

FAST INTERFERENCE LOCALIZATION
TO PROTECT
GLOBAL NAVIGATION SATELLITE SERVICE OPERATIONS

A DISSERTATION SUBMITTED
TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Adrien Louis Henry Perkins

March 2020

© 2020 by Adrien Perkins. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/kz021dj7117>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

J Powell, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Stephen Rock, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mykel Kochenderfer

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Global Navigation Satellite Systems (GNSS) have become an invisible utility at the core of many modern industries. It is responsible for the safe and reliable navigation of ships and planes, for providing timing required by communications infrastructure and the financial industry, and for being a driving sensor in the autonomous future. It is so critical that many nations, or groups of nations, around the world have spared no expense to develop their own independent systems to not rely on another nation's system. However, by the time the signal reaches the Earth's surface its power is well below the noise floor, and, like any radio system, it is susceptible to interference. This interference, whether intentional or not, can cause disastrous effects on those relying on the reception of GNSS effects. Therefore, the goal of the work presented in this thesis is to develop and demonstrate in real-world environments an autonomous system capable of localizing the source of GNSS interference on a time scale of minutes and hours instead of the days and even months it can take today. To accomplish this, this thesis presents contributions to the hardware required to enable rapid localization from an unmanned aerial vehicle and to the algorithms needed to operate day or night with or without GNSS for navigation.

Previous work on rapid signal localization from an unmanned aerial vehicle relied on three core elements: a sensor to provide a spatial view of the signal strength, an algorithm to observe direction of arrival, and a set of algorithms to estimate the source and plan a path to collect the best measurements possible. This thesis mimics the same flow but provides significant improvements to each step in the process and additionally adds capabilities of navigating in a GNSS denied environment potentially posed by a GNSS interference source. Each improvement to the process is not just theoretical in nature, each system is implemented and tested in real-world environments.

This thesis presents a shift from a bulky physical rotation modality to provide a spatial view of the signal strength (known as a received signal strength pattern) to an electronic rotation modality using a lightweight antenna array, significantly increasing the rate of bearing observations possible

– and therefore data into the estimation algorithm – and de-coupling the spatial sensing from the vehicle platform.

To observe direction of arrival from a spatial view, this thesis presents a robust algorithm that enables capabilities in real-world environments previously not possible. The algorithm developed represents a given measured received signal strength pattern by the “lobes” present within the pattern without the need for modules of the antenna itself. The resulting representation enables the algorithm to perform even in the harshest of real-world environments. To demonstrate these capabilities and the versatility of this algorithm, the algorithm is implemented on thousands of received signal strength patterns generated onboard an unmanned aerial vehicle in harsh environments typical of real-world deployment.

Finally, this thesis presents the development of a test platform capable of 24/7 operations in a GNSS denied environment using an infrared camera. This test platform enables flight testing of the entire system in real-world environments to better inform the system’s capabilities and expected performance outside of a laboratory environment. These real-world tests, comprising of nearly 20 successful GNSS independent, autonomous signal source localizations, are crucial in informing and demonstrating the capabilities of a system as it would apply in non-laboratory conditions, adding an additional dimension over a merely theoretical analysis of different sensors and algorithms required.

Acknowledgments

The journey to complete my Ph.D. has not been a journey that I have walked alone; I have been supported by numerous advisors and mentors and, of course, friends and family along the way. While there are too many to enumerate, there are a handful that deserve a special mention.

First and foremost, my time at Stanford would have ended much earlier had it not been for Per Enge's GPS class and his musing of a UAV based approach to the problem of GPS interference – musing which turned into the focus of my Ph.D. research. I am truly grateful to have been advised by Per before he passed; he has left a lasting impact on me. This is most definitely visible in the research I present in this thesis. He encouraged “doing”, going beyond the realm of theory and simulation to the world of practice. He instilled in me the importance of what can be learned in implementation and facing the challenges the real world imposes, ultimately keeping my ideas and work grounded in reality. His mentorship taught me to go well beyond the familiar, into the truly unknown, to learn and be a better researcher, and a better engineer. Never would I have thought that I would be diving into the world of electrical engineering and building a phased array antenna, but many years later, that is exactly what I did.

I am grateful to Professor J. David Powell for stepping in during a time of need and providing the support and encouragement to “get out the door” in the last few years as his “last student”. Throughout my journey, the Stanford GPS Lab was always there. The researchers of the lab, Dr. Todd Walter, Dr. Juan Blanch, Dr. Eric Phelts, Dr. Sherman Lo, and Dr. Yu-Hsuan Chen, were always accessible and entertained all sorts of questions. They fostered a community that was approachable and supportive – academically and personally, even in the most unexpected of times – that made coming into lab every day truly enjoyable.

With research so heavily based on hardware and testing, results would not have been possible without the support of all those who assisted in the development of hardware and the various flight tests over the years. On the hardware front, Dr. Yu-Hsuan Chen was instrumental in helping bring

my radiofrequency designs to reality with his incredible skills in the development of printed circuit boards and all things hardware in our lab. When it comes to flight testing, you'd think the hardware is the hardest part, but over the years I have learned the paperwork is a challenge in and of itself. For that, I must thank Dr. Sherman Lo who patiently and courageously took care of all the red tape and coordinated with various government entities to enable flight testing, allowing me to focus my efforts on my research. Additional thanks goes to Wei Lee, a graduate of the GPS lab now at Edwards Air Force Base, who worked hard to enable multiple testing campaigns down at Edwards. Last but not least, special thanks goes to my lab mates – Tyler Reid, Kazuma Gunning, Andrew Neish, and Fabian Rothmaier – each of whom I managed to convince that flight testing was all fun and games only for them to stay up through the night in various deserts from New Mexico to Southern California.

I would be remiss if I did not mention those who helped bring to life the very first version of the UAV discussed in this thesis – Dr. Louis Dressel, Mark James, and James Spicer. In particular, Dr. Louis Dressel who not only “co-struggled” in the early days of flight testing as we together learned the ropes of what flight testing truly meant but also had an incredible work ethic that I have strived to match over the years.

The Aero Astro department at Stanford has been a source of incredible teachings, mentorship, and resources, with a faculty that was always there for support along the way. In particular, I want to thank Professor Juan Alonso who provided me with a much-needed outlet for my passion for drones by allowing me to assist with the AA241x class, which provided an incredibly rewarding experience watching and helping students build drone systems from the ground up.

While my days may have been focused on research, I have my friends to thank to make sure my weekends – and sometimes weekdays – were not. From exploring the hills of San Francisco to skiing the slopes of Tahoe, from the vineyards of Napa to the coast of Monterrey, and from competing for the fields of Catan to terraforming Mars, you have all helped me create great memories and lasting friendships. I'm sure my Ph.D. would not have taken as long if it wasn't for “y'all”, but, in the end, I'm glad it did.

I have my lab to thank for the Position, Navigation, and Timing in my research life, but I have my family to thank for the Position, Navigation, and Timing in my personal life. Where I am today, my “Position”, is in no small part due to my parents; as a product of two engineers, I've grown up with an engineering mindset instilled in me. “Navigation” is easy – throughout my life my family has helped me navigate many a tough decision and, looking back, I thank them for steering me in directions I may not have gone on my own. And, finally, in terms of “Timing”, they have provided

a (much needed if you ask them) ever-present clock throughout my Ph.D. with the periodic, gentle reminders of just how long I've been at Stanford.

Finally, this work would not have been possible without the financial support provided by the Federal Aviation Administration (FAA) and the Stanford Center for Position Navigation and Time (SCPNT).

Contents

Abstract.....	v
Acknowledgments.....	vii
Chapter 1 Introduction	1
1.1 Motivation.....	2
1.2 How to Locate a Signal.....	6
1.2.1 Received Signal Strength	6
1.2.2 Time Difference of Arrival.....	7
1.2.3 Direction of Arrival	7
1.3 Previous Work	8
1.4 Contributions	9
1.5 Organization.....	10
Chapter 2 Preliminaries.....	13
2.1 Frequency of Operation	13
2.2 Jammers	14
2.3 Flight Test Platform	15
2.4 Flight Test Locations	17
2.5 Antenna Terminology	20
Chapter 3 Signal DoA Sensor for a UAV.....	23
3.1 Chasing the Perfect Sensor	23
3.2 Phased Array Theory	26

3.2.1	Array Geometry.....	27
3.2.2	Weighting	30
3.2.3	Element Response	32
3.3	The Antenna.....	34
3.3.1	Design.....	34
3.3.2	Schematic / Physical Implementation	37
3.3.3	Calibration.....	39
3.3.4	Command and Control	41
3.4	Performance	43
3.4.1	Steering Performance	44
3.4.2	Gain Pattern.....	45
3.5	Conclusion	46
Chapter 4	Algorithm for Determining DoA.....	49
4.1	Related Work and Motivation.....	49
4.1.1	Maximum Value.....	50
4.1.2	Cross-Correlation	51
4.1.3	Other Methods.....	53
4.2	Real-World Effects	53
4.3	Lobe Detection Based DoA	55
4.3.1	Algorithm	55
4.3.2	Integrity Checks	58
4.4	Performance	60
4.4.1	Physical Rotation.....	60
4.4.2	3-element Array.....	64
4.4.3	Multiple Sources	67
4.5	Conclusion	80
Chapter 5	Sensing Platform Development.....	83
5.1	JAGER System Overview.....	83
5.2	Navigating in a GNSS Denied Environment	85
5.2.1	Related Work and Motivation	86
5.2.2	IR Monocular Visual Odometry.....	90
5.3	Localizing the Signal Source	95

5.3.1	An Introduction to Bearing Only Localization.....	96
5.3.2	Estimator	97
5.3.3	Path Planner.....	100
5.4	Simulated System Performance	103
5.5	Conclusion	108
Chapter 6	Live Demonstration.....	111
6.1	Test Setup	111
6.1.1	The JAGER System.....	112
6.1.2	The Environment.....	114
6.1.3	The Time of Day	115
6.2	Example Missions.....	116
6.2.1	Mission 1	117
6.2.2	Mission 2	121
6.2.3	Wrong Initialization	124
6.3	Flight Test Overview	125
6.4	Conclusion	127
Chapter 7	Conclusion.....	128
7.1	Summary and Contributions	128
7.2	Directions for Future Work.....	130

List of Tables

Table 2.1: Overview of conditions for each of the flight testing locations used.....	17
Table 4.1: Results of performance of each DoA determination method on the physical rotation dataset showing both the standard deviation of the error, σ_{error} , and the probability of an outlier measurement, <i>Poutlier</i>	61
Table 4.2: Results of performance of each DoA determination method after removing rotations from nearly directly overhead of the signal source.	62
Table 4.3: Results of performance of each DoA determination method on the 3-element phased array RSS patterns.	65
Table 4.4: Summary of key results from 2-source simulations for each array size	81
Table 6.1: Overview of the mission performance for flights in each of the starting areas.	117
Table 6.2: Overview of the performance of signal localization for all successful localizations..	126
Table 6.3: Overview of the performance of the IR navigation system during all missions.....	126

List of Figures

Figure 1.1: Approximate area of GNSS outage on approach to Hayward Executive Airport (HWD).	5
Figure 1.2: 50nm radius circle representing the FAA area of interest for possible location of the interference source at Hayward Executive Airport (HWD).	5
Figure 2.1: Signal source used for flight testing on a tripod equipped with surveying equipment to localize the location of the signal source.....	15
Figure 2.2: Side view of DJI S1000 equipped with a directional antenna (bottom mounted) and companion computer (top-mounted).	16
Figure 2.3: View of Stanford University flight testing location.	18
Figure 2.4: DJI S1000 in (a) the hill environment at Camp Roberts and (b) the semi-urban environment at Camp Roberts.	18
Figure 2.5: DJI S1000 at White Sands Missile Range at night with visual aids for pilot control..	19
Figure 2.6: DJI S1000 at Edwards AFB at sunset before taking off for flight testing	20
Figure 2.7: Example gain pattern from antenna manufacturer.	21
Figure 2.8: Directional antenna for the gain pattern.	21
Figure 3.1: Nominal RSS pattern for the directional antenna.....	24
Figure 3.2: Example element configurations for antenna arrays: (left) linear array, (center) rectangular grid planar array, and (right) hexagonal grid planar array.....	28
Figure 3.3: Example element layout for a 3-element array on a rectangular grid.....	28
Figure 3.4: Array response for 3-element array on a rectangular grid with (a) isometric view and (b) head-on view.....	29

Figure 3.5: Array factor for variable inter-element spacing for a uniform linear array with (a) half-wavelength spacing, (b) quarter-wavelength spacing, and (c) one wavelength spacing.	30
Figure 3.6: Azimuth angle for the possible steering direction of a linear array.....	31
Figure 3.7: Spherical coordinate system demonstrating the polar (θ) and azimuth (ϕ) angles that a planar array in the XY plane can be steered.	32
Figure 3.8: Patch antenna element response (yellow) and array factor (blue) for a steering command to 80° for a 7-element hexagonal grid array.	33
Figure 3.9: Response of antenna array steered to 80° with a patch antenna with a 60° beamwidth demonstrating the inability to steer to the desired angle due to the single element response.	33
Figure 3.10: View of a 19- (blue), 7- (orange), and 3- (green) element array based on a hexagonal grid.	35
Figure 3.11: Layout of elements in 3-element phased array antenna with element numbering.....	35
Figure 3.12: (a) isometric and (b) head-on view of the array response for the 3-element antenna array with a hexagonal grid.	36
Figure 3.13: Performance of the steering capabilities of the antenna with (a) an analysis of the ability to set ϕ and (b) the ability to set the θ based on measurement from a vector network analyzer.	37
Figure 3.14: Schematic of the 3-element phased array antenna.....	38
Figure 3.15: Detailed schematic of RF components for each antenna element within the 3-element phased array antenna.	38
Figure 3.16: Example nonlinear effects of a board-level bias in the phase on a single element, showing the error between the desired steering angle (a) ϕ and (b) θ and the resulting steering angle with no bias (blue) and with the 10° bias present (orange).....	40
Figure 3.17: Output of power detector with varying window sizes for averaging signal strength samples to handle white noise induced by the RF components on the phased array antenna PCB.	41
Figure 3.18: Diagram of the connectivity within the phased array antenna demonstrating the elements controlled by the microprocessor.	42
Figure 3.19: Example RSS patterns generated with a pulse jammer with a dwell time of (a) 0.5ms and (b) 1ms.....	43

Figure 3.20: Performance of steerability as a function of the desired θ angle, the polar angle, in (a) the resulting error in ϕ , the steered azimuth angle, and (b) the resulting error in θ , the steered polar angle.	45
Figure 3.21: Example of the three possible resulting RSS patterns from the 3-element phased array antenna with (a) a single unpronounced lobe, (b) a single strong lobe and two large side lobes, and (c) one main lobe and strong back lobe	46
Figure 4.1: Example RSS pattern and true DoA (red dashed line) with large side lobes resulting in an incorrect DoA by the maximum value method (orange line).	51
Figure 4.2: RSS patterns generated with a physically rotated directional antenna at (a) close to 100m from source and at (b) the edge of the effective range of the directional antenna showing the true DoA (red dashed line) and cross-correlation computed DoA (orange line).	52
Figure 4.3: Nominal RSS pattern for the directional antenna.	53
Figure 4.4: Example challenging RSS patterns from real-world environments with corresponding true DoA as the red dashed line: (a) at the edge of the range of the antenna, (b) with large side lobes, and (c) corrupted by environment noise (such as from reflections).	54
Figure 4.5: Example detailed view of the LDB-DoA algorithm depicting the maximum value point in the RSS pattern (blue x), the crossing points 3dB below the maximum value (orange x), and the resulting detected lobe (yellow pie slice) in addition to the true DoA (red dashed line). .	56
Figure 4.6: Result of the LDB-DoA algorithm on RSS patterns from challenging environments with (a) a pattern at long distance resulting in two detected lobes and therefore two possible DoA values, (b) a pattern with large sidelobes resulting in three detected lobes and three possible DoA values, and (c) a pattern corrupted by noise resulting in six detected lobes and six possible DoA values.	57
Figure 4.7: Example noise corrupted RSS pattern with multiple lobes detected in an impossible configuration for the known antenna characteristics.	59
Figure 4.8: Histogram of the distance from the source for the dataset of RSS patterns generated by physical rotation.	61
Figure 4.9: Effectiveness of using the cross-correlation coefficient as an integrity check of the measured DoA.	63
Figure 4.10: Effectiveness of using the custom integrity check based on a model of the antenna.	63
Figure 4.11: Histogram of the distance from the source for the dataset of RSS patterns generated by the 3-element phased array antenna.	64

Figure 4.12: Distribution of DoA values measured using the cross-correlation (blue), the maximum value (red), and the LDB-DoA (yellow) method on the RSS patterns generated by the 3-element phased array antenna.	66
Figure 4.13: Example RSS pattern with large back lobe causing the cross-correlation (blue) and the maximum value (red) estimates to point in the wrong direction (true DoA is the dashed line) while the LDB-DoA algorithm presents two possible options.	66
Figure 4.14: Theoretical gain pattern of unsteered (a) 3-, (b) 7-, and (c) 19-element arrays.	68
Figure 4.15: Simulated RSS patterns for the 3-element phased array antenna with sources at (a) $\phi_{s1}, \phi_{s2} = (0^\circ, 90^\circ)$ and (b) $\phi_{s1}, \phi_{s2} = (90^\circ, 270^\circ)$ demonstrating the distortion effect due to large side lobes.	69
Figure 4.16: Simulated RSS patterns for two sources at $\phi_{s1}, \phi_{s2} = (0^\circ, 90^\circ)$ for array sizes of (a) 7-elements and (b) 19-elements.	70
Figure 4.17: RSS patterns for (a) 3-element, (b) 7-element, and (c) 19-element arrays with sources at $(\phi_1, \phi_2) = (0^\circ, 40^\circ)$	71
Figure 4.18: Bearing simulation results for a 3-element antenna with 2 interference sources at (0, -400m) and (0, 400m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source.	73
Figure 4.19: Example simulated RSS pattern with two sources such that the power difference between the two results in only one signal being detected by the LDB-DoA algorithm.	74
Figure 4.20: Bearing simulation results for a 7-element antenna with 2 interference sources at (0, -400m) and (0, 400m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source.	75
Figure 4.21: Visibility of each signal source using a 19-element phased array antenna.	76
Figure 4.22: Bearing simulation results for a 3-element antenna with 2 interference sources at (0, -200m) and (0, 200m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source.	77
Figure 4.23: Example simulated RSS pattern demonstrating two sources with an angular separation of less than $2BW$ resulting on the LDB-DoA algorithm to only detect a single lobe.	78
Figure 4.24: Bearing simulation results for a 7-element antenna with 2 interference sources at (0, -200m) and (0, 200m) showing: (a) the ability to observe a bearing to each source, (b) the	

number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source.....	79
Figure 4.25: Visibility of two sources with the 19-element phased array antenna.	80
Figure 5.1: Schematic of full JAGER system.	84
Figure 5.2: Images captured from the same location at Edwards Air Force Base at 11 pm with (a) a regular camera and (b) an IR camera.....	87
Figure 5.3: IR camera pictures from (a) 11 pm and (b) 9 am showing the difference in color intensity based on time of day and heat of objects.....	88
Figure 5.4: Pipeline for IR images to optical flow to the final state estimate of the EKF onboard the JAGER system.	91
Figure 5.5: IR image captured at Stanford University (a) without the detected features and (b) with the detected features.	92
Figure 5.6: Flight test results of IR optical flow with (a) the flight path flown and (b) the true and optical flow velocities throughout the flight for a commanded speed of 5m/s during each straight segment with pauses at each vertex of the flight path.	93
Figure 5.7: Performance of optical flow based navigation system with (a) GNSS (blue) and estimated (red) flight path and (b) 2D position error between the estimated and GNSS position throughout the flight.....	93
Figure 5.8: Performance of the optical flow system showing the GPS (“true”), measured, and EKF estimated velocity in the (a) North and (b) East directions during a night flight at Edwards Air Force Base using a 9Hz FLIR Vue Pro camera.....	94
Figure 5.9: Performance of flight at night at Edwards Air Force Base using a 9Hz FLIR Vue Pro camera with (a) the GPS and the estimated flight path and (b) the estimated position error as a function of the distance traveled.....	95
Figure 5.10: Depiction of the initialization step of the Gaussian Sum Filter given a single DoA measurement.....	98
Figure 5.11: Process flow for pruning and updating the bank of GSF used to handle the tracks for clutter management.	100
Figure 5.12: Example subset of the family of curved trajectories possible based on changing the approach angle to the signal source 1km from the takeoff position with 90° (dot), 80° (dash), and 60° (line) approach angles.	101
Figure 5.13: Estimated signal source position error for three example curved approaches: 90° (dot), 80° (dash), and 60° (line).	102

Figure 5.14: Illustration of the 3 phases of flight throughout a mission managed by the localization system.....	103
Figure 5.15: Error in the estimated interferer's position in simulation flying the three different trajectories (circle [dot], 80° [dash] and 60° [line]) with and without GPS for an interferer 1km away.....	104
Figure 5.16: Error in interferer's estimated position in simulated flights of all 3 trajectories without GPS compared to standoff strategy at a distance of 500m	105
Figure 5.17: Error in interferer's estimated position in simulated flights of all 3 trajectories without GPS compared to standoff strategy at a distance of 1500m.	105
Figure 5.18: Estimated vehicle position error in simulated flights with varying velocity measurement performance for a 1km circle flight while flying at 5m/s.....	107
Figure 5.19: Error in interferer's estimated position in a simulated 60° inward flight trajectory at a distance of 1km with varying velocity measurement noise.....	107
Figure 6.1: Side view of DJI S1000 equipped with navigation and localization systems; the circled items, from left to right, are: the 3-element phased array antenna, 1D LIDAR, gimbal for camera stabilization, FLIR Boson IR camera, and the Pixhawk autopilot.	113
Figure 6.2: Screenshot of accompanying ground station showing the position of JAGER and the current estimate (green oval) for the source location.	113
Figure 6.3: Overview of flight testing area for demonstration at Edwards Air Force Base. In red is the full flight test area and the yellow and orange regions highlight regions of unique levels of interference as a result of the presence of the large hangar (labeled Hangar).	114
Figure 6.4: Example RSS pattern from different parts of the flight test environment with: (a) from the orange zone where the reflection causes a distortion in the pattern, (b) from a non-interfered part of the flight test area, and (c) from the yellow zone where the reflections from the hangar caused destructive noise.	115
Figure 6.5: Example IR images over a similar region in the flight area (a) around 9 am and (b) around 11 pm.....	116
Figure 6.6: Frames of the Mission 1 flight showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) at the initialization of the GSF for the signal source occurred, (b) when the filter internally determined the source to be localized, and (c) at the final state of the estimate before returning to manual control.	118

Figure 6.7: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight.	119
Figure 6.8: 2D position error of the estimated position of the vehicle throughout the flight.....	120
Figure 6.9: Overview of the velocity performance throughout the flight with the Optical Flow (OF) (blue x), EKF estimate (yellow), and GPS (red) velocities for the North and East directions.	120
Figure 6.10: Frames of the Mission 2 flight with no integrity checks showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) the initialization of the GSF for the signal source occurred, (b) begins to convert towards hangar door due to interference region, (c) the vehicle leaves the regions of interference and the estimate moves away from the hangar door, and (d) the final state of the estimate.	121
Figure 6.11: Frames of Mission 2 with integrity checks showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) the estimate covariance remaining large due to the discarding of noisy measurements, (b) the estimate remaining away from the hangar as the vehicle leaves the regions of interference, and (c) the final state of the estimate.	122
Figure 6.12: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight with no integrity checks.	123
Figure 6.13: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight with integrity checking.....	124
Figure 6.14: Two example flights showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself where the initialization of the estimate failed with (a) the estimate immediately converging close to the initial position due to noisy RSS patterns and with (b) the estimate converging on DoA measurements from the non-main lobe (e.g. back lobe).....	125

Chapter 1

Introduction

Global Navigation Satellite Services (GNSS), or more colloquially known by the name of the United States system – the Global Positioning System (GPS) – is an invisible utility relied upon by the public and many different industries. Over the past 40 years since the first GNSS system went online, GNSS has become a critical component of industries such as transportation, finance, telecommunications, and the power grid. Beyond simply guiding smartphone users to their final destination, GNSS enables numerous core functionalities that drive modern society, from facilitating vital safety for aviation and maritime navigation to providing the critical timestamps on all financial transactions [1]. In fact, GNSS has become so critical that the US Department of Homeland Security has classified GPS as critical infrastructure for 13 of the 16 sectors of the US critical infrastructure [2].

All GNSS systems, at their core, rely on the ability to receive a fairly weak signal from over 20,000 km away in space, making it inherently fragile to radio frequency interference (RFI) [3]. As a result, there has been a constant battle to continue to protect, toughen, and augment these signals to ensure integrity for its users since the first deployment of GPS in 1995. For example, these efforts include augmenting GNSS systems with other sensors on the user's device to detect and mitigate the effects of interference; fortifying the receiver in terms of algorithms and hardware itself to further reduce the effect an interfering signal has on the GNSS system; and protecting the frequency spectrum used by GNSS systems through regulation to safeguard the weak signal from encroachment or interference from other users [4]. However, while all these efforts have greatly improved GNSS receiver performance towards interference, especially unintentional interference

from the ever-increasing number of radio frequency (RF) based technologies, there is still always the threat of an interferer that is too strong to be handled via improved receiver technology alone.

Another mechanism to protect GNSS based operations is, therefore, necessary to handle such cases. This means of protection must be twofold: mitigating the effect of the operations through backup systems and reducing the time that an interferer is interfering by rapidly localizing the source. For example, in aviation, GNSS RFI has been known to be a possible source of problems for aviation since the early days of the Federal Aviation Administration's (FAA) incorporation of GNSS into systems ranging from terminal operations to en-route operations [5]. Early efforts from the FAA outlined a strategy to protect the spectrum around GNSS, increase the robustness of receivers to sources of interferences, and develop a set of systems to localize RFI in real-time at major aviation hubs [5]. While the FAA recognized that real-time RFI localization is a necessary facet in mitigating impairment to aviation systems induced by GNSS interference, a rapid and reliable system capable of this desired RFI localization in real-time is still not employed today. This thesis provides a novel solution to this ongoing problem by presenting the design, development, and demonstration of an autonomous system capable of rapidly localizing a GNSS interference source via the rapid deployment of the autonomous system, as well as rapid localization once the autonomous system is deployed.

1.1 Motivation

Despite the great advances in protecting and augmenting GNSS signals, RFI will always need to be contended with, especially since it is not always intentional; it can be accidental interference or uninformed interference [6]. Accidental interference is simply RFI in the GNSS bands which typically results from malfunctioning equipment. For example, television (TV) towers such as one in Artarmon, Australia or unlicensed TV stations in Northern Italy run by the Mafia induced unintentional jamming [7], [8]. On a much smaller scale, certain USB3 devices are known to interfere with GNSS systems [9]. Uninformed interference is the result of users intentionally interfering with GNSS who's results caused more harm than intended [6]. As the use of GNSS has grown over the years to applications such as asset tracking, the category of uninformed interference has grown considerably [10]. For example, a Personal Privacy Device (PPD), used to interfere with GNSS within the scope of a vehicle, in a truck out front of the London Stock Exchange was found to have impacted the ability to receive GNSS signals at the Exchange [11]. In Germany, PPDs have been detected along almost all major highways in the country [12]. In the US, one of the most notable incidents of uninformed interference was a PPD traveling on the New Jersey Turnpike near

Newark Airport found to impact airport systems [13]. The extreme case, intentional interference, has typically been seen from state-level actors such as along the demilitarized zone between North and South Korea [11], in the Arctic Sea near Russian ports [14], [15], and in the Mediterranean [16].

The response to these incidents has varied widely over the years both in the method of responding and in response time. An example of an unintentional, accidental interference and associated response occurred in Monterrey bay in 2001. The Monterey Bay Aquarium Research Institute (MBARI) operated a differential GNSS service that provided correction information to ships in Monterrey bay to enable their determination of a more accurate position. In April 2001, the service was interrupted due to interference disabling the base station GNSS receiver resulting in degraded navigation performance [17]. After several weeks of manually searching the neighboring marina with a directional antenna and a spectrum analyzer, dedicated MBARI researchers were able to track down the source of the interference – the culprit was an active TV tuner on board one of the boats that, when powered on, “leaked” radiation in the GPS spectrum, effectively jamming the signal from MBARI [17]. Eventually, it was determined that there was not just a single culprit, and in fact, there were three different tuners onboard three different ships in the marina causing interference. Most notably, one of the more difficult sources to find was a tuner that would only interfere during the cooler temperatures at night, requiring the search party to work through the night to localize the source [17]. This example motivates the desire for an automated system that can more effectively localize GNSS interference, day or night. This example further motivates the desire for automatically determining a much smaller manual search space to enable more rapid disablement of an interference source by personnel.

Another high-profile case of unintentional intentional jamming was in 2007 in San Diego harbor when a United States (US) Navy ship left its GNSS jammer powered on after returning from exercises off the California coast [18]. It took nearly 2 hours for Navy personnel to identify that they themselves were the source of the interference, and to disable it. It then took a full 72 hours for the United States Coast Guard Center of Excellence for Navigation and supporting agencies to localize the source as there was no communication with the Navy personnel. In the 2 hours that the jammer was active, countless systems failed. This included crucial systems whose failure induced a heavy impact, such as telephone switches, cellular phone operations, and even the paging system at the local hospital [1]. This incident further demonstrates how critical GNSS is to a wide range of infrastructure, even 2 hours of GNSS interference can be long enough to induce deleterious impact to pivotal systems. This example further motivates the need for a system that is capable of

rapidly localizing GNSS interference to mitigate the impact of interference on critical infrastructure.

In the UK, the risk of interference to infrastructure is enough, identified as \$1B USD per day, that major efforts have started to put in place dedicated backup systems in case of GNSS interference [19], [20]. For example, the General Lighthouse Authorities (GLA) of the United Kingdom (UK) has invested in setting up an e-Loran system as a backup due to the high value of traffic in the Channel [21]. Onboard a particular ship during GNSS interference testing, not only did core navigation systems fail but so did safety systems such as the Automatic Identification System [21]. On land, Differential GPS (DGPS) and signal lights were also susceptible to interference, potentially creating pandemonium for nearby ships, especially in more crowded areas such as the Channel [21].

These jammer localization deficiencies of these historical examples remain unresolved in today's policies for responding to GNSS interference. As recent as early 2019, the Federal Aviation Administration (FAA) demonstrated the modern response to an interference incident at Hayward Executive Airport, located just across the bay from Stanford University. After several reports of interference on approach to the airport, in the region depicted in Figure 1.1, the FAA's response was to issue a Notice to Airmen (NOTAM) that indefinitely "closed" the GPS based approach to the Hayward Executive Airport until such a time that the interference source has been localized. Furthermore, the FAA issued a region of interference that spanned a radial distance of 50nm (nautical miles) from Hayward Airport as the search area for the interference source, depicted in Figure 1.2. This search area is so large that it encompasses three major international airports (where none of which had reported interference) and a range of quite difficult terrain for interference to travel through. The FAA's outdated policy for response for such interference incidents further motivates the need to rapidly localize the source of GNSS interference, and the ability to do so while operating in that same GNSS denied environment. This example also exemplifies the desire for such a system to be easily deployable at an airport.

The core mission targeted by the system discussed in this thesis aims to resolve the deficiencies of modern responses to GNSS outages as illustrated in the previous examples. To address the motivations for improvement prompted by these examples, the system should be capable of localizing the source of GNSS interference in a time scale of minutes and hours instead of the days and months required to localize these sources today. The autonomous system should further allow for rapid disablement of the interference source by providing sufficient localization accuracy such that a final search space can be determined by the autonomous system that is small enough to enable efficient, manual localization of the exact position of the source. There is a further desire for this

autonomous system to undergo 24-hour deployment and to be capable of localizing the source of interference in both daytime and nighttime conditions. Furthermore, as GNSS cannot be relied upon when such an interference source is present, this autonomous system must be capable of localizing the inference source in a GNSS denied environment. Finally, the autonomous system should be capable of being instituted and deployed in an airport environment, as GNSS interference at an airport can have a significant impact on both the aviation industry as well as civilian safety.



Figure 1.1: Approximate area of GNSS outage on approach to Hayward Executive Airport (HWD).

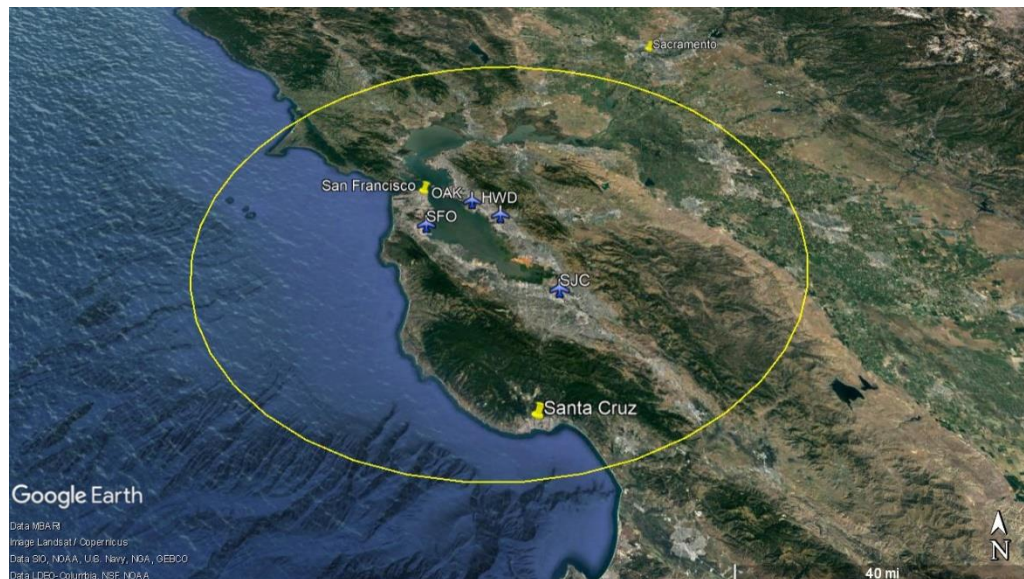


Figure 1.2: 50nm radius circle representing the FAA area of interest for possible location of the interference source at Hayward Executive Airport (HWD).

1.2 How to Locate a Signal

When it comes to the task of signal localization, over the years, there has not been a single method that serves as the be-all, end-all of localization methods. More importantly, signal localization is historically an area of active interest and research with methods continually evolving as technologies and capabilities evolve. While it's impossible to enumerate the countless methods of tackling the general challenge of signal localization today, they can be broadly broken down into three different categories: Received Signal Strength (RSS) based methods, Time Difference Of Arrival (TDOA) based methods, and Direction of Arrival (DoA) based methods.

1.2.1 Received Signal Strength

Fundamentally, this technique relies on the assumption that, for a source of known power, given the RSS value of a signal, the range to the signal can be estimated. If the power of the source is not known, the difference of the signal strength as the agent moves can be utilized to handle this case [22]. However, this measurement is only indicative of the range to the signal source: to localize the source, a localization algorithm is needed along with a strategy to move through the environment to collect the measurements needed.

Mathematically, the RSS value as a function of distance theoretically follows a $\frac{1}{r^2}$ law, enabling the range to be determined from the RSS value [3]. However, environmental effects such as reflections of the signal reaching the receiver or attenuation through walls and vegetation greatly affect the measured RSS value, thus disrupting the reliance on RSS to calculate range [23], [24].

