

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Efficient Drone-based Radio Tracking of Wildlife

Permalink

<https://escholarship.org/uc/item/4574s85j>

Author

Hui, Nathan

Publication Date

2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Efficient Drone-based Radio Tracking of Wildlife

A thesis submitted in partial satisfaction of the
requirements for the degree
Master of Science

in

Electrical Engineering (Intelligent Systems, Robotics, and Control)

by

Nathan Hui

Committee in charge:

Curt Schurgers, Chair
Ryan Kastner
Michael Yip

2019

Copyright
Nathan Hui, 2019
All rights reserved.

The thesis of Nathan Hui is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

DEDICATION

To the scientists that dedicate much blood, sweat, and tears to the furthering of
our understanding of the Earth.

EPIGRAPH

We have become, by the power of a glorious evolutionary accident called intelligence, the stewards of life's continuity on earth. We did not ask for this role, but we cannot abjure it. We may not be suited to it, but here we are.

—Stephen Jay Gould, *The Flamingo's Smile: Reflections in Natural History*

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Abbreviations	viii
List of Figures	ix
List of Tables	x
Acknowledgements	xi
Abstract of the Thesis	xii
Chapter 1	
Introduction	1
1.1 Background	1
1.2 Small Unmanned Aerial Systems (SUAS)	3
1.2.1 Types of Small Unmanned Aerial Systems (SUAS)	3
1.2.2 Control of SUAS	4
1.3 Using Drones in Wildlife Tracking	5
Chapter 2	
Tracking Iguanas using Drones	9
2.1 2015 Dominican Republic	10
2.2 2015 Dominican Republic and 2016 Cayman Islands	11
Chapter 3	
Payload System Overview	15
3.1 Wildlife Transmitters	16
3.2 Architectural Overview	18
3.3 Dipole Antenna	20
3.4 Radio Receiver	21
3.5 On-board Computer	23
3.6 Signal Processing	24
3.7 Location Estimation	26
3.8 Visualization	28
Chapter 4	
Flight System Overview	30
4.1 Flight System	30
4.2 Path Planning	32

Chapter 5	Experimental Results	36
	5.1 Operational Details	36
	5.2 Estimation Precision and Accuracy	38
	5.3 Estimation Speed - 2017	38
	5.4 Effective Range	39
Chapter 6	Continuing Improvements	42
	6.1 Double Bazooka Antenna	44
	6.2 On-board Computer	45
	6.3 Signal Processing	45
	6.4 Location Estimation	48
Chapter 7	Conclusion	50
	7.1 Comparison to Existing Platforms	50
Bibliography	53

LIST OF ABBREVIATIONS

<i>C. cornuta</i>	<i>Cyclura cornuta</i>
<i>C. n. caymanensis</i>	<i>Cyclura nubila caymanensis</i>
<i>C. ricordi</i>	<i>Cyclura ricordi</i>
E4E	Engineers for Exploration
FFT	Fast Fourier Transform
GIS	geospatial information system
GPIO	general-purpose input/output
ICR	Institute for Conservation Research
LNA	low noise amplifier
RC	radio controlled
RDF	radio direction-finding
RF	radio frequency
SDR	software defined radio
SNR	signal-to-noise ratio
SUAS	small unmanned aerial systems
UI	User Interface
VHF	Very High Frequency

LIST OF FIGURES

Figure 2.1:	2016 Estimate Visualization	12
Figure 2.2:	2015 and 2016 Radio Tracking Drone	14
Figure 3.1:	2017 Radio Tracking Drone	16
Figure 3.2:	Holohil BD-2 Wildlife Transmitters, in detail and attached to <i>Cyclura cornuta</i>	17
Figure 3.3:	Test Transmitter Configuration	18
Figure 3.4:	System Diagram	19
Figure 3.5:	Software v2.1 Signal Detector	24
Figure 3.6:	Zoomed in Waterfall Plot, showing pings	25
Figure 3.7:	Detection of Low SNR Pings	26
Figure 3.8:	Spatial Precision Distributions	28
Figure 3.9:	Mission 77 Results, 23 August 2017	29
Figure 4.1:	Lawnmower Flight Pattern	33
Figure 5.1:	2017 Certainty Distribution of Estimation Error	39
Figure 5.2:	2017 Deployment Operational Pace	40
Figure 5.3:	2017 Deployment Mission Time Distribution	41
Figure 5.4:	Mission 25 Results, 19 August 2017	41
Figure 6.1:	2019 Radio Tracking Drone	43
Figure 6.2:	Dynamic Threshold	46
Figure 6.3:	Software v3.0 Signal Detector	47

LIST OF TABLES

Table 2.1:	Hardware Version Comparison to 2015	10
Table 2.2:	Software Version Comparison to 2015	10
Table 2.3:	Deployment Summary to 2015	11
Table 2.4:	Hardware Version Comparison to 2016	12
Table 2.5:	Software Version Comparison to 2016	12
Table 2.6:	Deployment Summary to 2016	13
Table 3.1:	Hardware Version Comparison to 2017	15
Table 3.2:	Software Version Comparison to 2017	16
Table 5.1:	Deployment Summary to 2017	36
Table 6.1:	Hardware Version Comparison to Present	43
Table 6.2:	Software Version Comparison to Present	44
Table 6.3:	Deployment Summary to Present	44
Table 7.1:	Comparison of Performance	50

ACKNOWLEDGEMENTS

I would like to acknowledge the students, staff, and faculty associated with the UC San Diego Engineers for Exploration for providing technical assistance and support to this project. In particular, many thanks to Dr. Ryan Kastner, Dr. Curt Schurgers, Eric Lo, Eric Tran, Emmanuel Gharehbeekloo, Matthew Epperson, Brian Baxter, Katlin Dahn, Anthony Koutroulis, Jacob Domine, and Daniel Webber for their support and assistance in various aspects of this project.

I would also like to acknowledge our collaborators with the San Diego Zoo Institute of Conservation Research for their support and patience as we develop and deploy this project. In particular, a huge thanks to Dr. Glenn Gerber, Dr. Jen Moss, and Dr. Stesha Pasachnik.

Chapter 1, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 3, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 5, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 7, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

ABSTRACT OF THE THESIS

Efficient Drone-based Radio Tracking of Wildlife

by

Nathan Hui

Master of Science in Electrical Engineering (Intelligent Systems, Robotics, and Control)

University of California San Diego, 2019

Curt Schurgers, Chair

Radio telemetry is a critical technique in conservation ecology, particularly for studying the movement and range of individuals and populations. Traditionally, most radio telemetry work is done using handheld directional antennae by using either direction-finding and homing techniques, or radio-triangulation techniques. Over the past couple decades, efforts have been made to utilize aerial vehicles to make radio telemetry tracking more efficient, or cover more area. However, many these approaches require the use of manned aircraft and specialist skill sets. The proliferation of small unmanned aerial systems (SUAS) with high reliability and ease of use, as well as recent development and application of robotic sensing and estimation, opens up the possibility of leveraging SUAS to conduct radio telemetry studies. In this thesis, I present the

results of five years of development as well as the testing and deployment of a drone-based radio-telemetry tracking system that is able to track multiple targets simultaneously while operating in field conditions as part of a field expedition.

Chapter 1

Introduction

1.1 Background

Animal movements are an important metric in biology and ecology studies, which focus on protecting species and environments from excessive human impact. Such information is important for understanding migration and dispersal patterns, species interactions, and range models, among other aspects. Biologists and ecologists have developed a variety of techniques for tracking animals, including camera traps, isotope tracking, and radio tracking[1, 2, 3, 4, 5]. Of these, radio tracking can provide the most information about the animals location, as they allow the scientists to measure the precise location of the animal across their entire range.

A number of different tracking tags have been developed, including satellite tags, GPS tags, and Very High Frequency (VHF) tags. Satellite tags maintain a satellite link, which allow the tag to either upload GPS data, or provide the satellite with a radio beacon which can be localized. Such tags can be tracked remotely, without the need to deploy staff into the field. GPS tags will typically record GPS data, which then needs to be retrieved, either physically or through short-range wireless transmissions. While these do not automatically upload data, they are much smaller and typically have a longer life, as they do not use as much power. VHF tags only provide

a radio beacon for tracking, and so require field staff to manually track the tag.

The size of many animals, however, often limits the techniques that may be utilized. For instance, satellite and GPS tags are not yet universally applicable to movement studies, as these tags are often too large to fit to subject animals. As a result, tracking using VHF tags has become a very popular technique in ecological and conservation studies [6, 7].

Wildlife radio telemetry for tracking the exact location of an individual with VHF tags typically uses one of two methods - radio direction-finding (RDF), or triangulation. The simplest tracking technique is radio direction-finding (RDF). This technique uses a highly directional antenna to determine the direction of arrival of the transmitted signal. Typically, researchers will use a Yagi antenna, which can “hear” a transmitter when pointed within 2-10° of the source. Biologists can then start moving towards the source of the transmission to “home” in on the transmitter, and thus the tagged animal, using multiple signal detections along their path to realign to and fine tune the direction of arrival. However, this technique requires radio frequency (RF) line of sight to the transmitter, and thus the animal, which may not be feasible in heavy terrain such as slot canyons or hills. Additionally, RF reflections off the side of hills and canyons may confuse the signal, leading to incorrect inferences as to the direction of arrival of the transmitted signal. Furthermore, because users need to walk in the direction of signal in order to find the transmitter, this technique can become infeasible if the terrain or vegetation in the direction of arrival prevents passage [8, 9].

Another popular tracking technique is triangulation, which uses at least three bearing measurements to estimate the transmitter location. Because this technique only uses three separate measurements, ideally from different sides of the transmitter, this eliminates the need to move on foot through dense vegetation or difficult terrain. However, this technique still suffers from the limitations imposed by the requirement of RF line of sight. Furthermore, because we can only determine the direction of arrival with the Yagi antenna and analog receiver to within 2-10°, depending on the antenna and operator, we incur a large accuracy penalty if signal detections are

taken at significant range [8, 9].

Both of these techniques are conducted using a directional antenna and radio receiver, and require the user to be mobile. However, these methods are not ideal in environments which have dense foliage, impassable terrain, or are otherwise environmentally or technically hazardous to field staff [8].

1.2 Small Unmanned Aerial Systems (SUAS)

Over the past few decades, unmanned aerial systems, commonly known as drones, have become more capable and accessible to the general public. In particular, small unmanned aerial systems (SUAS) weighing less than 55 lbs have been growing in popularity and availability. The short learning curve and relative inexpensiveness of these systems make them very attractive to utilize as low-cost sensor platforms, as they can enable scientists to deploy sensors on a much larger scale.

1.2.1 Types of Small Unmanned Aerial Systems (SUAS)

SUAS can be classified into two categories: fixed wing, and rotary wing. Rotary wing systems can further be subcategorized into monorotor and multirotor systems. Of these three categories, the most popular SUAS have been multirotor systems, and are now one of the most ubiquitous “drones”.

