the continuous time-domain, the received sampled signal is now considered [46]:

$$v_{RX}[n] = v_{RX_s}(t) = v_{RX}(t) \Big|_{t=nT_s} \quad (3.110)$$

where $T_s = \frac{1}{f_s} = 40$ ns denotes the sample period after a decimation of 5 and n the sample index. The interferometer is fixed to receive from within 12.5 to 25 MHz in accordance with the design requirements (Section 3.1). Due to the discretisation, one needs to additionally take aliasing into account. A bandpass filter is placed before the RF inputs to suppress any unwanted spectral content (including noise) outside this frequency range. The utilised filter is elaborated in Section 3.4.12.

To anticipate the effects of sampling the continuous RF signal, especially with regards to the frequency-domain, the expression for $v_{RX}[n]$ is restructured. An impulse train sequence, which can be thought of as a series of discretised sampled values, is denoted [46]:

$$d(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_s). \quad (3.111)$$

This is useful as it allows the received sampled signal to be written as a sum of weighted Dirac-delta functions [46]:

$$v_{RX_s}(t) = v_{RX}(t) \times d(t) = \sum_{n=-\infty}^{\infty} v_{RX}(nT_s)\delta(t - nT_s). \quad (3.112)$$

Since $d(t)$ is periodic, it can be expressed as a Fourier series [46]:

$$d(t) = \sum_{k=-\infty}^{\infty} c_k e^{jk2\pi f_s t} = \sum_{k=-\infty}^{\infty} \frac{1}{T_s} e^{jk2\pi f_s t} \quad (3.113)$$

as:

$$c_k = \frac{1}{T_s} \int_{-T_s/2}^{T_s/2} d(t) e^{-j2\pi f_s t} dt = \frac{1}{T_s} \int_{-T_s/2}^{T_s/2} \delta(t) dt = \frac{1}{T_s}. \quad (3.114)$$

The Fourier transform of $v_{RX_s}(t)$ is the result of a convolution which yields [46]:

$$V_{RX_s}(f) = V_{RX}(f) * D(f) = V_{RX}(f) * \frac{1}{T_s} \sum_{k=-\infty}^{\infty} \delta(f - kf_s) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} V_{RX}(f - kf_s). \quad (3.115)$$

Therefore, by sampling the real signal, its spectral content is implicitly replicated on integer multiples of f_s . This property is used to demodulate the signal towards baseband. In the digital domain, this result is acquired using the Discrete Fourier Transform (DFT) of length N [46]:

$$V_{RX}[k] = \sum_{n=0}^{N-1} v_{RX}[n] e^{-j\frac{2\pi}{N}kn} \quad , \quad 0 \leq k < N - 1 \quad (3.116)$$

which is considered a cyclic function, with period N , since:

$$V_{RX}[k + rN] = \sum_{n=0}^{N-1} v_{RX}[n] e^{-j\frac{2\pi}{N}(k+rN)n} = \sum_{n=0}^{N-1} v_{RX}[n] e^{-j\frac{2\pi}{N}kn} e^{-j2\pi rNn} = V_{RX}[k] \quad (3.117)$$

as $e^{-j2\pi rNn} = 1$ and $r \in \mathbb{Z}$. Attention is drawn to the argument contained in the exponential $\frac{2\pi}{N}k = \omega$. Here the frequency domain is represented in terms of radians per second. A particular frequency of interest can be extracted by considering $f = \frac{\omega}{2\pi} f_s$. The highest representable frequency, as governed by the Nyquist rate, occurs at π rad/s, or equivalently $f = f_s/2$ Hz. Frequencies above $f_s/2$ are aliased into the principal range $[-f_s/2, f_s/2]$ Hz or $[-\pi, \pi]$ rad/s.

Figure 3.28 illustrates the equivalent, where aliasing is shown to be used advantageously for demodulation. The topmost axes depict the spectrum of the received signal before it has been sampled. After sampling at 25 MHz, the rate at which the PS can record data to the SD card, the spectrum is wrapped and repeated. The letters ‘P’ and ‘N’ are used to emphasise the orientation of the spectral content. As previously mentioned, bandpass filters are used to suppress anything outside the range 12.5-25 MHz to prevent ghosting or undesirable wrapping.




Figure 3.28: Graphical interpretation of the demodulation process

The lowest panel depicts the desired baseband signal. To extract it, several steps are required. The first is to convert the received signal to its analytical form after it has been sampled at 25 MHz. The Hilbert transform of $v_{RX}(t)$ is used to do this [35]:

$$\hat{V}_{RX}(f) = \begin{cases} e^{-j\frac{\pi}{2}} V_{RX}(f) & f > 0 \\ e^{j\frac{\pi}{2}} V_{RX}(f) & f < 0 \end{cases} \quad (3.118)$$

$$q_{RX}(t) = v_{RX}(t) + j\hat{v}_{RX}(t) \quad (3.119)$$

$$Q_{RX}(f) = [V_{RX}^+(f) + V_{RX}^-(f)] + [V_{RX}^+(f) - V_{RX}^-(f)] = 2V_{RX}^+(f). \quad (3.120)$$

Thereafter, $P_{RX}(f)$ is obtained by down-converting. The frequency offset required to acquire the baseband signal needs to take the aliasing shift into account:

$$f_a = -\left(\left(f_c\right)\right)_{-f_s/2} \quad (3.121)$$

$$P_{RX}(f) = Q_{RX}(f + f_a) \quad (3.122)$$

$$p_{RX}(t) = q_{RX}(t)e^{-j2\pi f_a t}. \quad (3.123)$$

To further correct for the aliasing, $p_{RX}(t)$ is conjugated, reflected about the y-axis and low-pass filtered. The first operation is required to compensate for the phase polarity of $p_{RX}(t)$ as its spectral content is initially derived from the negative frequencies of $v_{RX}(t)$ (see Figure 3.28). Recall that any real signal is conjugate symmetric $X(-f) = X^*(f)$ and hence this is needed. Similarly, the reflection about the y-axis is required to correct for the orientation of ‘P’ and ‘N’. This step is strictly not required as the LFM chirp maintains a symmetrical box-like spectrum. For transmitted signals with asymmetrical spectrums, this would be needed. The low-pass filter applied to $p_{RX}(t)$ is simply to discard any undesired frequencies outside the bandwidth of the transmitted signal $|f| \leq B/2$.

The Red Pitaya is configured in this demodulation mode to only capture every 5th ADC value and write it to the SD card. All the DSP operations required to bring the RF signal to baseband are post-processed in Matlab. Figure 3.29 shows the 400 order equiripple low-pass filter utilised in the last step - where over 200 dB of suppression is applied to the undesired frequencies. As mentioned previously, its passable band is made adjustable to dynamically fit the bandwidth of the transmitted signal.




Figure 3.29: Filter response used for demodulating the real sampled signal

Figure 3.30 illustrates the internal design of both the I/Q quadrature demodulator (discussed in Section 3.4.6) and the real sampling Hilbert-based demodulator. The flow of data from the ADCs is passed through either of the two and transferred to RAM, under the control of the multiplexors. The single demultiplexer prevents both demodulators being enabled concurrently. Note that one can also optionally choose to bypass the ADCs and plug a sinusoidal stream of 16-bit values into the system for testing as discussed in Section 3.4.4.

The block labelled ‘Dynamic-Decimator’ serves two purposes. The first is to decimate the signal:

$$v_{RX_{dec}}[n] = v_{RX}[nM] \Big|_{M=5}. \quad (3.124)$$

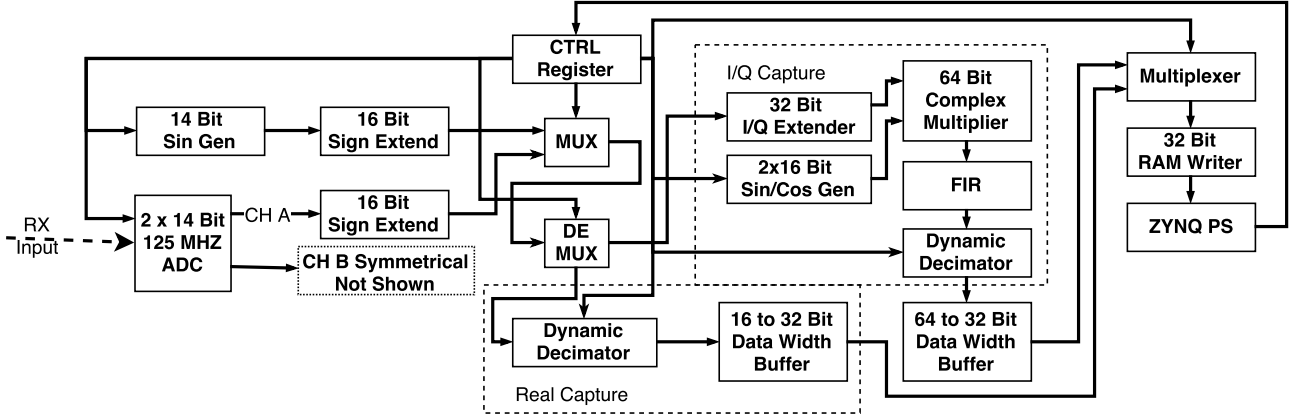


Figure 3.30: Receive-only components of the interferometer FPGA design

The second task of the Dynamic-Decimator is to drop or ignore the first X samples. Since the DAC requires additional clock cycles before it begins generating the transmitted waveform, the first few ADC samples need to be dropped to compensate. Furthermore, it is used to offset the group-delay of the FIR filter used for the I/Q demodulator. The parameter X is adjusted through the CTRL-Register. In reality, the CTRL-Register consists of multiple registers, which is expanded in Table A.8.2.

The ADC generates a 14-bit value, which is sign extended to 16-bits. The actual voltage at the input is simply translated:

$$V_{input} = V_{pp} \left(\frac{\text{ADC Value}}{2^{14}} \right) \Big|_{V_{pp}=2}. \quad (3.125)$$

Since the Ram-Writer only supports 32-bit inputs, two 16-bit ADC samples are concatenated and written to memory before they are interpreted in Matlab via a bitstream file. The effective steady-state data rate is $32 / 2 \times 25 \times 10^6 = 0.3725$ Gbps. Therefore, a 2 MB recording will contain 41.94 ms of sampled data.

This discussed form of demodulation is valid provided the entire spectral content of the transmitted pulse remain within 12.5-25 MHz. More formally stated:

$$12.5 \times 10^6 + \frac{B}{2} \leq f_c \leq 25 \times 10^6 - \frac{B}{2} \quad , \quad |B| \leq 12.5 \times 10^6. \quad (3.126)$$

The monochromatic sinusoidal test input was used to confirm the intended behaviour of the Hilbert-based demodulator. In this case, a $\text{rect}(\frac{t}{T})$ function appears at the output after all DSP operations are performed (as expected from Equation 3.66).

It is worth mentioning that implementing a digital Hilbert filter on the FPGA itself was considered, in exchange removing any off-board processing. This was avoided in order to proportion more DSP slices for the I/Q demodulator.

3.4.6 HF Quadrature Demodulation

The HF quadrature demodulator, in comparison, translates the RF signal on the Red Pitaya directly. Its structure is derived from the definition of the received analytical signal [35] (see Figure 3.2):

$$P_{RX}(f) = 2V_{RX}^+(f + f_c) = 2V_{RX}(f + f_c)H_{LPF}(f) \quad (3.127)$$

$$p_{RX}(t) = 2v_{RX}(t)e^{-j2\pi f_c t} * h_{LPF}(t) \quad (3.128)$$

$$p_{RX}(t) = 2v_{RX}(t) \times [\cos(2\pi f_c t) - j \sin(2\pi f_c t)] * h_{LPF}(t) \quad (3.129)$$

The transfer function $H_{LPF}(f)$ defines a low-pass filter of width $B/2$. Equation 3.129 is implemented and illustrated in Figure 3.30. A single complex-multiplier is used per channel. It is designed to accept two input arguments of 32-bits. The lower 16-bits represents the real value, while the remaining upper 16-bits designates the imaginary value. Mathematically, the output product is produced [47]:

$$p_r + jp_i = (a_r + ja_i) \times (b_r + jb_i) = a_r b_r + ja_r b_i + ja_i b_r - a_i b_i \quad (3.130)$$

The first 15-bits (emphasising not 16) of the sign extended ADC value is used to represent a_r . The remaining bits for a_i are zero filled. The effective product simplifies:

$$p_r + jp_i = (a_r + j0) \times (b_r + jb_i) = a_r b_r + ja_r b_i \quad (3.131)$$

An additional DDS IP core is used to generate an orthogonal cosine and sinusoid which is fed into the multiplier to represent b_r and b_i respectively. The output results in a 64-bit number, of which the first 32-bits represents the real part (p_r) and the latter 32-bits the imaginary part (p_i). No overflow in multiplication can occur since the ADC value is represented with 15-bits and the sinusoids 16-bits (the largest value can be represented as $15 + 16 + 1 \leq 32$ bits). The frequency of the sinusoids fed into the multiplier can be adjusted from the PS.

The DDS core dedicated to generating the unit-test signal (Section 3.4.4) uses 16-bits to represent its phase increment argument, whereas the DDS core for the quadrature multiplier uses 32-bits. Its lowest generatable frequency, and largest error, is approximately $\Delta f = 0.0291$ Hz (Equation 3.109). More bits are used to improve the output frequency precision. By default, each quadrature sinusoid starts with zero phase shift. This is also adjustable from the PS.

An alternative multiplication method involving the COordinate Rotation DIgital Computer (CORDIC) algorithm was also considered. This method implements pseudo-multiplication through a process of iterative additive approximations. It is desirable as it can reduce the number of logic gates and DSP slices required. Chen *et al.* develop a modified version of the algorithm to demodulate AM/FM signals [48]. This was not implemented, although explored, as capacity issues related to the multipliers were not restricting on the FPGA design.

The output of the multiplier is fed into a dual-channel FIR filter, as per Equation 3.129. Each channel filters the real and imaginary samples respectively. The filter serves two purposes: one to decimate the sample rate by 2, and the other to extract the positive spectrum in the received signal. Since the ADCs precede a 12.5-25 MHz analogue filter, no aliasing issues are introduced by doing this as no unsuppressed frequencies above $125/4 = 31.25$ MHz can wrap into the desired range.

The decimation allows for a higher order filter to be fitted onto the FPGA as there is a strong correlation between space and speed. At 125 MHz, only 10-tap filters can be generated before all 80 DSP slices are utilised. However, at 25 MHz, one can support 200-tap filters.

A 20th order equiripple low-pass filter of 10 MHz cut-off is used and operated at $125/2 = 62.5$ MHz. Its frequency response is illustrated in Figure 3.31. The phase linearity over the passband is required to preserve the argument information needed for the interferometric measurements. Only 30 dB of suppression is acquired above the cut-off frequency. During post-processing, the same filter illustrated in Figure 3.29 is used to further clean the baseband signal.




Figure 3.31: Low-pass filter frequency response used in the design of the I/Q demodulator

The output of the 64-bit FIR filter is subsequently further decimated by 10. This is to cater for the slower write-speed of the PS and SD card. Finally, the samples are width converted into two 32-bit samples via a FIFO queue which are chronologically written to a binary file via the Ram-Writer and AXI HP interface.

The overall decimation rate of the I/Q demodulator is $2 \times 10 = 20$. It is interesting to

note, however, that the effective data rate being recorded to file is no different to the previously discussed demodulation method: $64 \times 6.25 \times 10^6 = 0.3725$ Gbps. A single complex sample is represented with 64-bits, whereas the Hilbert-based demodulator uses 16-bits per data point.

The I/Q demodulator is considered more advantageous as it can be used to receive RF signals up to 125 MHz, provided the front-end 12.5-25 MHz filter was removed. One would also need to use a single rate low-pass filter instead of the filter shown in Figure 3.31. Additionally, a higher speed SD card is needed to cater for the larger data rate. These are considered minor adjustments. In summary, the I/Q method lowers the minimum sampling rate in exchange for larger sample sizes.

The two discussed demodulation methods are considered suitable for capturing the transmitted HF signal. Although strictly not necessary, both designs are embedded on the FPGA for flexibility and usability in future projects.

3.4.7 Transmit DAC Buffer Design

Thus far, details of how the RF signal is down-converted to baseband and how the samples are written to the file system on the Red Pitaya have been discussed. Complementary to receiving, a means to instruct the DAC to generate a transmittable waveform is required. In this section, the custom Verilog buffer used to do this is examined.

Similarly to the ADC, the DAC is simplified to accept a signed 14-bit value, and to generate an output voltage on the RF SMA connectors:

$$V_{out} = V_{pp} \left(\frac{\text{DAC Value}}{2^{14}} \right) \Big|_{V_{pp}=2} \quad (3.132)$$

The DAC produces a maximum of 10 dBm for a sinusoidal or LFM chirp, as the characteristic impedance is 50 Ω :

$$V_{RMS} = \frac{V_{pp}/2}{\sqrt{2}} \quad (3.133)$$

$$P_{out} = 10 \log_{10} \left(\frac{V_{RMS}^2}{Z} \right) \Big|_{Z=50} = -20 \text{ dBW} = 10 \text{ dBm} \quad (3.134)$$

Passing the value $(2^{13} - 1)$ and -2^{13} into the DAC produces +1 and -1 volts respectively. As a result, one can generate the transmittable waveform by effectively injecting an array of values into the DAC which are parsed at 125 MHz.

Figure 3.32 illustrates the overall principal of the buffer designed to store the transmitted signal. It is situated on the PL and can either be written to, or read from, at any given

time. Before populating the buffer, it is placed into a reset condition, where the read and write pointers are set to -1. In this state, the buffer will only provide zeros to the DAC.

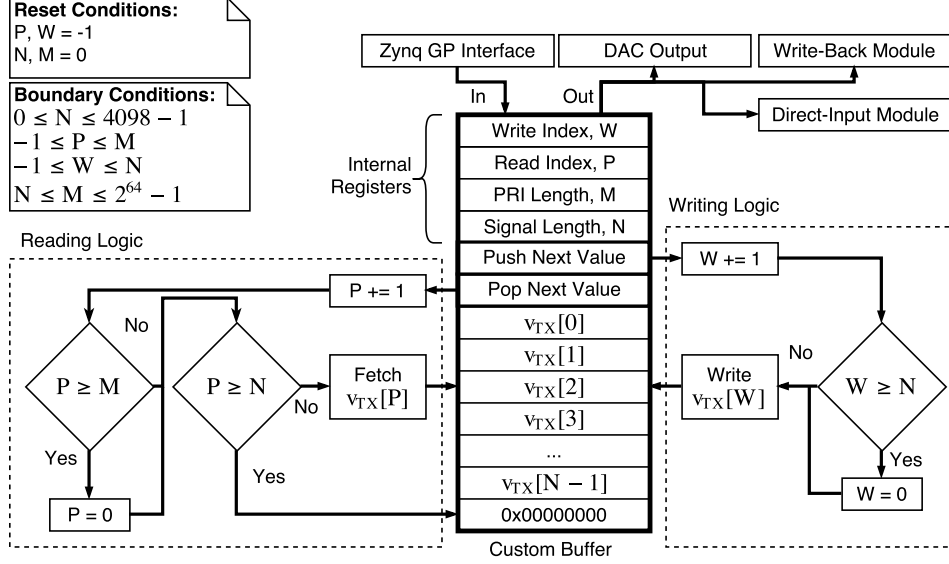


Figure 3.32: Custom cyclic transmitter buffer

The signal duration (in terms of the number of 125 MHz samples) is placed into register N . The buffer is limited to 4098×16 bits or 8 KB and can store a maximum signal duration of $32.78 \mu\text{s}$. This translates to a signal length of 4.92 km and is considered ample for the purposes of this study. One can extend the size to 16 KB, but that is not done in favour of saving resources. The PRI register, M , is appropriately updated to store the number of samples required before the read register is reset. This designates when to re-trigger the transmit signal and when to synchronise other daisy-chained Red Pitayas.

The RF definition of the transmitted signal is discretised and written to the buffer over the GP interface, 16-bits at a time. With every insertion, the write index, W , is incremented. It will overflow if $W \geq N$. In the case of the LFM chirp, it is defined and stored in the buffer (derived from Equation 3.2):

$$v_{TX}[n] = 2^{13} \times \text{Arect}\left(\frac{nT_s}{T}\right) \cos\left(2\pi\left[f_c nT_s + \frac{1}{2}k(nT_s)^2\right]\right)\Big|_{A=1}, \quad 0 \leq n \leq N \leq 4097 \quad (3.135)$$

where $T_s = 8$ ns in this context. Note that the signal has been explicitly scaled to utilise the full dynamic range of the DAC. All subsequent equations will assume this is implicitly done. If one transmits the chirp as a Continuous Wave (CW) by fixing $M = N$,

one expects to measure 10 dBm at the output. The total energy contained within a single pulse, accounting for the 50 Ω output:

$$E_{TX} = \frac{A^2 T}{2 Z} \Big|_{Z=50\Omega, A=1} \quad (3.136)$$

is found to be $T/100$ J. When quantifying these measurements through the receiving chain, one expects to observe twice the power and energy within the CW and pulsed waveforms due to the analytical representation used in the DSP algorithms (Equation 3.7).

If $M > N$, then the buffer will simply produce zeros for the duration of samples which do not contain the transmitted signal ($P \geq N$). After the PRI has expired, the read index is reset, and a new pulse is generated.

This flexibility in design allows the DAC to be utilised as an Arbitrary Waveform Generator (AWG), where the output signal is simply defined within a C array which is imported into the PL during run-time. This ease of use allows for the Red Pitaya to be utilised for many other unrelated applications.

3.4.8 Write-Back & Direct-Input Modules: Features of the DAC Buffer

The output of the buffer is additionally connected to the ‘Write-Back’ and ‘Direct-Input’ modules over the AXI4-Stream protocol (see configuration in Figure A.8.6 and link to Figure 3.32). The former module feeds the RF DAC values directly into the receiving chain - bypassing the ADC. This is done for the first PRI only. Thereafter, the real ADC values are fed into the system. The Direct-Input module on the other hand continuously bypasses the ADC and plugs a copy of the buffer stream into the chosen demodulator.

Figure 3.33 illustrates the received signal captured directly from the Red Pitaya which incorporates numerous of the discussed elements. Channel A is represented in blue, while channel B in red. The top half shows the argument of $p_{RX}(t)$, while the lower half depicts its real and imaginary components. The magnitude envelopes with the dashed line. The PRI and pulse length, T , is fixed at 50 and 20 μ s respectively. The LFM signal contains a bandwidth of 3 MHz.

The Hilbert-based demodulator is used to derive this result. A 50 Ω terminated 7 cm coaxial cable connects the output DAC to channel A’s ADC via a series 3 dB attenuator. Channel B is left disconnected.

The first received pulse is shown to be identical in both channels. This is because the Write-Back module has bypassed the ADC in order to record an internal copy of the pulse. This module allows one to capture a clean version of the transmitted signal without any

externally introduced noise. One can use this to derive a matched filter by inspecting the signal present in the first PRI.




Figure 3.33: Received signal from loop-back test in channel A (Blue) and B (Red)

The Direct-Input module is applied to channel B, where its received signal is sourced directly from the DAC buffer. It suffers zero distortions or losses due to the front-end. Recall that in reality, channel B's SMA input is left disconnected.