Early works in WiFi localization identified the difficulties in estimating range directly from the RSS measurements. These works explored the use of more complex mapping-based approaches, creating detailed RSS maps of the environment. These maps were then utilized to localize the source from those maps [25]–[28]. While effective, these methods require either prior knowledge of the transmit power of the signal source or extensive maneuvering and mapping of the environment to generate an estimate. Instead of using a single sensor, [29] identified the ability to use a high density of multiple sensors to more effectively localize a signal source with RSS measurements. RSS based methods, therefore, perform admirably in areas where crowdsourcing is possible, for example, via cell phones [29]. However, as such crowdsourcing is not always possible or quickly exploitable, reliance on RSS based methods in this manner does not enable an easily deployable solution, especially in an environment such as the airport. Furthermore, even with a high number of sensors, the general performance of a signal strength based method tends to only be on the order of 100s to 1000s of meters [22].

1.2.2 Time Difference of Arrival

This next method relies on using multiple highly-calibrated and time-synchronized receivers to measure the TDOA of a signal of interest between each of the receivers. From this time difference of arrival between the pairs of receivers (or between all the receivers), hyperbolic lines of position can be calculated from each receiver pair, resulting in the localization of the signal source [22].

TDOA based methods, therefore, require a collection of calibrated and time-synchronized receivers. While the use of a TDOA system can produce superior results, establishing the required infrastructure of these calibrated and time-synchronized receivers is complex and can induce significant costs [30]. Furthermore, this approach is once again not the most suitable for deployment at an airport, especially given the difficulties of instituting an entirely new infrastructure at an airport. It is worth noting that, while TDOA methods are by far the most accurate of the three categories discussed here with typical performance in the 10s of meters, TDOA methods are also by far the most complex [30].

1.2.3 Direction of Arrival

Direction of Arrival (DoA) based localization relies on measuring the relative bearing to, or the direction of arrival of, the signal source, and using a collection of those measurements to localize the source. Similar to the RSS methods, DoA methods require not only the DoA measurement itself but further requires both an estimation scheme to determine the location of the signal source given a series of measurements, as well as a path planning strategy to maneuver through the environment to collect measurements. Because the range to the signal source is unknown with direction of arrival measurements, one limitation of this method is the requirement of either multiple sensors that are angularly positioned around the signal source or maneuvering of the receiver around the signal source to collect measurements from multiple, angularly different locations [31]. However, this limitation can be viewed as a benefit of this method as it provides flexibility in the implementation of a DoA based approach to signal localization.

Because direction of arrival is fundamentally not a simple measurement, this is a very interesting category as there are many different techniques for determining the direction of arrival of a signal. Methods range from physically rotating a directional antenna to find the direction of peak power to using complex antenna arrays to compute the DoA from the properties of the incoming waveform from the antenna elements within the array [32]–[36]. With each different approach to determining the DoA of the signal comes a different strategy for localization, which includes usage of different estimators and different path planning strategies. Strategies range from

simply aiming to “home in on” the signal source to more greedy, information-theoretic trajectory planning methods [33], [34], [37], [38].

On the whole, due to the range of different techniques to generate the DoA measurement, it is difficult to make a sweeping statement on the accuracy as it can range from 10s of meters to 1000s of meters based on the specific implementation [22].

1.3 Previous Work

The FAA’s first version of a system for localizing GNSS interference was comprised of fixed ground elements, mobile ground elements, and an airborne element, with the general strategy to get coarse localization from the airborne element and to refine the localization using a combination of the fixed and mobile ground elements [5]. Their airborne element, known as the Aircraft RFI Localization and Avoidance System (ARLAS), consisted of a small, manned aircraft with a patch GPS antenna on the roof. When the aircraft rolled, the hemispherical gain pattern atop the aircraft would “illuminate” the ground. If RFI was detected during the roll, it would imply the source was somewhere within the illuminated region. By performing a succession of rolling maneuvers, the interference source could be localized [39]. While capable of localizing an interference source, a solution like ARLAS is expensive because it requires a human pilot and time to mobilize and deploy. ARLAS also suffers from the coupling of the vehicle and sensor dynamics. In order to detect RFI, the aircraft had to roll, which alters the flight trajectory. This coupling can result in compromises between trajectory and sensing, leading to longer search times.

While ARLAS is the main endeavor by the FAA to localize GNSS interference, the WiFi and wildlife tag tracking communities have both been involved in developing reliable and mobile methods of signal localization. Exploring the state of the art of Unmanned Aerial Vehicle (UAV) based methods began with the work by Gabe Hoffman at Stanford University in 2008 to develop a UAV capable of localizing an avalanche beacon [38]. While successful, Gabe’s work was limited in scope in the size of the world (a 9m x 9m area). Since then, there have been DoA based methods of localization starting with ground robots physically rotating a directional antenna and quickly moving to airborne platforms [32]–[34]. Airborne systems ranged from highly customized vehicles designed to constantly rotate to more off the shelf platforms employing a stop and rotate strategy to make DoA measurements [33], [34]. The wildlife radio tag community’s early efforts embraced the stop and rotate style modality, employing it on several different systems in New Zealand and the United States to demonstrate tracking and localization of different avian species [37], [40]–[43]. However, all these methods are hampered by the speed, or really lack thereof, when it comes

to the stop and rotate modality taking 30 seconds [41] or even 45 seconds [37] to execute a rotation. Additionally, most of these methods have only been demonstrated to work within a small world, some even constrained within an indoor lab environment, which is certainly not the most representative of real-world environments. As a result, the research has gone in two directions: a first option was to simplify the sensor and rely on more complex path planning and estimation to leverage the vehicle's maneuverability to localize a source [44], and the second option was to add a little complexity to the sensor and improve each of the elements in the localization chain. The work presented in this thesis focuses on the second option.

1.4 Contributions

This thesis presents the design, construction, and demonstration of an autonomous system capable of localizing an interference source in a time scale of minutes and hours instead of days and months it takes today. To achieve this end goal, an autonomous UAV was built comprised of a signal localization system and a GNSS denied navigation system to support 24/7 operations.

Lightweight Sensing

The first contribution of this thesis is directed at improvements of the core element in the signal localization pipeline – the sensor used. A UAV based system that can meet the needs for rapid deployment and autonomous operations requires a sensor that is lightweight (i.e. portable) and agnostic to the type of vehicle being used to increase the versatility of the sensor. Additionally, to support the rapid localization requirements, the sensor must be computationally lightweight and must operate quickly. To address these needs, the first contribution presented is the creation of a computationally and physically lightweight, electronically steered 3-element antenna array that provides high-rate RSS patterns onboard a Commercial Off-The-Shelf (COTS) UAV.

Algorithm for Direction of Arrival Determination and Localization

The second step in the chain of localization when using DoA measurements is the ability to robustly determine the DoA from an RSS pattern. The algorithm must be capable of performing on RSS patterns generated in real-world environments that contain the impact of reflected signals and other noise in the environment. The second contribution of this thesis presents the development of a robust algorithm for determining the DoA of a signal source from an RSS pattern: an algorithm termed Lobe Detection Based Direction of Arrival (LDB-DoA). LDB-DoA not only outperforms previous state-of-the-art methods in real-world environments but also provides a novel

parameterization of the underlying RSS pattern without the need for a calibration pattern. Additionally, LDB-DoA is agnostic to the number of signal sources in the environment, enabling its utilization for scenarios beyond the single interferer case where no other state-of-the-art method is capable of operating. These DoA measurements provide the necessary information to both the localization filter and the path planning strategy to successfully localize a signal source to within 20m within the flight time of a UAV.

Sensor Platform Development and Deployment

The third contribution of this thesis demonstrates the real-world capabilities of this GNSS RFI localization system. Real-world operations include operating in both real-world environments and in real-world times. As a result, a navigation system capable of 24/7 operation based on infrared (IR) imagery was developed and demonstrated in both day and night conditions. This contribution also presents several demonstrations of the fully-functioning autonomous system combining the efforts of the phased array antenna and the direction of arrival determination algorithm to successfully localize a simulated interference signal in a semi-urban airport-like environment in both day and night conditions. The flight test platform enables the collection of measurements in real-world environments ranging from environments subject to commercial jamming to environments that simulate airport conditions. Furthermore, the test platform, utilizing IR-based navigation, enables the full demonstration of a signal localization mission in a GNSS denied environment using the three-element phased array antenna created in this thesis as the sensor along with the DoA algorithm developed in this thesis.

1.5 Organization

This thesis is broken down as follows. Chapter 2 presents some preliminaries to familiarize the reader with some of the core assumptions at play throughout the thesis along with presenting some of the fundamental hardware used for all the data collection. From there, Chapter 3 explores the first contribution of this thesis: the development of a lightweight 3-element phased array antenna capable of generating high-rate RSS patterns onboard a small COTS UAV. Chapter 4 presents the second contribution of this thesis: the Lobe Detection Based DoA (LDB-DoA) algorithm that is robust to real-world environmental noises and that outperforms state of the art DoA algorithms given an RSS pattern. Chapter 4 also presents the first of the real-world flight test data in demonstrating the full extent of the capabilities of the LDB-DoA algorithm developed as part of this thesis. The final contribution of this thesis is broken down into Chapter 5 and Chapter 6.

Chapter 5 presents the development of the system itself, including the IR based GNSS independent navigation system that enables 24/7 operations and the localization and path planning algorithms use onboard to localize an interference source. Additionally, simulation results are presented to demonstrate the theoretical capabilities of the system in a large-scale environment. The third contribution is completed in Chapter 6 with the presentation of a real-world, real-time demonstration of the full autonomous system operating without GNSS, leveraging the IR based navigation system to localize a signal source using the 3-element phased array sensor as the core sensor and the LDB-DoA algorithm as the algorithm to provide DoA measurements to the localization filter. Finally, Chapter 7 provides concluding remarks and parting words on this thesis and discusses a few directions for future research.

Chapter 2

Preliminaries

This chapter serves to present some necessary background and supporting information that is used throughout the rest of the thesis. Section 2.1 discusses fundamental information regarding the frequencies used throughout the testing as legal requirements limited the feasibility of using GNSS frequencies. Section 2.2 presents the types of interference for which the system is investigated to familiarize the reader with details on the typical interference threats that face GNSS today. Sections 2.3 and 2.4 address the flight testing aspects of this thesis by presenting an overview of the experimental flight test platform used for data collection and demonstration along with the environments in which those flights were performed. Finally, Section 2.5 provides foundational background in antenna terminology used throughout the thesis.

2.1 Frequency of Operation

Throughout the introduction, the emphasis was placed on the localization of GNSS interference sources, which would presume sources in the L-band, for example at 1.575GHz where the most widely used GPS L1 signal sits. However, the critical nature of GNSS as explained in the introduction presents a unique difficulty in demonstrating and testing systems for localizing GNSS interference as one cannot “willy-nilly” set up a GNSS jammer for testing. Thankfully, outside of the specific geometry required in an antenna array discussed further in Chapter 3, the exact frequency of the target signal to localize does not impact the strategy or algorithms in any way. As a matter of fact, most of the algorithms presented in this thesis are frequency agnostic. Therefore, to enable faster-paced development, truly allowing for real-world results to shape the strategy and

algorithm, the antenna developed, and signal sources used are tuned for the 2.45GHz spectrum – a slice of spectrum for which rules and regulations are significantly more favorable for testing purposes.

Throughout this thesis, nearly all results presented are results using systems tuned for 2.45GHz. However, there are a handful of results discussed from flight tests performed at White Sands Missile Range that used a real commercial GNSS that provided an understanding of the type of environment – in terms of the characteristics of the degradation and loss of GNSS – the system must be capable of operating in when it comes to truly localizing a GNSS interference source.

2.2 Jammers

The target frequency used throughout the testing may not be the same frequency as a true GNSS interference source, however, effort was made to ensure that the type of signal is still representative of a typical GNSS interference source. For that, the source used for testing is designed to mimic the characteristics of Personal Privacy Devices (PPDs) – arguably the most popular of the interference sources for GNSS systems. The large majority of PPDs employ a “chirp” signal structure, that is at any given moment in time the jammer is broadcasting a single tone, however, the frequency of that tone is swept through a frequency range, effectively jamming a larger bandwidth [45]. A typical PPD sweeps over a 10-20MHz wideband ($L_1 \pm 5 - 10MHz$) in a period of about 10 microseconds [45]. On the antenna side of things, a typical handheld PPD that can be purchased online comes equipped with a 2dBi omnidirectional dipole antenna. In terms of power, the most common power level found in PPDs online is 0.5 W per signal.

To recreate this signal structure, a Universal Software Radio Peripheral (USRP) B210 Software Defined Receiver (SDR) was used to generate a signal with a bandwidth of 10MHz centered on 2.45 GHz at a power level of 0.5W. A 2dBi omnidirectional dipole antenna is attached to the USRP B210 to complete the ensemble of mimicking a PPD. While this source is used for most of the results presented in this thesis, Edwards Air Force Base, the location of the final full system demonstration, did not allow the use of a custom-built transmitter as the signal source. Therefore, the final demonstration flights are performed with a commercially available video transmitter broadcasting a 1W signal at 2.45GHz with a bandwidth of 10MHz connected with a 2dBi omnidirectional dipole antenna. Bench testing of the antenna developed as part of the first contribution, discussed in more detail in Section 3.4, demonstrated that there was no significant difference between the antenna’s performance with the chirp style signal and the non-chirp style signal structure of the commercially available video transmitter.



Figure 2.1: Signal source used for flight testing on a tripod equipped with surveying equipment to localize the location of the signal source.

To finalize the setup of the signal source, the transmitter, be it the USRP B210 or the video transmitter, is mounted on a platform, such as the one depicted in Figure 2.1, alongside a GNSS receiver set to survey the location of the signal source. During flight testing, the signal source is kept static for periods of at least half an hour at a time which enables, at worst, meter-level precision on the location of the signal source.

2.3 Flight Test Platform

A critical component of this thesis is the ability to collect data outside of a laboratory environment. As a result, this means that a flight test platform is required as the foundational element to carry the sensors, equipment, and final system for the work presented in this thesis. While Chapter 6 provides a more detailed explanation of the implementation of the final system as part of the third contribution of this thesis, this section presents a brief overview to familiarize the reader for the platform used when discussing results based on flight test data.

While many parts of the overall system is designed to be deployable on a wide range of different types of UAVs, for all the testing a multirotor style UAV is used as it is capable of operating in multiple different modalities – namely an ability to stop and rotate in place or hover in place to collect measurements over a long period of time – that are desired for collecting the best data. The vehicle itself is a DJI S1000 Spreading Wings octocopter, shown in Figure 2.2, that weighs 4.2kg

empty and has a total takeoff weight of up to 11kg that provides capabilities for nearly 7kg of payload. Designed for carrying movie quality camera equipment, the DJI S1000 is more than capable of carrying large amounts of research equipment in both the weight capacity and the space on the vehicle itself. In the example configuration in Figure 2.2 a directional antenna is seen mounted below the airframe – the standard mounting location for all antenna throughout this thesis – and a computer is mounted on the top platform – again the standard mounting location for any computer. Fully loaded the vehicle’s average flight time, weather conditions dependent, is about 10 minutes, while the hover flight time is about 15 minutes. The vehicle is capable of forward flight of up to 20m/s of forward flight, however, most of the autonomous flights performed throughout this work limit the speed to 5m/s.




Figure 2.2: Side view of DJI S1000 equipped with a directional antenna (bottom mounted) and companion computer (top-mounted).

For the command and control of the vehicle, including providing autonomous capabilities, the vehicle is equipped with the Pixhawk autopilot, running the PX4 flight stack, a commonly used configuration for research applications [46]. Additionally, a companion computer is typically mounted alongside the autopilot to run the research systems within the Robot Operating System (ROS) framework [47]. While the Pixhawk autopilot is capable of fully autonomous operations alone, most flight testing occurred with the Pixhawk executing velocity commands generated by

the companion computer, allowing the systems onboard the companion computer to make the high-level decisions throughout flight testing. To create the GNSS denied environment desired for testing and demonstration of the system, the Pixhawk allows for a software disabling of the use of GNSS by the control system. The position is still logged throughout the flight to serve as a “truth”, however, none of the control systems onboard the vehicle have access to that position information.

2.4 Flight Test Locations

Flight tests were performed at a variety of different locations. For completeness, this section provides a brief summary of each of the different flight test environments for which data is presented in this thesis. A summary of the conditions for each of the flight test locations is presented in Table 2.1. The variety of environments in which test flights were performed has enabled an unprecedented amount of information for the performance of this kind of system in the real-world.

Table 2.1: Overview of conditions for each of the flight testing locations used.

	Area (km ²)	Environment	Time of Day
Stanford University	0.10	Open sky	day
Camp Roberts	0.25	Open sky + semi-urban	day
White Sands Missile Range	0.47	Open sky	night
Edwards AFB	0.15	Semi-urban	Day + night

Stanford University

One of the main flight test locations for which many RSS patterns are generated using both a physically rotated directional antenna and the 3-element phased array antenna is on Stanford University grounds. The test location is an area known as Coyote Hill that contains an open environment and a gentle rolling hill, depicted in Figure 2.3. All flight tests performed at this location was at 2.45GHz with the simulated chirp signal and with the video transmitter. Effectively this testing location was used to ensure that everything was working smoothly before doing larger tests at other locations. The IR camera was also tested at this location during daytime operations.





Figure 2.3: View of Stanford University flight testing location.

Camp Roberts

Camp Roberts was one of the earliest testing locations and was mainly used before flight testing at Stanford University was approved. It provided a much larger testing environment enabling longer range testing. It also provided two different environments: an open sky environment with rolling hills (Figure 2.4a), and a simulated urban environment populated with several multi-story buildings (Figure 2.4b). All the flights performed at Camp Roberts were with the physically rotated directional antenna tuned for the 2.45GHz spectrum. No work with either the 3-element beam-steering antenna or the IR camera was performed here.



(a)



(b)

Figure 2.4: DJI S1000 in (a) the hill environment at Camp Roberts and (b) the semi-urban environment at Camp Roberts.

White Sands Missile Range

White Sands Missile Range (WSMR) is the only flight test location where the tests involved a real GNSS jammer. For these tests, the vehicle was equipped with a highly directional GPS antenna (seen below the UAV in Figure 2.5) and the physical stop and rotate modality was used for all the flights. This test location enabled understanding the effects of an off-the-shelf PPD on a UAV and the impact it has on the signal localization mission. While all testing at WSMR was performed at night, the IR camera was not used due to regulatory requirements.



Figure 2.5: DJI S1000 at White Sands Missile Range at night with visual aids for pilot control.

Edwards Air Force Base

Edwards Air Force Base (AFB) is where all the full system tests took place. The environment presented a similar environment to what would be expected at an airport with mostly open space as seen in Figure 2.6. Due to legal requirements, the video transmitter was the source used for all the tests at this location. All the tests performed here included the use of both the IR camera and the 3-element phased array antenna. No tests were performed with the stop and rotate

implementation. In addition to Stanford University for validation testing, this is the only test location where the full flight stack was flown and tested to be able to demonstrate the capabilities of this system. A more detailed description of the environment and some of the more interesting features of the environment are discussed in Chapter 6 in the context of the demonstration of the system that is presented.



Figure 2.6: DJI S1000 at Edwards AFB at sunset before taking off for flight testing

2.5 Antenna Terminology

To assist in understanding many of the plots and some of the terminology used related to antennas themselves, this section briefly describes some of the more important terms and properties of antennas.

The gain pattern, or the antenna response, is the pattern of the antenna's spatial response. This is typically provided by the manufacturer of the antenna and is typically determined through analysis in an anechoic chamber where the response of the antenna can be directly measured. However, it is also possible to determine an approximation of the response, or gain pattern, in simulation environments. An example gain pattern is shown in Figure 2.7, where the “slice” of the gain pattern in both the vertical and horizontal direction are shown. This can be interpreted as a spherical conical shape that “emanates” from the tip of the directional antenna depicted in Figure 2.8




Figure 2.7: Example gain pattern from antenna manufacturer.



Figure 2.8: Directional antenna for the gain pattern.

When looking at a gain pattern, the pattern has a main “lobe” to it, which is a main beam element to the directional gain pattern. This lobe has a width, known as a beamwidth, that is defined by the width of the beam at the half power point in the lobe – the half power point being 3dB below the maximum value of the lobe.

Without being in an anechoic chamber with a signal source with known properties, it is very difficult to recreate the actual gain pattern of the antenna. Instead, it is possible to create the Received Signal Strength (RSS) pattern, that is, a spatial representation of the measured response from the antenna itself. These RSS patterns are the foundational element from which the work in this thesis determines the direction of arrival to the signal and therefore enables the localization of the signal source. As RSS patterns are based on RSS measurements a final key parameter that helps

define an RSS pattern is the pattern density – the number of measurements that make up the RSS pattern.

Pattern density is the number of RSS measurements per degree of the pattern – effectively the spatial density of the pattern. The higher the spatial density of a pattern, the more accurately the direction of arrival can be observed. The pattern rate is the rate at which RSS patterns can be generated by the sensor – this relates to the maximum rate that the system can provide bearing measurements needed for signal localization. The higher the pattern rate, the simpler the localization scheme, both the estimator and the path planner, can be. Beamwidth is the beamwidth of the main lobe of the antenna – it represents the directionality of the antenna. A smaller beamwidth results in a more accurate direction of arrival estimate. Sidelobe suppression is a measure of the difference in the RSS value between the main lobe and the side lobe of a pattern. Greater side lobe suppression – smaller sidelobes – reduces the amount of incorrect direction of arrival observations caused by sidelobes. These metrics are used to compare the performance of each of the methods to generate an RSS pattern. The beamwidth of an antenna is defined as the width, in degrees, between the two points 3dB below the maximum gain of the pattern.

Chapter 3

Signal DoA Sensor for a UAV

At the core of the system presented in this thesis is the spatial sensing of the signal strength of a broadcast signal to generate a Received Signal Strength (RSS) pattern. This chapter presents the first contribution of this thesis: the development of a 3-element phased array antenna capable of generating high rate RSS patterns. Section 3.1 presents related work in the domain of techniques for generating an RSS pattern and in the domain of the development of phased array techniques. To assist in the understanding of the development of the antenna, Section 3.2 provides a primer on phased array antenna theory. From that work, and from understanding the mission requirements of this system, Section 3.3 presents the specifics of the three-element phased array developed and presents the first contribution of this thesis. The next section provides an analysis of the performance of the antenna.

3.1 Chasing the Perfect Sensor

Before exploring the antenna array developed for the work in this thesis, a background of related work and other core sensing modalities that have been used for DoA based signal localization are provided. In particular, the focus of the work presented and discussed in this section relates to the sensor that is used as the building block to determine DoA. For background on actual DoA algorithms, explore the Related Work section of Chapter 4.

The first, and “simplest” approach to DoA includes physically rotating a directional antenna to create a spatial representation of the received signal strength. For example, physically rotating an antenna with the manufacturer provided gain pattern in Figure 2.7 and measuring the signal strength

throughout the rotation results in the RSS pattern in Figure 3.1. Notice that the patterns look very similar, and the fidelity to which the RSS pattern matches the gain pattern of the function is a function both of the density of signal strength measurements made spatially and the relative polar angle of the incident signal [32], [48].




Figure 3.1: Nominal RSS pattern for the directional antenna.

Physically rotating a directional antenna involves some level of mechanical complexity, but the computational simplicity makes it a very appealing approach to mobile platforms that may have limited compute power. For this reason, approaches that used physically rotating antennas were some of the first methods utilized for DoA determination on mobile robots. For example, on ground robots, such an antenna has been physically rotated with specialized servo motors [32], [49], [50]. However, these servo motors that enabled the antenna’s physical rotation have proven to be prohibitively costly in weight to be implemented on an airborne system. To further reduce the mechanical complexity and weight, [33] removed the dedicated servo for rotating the antenna and increase the mobility by developing a custom-built UAV that rotates as it flies to provide the rotation to the directional antenna onboard the UAV. A less custom solution was made in [34], where a standard Parrot drone’s control software was modified to continually rotate throughout the flight to provide the rotating platform. Stop and rotate approaches, where the UAV pauses its flight pattern to rotate, were later implemented by several works for wildlife tag tracking and, in parallel, in early efforts by the author of this thesis alongside Louis Dressel at Stanford University, to further simplify the complexity [37], [41], [51]. As another example of leveraging a vehicle’s dynamics in this regard, [52] uses a fixed-wing aircraft and uses the vehicle dynamics to move the sensor. To reduce the time to complete the necessary rotation, [49], [53] both employed a sensor system with

two and four sensors, respectively, to reduce the amount of rotation required by the system to enable a DoA measurement.

The mechanical nature of these rotations induces major limitations to methods relying on the physical rotation of the antenna. The first drawback is the mechanical speed of the rotation. In all of these implementations, the rotation speed is slow and is the main limiting factor for the overall performance of the system [54]. In particular, slow rotation speed causes a significant portion of the flight time of a small UAV to be spent rotating to make a measurement. The slow rotational speed introduces additional problems in detecting interference induced by a time-varying signal such as a signal that turns on and off periodically or varies its power over time. Such a signal may not be appropriately measured during a portion of a slow physical antenna rotation due to the signal power being low or off entirely. This can disrupt the generation of RSS patterns with high fidelity, as a resulting RSS pattern generated in this case may have missing signal strength measurements at certain angles, which can lead to determining an incorrect DoA.

A second drawback is caused by the coupling of the vehicle's motion to generate a measurement (the RSS pattern). Many other methods, such as those in [33], [34], [37], [52], require the UAV to maneuver in specialized ways in order to even collect a measurement. This necessitates that the localization mission must accommodate two different types of maneuvering: appropriate maneuvering for positioning of the UAV to collect the measurement, as well the additional required and specialized rotational maneuvering utilized by the UAV to actually capture the measurement. Rather than relying on the UAV's motion as a whole to achieve both positioning and measurement collection, these two requirements can be decoupled by utilizing the UAV as merely a platform that carries the sensor. This decoupling can enable greater accuracy to be independently achieved by the UAV in its positioning and by the sensor in its measurement collection, enabling the UAV to execute the best mission possible.

To avoid the cost of the stop and rotate modality, [55] uses a combination of a directional and an omnidirectional antenna, enabling an approximate DoA measurement to be made from a single measurement of signal strength. However, while this method significantly reduces the mechanical complexity and reduces the impact of time-varying effects, the resulting signal measurement can be inconclusive. In particular, the combined use of the directional and omnidirectional antenna can result in an ambiguity of the left/right side of the main lobe of the RSS pattern due to the symmetry in the gain pattern [55].

The next class of methods shifts from a mechanical approach to an electrical approach. Using an antenna array and altering the phase and amplitude of the different elements within the array – creating a phased array antenna – opens a vast collection of work in the field of determining the

DoA of an incident signal. The simplest method, known as conventional beamforming, is built on the same principle as the physically rotated antenna, with one major difference: the rotation is performed by adjusting the phase difference between the elements within the array rather than via a physical rotation. As a result, the limit of the mechanical speed is substituted for the limit of computational speed, which is typically orders of magnitude faster [54].

While the conventional beamformer overcomes the greatest limitation of the physical rotation-based method, researchers have continued to develop techniques using antenna arrays to develop a further class of techniques known as sub-space techniques. These techniques aim to adaptively change the phase and the amplitude of the incoming signals to optimally point the beam towards the signal while simultaneously working to reduce the overall noise in the system, resulting in very detailed peaks in the direction of the signal source. Using these methods, a 7-element phased array antenna was mounted onboard a container ship to localize interference sources near the Suez Canal [56]. It has further been demonstrated successfully on a vehicle-mounted solution [36]. However, this added performance comes at a steep cost of computational complexity. While a shipping container and minivan were suitable vehicles for integrating the hardware necessary to handle these computation requirements, the computational complexity renders these methods prohibitive for implementation via UAVs. [57].

As illustrated, none of the state-of-the-art solutions presented are suited to the rapid signal localization mission requirements put forth for this project, and a new solution is therefore required. Thus, the work discussed in this chapter introduces a novel sensor capable of overcoming the major shortcomings of existing approaches. In particular, a light-weight sensor that can be being carried on board a UAV and that is capable of generating RSS patterns at a variable rate to overcome the inherent weaknesses of a rotation-based modality to time-varying effects is presented. As the core element to a localization pipeline, ensuring a high enough rate of information for use by the localization algorithm further enables the use of tried and true, real-world tested, localization algorithms, which overcomes the potential weakness of a bespoke localization algorithm that has never seen real-world experience.

3.2 Phased Array Theory

This section serves to provide some of the necessary background on phased array antenna theory to be able to understand the antenna designed and developed as part of this thesis. An antenna array is an array of antenna elements configured in a specific geometry that results in improved performance over just a single element. For example, antenna arrays can be used to narrow the

beamwidth of a signal or increase the overall sensitivity of the system. Both antennas and antenna arrays are characterized by their response patterns or their gains patterns. For an antenna array, the response of an antenna array is given by

$$E(\theta, \phi) = R(\theta, \phi)AF. \quad (3.1)$$

In this equation $E(\theta, \phi)$ is the resulting array response, $R(\theta, \phi)$ is the response of a single element of the array, assuming that each element is the same within the array, and AF is a term called the array factor. The array factor encompasses the effect of the geometry of the array and the weighting (e.g. the effect of inline elements such as phase shifters) on the array response itself and is given by

$$AF = \sum_i^N w_i e^{j\mathbf{k} \cdot \mathbf{r}_i}, \quad (3.2)$$

where w_i is the weighting vector applied to the element, described later, and $\mathbf{k} \cdot \mathbf{r}_i$ represents the geometry of the antenna array with

$$\mathbf{k} = \frac{2\pi}{\lambda} [\sin(\theta) \cos(\phi) \quad \sin(\theta) \sin(\phi)], \quad (3.3)$$

and

$$\mathbf{r}_i = [x_i \quad y_i], \quad (3.4)$$

where x_i and y_i define the position of the element within the array.

From these equations, it is apparent that the response for an antenna array, with identical elements used within the array, is a function of three elements: the response of a single element, the weighting vector, and the geometry of the array. The effects of each of these three elements is made most clear when presented in the reverse order, building up from the geometry of the array through the weighting vector to the impact of the response of a single element. Therefore, the following sections provide more detailed information on the importance and necessary design considerations required for the “perfect” array.

3.2.1 Array Geometry

Antenna arrays largely fall into two categories of array geometries: linear arrays, where all the elements are in line, such as in Figure 3.2 (left), or planar arrays, where all the elements are laid out on a 2D surface, such as in Figure 3.2 (center) and (right) [58]. The arrangement of the antenna elements within the array not only changes the resulting response pattern but also changes the steerability of the antenna, as the weights are applied to the elements and therefore are affected by

the geometry. Therefore, the geometry of the array is the natural starting point for understanding the basics of antenna arrays.




Figure 3.2: Example element configurations for antenna arrays: (left) linear array, (center) rectangular grid planar array, and (right) hexagonal grid planar array.

One of the main motivations for using an antenna array is to be able to have higher gain and narrower beamwidth (more directionality) than what a single element can provide because, when placed correctly, the array factor combines the effects of each element creating a narrower beamwidth [59]. Additionally, the actual configuration of the elements, say in a planar array, can change the resulting symmetry of the array gain pattern. For example, a layout such as shown in Figure 3.3 results in the gain pattern shown in Figure 3.4. Notice that with an array geometry that is not rotationally symmetric, the resulting gain pattern does not have rotational symmetry.




Figure 3.3: Example element layout for a 3-element array on a rectangular grid.

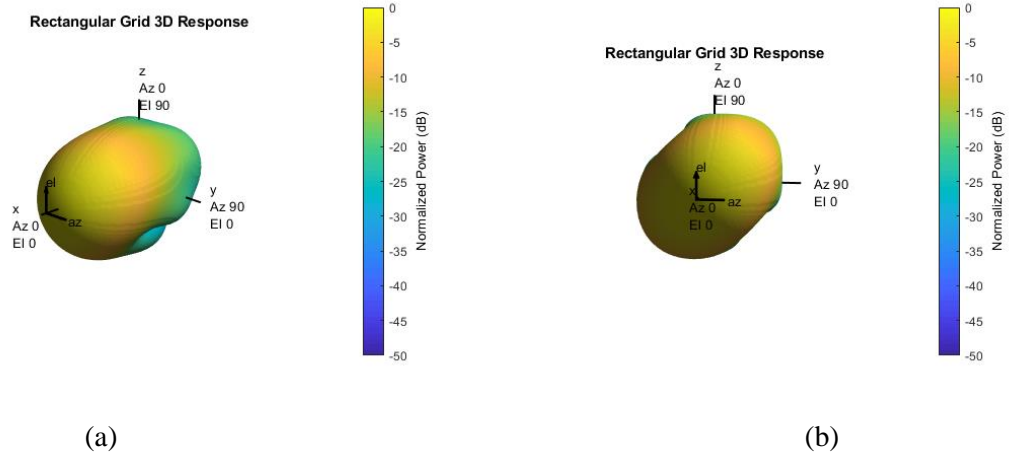


Figure 3.4: Array response for 3-element array on a rectangular grid with (a) isometric view and (b) head-on view.

The layout of the elements themselves within the classification of a planar array can differ from array to the next, resulting in different array responses. For example, there can be rectangular arrays, where the elements are laid out on a rectangular grid (Figure 3.2 (center)), triangular grid arrays, such as the hexagonal grid in Figure 3.2 (right), where the elements are laid out on a triangular lattice, or circular arrays, where the elements are laid out in a circular configuration. The packing efficiency of a triangular lattice, such as the hexagonal grid, has the best packing efficiency of any planar array geometry, making it a popular lattice used in arrays today [59], [60], [61]. For each of these array geometries, the key defining element is the spacing between the elements, or the inter-element spacing, in the x and y directions of the grid, d , with a scale of interest being on the order of the wavelength of the frequency of interest, λ . This inter-element spacing also affects the characteristics of the antenna array gain pattern. With a spacing that is too large, as the array is steered, grating lobes will occur providing false positives, or clutter, in the resulting RSS pattern [58]. This inter-element spacing also affects the beamwidth and gain of the antenna array [58]. These effects are demonstrated in Figure 3.5 with Figure 3.5a showing the array factor for a spacing of half a wavelength, $d = \frac{\lambda}{2}$, with a single narrow peak in the center and peaks again at $\pm 180^\circ$. If the spacing is increased to $d = \lambda$, such as in Figure 3.5c, the central peak is a lot narrower – that is the beamwidth has narrowed – however the secondary peak has moved to $\pm 90^\circ$. These secondary peaks, known as grating lobes, result in undesired false positives if being used for an application such as determining DoA from the RSS pattern. Finally, if the elements are brought closer together, such as the $d = \frac{\lambda}{4}$ spacing in Figure 3.5b, there are no grating lobes but the main beam widens. The

threshold for when grating lobes appear is at $d = \frac{\lambda}{2}$, therefore, the typical chosen value for the spacing is

$$d \leq \frac{\lambda}{2}. \quad (3.5)$$

With setting $d = \frac{\lambda}{2}$, the array can provide the narrowest beamwidth possible without grating lobes which makes this a commonly used choice for the inter-element spacing of an array [58].