Fixed wing SUAS have the same advantages as manned fixed-wing aircraft - they can typically fly faster, further, and longer than other aircraft. However, fixed wing SUAS also require significant area to operate in, as they require a “runway” for takeoff and landing. Although there exist several fixed wing SUAS that do not require a physical runway to operate from, they still require a similar amount of area to launch from and recover into. Fixed wing SUAS can also be capable of carrying large payloads. Because fixed wing aircraft generate the majority of their

lift from forward motion, they are, in general, more efficient in terms of energy per mass carried. They are also generally less expensive than other SUAS, as they have fewer electronic and moving components.

Monorotor SUAS are analogous to conventional manned helicopters. They typically use one main rotor for lift, and a second smaller rotor for yaw control. Some variations may use a second rotor coaxial to the main rotor for yaw control. Like their manned counterparts, monorotors do not require a runway to operate from - they can take off and land vertically, and can easily maneuver around obstacles. However, while monorotors are more maneuverable, they generally have slightly less payload capacity and flight time, as they put more of their energy into staying in the air. Monorotor aircraft are typically the most expensive SUAS, as they are mechanically complex.

Multirotor aircraft, also known as multicopters, are the most well known SUAS today. A large majority, known as quadcopters, consist of four rotors, placed so their thrust vectors point down. Other variants may use three, six, or eight rotors, and may also stack rotors vertically. Similar to monorotors, multirotors do not require runways to operate from. They have similar maneuverability to monorotor aircraft, but are significantly cheaper and easier to maintain due to their mechanic simplicity. However, unlike monorotors, which use a rotary wing for lift, multirotors primarily use vectored thrust for lift, which results in significantly reduced payload capacity and flight time. Multirotors today are typically some of the cheapest options due to their relative simplicity and ubiquity. For this reason, we have chosen a quadcopter for our current mission profile, discussed in Chapter 4, as it best fits our mission profile, and reduces the cost and complexity of our system.

1.2.2 Control of SUAS

Although radio controlled (RC) aircraft have been around since the mid 1900s, drones in the sense of a self-contained automated aircraft have only really come about in the past decade.

This is in large part due to the development of suitable autopilots that are small, cheap, and reliable enough to control RC aircraft.

SUAS control typically comes in one of three modes: computer stabilized, semi-autonomous, and fully autonomous. Computer stabilized modes allow the pilot the most amount of control over the aircraft, and simply attempt to make the aircraft follow the pilot inputs. Fully autonomous modes enable the pilot to become a systems manager, simply telling the SUAS where to go. Semi autonomous modes make up the spectrum in between, where the computer may control certain aspects of the flight envelope, whether it be maintaining altitude and position, or simply providing a keep-away zone around the aircraft.

1.3 Using Drones in Wildlife Tracking

Drones provide an easy solution to avoid hazards or impediments due to terrain and foliage - they escape the environment in which these hazards or impediments exist. In general, dense foliage or impassable terrain only affect land mobility - they do not affect aerial mobility. Thus, flying sensors over these areas provide an avenue for avoiding the obstacles presented by foliage and terrain, which are common obstacles in many biological and ecological field sites. In addition, flying sensors high above the ground can increase the effective range of those sensors, as they have a larger view of the survey area. This makes surveys done with flying sensors much more efficient.

Some scientists have used manned aircraft to conduct wildlife tracking surveys. In general, these surveys are conducted by attaching directional antennae to the outside of a fixed-wing aircraft and flying radio direction-finding (RDF) missions. Much like conducting radio direction-finding (RDF) tracks on foot, this requires RF line of sight, however, this is much easier to manage because the receiver is now in the air, and so has RF line of sight to far more places on the ground. In addition, because the aircraft is flying, the presence of dense vegetation does not impact the

ability of the field researcher to conduct the tracking, only affecting the ability of the researchers to actually do visual verification of the target [8, 9].

The primary issue with manned aircraft tracking surveys is the increased cost and logistics needed to support such surveys. Fixed wing aircraft surveys require pilots, aircraft, parking spaces, fuel, and maintenance, among other logistics and support items. Survey time would be limited to the aircraft's loiter time minus transit time from the support airfield to the survey area, potentially making manned aircraft surveys in particularly remote areas infeasible from a cost and logistics standpoint[8].

Since 2015, many researchers have developed low-cost drone based tracking systems that can avoid the logistical issues mentioned earlier. The advantages of using drones is that they can move across the survey area faster than ground surveys and with more precision than manned aircraft. They can make more informative and precise measurements than a human with an analog receiver, which allows us to use more information to generate a more precise estimate. In addition, as mentioned in earlier, flying sensors often increases their effective range. In the case of a radio receiver, flying the sensor allows RF line of sight to more areas on the ground, which reduces the impact of terrain and vegetation on the survey.

These systems generally approach the estimation problem in one of two ways - range based [10, 11, 12, 13] or bearing based [14, 15, 16, 17]. There are some approaches that take both into account [18, 19]. In general, range based systems tend to be easier to build, as they use a simpler antenna configuration. Bearing based systems require complex antenna configurations, and are generally physically larger.

Bearing based estimation seeks to emulate the traditional terrestrial approaches with a drone. These systems typically take one of two approaches to determine the bearing to the transmitter: using either a highly directional antenna and rotating to detect the signal, or an antenna array that can measure the direction of arrival of the signal. Because of this, bearing based estimation typically requires complex antenna or receiver configurations, and are typically

physically larger. This affects the suitability of such systems for field use, as larger systems are more difficult to utilize in the field, and often require more maintenance, as there are more components that can break.

Range based estimation utilizes the precision with which we can measure the signal strength of the received signal and the relationship between received signal strength and distance to the transmitter. This can use a much smaller and simpler antenna configuration, which increases its usability and robustness while decreasing the overall cost and complexity of the system. Some of these range based approaches also take the directionality of the antenna into account by mapping the directionality of the antenna to the range estimates [18, 19].

Many of the recent bearing based and range-based systems use estimation approaches such as particle filters [12, 18, 20], grid filters [17, 16], and Kalman filters [10]. These are all variations of Bayesian estimation, and rely on having an accurately characterized probability distribution of the source of noise from observations and sensors. These take advantage of knowing the probability of a particular measurement to determine the probability of a particular estimate being correct.

Few, if any, of the recent prototype drone tracking systems have been robustly tested under field conditions. A survey of the most mature systems (Cliff, Dressel, and Nguyen) shows that the majority of the systems have less than 20 field trials. Of these, the most precise system estimates the location to a 5 m cell, but only gives a 50 % certainty that the transmitter is located within that cell. The remaining two systems (Cliff and Nguyen) generate estimates to within 20 m. Both Dressel and Nguyen quote localization times of less than 5 min for trial flights; however, these were done with the copter starting within detection range of the transmitter, and not in field conditions, where often, the drone will start far from the subject animal [19, 16, 20].

In this thesis, I present the development and deployment of a robust and rugged drone system for radio tracking wildlife that is field deployable, and that outperforms existing aerial systems used for wildlife radio tracking. In particular, I will focus on the design and fielding

of the system deployed in 2017, which is a simple range-based estimator that uses a software defined radio (SDR) to record data that we then postprocess to generate the estimate. This system performs well during field tests where we show it to generate estimates with 16 m precision in 10 minutes of flight and processing time, which is on par with current drone-based tracking systems and far better than traditional ground-based tracking methodologies.

Chapter 2 discusses the design choices and results of previous design iterations and deployments, and highlights the specific design implementations that were attempted along with their drawbacks as found during field deployments. Chapter 3 discusses the structure, design choices, and development of the sensor payload in the 2017 design. Chapter 4 discusses the flight system selection and design choices for the 2017 deployment. Chapter 5 presents the quantitative and qualitative results of the 2017 deployment. Chapter 6 presents the goals and design considerations for work to be continued from the 2017 design. Finally, Chapter 7 compares the performance of the 2017 design and the system architecture to existing designs and systems.

Chapter 1, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 2

Tracking Iguanas using Drones

This project is a collaboration between the San Diego Zoo Institute for Conservation Research (ICR)¹ and UC San Diego's Engineers for Exploration (E4E)². The San Diego Zoo ICR is an internationally renowned research center focused on conservation research in a wide variety of ecosystems around the world. E4E is a student-run engineering group on campus that assists scientists in conducting field research by leveraging technology to automate or enhance those surveys.

In 2012, Dr. Glenn Gerber and Dr. Stesha Pasachnik from the ICR approached E4E with the idea to develop a drone to track iguana hatchlings through difficult terrain on various Caribbean islands, including the Dominican Republic and the Cayman Islands. The specific challenge the scientists faced in the Dominican Republic during the 2015 deployment was that the terrain consisted of limestone valleys full of cactus, which severely impeded movement on foot to follow the hatchlings. During the Cayman Islands deployments in 2016 and 2017, the challenge faced was that the areas the iguanas preferred to move into consisted mostly of thick vegetation or swamps, making tracking difficult. On each of these deployments, the scientists' objective was to track the subject animals twice a day, for as long as two weeks, in order to determine the way

¹<https://institute.sandiegozoo.org/>

²<http://e4e.ucsd.edu>

that the iguana hatchlings dispersed from their nesting sites to their home ranges.

The San Diego Zoo scientists tracked iguanas on the islands using small VHF radio tags that periodically transmit a short pulse. Each radio tag transmitted on a unique frequency. The scientists then use a combination of RDF and triangulation methods to track the animals over several days, often tracking each individual as much as twice a day to ensure good scientific data. These tracks were then compiled into maps showing the dispersion of the hatchlings as they moved from their nest sites into their home ranges.

We decided to develop a drone system that would measure the difference in received signal strength as we flew around the radio tag. Since we know that the strength of the received signal should decrease as a function of distance to the transmitter, we can use the way the signal strength decays to determine the location of that transmitter.

As we developed this system through the E4E program, we took as many opportunities as were available to test our system locally, as well as abroad during field expeditions. As we gained more experience with the design, we continuously found areas for improvement, which we iterated on. We discuss the 2015 and 2016 deployments and designs below to provide context for the development and deployment of the 2017 version, described in Chapters 3 and 4. Current developments are described in Chapter 6.

2.1 2015 Dominican Republic

Table 2.1: Hardware Version Comparison to 2015

Version	Year	Airframe	Computer	Receiver	Antenna
Rev A	2015	Tarot 600	BeagleBone Black	RTL-SDR	Dipole v1

Table 2.2: Software Version Comparison to 2015

Software Version	Signal Detector	Estimator	Visualization	Processing
v1.0	FFT & Loudest Signal	Average	Heatmap	Post-process

Table 2.3: Deployment Summary to 2015

Year	Location	Duration	Hardware	Software
2015	Dominican Republic	10 days	Rev A	v1.0

In the summer of 2015, we had an opportunity to deploy our system to the Dominican Republic to support Dr. Stesha Pasachnik (San Diego Zoo Institute for Conservation Research) in studying the dispersal patterns of *Cyclura cornuta* hatchlings. This initial system (hardware revision A, software v1.0) built on algorithms developed by dos Santos et al [13]. This consisted of a custom Tarot 600 airframe controlled with an ArduPilot autopilot, with a payload comprised of a BeagleBone Black³ recording data from an RTL-SDR⁴ paired with a tuned dipole antenna. We then calculated the transmitter locations from the recorded data offline using algorithms developed by dos Santos et al in [13]. Tables 2.1, 2.2, and 2.3 summarize the system configuration and deployment.