The 2nd, 3rd and 4th pulse of channel A shows the actual measured voltage being received without influence from the Write-Back module. The loss in 3 dB of power due to the attenuator results in an expected peak voltage, $V_p = V_{pp}/2$, of:

$$V_p = \sqrt{2 \times 10^{(-20-3)/10} \times Z} \Big|_{Z=50} = 0.7079 V \quad (3.137)$$

This result is derived by rearranging Equation 3.134 and substituting:

$$P_{rec} = P_{out} - 3 \quad (3.138)$$

The phase in the received signal of channel A is slightly offset from the zero-delayed version present in channel B. This is due to the time taken for the transmitted signal to traverse the 7 cm coaxial cable and is to be expected (recall that the phase is dependent on R_i , Equation 3.46).

The Direct-Input module is useful as it allows one to analyse for distortions within the system that are not necessarily related to the ADC, DAC, analogue filters, amplifiers nor antenna. It was used extensively during the developmental phase of this project, and used to verify expected results. It is clear that channel B's response matches that simulated in Figure 3.3.

It is important to recognise that the signal shown in Figure 3.33 is largely unprocessed. It only represents the RF received signal, which has been translated to baseband. When determining phase differences for an AOA measurement, matched filtering and integration is performed to produce and improve the SNR in $m(t)$ (Figure 3.8).

3.4.9 Daisy-Chaining The Red Pitaya

To increase the number of receive channels, additional logic was implemented on the FPGA to enable multiple Red Pitayas to be daisy-chained. Each SATA connector located on the device serves two differential lines, which can be used to both accept and generate a clock signal.

If a given Red Pitaya is marked as a master device, it provides the clock. Only one master should be present in any given configuration. The master also supplies the RESETN (active low) signal, which holds all the coupled slaves in reset or idle until transmission has begun.

Each Red Pitaya can be independently adjusted to activate some clock cycles later after RESETN has transitioned high. This is used to compensate for the time taken for the physical reset signal to pass through all the devices. Each device is synchronised to start recording at an identical time.

The order of initialisation is fairly important and is illustrated in Figure 3.34.

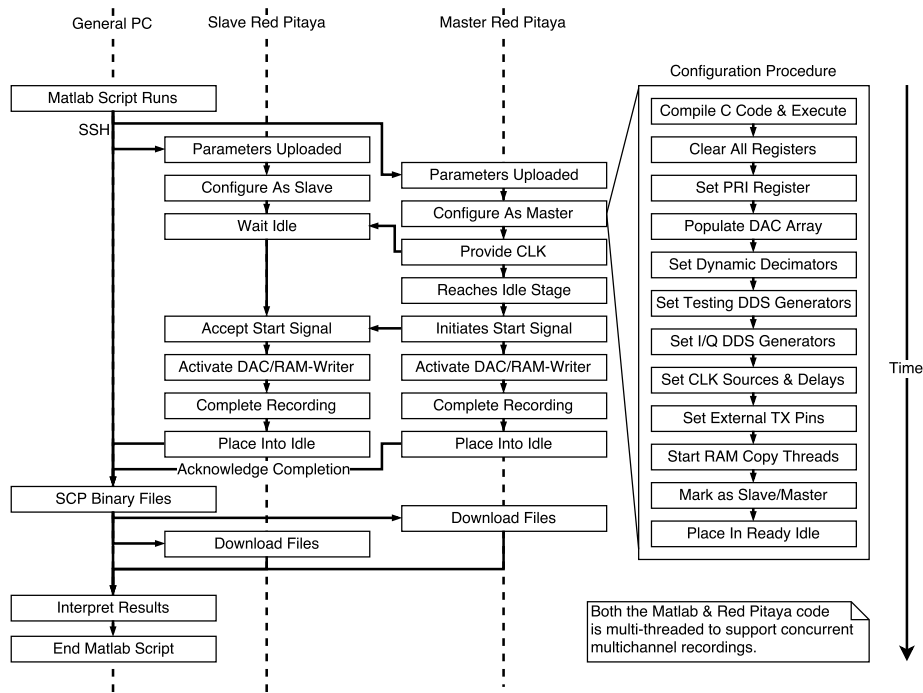


Figure 3.34: Timing sequence of the Red Pitaya initialisation when daisy-chaining

The master device is always configured last as it provides the global RESETN line. If this is not the case, a given slave will not ever receive the high-edge (as it is edge triggered) of the reset line and will perpetually hang. The PL of each slave is clocked by the master's 125 MHz crystal.

A Matlab script is used to configure each device over a Secure Shell (SSH) connection and to ensure the master is activated last. Each Red Pitaya is preloaded to utilise a chosen transmit signal, PRI, record length, demodulation mode and carrier frequency. The illustration also depicts the order in which the registers in the PL are initialised over the GP interface from the supporting C program residing on the devices.

Figure 3.26 illustrates a more global overview of how these devices are connected.

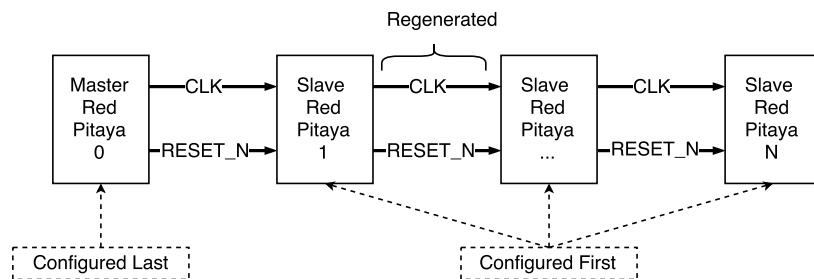


Figure 3.35: Daisy-chain configuration of the Red Pitaya

All the clock signals and RESETN lines are regenerated with each connected device, making physical scaling more realisable. One can choose to define all the devices as slaves, and provide a common clock and external RESETN trigger if they wish.

The illustration in Appendix A.5 is provided to illustrate how the Red Pitaya is used in a synchronised phased array configuration. The identical FPGA design from this project is adopted in a related study on HF surface wave propagation [49].

3.4.10 Switching the Receiver Off During Transmission

The Red Pitaya is configured in the FPGA to drive the general purpose DIO0_N (G18) pin on the start of every PRI or pulsed transmission. The output of this pin can be optionally adjusted in time. Furthermore, the duration for which it is driven high can be adapted programmatically. Figure 3.36 explains the basic principle.

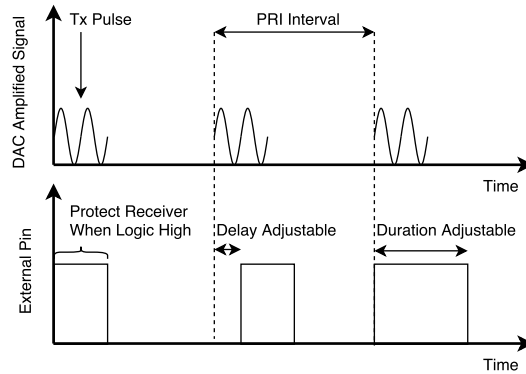


Figure 3.36: External pin can be used to drive a protecting circuit

This pin is used for two purposes. The first is for debugging, while the second is for protecting the receiver during transmission. Nominally the interferometer trials do not require this, as the receiver is designed to handle the full 10 dBm being ejected. If, however, a large amplifier on the transmitter is used on a close proximity bi-static antenna configuration, the receiver hardware may saturate or alternatively be damaged.

For this reason, an external pin is provided on the device to trigger an analogue circuit to redirect all the received power away from the ADCs while transmitting. This pin is denoted as the ‘Is Transmitting Flag’ in Figure A.8.6. It is simply provided as an additional feature.

3.4.11 Discussion on the Hardware Flexibility

Each of the registers outlined in Table A.8.2 can be used in interesting combinations to generate different modes of operation. Some of these include:

- **Arbitrary Waveform Generator (AWG):** The DAC can be configured to operate as an AWG, provided the period of the waveform fits within 4098 samples. In this configuration, the receive chain is simply not used.
- **Closed-Loop Model Testing:** Alternatively, the DAC can be utilised to inject a known signal into a system. The RF front-end can then be used to receive the output and model its transfer function or response.
- **Continuous Wave Radar:** As the DAC buffer can be defined to operate in a cyclic manner, the device can be utilised to operate a Continuous Wave (CW) or Frequency Modulated CW radar mode.
- **Pulse Doppler Radar:** The DAC and ADC can be utilised to form a pulsed-Doppler radar system. The interferometer makes use of this mode.

- **Low Probability of Intercept Radar:** The DAC can be configured to adopt a new waveform between each PRI. For this reason it can be utilised as a low probability of intercept radar as each successive pulse is less predictable.
- **Active Phased Array:** As multiple Red Pitayas can be synchronised, they can be used to receive on numerous antennas in an active phased array configuration. See Appendix A.5 for an example.

3.4.12 Front-end Filter Design

To prevent issues associated with aliasing and sampling $v_{RX}(t)$, a 12.5-25 MHz bandpass filter is required to precede the ADC inputs. This forms the first stage to down-converting the RF signal, as shown in the explanation of Figure 3.28.

A 6th order Bessel filter is constructed from a series of inductors and capacitors. Its circuit diagram is shown in Figure 3.37. The physical equivalent is shown in Appendix A.7.




Figure 3.37: Circuit diagram of a 12.5-25 MHz anti-aliasing bandpass filter [49]

The magnitude and phase response of the filter is illustrated in Figure 3.38. It is clear that there is a strong correlation between the simulated filter and realised empirical response. The linear phase over the passband is essential to preserving the phase information within the received signal.




Figure 3.38: Simulated & practical measurement of a 12.5-25 MHz anti-aliasing filter [49]

The 3 dB half power points are shown to reside within 12.26-24.35 MHz, meeting the required specifications. The filter is also expected to have an insertion loss within the order of 1 dB over frequencies of interested, which is considered acceptable. The filter is designed to have a -36 dB per decade roll-off.

3.4.13 HF Antenna Design

Zubeida conducts an in-depth analysis on the design of an HF conical antenna for surface wave radar [50]. This study extends from this project, and adopts an identical antenna to be used for the two receivers and transmitter. Its schematic is shown in Figure 3.39. The actual physical antenna can be seen in Figure A.8.2.




Figure 3.39: Conical antenna design used for the interferometric trials [50]

The physical structure expands 3.45 m in height, and 1.5 m in diameter around its centre. It is designed to maintain an omni-directional directivity of 4.5 dBi or greater and to be suitable to operate over a wide bandwidth [50]. The spokes are provided for support and do not conduct. The radials are utilised to provide a sufficient ground-plane.

Figure 3.40 shows the empirically measured reflection coefficient (S11) of the antenna. This measures the ratio between the reflected power to the forward radiated power: $\rho = \frac{P_{-}}{P_{+}}$.




Figure 3.40: Reflection coefficient curve of the conical antenna [50]

The -5 dB points are shown to reside approximately between 13.6 and 19.98 MHz. Any frequencies transmitted between this band are considered to effectively radiate. Clearly, the most power will be transmitted at 16.5 MHz. The wide bandwidth of this antenna structure is considered desirable as multiple frequencies can be utilised to improve the standard deviation AOA algorithm (Section 3.3.2). No matching circuit was utilised on this antenna.

The reception of frequencies around 48 MHz is not considered problematic as the front-end 12.5-25 MHz filter is expected to mitigate these frequencies by over 30 dB.

3.5 Chapter Summary

The chapter began with an outline of the interferometric requirements. These were mainly specified to match the operating parameters of the SANAE radar as to maintain relevance and resemblance for this research. An extensive theoretical radar model was subsequently developed to anticipate and predict results. This catered for resolution, sensitivity to range-Doppler coupling, pulse compression, windowing, SNR improvements and output power predictions. Three AOA algorithms were developed, of which the standard deviation alternative was shown to successfully expand the maximum measurable AOA through frequency agility.

The theoretical radar model was then applied to the FPGA design, where each IP core on the Red Pitaya was amalgamated to form an effective transceiver. The developed register set, data paths, sample rates and synchronisation capabilities were discussed.

The chapter concluded with a brief elaboration on the front-end analogue anti-aliasing filter and wide-bandwidth conical antenna design. Thus far, all the necessary theory and practical requirements needed to construct and operate a fully coherent pulsed-Doppler radar system which can be used for interferometry has been covered.

Chapter 4

Experiments and Results

Two categories of experiments are described in this chapter to assess the accuracy of the interferometer. The first deals with tasks dedicated to evaluating measurable and verifiable results within a lab environment. It does not utilise any antennas.

The second category, detailed in Section 4.2, consists of an expansive set of outdoor trials which are used to evaluate the performance of the three AOA algorithms (naïve, STDEV, CRT) mentioned in Section 3.3. Several parameters are systematically adjusted, where an analysis of the results is presented in Chapter 5. Images of the outdoor trial are shown in Appendix A.8.

Section 4.3 includes an additional task that attempts to detect moving aircraft (such as a large Boeing 737 jet) so that the HF interferometer developed in this project could be used as a pulse-Doppler radar. This task forms only a minor component of the project but preliminary results are included as they demonstrate the potential versatility of the interferometry, and provide a basis for further investigation to refine the overall system.

4.1 Pre-Field Test Trials

To validate the measuring capabilities of the Red Pitaya, several basic tests are defined. Each is identified by a unique code, starting with PFT000. The prefix PFT stands for ‘Pre-Field Test’.

4.1.1 General Closed-Loop Testing on the Red Pitaya

Tests PFT000 to PFT019 are dedicated to calculating the received power, energy, target location and phase difference. Table 4.1 describes the lists of each chronologically performed action to derive each trial.

The centre column designates the main active register variables (see Table A.8.2). It is simply used to distinguish which signals are artificially fed into channels A and B of the transceiver, and which mode of demodulation is adopted. The unmentioned registers such as DAG, SOF, GEN and QUA are simply configured to allow for an RF LFM chirp to be transmitted and received correctly.

Several components are utilised:

- $2 \times 50 \Omega$ T-junction terminators which are attached to the receive channels on the Red Pitaya for impedance matching purposes (seen in Figure A.8.4).
- $1 \times$ RF power splitter, which is used to split the transmit signal from the DAC into two output channels which can be connected to the inputs on the RF transceiver.
- 3×15.4 cm 50Ω coaxial cables
- 2×18 m 50Ω coaxial cables
- 1×0.665 , 3.12 , 6.22 and 208.42 m 50Ω coaxial cable

For each trial, both the PRI and signal duration is fixed at $20 \mu\text{s}$. Therefore, each pulse immediately follows after the next. A bandwidth of 3 MHz is used with a primary carrier frequency of 17 MHz. No tapering is applied to the transmitted signal. The Write-Back module is additionally disabled. The average power and energy in each received PRI is recorded by computing:

$$P_y = \frac{1}{ZN} \sum_{i=0}^{N-1} y[n]y^*[n] \quad E_y = P_y \times \text{PRI} \quad (4.1)$$

where $Z = 50 \Omega$. The value for $N = f_s \times \text{PRI}$ is respectively 500 and 125 provided the Hilbert-based and I/Q demodulator is used. In addition, the estimated AOA is computed using the algorithms outlined in Sections 3.3.1 and 3.3.2. The pseudo antenna distance, D , is assumed to be 10 m. In the outdoor trials, this value changes to 30 m. Two auxiliary carriers of 18 and 19 MHz are used for the latter algorithm. The phase difference between channels A and B at the peak output of the matched filter is quoted in metres (see Figure 4.3). This is stipulated using one-way propagation, and is derived from Equation 3.46:

$$\Delta R = -\frac{\Psi\lambda}{2\pi}. \quad (4.2)$$

Table 4.2 lists the results from all the closed-loop measurements. Figures 4.1 to 4.3 illustrate the output of PFT001, before and after integration. Trials PFT002 to PFT019 all have similar figures - but they are not included as Table 4.2 contains all the necessary information. Each recording is 2 MB large, and consists of 41.94 ms of sampled data per channel.

Table 4.1: Outline of general pre-field test trials

Experiment Code	Main Active Register Variables	Description
PFT000	DIRECT_INPUT_ENABLE_A DIRECT_INPUT_ENABLE_B HILBERT_DEMOD_ENABLE	The transmitted signal is fed to each receive channel internally through the FPGA. This bypasses the external ADCs and DAC.
PFT001	DIRECT_INPUT_ENABLE_B HILBERT_DEMOD_ENABLE	A 66.5 cm coaxial cable is connected from the DAC output directly to the ADC on channel A via the 50 Ω junction terminator.
PFT002	DIRECT_INPUT_ENABLE_A HILBERT_DEMOD_ENABLE	The 66.5 cm coaxial cable connected to channel A is switched to channel B.
PFT003	HILBERT_DEMOD_ENABLE	The 66.5 cm coaxial is removed. A 15.4 cm cable is then connected from the DAC output to the splitter input. In addition, the splitter outputs are connected to channels A and B with two identical 15.4 cm coaxial cables.
PFT004	HILBERT_DEMOD_ENABLE	The inputs to channels A and B are reversed. Thus, the two 15.4 cm cables extending from the splitter's outputs are swapped.
PFT005	HILBERT_DEMOD_ENABLE	The two 15.4 cm coaxial cables are once again reversed. The 15.4 cm cable connecting the transmitter to the splitter is replaced with the 208.42 m cable.
PFT006	HILBERT_DEMOD_ENABLE	The 15.4 cm cable connecting the splitter to channel A is replaced with a 6.22 m coaxial cable.
PFT007	HILBERT_DEMOD_ENABLE	The 6.22 m and 15.4 cm cables are reversed, switching the inputs to channel A and B.
PFT008	HILBERT_DEMOD_ENABLE	The channel inputs are once again reversed. The 6.22 m cable is extended to 9.34 m with extension of a 3.12 m coaxial cable.
PFT009	HILBERT_DEMOD_ENABLE	The 9.34 m and 15.4 cm cables are reversed, switching the inputs to channel A and B.
PFT010 - PFT019	-	This is a repeat of PFT000 to PFT009, except that the I/Q demodulator is utilised by setting HILBERT_DEMOD_ENABLE low.

Table 4.2: Measured metrics obtained from lab-environment trials

Experiment Code	CHA Power [dBm]	CHB Power [dBm]	CHA Energy [J]	CHB Energy [J]	CHA Energy [J]	CHB Energy [J]	Naïve AOA [°]	STDEV AOA [°]	Target Range [m]	ΔR [m]
PFT000	13.0077	13.0077	4.00E-07	4.00E-07	4.00E-07	4.00E-07	0	0	0	0
PFT001	12.3527	13.008	3.44E-07	4.00E-07	4.00E-07	4.00E-07	-	-	-	-
PFT002	13.008	12.511	4.00E-07	4.00E-07	3.57E-07	3.57E-07	-	-	-	-
PFT003	9.3033	9.4182	1.70E-07	1.70E-07	1.75E-07	1.75E-07	0.0579	0.0585	-1.991	0.0101
PFT004	9.3	9.4148	1.70E-07	1.70E-07	1.75E-07	1.75E-07	0.0603	0.0608	-1.991	0.0105
PFT005	2.5028	2.6185	3.66E-08	3.66E-08	3.66E-08	3.66E-08	0.0589	0.0593	205.034	0.0103
PFT006	1.573	2.5886	2.87E-08	2.87E-08	3.63E-08	3.63E-08	-33.5659	38.4785	212.997	-5.529
PFT007	2.6172	1.8765	3.65E-08	3.65E-08	3.08E-08	3.08E-08	33.7194	-38.321	205.034	5.5513
PFT008	1.2954	2.7662	2.70E-08	2.70E-08	3.78E-08	3.78E-08	-13.4792	70.307	216.978	-2.3309
PFT009	2.5774	1.3399	3.62E-08	3.62E-08	2.72E-08	2.72E-08	13.5974	-69.9751	205.034	2.351
PFT010	13.1109	13.1109	4.09E-07	4.09E-07	4.09E-07	4.09E-07	0	0	0	0
PFT011	12.434	13.1109	3.50E-07	3.50E-07	4.09E-07	4.09E-07	-	-	-	-
PFT012	13.1109	12.5921	4.09E-07	4.09E-07	3.63E-07	3.63E-07	-	-	-	-
PFT013	9.4907	9.6031	1.78E-07	1.78E-07	1.83E-07	1.83E-07	0.0585	0.0585	0	0.0102
PFT014	9.4869	9.6025	1.78E-07	1.78E-07	1.83E-07	1.83E-07	0.0606	0.0608	0	0.0106
PFT015	2.858	2.9736	3.86E-08	3.86E-08	3.97E-08	3.97E-08	0.059	0.0588	205.034	0.0103
PFT016	1.9561	2.9119	3.14E-08	3.14E-08	3.91E-08	3.91E-08	-33.6159	38.4676	212.997	-5.5362
PFT017	2.8522	2.1472	3.86E-08	3.86E-08	3.28E-08	3.28E-08	33.7675	-38.3193	205.034	5.5582
PFT018	1.5432	3.0008	2.85E-08	2.85E-08	3.99E-08	3.99E-08	-13.5414	70.1314	216.978	-2.3415
PFT019	2.8813	1.6578	3.88E-08	3.88E-08	2.93E-08	2.93E-08	13.6919	-69.7527	205.034	2.367




Figure 4.1: Received baseband signal from trial PFT001




Figure 4.2: Match filtered signal from trial PFT001




Figure 4.3: Processed and integrated signal from trial PFT001

4.1.2 Characteristics of the 208.42 m coaxial cable

Measurements PFT020, PFT021 and PFT022 are dedicated to characterising the 208.42 m cable which is used in the trials of Section 4.2. Figure 4.4 illustrates the anticipated insertion loss along the cable by using a network analyser. It is clear that one expects to observe a loss of roughly 6.5 dB when transmitting through the line.




Figure 4.4: Power loss measured along 208.42 m coaxial cable for trial PFT020

According to the manufacturers of the transmission line, the EM propagation velocity through the cable is 0.83 times the speed of light [51]. To evaluate its length, the Red Pitaya DAC is connected to the splitter input via a 15.4 cm coaxial cable. A second 15.4 cm cable connects one end of the splitter output to channel 1 on the Agilent Infiniium MS09104A oscilloscope. The 208.42 m cable connects channel 2 on the oscilloscope to the other unused output of the splitter.

The received signal in each channel on the oscilloscope is plotted in Figure 4.5. The top panel shows the received signal travelling along the short 15.4 cm cable, while the lower shows the delayed RF pulse travelling along the longer cable. The difference in time is empirically measured to be $3.356 - 2.519 = 0.837 \mu\text{s}$. Figure 4.6 depicts an equivalent test, except that a Hamming window has been applied to the transmitted chirp.

The cable is therefore assumed to be approximately:

$$R = c \times 0.837 \times 10^{-6} \times 0.83 + 0.154 = 208.4228 \text{ m} \quad (4.3)$$

by taking the difference in times of arrival, propagation speed and the shorter cable length into account.