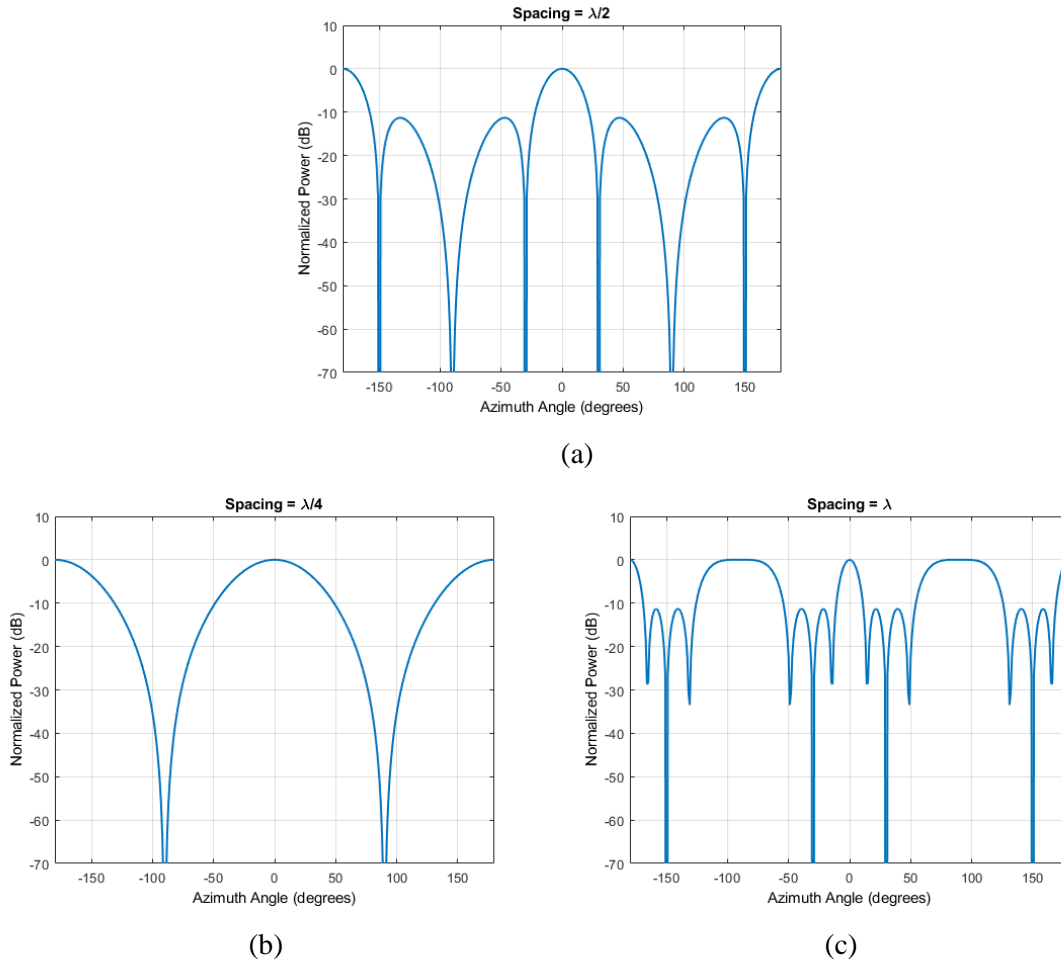


Figure 3.5: Array factor for variable inter-element spacing for a uniform linear array with (a) half-wavelength spacing, (b) quarter-wavelength spacing, and (c) one wavelength spacing.

3.2.2 Weighting

While the geometry defines the static component of an antenna array's response, the weighting component can create a dynamic element of the antenna array's response. Mathematically the

weighting term is the w_i in Equation 3.2, but, more intuitively, the weighting vector is a change of the amplitude and phase of the signal for each of the elements within the antenna. For example, adding a set of phase shifters in line with each antenna element – creating a phased array antenna – within the array allows the array gain pattern to be shaped as desired. The phase shifters apply a phase to the output from an element, creating a weight, which effectively “steers” the beam [61]. While the weights applied to each element steer the antenna, the physical limitations of what steering is possible is dictated by the array geometry. For example, for a linear array, phase shifters can be used to steer the beam in the azimuth direction as depicted in Figure 3.6. For a planar array, changing the phase of the elements can allow the antenna to be steered in both polar and azimuth angles as depicted in Figure 3.7 [61]. More generally, there is a set of closed-form equations that relate the desired steering angle to the necessary phase shifts. For a desired steered angle (θ_s, ϕ_s) the weighting term is as follows,

$$w(\theta_s, \phi_s)_i = e^{j\psi_i} \quad (3.6)$$

where ψ_i is the phase shift that is a function of the desired steered angle and is given to be

$$\psi_i = \frac{2\pi}{\lambda}(x_i \sin(\theta_s) \cos(\phi_s) + y_i \sin(\theta_s) \sin(\phi_s)). \quad (3.7)$$

It can be seen here that the equations are very much non-linear and that introduces some challenges when it comes to calibrating a phased array antenna as it is important to ensure that there are no additional biases (of phase shifts) added to the signal due to the design of the antenna itself resulting in nonlinear distortion of the true steered angle. It can also be seen that using a weighting scheme like this allows precomputing the phase shifts for a given steering command since there is a direct relationship between steering command and phase shift.




Figure 3.6: Azimuth angle for the possible steering direction of a linear array.




Figure 3.7: Spherical coordinate system demonstrating the polar (θ) and azimuth (ϕ) angles that a planar array in the XY plane can be steered.

There are many different strategies for calculating the weights to be applied to the system, but the focus of this thesis is a data-independent beamformer [62]. That means that the weighting is only the phase and can be pre-computed. Other methods, including adaptive beamforming, change the applied weights in real-time based on the incoming signal characteristics but require prohibitive amounts of computation making them impossible to be flown on a UAV [57].

3.2.3 Element Response

With an understanding of the effects of the geometry and the weighting vector, the effects of the response of a single element can be understood. As seen in Equation 3.1, the array factor multiplies its effect to the element response itself, therefore the element response itself is important in the ability to steer different angles. For an illustrative example, Figure 3.8 shows, in blue, the array factor computed for a 7-element hexagonal grid phased array antenna steered to 80° and shows, in yellow, the element response. If the element was an isotropic element, meaning having a value of 1 at all angles, the resulting response would be that of the array factor directly. However, Figure 3.9 shows the result of that array factor applied to the patch element. Notice that the gain is significantly reduced, and the peak is not at 80° due to the element's significantly reduced response near 80° . Here it is clear that as the beamwidth of a given element is reduced, the effect that the array factor can have on the resulting pattern is degraded as the weak response of each element on those extreme angles has a much greater impact than the increase in the response of the array factor. For this reason, antenna arrays typically are created with patch antennas as they provide some of the largest beamwidths of any antenna type. This physical limitation on the maximum polar angle that can be successfully steered is an important property when it comes to the design and operation of a phased array antenna.

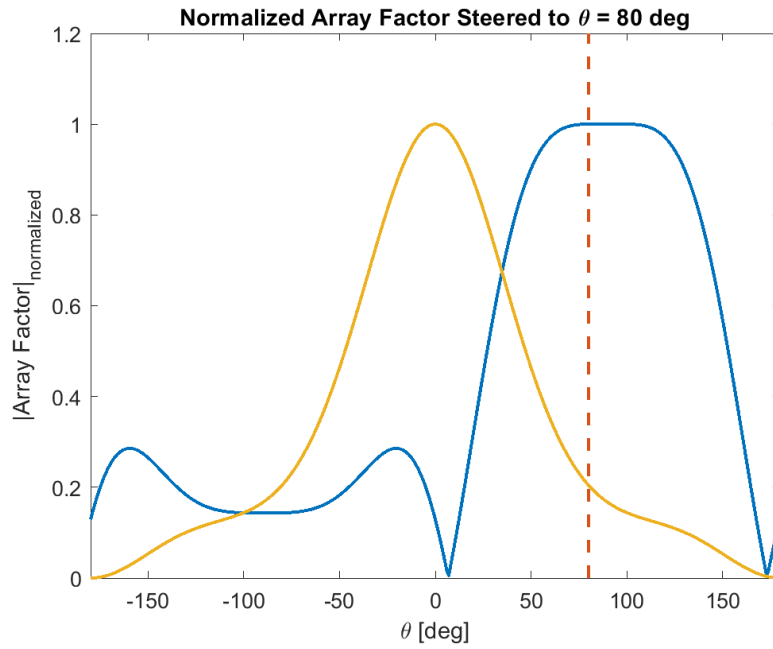


Figure 3.8: Patch antenna element response (yellow) and array factor (blue) for a steering command to 80° for a 7-element hexagonal grid array.

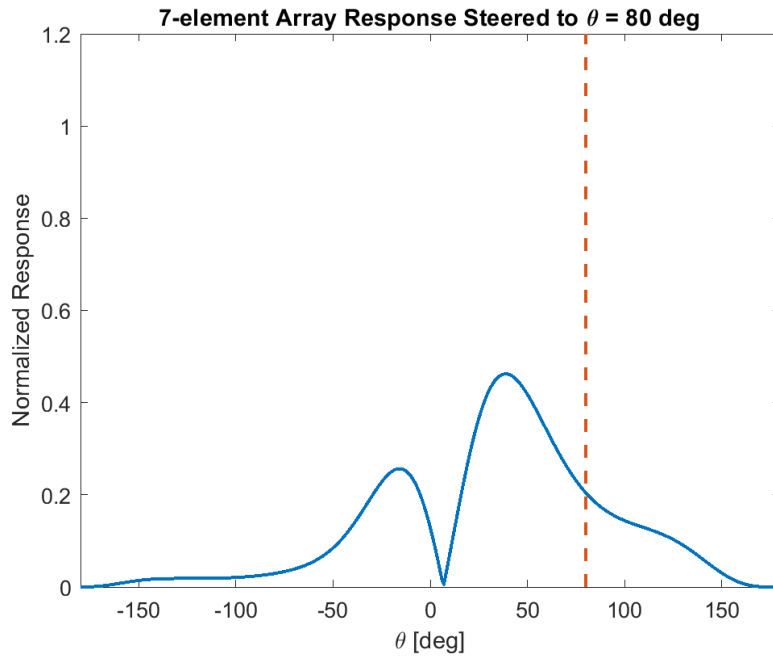


Figure 3.9: Response of antenna array steered to 80° with a patch antenna with a 60° beamwidth demonstrating the inability to steer to the desired angle due to the single element response.

3.3 The Antenna

This section presents the first contribution of this thesis: the development of a computationally and physically lightweight electronically steered antenna that provides high rate RSS patterns onboard an off-the-shelf drone. This section focuses on applying the theoretical antenna array design elements specifically to design an antenna suited for the localization need of the system. Additionally, it discusses specifically how the antenna was built on a Printed Circuit Board (PCB) and calibrated to ensure proper operations. Finally, it presents an overview of the core command and control of the antenna that enables the high-rate measurements along with the ability to handle time-varying environmental effects.

3.3.1 Design

With one of the main goals of the antenna developed in this thesis being a small, lightweight antenna that can easily be flown on a UAV, the antenna developed contains only 3 elements – the fewest number of elements to be able to make a planar array and therefore steer the main beam in both polar and azimuth angles. Given the antenna only contains 3 elements, the rectangular grid is not a possible choice since, with only 3 elements, the result is a right triangle configuration. Figure 3.4 shows the 3D response pattern for that configuration and it lacks symmetry making it an unideal configuration for the signal localization mission. However, hexagonal grid arrays, which are very popular due to their space efficiency and rotational symmetry, are typically found in 7-, 19-, etc. element sizes. Therefore, to create the 3-element configuration, the 3-element subset depicted in green in Figure 3.10 was selected [59].




Figure 3.10: View of a 19- (blue), 7- (orange), and 3- (green) element array based on a hexagonal grid.




Figure 3.11: Layout of elements in 3-element phased array antenna with element numbering.

Despite only using 3 elements, the inter-element spacing remains the same as what is typically found for a hexagonal grid array. The spacing of the elements for a hexagonal grid array in the X and Y directions, depicted in Figure 3.11, is

$$d_y = \frac{\sqrt{3}}{2} d_x, \quad (3.8)$$

$$d_x = \frac{\lambda}{2}. \quad (3.9)$$

More specifically, using the numbering of the elements from Figure 3.11, the element positions for the 3-element version of the antenna are as follows,

$$\mathbf{r} = \begin{bmatrix} r_0 \\ r_1 \\ r_2 \end{bmatrix} = \frac{1}{4}\lambda \begin{bmatrix} 1 & -\sqrt{3} \\ -1 & -\sqrt{3} \\ 0 & 0 \end{bmatrix} \quad (3.10)$$

For comparison, Figure 3.12 shows the 3D response pattern for the selected hexagonal grid layout. Note that in this case, the main beam has rotational symmetry which is desired when it comes to rotating the beam in an azimuth direction to generate a directional RSS pattern. Additionally, there is a desire to have a large beamwidth in the polar direction to make sure the antenna is equally sensitive to signals on the horizon and signals close to the vehicle to avoid missed detections due to orientation.

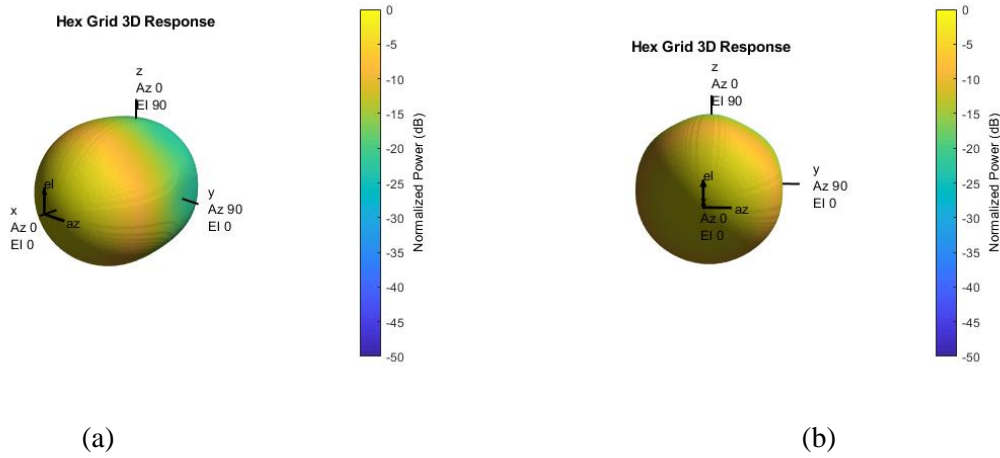


Figure 3.12: (a) isometric and (b) head-on view of the array response for the 3-element antenna array with a hexagonal grid.

Now that the geometry is set, defining the static component of the array factor, the next piece is to dive into the weighting array. Placing one of the elements (element 2) on the origin helps to simplify the weighting vector, and therefore the number of calculations required onboard the computer that runs the phased array antenna, as it results in that element to always have zero degrees of phase shift (i.e. $\psi_2 = 0$). With this simplification made, the transformation from the spherical coordinates for the desired steering angles to the phase shifts can be simplified to the following set of equations [61]:

$$\psi(\theta, \phi) = \begin{bmatrix} \psi_0 \\ \psi_1 \\ \psi_2 \end{bmatrix} = \begin{bmatrix} -\pi \sin(\theta_s) \left(\cos(\phi_s) + \frac{1}{2} \sin(\phi_s) \right) \\ \pi \sin(\theta_s) \left(-\cos(\phi_s) + \frac{1}{2} \sin(\phi_s) \right) \\ 0 \end{bmatrix} \quad (3.11)$$

This selected geometry results in an antenna with rotational symmetry and a large beamwidth in both the polar and azimuth directions. Specifically, for a standard patch antenna with an 80°

beamwidth selected as the element, the theoretical beamwidth in each direction is 60° and 60° , respectively.

The last element of the design is understanding the physical limitations that only having 3-element poses on the steerability of the antenna. Figure 3.13 plots out the commanded polar angle (θ) and the theoretically possible polar angle (θ) given the geometry which demonstrates some of the limitations that need to be accounted for in this antenna; specifically, the fact that the maximum steerable polar angle is about 30° . The geometry, for the most part, does not affect the performance of the azimuth angle (ϕ), shown in Figure 3.13. Note that, despite these limitations, the beamwidth of 60° means that signals on the horizon are not greatly attenuated, making this an ideal configuration for the signal localization mission.




Figure 3.13: Performance of the steering capabilities of the antenna with (a) an analysis of the ability to set ϕ and (b) the ability to set the θ based on measurement from a vector network analyzer.

3.3.2 Schematic / Physical Implementation

For the antenna to be able to measure the RSS value, the combined signal of all the elements must go to a power detector. Unfortunately, when building an antenna array, each of the elements cannot simply be directly connected to a power combiner; additional components are needed throughout the entire process to control the phases and ensure that the antenna is tuned to the desired frequency band – in this case, the 2.4GHz spectrum, as discussed in Chapter 2. Each antenna element passes through a set of different RF components placed on a Printed Circuit Board (PCB) to ensure the antenna is tuned to the proper frequency, is controllable, and combines the signal strengths from each element properly to be able to measure the RSS value for the commanded angle. The overall

schematic of the antenna is shown in Figure 3.14 and Figure 3.15. At a high level, each of the element traces is combined in a power combiner and combined completely into a power detector.

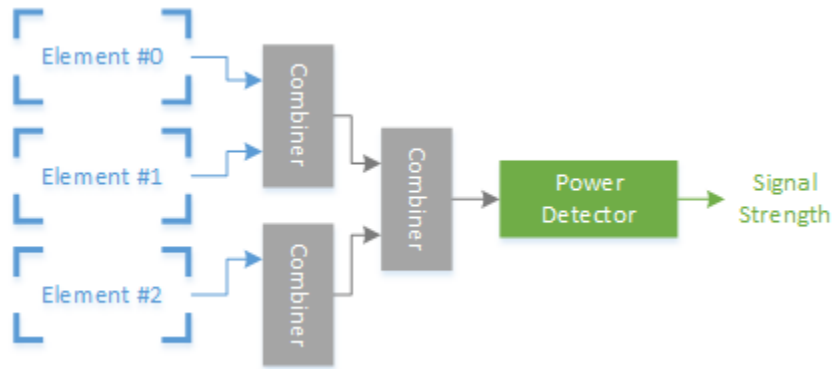


Figure 3.14: Schematic of the 3-element phased array antenna.




Figure 3.15: Detailed schematic of RF components for each antenna element within the 3-element phased array antenna.

As detailed in Figure 3.15, the signal from each antenna element passes through three key components before being combined in the power combiners: a Surface Acoustic Wave (SAW) filter, a Low Noise Amplifier (LNA) and a phase shifter. The 15MHz wide SAW filter ensures that only signals within 15MHz of 2.45GHz, the frequency of interest, pass through the next stages of the antenna, removing out of band noise from affecting the readings of the power detector. For a GNSS version of this antenna, the SAW filter would be centered on the GPS L1 frequency where an interference device typically broadcasts [45]. The 20dB LNA boosts the power level of the incoming signal to increase the sensitivity, and therefore the operating range, of the antenna. The 8-bit phase shifters provide the necessary control over the phase on the incoming signals to successfully steer the resulting beam in the desired two dimensions. Each 8-bit phase shifter breaks down the 360° of possible phase shifts into 256 discrete values, meaning that each phase shifter can impart a phase shift on the signal passing through it in step sizes of about 1.4° each. Each of

these three feeds are then combined in a set of power combiners – which effectively sum the signals from the three elements – before finally going to the RF power detector. The power detector simply measures the signal strength – or the amplitude – of the incoming signal, proving the core measurement to eventually get a bearing observation.

In addition to the signal strength value, to create an RSS pattern, the commanded steering angle also needs to be known at the time the RSS measurement is made. To handle this, and to control the phase shifters themselves, the board also has an ST Microelectronics STM32L4 ARM Cortex microprocessor. This chip carries out all the math needed to calculate the required phase shifts, executes the phase shifts at the desired time intervals, and handles aligning a signal strength value with the commanded phase shift. The upcoming Command and Control section provides much more detail on the critical elements of the command and control used on the antenna to ensure rapid rotation and reliable operation of the antenna.

One of the biggest challenges of building an RF system is that each element in the chain adds additional noise in both amplitude and phase that needs to be accounted for to ensure proper operation. Therefore, after the antenna array is assembled, it is important to calibrate the system, a process that is described in the following section.

3.3.3 Calibration

When placing the components on the PCB itself, careful consideration needed to be made to ensure the signal from each antenna element is affected similarly. However, even when being very careful with the design of the PCB, each path from an antenna element to the signal strength detector is not identical resulting in small biases between each antenna. While the design minimized possible sources for those biases, the biases still needed to be determined for accurate control of the antenna beam with the phase shifters. Those biases are a constant offset in the phase domain, so the equations for the steering of the antenna need to be modified as follows

$$\hat{\psi} = \begin{bmatrix} \hat{\psi}_0 \\ \hat{\psi}_1 \\ \hat{\psi}_2 \end{bmatrix} = \begin{bmatrix} \psi_1 \\ \psi_2 \\ \psi_3 \end{bmatrix} + \begin{bmatrix} b_0 \\ b_1 \\ b_2 \end{bmatrix}. \quad (3.12)$$

The true phase shift applied, $\hat{\psi}$, to each of the antenna elements is the desired phase shift, ψ , plus the bias from the board itself that is different for each antenna element as a result of small differences in the construction of the board. Therefore, to make sure that the true phase shift is the

desired phase shift, those biases, $\begin{bmatrix} b_0 \\ b_1 \\ b_2 \end{bmatrix}$, need to be computed.

This bias, while constant in the phase shift domain, creates a non-linear effect in the spherical coordinate system for which the steering is defined in. To better understand the effects of this non-linear transformation on the pointing angle, consider a constant board-level bias of 10° on a single antenna element (i.e. let $b_0 = 10^\circ$). The reverse transformation made from the phase shift values to the spherical coordinate results in a non-linear pointing error depicted in Figure 3.16. In this figure, when the bias is present (orange line) the error between the desired and resulting steering angle is a non-linear function of the desired steering angle. This results in matching a measured signal strength with an incorrect pointing angle, distorting the resulting gain pattern and causing errors in bearing observations. Because these effects are detrimental to being able to adequately steer the main beam where desired, it is important to calibrate the antenna itself by determining these board-level offsets.




Figure 3.16: Example nonlinear effects of a board-level bias in the phase on a single element, showing the error between the desired steering angle (a) ϕ and (b) θ and the resulting steering angle with no bias (blue) and with the 10° bias present (orange).

A vector network analyzer was used to measure the phase bias of the line for each of the antenna elements. A vector network analyzer generates a signal with a known phase and amplitude and connects to both ends of the path to analyze and when it receives the signal that was generated it determines the phase offset of the signal due to the path the signal took. By directly measuring this value, the equations for determining the necessary phase shifts to achieve a specific pointing angle can be adjusted to account for the bias terms in Equation 3.10 with their measured values.

The second major component to calibrate on the antenna board is the power detector itself. The RF components for each antenna element used to combine the signals and to control the antenna array introduce white noise in both the amplitude and the phase of each signal. As a result, the controller on the antenna is designed to average over several samples to reduce the effects of the white noise on the signal strength readings. Figure 3.17 shows the standard deviation of the

measurements as the averaging window changes from having a window of a single sample to a window of 100 samples. From this collected data, it can be seen that a larger window size helps in reducing the effects of the noise induced by the board itself.




Figure 3.17: Output of power detector with varying window sizes for averaging signal strength samples to handle white noise induced by the RF components on the phased array antenna PCB.

3.3.4 Command and Control

The STM32L4 microprocessor onboard the antenna is responsible for setting the desired phase of each phase shifter and of reading the value of the analog signal strength detector. In order to achieve the fastest speeds possible on the phased array platform, the computation costs were reduced as much as possible. Mathematical operations were reduced when possible to the minimum required by leveraging some pre-computed values to ensure that executing mathematical operations is not a limiting factor on the speed of rotation possible by the antenna.

Unlike any other rotation-based method of generating an RSS pattern, the controller being responsible for both rotating the beam and measuring the signal strength values provides an unprecedented level of control over time. The diagram in Figure 3.18 helps to visualize the control that the microprocessor has on the system. The microprocessor on the phased array antenna has direct control of how often the antenna steps through the desired phase shifts. Therefore, the true inputs to the phased array antenna are the step size in the azimuth direction (controlling the density




Figure 3.19: Example RSS patterns generated with a pulse jammer with a dwell time of (a) 0.5ms and (b) 1ms.

The second major advantage of the antenna being able to electronically rotate is the fact that the azimuth angles of the specific RSS measurements are independent of external environmental effects. When the stop and rotate modalities are rotated by the UAV platform, the quality of the azimuth angles of the RSS measurements are a function of the control system onboard the UAV. This means that for the stop and rotate modalities, environmental factors such as wind, can affect the rotation rate of a multirotor UAV and results in varying performance of RSS pattern generation [48].

3.4 Performance

There are several important metrics to evaluate to determine the performance of the phased array antenna designed and built in this chapter. Of most interest is the antenna's ability to create directional RSS patterns from which a direction of arrival can be determined and that the process can happen very quickly. However, before assessing the ability to measure a DoA, the antenna's ability to create reliable patterns must be assessed. Specifically, how well does the commanded angle match a given direction of arrival measure given the geometry of the antenna and the use of 8-bit phase shifters that have a limited resolution. To evaluate the steering performance a vector network analyzer was used to determine the actual phase shift set given a commanded phase shift.

From there, the second point of interest is to ensure that the patterns have directionality and that they are generated at a reasonably fast rate. This is all evaluated on a test platform that allowed

placing the antenna array in different configurations to a signal source allowing to test the antenna's performance for a wide range of incident angles on the source. Additionally, the performance evaluation process ensures understanding the beamwidth and other parameters of the antenna.

This section presents the results of those bench tests to analyze the “bare necessities” of the capabilities needed from this antenna. In later chapters, the antenna's performance will be evaluated in a variety of different environments to further examine its capabilities of providing direction of arrival measurements at a suitably high rate onboard a UAS.

3.4.1 Steering Performance

Analyzing, and understanding, the steering performance of the antenna ensures minimizing sources of DoA errors due to the construction of the board and helps to properly characterize the expected performance to include that information into the filter. The phased array antenna board's performance is driven by two key elements: the geometry and the phase shifters being used to control the phase of each of the elements. The geometry has been discussed at length in the previous section and won't be discussed more since, at this point, the geometry can't be changed. One of the purposes of this performance evaluation is understanding how well the board was built (the geometry was adhered to) and understanding the characterization of the build of the antenna itself. The phase shifters have been discussed briefly, but these performance results help to provide a better understanding of the impact the phase shifters have on the expected performance of the antenna's true purpose – determining DoA to a signal source.

On the board, the phase shifts are set using an 8-bit phase shifter. Each 8-bit phase shifter breaks down the 360° of possible phase shifts between each antenna element into 256 discrete values, limiting the achievable set of phase shifts and therefore limiting the exact steering angles the phased array antenna can achieve. The limitation is mostly dictated by the polar angle, with steering accuracy in azimuth degrading below a polar angle of $\theta = 10^\circ$ and above a polar angle of $\theta = 65^\circ$, as shown in Figure 3.20a and Figure 3.20b. This is just due to the fact that, looking at Equation 3.11 again, for small values of θ , the value of the phase shift (ψ) changes only a little even for large changes in ϕ , and with the precision of the phase shifters being limited to about 1.4° between steps, the step size of the phase shifter is too coarse to achieve proper steering.

Being able to understand and know the limitations allows the mission profile to be designed in the most successful possible way. It is important to note that none of these constraints make the mission impossible.




Figure 3.20: Performance of steerability as a function of the desired θ angle, the polar angle, in (a) the resulting error in ϕ , the steered azimuth angle, and (b) the resulting error in θ , the steered polar angle.

3.4.2 Gain Pattern

Being only a 3-element array, the resulting gain pattern of the antenna has a large main lobe beamwidth and can have large sidelobes compared to a highly directional antenna. When looking at an RSS pattern at a given steered polar angle, the antenna can result in three different possible patterns, depicted in Figure 3.21. This is a property of the antenna that can be leveraged by the estimator.

Which of the three possible gains patterns created in a rotation is a function of the relative location of the antenna with respect to the signal source. If the antenna is overhead of the signal source, the directionality of the antenna is lost, resulting in very low directional RSS patterns such as the one in Figure 3.21a. When not overhead of the signal source, the pattern generated is a function of the incident angle – the direction of arrival to – the signal source. The pattern can take on the shapes depicted in Figure 3.21b and Figure 3.21c, each having its own benefits and drawbacks. The pattern in Figure 3.21b has one main lobe and two small sidelobes with about 5dB of suppression. While the pattern in Figure 3.21c only has a single sidelobe in the back of the pattern but has less suppression of only 3dB. While the pattern in Figure 3.21c may be advantageous due to only having a single sidelobe, the increased suppression of Figure 3.21b can also be beneficial. As a result, development in future chapters focus on both types of patterns that demonstrate directionality.

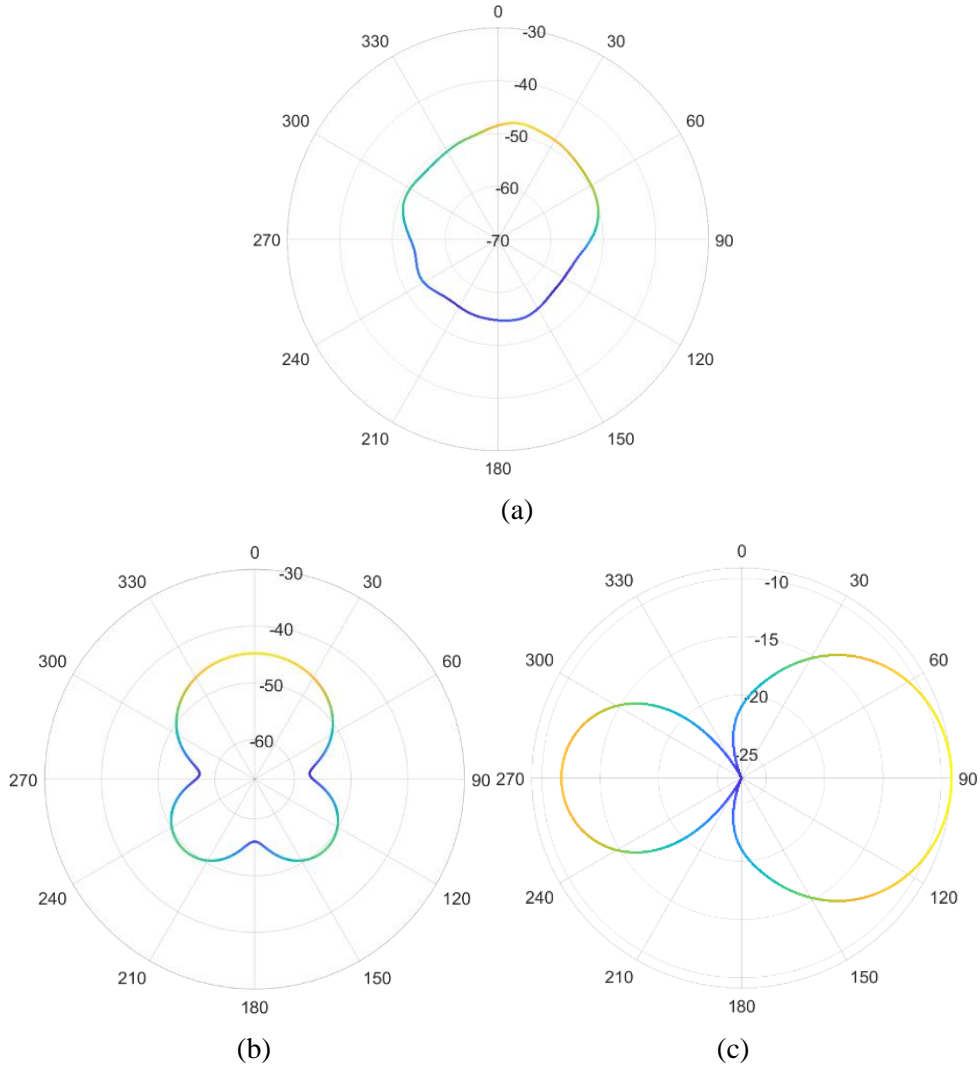


Figure 3.21: Example of the three possible resulting RSS patterns from the 3-element phased array antenna with (a) a single unpronounced lobe, (b) a single strong lobe and two large side lobes, and (c) one main lobe and strong back lobe

3.5 Conclusion

This chapter presented the first contribution of this thesis: the design and development of a computationally lightweight three-element phased array antenna capable of COTS UAV based operations. Shifting from the physical rotation modality to an electronic rotation modality increased the rate of RSS pattern generation from 40 seconds to a mere 0.37 seconds, a near hundredfold improvement in speed, for an equally detailed RSS pattern generated. Additionally, the antenna's command inputs of both angular step size and time means that the antenna is not just capable of setting a user-defined rotation rate but is also capable of being configured to adapt to time-based

environmental concerns, such as a pulsed jammer. Finally, the antenna provides a self-contained rotation scheme, meaning that the vehicle's motion is no longer required for the generation of measurements (RSS patterns) resulting in the sensor being completely decoupled from the vehicle. This will allow the path planning scheme presented in Chapter 5 to focus on the best movement of the vehicle within the environment instead of having to factor in the cost of making an RSS measurement.

Chapter 4

Algorithm for Determining DoA

From a given RSS pattern, the next step in the process for successful localization of a signal source is turning that RSS pattern to a more usable measurement – the Direction of Arrival (DoA) of, or bearing to, the RFI source. To determine the DoA from an RSS pattern, this chapter presents the second contribution to this thesis: a robust algorithm capable of determining the DoA of a signal in real-world conditions that parameterizes the underlying RSS pattern in such a way that it provides an unprecedented level of information for an estimation algorithm. Section 4.1 provides a summary of previously developed methods for determining bearing from an RSS pattern and discusses the shortcomings of these methods when it comes to more real-world scenarios. Section 4.2 discusses the real-world effects that need to be handled by an algorithm that determines bearing from a given RSS pattern. Given the understanding of the shortcomings of existing DoA determination methods, Section 4.3 presents the second contribution of this thesis: a lobe-detection based DoA algorithm. Finally, the performance of the lobe-detection based DoA algorithm compared to existing methods is presented in Section 4.4 based on real-world flight test data collected in a variety of environments that simulate possible environments for which a system might be deployed.

4.1 Related Work and Motivation

Fundamental to a bearing only localization mission is the ability to determine the DoA of, or the bearing to, the signal source. The state of the art when it comes to determining DoA from an RSS pattern broadly falls into one of two categories: maximum value based methods and cross-

correlation methods. For completeness, this section also covers some other more specialized methods that do not fall into either category.

4.1.1 Maximum Value

This first family of methods is by far the simplest of the methods described. As implied by the name, this family of methods simply takes the DoA to be the signal source to be the azimuth direction for which the RSS pattern has a maximum value. The best example of the simplest approach is the parrot AR drone that continually rotated, where the magnetometer-based azimuth angle of the maximum value for a given rotation is directly used as the DoA [34]. In the case of the “maple seed drone”, due to a limited pattern density, the results of several rotations are combined together in a set of azimuth bins with the center of the maximum value bin being used as the DoA [33]. In the flight tests of the maple seed UAS, across 2671 bearing estimates, 613 (or 22.9%) of the estimates are considered to be outliers (missed detections) [33]. For the parrot AR drone research, the specified missed detection rate (or probability of measuring an outlier) is 19% and the standard deviation of the estimate is 14° [34].