We chose to use a custom airframe for this system because of a lack of robust and mature commercially available quadcopters in 2015. We selected the BeagleBone Black and RTL-SDR as the core of our payload because of their cost and availability.

During this deployment, we found that the BeagleBone Black did not perform well enough in that environment due to thermal issues. We found that the BeagleBone Black would overheat, reduce its operating frequency, start losing samples because it could not keep up, and eventually crash. We also found that the implementation of the post-processing was flawed, as it would routinely lose samples and would miscalculate the transmitter location estimate.

2.2 2015 Dominican Republic and 2016 Cayman Islands

To address the problems encountered in the first 2015 Dominican Republic deployment, we rebuilt the system. We replaced the BeagleBone Black with a Raspberry Pi 3 and refined the

³<https://beagleboard.org/black>

⁴<https://www.rtl-sdr.com/>

Table 2.4: Hardware Version Comparison to 2016

Version	Year	Airframe	Computer	Receiver	Antenna
Rev A	2015	Tarot 600	BeagleBone Black	RTL-SDR	Dipole v1
Rev B	2016	Tarot 600	Raspberry Pi 3	RTL-SDR	Dipole v1

Table 2.5: Software Version Comparison to 2016

Software Version	Signal Detector	Estimator	Visualization	Processing
v1.0	FFT & Loudest Signal	Average	Heatmap	Post-process
v2.0	FFT & Loudest Signal	Manual	GIS / Heatmap	Post-process

post-processing software. This software took the recorded signal and isolated the transmitter frequency using Fast Fourier Transforms. Since the radio tag transmits pulses at regular intervals, we then identified the pulse as the loudest amplitude in each ping period. We then plotted that amplitude at the position the drone was in at that moment, and estimated the position as the “hottest” part of the map, as shown in Figure 2.1. These changes are summarized in Tables 2.4, 2.5, and 2.6.

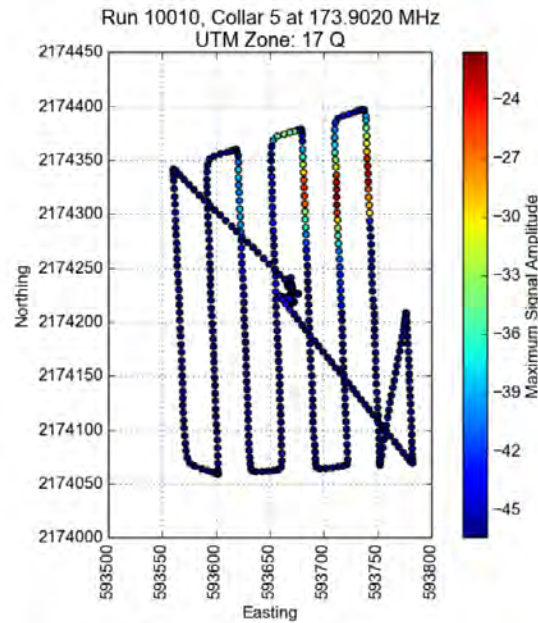


Figure 2.1: 2016 Estimate Visualization

We were able to redeploy this new system (hardware revision B, software v2.0) in the

Table 2.6: Deployment Summary to 2016

Year	Location	Duration	Hardware	Software
2015	Dominican Republic	10 days	Rev A	v1.0
2015	Dominican Republic	10 days	Rev B	v1.0
2016	Cayman Islands	10 days	Rev B	v2.0

fall of 2015 to the Dominican Republic to support Dr. Stesha Pasachnik’s survey of *Cyclura ricordi* hatchlings, and again in the summer of 2016 to the Cayman Islands to support Dr. Jen Moss’s (Welch Lab, Mississippi State University) survey of *Cyclura nubila caymanensis* adults. This system, shown in Figure 2.2, was comprised of a Raspberry Pi 3 recording data from an RTL-SDR paired with a tuned dipole antenna, and data post-processed on a laptop, as described above.

Over the two deployments, we found that the estimation of the transmitter position was not very robust. Estimating the transmitter position by simply looking at a plot was generally slow, and involved some amount of guesswork. Additionally, trying to do the estimation by hand on top of sorting the rest of the day’s tracking data was not operationally feasible for a full scale deployment of the system. We also found that the Tarot 600 airframe was not suited to field use, as it did not fly well and was too bulky for use and transport in the field. Lastly, we found the detection range of the system was particularly low, at around 30 m.



Figure 2.2: 2015 and 2016 Radio Tracking Drone

Chapter 3

Payload System Overview

Based on lessons learned from the 2015 and 2016 deployments, we made several changes to address these changes in preparation for the 2017 Cayman Islands deployment. We developed a model based estimator to determine the transmitter position based on the pings that we received, which will be explained in Section 3.7. We also switched from the Tarot 600 airframe to the 3DR Solo, which we explain in Section 4.1. Finally, we changed out the RTL-SDR and Raspberry Pi for a USRP B200mini and Intel Joule, which we explain further in Sections 3.5 and 3.4. These changes are summarized in Tables 3.1 and 3.2.

The 2017 drone tracker is comprised of two independent systems - the flight system, and the sensor payload. We elected to make this separation to permit the future move to different flight platforms, and to ensure that failures in the sensor payload do not propagate into the flight platform and cause further issues. In this chapter, we explain the design choices and composition of the payload system.

Table 3.1: Hardware Version Comparison to 2017

Version	Year	Airframe	Computer	Receiver	Antenna
Rev A	2015	Tarot 600	BeagleBone Black	RTL-SDR	Dipole v1
Rev B	2016	Tarot 600	Raspberry Pi 3	RTL-SDR	Dipole v1
Rev C	2017	3DR Solo	Intel Joule	USRP B200mini	Dipole v2

Table 3.2: Software Version Comparison to 2017

Software Version	Signal Detector	Estimator	Visualization	Processing
v1.0	FFT & Loudest Signal	Average	Heatmap	Post-process
v2.0	FFT & Loudest Signal	Manual	GIS / Heatmap	Post-process
v2.1	FFT & Loudest Signal	Model-fit	GIS	Post-process



Figure 3.1: 2017 Radio Tracking Drone

3.1 Wildlife Transmitters

We are tracking VHF wildlife transmitters, in particular, the Holohil BD-2 and PD-2 transmitters¹. Figure 3.2 shows the BD-2 transmitter in detail, and attached to a *C. cornuta* hatchling in the Dominican Republic. These transmitters are all very small, in general less than 2 g. As a result of their size and weight, they do not have a very high output power, on the order of 1-10 mW. These are typically configured to transmit a 10-20 ms pulse on a specific frequency every 1-2 s in a way that maximizes their battery life over the intended field life.

One of the challenges of working with these transmitters is the variability in transmission characteristics. Because of the simplicity of these devices, the transmit frequency will vary as the battery voltage decreases and the temperature varies. Field notes suggest a drift due to battery depletion of as much as 1 kHz over a period of 1 week. Lastly, as the transmitter ages, the transmission power decreases, which makes these transmitters more difficult to detect.

We chose to use these exact transmitters to test the performance of the tracking system,

¹<http://www.holohil.com/>




Figure 3.2: Holohil BD-2 Wildlife Transmitters, in detail and attached to *Cyclura cornuta*

as they are more representative of the expected power output than handheld radios or a SDR analogue. In addition, by using these transmitters, we can see the effect of directionality and occlusion from the orientation and location of the transmitter, as well as the effectiveness of our signal processing and detection.

We primarily tested with a set of transmitters with internal batteries obtained from the San Diego Zoo prior to the 2017 Cayman Islands deployment, as those were the transmitters used in the Cayman Islands study. After that deployment, we requested a pair of transmitters that could be attached to external batteries from Holohil, as shown in Figure 3.3. This allowed us to more directly control the output power, and eliminate the variable of transmitter battery age from testing.




Figure 3.3: Test Transmitter Configuration

3.2 Architectural Overview

We have chosen to use received signal strength to drive our estimation of the transmitter location as it is the simplest and potentially yields the more robust system. Fundamentally, we can measure two pieces of information about the signal pulse: the direction of arrival, and the signal strength. With either of these pieces of information, we can estimate the location of the transmitter by using triangulation or trilateration techniques.

As mentioned in Section 1.3, determining the direction of arrival of a signal typically requires either a directional antenna or an antenna array. However, because most radio tags are in the 1-2 m band, the antenna for the receivers is rather large, forcing the system to be physically large as well. This is infeasible for deployment on a drone intended for use in austere environments, as the size and weight of such an antenna system would severely impact the drone's range, and the size of the antenna would make transporting the drone to its launch point impractical.

Determining the signal strength of a ping simply requires an analog to digital converter. Determining the distance from the ping to the transmitter then requires knowing the gain in the system, the directionality of the antenna, and the transmit power of the transmitter, as discussed later in Section 3.7. However, we know that the received signal strength of a transmission decays proportional to the distance to the transmitter. This is independent of the directionality of the

antenna, transmission power, and system gain. In addition, this approach benefits from having an antenna with as little directionality as possible, which aligns with the constraint that this system should be as small as possible so as to be field friendly.

To accomplish this, we break the payload system down into several subsystems: antenna, low noise amplifier (LNA), software defined radio (SDR), GPS/compass, on-board computer, data storage, processing pipeline, and visualization tools. We show the overall system architecture in Figure 3.4.

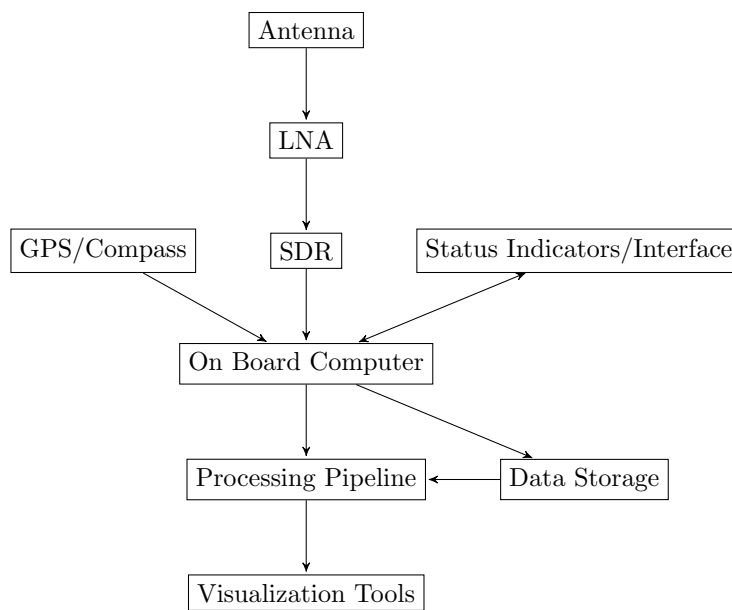


Figure 3.4: System Diagram

In this system, the antenna receives the RF ping from the transmitter. This signal is amplified by the LNA, which is then amplified yet again and digitized by the SDR. The on-board computer records this digital signal from the SDR, along with GPS and heading information from the GPS/compass unit, and stores all of this into external storage. We then feed all of the recorded data into our processing pipeline, which detects all of the pings, then estimates the location and certainty of the transmitter location. We can then visualize the information in geospatial information system (GIS) software. These subsystems are explained in further detail below.