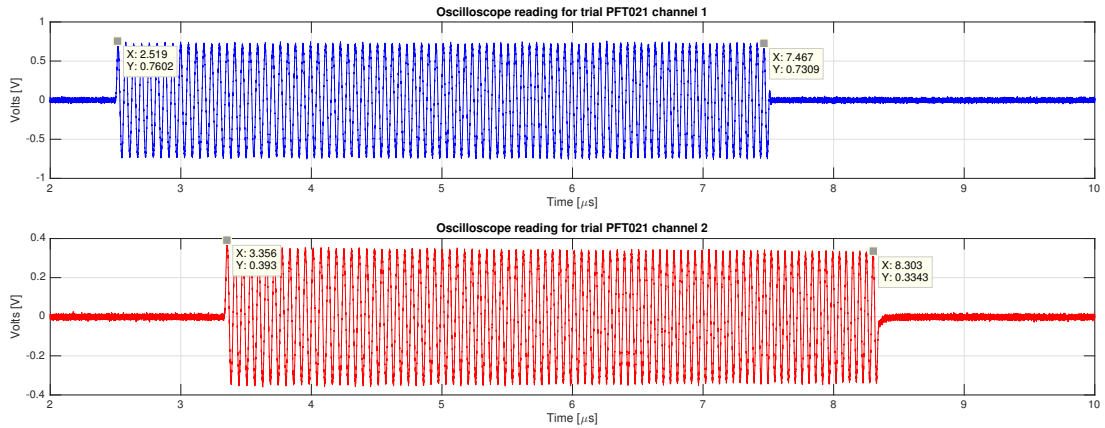


Figure 4.5: LFM chirp pulse delay test on a 208.42 m coaxial cable for trial PFT021




Figure 4.6: Hamming LFM chirp delay test on a 208.42 m coaxial cable for trial PFT022

4.1.3 Phase delays on the 18 m cables

Two 18 m long coaxial cables are used to connect the conical antennas to the input channels on the Red Pitaya in the field experiment conducted in Section 4.2. It is important that no external phase difference is induced between each antenna and the RF ADCs.

Trial PFT023 is conducted by connecting the 15.4 cm coaxial cable from the DAC output to the splitter input. The two 18 m cables are then attached to channels 1 and 2 of the oscilloscope, with the other ends connected to the splitter outputs.

A signal is emitted from the transceiver, bisected at the splitter and received by the oscilloscope after travelling down the two 18 m coaxial cables. The result is illustrated in Figure 4.7. It can be seen that both signals remain in-phase.




Figure 4.7: The split signal travelling along the two 18 m cables for trial PFT023

4.1.4 Synchronisation Test using Two Red Pitaya Devices

Two Red Pitayas are configured in a master and slave mode as discussed in Section 3.4.9. Trial PFT024 is conducted by connecting each DAC output of each device to channels 1 and 2 on the oscilloscope respectively. A tapered 20 μs pulse is ejected from each device using the synchronisation logic as shown in Figure 3.34. The received signal at the oscilloscope is illustrated in Figure 4.8.




Figure 4.8: Received pulse from two independent Red Pitaya devices for trial PFT024

4.2 Field Test Trials

Attention is now drawn to the outdoor experiments conducted to determine an AOA using the three algorithms developed in Section 3.3.

4.2.1 Background & Utilised Interferometric Geometry

A geometrical layout is shown in Figure 4.9 to measure various azimuthal angles. Both the primary receiving antenna, RX_A , and transmitting antenna, TX_A , remain stationary and are aligned through point J across 208.42 m. This distance matches the cable length used in Section 4.1.2.

The secondary receiving antenna, RX_B , is rotated from A to J and AA to JJ throughout the field test to adjust the AOA. Each adjacent location is separated by 10° . While Figure 2.10 suggests relocation of the target transmitter, this layout adjusts the secondary antenna along a $r = 30$ m circumference. This is practically much easier than moving the transmitting antenna over a 208.42 m arc and keeping RX_B stationary. The same phase difference phenomenon is achieved.




Figure 4.9: Geometrical layout of the interferometric field test

A total of 20 antenna positions exist in the test configuration. A protractor was used to construct a similar layout which was used to position all the elements. Appendix A.8 illustrates this and various other relevant images of the experiment construction.

The outdoor trials were conducted on 28 April 2017 on the field of Edgemoad High School ($33^\circ 52' 17.5''\text{S}$, $18^\circ 33' 04.1''\text{E}$). The weather had clear skies, with a maximum wind speed of 13 km/h and peak temperature of 20°C . Due to time constraints and logistical permissions, measurements were only taken at A-J and AA. Nevertheless, the unused 4th quadrant positions were compensated for by exchanging the inputs to channels A and

B on the Red Pitaya when operating in the first quadrant. This effectively models an equivalent mirrored phase delay, as if the antennas were arranged from BB to JJ.

4.2.2 Methodology

The steps required to repeat this experiment is as follows:

1. Construct antenna RX_A atop a grid layout as shown in Figure 4.9. This grid is used to demarcate the 10° incremental positions (See Figure A.8.1).
2. Build antenna RX_B at position A, 30 m away from RX_A .
3. Erect the transmit antenna, TX_A , 208.42 m due North from RX_A .
4. Connect RX_A and RX_B to channel A and B of the Red Pitaya via the two 50Ω T-terminated SMA adapters. Precede each terminator with a 20 dB LNA and 12.5-25 MHz filter. The amplifiers improve the usable dynamic range on the ADCs.
5. Connect TX_A to the DAC output using the 208.42 m coaxial cable.
6. Perform a simple transmit and receive test to ensure the generated beacon signal can be detected. One should measure an AOA of 0° .
7. Conduct the arrangement of trials as outlined in Section 4.2.3.
8. Move RX_B to position B, repeat step 7, followed by relocating and testing at positions C, D, E and so on until all the positions have been visited.

4.2.3 Default Operating & Test Parameters

To gather a variable set of data to be utilised for a discussion on the efficiency of the three AOA algorithms, various transmission parameters are used. At every test location, the following actions are performed:

1. Clear the FPGA memory of any residual data.
2. Apply the default parameters as identified in Table 4.3.
3. Conduct a test measurement, ensuring the transmitted signal can be detected and that the measured AOA appears logical.
4. Record 2 MB of the beacon signal (per channel) and store it to disk.
5. Incrementally adjust a parameter of interest. Repeat step 4 until all the variable combinations have been exhausted.
6. Exchange the inputs to channels A and B on the Red Pitaya and repeat steps 1 to 5.

It is important to emphasise that the same set of parameters are tested at each location. Different combinations of carriers, window types and modes of demodulation are assessed.

For every position, 88 unique measurements are taken. As 10 locations were visited, a total of 880 readings were produced. This amounts to 3.44 GB of sampled data.

Each recording was catalogued and identified by a unique identifier. Two example IDs are FM00E and FM00EE. The prefix FM stands for ‘Field-Measurement’, while the suffix E indicates that the recording was conducted at position E. The double suffix EE indicates that the readings were taken with channels A and B on the Red Pitaya reversed.

In addition, each recording is associated with a metadata file containing the transmission and reception details. The metadata for trial FM00E is shown below and contains vital information such as the utilised PRI, sample rates and register values. All the data is well organised and made publicly available. Access information is provided in Appendix A.1. This was managed with the intention of making the data available to other related fields of study.

Listing 4.1: Example metadata entry for recording FM00E

Initiate Time (UTC)	: 2017-04-27 09:33:54
Comment	: DRY-FM00E
Sample Rate (Hz)	: 25000000
Decimation	: 5
Operating Mode	: Real
Buffer Size Per Channel (MB)	: 2
PRI Duration (us)	: 1000.000
TX Delay N Samples	: 0
Start Offset	: 16
TX Window	: Boxcar
Left N TX Switch Gap	: 0
Right N TX Switch Gap	: 0
DAC Bit Length	: 14
Number DAC Samples Generated	: 2500
DAC Rate	: 125000000
Start Frequency (Hz)	: 15500000.000
End Frequency (Hz)	: 18500000.000
Bandwidth (Hz)	: 3000000.000
Chirp Duration (us)	: 20.000
Chirp Time Bandwidth	: 60.000
Start PRI N Samples Dropped	: 0
Gen Channel A Frequency (Hz)	: 999450.688
Gen Channel B Frequency (Hz)	: 999450.688
Gen Phase A (degree)	: 0
Gen Phase B (degree)	: 90
Quad Frequency (Hz)	: 17000000.000
Quad Phase Offset (degree)	: 0.000
CFG REG	: FF05000009C4000D
GEN REG	: 4000020C0000020C
PRI REG	: 1E847
QUA REG	: 22D0E560
GAL REG	: 0
GAR REG	: 0

Table 4.3 lists a combination of the default and varied parameters used to conduct the measurements at each visited location. The default values were used to check the transmitting and receiving capabilities of the arrangement. This was taken as a precautionary step to ensure that the relocation of antenna RX_B had not disrupted any reception capabilities. Once an interpretable and logical AOA had been verified, the varied parameters were used.

Table 4.3: List of default and varied parameters used at each visited test location

Parameter	Default Value / Type	Varied Value / Type	Comments
Carrier Frequency	[17, 18, 19] MHz	[17, 17.5, 18.00, 18.01, 18.02, 18.03, 18.04, 18.05, 18.5, 18.6, 19] MHz	Several carrier frequencies are used to unwrap the ambiguous AOA. Both expansive and closely spaced frequencies are utilised. The use for this is discussed in Section 3.3. For the naïve algorithm, only 17 MHz is used to determine the AOA with regards to the default values.
Windowing	Boxcar	[Hamming, Boxcar]	For each carrier frequency, both a Hamming and Boxcar tapering function on the transmitted LFM signal is used.
Demodulation Mode	Hilbert-Based Demodulator	[I/Q Demodulator, Hilbert-Based Demodulator]	Both the Hilbert-based and I/Q demodulators are used for each carrier frequency and each window function. The design of each is addressed in Sections 3.4.5 and 3.4.6.
LFM Bandwidth	3 MHz	3 MHz	The bandwidth is fixed at 3 MHz for each recording. This results in a one-way range resolution of 89 m.
PRI	1000 μ s	1000 μ s	The PRI is set to produce a large two-way unambiguous range of 150 km. This parameter is not critical to the performance.
PRF	1000 Hz	1000 Hz	The PRF determines the maximum measurable Doppler. In this case, one can measure targets travelling radially unambiguously up to ± 4411.76 m/s using the Hilbert-based demodulator with a carrier of 17 MHz.
Pulse Duration	20 μ s	20 μ s	This duration results in a time-bandwidth product of 60.
Number Pulses	40	40	For each recording, 2 MB or 41.94 ms of data is captured per channel. This translates to 41 received PRIs. Since the first PRI is dedicated to the Write-Back module, only 40 intervals can be used for coherent integration. See Section 3.2.7.
Transmit Power	10 dBm	10 dBm	No amplifiers are used to bolster the signal at the output of the DAC.

4.2.4 Estimated Power Budget

Table 4.4 lists a power budget derived for the experiment. It is based on Equation 3.12 which has been adapted for one-way communication:

$$P_{ar} = \frac{P_{at}G_tG_r\lambda^2}{(4\pi)^2R^2L_sL_c} \quad (4.4)$$

The average transmit power, $P_{at} = \frac{T}{PRI} \times P_t$, is used instead of the peak transmit power. This is intentionally adopted as the average received power, P_{ar} , in each channel is measured by the Red Pitaya and quoted in Table 4.5. The transmit duty cycle is 2%.

Table 4.4: Estimated power budget for field experiment

Description	Symbol	Unit	Positive Contribution	Negative Contribution
Average Transmit Power	P_{at}	dBW	-36.990	
Total Antenna Gain	$G_t = G_r$	dB	10.000	
208.42 & 18 m Cable Loss	L_c	dB		6.483 + 1.5
Range Loss	R^2	dBsm		46.379
$(4\pi)^2$	-	dB		21.984
Wavelength Squared	λ^2	dBsm	24.927	
Additional System Losses	L_s	dB		1
Total	-	-	-2.062	77.346

Total Signal Power Before 20 dB LNA		dBW		-79.408
Total Signal Power After 20 dB LNA		dBm		-29.408
Total Noise Power Before 20 dB LNA (ITU Model)		dBW		-80.450
SNR Before 20 dB LNA		-		1.042
SNR After Integration & Pulse Compression		-		34.844

The power levels before and after the utilised 20 dB low noise amplifier are also quoted. The background noise power is estimated from the International Telecommunication Union Radiocommunication Sector (ITU-R) P.372-7 [52]. Figure 4.10 illustrates the assumed noise PSD.




Figure 4.10: PSD of background noise power derived from the ITU-R P.372-7 [52]

It is important to emphasise that this budget is simply used as a model and does not represent an accurate measurement of the received power or SNR. Elements such as the antenna gain, background noise, range degradation factor and system losses could not be measured accurately before the trials. Values were logically estimated to act as a preliminary check one would be able to conduct the experiment over the scale of 208.42 m given the operating parameters.

Section 4.2.5 details a series of illustrations of the received signals recorded at each location before and after integration. The measured AOA, received power and perceived target location is listed in Tables 4.5 and 4.6.

4.2.5 Empirical Results

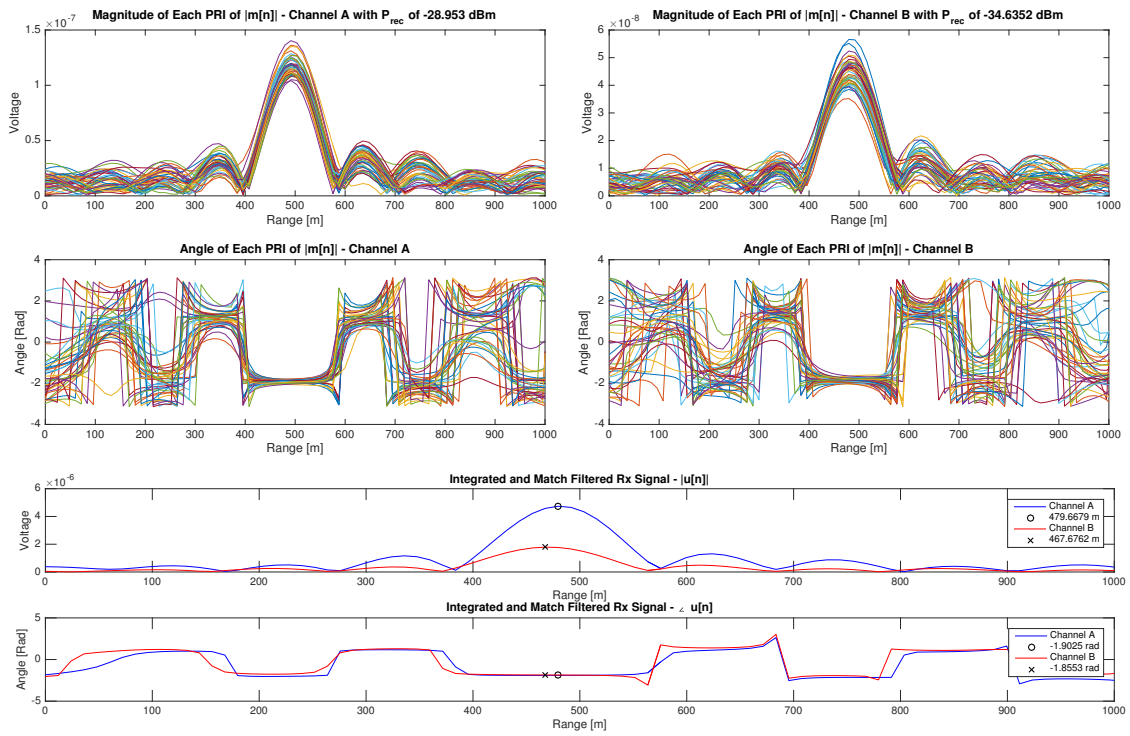


Figure 4.11: Trial FM00A with default parameters. Only $f_c = 17$ MHz is shown.

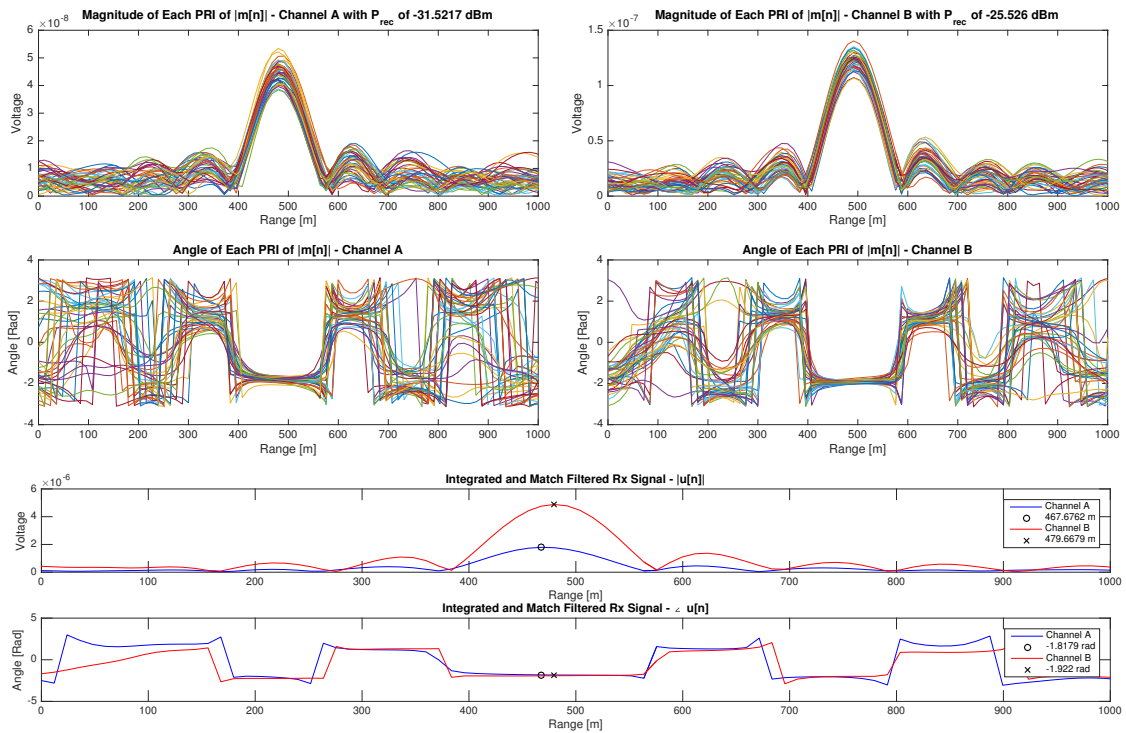


Figure 4.12: Trial FM00AA with default parameters. Only $f_c = 17$ MHz is shown.

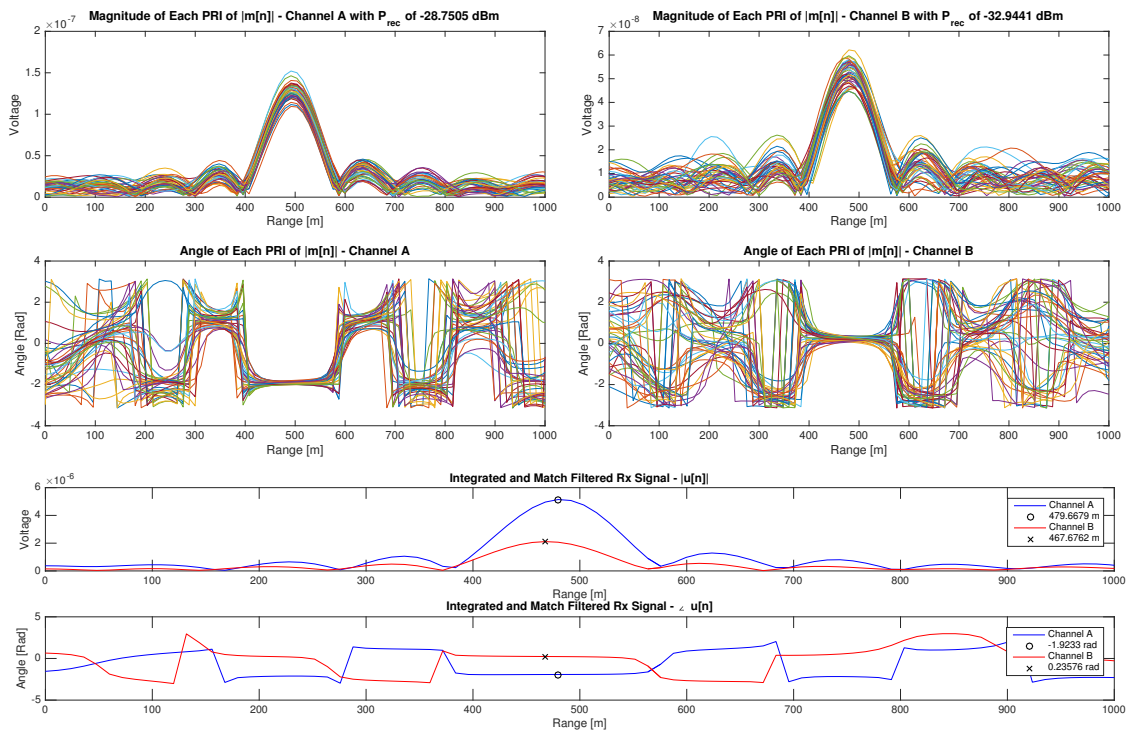


Figure 4.13: Trial FM00B with default parameters. Only $f_c = 17$ MHz is shown.

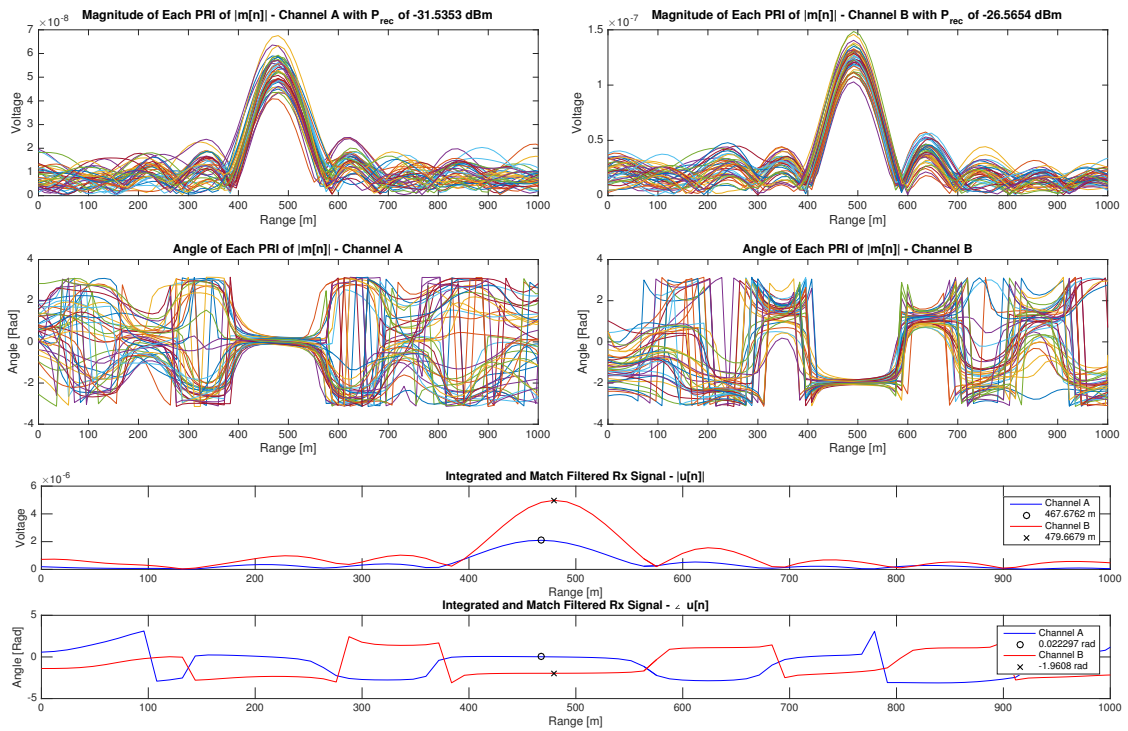


Figure 4.14: Trial FM00BB with default parameters. Only $f_c = 17$ MHz is shown.