This maximum value method has been modified by several other researchers to fit their customized setups. For example, using multiple omnidirectional antennas and determining the DoA to be in the direction of the antenna that had the higher of the two RSS values in the direction that results in the highest RSS value [49]. A “happy medium” example is a rotating directional antenna that uses the same binned maximum value method, but instead of just taking the maximum value bin, it averages the azimuth angles of the set of bins above a given threshold [50].

While the maximum value method does not require a prior knowledge of the exact antenna pattern – it merely requires that the antenna is directional in nature – and fundamentally only requires a handful of measurements to return a result (albeit a possibly bad result) it is susceptible to missed detections in cases with environmental noise, weak maximum RSS value, antennas with large side lobes, or environments with multiple sources. This weakness is illustrated in Figure 4.1 where an RSS pattern from the 3-element phased array antenna, with 2 large side lobes, results in the maximum value method returning an incorrect DoA value. This also highlights the second weakness of this method: the only piece of information returned is the DoA meaning there is no information for an estimation algorithm to use to determine the quality of the measurement.




Figure 4.1: Example RSS pattern and true DoA (red dashed line) with large side lobes resulting in an incorrect DoA by the maximum value method (orange line).

4.1.2 Cross-Correlation

The cross-correlation method compares the measured RSS pattern with a known “truth” pattern for the antenna. Cross-correlation has been used to determine bearing on a rotating robot in [32] and [37] with the following equation

$$c(\gamma) = (R_n * R_t)[\gamma] = \sum_{\alpha=0}^{360^\circ} R_n(\alpha) R_t(\alpha + \gamma). \quad (4.1)$$

The truth pattern R_t is shifted by some angle γ . At every angle α in the measured pattern, the value of this shifted pattern is multiplied by the value of the normalized, measured pattern, R_n [32]. The cross-correlation coefficient c is the sum of all such products for a single shift γ and is computed for every possible shift γ . The shift yielding the highest cross-correlation coefficient is taken to be the DoA to the signal source.

Using a known “truth” pattern resulted in better performance in [32], but does also result in a drawback of needing to know the “truth” pattern. To generate the “truth” pattern, one method is to put the antenna in an anechoic chamber to measure the RSS pattern in a nominal environment. Getting the “truth” pattern is not trivial with some researchers using the advertised gain pattern from the manufacturer and others attempting to generate their own in an anechoic chamber [32], [65]. Fundamentally, cross-correlation fundamentally relies on the RSS pattern resembling the calibration pattern, but as a result, it fails when that assumption no longer holds.

The cross-correlation method does have an added benefit of returning the maximum cross-correlation value that can be used as an indicator of confidence in the measurement – at least for the case of a single signal source. In [32] the tests of the system resulted in 18% of the measurements being classified as missed detection using their integrity check based on the cross-correlation value computed and the rest of the measurements have a standard deviation of 4.4° for a static signal source.

However, the cross-correlation method does suffer a significant drawback: it requires the knowledge of a “truth” pattern and assumes that the measured pattern resembles the expected “true” pattern. In a real-world environment, this may work for patterns that do resemble the expected pattern, say when sufficiently close to the signal source such as in Figure 4.2a, but if the sensor is far away from the signal source, an RSS pattern can still be generated but it will no longer resemble the calibration pattern due to significant portions of the antenna’s gain not being great enough to receive the signal. The result of such a pattern, depicted in Figure 4.2b, can be distorted enough such that the cross-correlation method is unable to determine the signal strength as the pattern no longer resembles the true pattern (Figure 4.3).




Figure 4.2: RSS patterns generated with a physically rotated directional antenna at (a) close to 100m from source and at (b) the edge of the effective range of the directional antenna showing the true DoA (red dashed line) and cross-correlation computed DoA (orange line).




Figure 4.3: Nominal RSS pattern for the directional antenna.

4.1.3 Other Methods

Not all the methods fit exactly into those two families of algorithms due to the unique nature of the sensor or target signal source. Most of these more unique cases are cases where the signal being localization is a known signal (e.g. wildlife tracking tag) with known characteristics (e.g. known signal power, known signal structure – e.g. a pulsed signal –, etc) [41], [42].

Some methods are very specialized to their target signal type, for example tracking wildlife tags that have short audio pulses with multiple directional antennas allows determining the approximate DoA by banking the aircraft [43]. Yet another method is to go with more “old school” direction finding techniques which is the idea of homing in on the signal using the null from an antenna which results in needing a fairly complex flight path due to the fixed null direction on the plane [52]. Another fairly unique work that has a single element that changes its radiation pattern to do some low key beamforming, though only in a single direction [66].

4.2 Real-World Effects

None of these state-of-the-art methods is anywhere near reliable enough for a real-world application. While they’ve demonstrated promising performance in a laboratory environment, when faced with challenging RSS patterns, each of the methods break down and are unable to cope with the information provided. Furthermore, both broad categories have made a critical assumption: the assumption that there is only a single source present in the environment.

Outside of the laboratory environment it is common for an RSS pattern to no longer resemble the expected theoretical pattern due to noise in the environment. Example RSS patterns generated in a real-world environment are shown in Figure 4.4. Figure 4.4a illustrates the effect that the distance to the signal source can have on the resulting RSS pattern from an antenna with the nominal pattern depicted in Figure 4.3. Figure 4.4b and Figure 4.4c illustrate RSS pattern created from measurements from a 3-element phased array antenna and measurements from very noisy environments, respectively. Given these are just a subset of example RSS patterns that can be found in real-world environments, it is imperative that an algorithm to determine the DoA of a signal can adequately handle these types of patterns.

As is seen in the following sections, the DoA algorithm presented in this chapter is capable of overcoming these real-world challenges and is agnostic to the exact antenna making the measurement (provided it is a directional antenna) and agnostic to the number of sources in the environment, making it a considerably more robust algorithm for signal localization efforts in the real world. Furthermore, the algorithm presented parameterizes the underlying RSS pattern in such a way that it can be used as a more powerful integrity check than is even possible with the cross-correlation method. Results are presented for the performance of the lobe-detection based DoA algorithm on a rotated directional antenna and the electronically steered antenna. This partially serves to demonstrate the versatility of the algorithm presented in this chapter and serves to demonstrate a more fair comparison of the methods.

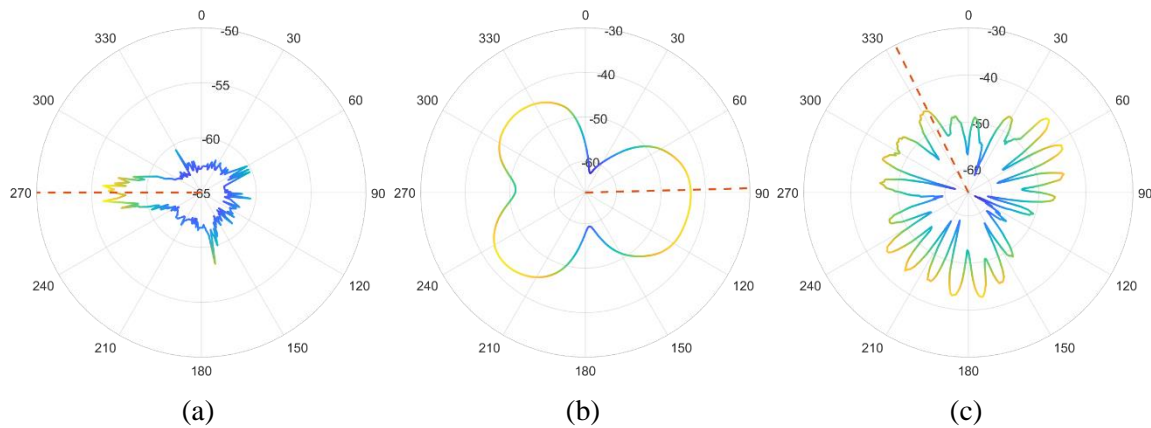


Figure 4.4: Example challenging RSS patterns from real-world environments with corresponding true DoA as the red dashed line: (a) at the edge of the range of the antenna, (b) with large side lobes, and (c) corrupted by environment noise (such as from reflections).

4.3 Lobe Detection Based DoA

Fundamentally, the other algorithms for DoA determination presented in Section 4.1 can be viewed as a very simple parameterization of an RSS pattern with just a single value, the perceived DoA to the signal source. The second contribution to this thesis – the Lobe Detection Based DoA algorithm (LDB-DoA) – instead parameterizes the RSS pattern by lobe – that is, this method looks for the lobes of the pattern and from there returns the beamwidth of each lobe present and the bearing that the lobe corresponds to. This parameterization not only provides better performance in determining the DoA of a signal, but it also provides additional information on the shape of the underlying RSS pattern to be used for integrity checks. This section describes the method in more detail. First, it presents the algorithm itself and then it presents some of the additional theoretical benefits that this algorithm provides.

4.3.1 Algorithm

This algorithm leverages the properties of a directional antenna that help support the definition of the beamwidth – that is the sharp roll-off that occurs at 3dB below the maximum gain of an antenna. Therefore, instead of needing a calibration pattern, this algorithm can simply find the location of those points where the lobe starts to drop off quickly. In doing so, the algorithm makes no assumptions on the number of lobes to look for, and therefore on the number of DoAs to return. Effectively the problem becomes a problem of finding where the lobes of the pattern are instead of simply trying to determine the DoA that the pattern is observing.

The algorithm breaks down into the following 3 steps given an RSS pattern:

1. Find the maximum RSS value within the RSS pattern (blue X in Figure 4.5);
2. Find every point in the pattern that crosses the threshold value set to be 3dB below the maximum RSS value found in step 1 and save the corresponding θ value for the measurement (region bounded by orange Xs in Figure 4.5);
3. For each pair of crossing points found in steps 2, compute the mean of the crossing points and save that value as the DoA (yellow pie slice in Figure 4.5).

The result of the algorithm is that each RSS pattern is parameterized into a set of crossing points of 3dB below the maximum value of the RSS pattern and computed DoA value for each interpreted lobe in the pattern. These crossing points can also be used to compute the beamwidth of the lobe which is additional metadata that can be used for integrity checking of the underlying RSS pattern and is discussed further in the next section.




Figure 4.5: Example detailed view of the LDB-DoA algorithm depicting the maximum value point in the RSS pattern (blue x), the crossing points 3dB below the maximum value (orange x), and the resulting detected lobe (yellow pie slice) in addition to the true DoA (red dashed line).

The unique aspect of this algorithm is the resulting parameterization of the RSS pattern. The fact that the entire RSS pattern is distilled down to the angle of the center of each lobe in the pattern (the DoA represented by each lobe) along with the beamwidth of the lobe itself. Looking at the comparative methods, the maximum value method distills the entire RSS pattern into a single value, the direction of the maximum value within the RSS pattern, and the cross-correlation method provides both the DoA along with the cross-correlation value that can be interpreted as a value of how closely the RSS pattern resembles the calibration pattern. However, of note, is the fact that both comparative methods are limited to assuming there is only a single DoA in the entire pattern. Furthermore, there is limited information about the overall shape of the RSS pattern that is returned. Effectively, the information is used to determine a DoA and not passed along, while the LDB-DoA method developed in this thesis retains information about the overall shape of the RSS pattern to be passed on to the localization algorithm.

Figure 4.6 visually shows the resulting parameterization on three example patterns from different environments. For each figure, each of the lobes, as detected by the algorithm is highlighted with a cone to represent the lobe itself (given the crossings for the edges) and the solid line down the center of the cone represents the DoA for that lobe. When it comes to a pattern with multiple possible lobes, such as the one in Figure 4.6b, the algorithm correctly identifies that

multiple lobes exist within the pattern. On the extreme, in Figure 4.6c, the noise in the RSS pattern results in 6 different lobes to be identified.

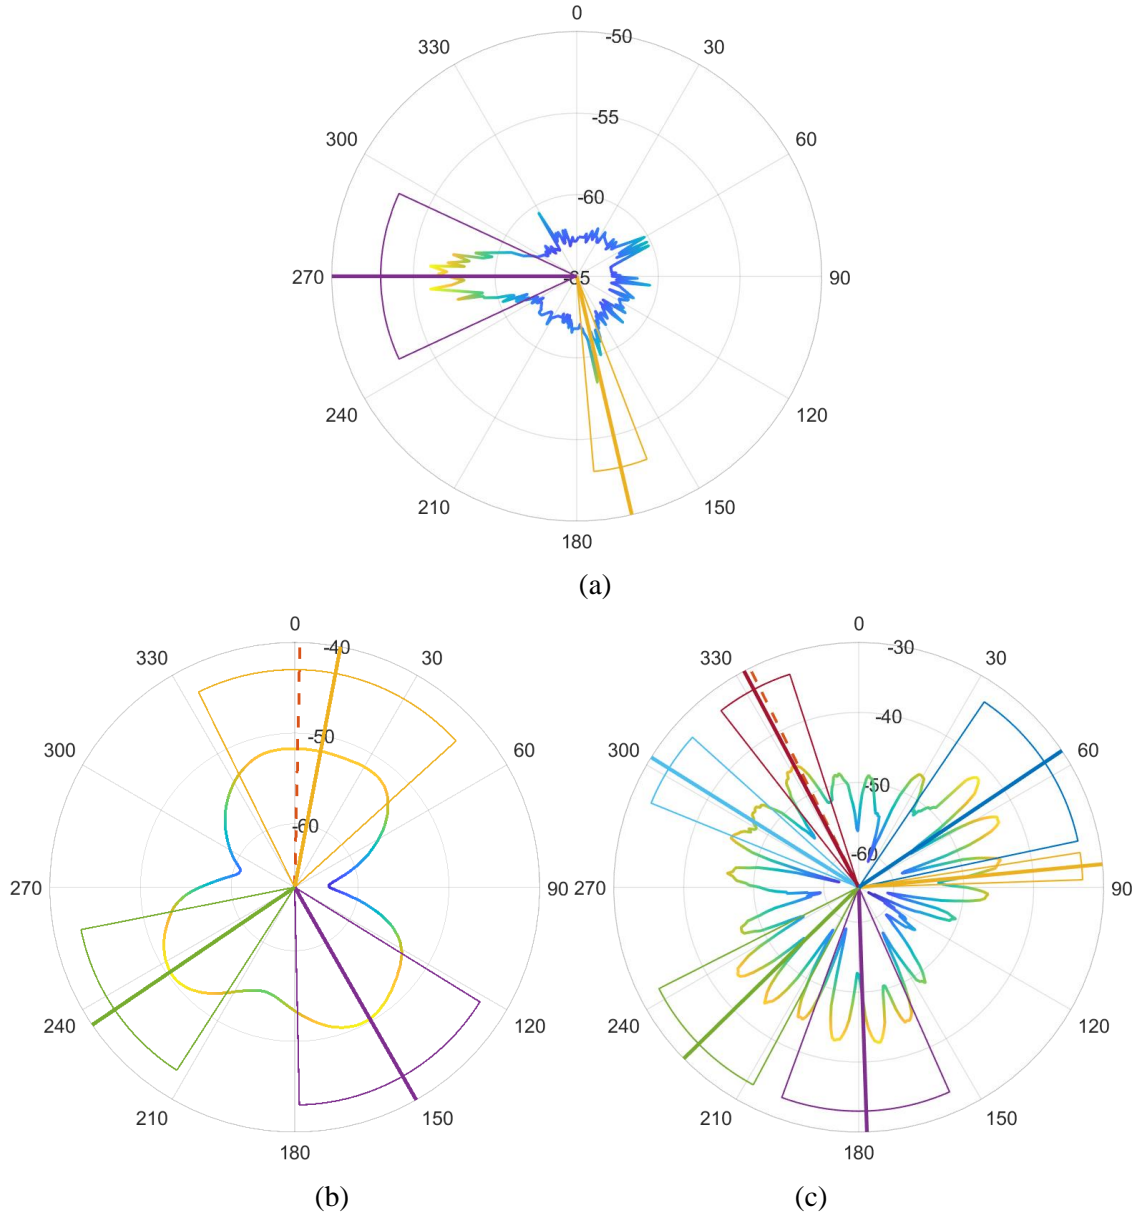


Figure 4.6: Result of the LDB-DoA algorithm on RSS patterns from challenging environments with (a) a pattern at long distance resulting in two detected lobes and therefore two possible DoA values, (b) a pattern with large sidelobes resulting in three detected lobes and three possible DoA values, and (c) a pattern corrupted by noise resulting in six detected lobes and six possible DoA values.

While these results can be directly used within an estimation framework, provided the framework has been appropriately tuned, the second component of this contribution – the integrity checks that become possible due to this parameterization of the RSS pattern – allows fine-tuning control of how observations are used within the estimation framework. This extra information returned by LDB-DoA (i.e. the shape and location of the lobes within the pattern) can be used for integrity checking of the validity of the underlying RSS pattern for a given DoA measurement and is discussed further in the next section.

4.3.2 Integrity Checks

The LDB-DoA method returns information that provides a general view of the overall shape of the underlying RSS pattern enabling the addition of a possible integrity check of the DoA measurement that can be based on a user-defined model of the antenna from which the RSS pattern is generated. The parameterization of the RSS pattern in the way described in the previous section additionally provides an avenue for being able to compute a level of confidence in both the estimator and the quality of the RSS pattern itself, providing more information than even the cross-correlation method. To compute a metric on the quality of the RSS pattern itself, some knowledge on the antenna properties are needed. Since this is the case, two different example integrity checks are described below: one for a physically rotated directional antenna and one for the 3-element phased array antenna. These are just example integrity checks that are possible, part of the real power of this algorithm's parameterizing the RSS pattern in this way is that these examples are even possible. This is a piece that needs to be modified and customized for different applications, much in the same way that an estimator is typically tailored for a given situation.

For a physically rotated directional antenna, an example simple integrity check assesses the number of lobes in the pattern and the beamwidth of those lobes. The underlying model for a simple directional antenna is that there should only be a single lobe present in the pattern. However, collected data has shown that there may be secondary lobes detected, but those secondary lobes always have a beamwidth less than 10° . Therefore, the integrity check ensures that there is either just a single lobe or that any side lobes are less than 10° . The performance of these integrity checks applied to real data sets is discussed further in the next section.

For the 3-element phased array antenna, it is known that, for a single signal source, it takes on one of the three shapes shown in Figure 3.21. Therefore, the model used for the check is a little more complex but breaks down as follows:

- With 2 measured DoAs (and therefore 2 possible main lobes) the possible underlying RSS pattern is either the 2-lobe pattern (with lobes 180° apart) or the 3-lobe pattern (meaning the 2 lobes are $\pm 120^\circ$ apart).
- With 3 measured DoAs the underlying pattern is most likely the 3-lobe pattern, meaning that the DoAs should maintain a spacing of 120° .

Using this information, the resulting list of DoAs and beamwidths can be used to confirm that the pattern has indeed taken on one of those shapes. For example, Figure 4.7 demonstrates a pattern that has not followed those rules and therefore should not be used by the estimator as it is likely that the underlying pattern is corrupted by noise and therefore the DoA values are not reliable.

As a result, different models can be used as desired within the localization filter to get the best performance possible by removing DoA measurements that come from potentially unreliable underlying RSS patterns. Furthermore, since measurements come at such a high rate, due to the rapid rotation speed of a phased array antenna, the entire system can tolerate the loss of DoA measurements allowing this integrity check to be tuned as desired.




Figure 4.7: Example noise corrupted RSS pattern with multiple lobes detected in an impossible configuration for the known antenna characteristics.

4.4 Performance

Before integrating this observation scheme into the localization scheme, it's best to evaluate the performance of this algorithm compared to the state-of-the-art algorithms presented in Section 4.1 in several different environments and configurations. In addition to providing a compelling comparison with the existing state of the art methods, this evaluation also provides much-needed information for the design of the localization filter. This section presents the results of the lobe-detection based DoA algorithm presented in Section 4.3, an implementation of the maximum value algorithm in [33], and an implementation of the cross-correlation algorithm in [32] applied to RSS pattern generated by a directional antenna in a stop and rotate scheme and applied to RSS pattern generated by the three-element phased array antenna developed in Chapter 3. For both antenna types/strategies, the evaluation presented in the following sub-section is based on RSS patterns generated across a variety of different environments. The exact flight test days and location varied slightly between each antenna type and therefore additional details on the RSS pattern dataset used is presented in the following sub-sections. These two different types of strategies are analyzed in part to demonstrate the versatility of the DoA by lobe-detection algorithm and in part to provide a fair comparison against the methods presented in Section 4.1 as they were developed for a strategy of rotating a directional antenna.

In addition to presenting the results for both the physically rotated directional antenna and the 3-element phased array antenna, results are also presented for simulated results within a multi-source environment. This additional section only provides results for the performance of the LDB-DoA algorithm as none of the comparative methods can support an environment with multiple sources. These simulated results provide a unique perspective of the versatility of using the LDB-DoA algorithm for localization missions in more complex environments than what was previously capable.

4.4.1 Physical Rotation

This first evaluation set provides the most direct comparison with previous methods for determining DoA from an RSS pattern as there is only a single signal source and the RSS pattern is generated with a physically rotated directional antenna. The RSS patterns are generated using a directional antenna and a stop and rotate strategy like that used in [37]. The dataset used for this comparison contains nearly 200 physical rotations from a 60° beamwidth directional antenna and a 30° beamwidth directional antenna. The rotations were performed at Camp Roberts (in both the open and semi-urban environments) with the 60° beamwidth directional antenna tuned for 2.45GHz and

White Sands Missile Range with the 30° beamwidth directional antenna tuned for 1.575GHz (GPS L1 frequency) at a range from the signal source varying from directly overhead to nearly 800m away distributed as shown in Figure 4.8.




Figure 4.8: Histogram of the distance from the source for the dataset of RSS patterns generated by physical rotation.

For every rotation in this data set, each method is applied to determine a DoA from each measured RSS pattern. This resulted in the overall bulk statistics shown in Table 4.1. This table shows the value of the standard deviation of the error, σ_{error} , and the probability of a measurement being an outlier of that distribution, $P_{outlier}$. This probability of a measurement being an outlier is the probability that the calculated DoA is two standard deviations outside of the expected distribution and is used to help classify the characteristics of the RSS pattern combined with the DoA method.

Table 4.1: Results of performance of each DoA determination method on the physical rotation dataset showing both the standard deviation of the error, σ_{error} , and the probability of an outlier measurement, $P_{outlier}$.

Method	σ_{error}	$P_{outlier}$
Max Value	15°	0.08
Cross Correlation	11°	0.10
Lobe Detection Based DoA	12°	0.05

One of the key pieces of information gleaned from these extensive flight tests is the fact that all the methods break down when nearly directly overhead of the signal source and the RSS pattern has lost the directionality needed for these methods to work – that is when the measured RSS pattern resembles that of Figure 3.21a. Since the antenna is mounted at a fixed polar angle and is unable to rotate in the polar direction, when directly overhead of the signal source the directionality of the antenna is lost (at any given azimuth angle, the gain of the antenna is identical due to the geometry and the beamwidth of the antenna). Adjusting the results to reflect a possible flight strategy that attempts to avoid flying directly overhead of the signal source (for the 100m flight altitude used in these tests that distance was deemed to be 58m away horizontally from the source [48]) returns statistics in Table 4.2.

Table 4.2: Results of performance of each DoA determination method after removing rotations from nearly directly overhead of the signal source.

Method	σ_{error}	P_{outlier}
Max Value	15°	0.00
Cross Correlation	11°	0.04
Lobe Detection Based DoA	12°	0.00

In this case, both the maximum value method and the LDB-DoA method did not have any significant outliers in the measurements – that is, both algorithms are able to accurately determine the signal DoA from all the RSS patterns. However, the cross-correlation method still fails 4% of the time. This 4% of the time are the patterns that are generated at a large distance from the signal source due to the fact that as the distance between the source and the vehicle increases, the shape of the RSS pattern degrades – that is, it no longer resembles the calibration pattern (or the expected RSS pattern shape from the true gain pattern).

For the complete study, it is important to understand the ability to potentially remove that 4% of outliers from the cross-correlation method as the cross-correlation method allows an integrity check by using the cross-correlation coefficient. For this dataset, the integrity check presented in Section 4.3.2 for the LDB-DoA algorithm is used and the cross-correlation coefficient based integrity check from [32] is also used. The results of those integrity checks are seen in Figure 4.9 and Figure 4.10. For the cross-correlation coefficient method, Figure 4.9 shows the computed coefficients for all the RSS patterns in the dataset with those considered as outliers circled in orange. The threshold used is the threshold value from [32] of 0.3. With that threshold value, while almost all the outliers are properly removed, so are 12% of the non-outlier measurements.




Figure 4.9: Effectiveness of using the cross-correlation coefficient as an integrity check of the measured DoA.

For LDB-DoA, Figure 4.10 shows the computed integrity check value using the method presented in Section 4.3.2 for the directional antenna and the outliers are circled. Using a threshold of 0.9 results in all but one of the outliers being rejected and only 2% of the non-outlier measurement being rejected.




Figure 4.10: Effectiveness of using the custom integrity check based on a model of the antenna.

The integrity check possible with the LDB-DoA algorithm greatly outperforms the integrity check using the cross-correlation coefficient in the sense that there are significantly fewer good measurements discarded which is detrimental to a localization approach using a physically rotated antenna due to the time cost of a rotation. Since the time cost of a rotation is so high, having to throw away a perfectly good measurement is a steep penalty.

4.4.2 3-element Array

A similar analysis is also performed with a dataset of over 2,000 RSS patterns generated by the 3-element phased array antenna. For these results, RSS patterns were generated at Stanford University and Edwards Air Force Base at a range of distances from directly overhead to 250m away as depicted in Figure 4.11. In this case, the furthest distance is much lower than the furthest distance for the physical rotation case because the gain of the 3-element phased array antenna is a lot less than the gain of the directional antenna used in the previous section. Therefore, despite the smaller range, this dataset is composed of similar types of RSS patterns (i.e. there are degraded patterns in this dataset).




Figure 4.11: Histogram of the distance from the source for the dataset of RSS patterns generated by the 3-element phased array antenna.

For these results, is important to remember that the RSS pattern from this antenna can take on one of several different types of shapes based on the incident angle of the signal source. As a result, for the cross-correlation method, instead of using a specific one of the possible RSS patterns, the calibration pattern used was that of a single lobe which has been shown to perform well for irregular

RSS patterns [32]. Once again, the statistics on the performance of the different methods are shown in Table 4.3.

Table 4.3: Results of performance of each DoA determination method on the 3-element phased array RSS patterns.

Method	σ_{error}	P_{outlier}
Max Value	19.8°	0.28
Cross Correlation	26.5°	0.33
Lobe Detection Based DoA	19.0°	0.14

Immediately, this table shows that across the board the standard deviation of the error on this antenna is larger than it was for the stop and rotate directional antenna presented in the previous section. This is a result of the beamwidth on the 3-element phased array antenna being larger than the beamwidth of the directional antenna, and, in general, the ability to determine the DoA is a function of the beamwidth of the antenna [32]. Additionally, the probability of an outlier has also increased, quite significantly in the case of the cross-correlation method, as a result of the multiple lobes that can be present in the RSS pattern. Looking at the results a little differently, Figure 4.12 shows the distribution of the DoA error for each of the methods across the entire dataset. Of note is that the distributions for the cross-correlation method, CC, and the maximum value method, Max, exhibit fatter tails than the LDB-DoA distribution. Furthermore, the cross-correlation and maximum value distributions have a slight peak near the 180° error mark.

The increase in the error out towards 180° is an artifact that both the maximum value and the cross-correlation methods inherently assume that there can only be a single lobe in the pattern – and therefore only a single source in the environment. In the example pattern in Figure 4.13, the dashed red line shows the true DoA, the dark red line shows the maximum value DoA, the blue line shows the cross-correlation DoA, and the purple and yellow lines show the lobe detection DoA values along with the beamwidth as determined by the algorithm. Since the maximum value and the cross-correlation methods assume only a single option, both have mistakenly identified the strong back-lobe as the true lobe. LDB-DoA, as it makes no assumption on the number of possible lobes, returns, in the case of the RSS pattern in Figure 4.13, two possible options for the DoA and as a result can significantly outperform the other two methods.




Figure 4.12: Distribution of DoA values measured using the cross-correlation (blue), the maximum value (red), and the LDB-DoA (yellow) method on the RSS patterns generated by the 3-element phased array antenna.




Figure 4.13: Example RSS pattern with large back lobe causing the cross-correlation (blue) and the maximum value (red) estimates to point in the wrong direction (true DoA is the dashed line) while the LDB-DoA algorithm presents two possible options.

When multiple lobes are present, LDB-DoA can successfully determine the bearing for each of the lobes and provides a lot more information for a given measurement than the other methods. Additionally, since there is no assumption on the exact pattern shape (just the assumption that the

antenna gain pattern follows standard gain pattern laws for a directional antenna), it does not matter how much the pattern deteriorates from the theoretical pattern, LDB-DoA is still capable of providing an accurate bearing to the source.

When it comes to using an integrity check with these patterns, the second check that is presented in Section 4.3.2 is used. For these patterns, the integrity check ends up being conservative with a lot of measurements discarded, but due to the much higher rate of RSS patterns, and therefore DoA measurements, the estimation algorithm (presented in Chapter 5) is still able to perform well. Once again, the cross-correlation method does have the computed cross-correlation coefficient that can be used for integrity checks. However, once the RSS pattern starts to have multiple lobes, the effectiveness of that coefficient is deteriorated and cannot be used directly.

4.4.3 Multiple Sources

The previous two sections compared the LDB-DoA method against the maximum value and cross-correlation methods in single-source real-world environments. This section looks towards a completely different environment for which the maximum value and the cross-correlation methods are unusable – multiple sources. The dataset used throughout this sub-section contains RSS patterns generated in a simulated environment using models of the behavior of the phased array antenna.

As a simulated environment is used, this dataset contains RSS patterns generated not only using the model of the 3-element phased array antenna but also models of a 7- and 19-element array. These larger arrays provide a unique context in understanding the possibilities of the LDB-DoA algorithm in more complex environments. This helps to illustrate some of the changes in RSS pattern characteristics that are driven by the size of the antenna array. Looking at the RSS patterns for each of the 3-, 7- and 19-element arrays in Figure 4.14, it can be seen that as the number of elements in the pattern increases (array gets larger) the beamwidth narrows and the sidelobe suppression increases. These characteristics can be very desirable when exploring environments with multiple signals, hence the comparison against these larger antennas. As the number of elements increases from 3- to 19-elements, the beamwidth of the main beam narrows from 60° to 20° and the maximum gain of the antenna increases by almost 8dB. Additionally, as the number of elements increases, the sidelobe suppression – the difference between the main lobe and the large side lobe – also increases. For example, the sidelobe suppression for the 19-element array is 18dB while the sidelobe suppression is only about 5dB for the 3-element array. Three different sizes are depicted to illustrate that moving from the 3-element to the 7-element array nets a much bigger change in performance than moving from the 7-element to the 19-element array (namely the 7-element array's side lobe suppression is 15dB and provides 3dB higher gain). This potentially

motivates the desire for moving from a 3-element to a 7-element array, but the move all the way to 19-elements may be unnecessary in most cases.

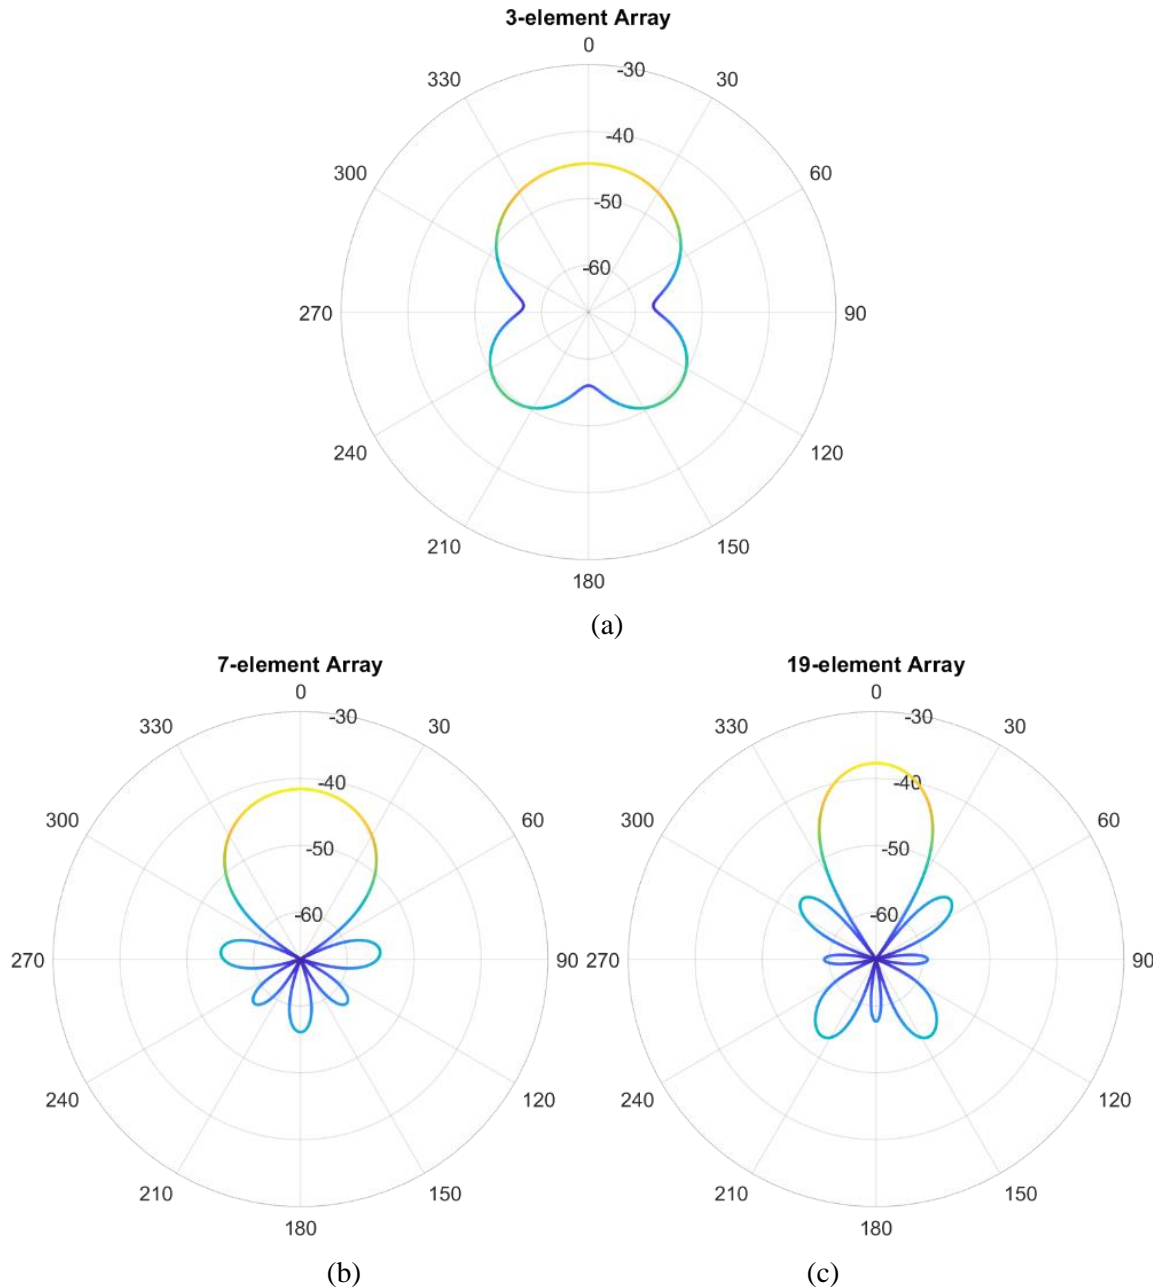


Figure 4.14: Theoretical gain pattern of unsteered (a) 3-, (b) 7-, and (c) 19-element arrays

Multiple sources in an environment, using an RSS pattern based approach for DoA, has more challenges than just the fact that there can be multiple DoAs. The two key additional challenges are the fact that side lobes of an antenna gain pattern can distort the RSS pattern in such a way that

the RSS pattern does not point in the true source direction and the fact that the beamwidth of the antenna gain pattern can result in multiple sources to not be identifiable as the resulting RSS pattern blends the result of the multiple signals into a single lobe. For all the patterns, the exact same LDB-DoA algorithm is used to further demonstrate the versatility of this algorithm.