3.3 Dipole Antenna

Our system uses the change of received signal strength as the drone moves around the transmitter to determine the location of the transmitter. In order for this to be effective yet simple, we need an omni-directional antenna so that the received signal strength is independent of the orientation of the survey platform.

We elect to use a dipole antenna because of its near omni-directional characteristics. In its horizontal orientation, with the nulls pointed to the horizon, the projected beam pattern on the ground is nearly circular, which eliminates attenuation of the target signal due to directionality. This allows us to treat the received power of the transmit signal as a direct proxy for the distance between the transmitter and receiver. This simplification allows for a much simpler computational model, and opens up the possibility of using this exact sensor payload on a fixed wing aircraft, where we cannot as easily control the orientation of the antenna with respect to the flight path of the aircraft.

In addition to its omni-directional characteristics, the dipole is known to have a relatively wide bandwidth, which is ideal for our purposes because it allows us to receive a wide variety of signals without significant losses. Since the radio tags are assigned unique frequencies for identification, this can result in a significant range of frequencies that need to be received, which is achievable with a dipole antenna.

This antenna is also extremely simple to design and build. This is important because of the anticipated field use, where we expect that at some point, it will be necessary to repair the antenna in the field. Having a simple design with few critical dimensions allows for a simpler repair process.

Additionally, because of the simplicity of the dipole antenna, the resulting design is lightweight, which permits its integration onto small drone platforms. This is particularly important due to the requirement that this be operable in field conditions, where the SUAS will

often be transported in a vehicle with a lot of additional equipment, or on foot through dense foliage.

In the 2017 deployment, the radio tags were transmitting on frequencies around 172 MHz. We first built a half-wave dipole for that band, and later refined the design² [21]. This was then encased in a fiberglass and 3D printed enclosure to protect from damage during transport and flight.

We initially mounted the antenna under the motor pods as shown in Figure 3.1 to maximize the vertical clearance of the drone as it landed. However, field tests suggested that the proximity to the high current wiring in the motor pods induced noise in the signal, and so we moved the antenna to the legs of the drone, which helped reduce some of the electrical noise.

3.4 Radio Receiver

In order to detect the signal from the dipole antenna, we need to amplify the signal, particularly because the radio tags are very small and do not transmit with a lot of power. To do this, we added the LNA4ALL low noise amplifier (LNA)³ to amplify the signal. We selected the LNA4ALL as it provided a clean and packaged LNA in the VHF band with a noise figure on par with other prepackaged LNAs.

In order to convert the analog RF signal to a digital signal we can process, we use a software defined radio (SDR). The SDR operates by first amplifying the induced voltage in the antenna to measurable levels, then shifting the signal from the radio frequency f_c to baseband (0 MHz). This signal is then sampled at a sampling frequency f_s in both the real (in-phase) and complex (quadrature) components to generate the digital IQ signal that is then passed to the onboard computer for processing.

²Acknowledgements to Daniel Webber for his work during the 2016 UCSD E4E Research Experience for Undergraduates

³<http://lna4all.blogspot.com/>

We specifically use a SDR as opposed to the conventional analog receiver that is commonly used in radio tracking because the SDR allows us to receive multiple frequencies simultaneously, which opens up the possibility of tracking multiple transmitters during any given survey. In addition, the SDR allows us to control the center frequency, sampling rate, gain, and resolution via software commands. This allows us to rapidly reconfigure the system to scan for different types of transmitters (VHF or UHF) without making any hardware changes, whereas using a conventional receiver would require a new receiver for VHF and UHF, and even potentially different receivers for different frequencies within VHF/UHF ranges.

We previously used the RTL-SDR TV tuner dongle as an SDR because it was an inexpensive and physically robust option, however, we felt that we could get significantly better signal performance by moving to a purpose built SDR. The Ettus Research USRP B200mini⁴ SDR boasts a better built-in LNA, higher resolution, and a lower noise figure than the RTL-SDR. This is advantageous in many ways: the better LNA allows us to pick up fainter signals, the higher resolution allows us to resolve those fainter signals, and the lower noise figure results in a higher overall signal-to-noise ratio (SNR). The USRP also appears to be built sufficiently well that it does not present a very large DC spike. This is an issue with almost all SDRs where the received signal will always have some amount of DC bias due to mismatch in the in-phase and quadrature sampling converters.

We chose to set the center frequency f_c to 172.5 MHz and the sampling frequency f_s to 2 MS s^{-1} . Since we were only recording data to disk, our system was I/O bound. We selected $f_s = 2 \text{ MS s}^{-1}$ and $f_c = 172.5 \text{ MHz}$ so that the bandwidth of the SDR covered the frequencies of the transmitters we were interested in and avoided the DC spike. The SDR receives complex signal data, which allows us to determine the sign of frequencies, so the signal bandwidth is equivalent to the sampling frequency, thus allowing us to be able to receive frequencies from 171.5-173.5 MHz.

⁴<https://www.ettus.com/all-products/usrp-b200mini/>

3.5 On-board Computer

In order to record the RF and position data for postprocessing, we need a computer that can interface with the USRP SDR and the GPS unit. We previously used a Raspberry Pi 3, which is a ARM core single board computer. However, the USRP Hardware Driver library requires an x86 processor, and the USRP requires a USB 3.0 connection.

For the 2017 revision of the payload hardware, we elected to use the Intel Joule⁵, which is a compact x86 single board computer, as the onboard computer. The Intel Joule comes with a breakout board that exposes USB 3.0, hardware serial, general-purpose input/output (GPIO), and external storage. This allows us to interface directly with the USRP, GPS unit, and external storage. In addition, the GPIO allow us to interface with a custom User Interface (UI) Board, which provides a way for the user to signal the computer to start recording and to check the status of the various systems on board.

The Joule is configured to start a suite of monitoring software on boot, which allows it to monitor and manage the status of the various sensors and subsystems and display those on the UI Board. This software suite is also responsible for starting the recording software and marshalling the data into the appropriate locations on an external storage device.

The USRP streams the complex IQ data to the Intel Joule over a USB 3.0 connection. We use the USRP Hardware Driver library to unpack the stream and record the raw data directly to disk as sequential complex integers. We cap each file at 64 MB of data, or 1024 frames of 16384 samples (8 s of data), to preserve data integrity in the event that the payload loses power or the software decides to crash.

In order to synchronize the pings as recorded by the drone to the position where the drone recorded the ping, we need to also record some metadata. We record the timestamp of the first sample, sampling rate, center frequency, and SDR LNA gain as a metadata file alongside the RF data. The single timestamp for the first sample is enough to synchronize the data streams, as the

⁵<https://ark.intel.com/content/www/us/en/ark/products/96421/intel-joule-550x-compute-module.html>

clock for the SDR has a drift on the order of 1 ppb, and so is not expected to cause significant drift over the course of the mission time.

The onboard computer also records GPS data to a file. We record the local timestamp, GPS timestamp, position, velocity, and heading from the GPS into a comma separated values file alongside the RF data and metadata files.

3.6 Signal Processing

In order to estimate the location of the transmitter, we need to first identify and measure the individual pings as recorded by the payload computer. We chose to design the system to post-process the data for two reasons: first, the existing signal processing was fairly resource intensive, and would require either significant hardware or redesign to make real-time, and second, we wanted to record all of the signal data to use as test data for future iterations.

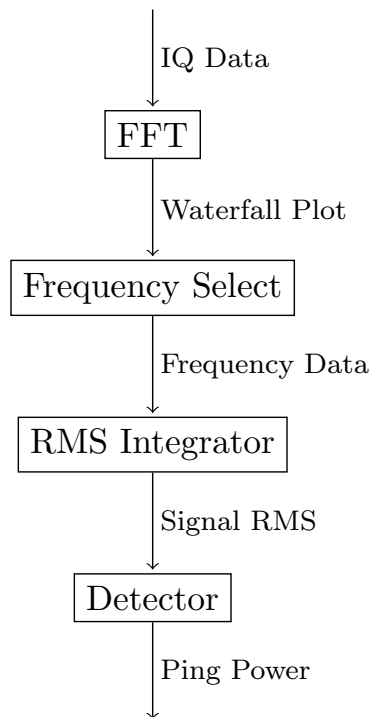


Figure 3.5: Software v2.1 Signal Detector

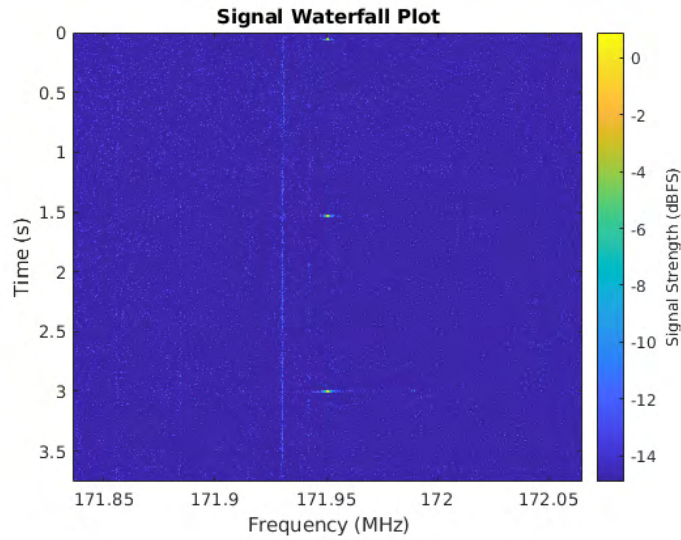


Figure 3.6: Zoomed in Waterfall Plot, showing pings

To process the radio frequency data, we first extract the specific frequencies that we know transmitters to be on. We can think of a Fast Fourier Transform (FFT) of length n as a bank of n adjacent band-pass filters. Each filter has a passband centered at $f_c + \left(\frac{i}{n}f_s\right)$ for the i th filter with width $\frac{f_s}{n}$, where i is in $\left[\frac{-n}{2}, \frac{n}{2}\right]$. We start by running a FFT over the entirety of the data to convert the time domain signal to a waterfall plot (time vs frequency vs amplitude), a sample of which is shown in Figure 3.6. We do an FFT of 4096 elements, which gives us a frequency resolution of roughly 500 Hz per bin. Quantitative tests show that the transmitted signal has a bandwidth of less than 500 Hz, which results in us being able to positively identify each ping to the correct frequency. The resulting signal also has a much higher SNR as it is rejecting the remaining frequencies.