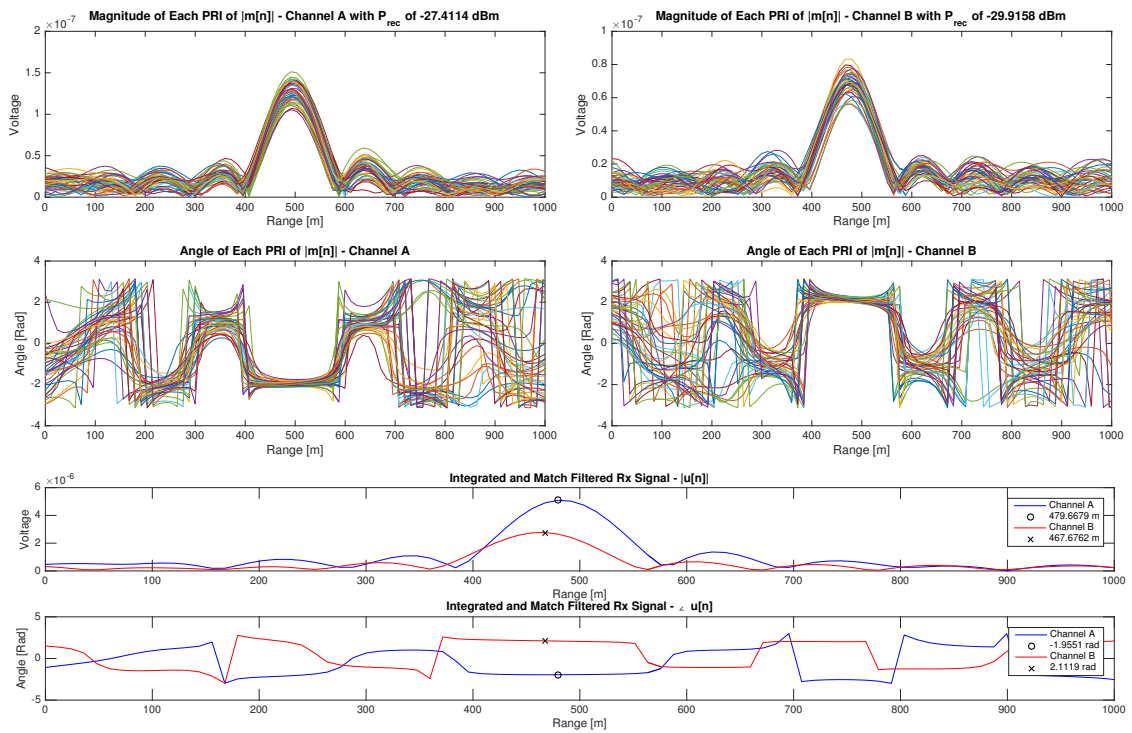


Figure 4.15: Trial FM00C with default parameters. Only $f_c = 17$ MHz is shown.

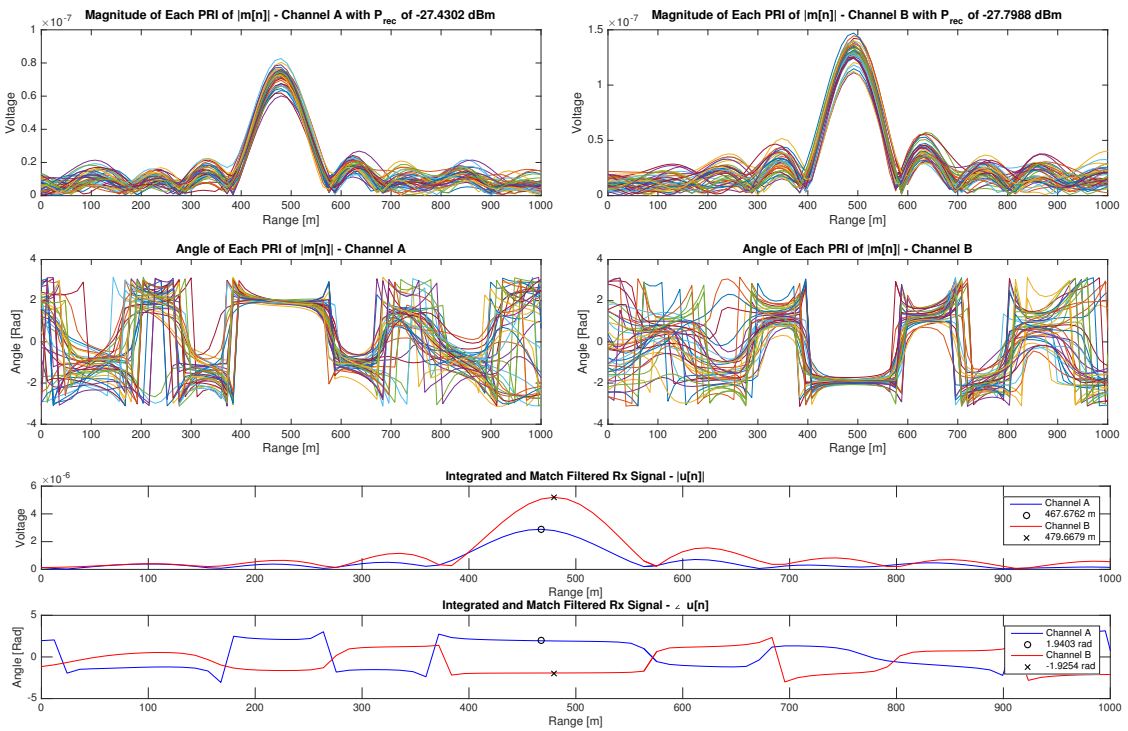


Figure 4.16: Trial FM00CC with default parameters. Only $f_c = 17$ MHz is shown.

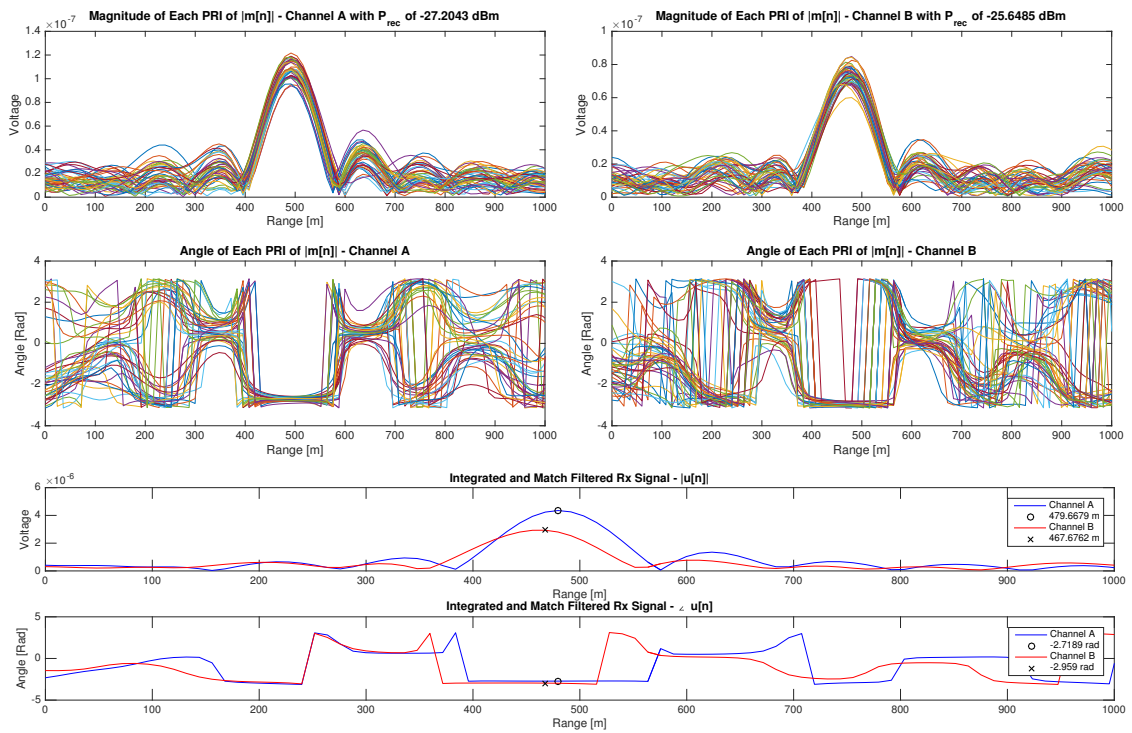


Figure 4.17: Trial FM00D with default parameters. Only $f_c = 17$ MHz is shown.

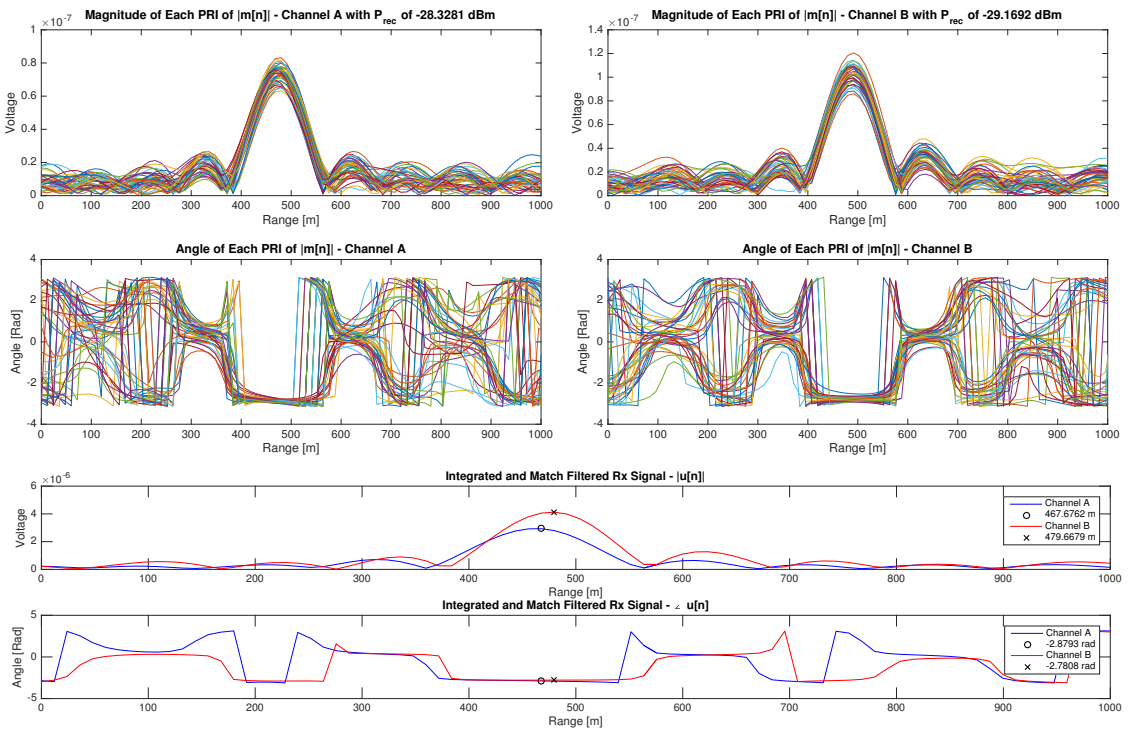


Figure 4.18: Trial FM00DD with default parameters. Only $f_c = 17$ MHz is shown.

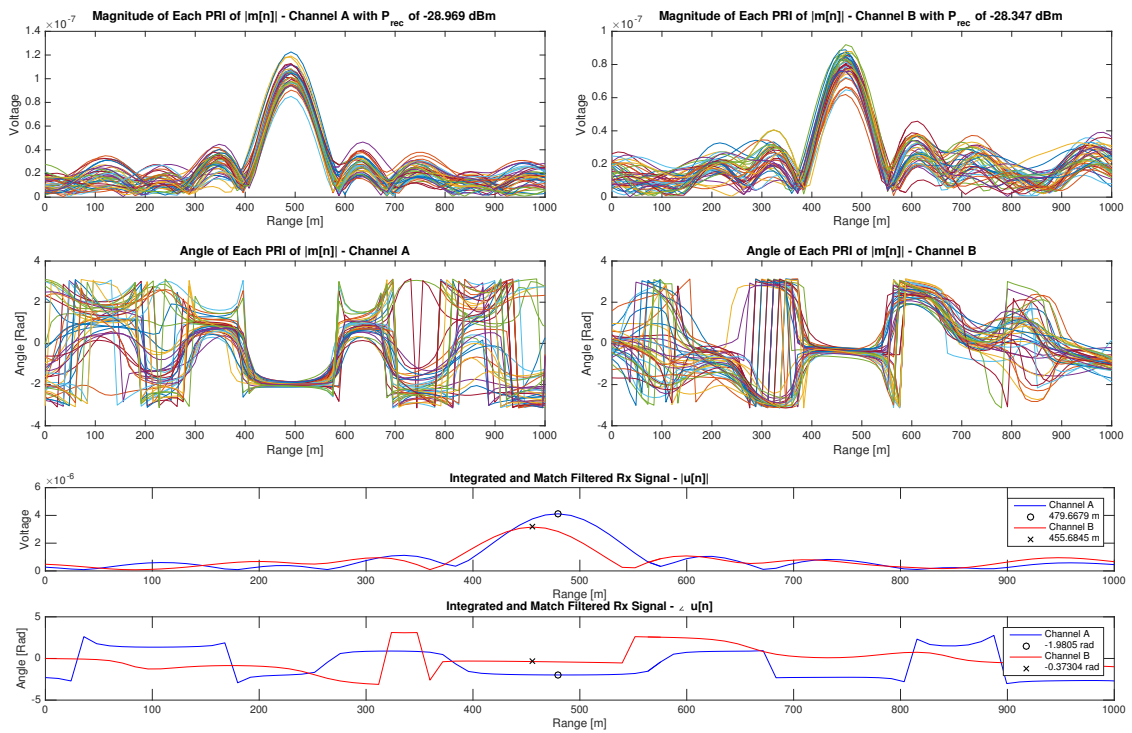


Figure 4.19: Trial FM00E with default parameters. Only $f_c = 17$ MHz is shown.

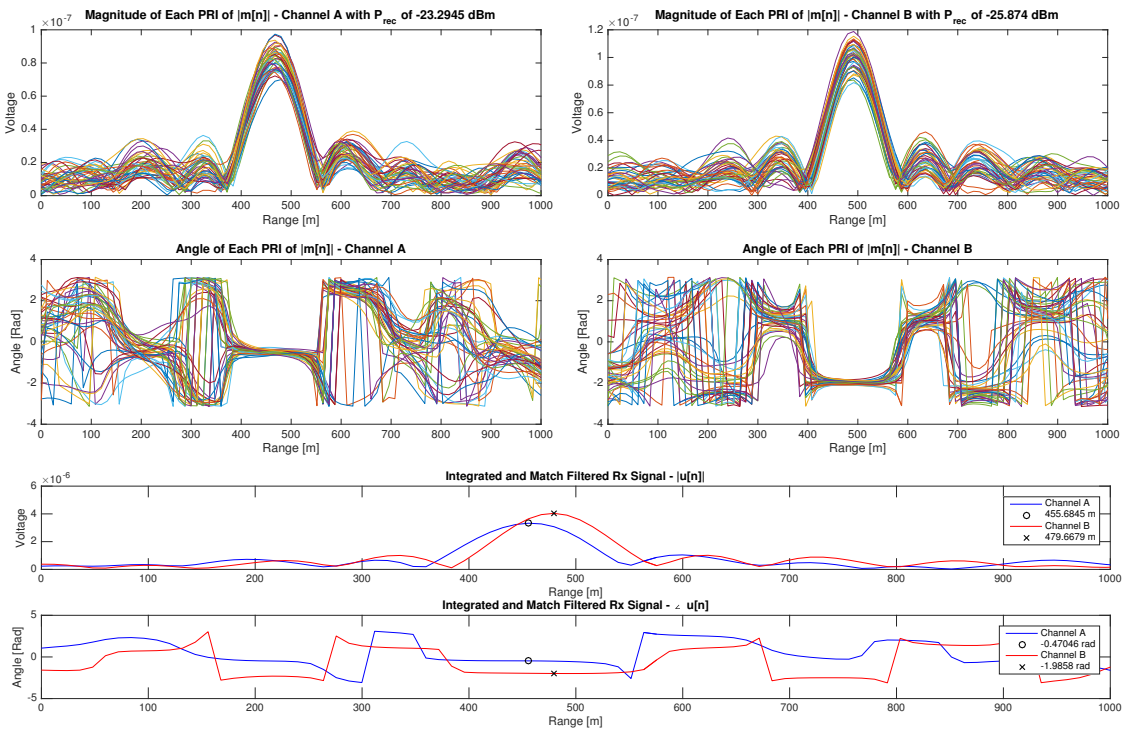


Figure 4.20: Trial FM00EE with default parameters. Only $f_c = 17$ MHz is shown.

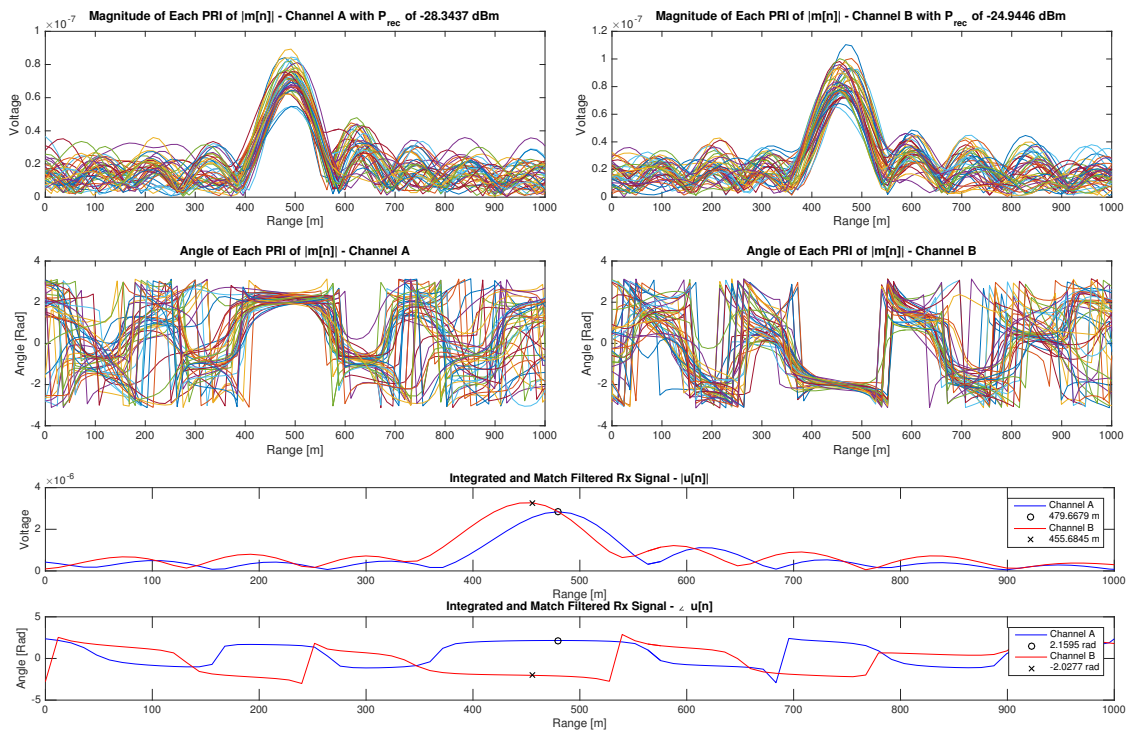


Figure 4.21: Trial FM00F with default parameters. Only $f_c = 17$ MHz is shown.

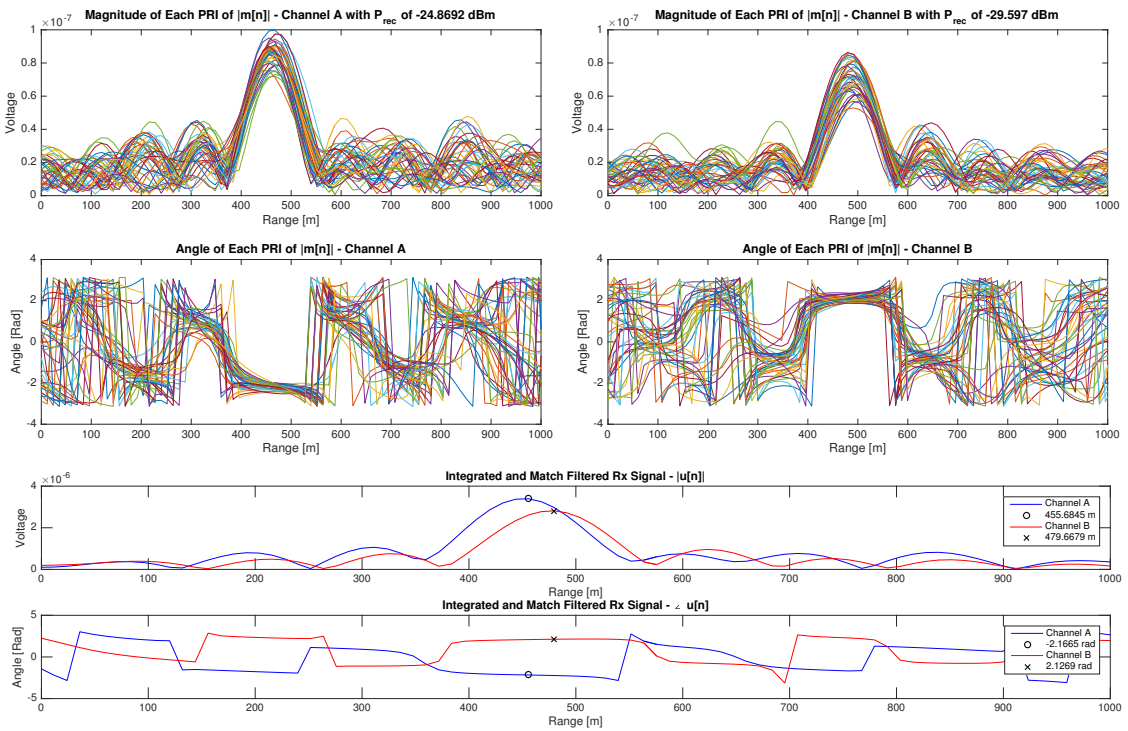


Figure 4.22: Trial FM00FF with default parameters. Only $f_c = 17$ MHz is shown.