The large side lobes of the 3-element phased array antenna can have large effects on the resulting RSS pattern when multiple interference sources are present, creating biases in the lobe directions depending on the incident angle of the source, ϕ_{source} . This effect is illustrated in Figure 4.15a which shows the resulting RSS pattern from having two interference sources at $(\phi_{s1}, \phi_{s2}) = (0^\circ, 90^\circ)$. Due to the large side lobes for the 3-element array, the resulting lobes in the pattern does not point at $(\hat{\phi}_{s1}, \hat{\phi}_{s2}) = (0^\circ, 90^\circ)$, but rather at $(\hat{\phi}_{s1}, \hat{\phi}_{s2}) = (45^\circ, 120^\circ)$. This distortion based bias needs to be resolved either through using models of the theoretical antenna performance or larger arrays. While the bias is large for some incident angles, the 3-element array can also return an RSS pattern with no bias illustrated in Figure 4.15b when the sources are incident at $(\phi_{s1}, \phi_{s2}) = (90^\circ, 270^\circ)$.

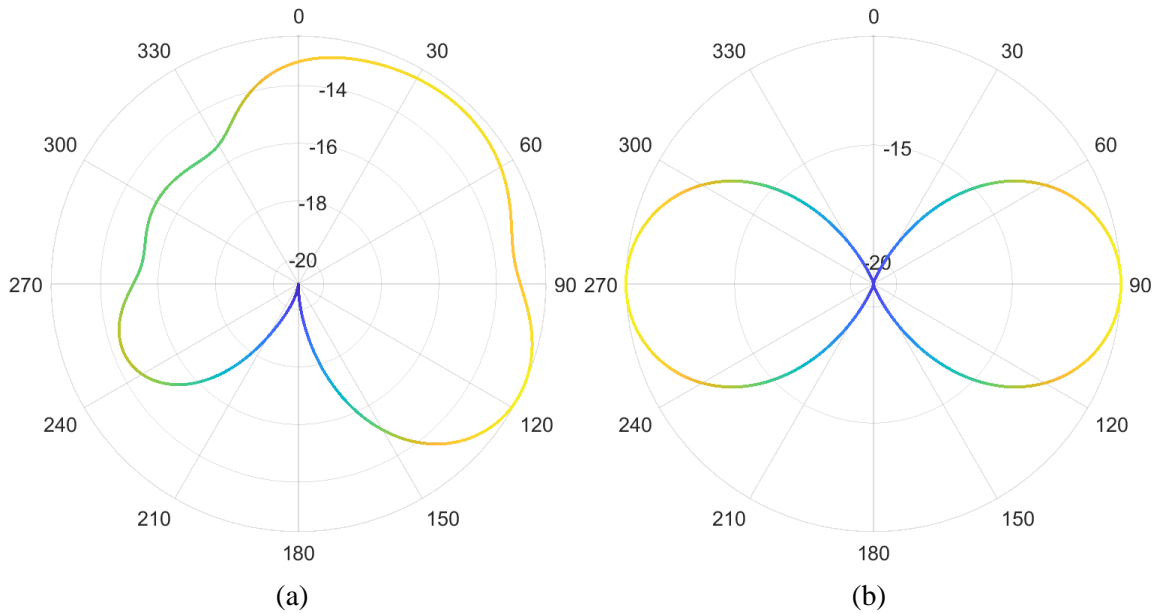


Figure 4.15: Simulated RSS patterns for the 3-element phased array antenna with sources at (a) $(\phi_{s1}, \phi_{s2}) = (0^\circ, 90^\circ)$ and (b) $(\phi_{s1}, \phi_{s2}) = (90^\circ, 270^\circ)$ demonstrating the distortion effect due to large side lobes.

As mentioned, one approach to reduce the distortion due to the side lobes is to use an array with a larger number of elements. Returning to the situation with 2 sources at $(\phi_{s1}, \phi_{s2}) = (0^\circ, 90^\circ)$, the resulting pattern for the 7- and 19-element arrays are depicted in Figure 4.16. The

sidelobe suppression of the 7-element array is already enough to remove all bias. The 19-element array, having a much smaller beamwidth, creates a set of even more defined lobes pointed in the direction of the sources and does not have any significantly large side lobes.




Figure 4.16: Simulated RSS patterns for two sources at $(\phi_{s1}, \phi_{s2}) = (0^\circ, 90^\circ)$ for array sizes of (a) 7-elements and (b) 19-elements.

The RSS pattern from an array can also result in biases when the sources are angularly close to each other. Specifically, when $(\phi_{s1} - \phi_{s2}) < 2BW$, where ϕ_{s1} and ϕ_{s2} are the incident angles of the two sources in the ϕ direction and BW is the beamwidth of the steered antenna. When the sources are that close to each other and have the same power level, the main lobe blends together, resulting in one larger main lobe pointed about halfway between the two sources. Figure 4.17 illustrates this case for $(\phi_{s1}, \phi_{s2}) = (0^\circ, 40^\circ)$ for the (a) 3-, (b) 7-, and (c) 19-element arrays. The 40° angular separation is less than $2BW$ for all three arrays so all the arrays have distorted main lobes with a central peak around 20° .

To evaluate the effectiveness and abilities of the different antenna array sizes, a $3\text{km} \times 3\text{km}$ world is simulated with two different sets of source locations (discussed separately in Source Set A and Source Set B). The metrics used for comparison are the ability to resolve a bearing to each single source, the number of clutter measurements, and the error of the bearing measurement due to the bias of the antenna. For this evaluation, a DoA is associated with a source if the error was less than 20° . In each simulation, the sources have the same power levels (1W) and the same

antenna gain (2dBi). The only losses accounted for in the simulation are freespace path losses. A theoretical model of each the 3-, 7- and 19-element array is rotated at 50 meters intervals over the entire world. For each position, the antenna is held at a constant altitude of 100 meters and the antenna is steered to $\theta_{steer} = 60^\circ$ and $\phi_{steer} = [0, 360]$ in one-degree intervals.

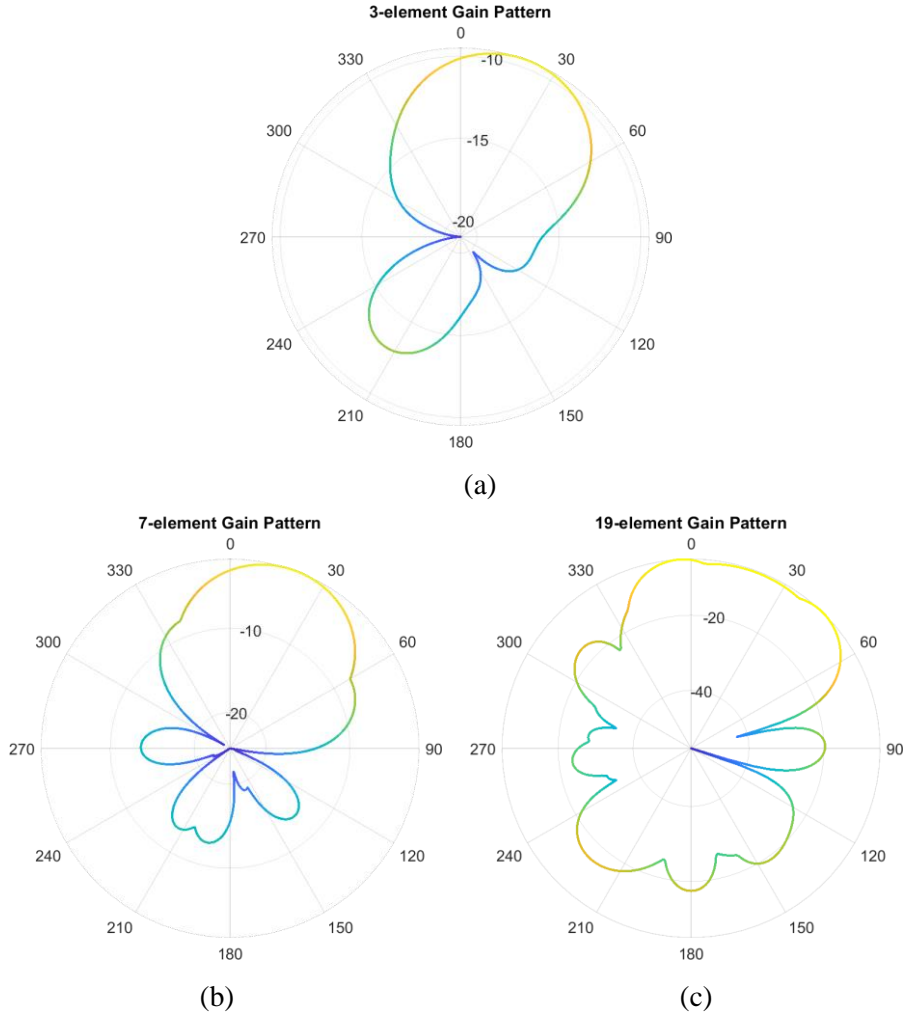


Figure 4.17: RSS patterns for (a) 3-element, (b) 7-element, and (c) 19-element arrays with sources at $(\phi_1, \phi_2) = (0^\circ, 40^\circ)$

The following presents the results for two different environmental conditions: source set A contains two sources placed at -400m and 400m from the origin along the East axis and source set B contains two sources placed at -200m and 200m from the origin along the East axis. The placement is such that, for source set A, the signal source goes out of range of the antenna before

the second condition of obscurity occurs – where $(\phi_{s1} - \phi_{s2}) < 2BW$ – and for source set B the sources are still in range of the antenna when that second condition of obscurity occurs.

4.4.3.1 Source Set A

Figure 4.18 illustrates the results for the 3-element array. Figure 4.18a shows the visibility of each source from each location; Figure 4.18b shows the number of clutter observations, or bearing measurements associated with neither source (i.e. a side lobe large enough to be detected as the main lobe); Figure 4.18c and Figure 4.18d show the bearing error in the extracted bearing as a result of a bias in the pattern. Immediately it can be seen that the 3-element antenna results in cluttered measurements due to the very large side lobes of the RSS pattern. Nearly 40% of the positions where the 3-element antenna was capable of detecting one of the signal sources resulted in a pattern with more than one detected lobe. Furthermore, those large side lobes also cause biases in the bearing results which are depicted in Figure 4.18c and Figure 4.18d. These biases result in errors that are up to 15° before the source becomes out of range of the sensor. Another notable result is the fact that the overlapping region does not maintain a circular shape that would be expected since the only loss considered is the freespace path loss. This is as a result of the fact that as one signal source becomes strong enough, the lobe to the weaker signal drops below the 3dB threshold of the maximum signal strength received causing it to not be detected, illustrated in Figure 4.19. In this example, only the larger lobe is selected as a possible bearing and the smaller lobe is missed. The second notable result is the interesting pattern in the clutter due to the distortion of the pattern is very similar to the pattern seen in the Source Set A results. This again just accentuates the difference in the RSS pattern given the incident angle of the source with some directions resulting in larger side lobes than others. This eventually can potentially be integrated into the flight controller to rotate the antenna in such a way that the signal source is incident on the antenna in a direction that minimizes clutter.

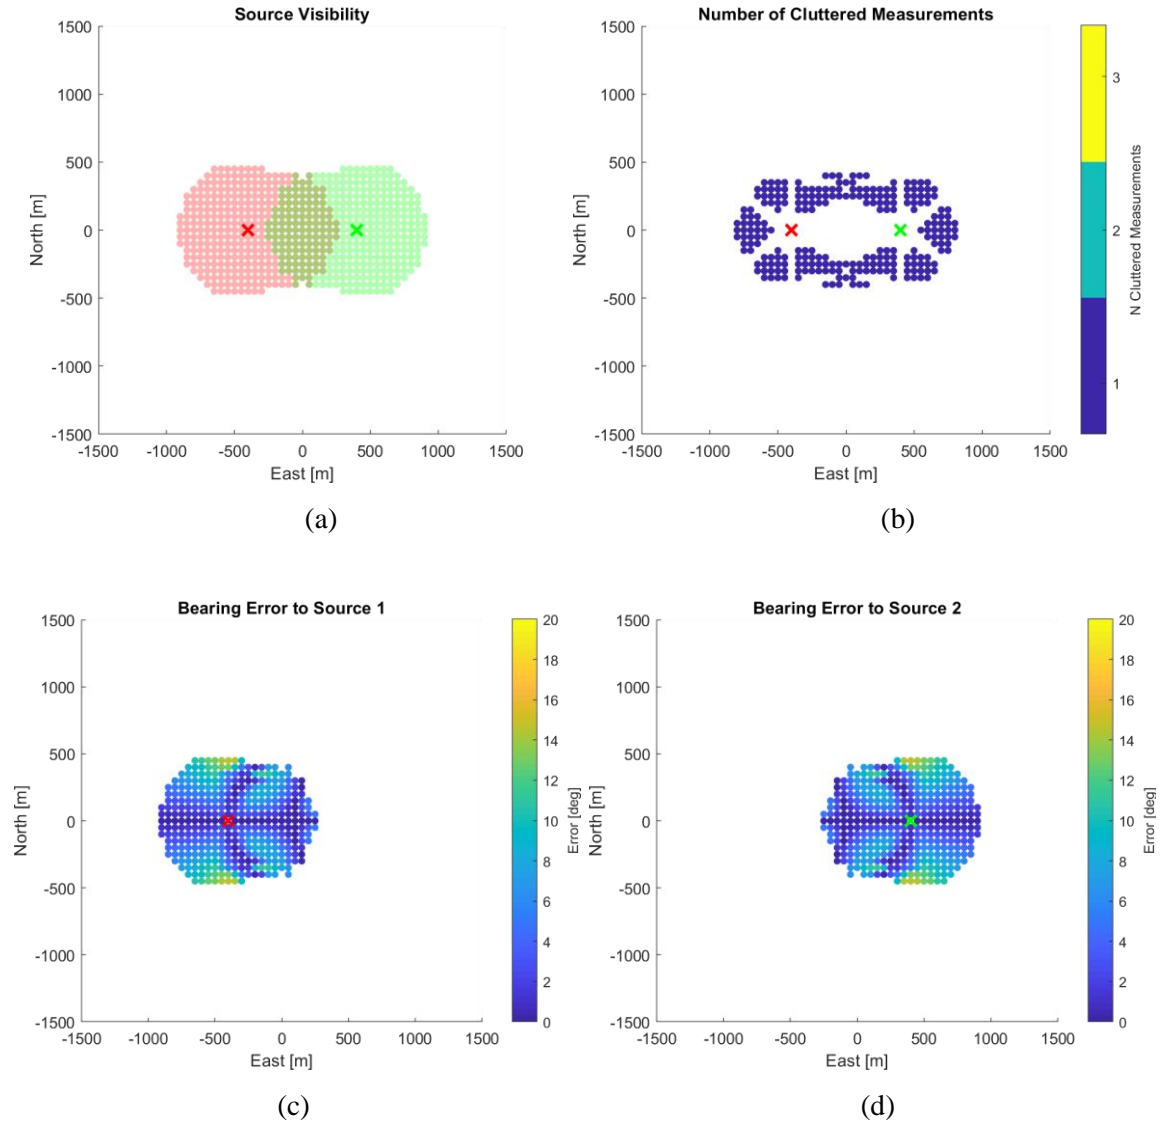


Figure 4.18: Bearing simulation results for a 3-element antenna with 2 interference sources at (0, -400m) and (0, 400m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source




Figure 4.19: Example simulated RSS pattern with two sources such that the power difference between the two results in only one signal being detected by the LDB-DoA algorithm.

The same results for the 7-element array are shown in Figure 4.20. For the 7-element array, there are no clutter measurements at any location on the grid. The 7-element side lobe suppression is much greater than that for the 3-element array resulting in much smaller side lobes in the created RSS patterns leading to the significant reduction in the number of clutter measurements. Looking back to the requirements for a localization algorithm and the challenges of data association, increasing the array size from 3 to 7 can greatly affect the overall performance due to the reduced amount of clutter. The smaller side lobes also reduce the distortion of the pattern and bias in the bearing measurements such that the worst-case error in the RSS patterns bearings is less than 5° . However, it can also be seen that the region of overlap, where both sources can be seen, is also significantly reduced. Since the 7-element array has a higher gain than the 3-element array, the range to the source before the weaker signal is no longer detectable is reduced, resulting in fewer measurement locations that can provide information to both sources. Finally, as expected, the higher gain of the 7-element array results in a larger range of visibility of each of the sources.

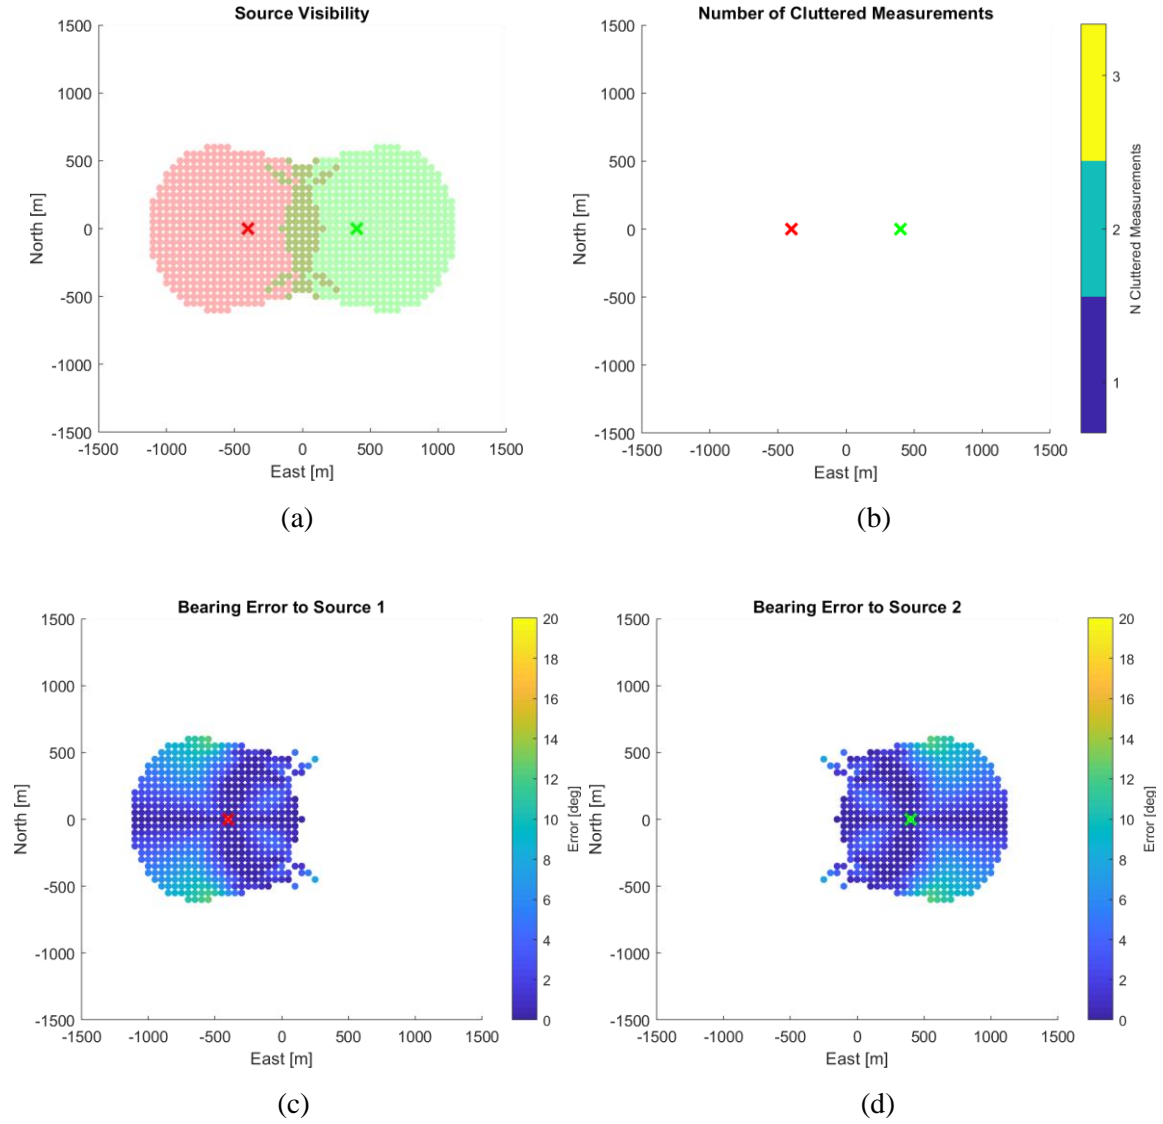


Figure 4.20: Bearing simulation results for a 7-element antenna with 2 interference sources at (0, -400m) and (0, 400m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source




Figure 4.21: Visibility of each signal source using a 19-element phased array antenna.

The 19-element array maintains very similar error performance as the 7-element array with the main difference being the additional range provided due to the higher gain, seen in Figure 4.21, showing the visible signals from the antenna. There is also a continued reduction in the central area for which both sources can be detected within the RSS pattern using the LDB-DoA algorithm.

4.4.3.2 Source Set B

Looking at the results for the 3-element array in Figure 4.22, the effects of the angular proximity of the sources is apparent in the visibility plot for each source (Figure 4.22a) as there is a region in the center that previously had overlapping measurements that is now marked in gray and only seen in the clutter. The gray marking indicates that a DoA was computed but did not meet the error threshold for association and therefore is not considered to associate with either source and is only clutter. For these simulations, a bearing error of greater than 20° was considered to not be associated with the source, which is why those measurements now appear in clutter and not in the visibility plot. An RSS pattern in this region looks similar to the RSS pattern depicted in Figure 4.23 where the main lobe is stretched and the bearing is estimated to be the midpoint of the two different sources. Finally of note is the much larger errors in the calculated DoA due to the strong distortion that results from angularly close sources.

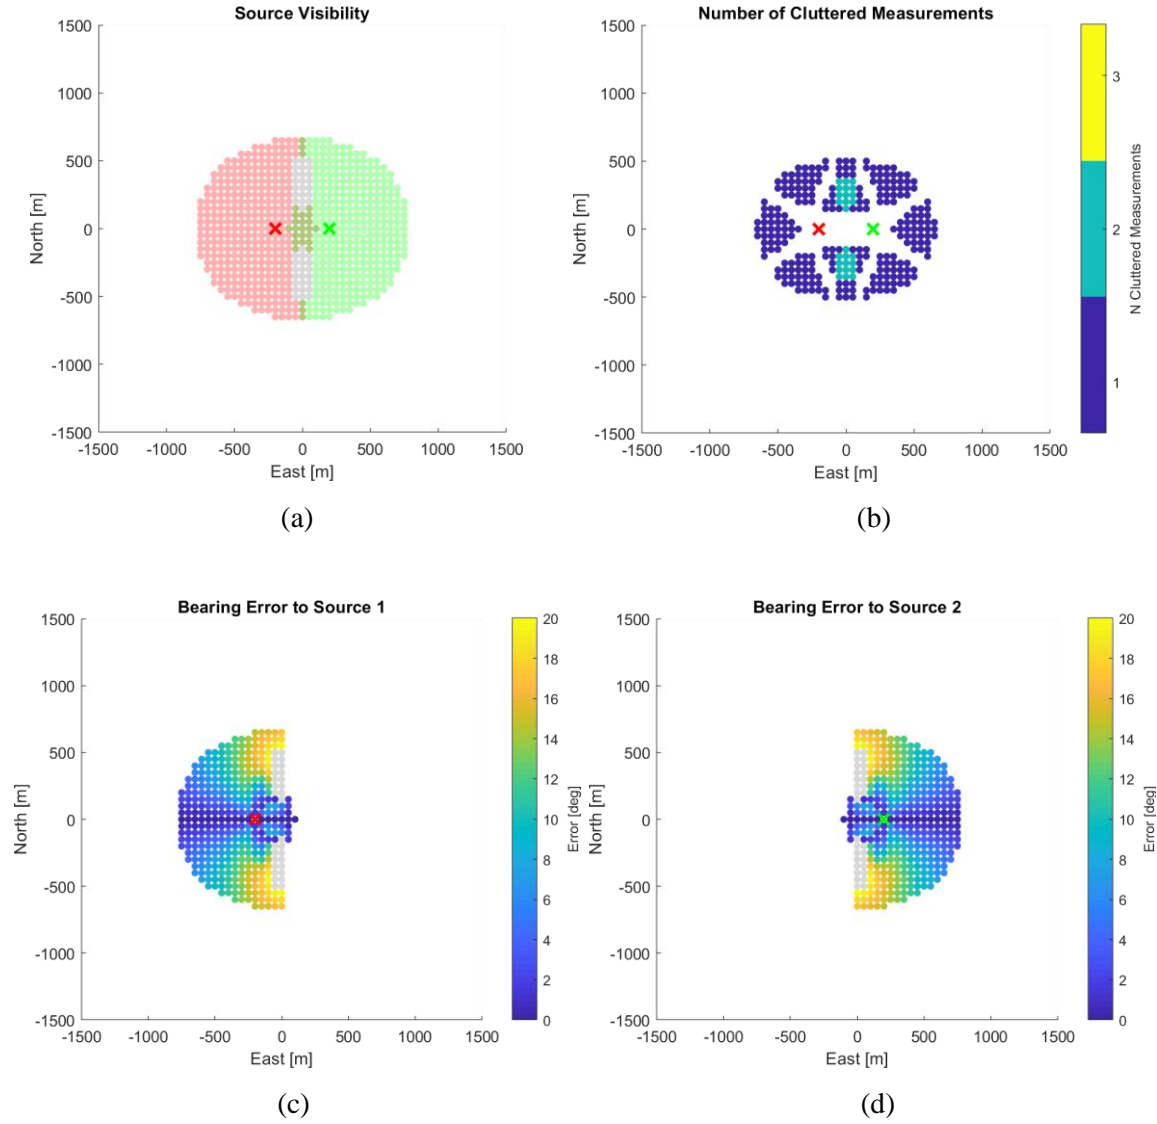


Figure 4.22: Bearing simulation results for a 3-element antenna with 2 interference sources at (0, -200m) and (0, 200m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source




Figure 4.23: Example simulated RSS pattern demonstrating two sources with an angular separation of less than $2BW$ resulting on the LDB-DoA algorithm to only detect a single lobe.

For the 7-element array, Figure 4.24, this source configuration does result in some clutter measurements due to the region that is affected by the sources being angularly close to each other. However, that region affected by the angular proximity is much smaller than the region for the 3-element array due to the narrower beamwidth and ability to support a smaller minimum angular separation. Just like with the 3-element array, there is an increase in the bearing error caused by the angular proximity in the sources. Once again, the 19-element array has a very similar performance as the 7-element array, with a much larger field of view, as depicted in Figure 4.25.

In the end, there are two main sources of possible bearing error: first due to biases of the pattern due to large side lobes, and second due to sources that are angularly close to each other. While the first source of error is controllable by increasing the number of elements to a 7-element array instead of a 3-element array, the second source of error is not controllable at the sensor level. However, knowing the existence of this error source can help inform the path planning algorithm to fly in the best trajectory possible to reduce the potential of this situation occurring.

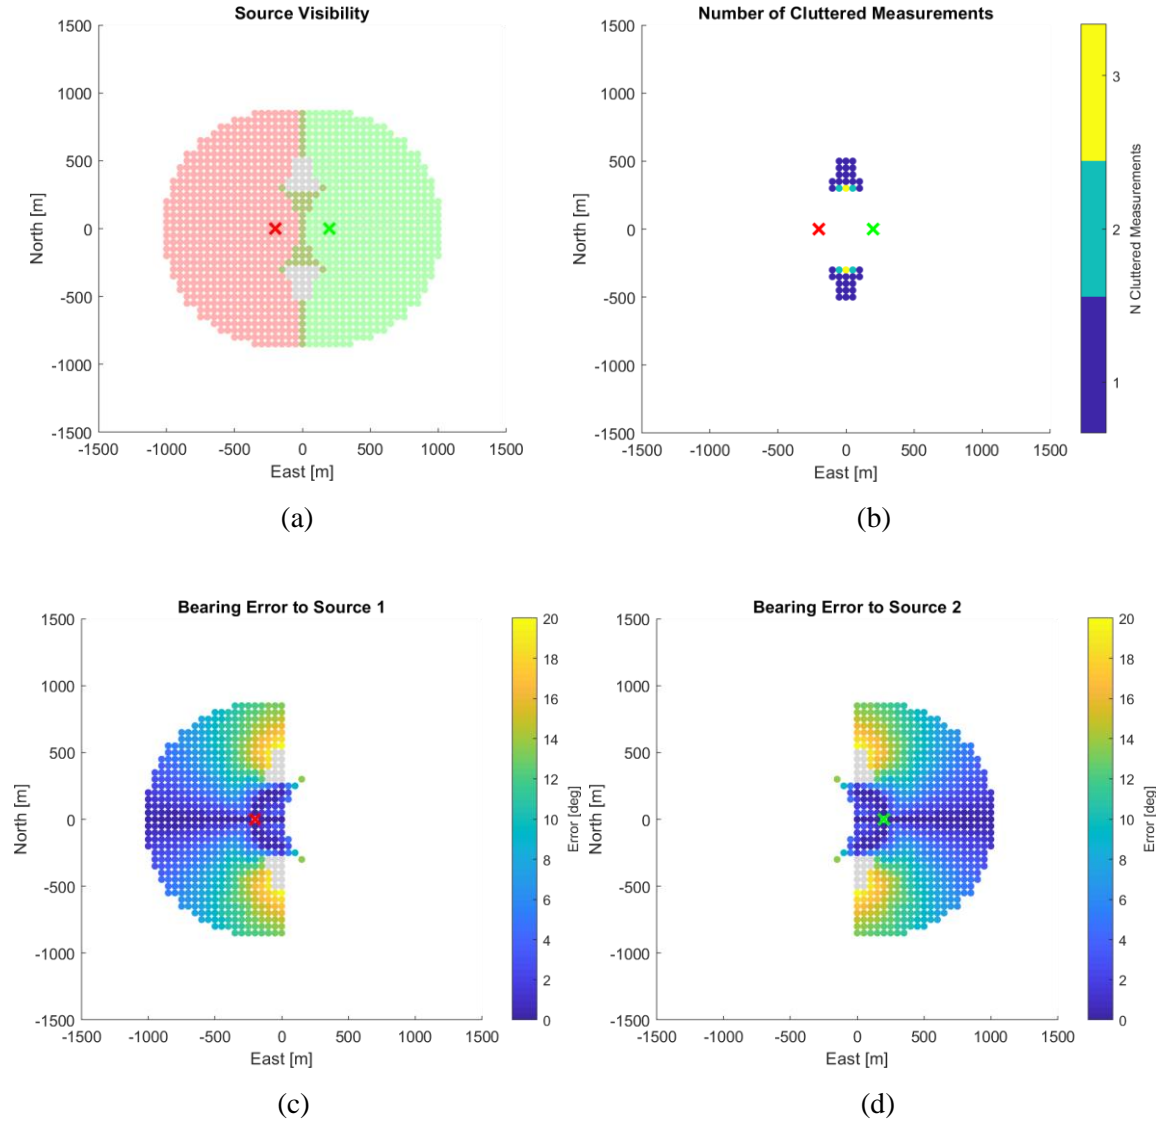


Figure 4.24: Bearing simulation results for a 7-element antenna with 2 interference sources at (0, -200m) and (0, 200m) showing: (a) the ability to observe a bearing to each source, (b) the number of bearing observations categorized as clutter, and (c) and (d) the observed bearing error for each source




Figure 4.25: Visibility of two sources with the 19-element phased array antenna.

4.5 Conclusion

This chapter presented the second contribution of this thesis: the development of LDB-DoA, a robust algorithm to determine DoA from an RSS pattern. The algorithm presented proved to outperform existing methods of DoA determination for both a physically rotated directional antenna and the 3-element phased array antenna in a wide variety of environments. Furthermore, the algorithm does not require a calibration pattern, or really any more information than the fact that the pattern was generated by a directional antenna. The unique parameterization that results from this algorithm means that the entire RSS pattern is not merely reduced to a single value (such as the max value method) or just a simple tuple that makes assumptions on the measured pattern maintaining the shape of the expected pattern (the cross-correlation method), but rather it captures the information on the lobes of the pattern and therefore the underlying shape of the pattern along with a DoA observation. This representation enables the use of integrity checks based on a user-defined model of the antenna that does not impact the LDB-DoA algorithm's performance or behavior on the system, unlike the cross-correlation method that inherently assumes a model of the antenna in order to be able to determine a DoA from an RSS pattern.

The LDB-DoA method has also been shown to successfully extract bearing information from RSS patterns due to multiple RFI sources and the performance of several array sizes has been characterized and evaluated. The simulation results presented are the first step in being able to make

an informed decision on the localization algorithm to use and the path planning method to use. Understanding the sensor is key to being able to get the best performance out of the sensor and get the best estimates of the RFI sources at the lowest cost (time to localization).

Comparing the results of the three different antenna sizes, summarized in Table 4.4, some conclusions of the combined performance of the LDB-DoA method and each array size can be made. Overall the largest design tradeoff that needs to be made is the tradeoff between array size (weight and power) and small biases/low clutter. Using the 3-element array, the localization algorithm needs to be very robust to clutter and a good model of the sensor should be used to account for the biases that are present based on the source locations. While some biases can be addressed by cleverly choosing the flight path (e.g. flying in such a way to maintain a signal at a specific incident angle), some of the errors will need to be accounted for in the localization algorithm. Using a 7-element array comes at a cost of size, weight, but provides significant reduction of clutter and bias errors and partially reduces the “no-fly area” that would result in the sources being too angularly close to each other. The 19-element array does not provide a significant enough advantage over the 7-element array in these metrics to be worth the significant additional monetary and weight costs.

Table 4.4: Summary of key results from 2-source simulations for each array size

	Max Bearing Bias (far)	Max Bearing Bias (close)	Min Source Angular Distance	Clutter
3-element	15°	20°	150°	High
7-element	5°	20°	70°	Low
19-element	4°	20°	30°	Low

Chapter 5

Sensing Platform Development

This chapter presents the first part of the final contribution of this thesis: the design and development of a flight test platform capable of 24/7 operations – the Jammer Acquisition with GPS Exploration and Reconnaissance (JAGER) platform. One of the major benefits of this flight test platform is its ability to provide real-world data and results to present throughout this thesis (as was seen in the previous chapters).

