Antenna Characterization for UAV Based GPS Jammer Localization

Adrien Perkins, Louis Dressel, Sherman Lo, & Per Enge
Stanford University

BIOGRAPHY

Adrien Perkins is a Ph.D. candidate in the GPS Research Laboratory at Stanford University working under the guidance of Professor Per Enge in the Department of Aeronautics and Astronautics. He received his B.S. in Mechanical Aerospace Engineering from Rutgers University in 2013 and his Master of Science in Aeronautics and Astronautics from Stanford University in 2015.

Louis Dressel is a Ph.D. candidate in the Aeronautics and Astronautics Department at Stanford University. There, he works on a joint project with both the Stanford Intelligent Systems Lab and the GPS Research Laboratory implementing control algorithms for a jammer-hunting UAV. He received his B.S. in Aerospace Engineering with a minor in Computer Science from Georgia Tech in 2013 and his Master of Science in Aeronautics and Astronautics from Stanford University in 2015.

Sherman Lo is a senior research engineer at the Stanford University GPS Laboratory. He is the Associate Investigator for the Stanford University efforts on the FAA evaluation of alternative positioning navigation and timing (APNT) systems for aviation.

Per Enge is a Professor of Aeronautics and Astronautics at Stanford University, where he is the Vance and Arlene Coffman Professor in the School of Engineering. Here, he directs the GPS Research Laboratory which develops navigation systems based on the Global Positioning System (GPS). He has been involved in the development of WAAS and LAAS for the Federal Aviation Administration (FAA). He has received the Kepler, Thurlow, and Burka Awards from the ION. He also received the Summerfield Award from the American Institute of Aeronautics and Astronautics (AIAA) as well as the Michael Richey Medal from the Royal Institute of Navigation. He is a fellow of the Institute of Electrical and Electronics Engineers (IEEE), a fellow of the ION, a member of the National Academy of Engineering, and has been inducted into the Air Force GPS Hall of Fame. He received his Ph.D. from the University of Illinois in 1983.

ABSTRACT

The increased integration of GPS in the airspace to improve efficiency and enhance safety comes with an increased risk of harm from a GPS jamming device. In an effort to protect the airspace of the future from the effects of GPS jammers by rapidly mitigating their effects, we are developing an unmanned aerial vehicle (UAV) based solution for the localization of a jammer that can operate at an airport. The development of this UAV solution is centered on two main areas: the localization of the jammer and navigation in a denied environment. In this paper, we examine developments in the first area, the localization of the jammer.

This paper examines the two primary components of our localization system: radio frequency (RF) detection and a navigation control scheme for localization. The RF detection component uses a directional antenna and the UAV's ability to rotate to determine a bearing to the jammer. The navigation control scheme selects a trajectory for making bearing measurements that allow rapid localization of the jammer.

The ability to make reliable bearing measurements is paramount to quickly localizing the jammer. The research in this paper analyzes our directional antenna and three bearing calculation methods: max, cross-correlation, and a modification of max leveraging the shape of the antenna's main lobe (referred to as max3 in this paper). Flight test data is used to analyze the performance of these methods at various distances from the jammer.

Localization is achieved by determining and moving to the best locations to make signal observations and merging these observations to get an accurate estimate of jammer location. In order to pick measurement locations that minimize localization time, we've formulated the problem of picking these locations as a partially observable Markov decision process (POMDP). This closed-loop controller determines the next measurement location given the information gathered from prior measurements. Flight tests with this control scheme were performed and are presented in this paper. The effect of bearing measurement quality on the decision-making is also shown.

INTRODUCTION

Whether malicious or unintentional, GPS jamming events have already proven to disrupt airports [1] and pose an increased risk to commercial aviation in the future. An important mitigation for this risk is the ability to rapidly locate and interdict the GPS jamming device.