Once we have the waterfall plot, we perform a sum over the square of the magnitude of the past 60 ms for the frequency we are interested in, as shown in the right image in Figure 3.7. Because the transmitted signal is 40-60 ms long, this is the root mean square measure of the signal power over that window. This makes the pings, which are otherwise not visible in the left image showing instantaneous power, visible in the right image.

Finally, to identify each ping, we simply identify the maximum power in each 1.5 s

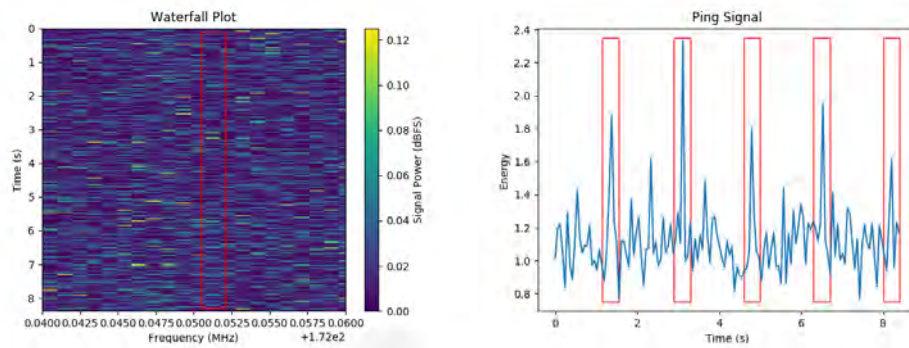


Figure 3.7: Detection of Low SNR Pings

window. Since the transmitted signal has a roughly 1.5 s period, we assume that the loudest signal in the frequency in each 1.5 s window is the ping, if we heard the signal at all.

This approach in general works, provided that there are no other transmissions on the same frequency, and that the transmitter frequency avoids the spike caused by the DC bias in the some SDRs (see Section 3.4 for more details).

3.7 Location Estimation

In order to localize the transmitter, we need to associate each ping with the location from which the drone heard said ping. Both the radio data and position data have local system timestamps (see Section 3.5), which allow us to synchronize the two data streams.

As we identify pings, we fit them to a signal propagation model. We know that the amplitude of a radio signal decays proportional to the inverse square of distance in a vacuum. In practice, the exponent for the path loss can range from 2 to 6, depending on the environment, so we elect to use the model shown in Equation 3.1, where L is the loss in dB, n is the path loss exponent, d is the path length, and C represents additional system loss.

$$L = 10n \log_{10}(d) + C \quad (3.1)$$

$$R = P - L = P - 10n \log_{10}(d) - C \quad (3.2)$$

$$k_1 = \frac{-1}{n} \quad (3.3)$$

$$k_2 = \frac{P - C}{n} \quad (3.4)$$

$$d = 10^{\frac{k_1 R + k_2}{10}} \quad (3.5)$$

$$d = \sqrt{(x_d - x_t)^2 + (y_d - y_t)^2 + (z_d)^2} \quad (3.6)$$

Since our measurement consists of the received signal strength R in dB and location (x_d, y_d, z_d) in meters, we need to rewrite Equation 3.1 with respect to the received signal strength and transmitter power P , shown in Equation 3.2. This is particularly important as the transmitter power varies, as mentioned in Section 3.1. To facilitate computation, we reparameterize Equation 3.2 with k_1 and k_2 , shown in Equations 3.3 and 3.4. This reparameterized model (shown in Equation 3.5) models the distance d in meters between the transmitter and the drone as a function of the two model parameters k_1 and k_2 , and the received signal power R in dB. We calculate d using Equation 3.6, where (x_d, y_d, z_d) is the drone location for that ping in meters, and (x_t, y_t) is the location of the tag on the ground in meters. This allows us to fully parameterize the model to best account for the variable transmitter power P and other system intrinsic while minimizing the computational complexity of the model.

Since we are solving not only for the transmitter location in two dimensions, but also for the RF signal parameters, we have a model with four parameters. In order to find a solution for these parameters, we need at least four measurements, more if we wish to characterize the accuracy of our estimate. We then use a non-linear least squares approach to find parameters that best fit the measurements.

Once we generate the model parameters, we still need to provide a measure of the precision

of the estimate. To do this, we overlay an estimate of the probability of the transmitter position for each ping. This consists of a normal distribution centered on the distance calculated by the model parameters for that ping with a standard deviation of 40 % of the distance, rotated around the ping's receive location. The hot spot resulting from the sum of all the pings represents a measure of the probability contour for the location of the transmitter. These distributions are shown graphically in Figure 3.8.

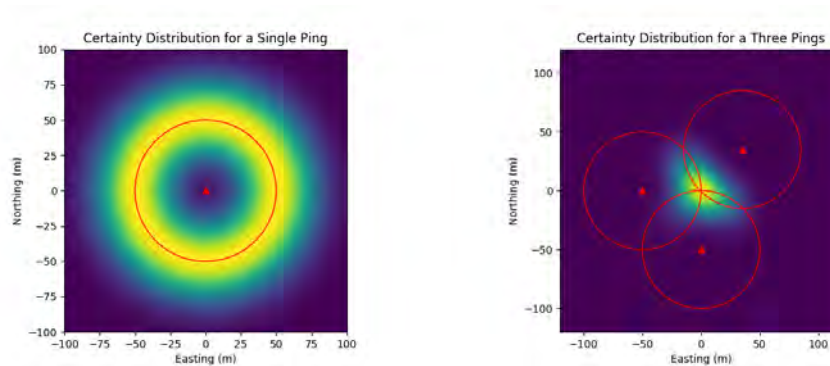


Figure 3.8: Spatial Precision Distributions

While other systems have used Bayesian estimation to solve for their estimate, we feel that using a non-linear least-squares solver keeps this system simpler. Bayesian estimation requires a prior characterization of the noise in the system, and in general is more computationally expensive, which can be unfeasible for deployment on small drones due to weight limitations.

3.8 Visualization

We can visualize the data as geospatial information system (GIS) data, plotting the heatmap of probable locations, as well as the location with the highest probability of location. This is provided as a GeoTIFF image containing the probability map, and an ESRI Shapefile containing the highest probability of location. This visualization and data format provides a nearly seamless method for the scientists to incorporate our measurements into their work. An example of the resulting data visualization is shown in Figure 3.9.

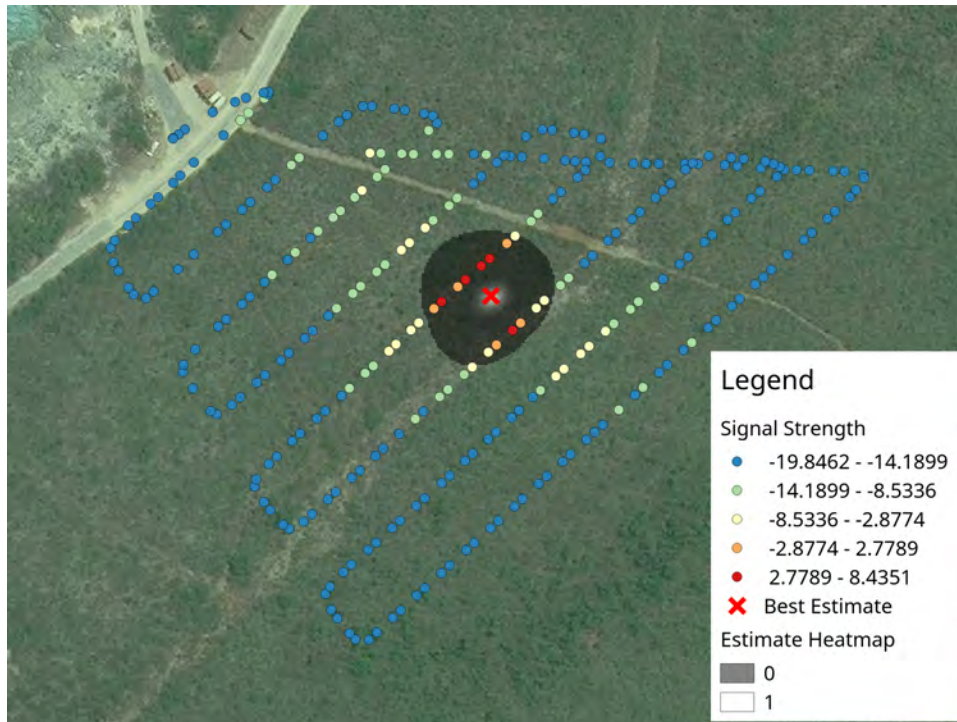


Figure 3.9: Mission 77 Results, 23 August 2017

Chapter 3, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Chapter 4

Flight System Overview

As mentioned in Chapter 3, the drone tracker is comprised of two independent systems - the flight system, and the sensor payload. We elected to make this separation to ensure that failures in the sensor payload do not propagate into the flight platform and cause further issues, and to permit the future move to different flight platforms. In this chapter, we explain the design choices and composition of the flight system for the 2017 deployment to the Cayman Islands.

4.1 Flight System

Our system is designed around a flight platform agnostic payload. This allows us to isolate development of the payload and the integration of the payload to a mobility platform. As mentioned in Section 1.2.1, there are multiple types of SUAS available, with different advantages to each. Uncoupling the design of the payload and the selection of a mobility platform permits us to tailor the mobility platform to suit the mission parameters. For this application, we are using the 3DR Solo¹.

Because we are operating in an austere environment, with limited open space, we cannot use fixed-wing aircraft. In addition, because our system post-processes the data to generate the

¹Discontinued in 2016

result, utilizing the long-range of a fixed-wing aircraft would result in a lot of wasted time if the system were unable to detect a transmitter. Because we are trying to design this system so that scientists and field researchers can easily use it, we cannot use monorotor aircraft, as they have too long of a learning curve.

In terms of available quadcopters, we can select from commercially available kits or custom built drones. Over the past few years, the two largest companies making drones were DJI and 3DR.

We elected to avoid custom quadcopters, such as the custom Tarot 600 used in previous deployments, as the control interface for these can be complex, with many interface features that could confuse users. In addition, many of the developmental features on custom drones are exposed on the controller, which is dangerous for field use, as less trained personnel might put the aircraft into a mode for which they are not trained. The 3DR Solo eliminates these by providing a controller which essentially eliminates any possibility of placing the aircraft into a non GPS controlled mode, and minimizes the number of switches and buttons relevant to the mission.

We elected to avoid using a DJI platform due to a lack of mature path-planning software, as well as a lack of integration facilities. All of the DJI consumer drones are designed to be monolithic systems that are ready to fly out of the box. As a result, they do not provide a simple way to program complex or nuanced flight patterns, nor do they provide a way to adjust their performance to accept a significant payload.

One of the better SUAS autopilots available is ArduPilot². ArduPilot is an open source autopilot in development since the early 2010s, and widely used in many developmental SUAS due to its stability, ease of integration, and capability. These qualities make it an excellent choice of autopilot to integrate with payload systems that require positional data.