Section 5.1 provides a detailed description of the complete system design to autonomously complete the signal localization mission. From there, Section 5.2 discusses the details of the IR based navigation system and puts it in the scope of the state-of-the-art when it comes to navigating in a GNSS denied environment. Section 5.3 then presents the complete localization system with a focus on the estimator used to localize the signal source and the path planner/concept of operation for the complete autonomous flight to execute the desired mission. Section 5.4 presents the simulated performance of the complete system in a large-scale environment. Finally, Section 5.5 provides some concluding words on the design and construction of the entire system.

5.1 JAGER System Overview

This prototype-level system, JAGER, is designed to demonstrate an end-to-end solution for a system that can be easily deployed in an airport environment to rapidly localize the source of GNSS interference. From takeoff to landing, the system must be able to autonomously operate without user intervention to localize the signal source using a navigation system that operates independently of GNSS. This means that the localization system must be capable of estimating the location of the

signal source given the series of DoA measurements, executing a desired trajectory for localization, and determining the signal source is deemed as “found” within the desired performance limits of the system.

JAGER includes two core custom sub-systems: the localization system responsible for localizing the signal source, and the navigation system responsible for enabling 24/7 operations in GNSS denied environments. Each of these custom systems is described separately, however, as depicted in Figure 5.1, these sub-systems work together to accomplish the mission, and must be able to operate with a limited amount of computation capability that is available on a UAV. In this case, these sub-systems were integrated utilizing an Intel Core-i3 Compute Stick that would be capable of being flown on much smaller UAVs than the DJI S1000 used in this research. JAGER also includes an off-the-shelf autopilot system responsible for executing the desired command and control for the vehicle in addition to these custom sub-systems.

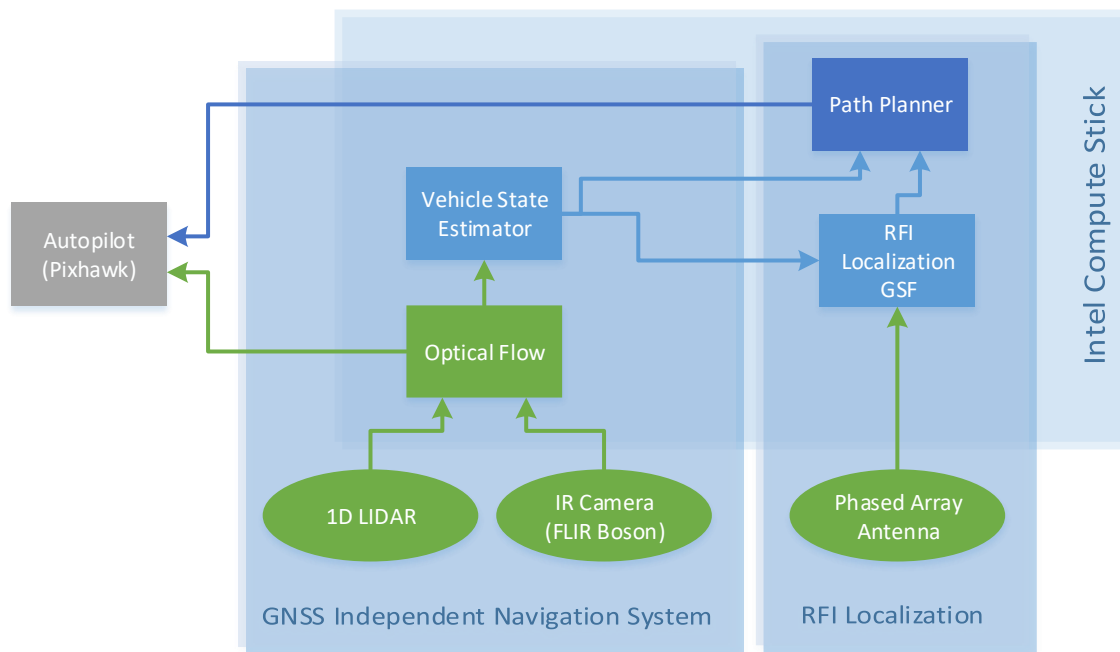


Figure 5.1: Schematic of full JAGER system.

The localization system revolves around the 3-element phased array antenna described in Chapter 3 and the Lobe Detection Based Direction of Arrival (LDB-DoA) algorithm described in Chapter 4. The final two pieces of the localization system, described in further detail in Section 5.3, are the estimator – responsible for estimating the location of the signal source – and the path planner – responsible for determining the vehicle’s trajectory in real-time to best localize the signal source.

The LDB-DoA algorithm described in Chapter 4 is applied to the RSS pattern created by the 3-element phased array antenna of Chapter 3 to create the DoA measurements for the estimator to estimate the location of the signal source. These estimates are used by the path planner, along with the current estimate of JAGER's position determined by the navigation system, to compute the desired velocity vector that results in JAGER's curved trajectory around the signal source. This desired velocity vector is then sent as a command to the Pixhawk autopilot, which handles the actual control of the drone, to actualize the flight path determined by the path planner.

Critical to the success of such a system is the ability to operate in a GNSS denied environment. While some domains are quite familiar with the challenge of operating without GNSS, such as underwater applications that have never had the benefit of GNSS, UAVs, and ground vehicles have often taken GNSS for granted as a core sensor in navigation systems. However, for this specific application, since the signal source of interest is a GNSS interference device, the luxury of using GNSS for navigation cannot be assumed.

This thesis, therefore, presents a GNSS independent navigation system, which is built around a gimbal-stabilized infrared (IR) camera to perform monocular visual odometry using optical flow – described in more detail in Section 5.2 – that is capable of operating in both daytime and nighttime environments. The optical flow computed from the IR imagery feeds directly into the Pixhawk autopilot, due to limitations of the PX4 firmware not being able to operate without either optical flow or a GNSS position, and into an internal estimator described in Section 5.2.2 that provides position information for both the estimator and the path planning portions of the localization system.

The autopilot used for the full system implementation is the same as the autopilot used during data collection flights presented in Chapter 2. The only difference is that the autopilot, for the complete JAGER system, utilizes velocity commands generated by the path planning system as the core command to be executed in flight.

5.2 Navigating in a GNSS Denied Environment

The first JAGER sub-system described in detail is the GNSS independent navigation system. The first sub-section presents a perspective on the world of GNSS independent navigation today along with related works using IR imagery for navigation. From there, the second sub-section presents the final algorithm implemented onboard the JAGER system to navigate in a GNSS denied environment using IR imagery from a gimbal-stabilized camera. Finally, the last sub-section presents a small sample of real-world flight test data on the performance of the optical flow

system's capability of measuring the velocity from the IR imagery. A more detailed analysis of the complete performance of the system in simulation is presented in Section 5.4 and an analysis of the performance of the system in real-world flights is presented in Chapter 6.

5.2.1 Related Work and Motivation

Researchers have been interested in the problem of GNSS independent navigation to handle a variety of environments for which GNSS is naturally impossible to receive. These environments include underwater environments and even urban centers where GNSS can be unreliable due to the urban canyon effects with a variety of different approaches [67]. Some research leverages the many signals of opportunity that exist in the terrestrial environment today including cell towers [68], television station radio signals [69], and even other satellites in orbit around Earth that broadcast a known signal structure [70]. However, such systems that merely rely on other types of broadcasted signals are still susceptible to interference, especially from a particularly nefarious attacker.

A group of more “self-centered” approaches to GNSS independent navigation relies on sensors such as cameras, light detection and ranging (LIDAR) sensors, inertial navigation systems (INS), and the like. However, many of these approaches have significant drawbacks. For example, an INS based approach tends to have prohibitive levels of drift over time or can be prohibitively expensive [67]. As another example, LIDAR based methods have a high computation cost requiring extensive computation resources which is problematic for real-time applications, especially onboard small UAVs [67]. Cameras, on the other hand, have proved to be a versatile sensor capable of a wide range of performance depending on how the imagery is used (e.g. with simultaneous localization and mapping – SLAM – where the imagery is used to build a map of the environment and localize the vehicle within that map or visual odometry) [67]. Using cameras to perform visual odometry (VO) provides a reasonable balance for complexity, cost, and performance needed, making it a compelling solution for a low cost, yet reliable, approach for GNSS independent navigation [71]. A camera-based method is therefore used to accomplish GNSS independent navigation in this thesis, and prior VO work is thus discussed in more detail in this sub-section.

VO is designed to recover the position of the vehicle incrementally by using the most recent information to update the position of the vehicle based on the distance traveled between the last known measurements. VO techniques themselves break down into two main categories: stereo VO and monocular VO. Quite simply, stereo VO leverages two synchronized cameras pointed in the same direction to resolve the depth of features within the image (much like human vision), while monocular VO uses only a single camera and therefore tends to be a cheaper and lighter solution with the added complexity requiring that the depth (or scale) ambiguity be resolved [72]. Stereo

VO requires more calibration than monocular VO and requires stringent synchronization between the cameras, resulting in a potential for larger errors over time if either of those conditions fails [73]. Despite the added challenge of the scale ambiguity, monocular VO has been a popular technique employed on anything from underwater robots [74] to ground robots to UAVs [67], [73] with position errors on the order of 1-5% of the distance traveled [67]. However, this heavily researched domain has been limited to the use of cameras that work in the visible spectrum, making them ideal for many common applications, but are inherently incapable of working during nighttime environments.



Figure 5.2: Images¹ captured from the same location at Edwards Air Force Base at 11 pm with (a) a regular camera and (b) an IR camera.

Expanding the operating window of vision-based methods to nighttime operations turns the attention towards leveraging the capabilities of infrared (IR) cameras to be able to “see” at night, as illustrated in Figure 5.2. While it may seem that working with IR images is nothing more than dealing with grayscale images, IR cameras and the imagery itself poses several critical challenges when it comes to adapting existing VO methods due to the fact that the camera is fundamentally sensitive to the temperature of the objects in view. As a result, the first uses of IR cameras for GNSS independent navigation used the IR imagery as a secondary sensor to augment regular vision camera imagery and other onboard sensors [75], [76].

The first challenge posed by IR cameras is their calibration. Common calibration techniques for cameras rely on the use of either a black and white chessboard pattern or a specialized pattern

¹ DISTRIBUTION STATEMENT A. Approved for public release; Distribution is unlimited 412TW-PA-19478

with colored circles. While these patterns work well for cameras sensitive to color, in the infrared domain, the entirety of a sheet of paper looks the same, resulting in the need to adapt existing patterns to be seen in infrared through the use of specialized metal versions of the patterns as in [77] or by creating a temperature difference between the squares of the chessboard on the paper using a heat lamp, as done in this thesis. IR cameras are already notoriously difficult to calibrate due to traditional calibration methods (colored circles or the chessboard pattern) being geared towards color or grayscale cameras that are sensitive to color, not temperature – basically to an IR camera, all traditional calibration methods look alike to the camera (a blank sheet of paper) [77].




Figure 5.3: IR camera pictures² from (a) 11 pm and (b) 9 am showing the difference in color intensity based on time of day and heat of objects.

The second challenge is the fact that an IR camera measures the temperature of objects in view, and, over time, objects' temperatures change. As the temperature of the environment changes, so does the resulting intensity in the image. For a concrete example, Figure 5.3a and Figure 5.3b show an overhead view of two aircraft on the taxiway at Edwards Air Force Base at 11 pm and 8 am, respectively. Notice how the concrete taxiway has drastically changed intensity and the aircraft in the bottom left of each frame also has drastically changed intensity. This change of intensity can be equated to lighting changes that are known to be a challenge for VO based methods [73]. As a result, there is a more limited time scale – on the order of minutes – for which feature matching of existing feature sets works best for IR imagery [78]. This means that attempting to do loop closure, as is done with Simultaneous Localization and Mapping (SLAM) based methods by using the built

² DISTRIBUTION STATEMENT A. Approved for public release; Distribution is unlimited 412TW-PA-19478

map of the environment, from the beginning to the end of even a fairly short 10-minute flight may not be possible due to the degradation in the quality of the features due to temperatures changes in the environment [78].

The third and final challenge posed by IR cameras is the fact that the temperature of the sensor itself changes over time. As the sensor's temperature changes, the camera needs to recalibrate via a process that is known as a nonlinear uniformity correction (NUC) of a flat field correction (FFC) that can change the intensity of the pixels within the images before and after the calibration. Furthermore, this calibration, on a camera such as the FLIR Boson, can take up to a second, during which time no information comes from the camera [79]. Before and after a NUC the intensity of an image can change drastically, which again poses a challenge to SLAM based methods that rely on building a map of the environment and recognizing when a vehicle has returned to a known position in the map for loop closure [72].

As a result of these challenges, using IR imagery by itself for navigation has only just begun to be explored. In 2015, [77] used a forward-facing stereo IR setup on a UAV to perform stereo visual odometry in both daytime and nighttime environments with a position error of about 1% of the distance traveled. To overcome the NUC outage period, this work thermally calibrated both cameras extensively to disable the camera's standard recalibration procedure and instead used the calibrated information to allow the cameras to continuously operate [77]. In addition to the complex calibration required for this approach, the cameras used were time-synchronized with GPS, making this a difficult approach to apply in a GNSS denied environment. In 2016, moving to a simpler scheme of monocular VO, [80] performed monocular VO with a forward-facing IR camera mounted on a ground robot. To handle the scale ambiguity, [80] developed a road estimation algorithm to detect the road within the imagery to aid in resolving the scale ambiguity. The second contribution of [80] is the development of a simpler approach to the NUC challenge. Instead of disabling the default NUC and extensively calibrating the camera, the authors decided to manually trigger a NUC based on a custom optimization algorithm that attempts to balance the need for recalibration based on image quality and the dynamics of the vehicle such that the resulting outage has minimal impact on the position estimation. Combining these efforts resulted in a position error of about 8% of the distance traveled in both daytime and nighttime environments [80].

The work presented in the following sub-section aimed at continuing to extend the application of IR imagery for GNSS independent navigation to monocular VO performed onboard a UAV. Additionally, the system developed is designed to reduce the reliance on extensive thermal calibration of the camera or the need to manually trigger NUC events by streamlining the process and allowing the camera to operate as designed to provide the best information possible.

5.2.2 IR Monocular Visual Odometry

The GNSS independent navigation system onboard JAGER performs monocular VO with a gimbal-stabilized, downward-facing IR camera. In addition to the IR camera, the system employs a 1D LIDAR mounted alongside the camera to provide precise height above ground measurements to resolve the scale ambiguity challenge faced by monocular VO.

The VO process flow, depicted in Figure 5.4, is broken down into two major steps after each image taken by the camera is stamped with the height information captured via the LIDAR. The first step includes employing an optical flow process, which is a common building block to implement VO, and includes effectively measuring the velocity of the vehicle using a series of IR images as optical flow measurements. The second step includes estimating the state of the vehicle (e.g. position and velocity) given these optical flow measurements.

One of the biggest challenges of monocular VO is scale ambiguity. Scale ambiguity can be resolved by using a 1D LIDAR to provide a precise measurement of the distance from the camera to the ground. However, this strategy includes generating a single distance measurement, corresponding to the distance of the ground plane at the center of the corresponding image captured by the IR camera. Thus, the integrity of these distance measurements breaks down when the ground plane captured in an IR image is at varying distances from the vehicle, as inappropriate distance measurement may be assigned to features. For example, this variation in the ground plane distance can be caused by uneven terrain. As another example, this can also be caused by the banking of the camera itself, as images taken by a camera that is not pointed directly at the ground will capture features of the ground plane at different distances from the vehicle in different portions of the image frame. The latter case is especially prevalent in systems that utilize a fixed camera whose orientation relative to the ground is dependent on vehicle dynamics. To ensure the camera is pointed directly down toward the ground plane to eliminate the variation in ground plane distances caused by camera banking, regardless of attitude motions of the vehicle itself, mechanical stabilization of the camera is implemented onboard JAGER via a gimbal stabilizer. This mechanical stabilization also presents the additional advantage of effectively removing noise (specifically attitude-based noise) that would otherwise need to be accounted for in estimation.

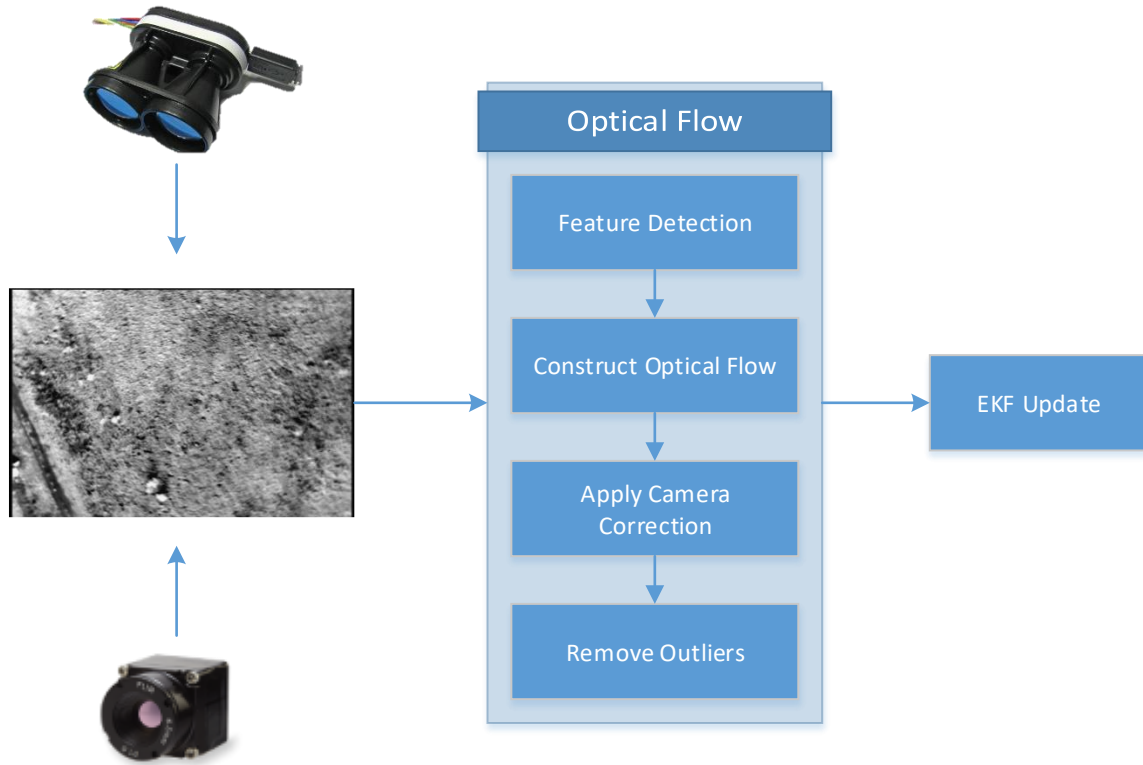


Figure 5.4: Pipeline for IR images to optical flow to the final state estimate of the EKF onboard the JAGER system.

The first step, employing the optical flow process, is broken down into several sub-steps. A Shi-Tomasi corner detector is applied to each image to extract a set of up to 500 features to be tracked using Lucas-Kanade optical flow. This feature extraction step is one of the most computationally intensive steps and, given the need to often extract new features, the threshold of 500 features was determined to be an appropriate balance between computation cost and performance in flight testing. Both the feature extraction and the optical flow are implemented using the OpenCV library [81]. The feature detection is implemented with ‘goodFeaturesToTrack’ which implements a Shi-Tomasi corner detector [82], and the optical flow is implemented with ‘calcOpticalFlowPyrLK’ which implements an iterative Lucas-Kanade optical flow with pyramids [81]. The computed flow field at this point is still in the camera frame and needs to be converted to the body frame of the vehicle by accounting for lens properties of the camera (camera correction) and the height above the ground measurement stamped to the image. Unfortunately, not all features are perfectly matched resulting in the optical flow field to have outliers that are then removed. Because optical flow can use a bundle of tracked features across multiple sequential images to

improve performance, the rate that optical flow velocities are generated and fed into the estimation algorithm is reduced to 10 Hz from the original 30 Hz rate in which the images are captured.

Finally, these velocity measurements are processed via an Extended Kalman Filter (EKF) that estimates the 2D position and 2D velocity of the UAV. The altitude is held accurately with a combination of the onboard barometer and LIDAR and therefore is not estimated. The EKF uses the optical flow determined velocity as a measurement and uses a constant velocity motion model in the prediction step to generate these estimates.

To adapt the system to the IR imagery, two key changes are made to traditional VO solutions: first, the time scale between new features being found within a frame is significantly reduced, and second, additional safeguards are put in place to detect a sensor recalibration and overcome the loss of information during that time period. For the system developed for JAGER, features are never tracked for more than 3 frames to ensure that the change in intensity within the frame does not greatly affect the ability to track the images. This high rate of new feature detection does limit the system to using only 500 Shi-Tomasi features to still allow for constant operation of the camera at nearly 30Hz.

With this breakdown, the analysis can be simplified to determining how well the optical flow system can measure the velocity of the vehicle in flight given the IR imagery. Flight tests performed on Stanford University campus provide an example of two extremes of the system's performance.




Figure 5.5: IR image captured at Stanford University (a) without the detected features and (b) with the detected features.

The grassy environment at the Stanford University test site provided features distributed throughout the image in the IR domain as illustrated in Figure 5.5. Consequently, for the test flight pattern shown in Figure 5.6a, these IR images resulted in good optical flow performance. For the

full test flight, shown in Figure 5.6a, the measured velocity can be seen in Figure 5.6b. Throughout the entire flight, the optical flow was capable of measuring the vehicle velocity utilizing the features extracted from the IR images, resulting in an ability to model the measurement noise as a Gaussian with zero mean and a standard deviation of 0.7m/s.

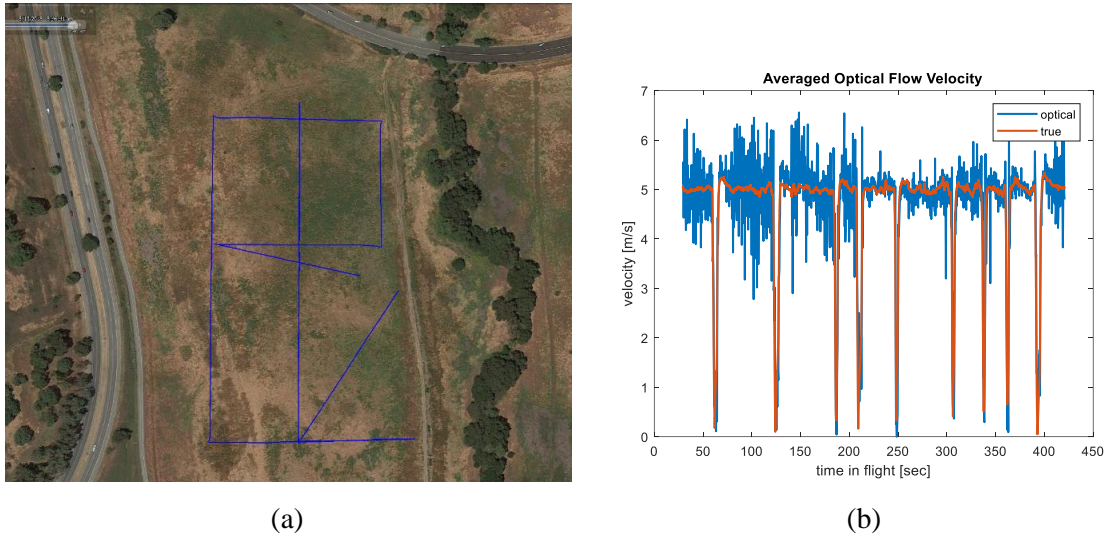


Figure 5.6: Flight test results of IR optical flow with (a) the flight path flown and (b) the true and optical flow velocities throughout the flight for a commanded speed of 5m/s during each straight segment with pauses at each vertex of the flight path.

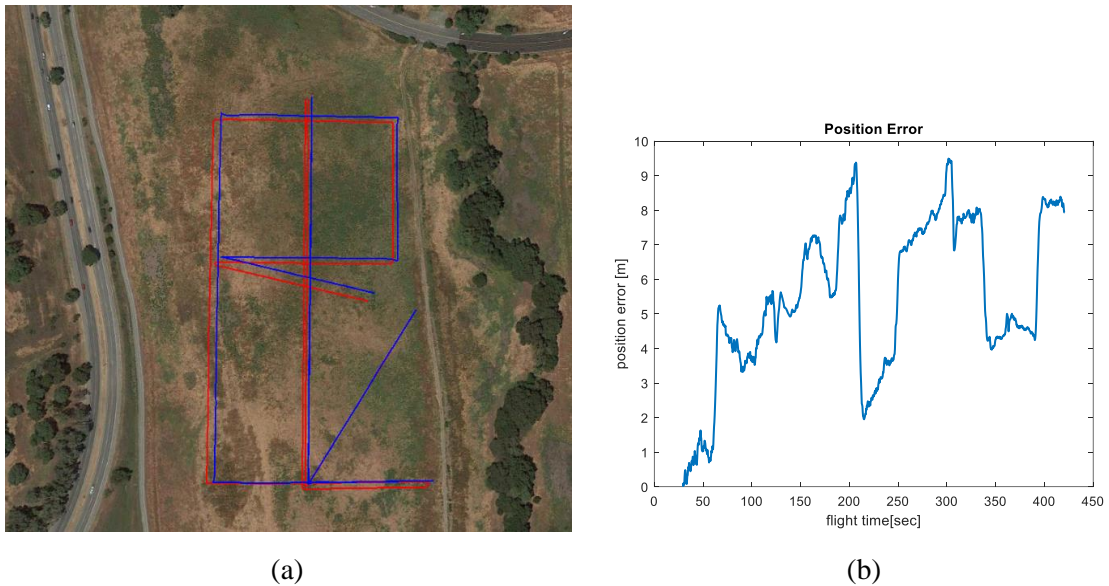


Figure 5.7: Performance of optical flow based navigation system with (a) GNSS (blue) and estimated (red) flight path and (b) 2D position error between the estimated and GNSS position throughout the flight.

For the same flight, the EKF was run using only the velocity measurements and a known initial position of the vehicle. Running the EKF resulted in the estimated position displayed as the red track in Figure 5.7a, next to the true position in the blue track. As expected, the estimated position does drift over the time of the flight. The 2D position error throughout the flight, calculated based on comparing the estimated position to the GPS position, can be seen in Figure 5.7b. Once again, a general drift trend can be noticed, though it is obscured by several large jumps in the error. These jumps are a result of the constant velocity assumption in the motion model that breaks down at each of the corners of the flight pattern, sometimes resulting in smaller errors, other times resulting in larger errors. Overall, for the duration of the 7.5-minute flight, JAGER flew a total distance of 1.8km and maintained a constant velocity of 5m/s through each of the straight sections of the flight. Throughout this flight, the largest error was $\sim 9\text{m}$, with the final position error being $\sim 8\text{m}$, or 0.4% of distance traveled.

On the other end of the spectrum of performance, the same monocular VO pipeline was used on a 9Hz FLIR Vue Pro IR camera during nighttime operations at Edwards Air Force Base in 2017 [83]. In this case, the optical flow's capability of measuring the velocity throughout the flight resulted in a standard deviation of the error of $\sim 1.5\text{ m/s}$ in each velocity component and can be seen in Figure 5.8. As a result, the performance of the pose estimation degraded and resulted in a position drift of about 6% of the distance traveled, as depicted in Figure 5.9. Given this range of possible performance seen during the flight testing, Section 5.4, further explores the impact of the optical flow performance as part of the complete localization mission.




Figure 5.8: Performance of the optical flow system showing the GPS (“true”), measured, and EKF estimated velocity in the (a) North and (b) East directions during a night flight at Edwards Air Force Base using a 9Hz FLIR Vue Pro camera.

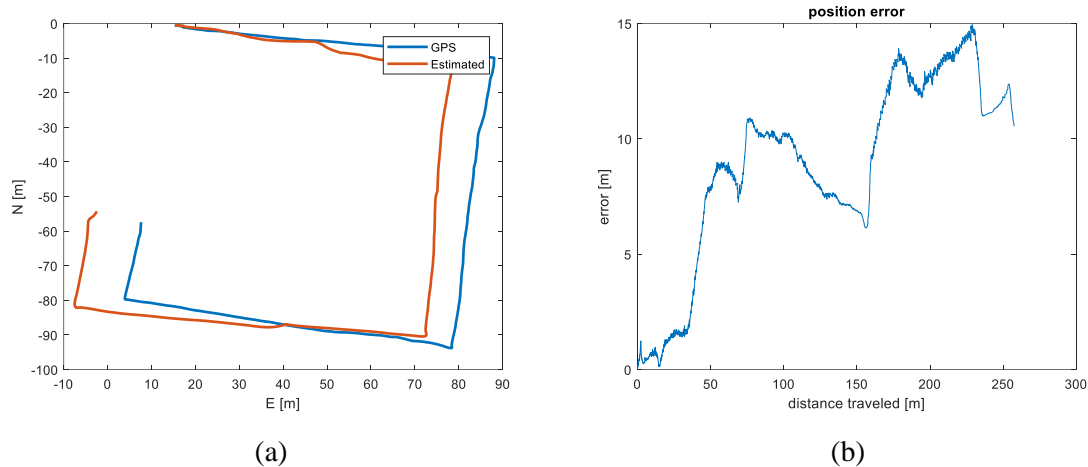


Figure 5.9: Performance of flight at night at Edwards Air Force Base using a 9Hz FLIR Vue Pro camera with (a) the GPS and the estimated flight path and (b) the estimated position error as a function of the distance traveled.

The challenge posed by the recalibration process, the NUC, of the FLIR Boson camera is handled with the assistance of the path planning system described in the next section. The FLIR Boson camera provides information for when a NUC is about to be performed and can thus be used to adjust the flight path of the vehicle for the corresponding outage period if desired. However, JAGER is capable of flying anywhere in the environment, provided the flight altitude is high enough, and therefore constraints can be made on the entire flight path to ensure that the trajectory remains smooth even during NUC outage periods. Therefore, no major changes in direction of the system that would result in unexpectedly large errors in the position estimate of the system during the period of the outage of the imaging information are employed in JAGER's flight path. Indeed, it was determined in the flight testing of the system that with the curved approach employed by the path planning system, discussed in more detail in the next section, the need to readjust the flight path was not necessary, at least on the scales evaluated in flight testing.

5.3 Localizing the Signal Source

The localization algorithm is comprised of two components working in tandem in order to localize the signal source to the required level as quickly as possible: an estimator and a path planner. The estimator is responsible for estimating the location of the signal source using the DoA measurements generated by the LDB-DoA algorithm applied to the RSS patterns created from the

3-element phased array antenna. The path planner is then responsible for directing the vehicle through the localization mission.

Before discussing the specific estimator and path planner used to create the localization system onboard JAGER in Sections 5.3.2 and 5.3.3, Section 5.3.1 briefly provides background on the broader bearing-only localization problem that has been extensively studied. As is discussed in the following section, there is a wide breadth of approaches to the bearing-only localization problem, with entire dissertations devoted to developing more “optimal” methods for localization. While a brief introduction is presented, the estimator and path planner used within JAGER’s localization system is not intended to be optimal – they are intended as an engineering solution to demonstrate JAGER’s capabilities in rapidly localizing a signal source given the use of the 3-element phased array antenna and the LDB-DoA algorithm presented in Chapter 3 and Chapter 4, respectively. It is important to note that DoA measurements are bearing measurements and the word bearing is used in the following sub-section to conform to the literature around this problem.

5.3.1 An Introduction to Bearing Only Localization

Bearing only localization has been a popular area of research for a long time and as a result, there are a myriad of different approaches to solve the problem. However, almost all methods fundamentally break down into two components: an estimator and a path planner. This is due to the limitation that a bearing measurement has an ambiguity on the range to the source and results in an infinitely long conical region where the source may lie. The result is that without “out-maneuvering” the target, the localization algorithm is incapable of narrowing the estimate with bearing only measurements [84]. Therefore, in addition to a localization algorithm, it is important to also have a path planning algorithm in place to ensure a geometric diversity of bearing measurements.

The range ambiguity and the non-linearity of angular error using bearing only measurements have led to quite a variety of different filters being used for the bearing only estimator. A natural and popular choice for the unique area of uncertainty created by a bearing only measurement is the use of either a grid-based filter or a particle filter [34], [37], [52], [85]. While the grid-based approaches work well in small environments, the computational challenge of scaling the grid to meet the needs of an airport environment renders them unideal in such conditions. Furthermore, in situations where the sensor – in this case, the 3-element phased array antenna – contains large amounts of cluttered measurements, the particle filter approach can become computationally intensive. While others have used more specialized filters that handle the angular error more effectively than a standard Extended Kalman Filter (EKF), such as the Shifted Rayleigh Filter

(SRF), others have contended to adapt the EKF in a variety of ways [86]–[88]. Some more popular adaptations of the EKF to the bearing-only problem is to use a weighted sum of Gaussians to effectively create a Gaussian Mixture Model (GMM), range parameterizing the EKF and representing the unique shape of the area of uncertainty created by a bearing measurement [89]–[93]. After the initialization, this approach relies on data association to appropriately update the weights of each of the EKFs in the set with subsequent measurements [89], [91]. Such an approach is computationally lightweight and thus lends itself well to being deployed on lightweight UAV platforms.

The path planning aspect of bearing-only localization can vary significantly with the rate of information provided by the system. For example, in the early stop-and-rotate based modalities for sensing described in Chapter 3, the time cost of pausing to create the RSS pattern and make a corresponding DoA measurement necessitated that the path planners onboard carefully accounted for the cost of making these measurements, which required the path planners to be tightly coupled to the estimator [65], [94]. On the other hand, as the rate of information to the estimator increases, the path planner can be increasingly decoupled from the estimator in response. At high rates of information to the estimator, such as via the use of the 3-element phased array antenna of Chapter 3, the path planner can thus be simplified to focus only upon creating a trajectory through the environment. More specifically, the path planner in these cases no longer needs to worry about the exact position of the vehicle within the world, relative to the estimated position of the signal source. Instead, the path planner serves to provide a trajectory (e.g. through velocity measurements) based on the most recent bearing measurements from the sensor. Previous research has worked to determine the optimal trajectories to maneuver during a signal localization mission that aims to optimize between the distance to the signal source and achieving an angular diversity of measurements [31], [95]. These optimal trajectories are the inspiration behind the path planner used in JAGER’s localization system.

5.3.2 Estimator

The biggest challenge faced by the estimator, beyond its use of only angular DoA measurements, is the fact that the RSS patterns from the 3-element phased array antenna are prone to resulting in cluttered measurements – extraneous DoA measurements that do not point towards a signal source – due to the strong side lobes. To address these additional challenges, work from multi-target tracking applications is leveraged as a strategy for dealing with this clutter. As a result, the estimator operates as not just a single filter, but rather a bank of Gaussian Sum Filters (GSF), with each GSF being a weighted sum of Gaussians representing a possible estimate effectively as a range

parameterized EKF given a set of bearing measurements [89]. This bank of GSF approach is computationally lightweight, rendering it even more ideal to be used onboard a computation resource-limited system. Using a bank of filters in this fashion not only handles the clutter but also enables this algorithm to seamlessly handle environments with multiple sources as easily as it handles environments with a single source, keeping the overall localization system onboard JAGER as versatile as possible and thus ensuring the system can operate in environments beyond the specific case demonstrated in Chapter 6. Leveraging filter designs from multi-target tracking also keeps the estimator as versatile as possible for further applications that involve localization of multiple sources or even moving sources – two types of applications that are made possible when combining the 3-element phased array antenna and the LDB-DoA algorithm previously developed in conjunction with this thesis. Even if only a single, static target exists, the cluttered measurements are still significant as the idea of associating a measurement from the most likely source versus from clutter, such as strong side lobes, is crucial in the implementation of the 3-element phased array antenna [96].