There are several challenges faced by a localization system, especially in the semi urban environments typically found in and around an airport. The system must be capable of reliably determining jamming direction and quickly localize the source. This paper examines both aspects. We demonstrate the capability in the field using a UAV as a mobile sensor platform to make measurements of the jamming signal at different positions dynamically chosen to optimally locate the source of the jammer.

In developing a localization algorithm, the measurements being made by the system can greatly impact the performance of that algorithm. In our jamming localization, the primary sensor for measurements is a directional antenna. Using this antenna, our multirotor platform Jammer Acquisition with GPS Exploration & Reconnaissance (JAGER) can measure the bearing to the jammer which is the main input into the localization algorithm [2]. Leveraging the behavior of a multirotor platform, JAGER is able to rotate a directional antenna at a fixed location to create a gain pattern for that antenna and determine bearing from that location, in the same nature as work done in [3]. With every signal strength measurement made, the system is able to match it with a heading from a magnetometer. In this paper we examine three different bearing calculation techniques from a gain pattern: max, cross-correlation, and max3.

The closed-loop navigation controller uses the gathered information to determine where to go next in order to most quickly localize the jammer. In this paper, the localization objective is modeled as a POMDP to determine the optimal route. The viability of this technique for locating the jammer source will be demonstrated through flight testing in a simulated environment.

Existing Solutions

Because of aviation's increasing reliance on GPS, the Federal Aviation Administration (FAA) has pursued technologies for radio frequency interference (RFI) detection and localization [4]. These technologies included ground-based systems and airborne systems like the Aircraft RFI Localization and Avoidance System (ARLAS) [5]. 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 be "illuminate" the ground. If RFI were detected during the roll, it would imply the source was somewhere in the illuminated region. By performing a succession of rolling maneuvers, the interference source could be localized.

However, a solution like ARLAS is expensive because it requires a human pilot. ARLAS also suffers from 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.

By using an agile, multirotor UAV, we hope to overcome these limitations. A multirotor UAV can rotate a gain pattern without having to alter its flight path. This solution also can be easier to deploy, have lower cost and provide faster response times than a manned system.

LOCALIZATION ALGORITHM

The need for antenna characterization is motivated by our localization algorithm, which requires a model of the sensor's performance. This section describes the localization algorithm used by the vehicle. The problem and some assumptions are presented, followed by a detailed formulation.

Problem Overview

Formally, the goal is to minimize the time to localize a jammer within some small region. The vehicle picks a trajectory leading to measurements that minimize this localization time. To improve performance, the algorithm is closed-loop; the vehicle updates its trajectory online with each measurement.

Because our vehicle is an agile, multirotor UAV, it can translate, climb, rotate, and make received signal strength indicator (RSSI) measurements at the same time. It is computationally difficult to reason over such a large input space. It is also difficult to coherently combine measurements taken at different vehicle positions, because the jammer antenna might not be omnidirectional. Furthermore, measurements made while the vehicle aggressively maneuvers can suffer from vehicle instability or vibrations.

To simplify the problem, we constrain the vehicle to a constant altitude and assume a single, stationary jammer. We also decouple the problem into two actions: the first is making an observation (determining bearing to the jammer by rotating the vehicle in place), and the second is taking action (determining and moving to the next location for the next observation).

This paper focuses on the first action: how accurately can bearing be estimated if the vehicle samples RSSI values

while rotating in place and how can those measurements affect the decisions of where to rotate next?

POMDP Formulation

The problem of choosing successive rotation locations has been formulated as a partially observable Markov decision process (POMDP). POMDPs are a principled approach to decision making and closed-loop control in stochastic domains [6]. Fully solving them is generally intractable [7], but there exist algorithms that provide good, approximate solutions. Our POMDP formulation has been presented in prior work [8], but a brief review is presented here.

At each time step, the problem can be described by a state $s \in S$, where S is the state space, or set of all possible states. To limit the size of the state space, the search area is split into a grid. The vehicle's position is constrained to grid cell centers. A state consists of four state variables: the vehicle x-index x_v, the vehicle y-index y_v, the jammer's x-index x_j, and the jammer's y-index y_j. The state variables x_j and y_j indicate the grid cell in which the jammer resides. At each time step, the state is only partially observable—the jammer's position is unknown.