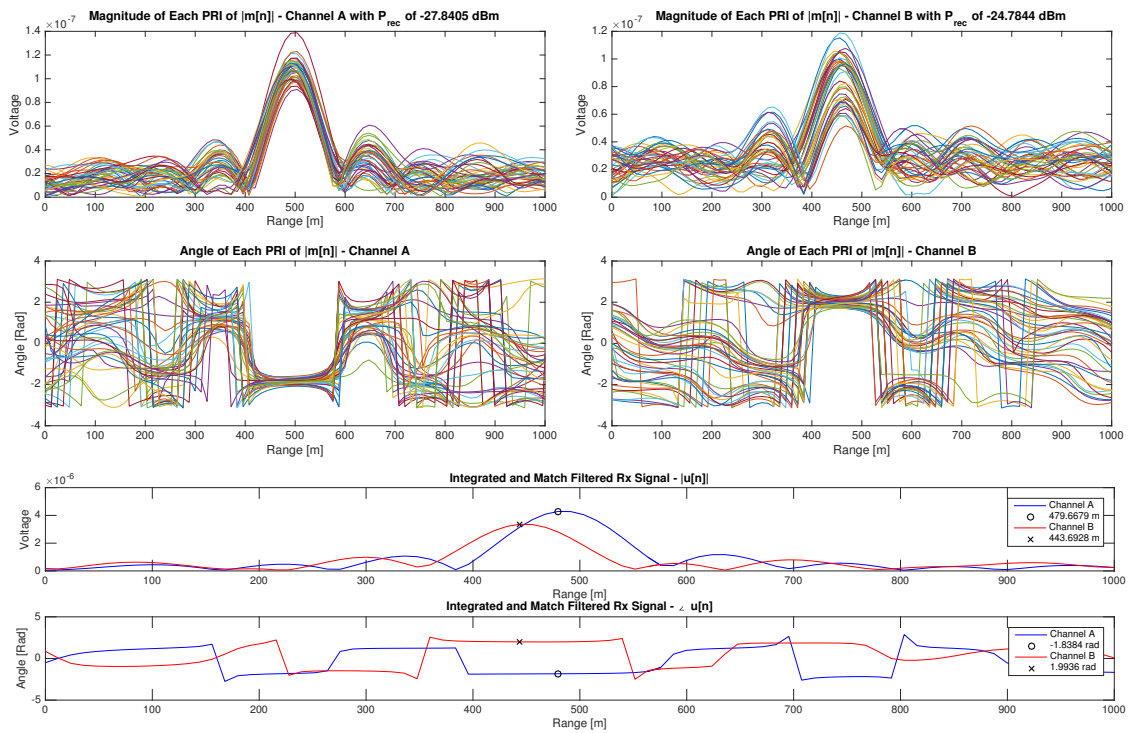


Figure 4.23: Trial FM00G with default parameters. Only $f_c = 17$ MHz is shown.

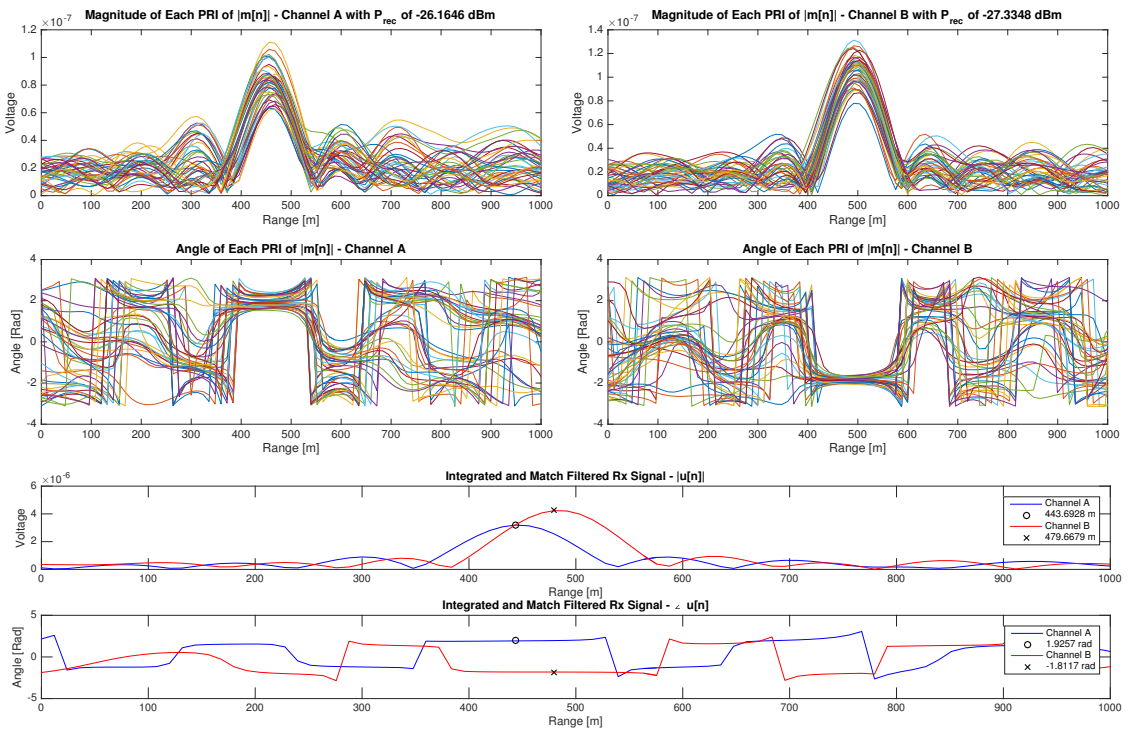


Figure 4.24: Trial FM00GG with default parameters. Only $f_c = 17$ MHz is shown.

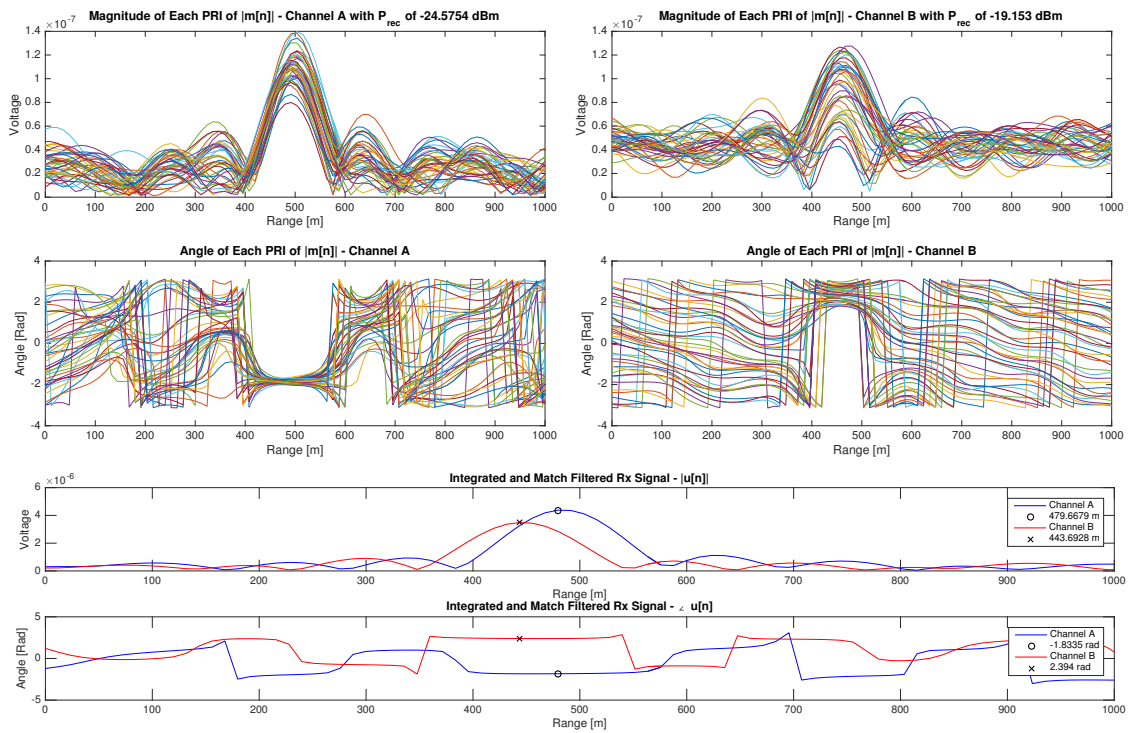


Figure 4.25: Trial FM00H with default parameters. Only $f_c = 17$ MHz is shown.

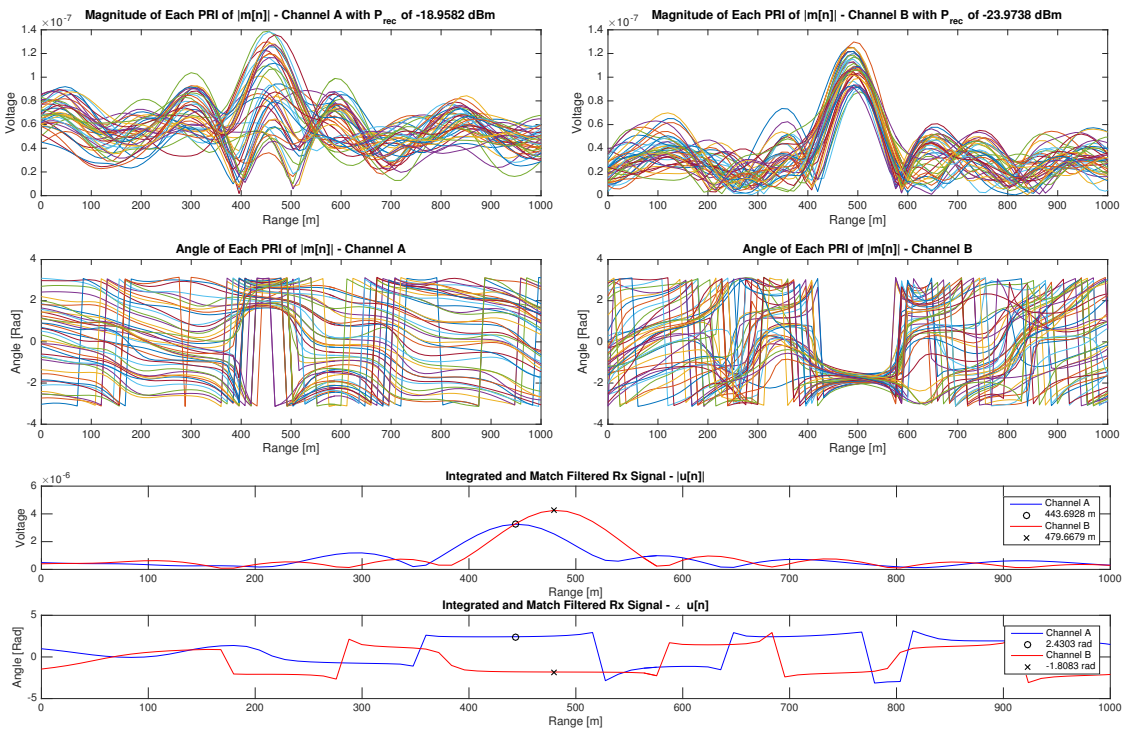


Figure 4.26: Trial FM00HH with default parameters. Only $f_c = 17$ MHz is shown.

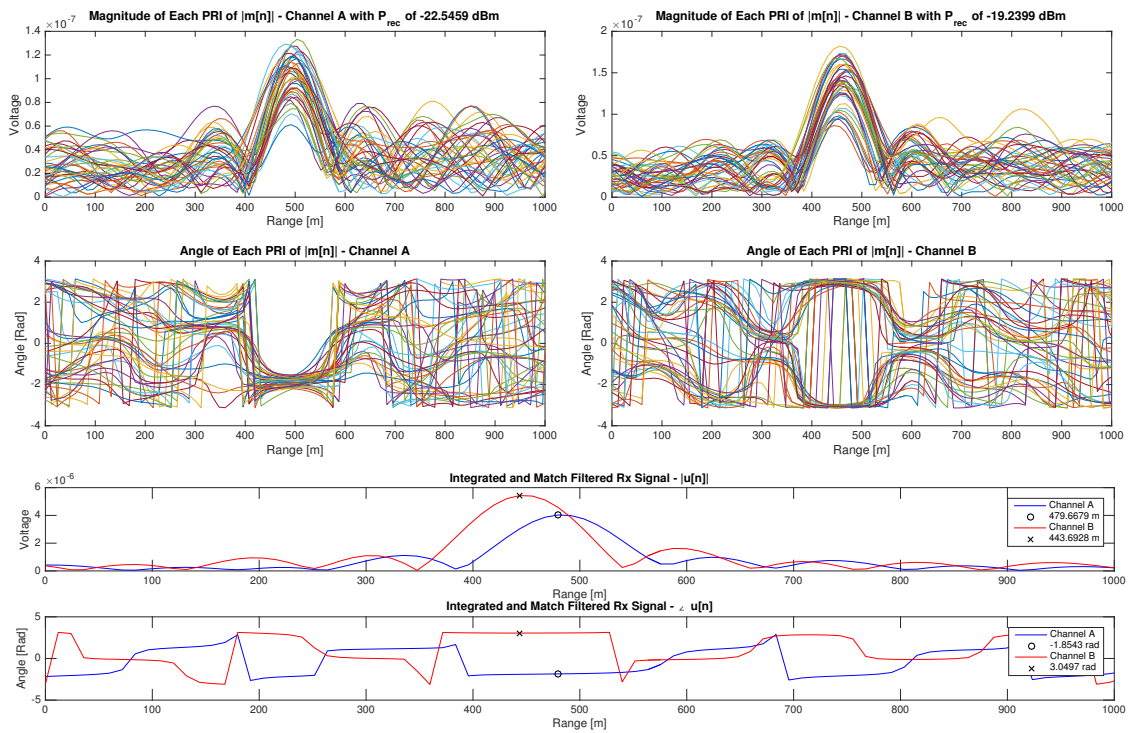


Figure 4.27: Trial FM00I with default parameters. Only $f_c = 17$ MHz is shown.

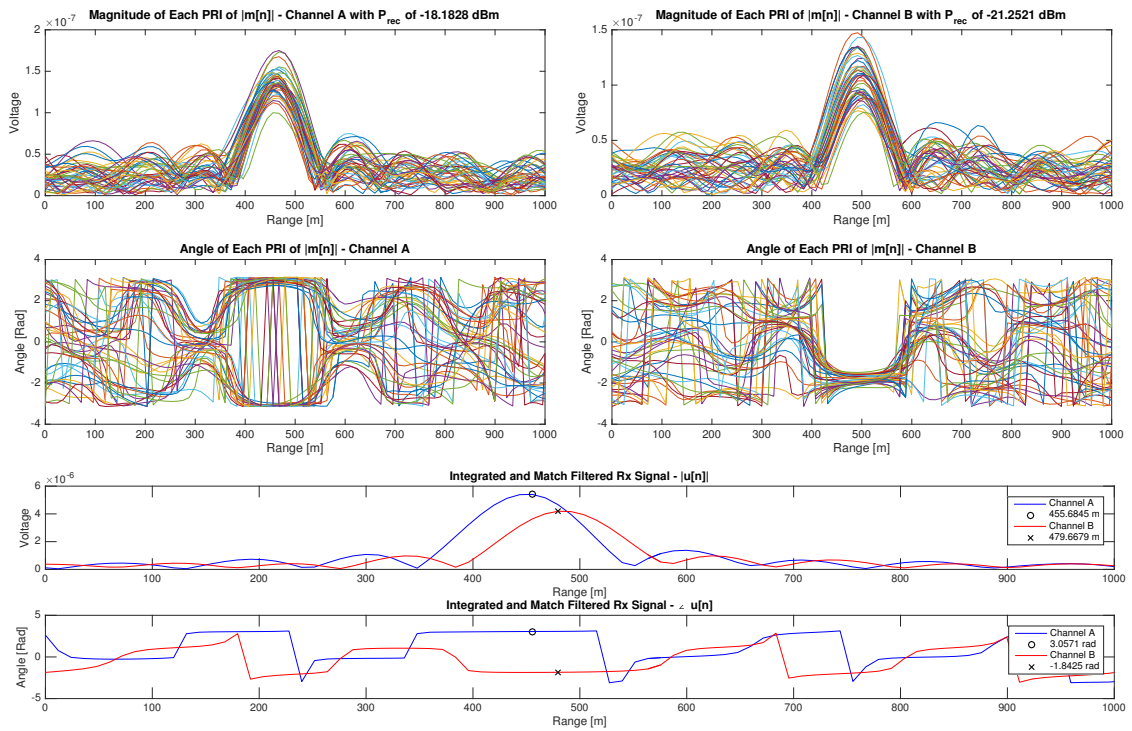


Figure 4.28: Trial FM00II with default parameters. Only $f_c = 17$ MHz is shown.

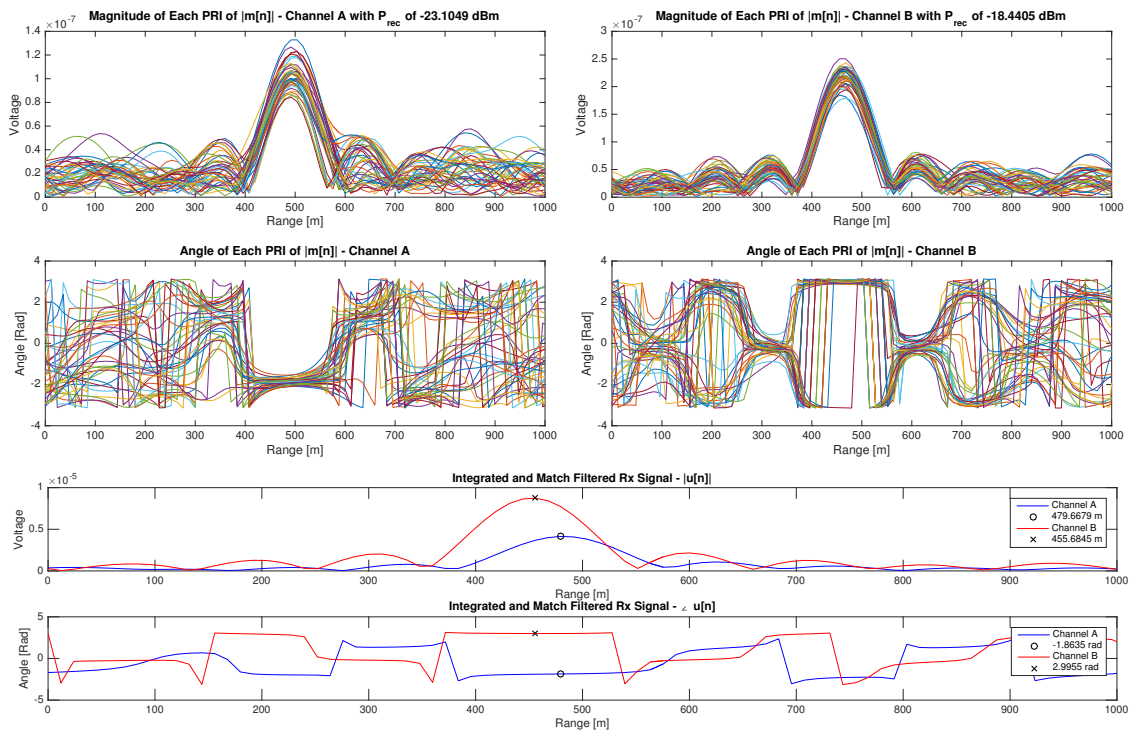


Figure 4.29: Trial FM00J with default parameters. Only $f_c = 17$ MHz is shown.

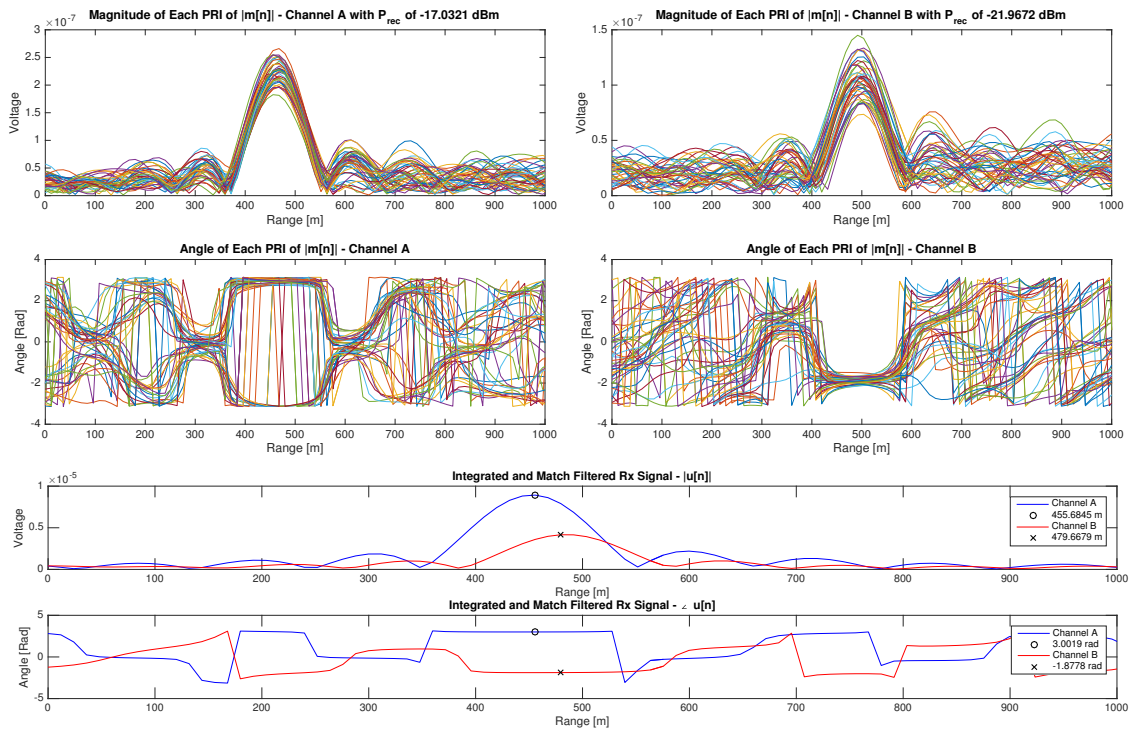


Figure 4.30: Trial FM00JJ with default parameters. Only $f_c = 17$ MHz is shown.

The predicted AOA for each algorithm is listed in columns 3, 4 and 5 of Table 4.5. The anticipated target range as seen by antennas RX_A and RX_B follows adjacently. The average power measured in each channel for $f_c = 17$ MHz is indicated by P_{RX_A} and P_{RX_B} . The standard deviation in received power for each channel is listed in the last two columns over the operated frequencies of 17, 18 and 19 MHz. All the results are based on the default values listed in Table 4.3.