The initialization of the filter is based on the initialization presented in [97] to represent the uncertainty with as few Gaussians as possible and is illustrated in Figure 5.10. With this initialization, there is an assumption placed on the bounds to the range of the sensor, both a minimum and maximum range. Benefits of this representation is that the number of filters grows linearly with the maximum range of the sensor, ensuring that the computational complexity does not grow out of control in larger environments. As a result, a very conservative bound, say 1km for the 3-element phased array antenna that only has a range of about 500m with a 1W interference source, initializes with only 7 filters and remains computationally lightweight.




Figure 5.10: Depiction of the initialization step of the Gaussian Sum Filter given a single DoA measurement.

As this approach employs a bank of filters – or a set of different tracks with each track being represented by a filter – to help handle the clutter and maintain versatility to future expansions with multiple signal sources, a track management scheme is required. The track management scheme used is illustrated in Figure 5.11. The track management scheme has two levels of classification for each track: likely filters to represent tracks that have often been updated and potential filters to represent tracks that have only been newly formed. For each possible DoA from an RSS pattern, the Mahalanobis distance – the distance between a point and a distribution – is calculated for each track. The most likely track, provided the distance is above a given threshold, is updated with the DoA measurement [98]. If no filter has an associated distance above the given threshold, a new GSF is initialized and added to the bank as a possible source location. Over time, the GSFs within the bank that have been continually updated are promoted to be considered as very likely locations. Those that have been very infrequently updated are pruned by being removed from the bank of filters, as they are assumed to have been created from cluttered measurements. Throughout both simulation testing and real-world testing, typically only a single GSF is promoted in a case of a single source, and only two GSFs are promoted in a case of two sources (larger number of sources are beyond the scope of this work).

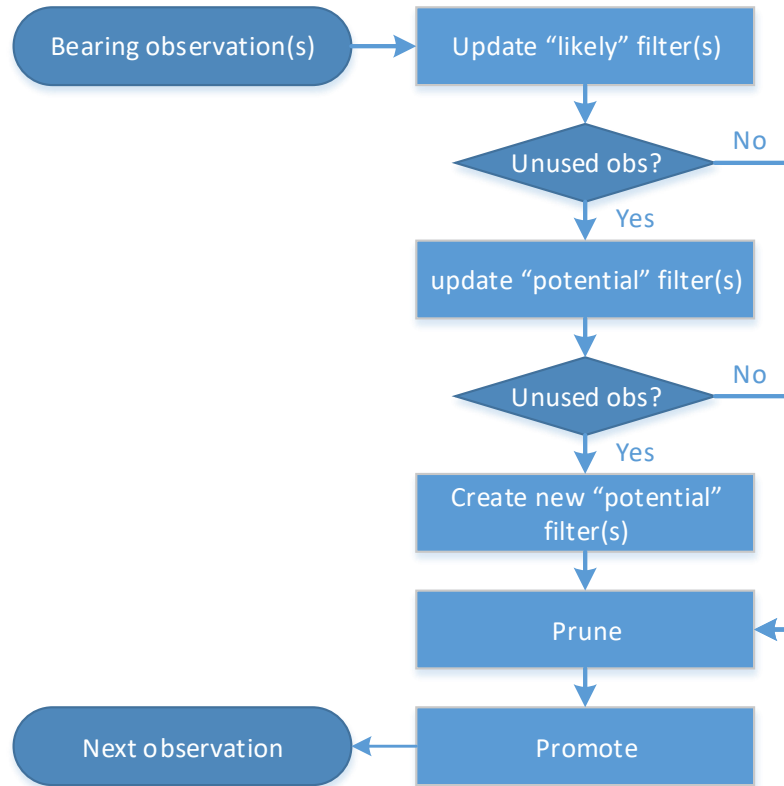


Figure 5.11: Process flow for pruning and updating the bank of GSF used to handle the tracks for clutter management.

5.3.3 Path Planner

The path planner employed as part of the localization system onboard JAGER is inspired by the curved optimal trajectories made possible via the high-rate DoA measurements provided by the 3-element phased array antenna [31], [95]. For a simpler approximation of these optimal algorithms for path planning, the paths can be viewed as a family of curves where the only variable that changes from one curve in the family to another is the approach angle to the signal source (i.e. the angular offset between the last measured DoA and the commanded flight direction). An example subset of curves of a family is shown in Figure 5.12 for a scenario with a signal source 1km away from the initial position of the vehicle. Note that the family has also been modified to ensure that a minimum standoff distance, in this case of 200m, is maintained due to the antenna's reduced effectiveness when nearly directly overhead of the signal source as discussed in Section 4.4. Each of the trajectories in the family are designed to rely on either the last measured DoA to the signal source or the bearing to the current estimated location of the interferer to determine the trajectory.




Figure 5.12: Example subset of the family of curved trajectories possible based on changing the approach angle to the signal source 1km from the takeoff position with 90° (dot), 80° (dash), and 60° (line) approach angles.

For the example subset of curves of the family shown in Figure 5.12, the circle approach is flying a trajectory of $\pm 90^\circ$ from the DoA or from the bearing to the estimated location and the inward spirals are achieved by flying $\pm 60^\circ$ or $\pm 80^\circ$ from the DoA or from the bearing. To ensure that the minimum distance is not violated, while these trajectories do not rely on the range to the signal source to navigate, they must rely on a range estimate to the signal source. For this reason, the family of trajectories does have a limitation of requiring a good enough estimate of the source before getting too close (or else the minimum distance might be violated).

This family of path planning trajectories is selected for this work for its simplicity and for the fact that the necessary velocity command can be generated for the vehicle in flight with knowledge of only the DoA measurement, or only the estimated DoA to the signal source. This makes it a very simple and lightweight path planning algorithm that provides an approximation to the more optimal trajectories possible. This family is also of interest as the circle approach (the outermost curve in Figure 5.12) maintains a constant offset to the signal source which can be used in scenarios where the vehicle's takeoff location is at a distance for which the sensor is able to "see" the signal and the GNSS system onboard the vehicle is not jammed, enabling a localization run without necessitating the reliance upon a GNSS denied navigation system [65]. While the circle trajectory has its advantages in terms of the potential use of a GNSS position throughout the flight [65], encouraging motion towards the signal source via one or more inner trajectories is ideal in more quickly localizing the signal source. This advantage is depicted in Figure 5.13 for the highlighted trajectories from Figure 5.12, with the inwardly curved trajectories more quickly localizing a signal

source to the 20m threshold, in an environment where the vehicle's position is known throughout the flight. As a result, simulation analysis, presented in the next section, is used to help determine which approach from the family of trajectories should be used in the demonstration flights.




Figure 5.13: Estimated signal source position error for three example curved approaches: 90° (dot), 80° (dash), and 60° (line).

In order to execute one of these curved trajectories, the path planning system does need either a single DoA measurement to use as a guide or a bearing to an estimated location of the signal source. Given the fact that the 3-element phased array antenna is prone to cluttered DoA measurements, the path planner needs to use the bearing to an estimated location of the signal source and therefore needs to allow the estimator to initialize a “likely” track to be used as a guide for the trajectory. Therefore, an additional step is added to the path planner to allow the estimator to initialize a “likely” track to be used as a guide. During this initialization phase, JAGER is commended to fly in a user pre-defined initial heading. Flying this straight-line trajectory set by the user gives the necessary delay required by the estimator to collect enough measurement to be able to promote a track to a “likely” track and be used as a guide. Once a “likely” filter becomes initialized, the path planner switches into a controlled localization phase that executes the curved trajectory based on this “likely” track.

The final addition to the path planner, to complete the entirety of the mission, is a trigger when the mission is deemed to be complete. This trigger, given by the estimator itself, can be used to change the path planner to execute a higher-level command, such as to return back to the initial

takeoff location or another predefined landing location. The complete illustration of these three defined steps is shown in Figure 5.14.




Figure 5.14: Illustration of the 3 phases of flight throughout a mission managed by the localization system.

5.4 Simulated System Performance

A simulated environment was used to simulate the theoretical capabilities of the complete JAGER system on a scale that might be seen at an airport environment – that is, with a signal source of varying distances from 500m to 1.5km away. This simulation not only provides insight for the configuration of settings for the flight demonstrations but also provides insight into the behavior on a scale that cannot be flight-tested due to the limited size of the flight space available. While theoretical in nature, the underlying models used in the simulation environment are derived from flight test results. More specifically, the model for the optical flow system is based on flight test data; the model for the 3-element phased array antenna is developed with MATLAB's phased array toolbox to mimic the behavior seen in flight testing, and the vehicle dynamics are modeled to mimic typical behavior seen throughout all the flight tests with the DJI S1000. The world setup for the simulation places an interference source at a range of distances from JAGER's initial takeoff position. Throughout the entire flight for all simulations, JAGER is in a GNSS denied environment with the only piece of known information being the takeoff position. To navigate and localize the signal source, JAGER only uses velocity measurements (from a model of the optical flow system) and RSS patterns (from a model of the 3-element phased array antenna). For each scenario simulated, 1000 Monte Carlo simulations were performed to generate the results presented.

To first evaluate the performance of the path planning trajectories, the optical flow model is set with noise characterized by a standard deviation of 1m/s and the path planner is configured to run each of the three different trajectories illustrated in Figure 5.12 – the circle, 80° inward, and 60° inward trajectories. Note that for the plots discussed with regards to these simulations, the results plotted with a solid line represent the 60° inward trajectory, the results plotted with a dashed line represent the 80° inward trajectory, and the dotted line represents the circular trajectory. To illustrate the impact of using an estimated position of the vehicle, each of the results for the different trajectories are plotted alongside a representative trajectory using GNSS throughout the flight.




Figure 5.15: Error in the estimated interferer's position in simulation flying the three different trajectories (circle [dot], 80° [dash] and 60° [line]) with and without GPS for an interferer 1km away.

For an interference source 1km away, Figure 5.15 compares the difference in the localization performance between the case of JAGER maintaining a GNSS position solution throughout the flight and the case of JAGER operating in a GNSS denied environment for all three trajectories. These results show that with an optical flow system providing velocity measurements having noise with a standard deviation of 1m/s, the system can yield interferer localization performances similar to that of the case where an interferer is found in a GNSS available environment. Most importantly, this means that JAGER can realize a near 3x improvement in time to reaching the 20m error threshold by being able to move towards the signal source. Looking at the final position error over the full span of the possible flight time, the GNSS denied case does not perform as well as the highest performing GNSS available case for each of the trajectories. This is to be expected as the

drift of JAGER's position estimate in the GNSS denied case, along with the increased uncertainty, begin to negatively impact the performance of the bearing only measurements. However, it is also important to note the 2 inward spiral trajectories in a GNSS denied environment do perform better than the standoff strategy's circle in a GNSS available environment after the full duration of the flight, meaning that while a standoff strategy might be possible and would allow the use of GNSS throughout the flight, it is still a better strategy to have a GNSS independent navigation system to allow JAGER to approach the signal source.




Figure 5.16: Error in interferer's estimated position in simulated flights of all 3 trajectories without GPS compared to standoff strategy at a distance of 500m




Figure 5.17: Error in interferer's estimated position in simulated flights of all 3 trajectories without GPS compared to standoff strategy at a distance of 1500m.

The improvements to the localization system provided by the different trajectories are also a function of the initial distance from the interference source. Therefore, the same simulations have also been run on scenarios with an interferer at an initial distance of 500m and 1.5km. These results are shown in Figure 5.16 and Figure 5.17, respectively. Note that in these cases, only the feasible GNSS-based strategy of the standoff circle flight is plotted, instead of plotting all three trajectories with GNSS.

In these figures, in every case, the 60° and 80° inward trajectory improves the time to the 20m error threshold. Furthermore, as the initial distance to the interferer increases, the relative benefit of the 60° inward path also increases. In the scenario with the interferer 500m away, the relative benefit of the 60° inward trajectory may not be worth the loss in performance over the full duration of the flight when compared to the standoff strategy. Namely, the time to achieve an estimated signal source position within the 20m threshold is very similar for all three trajectories, and the possible improvements in the estimated signal source position at the end of a typical flight of a UAV makes the standoff strategy potentially a better strategy – assuming the vehicle can successfully maintain a GNSS solution at that range. Therefore, if JAGER is already in a GNSS denied environment at 500m, the best strategy then shifts to either the 60° or 80° trajectories to achieve a slightly faster localization.

These results so far have assumed an optical flow system capable of measuring the vehicle's velocity to 1m/s. To evaluate the impact of the noise in the optical flow velocity measurements on the EKF's position estimate for the vehicle and therefore the ability to localize the signal source, another set of simulations was performed with a fixed path planning strategy and varying the optical flow velocity model. First, fixing the path planner to the simplest circle strategy and evaluating the navigation system's EKF performance results in Figure 5.18 for velocities with a standard deviation of 0.7 m/s, 1 m/s, and 1.5 m/s – ranging from the best to the worst performance seen in flight testing of the IR optical flow system on JAGER. For these simulations, JAGER is flown for the full 20-minute capability of the vehicle at a flight speed of 5m/s. As expected from velocity only measurements, an increase in the noise of the velocity measurements result in an increase in the drift of the filter. In addition to the increased drift, it is important to also note that the initial slope of the error increases as the velocity noise increases, which will impact the initialization phase of the estimator in the localization system.




Figure 5.18: Estimated vehicle position error in simulated flights with varying velocity measurement performance for a 1km circle flight while flying at 5m/s.




Figure 5.19: Error in interferer's estimated position in a simulated 60° inward flight trajectory at a distance of 1km with varying velocity measurement noise.

To evaluate the localization system's performance with the varying navigation system's performance with a signal source 1km away, the path planner is configured with the most aggressive trajectory – the 60° inward trajectory. The effects of increasing the velocity measurement noise pose two potential problems to the localization error: first, the final achievable error after the 20-minute flight will get worse due to the higher overall drift rate, and second, the initial error growth may be too fast for the localization filter to ever achieve the 20m error

localization target due to challenges in initializing the estimates. Both of these effects are shown in Figure 5.19, where it is apparent that velocity measurements with noise $>1.0\text{m/s}$ result in an inability to achieve the 20m error target, demonstrating that the 60° inward trajectory may be too aggressive in moving towards the signal source.

Based on these simulation results, the final configuration of the path planner must be selected with an understanding of the capabilities of the navigation system. These results demonstrated that if the optical flow system performs closer to the worst-case test sets – i.e. noise of 1.5m/s – a more conservative trajectory should be selected to ensure enough of a diversity of DoA measurements are collected for a well-initialized estimate for the signal source. On the other hand, having an optical flow system with performance closer to the best-case test sets – i.e. noise of 0.7 m/s – allows the system to operate more aggressively to close the distance between the vehicle and the signal source to enable faster localization of the signal source.

5.5 Conclusion

This chapter presented the first part of the third contribution to this thesis: the design and development of a system capable of localizing a signal source 24/7 in a GNSS denied environment. The focus of this first portion was the design and development of the system itself, with the overall system overview for JAGER – the name of the complete system – presented in Section 5.1. Additionally, details for the theory and the implementation of both the navigation system and the localization system were presented in Sections 5.2 and 5.3, respectively.

The GNSS independent navigation system extended downward-facing monocular VO onboard a UAV into the domain of IR imagery and was demonstrated in several different test environments both at Stanford University and Edwards Air Force Base.

With the myriad of approaches to the bearing-only localization problem, one option was selected to provide the most versatility for the complete system. The bank of GSF approach (inspired by multi-target tracking literature) allows the system to gracefully handle the cluttered measurements that result from the 3-element phased array antenna and provide the versatility to operate in not just environments with a single source, but potentially environments with multiple or even moving signal sources. Finally, the high-rate of DoA measurements that result from the 3-element phased array antenna allow the path planning system to be simplified – there is no need to optimize between making a measurement and traveling to get the best set of measurements – to a model with a family of curves inspired by more optimal approaches to bearing-only localization path planning.

The simulation-based performance of the system indicates, especially for a smaller initial distance to the signal source, that a circular or lightly inward spiral approach will be the best for the test environment described in the next chapter. The simulation also demonstrated that the system is capable of localization missions on a larger scale using the same type of sensor and vision-based system with the level of performance described in previous chapters.

The next chapter dives into the real-world demonstration of the complete system in a GNSS denied environment during both daytime and nighttime operations to rapidly localize a signal source. The next chapter also presents the final form of the system – that is, not just the paper design and paper schematic of the system, but the full, complete implementation of every component from the 3-element phased array antenna to the software needed to operate the vehicle.

Chapter 6

Live Demonstration

This chapter finally puts all the elements discussed in the previous chapters together to demonstrate the full capability of the system to autonomously localize a signal source using the 3-element phased array antenna developed in Chapter 3 combined with the Lobe Detection Based Direction of Arrival (LDB-DoA) algorithm presented in Chapter 4 in a real-world environment. The estimation for the signal source given the DoA measurements is done using the method presented in Chapter 5 and the entire system is operated without the use of GNSS, relying on the navigation system presented in Chapter 5 for positioning throughout the duration of the flight.

Section 6.1 sets the scene for the location and circumstances for the flight tests. The following section presents an overview of two example representative successful missions and an example of failure modes that the system encountered during the testing. Finally, Section 6.3 provides an overview of the results of the combined 35 flights flown during this live demonstration campaign.

6.1 Test Setup

This section provides additional details on the setup for the test campaign to demonstrate JAGER's capabilities in a real-world environment. In setting up for the tests, Section 6.1.1 provides a detailed overview of the hardware implementation of the JAGER system described in Chapter 5. The environment selected at Edwards Air Force Base for these real-world demonstrations is detailed in Section 6.1.2 with an explanation of some of the more challenging elements within the environment. Finally, as the navigation system is built on an infrared (IR) camera system, Section

6.1.3 discusses the three times of day selected to demonstrate the system's full capabilities of autonomously navigating in a GNSS denied environment.

6.1.1 The JAGER System

For the live demonstration, along with all pre-demonstration flight testing, all the systems were implemented to operate in real-time onboard the DJI S1000 vehicle. Figure 6.1 shows how all the components were mounted on the DJI S1000 airframe to create the final form of JAGER. JAGER is equipped with two major sub-systems, in addition to the off-the-shelf autopilot responsible for executing the command and control of the vehicle. The software for both the localization system and the GNSS independent navigation system are run within the Robot Operating System (ROS) framework onboard an Intel Core-i3 Compute Stick running Ubuntu 14.04. To create the GNSS denied environment to fully demonstrate the capabilities of the system, the autopilot's software is configured to disable the use of the GNSS antenna for anything outside of logging purposes to have a "truth" position to analyze JAGER's mission performance.

The heart of the localization system is the 3-element phased array antenna (developed and discussed in Chapter 3). As shown in Figure 6.1, the 3-element phased array antenna is mounted upside down on the underside of the vehicle to provide the best field of view from the ground to the horizon – the most likely location of the signal source. This mounting position also provides the necessary separation with the power distribution system onboard the DJI S1000 that causes self-interference to radio frequency (RF) components near it. The rest of the localization chain – the LDB-DoA algorithm (developed and discussed in Chapter 4) along with the estimator and path planner (developed and discussed in Chapter 5) – is implemented on board the Intel Core-i3 Compute Stick mounted on the top side of the vehicle.

The GNSS independent navigation system also contains both hardware and software onboard JAGER. The hardware consists of a mechanically stabilized FLIR Systems Boson 640 infrared (IR) camera with a 4.9mm lens providing a 95° horizontal field of view providing images at 30Hz alongside a Lightware SF 11/C one dimensional (1D) light detection and ranging (LIDAR). The images and the 1D LIDAR measurements are processed onboard the same Intel Core-i3 Compute Stick used for the localization system within the same ROS framework.

In addition to the vehicle itself, a companion ground station, shown in Figure 6.2, was also built and deployed to mimic the operation of the vehicle in a real-world situation. This ground station provided high-level control of the system (i.e. start and stop the localization mission) and enabled the operator to monitor the progress of the system throughout the flight and be notified when the mission was completed.




Figure 6.1: Side view of DJI S1000 equipped with navigation and localization systems; the circled items, from left to right, are: the 3-element phased array antenna, 1D LIDAR, gimbal for camera stabilization, FLIR Boson IR camera, and the Pixhawk autopilot.




Figure 6.2: Screenshot of accompanying ground station showing the position of JAGER and the current estimate (green oval) for the source location.

6.1.2 The Environment

The flight tests were performed at Edwards Air Force Base at the location shown in Figure 6.3. This environment was chosen because it possesses many of the potential features seen at an airport, specifically a taxiway and some off- taxiway vegetation, which is an example deployment location for such a system. Furthermore, the environment also presented some urban environment type challenges, namely a large metal hangar door that reflected the RFI source signal. This environment challenges both the localization scheme and the navigation scheme, enabling both the demonstration of the capabilities of JAGER and the understanding of the limitations of the system as it is designed. Localization is challenged by the variety of the environment (e.g. the hangar door) and navigation is challenged by the potentially feature-poor environments of large regions of similarly textured environments of an airport.




Figure 6.3: Overview of flight testing area for demonstration at Edwards Air Force Base. In red is the full flight test area and the yellow and orange regions highlight regions of unique levels of interference as a result of the presence of the large hangar (labeled Hangar).

The test area, depicted in dark red in Figure 6.3, is about 270m long and 175m wide. Within the test area, three different starting areas were used. Starting area A is about 80m North of the source, starting area B is about 100m East of the source, and starting area C is about 90m South of

the source. Throughout the test campaign, 35 missions were performed approximately starting from these three different starting locations.

Within this test area, the hangar provided two regions of different types of interference. Figure 6.3 highlights two different regions in orange and yellow. In the orange region, the reflected signal appeared stronger than the main signal from the source, resulting in a distorted RSS pattern with the lobe pointing towards the hangar door rather than towards the source itself, as seen in the pattern in Figure 6.4a (the dashed line represents the true bearing to the source and the solid cone represents the measured bearing). For comparison, Figure 6.4b shows an RSS pattern measured outside of these regions of interferences and has a lobe aligned with the true bearing to the source. In the yellow region, the interference resulted in the RSS patterns not having any distinct main lobe pointing in any direction, such as the pattern in Figure 6.4c. The effect and handling of these regions are discussed in more detail in discussing the results of example missions in Section 6.2.

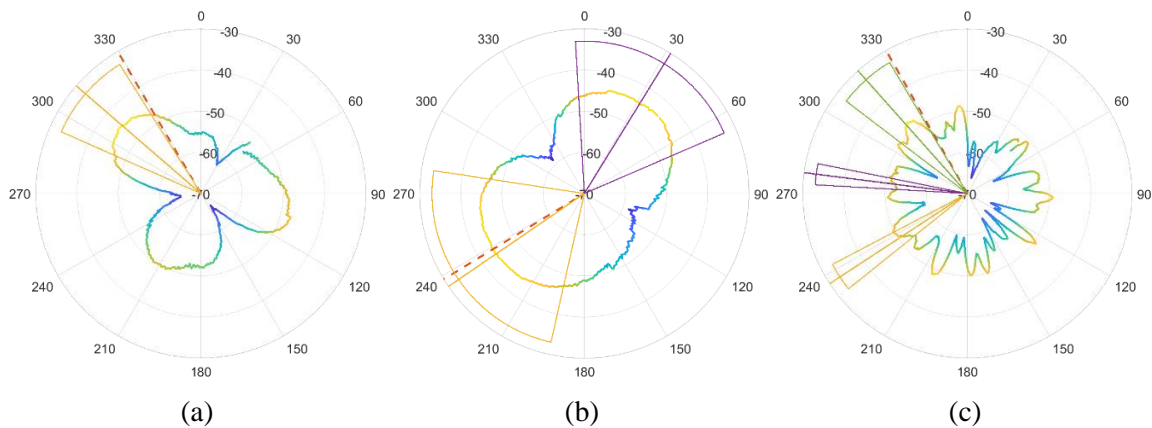


Figure 6.4: Example RSS pattern from different parts of the flight test environment with: (a) from the orange zone where the reflection causes a distortion in the pattern, (b) from a non-interfered part of the flight test area, and (c) from the yellow zone where the reflections from the hangar caused destructive noise.

6.1.3 The Time of Day

Flights were performed in three different times of day, each providing a very different environment and challenge for the GNSS independent navigation system: late in the night (at least an hour after sunset), early in the morning (about an hour before sunrise), and late in the morning. The late-night flights provided an environment to demonstrate the capabilities of an IR camera over a vision camera as it is a time where a vision camera is not able to work successfully.

Even though the IR camera is capable of operating at night, daytime flights were flown to assess the difference in performance between day and night conditions, as daytime conditions do generally offer better features, even in IR. Figure 6.5 shows IR images from above the side of the taxiway taken during the day (left, a) and at night (right, b). There are more details in the day image (a) than in the night image (b) with more of the vegetation and even the runway providing sharper edges to features within the image that help the performance of feature detection algorithms. For both images, the camera is in a “white hot” configuration, meaning that the white regions of the image are the hottest and the black regions of the image are the coolest. It can be seen in these images that this same terrain changes in what is relatively the warmest and coolest region of the image from day to night. This transition happens fairly quickly in the day to night transition but happens fairly slowly in the night to day transition, which motivated the inclusion of the early morning time period for the flight tests.




Figure 6.5: Example IR images³ over a similar region in the flight area (a) around 9 am and (b) around 11 pm.

6.2 Example Missions

The 35 total missions flown from the three different starting locations performed at different times of day attempted to mimic the sequence of events that would occur in a real-world deployment from those three different starting locations. By time, the breakdown of missions was: 18 were flown at night, 10 were flown in the early morning, and 7 were flown in the late morning. For each

³ DISTRIBUTION STATEMENT A. Approved for public release; Distribution is unlimited 412TW-PA-19478

time of day, the flights were fairly evenly distributed across the starting locations. At the beginning of each mission, the operator initiated the autonomous localization mission from within the target starting location. Each mission, lasting at most a minute and a half, ended in one of three ways: with the successful localization of the signal source, with an incorrect initialization of the signal source location estimate, or with the system never initializing an estimate. In the latter two cases, the operator recognized the error and immediately ended the mission to quickly restart a new mission from the same starting location. Of the 35 missions flown, 19 resulted in successfully localizing the source, 15 flights initialized to the wrong location due to clutter in the measurements, resulting in an inability to properly localize the source on the first try, and 1 flight never initialized, also resulting in an inability to localize the source on the first try. However, despite some flights mis-initializing the estimator on the first try, the system could localize the signal source for every one of the starting locations. A more specific breakdown of the number of successful localizations and failed initializations for each starting point is recapped in Table 6.1.

Table 6.1: Overview of the mission performance for flights in each of the starting areas.

	Successful Localization	Wrong Initialization	Never Initialized
Starting Location A	8	2	0
Starting Location B	8	7	1
Starting Location C	3	6	0

Two example successful missions are highlighted in this section, with the first mission being a representative example of most of the successful missions. The second successful mission highlighted in this section provides additional commentary on the impact of the ability to perform integrity checks on the underlying RSS pattern thanks to the unique parameterization provided by the LDB-DoA algorithm from Chapter 4. The 15 missions that resulted in an incorrect initialization of the signal source estimate were all the result of the same two phenomena that are described in Section 6.2.3.

6.2.1 Mission 1

The first mission discussed starts from location B in a counterclockwise direction at 10:45 pm. Most of the successful missions followed the same main elements as this flight as presented here. All the phases of the flight are laid out in Figure 6.6. The first step, shown in Figure 6.6a, is the end of the predefined straight line for the initialization of the estimate, in this case, a heading of 80°

was chosen. Once an estimate reached a high enough confidence level to be initialized as the likely source estimate, the vehicle flies an arc around the source until the covariance of the estimate drops below a given threshold, as shown in Figure 6.6b, and the filter considers the source as found. While the real-time display of the estimate allowed knowing when JAGER deemed the source as found, during these tests JAGER was directed to continue to fly to both collect more data and reduce the covariance size further, to the final state, shown in Figure 6.6c.




Figure 6.6: Frames of the Mission 1 flight showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) at the initialization of the GSF for the signal source occurred, (b) when the filter internally determined the source to be localized, and (c) at the final state of the estimate before returning to manual control.

Looking at the full-highlighted path in Figure 6.6c, the expected circular trajectory does not have a constant radius of curvature. This change in radius of curvature is a result of the path planner algorithm using the mean of the estimate for the center of the circle. Therefore, in the early stages

of the controlled localization, when the Gaussian Sum Filter (GSF) of the estimated location of the signal source is still comprised of many evenly weighted estimates (e.g. the state in Figure 6.6a), the estimated position of the source is at a point much further out. As more measurements come in and the mean of the estimate moves closer to the true source location, the radius of curvature decreases, resulting in the tighter circle seen in the later portion of the flight.




Figure 6.7: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight.

Figure 6.7 shows the covariance of the estimate and the true position error in the estimate for the source location throughout the flight. Note that the covariance and error is plotted against the angle traveled around the source rather than time as plotting against angle traveled normalizes the results of different starting distances from the source and removes some distortion that results from JAGER not flying at a perfectly constant velocity. Once JAGER has flown enough around the source, the covariance of the estimate quickly drops. The threshold used to consider the source as found is when the 95% confidence ellipse is less than 1000 m^2 , based on previous simulation results of the system, which in this case is triggered at about 30° traveled around the source. In simulation, the more typical result is closer to 100° around the source, and as is discussed later, this mission is atypically quick in localizing the source. However, all missions follow a similar trend with the covariance, and position error, quickly dropping after enough of an angle around the source has been traveled.

On the navigation front, the error in the estimated location of JAGER compared to the GNSS position of JAGER is plotted in Figure 6.8. Additionally, the velocity, as measured by the optical

flow system and as estimated by the EKF compared to the GNSS velocity measurement, is plotted in Figure 6.9. The EKF for the state of the vehicle assumes a constant velocity, which is not the case with the path planning strategy having to take the turn from the predefined heading to the circle trajectory, and results in a sharp jump in position error. Overall the error growth over time, for this flight, was 6% of the distance traveled, which was more typical for the flights at night.




Figure 6.8: 2D position error of the estimated position of the vehicle throughout the flight.

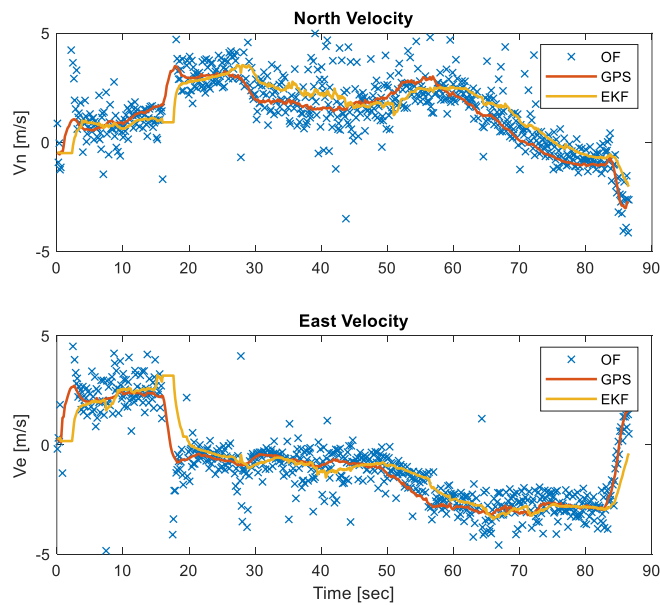


Figure 6.9: Overview of the velocity performance throughout the flight with the Optical Flow (OF) (blue x), EKF estimate (yellow), and GPS (red) velocities for the North and East directions.

6.2.2 Mission 2

The second mission goes counterclockwise beginning at starting location C. Starting location C was subject to a much greater amount of interference from the hangar door than either of the other two starting locations. For this mission, there are two different sets of results, the first is the result of the missions exactly as flown, and the second is a result of rerunning the mission including an integrity check for each of the measured RSS patterns to exclude measurements from low confidence patterns from updating the localization estimate.

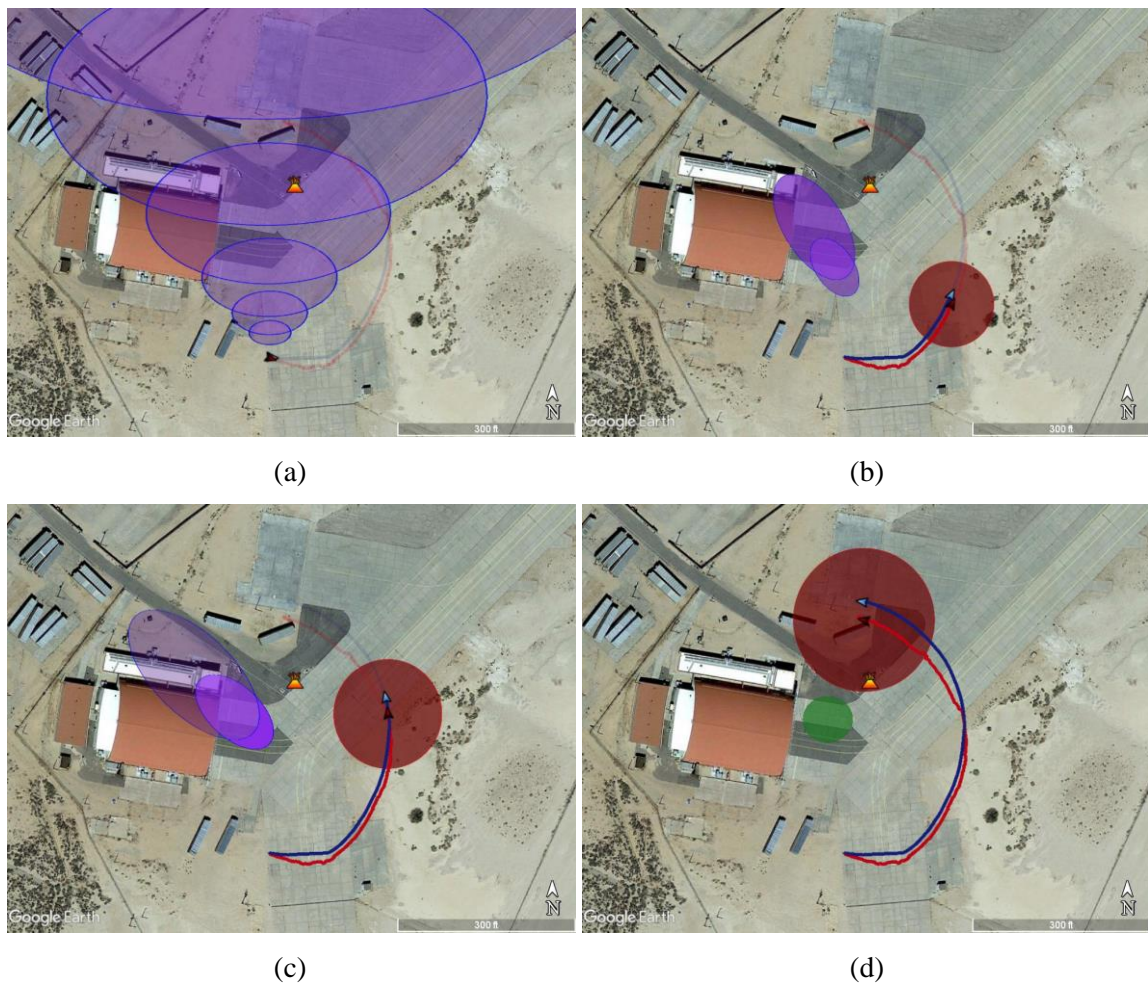


Figure 6.10: Frames of the Mission 2 flight with no integrity checks showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) the initialization of the GSF for the signal source occurred, (b) begins to convert towards hangar door due to interference region, (c) the vehicle leaves the regions of interference and the estimate moves away from the hangar door, and (d) the final state of the estimate.