Of the commercially available ready-to-fly solutions available, only the 3DR Solo used an ArduPilot autopilot and had the cleanliness to have a very short learning curve. Even though the

²<http://ardupilot.org/>

3DR Solo was just being discontinued, it was the best option for a drone with a mature and robust autopilot architecture with a very clean user interface.

In conjunction with the 3DR Solo, we use a MacBook Air with a Windows install to run Mission Planner³ as our primary mission control ground station and mission planning software. We selected Mission Planner as it is the most mature mission control and planning software available for the ArduPilot family of SUAS autopilots. The ground station connects to the 3DR Solo via the 3DR Solo's WiFi link, which is hosted on the 3DR Solo Controller. The WiFi link is used for flight platform control and telemetry.

4.2 Path Planning

As mentioned in Section 3.7, our system estimates the location of the transmitter by fitting the measurements of ping amplitude and receive location to the physics model relating signal strength to distance between transmitter and receiver. This is directly affected by the flight path the drone takes. Since our model relies primarily on the way the signal strength falls off from the transmitter, we need to measure the transmitter's signal strength at multiple ranges. However, since we do not know the transmitter's exact location prior to flying the mission, we need to maximize the efficiency of the flight path in terms of area covered.

Given the constraints above, our flight path must maximize its coverage of an area while still loitering above each point on the ground long enough to determine whether or not there is a ping at that location. In addition, the flight path must be time efficient, as the 3DR Solo only had a 12 minute flight time. As a result, we need to have some prior assumptions about where the transmitter might be, then efficiently cover that area with a flight pattern.

We develop our search area by first looking at where the transmitter might be. Our first choice is typically the transmitted animal's last known location - in this case, the subject species

³<http://ardupilot.org/planner/>




Figure 4.1: Lawnmower Flight Pattern

tends to stay in one place for a while, then just take off and start moving in a direction for a couple days, then stop moving. Typically, we found that if we set a search area centered around the animal’s last known location, and maximized the area covered, we would find the animal if it was not moving, or if it had only meandered around. In some cases, we would find that the animal had started moving significant distances - in these cases, we would find the bearing along which the animal was, drive to a convenient launch point, and set up a search area to search along that bearing as far as possible. This maximizes the coverage of the area that we think the animal is in.

The simplest and most efficient flight path for surveying a convex polygon is what is known as a lawnmower pattern. This pattern consists of evenly spaced lanes oriented parallel to the survey area’s longest axis, with the drone flying up and down each lane. An example of this can be seen in Figure 4.1.

We have selected our survey speed to be 5 ms^{-1} based on our understanding of the effective range of the payload, and the most efficient speed for the flight system. Based on an detection range of 30 m and a lane spacing of 30 m, if we fly at our minimum speed of 5 ms^{-1} , we can, in the worst case scenario of flying a line directly over the transmitter, with adjacent lanes just out of range of the transmitter, detect 8 pings. In the best case scenario of flying two lanes equidistant from the transmitter, we can expect as many as 13 pings.

While flying a closer lane spacing would generate more pings, it would also reduce the

efficiency of the survey flight, which could potentially cause us to not gather enough data to accurately identify the transmitter. However, flying a large lane spacing would introduce dead zones where we would not detect the transmitter, as it would be between two adjacent lanes.

In practice we found that the detection radius of the payload was sufficiently larger than 30 m, so it gave us room to increase the survey speed. We increased the survey speed up to 8 ms^{-1} , as that was the most efficient cruising speed for the copter. This allowed us to maximize our flight time, which in turn, increases our survey area.

Even though this method relies on knowing the detection radius of the system and specific transmitter, this method is tolerant of some uncertainty in detection radius. For a lawnmower pattern, if the lane widths are just slightly less than the detection radius, then we can guarantee that there will be a lane within the detection radius of the transmitter. Because the detection can only happen once every ping period of 1.5 s, then this reduces the guarantee. For a survey speed of 5 ms^{-1} , we can guarantee that if the transmitter is within 29.76 m of the lane, then its ping will be picked up.

In general, the survey speed is slow enough that we can guarantee more than four detections in the worst case scenario. We can, in principle, generate solutions for transmitters where we only detect them on one lane if the detection radius is actually smaller than we estimate. In theory, we can still generate an estimation for the actual location for the transmitter, but it will have some ambiguity as to which side of the survey line the transmitter is on. However, this is an acceptable tradeoff, as transmitters that have decayed to this state are likely also very difficult for the researchers to detect, and likely need to be replaced. This method is entirely tolerant of underestimation of detection radius, since a detection radius larger than our estimate guarantees that the transmitter will be heard across at least two lanes, and provides more data, thus improving the estimate.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author

is the primary investigator and author of this material.

Chapter 5

Experimental Results

5.1 Operational Details

In the summer of 2017, our team of researchers for UC San Diego worked with the San Diego Zoo’s Institute of Conservation Research and the Welch Lab at the Mississippi State University to deploy an autonomous drone-based radio transmitter tracker to assist with iguana studies on Little Cayman. We operated the tracking system alongside the scientists and collected performance and usability data on the system to determine whether or not the system was a viable tool for scientists to use in the field. The expedition involved 14 days of operations searching for 22 individual transmitters, with an average of 11 tracks per day, for a total of nearly 200 missions generating over 150 position estimates. Where possible, we validated tracks with standard tracking methodology, supplementing the data gathered by the scientists. These

Table 5.1: Deployment Summary to 2017

Year	Location	Duration	Hardware	Software
2015	Dominican Republic	10 days	Rev A	v1.0
2015	Dominican Republic	10 days	Rev B	v1.0
2016	Cayman Islands	10 days	Rev B	v2.0
2017	Cayman Islands	14 days	Rev C	v2.1

operational details are summarized in Table 5.1.

During the expedition, we often split into a drone team and a triangulation team, with each team taking half of the subject animals and switching midday. To determine our search area, we used the last known location of the animal, as well as its previous movement. We would then do a quick bearing check from the road to confirm that the subject animal was roughly in the expected direction, and at a reasonably correct distance. We would then program the search area and flight path, upload to the drone, and launch the mission. Since the quadcopter is capable of fully autonomous flight, all we would have to do is arm the aircraft and switch the drone into the autopilot mode. However, we sometimes manually launch and recover the drone, particularly where the takeoff area was somewhat crowded with foliage.

Once the drone landed, we would process the recorded data and generate the estimated location. The drone recorded the data directly to removable storage, so this was a simple matter of plugging the SD card into the laptop and copying the data to an external hard drive for redundancy purposes. We then used the custom processing software to analyze the data - this again was simplified so that the science team simply had to input the expected transmitter frequency and data location, and the software would generate the location estimate and probability maps automatically. We could then open the estimation and heatmap in GIS software to determine whether or not the mission had successfully located an iguana. The scientists would then take the GIS data and compile a time-series track of the iguana's movements over the course of several days. An example of the data visualization is shown in Figure 3.9.

The ease with which we could operate the drone was significant and reflected in the way the teams worked - the drone team often returned sooner than the manual tracking team, and with significantly more energy. The drone team was often able to stay in the supporting vehicle during flight operations, which made operations much more efficient. However, the time needed for post-processing meant that although the teams could conduct tracking in relative comfort, we spent a significant amount of time sitting around and waiting for things to happen. In addition,

due to the post-processing approach, we would not realize the drone had not detected a transmitter until 10 to 20 minutes after launching the drone, at which point, we would have to either re-fly the mission or fly a new search area. Yet another problem encountered was that, if the search area were large enough, much of that flight area would not yield information directly useful to estimating the location of the animal.

We found that our system, when operational, is on par in terms of operational pace with traditional radiotriangulation, which was the preferred method of tracking in the Cayman Islands. Our average track precision was also significantly better than the precision yielded by radiotriangulation. However, we also learned some important lessons that allowed us to further improve the system. During the 14 days of operations, we lost at least two days to airframe maintenance. These were needed to repair damage to exposed cables, connectors, and circuit boards. After two weeks of operations, we stopped using the drone because of approaching weather and irreparable damage caused to the LNA caused by continuous field use.

5.2 Estimation Precision and Accuracy

From a summary view of all missions and field results, we have a median certainty of 19 m, where we calculate certainty as the average distance of the 95 % confidence contour to the position estimate. This is significantly better than the 20-100 m precision of triangulation. The distribution of certainty is shown in Figure 5.1.

5.3 Estimation Speed - 2017

Over the 111 scientific flights flown, we flew a mean mission time of 5 min, with a total accrued mission time of 13.5 h. In general, postprocessing takes about the same time as the actual mission time; therefore the net time from launch to estimation of location is approximately

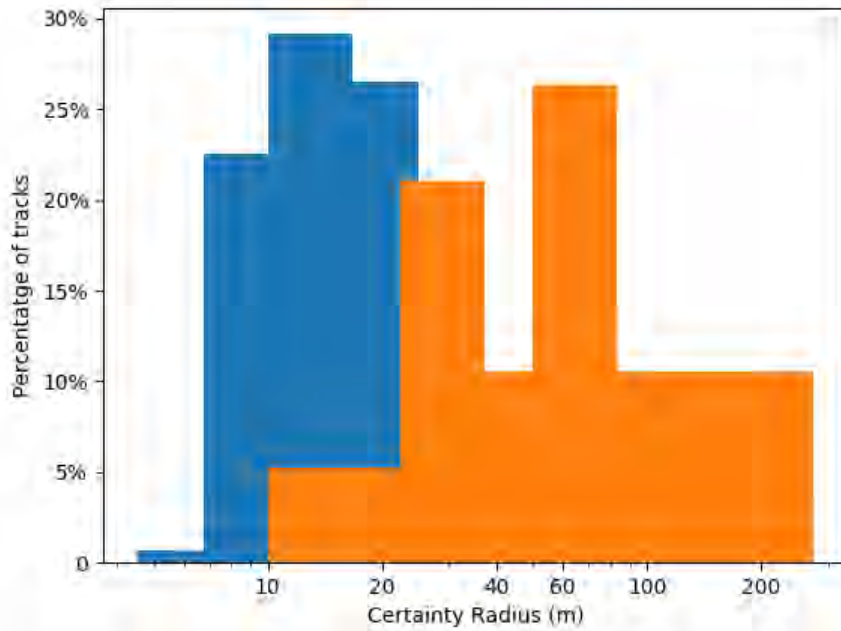


Figure 5.1: 2017 Certainty Distribution of Estimation Error

10 min. This is on par with the time needed to conduct a triangulation fix by hand. We show the distribution of mission times in Figure 5.3. Based on conservative estimated detection radius of 80 m, we estimate our missions covered a total of 14 km².

5.4 Effective Range

Based on previous field tests and experimentation, we expected a detection range of approximately 30 m. This would manifest as a concentration of pings around the estimated transmitter location with a radius of 30 m. During the course of the deployment, we found that the mean detection radius was approximately 150 m, with some transmitters being detectable from over 200 m away.