At each time step, the vehicle can take some action $a \in A$, the set of available actions. In our formulation, the vehicle can travel to any of the neighboring grid cells, rotate in place to make a bearing measurement, or simply hover, resulting in 10 possible actions. After taking action a from state s, the problem will transition to some state s'. This transition is specified by the transition function $T(s, a, s') = P(s' \mid s, a)$, which describes the probability of transitioning to some state s' from a given s and a. We've assumed the vehicle successfully moves to the desired grid cell, without any probability of deviating to other cells. The jammer is also assumed to be stationary, so the transitions are called deterministic.

At any time step, the state is unknown to the vehicle. Instead, it makes an observation $o \in O$, where O is the set of all possible observations. In our problem, these observations are the bearing measurements made when the vehicle rotates. To reduce the number of possible observations and computational complexity, the angular space is split into 36 10-degree bins, and measured bearings are rounded to the nearest bin. An additional "null" observation is included for a total 37 possible observations. The null observation is received when the vehicle does not rotate.

The POMDP formulation includes an observation model $Z(a, s', o) = P(o \mid a, s')$ describing the probability of making observation o is after taking action a and transitioning to state s'. This probability is a function of the bearing measurement quality. Prior to the work presented in this paper, it was assumed bearing measurements had zero-mean Gaussian noise with a 10-degree standard deviation. It was also assumed that if the vehicle rotated in the same grid cell as the jammer, it would receive the null measurement, because the space directly under the vehicle is outside the main lobe of its directional antenna. An updated observation model, using the characterization performed in this work, can be found in the section entitled *(Effect on Localization Algorithm)*.

Although the vehicle is unaware of the true state, it maintains a probability distribution over the state space, called a belief, denoted b. After taking an action and making a new observation, Bayes' law is used to update the belief. This updated belief is used in conjunction with policy π to determine the next action to take. A policy $\pi(b)$ maps beliefs to actions.

The belief $b \in R^{|S|}$ has a dimensionality equal to the number of states. The set of all beliefs, called belief space B, is large. Reasoning over this large space contributes to the general intractability of solving POMDPs exactly. However, efforts in recent years have developed approximation algorithms to compute near-optimal policies [9, 10]. This research uses SARSOP [11], which allows a policy to be computed offline and uploaded to the vehicle before a mission. The vehicle then relies on this policy to make decisions while in flight.

Generating a policy requires a reward model $R(s, a)$ that encourages the vehicle to perform certain actions. The function $R(s, a)$ returns a reward (positive value) or penalty (negative value) for taking action a from state s. In our formulation, we reward the vehicle when it stops in the grid cell containing the jammer. This encourages the vehicle to first find the jammer, which is our goal. We give penalties for movement and rotations to reflect the time taken to perform them. These penalties depend on grid size and rotation rate. If the vehicle travels at 10 m/s and grid cells are 10 meters wide, it will take roughly one second to travel to adjacent grid cells along the cardinal directions and $\sqrt{2}$ seconds to travel to diagonal grid cells (time spent accelerating is currently ignored, for simplicity). If the vehicle rotates at 15 degrees / second, it completes a rotation in about 20 seconds. In this case we might assign -1 for moving in cardinal directions, $-\sqrt{2}$ for diagonal movement, and -20 for rotations.

EXPERIMENTAL SETUP

Our UAV based GPS jammer locator, JAGER, is built on a commercially available multirotor platform, the DJI S1000 that has been modified to be a test platform for various system, including the one used in this paper [2]. For control and navigation, the vehicle is equipped with a Pixhawk autopilot system running a custom version of the PX4 firmware [12]. The Pixhawk also has all of the sensors onboard to determine the attitude and position of

the vehicle at all times. The localization decisions are made on the flight computer which is an Odroid-U3 ARM based computer that communicates with the autopilot throughout the flight. Finally all the signal strength measurements are made with a directional yagi antenna connected to the RN-XV WiFly module. A schematic of this configuration and the flow of information can be seen in Figure 1.