Figure 6.11: Frames of Mission 2 with integrity checks showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself at several specific instances: (a) the estimate covariance remaining large due to the discarding of noisy measurements, (b) the estimate remaining away from the hangar as the vehicle leaves the regions of interference, and (c) the final state of the estimate.

As flown, before including the confidence information, the flight progressed as depicted in Figure 6.10. In Figure 6.10a, the initialize estimate of the filter is in the proper direction towards the source. As JAGER flies through the region of interference that caused distorted patterns aimed towards the hangar door (e.g. the pattern in Figure 6.4a), the estimate begins to converge to the hangar door, depicted in Figure 6.10b, instead of the source itself. Continuing through the region of interference that results in antenna patterns with low confidence (no discernable lobes or only cluttered bearing measurements as depicted in Figure 6.4a and Figure 6.4b), the estimate can be seen in Figure 6.10c to continue to converge towards the hangar door. Finally, by the time JAGER

reached a region of less information, the estimate slowly started to pull away from the hangar door, but the final result, shown in Figure 6.10d, does not capture the jammer within its error ellipse.

After experiencing these effects from the hangar door, the localization estimator was adjusted to exclude “low confidence” RSS patterns, such as the pattern in Figure 6.4c, where there is no discernable lobe. Rerunning the exact same mission with the exact same measurements (in post-processing), resulted in a better estimate throughout the flight, depicted in Figure 6.11. Due to measurements being excluded early in the flight due to interference, the covariance of the estimate remains large, visually seen in Figure 6.11a and Figure 6.11b, until JAGER reaches the region of less interference. As a result, the final estimate, shown in Figure 6.11c, converges more accurately on the true location of the signal source.

Comparing the covariance and the 2D error of the estimate for both of these runs, in Figure 6.12 for the original mission flight, and Figure 6.13 for the rerun removing “low confidence” measurements, it can be seen that the covariance initially stays a lot larger (i.e. does not converge quickly to the hangar door) for the rerun than it does for the original flight, resulting in the estimate to converge more accurately to the source by the end of the flight.




Figure 6.12: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight with no integrity checks.




Figure 6.13: Position estimate (blue) and 95% covariance area (orange) for the signal source location throughout the duration of the flight with integrity checking.

6.2.3 Wrong Initialization

Not all the missions resulted in correct initialization in this campaign. These wrong initializations resulted in the majority of failures to localize the signal source in the first try. The interference from the hangar door combined with the cluttered bearing measurements from the 3-element phased array antenna posed a challenge to the initialization scheme used by the localization filter, and the subsequent path planner, resulting in some of the missions to “fail”, which is to say that the localization algorithm for the signal source initializes to an incorrect estimate.

Figure 6.14 shows some example flight paths of wrong initializations. These two examples illustrate the two main types of wrong initialization. Figure 6.14a demonstrates a wrong initialization that converges almost immediately to a location very close to JAGER. This type of failure was commonly seen when starting in one of the regions of higher levels of interference from the hangar door. Figure 6.14b demonstrates a wrong initialization that has the estimate from the back lobe as the highest likelihood estimate for the signal source. This type of failure also occurred when starting in location A, with no interference, as it is a result of the cluttered measurements from the 3-element phased array antenna. Both examples illustrate the sensitive nature of the initialization scheme implemented for this test with the cluttered measurements that result from using the 3-element phased array antenna.



Figure 6.14: Two example flights showing the jammer (orange volcano icon), flight paths (estimated in red and GPS in blue), and the estimates for the signal source and the vehicle itself where the initialization of the estimate failed with (a) the estimate immediately converging close to the initial position due to noisy RSS patterns and with (b) the estimate converging on DoA measurements from the non-main lobe (e.g. back lobe).

Both failed initialization examples started from location B, however, as seen in Table 6.1 showing the overview of the performance by starting location, location C was the most challenging starting location. In starting location C, many of the missions failed due to an almost immediate convergence to a location close to the takeoff location (the failure depicted in Figure 6.14a) due to the strong interference from the hangar door. In starting location C, most of the measurements through the initialization phase of the flight were a combination of distorted patterns (see Figure 6.4a) and corrupted patterns (see Figure 6.4c) that caused the initialization to fail.

6.3 Flight Test Overview

Table 6.2 and Table 6.3 look at the localization and navigation results, respectively, more broadly for all the missions flown. The success of the source localization was not greatly affected by the time of day (i.e. the performance of the navigation algorithm), so the overview results shown in Table 6.2 are for all the combined missions. While Table 6.3 shows that the ability to localize the source was much more strongly correlated with the starting location and the quality of the initial set of RSS patterns (and therefore DoA observations), further highlighting the importance of the initialization phase of the estimator, Table 6.2 does not break down the results by starting location, but rather only includes the successful missions, because, when a mission was successful, there was not a significant difference in the mission statistics (i.e. how the mission played out).

Table 6.2: Overview of the performance of signal localization for all successful localizations.

	Min	Max	Average
Time to Localization	19 sec	98 sec	45 sec
Angle to Localization	28°	123°	52°
Found Source Error	5 m	51 m	24 m
Final Source Error	4 m	21 m	13 m

Table 6.3: Overview of the performance of the IR navigation system during all missions.

	Day	Night
Velocity Noise	0.1 rad/s (2 m/s @ 20 m AGL)	0.23 rad/s (4.6 m/s @ 20 m AGL)
Drift Rate [% of distance traveled]	4%	6.2%

For the source localization, simulation results showed that when using DoA-only measurements, JAGER would need to fly about 100° around the source before successfully localizing the source. For these flight tests, Table 6.2 shows that the angle traveled turned out to be lower than expected, with only traveling on average of 52° around the source before localization. The threshold for the algorithm calling the source “found” was set to target a 2D error of <20m, however for these tests that threshold resulted in an average error of 24m. This is a tunable parameter that can, and should, be adjusted for future missions to better reflect the performance with real-world data. The mission that resulted in a 51m error was a mission that started in location C and rejected many of the measurements, resulting in an overall worse performance, yet is still considered a success as the system executed as desired (encircled the estimate) with the exception that the trigger for being considered “found” tripped too early.

The performance of the navigation system did have a difference between day and night operations, with much better performance during the daytime hours (post sunrise) than at night (night includes the pre-sunrise flights). For these tests, both the early morning and late-night flights performed very similarly and are therefore both combined into the “night” statistics. As mentioned above, despite the difference in the performance of the navigation algorithm, the final source localization performance was not greatly affected. This is most likely due to how quickly the source was found in all these flight tests (with the longest taking only 98 seconds) and therefore the short amount of distance traveled before the source was found. For a source that is much further away

than these tests, there would start to be a noticeable difference in the localization performance between the day and night times.

6.4 Conclusion

This chapter presented the final, real-world demonstration of the complete system capable of localizing a signal source in a GNSS denied environment in a matter of minutes as part of the third contribution of this thesis: the design, development, and demonstration of a flight test platform capable of 24/7 signal localization operations. The environment selected for these tests was able to demonstrate the performance of the system in both an open environment (location A) and in more strenuous urban-like environments (locations B and C). As depicted in Table 6.1, not every mission resulted in a successful localization due to the challenges posed by the urban environment on the 3-element phased array antenna, however, given multiple tries, the system was capable of localizing the signal source to within 20m from every starting location tested.

These flights tests served to not only demonstrate the system's capabilities using every element developed within this thesis from the 3-element phased array antenna of Chapter 3 to the LDB-DoA algorithm from Chapter 4 to the completing touches of the localization system and GNSS independent navigation system developed in Chapter 5 but also provide important lessons learned on the operating considerations that need to be made with this system when deployed at airport environments. On the navigation front, the IR camera-based system demonstrated to be a feasible system capable of providing vehicle positioning necessary to localize the source during the day, at night, and in the early morning for the configuration tested; however, these tests demonstrated that for larger-scale flights, the constant velocity assumption of the EKF may begin to degrade the quality of the position estimate before the signal is successfully localized. On the localization front, the importance of the initialization scheme of the estimator was highlighted, making it important to account for the environment around the system at the beginning of the flight as it has a much greater impact on the overall performance than any other portion of the flight.

Overall, this chapter served to demonstrate that this end-to-end system, built on the three contributions presented in this thesis working together, has the potential to meet the needs of the GNSS interference localization mission on a time scale well below any previously seen.

Chapter 7

Conclusion

7.1 Summary and Contributions

This thesis presented the design and development of key systems required to realize an autonomous system capable of rapidly localizing a signal source in a time scale of minutes and hours in a GNSS denied environment. Additionally, an end-to-end, prototype-level version of the system was constructed to demonstrate the capabilities of each sub-system and the complete system in real-world environments. Furthermore, the system developed in this thesis does not solely reside in the theoretical domain – this system has been tested numerous times in the field, demonstrating a path towards developing a deployable system to save countless hours and keep airports safe and operational. As more and more critical applications across multiple industries rely on GNSS, this platform and method has applications beyond aviation alone and can be implemented to protect position, navigation, and time for GNSS operators in a myriad of applications.

Lightweight Sensing

Chapter 3 presented a new, small form-factor, electronically steered antenna that provided high-rate RSS measurements in a self-contained package capable of being mounted on a small UAV. The small form-factor results in this antenna's ability to be easily deployed on a variety of existing COTS UAVs, greatly reducing the overall system complexity. In addition to the weight savings of using a 3-element phased array antenna, the electronic rotation of the antenna enables the sensing requirements to be decoupled from the UAV itself – more specifically, the UAV no longer needs to rotate or otherwise maneuver in order to create a Received Signal Strength (RSS) pattern,

allowing the UAV to focus on the simpler task of being a mere vessel for carrying a sensor through the environment. Section 3.3 described the specific implementation of the antenna with the command and control that enables the antenna to rotate a beam at a rate of 3Hz, a full hundred-fold faster than the 30 seconds required to physically rotate a directional antenna to create the same measurement [37], [65]. Furthermore, the antenna's ability to control both the steering angle and the measurement of the signal strength required to create the RSS pattern was shown in Section 3.3.4 to allow the antenna to handle temporal noise and events in the environment that rotation based methods have previously been unable to cope with.

Algorithm for Direction of Arrival Determination and Localization

Chapter 4 developed the Lobe Detection Based Direction of Arrival (LDB-DoA) algorithm – a robust algorithm for measuring the direction of arrival (DoA) from an RSS pattern generated in a variety of real-world environments. This computationally lightweight algorithm is designed to be agnostic to the antenna that generated the RSS pattern, unlike comparative cross-correlation based methods that require a knowledge of the theoretical pattern from the antenna in order to make the same DoA measurement. Furthermore, the algorithm makes no assumptions on the number of possible lobes in the resulting RSS pattern, meaning that the algorithm can reliably operate on RSS patterns generated from multiple sources as demonstrated in the analysis of a two source world in Section 4.4.3.

Critical to the development of a system capable of operating in real-world environments is an algorithm's robustness to noise that can be found in real-world environments. Section 4.4 put the LDB-DoA algorithm through its paces on both a physically rotated directional antenna and the 3-element phased array antenna of Chapter 3 and demonstrated that the LDB-DoA algorithm outperforms state of the art algorithms for determining DoA from an RSS pattern. Finally, the unique parameterization – detecting and returning information on the lobes of the RSS pattern itself – presented by the LDB-DoA algorithm, provides information on the underlying RSS pattern that can be used as an integrity check on the measurement to determine whether or not to discard given measurements due to noise during a localization mission.

Sensor Platform Development and Deployment

Finally, Chapter 5 and Chapter 6 focused on completing the entire system to rapidly and autonomously localize an interference source in a GNSS denied environment any time of day or night. The design of the system, presented in Chapter 5, included the development of a self-contained GNSS independent navigation system capable of 24/7 operation and the development of

the estimator and path planner required to complete the localization system built on the 3-element phased array antenna and the LDB-DoA algorithm.

The navigation system detailed in Section 5.2 extends monocular visual odometry techniques to the infrared (IR) domain using a gimbal-stabilized IR camera. The system adapts an optical-flow based visual odometry technique to 30Hz IR imagery from a downward-facing, gimbal-stabilized FLIR Boson camera alongside a 1D LIDAR that provides precise height above ground measurements.

The localization system detailed in Section 5.3 leverages techniques from bearing-only multi-target tracking to create an estimator that relies on a bank of Gaussian Sum Filters that represent the possible estimates for the signal source. In conjunction with the estimator, a simple path planner inspired by the curved trajectories of more optimal paths for bearing-only localization commands JAGER – the name of the UAV based system developed in Chapter 5 – to fly in a curved inward trajectory throughout the flight to enable rapid localization of the signal source.

While simulation results are presented in Section 5.4, a full end-to-end demonstration of the system is presented in Chapter 6 during flight trials at Edwards Air Force Base. The flight trials pushed the limits of the system through operations in challenging real-world environments – presenting noisy RSS patterns to demonstrate the robustness of the LDB-DoA algorithm – at different times of day and night to further demonstrate the capabilities of the IR-based navigation system. Through 35 different missions, the flight trials demonstrated JAGER’s ability to localize a signal source 100m away from the initial position in an average time of just 45 seconds with an average position error of just 24m. The navigation system was demonstrated to have a drift rate of 4% of the distance traveled during the day and 6% of the distance traveled at night, and further proved to be suitable in enabling the localization mission at all times of day and night.

7.2 Directions for Future Work

As with all systems, especially those at the prototype-level of development presented in this thesis, almost all elements of the chain present areas for potential improvements. The system presented in this thesis demonstrates the possibilities of what this design is capable of, and those possibilities can only be improved.

Generating RSS Patterns and DoA Measurements

While legal regulations limited the scope of the physical development to only a 3-element array, in discussing the LDB-DoA’s capabilities in environments with multiple sources in Section 4.4.3 it

was seen that increasing the size to even just a 7-element array can provide great improvements on the amount of clutter and resolution of DoA measurement and future developments can aim to leverage those larger array sizes. Additionally, as technology continues to miniaturize computation power, especially with the prevalence of Field-Programmable Gate Arrays (FPGAs) on Software-Defined Receivers (SDRs), the analog circuitry used in the development of the antenna in this thesis can be moved into the digital domain, potentially providing yet more orders of magnitude improvement on the time to create an RSS pattern and potentially improving the accuracy of the DoA measurements from those patterns.

Using a planar array allows a beam to be steered in two dimensions – a fact that was leveraged to rotate the beam of the 3-element phased array, albeit only in the azimuth direction as the angle was held constant in the polar direction. While the beamwidth of the 3-element array made it such that the polar direction did not contain a wealth of information, changing the size of the array, or even the sensing modality to look for “null” instead of beams with a null-steering version of the antenna, can result in that second dimension – that polar angle –containing more information. While this will increase the necessary search space for generating a good DoA measurement, this can also yield a two-dimensional DoA measurement that can aid in potentially speeding up the localization or enabling localization in a third dimension.

Airframe

While all the flight-testing and demonstration presented in this thesis is built on a multirotor vehicle, that does not need to be the case in the future. The antenna developed in Chapter 3 explicitly decouples the sensing modality from the vehicle, allowing for a variety of different types of vehicles to potentially be used in future developments of this system. The lightweight sensor can just as easily be mounted on a fixed-wing UAV that can much more quickly travel through the environment to execute the localization system. While that may add challenges to the rapid deployment of the vehicle, the growth of transitioning vehicles can also be explored to reduce the deployment challenges.

Estimator and Path Planner

Finally, the estimator and the path planner used in this thesis are an engineering solution to the challenge. As was briefly discussed in Section 5.3.1, there is a wealth of existing and ongoing research in the domain of bearing-only localization, and therefore a potential wealth of room for improvement on the estimator and path planning used. Research continues on bearing-only localization methods with the development of more robust algorithms that account for uncertainties

in the vehicle's position over time, or in even using the unknown signal source as a feature in the environment [99]. For example, research in bearing-only simultaneous localization and mapping (SLAM) techniques may enable the inclusion of the DoA measurements to the signal source into the navigation system to effectively cap the drift of a velocity only based navigation system by providing a key feature in the environment [86], [100]–[102].

Bibliography

- [1] J. Carroll, J. A. Volpe, and M. K. Montgomery, “Global Positioning System Timing Criticality Assessment - Preliminary Performance Results,” in *Proceedings of the 40th Annual Precise Time and Time Interval Systems and Applications Meeting*, 2008, pp. 485–506.
- [2] “National Space-Based Position, Navigation, and Timing Advisory Board, 15th Meeting Minutes,” Annapolis, MD, 2015.
- [3] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance second edition*, Second Edi. 2006.
- [4] B. Parkinson, “Assured PNT for Our Future: PTA,” *GPS World*, Sep-2014.
- [5] M. Geyer, P. D, and R. Frazier, “FAA GPS RFI Mitigation Program,” *Antenna*, no. September 1999, pp. 14–17, 2000.
- [6] S. Pullen, G. Gao, C. Tedeschi, and J. Warburton, “The impact of uninformed RF interference on GBAS and potential mitigations,” *Inst. Navig. Int. Tech. Meet. 2012, ITM 2012*, vol. 1, pp. 780–789, 2012.
- [7] A. T. Balaei, B. Motella, and A. G. Dempster, “GPS interference detected in Sydney-Australia,” *Electronics*, vol. 6257750, pp. 2–12, 2007.
- [8] D. Last, “GNSS: The Present Imperfect,” *Insid. GNSS*, 2010.
- [9] S. Navigation, “USB3 Interference.” [Online]. Available: <https://support.swiftnav.com/customer/en/portal/articles/2937832-usb3-interference%0D>.
- [10] C. Curry, “Sentinel Project Report on GNSS Vulnerabilities,” 2011.
- [11] “GPS Jamming - Out of sight,” *The Economist*, Jul-2013.
- [12] M. Stanisak, K. Von Hünerbein, and U. Bestmann, “Measured GNSS Jamming Events at German Motorways,” in *POSNAV 2016*, 2016, no. July.

- [13] J. C. Grabowski, "Personal Privacy Jammers : Locating Jersey PPDs Jamming GBAS Safety-of-Life Signals," no. April 1, 2012, pp. 1–31, 2012.
- [14] A. Coultrup, "GPS Jamming in the Arctic Circle," *Aerosapce Security*, Apr-2019.
- [15] D. Goward, "Jammers at Dachas add to Russia's Ability to Silence GPS," *GPS World*, Jun-2019.
- [16] D. Goward, "GPS Disrupted for Maritime in Mediterranean, Red Sea," *GPS World*, Nov-2018.
- [17] J. R. Clynch, A. A. Parker, R. W. Adler, W. R. Vincent, P. McGill, and G. Badger, "The Hunt for RFI: Unjamming a Coast Harbor," *GPS World*, pp. 16–22, Jan-2003.
- [18] "Ubiquity and the Joint Navigation Conference," *GPS World*, 2007.
- [19] "UK Study Indicates Just How Costly a GNSS Disruption Can Be," *Inside GNSS*, 2017.
- [20] J. Coffed, "The Threat of GPS Jamming The Risk to an Information Utility," 2014.
- [21] A. Grant, P. Williams, N. Ward, and S. Basker, "GPS jamming and the impact on maritime navigation," *J. Navig.*, vol. 62, no. 2, pp. 173–187, 2009.
- [22] A. G. Dempster and E. Cetin, "Interference Localization for Satellite Navigation Systems," *Proc. IEEE*, vol. 104, no. 6, pp. 1318–1326, 2016.
- [23] Y. Morio, T. Inoue, T. Tanaka, and K. Murakami, "Worker posture recognition for understanding agricultural worker behaviors," *Eng. Agric. Environ. Food*, vol. 10, no. 3, pp. 208–222, 2017.
- [24] G. Mao, B. Fidan, and B. D. O. Anderson, "Wireless sensor network localization techniques," *Comput. Networks*, vol. 51, no. 10, pp. 2529–2553, 2007.
- [25] A. Wadhwa, U. Madhow, J. Hespanha, and B. M. Sadler, "Following an RF trail to its source," *2011 49th Annu. Allert. Conf. Commun. Control. Comput. Allert. 2011*, pp. 580–587, 2011.
- [26] J. Biswas and M. Veloso, "WiFi localization and navigation for autonomous indoor mobile robots," *Proc. - IEEE Int. Conf. Robot. Autom.*, pp. 4379–4384, 2010.
- [27] A. Awad, T. Frunzke, and F. Dressler, "Adaptive distance estimation and localization in WSN using RSSI measures," *Proc. - 10th Euromicro Conf. Digit. Syst. Des. Archit. Methods Tools, DSD 2007*, pp. 471–478, 2007.
- [28] M. Chesser, "TrackerBots : Autonomous UAV for Real-Time Localization and Tracking of Multiple Radio-Tagged Animals."
- [29] "J911: Fast Jammer Detection and Location Using Cell-Phone Crowd-Sourcings," *GPS World*, 2010.
- [30] A. Cidronali, S. Maddio, G. Giorgetti, and G. Manes, "Analysis and Performance of a Smart

- Antenna for 2,45-GHz Single Anchor Indoor Positioning,” *IEEE Trans. Microw. Theory Tech.*, vol. 58, no. 1, pp. 21–31, 2010.
- [31] S. E. Hammel, P. T. Liu, E. J. Hilliard, and K. F. Gong, “Optimal observer motion for localization with bearing measurements,” *Comput. Math. with Appl.*, vol. 18, no. 1–3, pp. 171–180, 1989.
 - [32] J. Graefenstein, A. Albert, P. Biber, and A. Schilling, “Wireless node localization based on RSSI using a rotating antenna on a mobile robot,” in *2009 6th Workshop on Positioning, Navigation and Communication*, 2009, pp. 253–259.
 - [33] S. Venkateswaran *et al.*, “RF source-seeking by a micro aerial vehicle using rotation-based angle of arrival estimates,” in *2013 American Control Conference*, 2013, pp. 2581–2587.
 - [34] J. T. Isaacs, F. Quitin, L. R. Garcia Carrillo, U. Madhow, and J. P. Hespanha, “Quadrotor control for RF source localization and tracking,” in *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2014, pp. 244–252.
 - [35] V. Krishnaveni, “Beamforming for Direction-of-Arrival (DOA) Estimation-A Survey,” *Int. J. Comput. Appl.*, vol. 61, no. 11, pp. 4–11, 2013.
 - [36] M. Trinkle and D. A. Gray, “Interference Localisation Trials Using Adaptive Antenna Arrays,” no. September, pp. 24–27, 2002.
 - [37] O. Cliff, R. Fitch, S. Sukkarieh, D. Saunders, and R. Heinsohn, “Online Localization of Radio-Tagged Wildlife with an Autonomous Aerial Robot System,” *Robot. Sci. Syst. XI*, 2015.
 - [38] G. M. Hoffmann, “Autonomy for Sensor Rich Vehicles: Interaction Between Sensing and Control Action,” no. September, 2008.
 - [39] E. M. Geyer, B. M. Winer, and R. A. Frazier, “Airborne GPS RFI localization algorithms,” *Proc. ION GPS*, vol. 2, pp. 1447–1456, 1997.
 - [40] F. Körner, R. Speck, A. H. Göktoğan, and S. Sukkarieh, “Autonomous airborne wildlife tracking using radio signal strength,” *IEEE/RSJ 2010 Int. Conf. Intell. Robot. Syst. IROS 2010 - Conf. Proc.*, pp. 107–112, 2010.
 - [41] K. Vonehr, S. Hilaski, B. E. Dunne, and J. Ward, “Software Defined Radio for direction-finding in UAV wildlife tracking,” *IEEE Int. Conf. Electro Inf. Technol.*, vol. 2016-Augus, pp. 464–469, 2016.
 - [42] C. G. Muller *et al.*, “Aerial VHF tracking of wildlife using an unmanned aerial vehicle (UAV): comparing efficiency of yellow-eyed penguin (*Megadyptes antipodes*) nest location methods,” *Wildl. Res.*, vol. 46, no. 2, p. 145, 2019.
 - [43] P. J. Seddon and R. F. Maloney, “Tracking wildlife radio-tag signals by light fixed-wing

- aircraft,” *Department Conserv. Tech. Ser.* 30, no. October, pp. 1–23, 2004.
- [44] L. K. Dressel, “Efficient and Low-Cost Localization of Radio Sources with an Autonomous Drone,” no. December, 2018.
 - [45] R. H. Mitch *et al.*, “Know your enemy: Signal characteristics of civil GPS jammers,” *GPS World*, vol. 23, no. 1, pp. 64–72, 2012.
 - [46] L. Meier, D. Honegger, and M. Pollefeys, “PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms,” in *2015 IEEE international conference on robotics and automation (ICRA)*, 2015, pp. 6235–6240.
 - [47] Stanford Artificial Intelligence Laboratory *et al.*, “Robotic Operating System.” .
 - [48] A. Perkins, L. Dressel, S. Lo, and P. Enge, “Antenna characterization for UAV based GPS Jammer localization,” in *28th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2015*, 2015, vol. 3.
 - [49] M. Singh and P. M. Khilar, “Actuating Sensor for Determining the Direction of Arrival Using Maximum RSSI,” vol. 3, no. 8, pp. 1–7, 2014.
 - [50] B. N. Hood and P. Barooah, “Estimating DoA From Radio-Frequency RSSI Measurements Using an Actuated Reflector,” *IEEE Sens. J.*, vol. 11, no. 2, pp. 413–417, Feb. 2011.
 - [51] A. Perkins, L. Dressel, S. Lo, and P. Enge, “Antenna pilots UAV,” *GPS World*, vol. 27, no. 2, 2016.
 - [52] B. Huber, “Radio Determination on Mini-UAV Platforms: Tracking and Locating Radio Transmitters,” 2009.
 - [53] M. Malajner, P. Planinsic, and D. Gleich, “Angle of arrival estimation using RSSI and omnidirectional rotatable antennas,” *IEEE Sens. J.*, vol. 12, no. 6, pp. 1950–1957, 2012.
 - [54] S. V. Schell and W. A. Gardner, “Higher-Resolution Direction Finding,” in *Handbook of Statistics, Vol. 10*, 1993, pp. 755–817.
 - [55] L. Dressel and M. J. Kochenderfer, “Pseudo-bearing Measurements for Improved Localization of Radio Sources with Multicopter UAVs,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6560–6565.
 - [56] G. E. Berz, “Tracking RFI: Interference localization using a CRPA : GPS World,” *GPS World*, 2017.
 - [57] A. Badawy, T. Khattab, D. Trinchero, T. El Fouly, and A. Mohamed, “A Simple AoA Estimation Scheme,” 2014.
 - [58] H. L. Van Trees, *Optimum Array Processing*, vol. 4. Institution of Engineering and Technology, 2013.
 - [59] H. L. Van Trees, “Planar Arrays and Apertures,” in *Optimum Array Processing*, New York,

- USA: John Wiley & Sons, Inc., 2002, pp. 231–331.
- [60] J. Hsiao, “Properties of a nonisosceles triangular grid planar phased array,” *IEEE Trans. Antennas Propag.*, vol. 20, no. 4, pp. 415–421, Jul. 1972.
 - [61] T. C. Cheston and J. Frank, “Phased array radar antennas,” *Radar Handb.*, pp. 1–7, 1990.
 - [62] B. D. Van Veen and K. M. Buckley, “Beamforming: a versatile approach to spatial filtering,” *IEEE ASSP Mag.*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
 - [63] A. Perkins, Y.-H. Chen, W. Lee, S. Lo, and P. Enge, “Development of a three-element beam steering antenna for bearing determination onboard a UAV capable of GNSS RFI localization,” in *30th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2017*, 2017, vol. 4.
 - [64] J. Knight, C. Cahn, and S. Nair, “A New Anti-Jamming Method for GNSS Receivers,” <https://Dokumen.Tips/Documents>, pp. 1–43, 2018.
 - [65] A. Perkins, L. Dressel, S. Lo, T. Reid, K. Gunning, and P. Enge, “Demonstration of UAV-Based GPS jammer localization during a live interference exercise,” in *29th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2016*, 2016, vol. 5.
 - [66] R. Pohlmann, S. Zhang, T. Jost, and A. Dammann, “Power-based direction-of-arrival estimation using a single multi-mode antenna,” *2017 14th Work. Positioning, Navig. Commun. WPNC 2017*, vol. 2018-Janua, pp. 1–6, 2018.
 - [67] M. O. A. Aqel, M. H. Marhaban, M. I. Saripan, and N. B. Ismail, “Review of visual odometry: types, approaches, challenges, and applications,” *Springerplus*, vol. 5, no. 1, p. 1897, Oct. 2016.
 - [68] J. Khalife, K. Shamaei, and Z. M. Kassas, “A software-defined receiver architecture for cellular CDMA-based navigation,” *Proc. IEEE/ION Position, Locat. Navig. Symp. PLANS 2016*, pp. 816–826, 2016.
 - [69] M. Rabinowitz and J. J. Spilker, Jr., “Synchronization Signals,” *Broadcast. IEEE Trans.*, vol. 51, no. 1, pp. 51–61, 2005.
 - [70] C. T. Ardito, J. J. Morales, J. J. Khalife, A. A. Abdallah, and Z. M. Kassas, “Performance evaluation of navigation using LEO satellite signals with periodically transmitted satellite positions,” *ION 2019 Int. Tech. Meet. Proc.*, pp. 306–318, 2019.
 - [71] D. Nistér, O. Naroditsky, and J. Bergen, “Visual odometry,” *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 1, 2004.
 - [72] D. Scaramuzza and F. Fraundorfer, “Tutorial: Visual odometry,” *IEEE Robot. Autom. Mag.*, vol. 18, no. 4, pp. 80–92, 2011.

- [73] S. Poddar, R. Kottath, and V. Karar, “Evolution of Visual Odometry Techniques,” 2018.
- [74] M. Dunbabin, J. Roberts, K. Usher, G. Winstanley, and P. Corke, “A hybrid AUV design for shallow water reef navigation,” *Proc. - IEEE Int. Conf. Robot. Autom.*, vol. 2005, no. April, pp. 2105–2110, 2005.
- [75] C. Aguilera, F. Barrera, F. Lumbreras, A. D. Sappa, and R. Toledo, “Multispectral image feature points,” *Sensors (Switzerland)*, vol. 12, no. 9, pp. 12661–12672, 2012.
- [76] A. Bharadwaj, A. Schultz, R. Gilabert, J. Huff, and M. U. De Haag, “Small-UAS navigation using 3D imager and infrared camera in structured environments,” *Proc. IEEE/ION Position, Locat. Navig. Symp. PLANS 2016*, pp. 599–606, 2016.
- [77] T. Mouats, N. Aouf, L. Chermak, and M. A. Richardson, “Thermal Stereo Odometry for UAVs,” *IEEE Sens. J.*, vol. 15, no. 11, pp. 6335–6347, Nov. 2015.
- [78] T. Mouats, N. Aouf, D. Nam, and S. Vidas, *Performance Evaluation of Feature Detectors and Descriptors Beyond the Visible*, vol. 92, no. 1. Journal of Intelligent & Robotic Systems, 2018.
- [79] FLIR, “FLIR Boson Thermal Imaging Core Datasheet.” 2019.
- [80] P. V. K. Borges and S. Vidas, “Practical Infrared Visual Odometry,” *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 8, pp. 2205–2213, Aug. 2016.
- [81] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s J. Softw. Tools*, 2000.
- [82] J. Shi and others, “Good features to track,” in *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, 1994, pp. 593–600.
- [83] F. Systems, “FLIR Vue Pro.” .
- [84] S. C. Nardone, K. F. Gong, and A. G. Lindgren, “Fundamental Properties and Performance of Conventional Bearings-Only Target Motion Analysis,” *IEEE Trans. Automat. Contr.*, vol. 29, no. 9, pp. 775–787, 1984.
- [85] B. Ristic and S. Arulampalam, *Beyond the Kalman filter : particle filters for tracking applications*. Artech House, 2004.
- [86] R. Munguía and A. Grau, “Concurrent initialization for bearing-only SLAM,” *Sensors*, vol. 10, no. 3, pp. 1511–1534, 2010.
- [87] J. M. C. Clark, R. B. Vinter, and M. M. Yaqoob, “Shifted Rayleigh filter: A new algorithm for bearings-only tracking,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 4, pp. 1373–1384, 2007.
- [88] S. Arulampalam, M. Clark, and R. Vinter, “Performance of the shifted Rayleigh filter in single-sensor bearings-only tracking,” in *2007 10th International Conference on Information Fusion*, 2007, pp. 1–6.

- [89] D. Mušicki, “Bearings only single-sensor target tracking using Gaussian mixtures,” *Automatica*, vol. 45, no. 9, pp. 2088–2092, Sep. 2009.
- [90] D. Musicki and R. Evans, “Measurement Gaussian Sum Mixture Target Tracking,” in *2006 9th International Conference on Information Fusion*, 2006, pp. 1–8.
- [91] N. Peach, “Bearings-only tracking using a set of range-parameterised extended Kalman filters,” *IEE Proc. - Control Theory Appl.*, vol. 142, no. 1, pp. 73–80, Jan. 1995.
- [92] D. L. Alspach and H. W. Sorenson, “Nonlinear bayesian estimation using gaussian sum approximations,” *IEEE Trans. Automat. Contr.*, vol. 17, no. 4, pp. 439–448, 1972.
- [93] A. C. Ozelc, “Tracking and Estimation Algorithms for Bearings Only Measurements Declaration of Originality,” Imperial College London, 2012.
- [94] H. Bayram, J. Vander Hook, and V. Isler, “Gathering Bearing Data for Target Localization,” *IEEE Robot. Autom. Lett.*, vol. 1, no. 1, pp. 369–374, 2016.
- [95] Y. Oshman and P. Davidson, “Optimization of observer trajectories for bearings-only target localization,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 35, no. 3, pp. 892–902, 1999.
- [96] R. P. S. Mahler, *Statistical Multisource-Multitarget*. 2007.
- [97] J. Sola, A. Monin, M. Devy, and T. Lemaire, “Undelayed initialization in bearing only SLAM,” in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2499–2504.
- [98] N. M. Kwok, Q. P. Ha, and G. Fang, “Data association in bearing-only SLAM using a cost function-based approach,” *Proc. - IEEE Int. Conf. Robot. Autom.*, no. May 2014, pp. 4108–4113, 2007.
- [99] R. M. Vaghefi, M. R. Gholami, and E. G. Ström, “Bearing-only target localization with uncertainties in observer position,” *IEEE Int. Symp. Pers. Indoor Mob. Radio Commun. PIMRC*, no. 3, pp. 238–242, 2010.
- [100] T. Lemaire, S. Lacroix, and J. Sola, “A practical 3D bearing-only SLAM algorithm,” in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2449–2454.
- [101] K. E. Bekris, M. Glick, and L. E. Kavraki, “Evaluation of algorithms for bearing-only SLAM,” *Proc. - IEEE Int. Conf. Robot. Autom.*, vol. 2006, no. May, pp. 1937–1943, 2006.
- [102] S. M. M. Dehghan and H. Moradi, “SLAM-inspired simultaneous localization of UAV and RF sources with unknown transmitted power,” *Trans. Inst. Meas. Control*, vol. 38, no. 8, pp. 895–907, Aug. 2016.