We can see an example of this in the Figure 5.4. In this mission, the drone was flying so that the payload antenna is parallel with the long legs of the flight pattern. The signal strength




Figure 5.2: 2017 Deployment Operational Pace

as recorded by the drone is colorized in the image, where the blue dots represent the “quietest” points, and the red dots represent the “loudest” points. We can see the green, which are the pings just above the noise floor, reaching out all the way to just beyond 150 m from the estimated position, marked by the red “X” and the center of the estimate heatmap.

Although we found our detection range to be, on average, 150 m, we did not change our path planning setup. On occasion, we found that our detection range could drop to as low as 70 m. In addition, we found that we were able to cover enough area as is without changing the lane spacing, and having the additional measurements simply increased the confidence of the estimation.

Chapter 5, in part, is currently being prepared for submission for publication of the material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

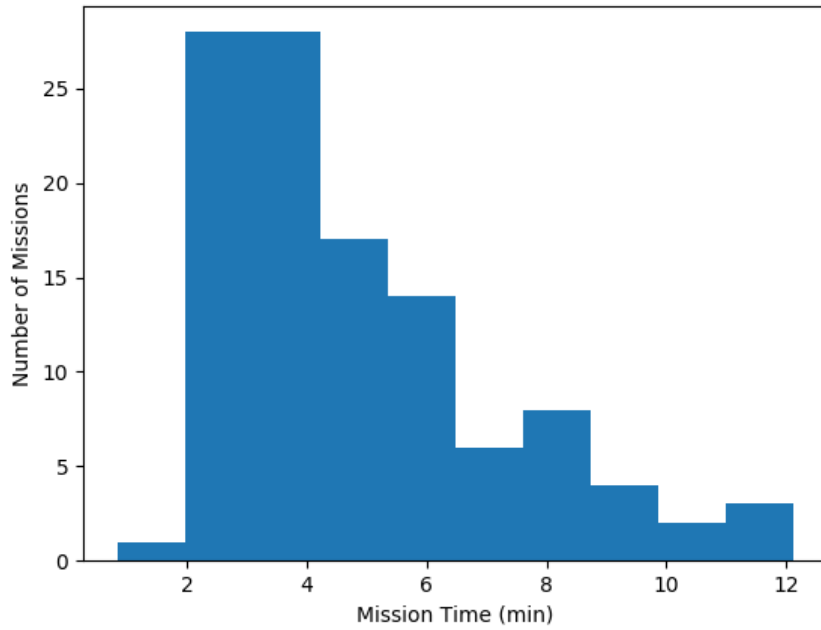


Figure 5.3: 2017 Deployment Mission Time Distribution

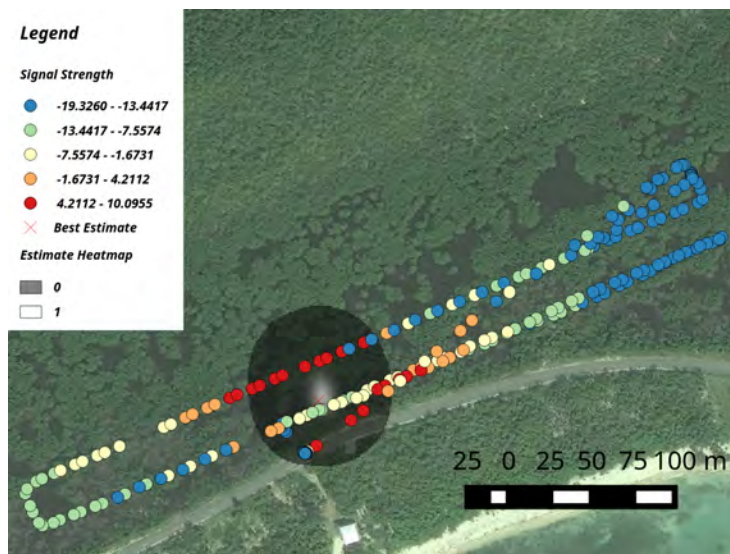


Figure 5.4: Mission 25 Results, 19 August 2017

Chapter 6

Continuing Improvements

Although the 2017 system has demonstrated that it is capable of tracking wildlife at a pace on par with scientists using traditional methodologies of radio-triangulation, there is room for significant improvement that will help make this system a much more valuable tool for field biologists. One of the primary ways this could be made better is by moving the signal processing and localization onto the drone for real-time processing. This would enable us to minimize the mission time by eliminating the post-processing wait, as well as minimize the flight time by eliminating the areas of the flight path after we have accrued enough measurements for a sufficient precision and avoided areas where we know we will not get any further information.

To address these changes, we are moving from a post-processing architecture to a real-time processing architecture in order to speed up the operational pace in the fourth version (hardware revision D, software v3.0) of the system. This change allows us to get feedback about the drone's estimation status as the drone is collecting data, so we can make the flight path more efficient while eliminating the post-processing wait. In order to achieve this, we are moving from the Intel Joule to the UP Core single board computer, and are rewriting the processing pipeline to be more efficient and robust. These systems are described further in Sections 6.3 and 6.2, and are summarized in Tables 6.1 and 6.2.

This fourth version, shown in Figure 6.1 is a work in progress and is anticipated to deploy to the Turks and Caicos Islands to support Dr. Glenn Gerber (San Diego Zoo Institute for Conservation Research) in studying boas sometime in August 2019, as summarized in Table 6.3.



Figure 6.1: 2019 Radio Tracking Drone

We also want to make the entire system physically more robust, as our operational pace suffered towards the end of the deployment as various components started to break. This is particularly important for drones being used for field research, as often, field researchers don't have the time or skill to repair broken components in the field.

In the following sections, we will discuss possible avenues and current work for upgrades to the 2017 system.

Table 6.1: Hardware Version Comparison to Present

Version	Year	Airframe	Computer	Receiver	Antenna
Rev A	2015	Tarot 600	BeagleBone Black	RTL-SDR	Dipole v1
Rev B	2016	Tarot 600	Raspberry Pi 3	RTL-SDR	Dipole v1
Rev C	2017	3DR Solo	Intel Joule	USRP B200mini	Dipole v2
Rev D	2019	3DR Solo	UP Core	USRP B200mini	Double Bazooka

Table 6.2: Software Version Comparison to Present

Software Version	Signal Detector	Estimator	Visualization	Processing
v1.0	FFT & Loudest Signal	Average	Heatmap	Post-process
v2.0	FFT & Loudest Signal	Manual	GIS / Heatmap	Post-process
v2.1	FFT & Loudest Signal	Model-fit	GIS	Post-process
v3.0	Match Filter	Model-fit	GIS	Real-time

Table 6.3: Deployment Summary to Present

Year	Location	Duration	Hardware	Software
2015	Dominican Republic	10 days	Rev A	v1.0
2015	Dominican Republic	10 days	Rev B	v1.0
2016	Cayman Islands	10 days	Rev B	v2.0
2017	Cayman Islands	14 days	Rev C	v2.1
2019	Turks and Caicos Islands	TBD	Rev D	v3.0

6.1 Double Bazooka Antenna

As we revised the system in preparation for the 2019 deployment, we tested a couple other antennas to evaluate their sensitivity and physical robustness. One of the antennas we evaluated was a double bazooka design, built out of RG174 cable for the 200 MHz band¹. The double bazooka antenna design is a derivation of a half-wave dipole, and theoretically has similar directional characteristics. This particular design is attractive because it is built of flexible materials, which, in addition to its omni-directional characteristics, makes it particularly suited for use in this system.

Since we are considering a more robust estimation pipeline, we will need to evaluate and characterize the directionality of this antenna. This information will be used to inform the estimation, which we will describe in Section 6.4, about the gain associated with the direction of arrival of the signal.

¹Acknowledgements to Brian Baxter from the UC San Diego Master of Advanced Study - Wireless Embedded Systems program for building this antenna

6.2 On-board Computer

For the 2019 deployment, our objective was to design a real-time tracking system. In order to achieve that, we need to be able to efficiently process data. In addition, Intel discontinued the Intel Joule single board computer in 2017, which resulted in a board that had little support and market.

We elect to use the UP Core ² as the onboard computer because it has an Intel Atom processor and a small form factor. We selected an Intel Atom processor in order to take advantage of the Single Instruction Multiple Data (SIMD) instruction set present on the x86 Atom architecture.

Due to the format of the peripherals on the UP Core, we needed to add a companion board to be able to provide a hardware user interface on the drone and receive pose information from the payload's GPS and compass module. We decided to build a custom board with an Atmel ATmega32u4, as that microprocessor has a native USB bus. This enables seamless integration with the UP Core, which has limited hardware serial lines.

This custom board aggregates information from the GPS and compass, and forwards it on to the UP Core over the USB serial link. It also provides a LED status indicator so that users can determine system status without needing to connect the ground station software, and breakouts to support external data storage for the UP Core.

6.3 Signal Processing

As mentioned earlier, one of our primary objectives is to implement real-time signal processing. We started by reexamining the approach we were taking, which used an FFT to identify the signal at the frequency we were targeting. One of the key impediments to using that approach in a real-time signal processing pipeline was that FFTs are expensive to compute, and our previous approach described in Section 3.6 computed the FFT for all samples recorded.

²<https://up-shop.org/up-core/271-up-core.html>

Our initial approach was to do away with the FFT and use a modular match filter. In theory, the match filter provides the best way to detect a particular signal. However, we found that while the match filter was able to run fast enough, it was not able to detect the transmitter signals accurately enough. This is probably due to the variability in the transmitter frequency, and because the received transmitter signal is not perfectly sinusoidal. We also found that the speed of this approach tended to break down as we tried to detect more transmitters simultaneously, as the complexity of this approach is $O(n)$ with respect to the number of transmitters.

We then attempted a different approach, where we tried to look at the energy being received. Each transmitted ping will cause some amount of energy to be received. This energy is proportional to the received power, so calculating the amount of energy received in the past pulse width would give a measure of how much energy could be attributed to the pulse. While this approach will fundamentally have a lower SNR, this would provide a foundation from which we could add filtering to increase SNR.




Figure 6.2: Dynamic Threshold

To detect the signal, we first calculate the squared magnitude (power) of the signal. We

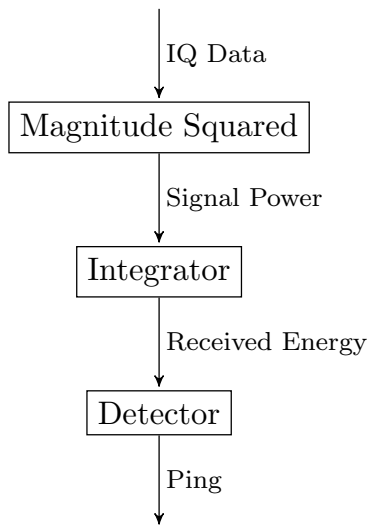


Figure 6.3: Software v3.0 Signal Detector

then integrate a 3.2 ms window to amplify the received signal, thus calculating the amount of energy received during that window. We then calculate a dynamic threshold based off the past 1.25 s as the median of the local peaks. This gives a very good approximation of the noise floor, as shown in Figure 6.2. We then compare the width of signals above the threshold to the expected width of the pings, and select those more than 75 % and less than 200 % of the expected ping width. This is shown graphically in Figure 6.3.