Table 4.5: Summary of results depicted in Figures 4.11 to 4.30

Code	AOA [°]	Naïve [°]	STDEV [°]	CRT [°]	Range RX_A [m]	Range RX_B [m]	P_{RX_A} [dBm]	P_{RX_B} [dBm]	STDEV(P_{RX_A}) [dBm]	STDEV(P_{RX_B}) [dBm]
FM00A	0	0.2018	0.3507	NaN	479.6679	467.6762	-28.953	-34.6352	0.12995	1.6523
FM00AA	0	-0.5447	-0.4699	-60.5	467.6762	479.6679	-31.5217	-25.526	2.6591	1.1968
FM00B	10	10.2619	10.4629	-50	479.6679	467.6762	-28.7505	-32.9441	0.22248	0.10226
FM00BB	-10	-9.4622	-9.5898	-69	467.6762	479.6679	-31.5353	-32.9441	0.5174	1.3292
FM00C	20	-11.0437	20.2718	-72.5	479.6679	467.6762	-27.4114	-29.9158	0.090191	4.798
FM00CC	-20	11.913	-19.6634	NaN	467.6762	479.6679	-27.4302	-27.7988	1.9654	0.16834
FM00D	30	-1.1937	29.9705	-61	479.6679	467.6762	-27.2043	-25.6485	1.8877	3.7335
FM00DD	-30	0.5997	-30.77	-59.5	467.6762	479.6679	-28.3281	-29.1692	0.10889	0.048608
FM00E	40	4.9488	40.5473	82	479.6679	455.6845	-28.969	-28.347	1.3368	0.0045965
FM00EE	-40	-5.7619	-40.8323	-52	455.6845	479.6679	-23.2945	-25.874	10.4009	0.092265
FM00F	50	10.5415	50.8656	59.5	479.6679	455.6845	-28.3437	-24.9446	1.0578	0.70938
FM00FF	-50	-10.5937	-50.9725	-60	455.6845	479.6679	-24.8692	-29.597	5.5223	2.686
FM00G	60	14.9312	60.2108	NaN	479.6679	479.6679	-27.8405	-29.597	1.0085	0.22634
FM00GG	-60	-14.3975	-59.4251	NaN	479.6679	479.6679	-26.1646	-27.3348	2.9038	0.25292
FM00H	70	-16.0376	69.9923	70	479.6679	443.6928	-24.5754	-19.153	0.073135	12.1589
FM00HH	-70	15.4687	-70.2943	NaN	443.6928	479.6679	-18.9582	-23.9738	2.0882	0.43885
FM00I	80	-11.8576	25.4203	73.5	479.6679	443.6928	-22.5459	-19.2399	0.73073	6.1148
FM00II	-80	11.5529	-25.5227	NaN	455.6845	479.6679	-18.1828	-21.2521	9.2759	2.4523
FM00J	90	-11.9243	25.5116	NaN	479.6679	455.6845	-23.1049	-18.4405	2.881	2.47
FM00JJ	-90	11.5642	-25.5451	-57	455.6845	479.6679	-17.0321	-21.9672	19.6008	0.33211

Table 4.6 lists the interpreted AOA using the standard deviation algorithm developed in Section 3.3.2. Both methods of demodulation are used, with varied carrier frequencies and windowing functions.

Table 4.6: Interpreted AOA using the STDEV algorithm. All values are in degrees [°].

Carrier Frequencies [MHz]		17.5, 19		18.01, 18.02, 18.03		18, 18.5, 19		18, 18.5, 19	
Window Function		Boxcar		Boxcar		Boxcar		Hamming	
Code	AOA	I/Q	Hilbert	I/Q	Hilbert	I/Q	Hilbert	I/Q	Hilbert
FM00A	0	29.987	31.624	32.455	0.428	1.104	31.377	2.021	31.591
FM00AA	0	-1.086	-31.541	-0.818	-0.461	-1.084	-31.540	32.471	-0.945
FM00B	10	41.059	44.411	8.567	10.525	11.497	44.053	11.492	43.854
FM00BB	-10	-44.211	-43.168	22.414	-9.631	-43.266	-42.940	-10.980	-42.703
FM00C	20	20.470	19.580	19.005	-11.514	19.341	20.083	19.994	60.579
FM00CC	-20	-20.360	-19.771	10.816	11.530	-57.779	-20.008	7.408	-19.936
FM00D	30	30.652	30.922	-0.413	-0.833	30.919	30.954	31.832	30.264
FM00DD	-30	-31.343	-30.967	0.867	0.549	-30.890	-31.154	1.682	-31.042
FM00E	40	42.804	40.653	6.848	-7.985	40.840	41.249	8.824	40.847
FM00EE	-40	-43.036	-40.576	23.001	-7.894	-41.159	-41.028	-4.592	-41.270
FM00F	50	-5.768	50.145	11.722	13.945	15.270	50.984	18.196	50.132
FM00FF	-50	-49.296	-50.391	-13.708	-13.918	-14.893	-50.864	-13.543	-51.392
FM00G	60	61.536	19.688	-14.000	-13.100	19.543	20.677	20.014	20.262
FM00GG	-60	-60.669	-60.402	12.517	-18.908	-19.247	-61.449	7.049	-60.545
FM00H	70	22.721	22.849	-10.999	-10.769	23.135	23.887	23.792	23.774
FM00HH	-70	-69.662	-69.388	8.941	10.673	-23.259	-23.533	-24.451	-23.753
FM00I	80	28.728	26.475	-6.146	-5.924	26.718	27.564	26.864	27.145
FM00II	-80	-26.846	-26.791	5.490	6.306	-26.992	-27.279	-27.121	-27.142
FM00J	90	25.704	66.584	-7.016	-6.663	24.248	24.641	65.610	24.689
FM00JJ	-90	-67.779	-67.126	6.786	6.569	-24.372	-67.592	-66.606	-67.085
Average MSE at all angles		924.29	744.93	3491.11	3104.72	1346.95	1018.37	1222.61	1054.99

4.3 Additional Experiment: Moving Aircraft Targets

As a fully coherent pulse-Doppler radar was developed for the interferometric trials, an auxiliary experiment was conducted in which moving aircraft were to be detected. No adjustments to the existing FPGA design were made. The arrangement used to detect the aircraft is shown in Figure 4.31.

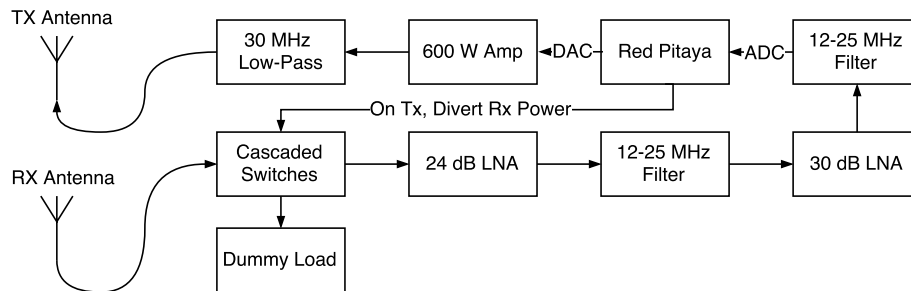


Figure 4.31: Overview of the construction used to detect moving aircraft

The output of the DAC was connected to a peak 600 W amplifier, which was operated at 400 W. Two of the conical antennas were used, one for transmit and one for receive. Four cascaded switches, amounting to 166.6 dBm of attenuation, were used to isolate the direct-path link between the antennas when transmitting. This was implemented to protect the Red Pitaya ADCs from receiving too much feed-through power (Section 3.4.10).

The output of the amplifier fed into a 30 MHz low pass filter to suppress unwanted harmonics. The filter then connected to the transmit antenna.

The receive antenna connected to the switches and a 12.5-25 MHz bandpass filter. The filter was sandwiched between a 24 and 30 dB low noise amplifier. A second identical bandpass filter followed, its output connected to the transceiver on channel A.

Exemplary detail on this arrangement can be found in research conducted by Burger [49]. Focus was not drawn to this task, as it did not relate to the scope of interferometry. What was relevant, however, is the application in which the developed FPGA design was used.

Figure 4.32 illustrates the range-Doppler map generated from this experiment. A centre frequency of 16.15 MHz, and bandwidth of 0.8 MHz was chosen. The pulse length was fixed at 20 μ s, with a PRI of 500 μ s. No windowing was applied, and the Hilbert-based demodulator was used. A map detailing the surrounding landscape is shown in Figure 4.33. The points labelled A-E on each of these figures correspond.




Figure 4.32: Range-Doppler heat-map generated from moving aircraft experiment

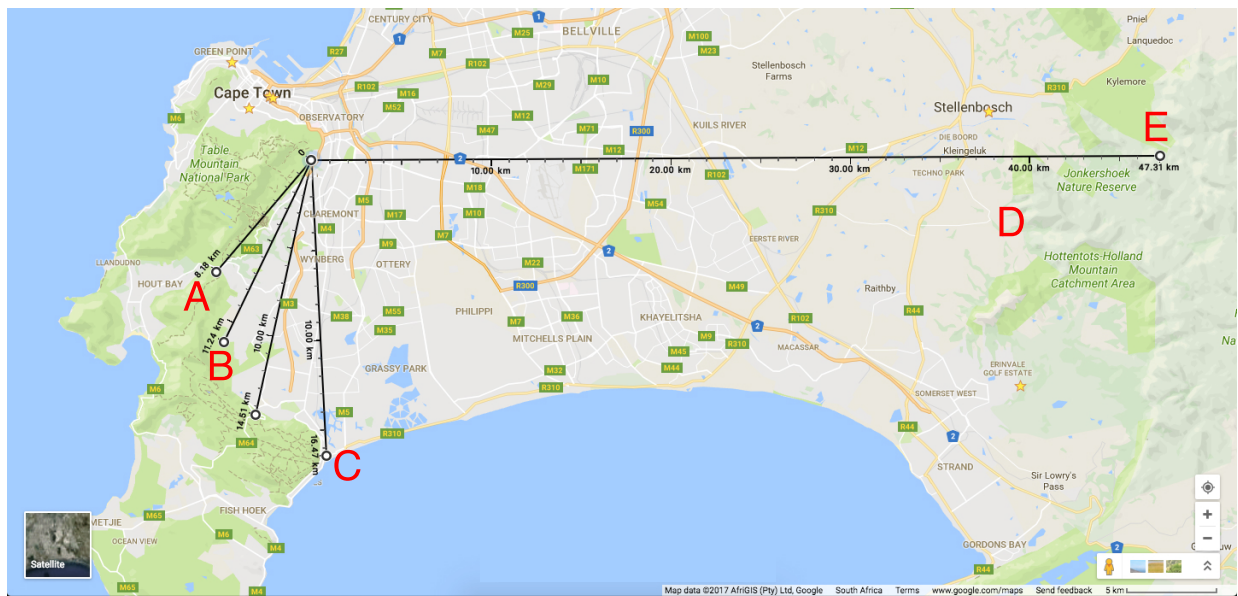


Figure 4.33: Geographical map of the illuminated area for the moving target experiment

Chapter 5

Discussion

While constructing the FPGA design on the Red Pitaya, concern was centred on the accuracy of the device. How could one confidently evaluate measured results in a field test where the transceiver has just been realised? Were all the binary samples correctly stored? Were the measured power, phase and amplitude correct?

The series of lab trials, designated from PFT000 to PFT024, were created to instil confidence in measured data before entering into the outdoor environment to perform the interferometric tests. Section 5.2 discusses the findings from these AOA experiments, while Section 5.3 provides commentary on the secondary moving target experiment.

5.1 Comments on the Pre-Field Tests

The first conducted test, PFT000, was arguably the simplest. It relied purely on mathematics and fed a digitised RF definition of the transmitted LFM chirp directly into the receiving chain on the Red Pitaya. No external influences, including noise, are introduced by doing so.

From Table 4.2, one observes that both receiving channels A and B identically measure 13.0077 dBm. This indicates that both the channels are internally symmetrical as both derive the exact same result.

In Section 3.4.7, it was shown that the Red Pitaya would produce a maximum output of 10 dBm which is the same amount of power being artificially fed into the receiving chains. This is approximately 3 dB less than that measured. This discrepancy is expected as the output power of the device is quoted using the RF definition of the LFM chirp, while the receiver uses analytical expressions. The analytical form measures twice as much power ($10 \log(2) = 3.01$ dB) and caters for this difference. An explanation is provided in Equation 3.7.

Trial PFT010 elicits an identical measurement, except it uses the I/Q demodulator instead of the Hilbert-based demodulator. Both channels measure 13.1109 dBm which further proves the existing symmetry in channel A and B. As the I/Q demodulator has an effective down-sampling rate of 20, its accuracy in power measurement degrades. Comparatively, the demodulator utilised for trial PFT000 only down-samples by 5.

Nevertheless, tests PFT000 and PFT010 are within 0.008 dB and 0.1 dB respectively of the correct value, which is considered adequate. In both these trials the measured energy in the received pulse is that anticipated $2 \times E_{TX} = 2 \times A^2T/100 = 0.4 \mu J$ (Equation 3.136). The additional factor of 2 caters for the analytical expression.

In trial PFT001, the full amplitude $2 V_{pp}$ LFM chirp is delivered from the DAC output into channel A over the 0.665 m coaxial cable. From Figure 4.1, one can observe the correct amplitude of roughly $\pm 1 V_p$ being measured. Each baseband chirp is correctly seen to span $20 \mu s$. The matched filtered signal shown in Figure 4.2 spikes on the start of every PRI, which is also $20 \mu s$ long. The matched filter responds with zero delay. Causality is not considered an issue since all the processing is performed post-capture. The output peaks with a value of $A^2T = 20 \mu J$ which is the total energy contained within the analytical form of the transmitted signal if one does not take the characteristic impedance of the line into account.

Figure 4.3 illustrates the output after $(41.94 \times 10^{-3}) / (20 \times 10^{-6}) - 1 = 2096$ PRI segments have been integrated. The peak amplitude resides at $0.04192 V$ as it is 2096 times larger than the non-integrated signal shown in Figure 4.2. The argument across the mainlobe of the magnitude output appears uniform, which is to be expected. This matches the result theoretically expressed and illustrated in Section 3.2.6 on pulse compression. The PSD of the integrated signal clearly resembles that shown in Figure 3.3, and spans 3 MHz. This is the same bandwidth used in the transmitted signal. Integrating across this bandwidth on channel B yields 13 dBm, which is near identical to the measured power stated for trial PFT001 in Table 4.2.

It is worth noting that channel A experiences a 0.66 dB loss from transmitting across the 0.665 m coaxial cable in comparison to the signal internally received through channel B. When the DAC output is directly fed into channel B over the same cable in trial PFT002, one measures a 0.497 dB loss indicating a small unwanted loss discrepancy of 0.16 dB between trials PFT001 and PFT002, which should be measuring the same thing. This difference is thought to be a result of using two different 50Ω T-junction terminators on each channel which are not identically matched. This tolerance is considered acceptable as it does not alter the received phase.

No angle or range measurements are taken for trials PFT001, PFT002, PFT011 and PFT012, as it makes little sense to compare a real transmitted and received signal with an artificially generated one.

Trial PFT003 and PFT013 transmit the LFM chirp out the DAC, through the splitter and into channels A and B across two identical 15.4 cm coaxial cables. In theory, both

channels are expected to receive the pulsed waveform at an identical time. Table 4.2 lists the phase difference between each channel for each trial to be 0.0101 and 0.0102 m respectively (Equation 4.2). This result is indicative that the correct phase difference is being measured as it is near zero as theory would suggest. In each case, this is interpreted as a naïve AOA of 0.0579° and 0.0585° respectively, as if the target resided at boresight. The error is within 60 millidegrees and is considered suitable.

With reference to the same trials, the use of the Hilbert-based demodulator measures the target to reside at -1.991 m, while the I/Q demodulator sees the target at 0 m. In reality, the signal has travelled a total of $2 \times 15.4 = 30.8$ cm from the DAC to either ADC. The one-way range bin size of each demodulator is $c \times \frac{5}{125 \times 10^6} = 12$ m and $c \times \frac{20}{125 \times 10^6} = 48$ m respectively. The incorrect perceived position of -1.991 m is simply due to sampling. As the test is being conducted over 30.8 cm, there is no realistic way to accurately place the target as it fits within a range bin size that is far too large. If one sampled more frequently, the measured target position would approach 30.8 cm. Since the interferometer relies on the phase which is constant for the width of a range bin, this is not considered an issue.

In trial PFT003, channel A and B measure 9.3033 and 9.4182 dBm respectively. Since the transmitted signal is passed through a splitter, one expects a loss of 3 dB or to measure half the transmitted power in each channel. From trial PFT001, a loss in the order of 0.66 dB was observed by simply transmitting out the DAC and over the 0.665 m coaxial cable. Therefore, one expects to measure around $13 - 3 - 0.66 = 9.34$ dBm using the transceiver in this arrangement, which is evidently the case.

Trial PFT004 is identical to PFT003, except that the two 15.4 cm cables feeding to the inputs of channels A and B are reversed. This has virtually no influence on the results, as one would expect. Channel B in trials PFT003, PFT004, PFT013 and PFT014 measures slightly more power. This suggests that the 50Ω T-junction terminator on channel B is slightly better matched. This has no influence on the measured phase difference and is not considered impactful. These 4 trials all measure the exact same AOA to within 1 millidegree, suggesting that both the I/Q and Hilbert demodulators are configured correctly to derive the same expected result.

Trial PFT005 differs from PFT003 by simply extending the common coaxial cable connecting the DAC output to the splitter input from 15.4 cm to 208.42 m. The measured phase difference, ΔR , changes by 0.2 mm between these two trials. This is expected since both the channels are still connected by two identical 15.4 cm cables. The phase difference should not change. The perceived target distance does, however, adjust to 205.034 m. Note that this value is short by 3.54 m of the correct length, even when accounting for the additional 15.4 cm coaxial. This discrepancy is due to sampling offset error and cannot unfortunately be avoided. A similar conclusion is derived for trial PFT015, which also sees the target at 205.034 using the I/Q demodulator.

Between Trial PFT013 and PFT015, the measured power in channel A dips from 9.4907

to 2.858 dBm with the addition of the 208.42 m cable. This is a difference of 6.63 dB. In trial PFT020, the loss along the long cable was measured to be 6.48 dB using a network analyser (Figure 4.4). These two estimations are considered to be in agreement. One can derive a similar conclusion by observing the power loss in channel B of trials PFT003 and PFT005.

Trials PFT021 and PFT022 further emphasise the loss of the 208.42 m cable to be in the order of 6.5 dB. One expects to measure a peak voltage of (adapted from Equation 3.137)

$$V_p = \sqrt{2 \times 10^{(-20-3-6.5)/10} \times Z} \Big|_{Z=50} = 0.335 \text{ V} \quad (5.1)$$

on the receiving side of the 208.42 m cable. This takes the 3 dB power division at the splitter into account. The amplitude of the received pulse is clearly of this magnitude in the lower panels of Figures 4.5 and 4.6. Three instruments involving an oscilloscope, network analyser and the Red Pitaya are all found to be in agreement when measuring the power loss at 17 MHz along the 208.42 m coaxial cable.

Trials PFT006 and PFT008 are dedicated to inducing a purposeful phase difference into the system by using different cable lengths. The coaxial cable connecting channel B is maintained at 15.4 cm, while channel A's is varied from 6.22 m to 9.34 m between each test. The phase difference in each trial is measured to be -5.529 and -2.3309 m respectively. Clearly phase wrapping has occurred as both phase differences are negative when they should be positive.

If one takes the difference of these two numbers, $-2.3309 + 5.529 = 3.1981$ m, they measure the physical differential distance of the cables used in the two trials. In this case, $9.34 - 6.22 = 3.12$ m. This error is perhaps fairly large, nearing 8 cm, however, it reduces substantially for smaller induced phase differences. Therefore, angles measured near 90° are more likely to produce incorrect results.

Note that the target range in trials PFT006 and PFT007 diminishes from 212.997 to 205.034 m. This is because the target range is calculated with reference to channel A. In PFT006, the 6.22 m cable is connected between the splitter and channel A. In PFT007, the 15.4 cm cable is applied to channel A, with channel B utilising the 6.22 m cable. This causes the perceived target position to diminish.

A similar result with reference to the target range is observed between trials PFT008 and PFT009, PFT016 and PFT017, and finally PFT018 and PFT019.

The difference in range between PFT006 and PFT008 is $3.981 \times \frac{2}{3 \times 0.83} = 3.197$ m, which also approximates the length of the additional 3.12 m cable used to extend the 6.22 m cable to 9.34 m. Note that this result has been scaled to account for the difference in propagation speed in the long 208.42 m cable and the shorter coaxial cables. This result

is near identical to that measured using the phase difference. Thus the 3.12 m cable length is equivalently estimated using the time of arrival and phase difference.

The naïve AOA in PFT006 and PFT007 switches polarity from -33.5659° to 33.7194° when channels A and B are reversed. This is expected since the phase difference toggles polarity. Ideally one would hope the magnitude of the perceived AOA to still be identical for symmetrical reasons. Nevertheless, they differ by 0.15° (in magnitude), which is still considered acceptable.

It is further interesting to observe how the naïve AOA measurement compares to the Standard Deviation (STDEV) based AOA algorithm. In trials PFT005, PFT006 and PFT008 the phase difference is artificially increased by extending the cable lengths used on channel A. The naïve algorithm measures the AOA to be 0.0589° , -33.5659° and -13.4792° respectively. Clearly the measured angle is aliasing. The STDEV algorithm perceives the AOA to be 0.0593° , 38.4785° and 70.307° . Theoretically one is meant to measure 0° , 37.3441° and 66.7029° . The STDEV algorithm is shown to successfully unwrap the AOA, but degrades in performance as the arrival angle approaches 90° , as predicted in Figures 3.22 and 3.23. A similar progression of angles were measured in the complimentary trials of PFT015, PFT016 and PFT018.

Trial PFT023 measures the time taken for a signal traversing from the DAC output into the splitter and along the two 18 m cables used to connect the receive antennas during the field experiment. Figure 4.7 shows that both channels on the oscilloscope receive the signal at the same time. No phase differences are induced by using these cables. This test was simply performed as a precautionary measure.

Each 18 m cable was measured to produce a loss of 1.5 dB using the network analyser. Taking this into account, and the 3 dB drop across the splitter, one expects to measure a peak value of $\pm 0.6 V$, which is approximately that illustrated.

Finally, PFT024 assess the synchronisation logic between two linked Red Pitayas. The master device commands the slave transceiver to begin transmitting at an identical time. Figure 4.8 illustrates that the two independently generated pulses on each device was received by the oscilloscope concurrently. No offset is evident. Both pulses reach $2 V_{pp}$ as no splitters or long coaxial cables are utilised. This test acts as evidence that both the devices are being clocked correctly and that each is capable of performing in unison.

This concludes the discussion on the pre-field tests. No obscure or unexplainable results were presented. All power, energy, range and phase measurements using the Red Pitaya were shown to be within reasonable accuracy of expected values. Both the I/Q and Hilbert-based demodulators operating on the FPGA were shown to derive strong correlated measurements with each other. This is to be expected as theoretically each demodulator is anticipated to perform identically.

5.2 Comments on the Outdoors Interferometry Experiment

Unlike with the Pre-Field Test (PFT) trials, the Field Measurement (FM) tests experience far more exposure to noise, power loss and inconsistencies due to the use of the antennas and non-lab environment. Figures 4.11 to 4.30 illustrate several useful metrics which are used to assess the measuring capabilities.

Each illustration is identical in format. The top two left and right panels indicate the received signal for each PRI in channels A and B respectively. Since the recording length is 41.94 ms, and the PRI 1 ms, a total of 41 full PRI segments are recorded. The first PRI is discarded as it is used by the Write-Back module. Therefore, all 40 PRI returns are superimposed and depicted. The lower 2 panels illustrate the result of integrating these PRIs. Note that the signal shown reflects that after it has been passed through the matched filter and been pulse compressed.