Because we are recording a wide spectrum of the radio frequency, and because the energy received in a signal is independent of the signal frequency, we are able to receive and detect multiple transmitter frequencies simultaneously with no additional computation. When we identify each ping in time, we also run an FFT on the underlying raw data associated with only the ping to determine the transmitter frequency. This lets us identify each ping’s frequency, and feed each transmitter’s data into its own model.

This approach is more efficient than methods such as running a bank of tuned filters or a 4096 element FFT over all the data because it does only time domain operations on the entirety of the raw data up until the frequency identification step, at which point only the data samples of

interest are put through the FFT. This reduces the amount of computation required, and allows us to pick up any transmitter without necessarily needing to know the target frequencies beforehand.

Although this approach was able to detect some pings, it suffered from a very low SNR because there was no filtering. One of our problems was that the software we have written is currently unable to leverage all of the compute power available on the UP Core. Once that problem has been resolved, we can use a FFT to isolate various portions of the receive band to improve the SNR of the signal that is being fed into the energy detector. This FFT would not need to be the 4096 elements used in the 2017 system, as described in Section 3.6, but could be a much more computationally smaller 128 or 256 element FFT that would help reject noise.

6.4 Location Estimation

In v3.0, we provide a more mathematically robust estimation of precision. We rewrite the original signal propagation model from Equation 3.1 in terms of the transmitter power P and additional losses k , as shown in Equation 6.1. This model explicitly models the transmit power of the transmitter with the parameters R being the received signal power, n is the path loss exponent, d is the distance between the drone and the transmitter, and k is inherent system losses. We assume the transmitter power P to have a Gaussian distribution with mean μ_P and variance σ_P . Additionally, once we have the directionality of the antenna from Section 6.1, we can use the heading of the drone to adjust the gain of the antenna in the direction of the estimate by replacing the measurement R with a mapping function $R(\theta)$, where θ is the bearing from the drone's orientation to the transmitter.

$$R = P - 10n \log_{10}(d) + k \tag{6.1}$$

$$P = 10n \log_{10}(d) + k - R \tag{6.2}$$

$$d = 10^{\frac{P-k+R}{10n}} \quad (6.3)$$

$$f_D(d) = \mathcal{G}(10n \log_{10}(d) + k - R, \mu_P, \sigma_P) \frac{10n}{\ln(10)d} \quad (6.4)$$

Once we have this, we can calculate the probability at each geographic point of its distance to the recorded pings. This is found by the form in Equation 6.4, where $\mathcal{G}(x, \mu, \sigma)$ is the probability of the value x drawn from a Gaussian distribution with mean μ and variance σ . Since we treat each ping as independent, the probability of the transmitter being at a particular location is simply the product of the probability of the transmitter at that location for each ping.

Chapter 7

Conclusion

Our 2017 system has a demonstrated capability of tracking multiple targets simultaneously. Our system can localize a transmitter to within 15 m in an average time of 10 min of launching. We demonstrate a mean effective detection range of 150 m. Our 2019 approach has the potential to drastically decrease the time to estimation of this system, but has some technical issues to be resolved before it can be meaningfully compared to existing systems.

7.1 Comparison to Existing Platforms

Table 7.1: Comparison of Performance

System	Platform	Avg. Precision	Avg. Detection Time	No. of Trials
Nguyen, et al	3DR Iris	22.7 m	135 s	16
Dressel, et al	DJI M-100	5 m	37 s	3
Cliff, et al	AT Falcon 8	55 m	600 s	8
2017 System	3DR Solo	19 m	300 s	152

We compare the 2017 system’s performance with the performance reported by Nguyen et al, Dressel et al, and Cliff et al in Table 7.1. Our system demonstrates performance that exceeds each of the other systems, and is also smaller and more user friendly.

In 2018, Nguyen et al fielded a range based particle filter system on a 3DR Iris¹ and conducted a series of 16 flights to validate tracking performance. These flights were conducted in a search area of 75 m by 300 m on two human animal analogues. Their system reported a best root mean square estimate error of 22.7 ± 13.9 m in 135 ± 53 s, or about 2.3 min [20]. Our flight platforms are very similar in size and weight. With our approach, we achieve a similar estimation precision, even though our reported mean mission times are longer. This is due to our mission times including the time to transit to the survey area, and our system not taking advantage of its 150 m detection range.

In 2018, Dressel et al fielded a bearing based histogram filter system on a DJI M-100² and conducted a series of 3 flights to validate tracking performance. These flights were conducted in a search area of 400 m by 400 m on a presumably static handheld radio as a radio collar analogue. Their system reported a fix when a 5 m by 5 m cell contained 50 % confidence of estimation, which occurred in 37 s [16]. It is difficult to compare the estimated error/precision, as we do not have a measure of the distribution of their estimation error. Although Dressel claims a faster time to estimate, we again point out that our reported mean mission times take into account time to transit to the survey area, and our system not being the absolute optimal in terms of leveraging its detection range. Our system is also significantly smaller, making it much more field friendly.

Lastly, in 2018, Cliff et al fielded a range and bearing based grid filter system on an Ascending Technologies Falcon 8³, and conducted a series of 8 flights to validate tracking performance on tagged swift parrots. Their system reported a best estimate precision of 55 m in 10 min [19]. Our approach is able to generate a far more precise estimate in a similar overall time, with a much smaller airframe. Again, if we modify our flight path to leverage the 150 m detection radius, our system will be significantly faster and more field friendly.

Chapter 7, in part, is currently being prepared for submission for publication of the

¹Discontinued in 2015

²<https://www.dji.com/matrice100>

³<http://www.asctec.de/en/uav-uas-drones-rpas-roav/asctec-falcon-8/>

material. Hui, Nathan; Lo, Eric; Gerber, Glenn; Schurgers, Curt; Kastner, Ryan. The thesis author is the primary investigator and author of this material.

Bibliography

- [1] A. M. Allen and N. J. Singh, “Linking movement ecology with wildlife management and conservation,” *Frontiers in Ecology and Evolution*, vol. 3, p. 155, 2016. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fevo.2015.00155>
- [2] B. G. Lascelles, P. R. Taylor, M. G. R. Miller, M. P. Dias, S. Opper, L. Torres, A. Hedd, M. Le Corre, R. A. Phillips, S. A. Shaffer, H. Weimerskirch, and C. Small, “Applying global criteria to tracking data to define important areas for marine conservation,” *Diversity and Distributions*, vol. 22, no. 4, pp. 422–431, 2016. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ddi.12411>
- [3] R. Kenward, *A manual for wildlife radio tagging*. Academic Press, 2007.
- [4] J. J. Millspaugh and J. M. Marzluff, *Radio tracking and animal populations*. Academic Press, 2009.
- [5] S. Harris, W. J. Cresswell, P. G. Forde, W. J. Trehwella, T. Woollard, and S. Wray, “Home-range analysis using radio-tracking data: a review of problems and techniques particularly as applied to the study of mammals,” *Mammal Review*, vol. 20, no. 23, pp. 97–123, 1990. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2907.1990.tb00106.x>
- [6] M. Wikelski, R. W. Kays, N. J. Kasdin, K. Thorup, J. A. Smith, and G. W. Swenson, “Going wild: what a global small-animal tracking system could do for experimental biologists,” *Journal of Experimental Biology*, vol. 210, no. 2, p. 181186, 2007.
- [7] R. Kays, S. Tilak, M. Crofoot, T. Fountain, D. Obando, A. Ortega, F. Kuemmeth, J. Mandel, G. Swenson, T. Lambert, B. Hirsch, and M. Wikelski, “Tracking animal location and activity with an automated radio telemetry system in a tropical rainforest,” *The Computer Journal*, vol. 54, no. 12, p. 19311948, 2011.
- [8] L. D. Mech, *Handbook of animal radio-tracking*. Univ. of Minnesota Press, 1986.
- [9] G. C. White and R. A. Garrott, *Analysis of Wildlife Radio-Tracking Data*. Academic Press, 2012.

- [10] A. Jensen and Y. Chen, "Tracking tagged fish with swarming unmanned aerial vehicles using fractional order potential fields and kalman filtering," *2013 International Conference on Unmanned Aircraft Systems (ICUAS)*, May 2013.
- [11] P. Soriano, F. Caballero, and A. Ollero, "Rf-based particle filter localization for wildlife tracking by using an uav," *INTERNATIONAL SYMPOSIUM ON ROBOTICS*, p. 239244, 2009. [Online]. Available: <https://www.tib.eu/en/search/id/BLCP:CN072750547/RF-based-Particle-Filter-localization-for-Wildlife/>
- [12] F. Korner, R. Speck, A. H. Goktogan, and S. Sukkarieh, "Autonomous airborne wildlife tracking using radio signal strength," *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2010.
- [13] G. A. M. D. Santos, Z. Barnes, E. Lo, B. Ritoper, L. Nishizaki, X. Tejada, A. Ke, H. Lin, C. Schurgers, A. Lin, and R. Kastner, "Small unmanned aerial vehicle system for wildlife radio collar tracking," *2014 IEEE 11th International Conference on Mobile Ad Hoc and Sensor Systems*, 2014.
- [14] K. Vonehr, S. Hilaski, B. E. Dunne, and J. Ward, "Software defined radio for direction-finding in uav wildlife tracking," *2016 IEEE International Conference on Electro Information Technology (EIT)*, 2016.
- [15] H. Bayram, N. Stefas, and V. Isler, "Aerial radio-based telemetry for tracking wildlife," *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [16] L. K. Dressel and M. J. Kochenderfer, "Efficient and low-cost localization of radio signals with a multicopter uav," *2018 AIAA Guidance, Navigation, and Control Conference*, 2018.
- [17] O. Cliff, R. Fitch, S. Sukkarieh, D. Saunders, and R. Heinsohn, "Online localization of radio-tagged wildlife with an autonomous aerial robot system," *Robotics: Science and Systems XI*, 2015.
- [18] S. S. Andr Posch, "Uav based search for a radio tagged animal using particle filters," 2009.
- [19] O. M. Cliff, D. L. Saunders, and R. Fitch, "Robotic ecology: Tracking small dynamic animals with an autonomous aerial vehicle," *Science Robotics*, vol. 3, no. 23, Oct 2018.
- [20] H. V. Nguyen, M. Chesser, L. P. Koh, S. H. Rezatofighi, and D. C. Ranasinghe, "Trackerbots: Autonomous unmanned aerial vehicle for real-time localization and tracking of multiple radio-tagged animals," *Journal of Field Robotics*, Jan 2019.
- [21] D. Webber, N. Hui, R. Kastner, and C. Schurgers, "Radio receiver design for unmanned aerial wildlife tracking," *2017 International Conference on Computing, Networking and Communications (ICNC)*, 2017.