The target appears to be consistently detected at 479.6679 m for all measurements which are derived from the channel that is physically connected to antenna RX_A . This translates to channel A for trials FM00A-FM00J, and to channel B for trials FM00AA-FM00JJ. Recall that the double suffix code in ‘FM00[AA-JJ]’ means that the measurement was taken with the inputs to channel A and B reversed.

The range is calculated using the speed of light in free-space, which does not take into account the true speed of propagation within the 208.42 m cable. A correction is therefore required as the signal first travels 208.42 m along the coaxial, shortly before travelling another 208.42 m over air back towards RX_A .

The apparent length of the 208.42 m coaxial is $0.837 \mu\text{s} \times c = 250.93$ m (adapted from Equation 4.3). Therefore, subtracting 479.6679 m by this amount yields 228.74 m. Reducing this further by 18 m to account for the additional coaxial used to connect the Red Pitaya to RX_A , one calculates the true one-way target distance to be 210.74 m. This is remarkably accurate, and measures within 2.32 m of the transmitter’s actual range. Allowance is also required as over 3.45 m of cable is used to stretch across the height of the conical antenna.

The target’s uncorrected range from the perspective of antenna RX_B at location A is 467.6762 m. As one moves towards position J, this amount reduces further to 455.6845 m and eventually to 443.6928 m. Since the range bin size is 12 m (with regards to the Hilbert-based demodulator), these changes in distance are only detected after RX_B has been rotated in increments of roughly 20-30°. One expects the distance to reduce since RX_B is brought closer to antenna TX_A (see Figure 4.9).

The next important factor to consider is the average received power. As antenna RX_B was relocated fairly often, its S11 performance varied. This is evident when one considers the standard deviation in the received power in Table 4.5. Observe that deviation in the channels connected to RX_B varies substantially in trials FM00C, FM00D, FM00EE,

FM0FF, FM00H, FM00I, FM00II, FM00J and FM00JJ. In the worst case, this deviation approaches 19.6 dBm. This is not desirable as it suggests that the received signal through RX_B was subject to varied reception performance. Relocating the antenna was a complicated task as it weighed in excess of 20 kg, was 3.45 m tall and had to be tied down with guide rope after every relocation. One can observe the deviation to steadily worsen from trials FM00A to FM00J. These were performed in chronological order. This implies that the antenna became more degraded with every repeated reconstruction.

Antenna RX_A on the other hand performed more reliably as it was not relocated. The power measurements produced from data received by it appears more uniform. In all cases, the standard deviation in average power is less than 2.881 dBm.

One can observe the stark contrast between the two receiving antennas by comparing Figures 4.11 and 4.25. The former measurement was taken at the start of the experiment, while the latter was taken after 7 relocations. Clearly each received fast-time echo in the first illustration is more consistent in channel B by comparison. This degradation correlates with the recorded standard deviation in received power through RX_B .

This variation in received average power does not strongly influence the AOA computations, except for in trials FM00I, FM00II, FM00J and FM00JJ.

The amount of received power is similar to that predicted by the link-budget presented in Table 4.4. It states to measure -29.408 dBm by the receiver which is within range of that being recorded in reality.

Since the receiving antennas are separated by 30 m, they are unable to interpret targets residing outside the range of $\pm 17.10^\circ$ at 17 MHz. A maximum separation distance of $\lambda/2 = 8.82$ m is needed in order to prevent angle aliasing. Clearly this arrangement forces it to occur.

If one only considers the AOA in the first quadrant, using the naïve algorithm, one detects the target angles from positions A-J to be respectively in degrees (Table 4.2):

[0.2018 10.2619 -11.0437 -1.1937 4.9488 10.5415 14.9312 -16.0376 -11.8576 -11.9243]

The first two measurements are correct, while the rest alias. This wrapping was predicted and shown in Figure 3.20. The last two readings corresponding to positions I and J are considered anomalies. They do not alias correctly as anticipated. This is due to the degradation of antenna RX_B . If one observes the data in channel B of Figures 4.27 and 4.29, one can see a large variation in comparison to the recordings taken at the other locations. This correlates to the discussion on the variable received power in antenna RX_B .

The Chinese Remainder Theorem (CRT) algorithm measures:

[NaN -50 -72.5 -61 82 59.5 NaN 70 -73.5 NaN]

The unit 'NaN' is used to represent that no solution could be found. This algorithm performs terribly and is unreliable. It has only one correct measurement at 70° . The

CRT method is found to be non-robust and unpredictable. This algorithm was predicted to break as any slight non-linearity results in a poor index and AOA mapping (Section 3.3.3).

The STDEV algorithm measures the AOA at positions A to J respectively:

[0.3507 10.4629 20.2718 29.9705 40.5473 50.8656 60.2108 69.9923 25.4203 25.5116]

There is strong evidence to suggest that the angle of arrival is correctly expanded for angles from 0° to 70° . As with the naïve algorithm, the last two measurements appear as anomalies. Both produce near the same answer of roughly 25.5° . Antenna RX_B was physically repositioned between these two measurements, which should have hypothetically adjusted the phase difference. In this case, one assumes the last two measurements to be different - if even incorrect. This result is theorised to be due to damage inflicted on antenna RX_B while relocating it and not on the STDEV algorithm. The large variation in received power through RX_B at positions I and J supports this argument.

For the 4th quadrant results, the naïve, CRT and STDEV algorithms measure the target to be residing at the following angles respectively, by row (Table 4.2):

-0.5447	-9.4622	11.913	0.5997	-5.7619	-10.5937	-14.3975	15.4687	11.5529	11.5642
-60.5	-69	<i>NaN</i>	-59.5	-52	-60	<i>NaN</i>	<i>NaN</i>	<i>NaN</i>	-57
-0.4699	-9.5898	-19.6634	-30.77	-40.8323	-50.9725	-59.4251	-70.2943	-25.5227	-25.5451

The naïve AOA aliases predictably with exception to the last two measurements. As with the first quadrant, the first two angles are not aliased as they reside within 17.10° . The CRT method fails to correctly measure any plausible AOAs. The STDEV algorithm unwraps the AOA properly up to -70° but fails to do so at positions II and JJ. Once again, this can be attributed to antenna RX_B .

In terms of quantifying the accuracy, the average Mean Square Error (MSE) of the naïve and STDEV algorithm is 1763.81° and 0.29° respectively, if one only considers positions A-H and AA-HH. The STDEV algorithm is therefore shown to be a vast improvement. Positions I, II, J and JJ are not considered in this evaluation as they appear to be genuine anomalies. The MSE of the CRT algorithm is practically meaningless due to its unpredictable performance.

The MSE of the STDEV algorithm was predicted to be within the order of around 0.1° for angles within 70° given an SNR of -5 dB (Figure 3.22). The measured MSE is slightly worse to that predicted. Given the real world environment, one expects the accuracy to degrade slightly from that modelled.

Table 4.6 expands on the STDEV algorithm by considering numerous different frequencies, window functions and demodulators. The MSE (this time catering for positions I, II, J and JJ) is used to derive conclusions.

The first noticeable result shows that the use of closely spaced carrier frequencies degrades performance. This was anticipated in Figure 3.23. Transmitting at 18, 18.5 and 19 MHz

performs considerably better than at 18.01, 18.02 and 18.03 MHz. The average MSE is more than twice as large in the latter case.

The use of a Hamming window was found to marginally influence the MSE. It improved the performance of the I/Q demodulator but worsened that of the Hilbert-based demodulator. The Hamming window is expected to degrade the range resolution in the received signal, but not overly effect the extrapolated phase difference. Hence this has little influence on the measured AOA.

Transmitting at 17.5 and 19 MHz yielded the best result in comparison to the other variations in Table 4.6. The large carrier gap improves the result. Utilising the three default carriers of 17, 18 and 19 MHz proved to be the best. As anticipated, the use of more carrier frequencies benefits measurement accuracy.

Finally, the average MSE is shown to be consistently lower when using the Hilbert-based demodulator. Given its greater sampling rate, one anticipates a more sensitive phase measurement and hence a better result. It also uses much higher order filters, reducing inaccuracies.

In comparison, the I/Q demodulator produces a quarter of the number of samples for any given reading. This inherently results in faster processing times when performing the interferometric algorithms. Therefore, the I/Q demodulator offers speed in exchange for accuracy.

5.3 Comments on the Moving Targets Experiment

The initial purpose of the additional experiment involving moving targets was to track aircraft in azimuth. This was intended to be a further extension of the interferometric trials. Unfortunately, however, the means to successfully do this could not be achieved as no moving targets could be consistently detected.

This can be attributed to a number of reasons. The most prominent of which is that not enough transmit power could be generated. The transmitter could not be upgraded before the deadline of this research given that the high-powered amplifier and filters had already been ordered from Russia, taken over 8 weeks to arrive, and cost in excess of \$575.

In response to this dilemma, emphasis was applied on the small-scale trials discussed in Section 5.2 which successfully managed to demonstrate that a stationary target could be detected in azimuth, well beyond the maximum unambiguous angle using the STDEV algorithm.

Nevertheless, striving to detect moving targets proved beneficial, as it enriched the developed theory in Sections 3.2.3 and 3.2.6.

Figure 4.32 illustrates the range-Doppler map captured from transmitting outwards from above the University of Cape Town (UCT). Figure 4.33 overlays an equivalent map of the illuminated surroundings.

The labels A-E correspond on each illustration. An integration of over 100 consecutive fast-time returns were used to improve the SNR. This, in combination to using a carrier of 16.15 MHz and PRF of 2 kHz, results in a maximum measurable unambiguous velocity of ± 92.878 m/s. The maximum unambiguous range is 75 km.

This maximum unambiguous velocity is low for aircraft which travel in the order of 200 m/s. As they continuously fly near tangential to the radar's LOS (given Cape Town's commercial flight path in relation to UCT), the radial velocity is expected to be much lower. Furthermore, if targets were detected and happened to alias, one would still see them - which one does not. Therefore, more echoed returns were integrated in exchange for a lowered maximum measurable velocity. Integrating only 40 pulses, with a maximum unambiguous velocity of 232.19 m/s did not realise moving targets either.

Interestingly, the configuration managed to detect the surrounding mountain ranges along the zero-Doppler line. Labels D and E correspond to the Jonkershoek reserve, while A-B and C map to the mountains around the Silvermine reserve. The dip in received power between A and B can be attributed to shadowing. The artefacts before 4 km and after 68 km are due to the diode switches used to protect the receiver from the high-powered amplifier.

In practice, far more work was dedicated to detecting moving targets than that documented due to the unsuccessful result. See Burger, however, for further research on detecting moving targets using the exact FPGA design utilised for the interferometer [49].

5.4 The Research Question

Returning focus to the research question presented in Chapter 1, it is evident that the azimuthal angle of a target can be detected using phase interferometry. Given two receiving antennas, an ambiguity in the AOA will always persist, provided they are separated further than half a wavelength of the carrier.

All relevant theory was developed in which a practical demonstration was shown to successfully and autonomously to expand a maximum unambiguous AOA from $\pm 17.10^\circ$ up towards $\pm 70^\circ$ using frequency agility. The average MSE across all angle measurements was empirically determined to be 0.29° .

Chapter 6

Summary & Conclusion

This research and project came about as a result of the interferometric errors associated with the geometrical antenna configuration of the SANAE IV SuperDARN radar. Given the large separation distance between the SuperDARN primary and secondary array being 100 m, in excess of several wavelengths, the system is unable to correctly measure targets at elevation angles greater than approximately 33° [2]. The aim of this project was to develop an HF interferometer demonstrator and to investigate techniques to accurately measure the AOA. This research was conducted in collaboration with SANSA.

Chapter 2 built the context of this work and discussed three relevant case studies on the calibration and evaluation of the SuperDARN interferometric capabilities. The most promising of which, involving the Buckland Park study [2], alters the generic antenna layout to resolve the AOA ambiguities. A study involving COTS Universal Software Radio Peripherals (USRPs) was also shared. It elaborated techniques used to measure target angles involving Phase Difference of Arrival (PDOA) corrections with unsynchronised transceivers. Background theory on HF propagation, refraction and particle detection in the ionosphere was also provided.

Chapter 3 developed a mathematical model describing the characteristics of a fully coherent pulsed-Doppler radar. Theory regarding the transmitted waveform, environment, moving target model, reconstruction filter, pulse compression, coherent integration and sidelobe reduction was discussed. Topics on range resolution, blind range, SNR improvements and demodulation were additionally developed. The theoretical content was kept relevant to the system under investigation, which was intended to lead smoothly into the digital implementation stage discussed in the latter half of Chapter 3.

Therefore, focus was given to the Red Pitaya, an FPGA hardware platform used to build the transceiver utilised for the interferometric experiments. Details regarding the internal AXI communication protocols, PS and PL synchronisation, register mappings, transmit DAC buffer design, Write-Back module and daisy-chaining capabilities were shared.

In addition, three algorithms were developed to unwrap the AOA. The first was considered a naïve measurement. It simply derived an angle from a given phase shift but did not account for any phase wrapping due to the result of a particular spatially separated antenna configuration. The Standard Deviation (STDEV) algorithm focused on autonomously unwrapping the AOA by varying the utilised carrier frequency and translating several aliased angles into a rectified estimation. The final algorithm relied on the Chinese Remainder Theorem (CRT).

Chapter 4 documented two major sets of experiments. The first group involved lab measurements and was used to validate the performance of the developed transceiver. It set out to test and successfully confirm that the Red Pitaya was capable of transmitting the correct waveform, measuring power, interpreting phase shifts and performing RF down-conversions. The second set of experiments contained an array of interferometric trials, in which the AOA produced from a beacon transmitter was calculated using the three algorithms. An additional sub-task was implemented to detect moving aircraft, however no targets were observed owing principally to insufficient transmit power. Nevertheless, the trial was valuable for establishing the system foundations.

The interferometric trials, discussed in Chapter 5, concluded that the Red Pitaya was capable of accurately down-converting a received RF signal within the range of 12.5-25 MHz using either its I/Q or Hilbert-based demodulator.

The naïve measurement of the AOA was shown to wrap consistently up to 70° before varying unpredictably. Similarly, the STDEV algorithm was able to successfully expand a hypothetical maximum unambiguous AOA of $\pm 17.10^\circ$ to $\pm 70^\circ$. It produced an average MSE of 0.29° between the target's true and measured angle. The CRT algorithm performed poorly, yielding 1 correct target angle from 20. It is simply unable to deal with non-linearities present in real-world measurements. Readings taken at angles $\pm 80^\circ$ and $\pm 90^\circ$ were considered degraded due to damage inflicted on the second receiving antenna, RX_B .

Overall, this dissertation demonstrates an application of the principles of interferometry. The developed theory governed the design of the FPGA hardware, in which the large majority of the practical experiments performed on it were shown to produce anticipated results. An algorithm involving frequency agility was shown to successfully expand the maximum unambiguous measurable AOA in azimuth.

6.1 Portability to SANAE IV SuperDARN Radar System

As a new separate system demonstrator to the SANAE IV radar was developed, it required far more theoretical and practical understanding. Working locally on this project as opposed to in Antarctica facilitated access to many more resources, sped production and lowered costs. However, as the results were not measured on the actual Super-

DARN radar, conjectures have to be made. These are proposed to be ported onto the SuperDARN firmware for testing.

In the local experiment, the maximum unambiguous AOA was shown to be successfully expanded from $\pm 17.10^\circ$ to $\pm 70^\circ$. This is a vast improvement, which could have been bettered had antenna RX_B performed more consistently. The AOA expansion was achieved by only adapting the carrier frequency.

It is therefore hypothesised from this demonstration, that one can unwrap a given ambiguous AOA by adapting the carrier frequency used by the SANAE IV radar. If a mathematical model to govern the received phase shift on the larger radar could be developed, as indicated by this research, it is assumed that targets could be resolved by unwrapping multiple aliased echoes. This was shown to be successfully done in azimuth, however, the principle can also be applied in elevation. The solution to the interferometric problem experienced by the SuperDARN radar is therefore proposed by an adjustment to its transmitted signal sequence.

This further needs to take into account that the local experiment relied on a single point target. In reality, the ionosphere consists of an abundance of charged particles which need to be modelled as an extended target.

Compensation for multi-frequency target detection also needs to be considered. Inherently, this also impacts the illuminated HF trajectory, refractive properties and time-delays. This would impact the performance of the discussed algorithms.

6.2 Recommendations

There are three areas of work which can be improved. The first is the moving targets experiment. Although the stationary transmitter used in the interferometric trials was successfully detected, moving aircraft were not. This limited insight into the AOA algorithms used for moving targets. It would have been interesting to model this.

The FPGA design can also be optimised. Currently the I/Q demodulator for each channel utilises a 20 order low-pass filter. This uses all the available DSP slices. If channel B were stripped from the FPGA design (freeing resources) and the low-pass filter on channel A was upgraded to 40 taps instead, it would result in a more precise demodulator. A second Red Pitaya could then be synchronised with 40 taps to act in replacement for channel B. Alternately, the Hilbert-based demodulator and sinusoidal DDS cores could be removed to reallocate usable logic. For now, the developed FPGA design acts as an all-round provider. Its performance can be fine-tuned to specific applications as mentioned in Section 3.4.11 by redistributing the usable logic.

The final suggestion is to test for targets in different arrangements. This work only considered the target residing within $\pm 90^\circ$. It is recommended that alternative AOAs which reside from behind the FOV should be catered for.

Bibliography

- [1] R. Greenwald, K. Baker, J. Dudeney, M. Pinnock, T. Jones, E. Thomas, J.-P. Villain, J.-C. Cerisier, C. Senior, C. Hanuise, *et al.*, “DARN/SuperDARN,” *Space Science Reviews*, vol. 71, no. 1-4, pp. 761–796, 1995.
- [2] E. Custovic, A. J. McDonald, J. Whittington, D. Elton, T. A. Kane, and J. C. Devlin, “New antenna layout for a SuperDARN HF radar,” *Radio Science*, vol. 48, no. 6, pp. 722–728, 2013.
- [3] G. Bryson, W. Bristow, and T. Parris, “An introduction to radar and the super dual auroral radar network.” http://superdarn.gi.alaska.edu/tutorials/SuperDARN_Radar_Fundamentals.pdf. Accessed: 2016-03-18.
- [4] “Radar overview: VirginiaTech SuperDARN.” <http://vt.superdarn.org/tiki-index.php?page=Radar+Overview>. Accessed: 2016-04-13.
- [5] G. Fabrizio, *High frequency over-the-horizon radar: fundamental principles, signal processing, and practical applications*. McGraw Hill Professional, 2013.
- [6] “Convection map overview: VirginiaTech SuperDARN.” <http://vt.superdarn.org/tiki-index.php?page=Conv+map+overview>. Accessed: 2016-04-16.
- [7] P. Ponomarenko, N. Nishitani, A. V. Oinats, T. Tsuya, and J.-P. S. Maurice, “Application of ground scatter returns for calibration of HF interferometry data,” *Earth, Planets and Space*, vol. 67, no. 1, pp. 1–9, 2015.
- [8] D. Agaba, “Calibration of a SuperDARN radar antenna by means of a satellite beacon,” Masters Dissertation, University of Cape Town, 2012. <https://open.uct.ac.za/handle/11427/13917>.
- [9] J. H. Reed, K. J. Krizman, B. D. Woerner, and T. S. Rappaport, “An overview of the challenges and progress in meeting the e-911,” *IEEE Communications Magazine*, 1998.
- [10] J. Gowans, “Design and implementation of an RFI direction finding system for SKA applications,” Masters Dissertation, Pending Publication, University of Cape Town, 2017.

- [11] M. Kim and N. Y. Chong, “RFID-based mobile robot guidance to a stationary target,” *Mechatronics*, vol. 17, no. 4, pp. 217–229, 2007.
- [12] M. I. Skolnik, *Radar Handbook*. 3rd ed., 2008.
- [13] “Ray-tracing: VirginiaTech SuperDARN.” <http://vt.superdarn.org/tiki-index.php?page=Ray-tracing>. Accessed: 2016-04-25.
- [14] T. Ida, “Crystal structure analysis,” October 2013.
- [15] D. P. Cilliers, “SANAE HF radar calibration,” tech. rep., Hermanus Magnetic Observatory, Hermanus, South Africa, 2007.
- [16] S. Milan, T. Jones, T. Robinson, E. Thomas, and T. Yeoman, “Interferometric evidence for the observation of ground backscatter originating behind the CUTLASS coherent HF radars,” in *Annales Geophysicae*, vol. 15, pp. 29–39, Springer, 1997.
- [17] R. Van Zyl, D. Visser, P. Cilliers, and B. Opperman, “ZACUBE-1 space weather mission: Characterize the SuperDARN HF radar antenna array at SANAE-IV,” *Space Weather*, vol. 11, no. 2, pp. 52–54, 2013.
- [18] F. S. Minko, “CubeSat mission design for characterising the dual auroral radar network (SuperDARN) field-of-view,” Masters Dissertation, Cape Peninsula University of Technology, 2013. <http://digitalknowledge.cput.ac.za/jspui/handle/11189/1315>.
- [19] K. Davies, *Ionospheric radio*. No. 31, IET, 1990.
- [20] B. G. Evans, *Satellite communication systems*, vol. 38. IET, 1999.
- [21] D. de Villiers and R. van Zyl, “ZACube-2: The successor to Africa’s first nanosatellite,” , SA AMSAT Space Symposium, 2015.
- [22] R. A. Greenwald, “Steering SuperDARN radars.” http://vt.superdarn.org/tiki-download_file.php?fileId=1060. Accessed: 2016-06-20.
- [23] H.-C. Chen, T.-H. Lin, H. Kung, C.-K. Lin, and Y. Gwon, “Determining RF angle of arrival using COTS antenna arrays: a field evaluation,” in *2012 IEEE Military Communications Conference*, pp. 1–6, IEEE, 2012.
- [24] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, “On the effectiveness of switched beam antennas in indoor environments,” in *International Conference on Passive and Active Network Measurement*, pp. 122–131, Springer, 2008.
- [25] K. Sayrafian-Pour and D. Kaspar, “Source-assisted direction estimation inside buildings,” in *IEEE Conference on Computer Communications*, 2006.
- [26] M. Maróti, P. Völgyesi, S. Dóra, B. Kusý, A. Nádas, Á. Lédeczi, G. Balogh, and K. Molnár, “Radio interferometric geolocation,” in *Proceedings of the 3rd international conference on Embedded networked sensor systems*, pp. 1–12, ACM, 2005.

- [27] Y. Kono, H. Hanada, P. Jinsong, Y. Koyama, Y. Fukuzaki, and N. Kawano, "Precise positioning of spacecrafts by multi-frequency VLBI," *Earth, planets and space*, vol. 55, no. 10, pp. 581–589, 2003.
- [28] C. M. Rose, "Method of detection and determining an angular location of frequency agile emitters," Nov. 6 2001. US Patent 6,313,794.
- [29] R. J. Lioio, G. E. Clayton, and R. A. Deaton, "Phase and time-difference precision direction finding system," Mar. 3 1998. US Patent 5,724,047.
- [30] R. L. Goodwin, "Ambiguity-resistant three-and four-channel interferometers," tech. rep., NAVAL RESEARCH LAB WASHINGTON DC, 1976.
- [31] S. Shepherd, "Elevation angle determination for SuperDARN HF radar layouts," *Radio Science*, vol. 52, no. 8, pp. 938–950, 2017.
- [32] A. G. Burrell, T. K. Yeoman, S. E. Milan, and M. Lester, "Phase calibration of interferometer arrays at high-frequency radars," *Radio Science*, vol. 51, no. 9, pp. 1445–1456, 2016.
- [33] R. A. Greenwald, N. Frissell, and S. Larquier, "The importance of elevation angle measurements in HF radar investigations of the ionosphere," *Radio Science*, vol. 52, no. 3, pp. 305–320, 2017.
- [34] P. A. Bernhardt, C. L. Siefring, J. F. Thomason, S. P. Rodriguez, A. C. Nicholas, S. M. Koss, M. Nurnberger, C. Hoberman, M. Davis, D. L. Hysell, *et al.*, "Design and applications of a versatile HF radar calibration target in low Earth orbit," *Radio Science*, vol. 43, no. 1, 2008.
- [35] A. Wilkinson, "Notes on radar/sonar signal processing: fundamentals," University of Cape Town, Radar Course Master Notes, 2016.
- [36] B. R. Mahafza, *Radar systems analysis and design using MATLAB*. CRC press, 2002.
- [37] M. A. Richards, *Principles of modern radar*. Raleigh, NC: SciTech Pub, 2010.
- [38] C. Leśnik, "Nonlinear frequency modulated signal design," *Acta Physica Polonica A*, vol. 116, no. 3, pp. 351–354, 2009.
- [39] Q. Wang, *Non-Linear Chirp Spread Spectrum Communication Systems of Binary Orthogonal Keying Mode*. PhD thesis, The University of Western Ontario, 2015.
- [40] S. Golomb and R. Scholtz, "Generalized barker sequences," *IEEE Transactions on Information Theory*, vol. 11, no. 4, pp. 533–537, 1965.
- [41] STEMLab, *Red Pitaya Hardware Specifications*, April 2014.
- [42] STEMLab, *FPGA memory map*, April 2016.
- [43] Xilinx, *AXI Reference Guide*, March 2011.

- [44] “AXI interfacing using DMA & AXI4-stream.” George Mason University, February 2016.
- [45] Xilinx, *DDS Compiler v6.0 - LogiCORE IP Product Guide*, November 2015.
- [46] F. Nicolls, “Introduction to DSP,” University of Cape Town, Course Notes, 2016.
- [47] Xilinx, *Complex Multiplier v6.0*, November 2015.
- [48] A. Chen, R. McDanell, M. Boytim, and R. Pogue, “Modified CORDIC demodulator implementation for digital IF-sampled receiver,” in *Global Telecommunications Conference, 1995. GLOBECOM'95., IEEE*, vol. 2, pp. 1450–1454, IEEE, 1995.
- [49] J. Burger, D. W. O’Hagan, and A. J. Wilkinson, “HF surface wave radar using low-cost components and SDR,” Masters Dissertation, Pending Publication, University of Cape Town, 2017.
- [50] Z. Patel, D. W. O’Hagan, and A. J. Wilkinson, “The design and build of an antenna for high frequency surface wave radar,” Bachelors Dissertation, University of Cape Town, 2016.
- [51] Huber+Suhner, *Coaxial cable: S_04162_B-60*, March 2009.
- [52] International Telecommunication Union, *Recommendation ITU-R P.372-7 - Radio Noise*, April 2003.
- [53] G. H. Hardy, J. E. Littlewood, and G. PĀşlyā, *Inequalities (Cambridge Mathematical Library)*. Cambridge University Press, 1988.
- [54] P. Demin, “Red pitaya notes.” https://github.com/pavel-demin/red-pitaya-notes/tree/master/cores/axis_ram_writer_v1_0, 2015. Accessed: 2016-11-01.

Appendix

A.1 Source Code Access

All of the source code developed for this MSc project was managed and maintained in four different GIT repositories. The description and URL of each is provided below. All access is made publicly available.

1. **General Matlab, C & Dissertation Repository:**
<https://Tighe-Barris@bitbucket.org/Tighe-Barris/masters>
2. **Vivado FPGA Project Repository:**
<https://Tighe-Barris@bitbucket.org/Tighe-Barris/redmasters>
3. **Red Pitaya Run-Time Code Repository:**
<https://Tighe-Barris@bitbucket.org/Tighe-Barris/redmasters-c>
4. **Raw Binary File Recordings:**
<https://Tighe-Barris@bitbucket.org/Tighe-Barris/masters-recordings>

A.2 Proof of Matched Filter Peak SNR Improvement

Recall from Chapter 3:

$$S_{n_b}(f) = \frac{\eta}{2} \quad (\text{A.2.1})$$

$$\overline{|n_o^2(t)|} = \int_{-\infty}^{\infty} |H_m(f)|^2 S_{n_b}(f) df \quad (\text{A.2.2})$$

$$m(t_d) = \int_{-\infty}^{\infty} Y_s(f) H_m(f) e^{j2\pi f t_d} df \quad (\text{A.2.3})$$

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} = \frac{\left| \int_{-\infty}^{\infty} Y_s(f) H_m(f) e^{j2\pi f t_d} df \right|^2}{\frac{\eta}{2} \int_{-\infty}^{\infty} |H_m(f)|^2 df} \quad (\text{A.2.4})$$

The Cauchy-Schwarz inequality states that [53]:

$$\left| \int_{-\infty}^{\infty} f(x) g(x) dx \right|^2 \leq \int_{-\infty}^{\infty} |f(x)|^2 dx \int_{-\infty}^{\infty} |g(x)|^2 dx \quad (\text{A.2.5})$$

which is valid for $f(x) = K g^*(x)$, where K is a constant. Therefore:

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} = \frac{\left| \int_{-\infty}^{\infty} Y_s(f) H_m(f) e^{j2\pi f t_d} df \right|^2}{\frac{\eta}{2} \int_{-\infty}^{\infty} |H_m(f)|^2 df} \leq \frac{\int_{-\infty}^{\infty} |Y_s(f)|^2 df \int_{-\infty}^{\infty} |H_m(f)|^2 df}{\frac{\eta}{2} \int_{-\infty}^{\infty} |H_m(f)|^2 df} \quad (\text{A.2.6})$$

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} = \frac{\left| \int_{-\infty}^{\infty} Y_s(f) H_m(f) e^{j2\pi f t_d} df \right|^2}{\frac{\eta}{2} \int_{-\infty}^{\infty} |H_m(f)|^2 df} \leq \frac{2}{\eta} \int_{-\infty}^{\infty} |Y_s(f)|^2 df \quad (\text{A.2.7})$$

$$\frac{|m(t_d)|^2}{\overline{|n_o^2(t)|}} \leq \frac{2E_y}{\eta}. \quad (\text{A.2.8})$$

This is the same result concluded and stated in Equation 3.58.

A.3 List of Common Window Functions

A list of commonly used discretised window function is provided in the Table A.3.1. Each function (excluding the Dolph-Chebyshev window) is assumed to be of length N , where $0 \leq n \leq N - 1$.

Table A.3.1: List of common window function definitions

Window Function	Definition
Rectangular	1
Hann / Raised Cosine	$0.5 - 0.5 \cos\left(\frac{2\pi n}{N-1}\right)$
Hamming	$0.54 + 0.46 \cos\left(\frac{2\pi n}{N-1}\right)$
Kaiser	$\frac{I_0(\pi\alpha\sqrt{1-(\frac{2n}{N-1})^2})}{I_0(\pi\alpha)}, \beta = \pi\alpha$
Dolph-Chebyshev	$\frac{1}{N} \sum_{k=0}^{N-1} \left[\frac{\cos\{N \cos^{-1}(\beta \cos(\frac{\pi k}{N}))\}}{\cosh\{N \cosh^{-1} \beta\}} \right] e^{j2\pi kn/N}$ $\beta = \cosh\left[\frac{1}{N} \cosh^{-1}(10^\alpha)\right]$

Note that $I_0(x)$ indicates the zero-th order modified Bessel function of the first kind.

The Dolph-Chebyshev is defined for $|n| \leq N/2$, where α is a tuning parameter.

A.4 Internal Diagram of Zynq IP Core

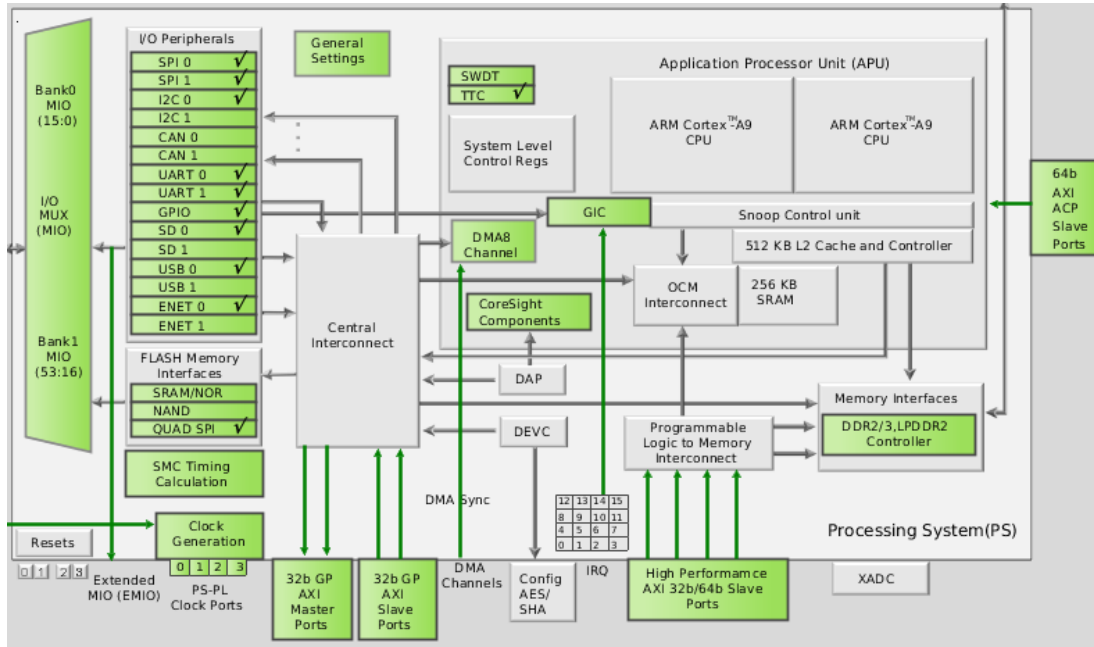


Figure A.4.1: Internal configuration of the Zynq IP Core

A.6 Vivado FPGA Design




Figure A.6.1: Final Vivado FPGA design (Only viewable through PDF zoom)

A.7 Front-End 12.5-25 MHz Filter




Figure A.7.1: Front-end 12.5-25 MHz filter [49]

A.8 Field Test Experiment Images



Figure A.8.1: Constructed fishing gut layout used to mimic Figure 4.9



Figure A.8.2: Antenna RX_A positioned atop the grid shown in Figure A.8.1



Figure A.8.3: Antenna RX_A with antenna RX_B 30 m in the distance past the umbrella.




Figure A.8.4: The Red Pitaya connected with two 50Ω T-terminators, each of which is preceded by a 20 dB low noise amplifier. The DAC output is connected to the transmit antenna, TX_A . The device is powered by a 2.1 A and 5 V external power bank. The Ethernet port connects to a PC which controls the device.




Figure A.8.5: The transmit antenna, TX_A , located roughly 200 m from RX_A .

A.8.1 The Overall Complete FPGA Design & Instruction Set

Each function of the FPGA design has been discussed in isolation. In this section, each component is incorporated and amalgamated to consider the device as a whole. The register sets used to control the behaviour of the digital transceiver is explained.

Figure A.8.6 illustrates the complete design of the interferometer, which acts as an extension of Figure 3.30. The receive chain previously discussed can clearly be distinguished. In addition, the Write-Back, Direct-Input, Daisy-Chaining and DAC Buffer modules have been included. The diagram only depicts the receive pipeline for channel A, as channel B is identical.

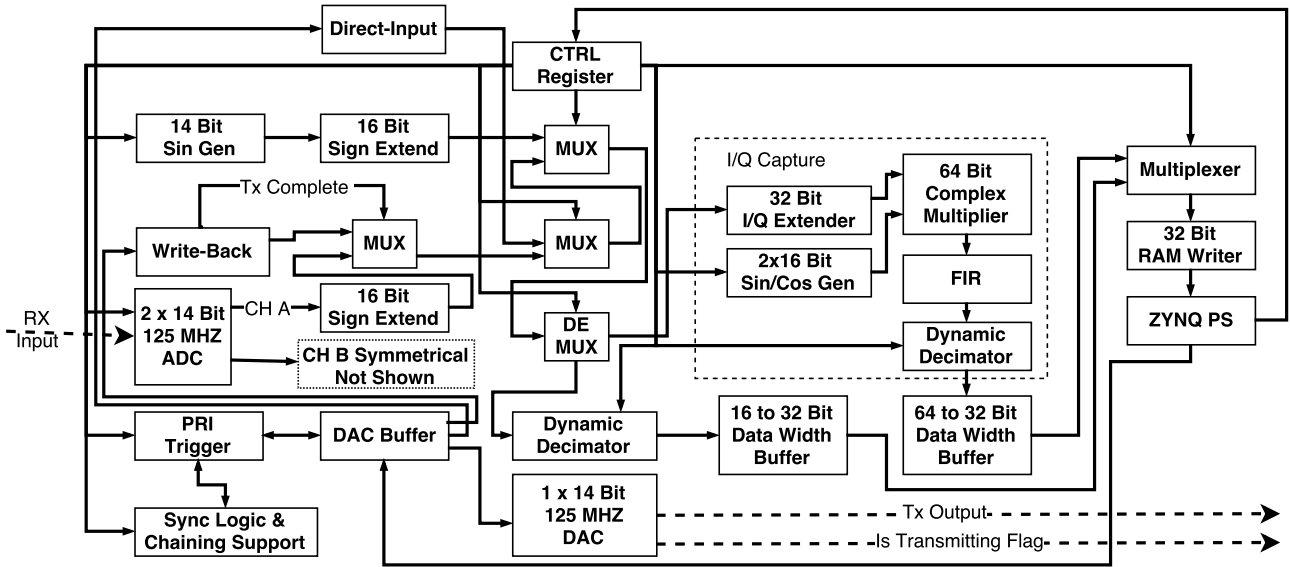


Figure A.8.6: Complete interferometer FPGA design

All the IP cores, with exception of the DDS, FIR, Sign Extenders and Data Width Buffers were custom developed. The Ram-Writer, DAC and ADC interpreting modules were modified from a previous project by Pavel to fit this application [54]. Online access to all the Verilog code is made available in Appendix A.1.

Table A.8.1 outlines the basic resource usage on the FPGA. It is clear that the project is not largely constrained, with exception of the DSP slices. The FIR filters for the I/Q demodulator consume the majority of these components which limit the number of filter taps available.

Table A.8.1: Resource utilisation on the 7z010clg400-1 FPGA device

Site Type	Usage	Available	Percentage
Slice LUT as Logic	3594	17600	20.4%
Slice LUT as Memory	1812	6000	30.2%
Slice Registers as Flip Flops	8529	35200	24.2%
Block RAM: RAMB36	20	60	33.33%
Block RAM: RAMB18	4	120	3.3%
DSP48E1	68	80	85%

The CTRL-Register block in Figure A.8.6 is expanded in Table A.8.2. The collective register set is used to control the FPGA behaviour from the PS. Each register is theoretically constrained to hold a maximum of 64-bits, although each can support up to 4 KB. An extremely sparse instruction set definition is utilised to encourage easy reading. This privilege is made affordable since the address space is not limiting.

A convention is adopted where each register is thought to consist of multiple register variables. Each variable maps to a subset of bits contained within a register, and is written in a bold font to avoid cause confusion.

The ‘Bit Slice’ column in Table A.8.2 denotes this mapping. For example, writing the value 0x3 (0b11) to the CFG register (residing at 0x40000000) will enable the system for I/Q demodulation, disable the ADCs and use a test tone, keep the buffer in reset, disconnect the Direct-Input and Write-Back modules, and mark the device as a master. In addition, the debug LEDs will all be off, no decimation will be applied and no initial recorded samples will be intentionally dropped before being written to the SD card. The test-tone frequency and phase supplied to the quadrature demodulator is specified by the value stored in the GEN register.

Any register variable name ending with the suffix ‘ENABLE’ requires a logic high to be asserted. Similarly, a suffix of ‘RESETN’ requires a logic low to be applied. Those ending with either ‘A’ or ‘B’ refer to each channel on the Red Pitaya. A suffix of ‘N’ requires the value being supplied to be in terms of the number samples being clocked at 125 MHz.

The accompanying C program residing on the Red Pitaya is written to support access to these registers via appropriately documented functions. A link to the repository containing this code is provided in Appendix A.1.

The CFG or configuration register primarily controls the status of the transceiver. It governs the behaviour of several components:

- **SYSTEM_ENABLE**: Writing a logic zero to this bit disables the transceiver. The RAM-Writer, DDS, FIR and buffer cores are held in a reset condition (and cleared of residual memory). In this state the other registers can be safely altered. It is inadvisable to adjust parameters during run-time as to avoid unintended results. This bit is used to properly configure a known state in the PL.
- **TEST_TONE_OVERRIDE_ENABLE**: Setting this bit high disconnects the ADCs, Write-Back and Direct-Input modules. It streams a sinusoidal test tone into the receiver as discussed in Section 3.4.4. The particular frequency and initial phase fed into each channel can be specified with the **PHASE_INC** and **PHASE_OFFSET** variables of the GEN or generation register (see Equation 3.109).
- **HILBERT_DEMOD_ENABLE**: Enabling this bit activates the Hilbert-based demodulator as discussed in Section 3.4.5. Otherwise, a logic zero enables the quadrature demodulator discussed in Section 3.4.6.

Table A.8.2: The Red Pitaya register set & bit mappings

Reg Code	Address	Bit Slice	Register Variable Name
CFG	0x40000000	0:0	→ SYSTEM_ENABLE
		1:1	→ TEST_TONE_OVERRIDE_ENABLE
		2:2	→ HILBERT_DEMOD_ENABLE
		3:3	→ BUFFER_FORCE_RESETN
		4:4	→ DIRECT_INPUT_ENABLE_A
		5:5	→ DIRECT_INPUT_ENABLE_B
		6:6	→ WRITE_BACK_ENABLE_A_B
		7:7	→ DESIGNATE_AS_SLAVE
		15:8	→ UNUSED
		31:16	→ PULSE_LENGTH_N
		47:32	→ PRE_DROP_SAMPLES_N
		55:48	→ DECIMATE_RATE
63:56	→ LED_ARRAY		
STS_A	0x40001000	31:0	→ RAM_WRITER_PTR_ADDR_A
STS_B	0x40002000	31:0	→ RAM_WRITER_PTR_ADDR_B
GEN	0x40003000	15:0	→ PHASE_INC_A
		31:16	→ PHASE_OFFSET_A
		47:32	→ PHASE_INC_B
		63:48	→ PHASE_OFFSET_B
SOF	0x40004000	31:0	→ START_OFFSET_N
PRI	0x40005000	63:0	→ PRI_LENGTH_N
DAG	0x40010000	15:0	→ DAG_VALUE
QUA	0x40011000	31:0	→ PHASE_INC_A_B
		63:32	→ PHASE_OFFSET_A_B
GAL	0x40012000	63:0	→ EXT_OFFSET_LEFT_N
GAR	0x40013000	63:0	→ EXT_OFFSET_RIGHT_N
SLA	0x40014000	63:0	→ SLAVE_OFFSET_N

- **BUFFER_FORCE_RESETN:** This bit is used to force the read and write registers of the DAC buffer (Section 3.4.7) into a known state such that it can be populated correctly.
- **DESIGNATE_AS_SLAVE:** Enabling this bit adjusts the PL to adopt an external clock and activate signal which is fed into the Red Pitaya through its SATA connectors. The intended purpose of this functionality is detailed in Section 3.4.9.
- **DIRECT_INPUT_ENABLE** and **WRITE_BACK_ENABLE** are used to

activate or disable the additional modules as explained in Section 3.4.8. Note that the Write-Back module is applied to both channels A and B, whereas the Direct-Input module can be toggled for each channel independently.

- **PULSE_LENGTH_N** defines the number of samples used to represent the transmitted pulse. For example, a 20 μ s pulse is represented with 2500 cycles at 125 MHz. This value is used to operate the DAC buffer, as shown in Figure 3.32. Any 16-bit number written to **DAG_VALUE** is automatically ‘pushed’ onto the DAC buffer.
- **PRE_DROP_SAMPLES_N** instructs how many initial samples should be dropped to align the transmitted signal with the received signal. This caters for the lag experienced between the DAC, ADC and filter group-delay as discussed in Section 3.4.5.
- **DECIMATE_RATE**, as the variable name suggests, defines the decimation rate. For the quadrature demodulation, a value of 10 is used to achieve an effective down-sample rate of 20 as the FIR filters already decimate by 2. For the Hilbert-based demodulator, a value of 5 is used. The reasoning and academic understanding for this was raised in Sections 3.4.5 and 3.4.6.
- **LED_ARRAY** illuminates a row of 8 on-board LEDs which can be used for debugging. A logic high represents a lit LED. The pattern displayed matches the stored 8-bit value.

The STS_A and STS_B registers represent the status or pointer locations of the RAM-Writers used for each channel respectively. The former register is illustrated in the explanation of Figure 3.27.

The SOF or start off-set register works similarly to the **PRE_DROP_SAMPLES_N** register variable. It defines how many many clock cycles should pass (after **SYSTEM_ENABLE** has transitioned high) before the ADCs activate. It is generally used to debug and align the received signal to synchronise with the DAC output.

The PRI register defines the number of 8 ns cycles which should pass before the DAC buffer is reset and a new pulse is transmitted. A PRI of 1 ms, for example, requires 125000 to be stored in this register. If one sets **PRI_LENGTH_N = PULSE_LENGTH_N**, then a CW signal is defined.

The phase and frequency of the sinusoids used in the quadrature demodulator are defined in the QUA register. The frequency is typically chosen to match the carrier used in the transmitted signal. This forms part of multiplying by $e^{-j2\pi f_c t}$ in order to baseband the received signal, as shown in Figure 3.5.

The GAL and GAR registers are used to set how many cycles should pass before the external trigger pin goes high and low respectively relative to the start of the PRI. This allows the external pin’s logic high duration and position to be adjusted as illustrated

and discussed in Figure 3.36.

Finally, the **SLAVE_OFFSET_N** register variable allows one to delay the activation of a given Red Pitaya by N clock cycles after the **SYSTEM_ENABLE** variable has transitioned high. It is used to synchronise multiple devices in a daisy-chain configuration. Typically the master device will trigger the slaves, but delay itself by 1 clock cycle - the equivalent time taken for which the slaves require to acknowledge the request travelling over the physical SATA cable. This requirement was detailed in Section 3.4.9.