Throughout the flight, the Odroid and the Pixhawk are in constant communication using the Mavlink protocol. This allows the Pixhawk to send real time attitude and position information to the data collection and localization algorithm onboard the Odroid. This same protocol was adapted to allow the localization algorithm to send commands to the Pixhawk. Furthermore, the Mavlink protocol is also used to send real time telemetry information to the ground which could also be modified to allow for real time monitoring of both the vehicle state as well as signal strength measurements and localization decisions.

With this setup, signal strength measurements are made at 2Hz throughout all phases of the flight. When a rotation is commanded, the vehicle rotates with a targeted rotation rate of 15 degrees / second. Given the small size of this payload, the flight time achieved during testing was 20 minutes on two 6 cell 8000mAh batteries.

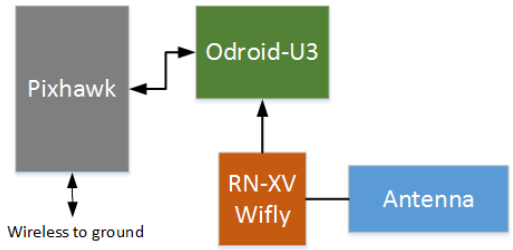

Figure 1: Schematic of components on UAV

Signal Source

Due to restrictions on being able to actively interfere with GPS signals, a 2.4GHz WiFi router was used as a proxy jammer for all our flight testing. The WiFi router was placed on the ground at a surveyed location. In these tests, GPS was used for navigation as we are still developing alternate and GPS jamming resistant navigation.

Antenna

A single the L-com HG2409Y yagi antenna was used for this experiment. This 2.4GHz WiFi antenna has a 60-degree beam width both horizontally and directionally as shown in Figure 2. As depicted in Figure 3, the antenna was mounted below the vehicle in order to have the clearest view to a ground based signal. Furthermore, the antenna is placed angled down at 30-degrees in order to have the main lobe of the antenna extend out to the horizon. This also leaves a cone underneath the vehicle with a weak signal that was aimed to be leveraged as an additional measurement of being able to get a null measurement when over the jammer.

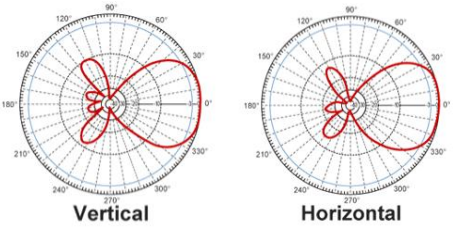

Figure 2: Directional antenna gain pattern from L-com datasheet [13]

Figure 3: Directional antenna mounted on underside of UAV

Measurements

Throughout the UAV's flight, the directional antenna is continually making RSSI measurements. In order to calculate bearing from a given location commanded by the localization algorithm, the vehicle simply rotates at that position and combines all of the RSSI measurements to form the antenna's gain pattern. This gain pattern can then be used to estimate the bearing of the signal source from that given position. In this paper that bearing calculation is done with three different methods: max, cross-correlation, and max3.

The directional antenna is used to provide measurements to the WiFly module which then outputs RSSI. The RSSI measurements from one full rotation is then used to determine the bearing to the jamming signal. Three different methods were used for the bearing determination.

The max method simply finds the maximum RSSI value in the measured pattern and uses that heading as the bearing to the jammer. This method has been used on a rotating UAV in prior research [3].

The cross-correlation method compares the measured pattern with the known "truth" pattern for the antenna. Cross-correlation has been used to determine bearing on a rotating robot in previous research [14]. 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. The cross-correlation coefficient c is the sum of all such products for a single shift γ. Equation (1) shows the process of cross-correlation.

$$c(\gamma) = (R_n * R_t)[\gamma] = \sum_{\alpha=0}^{360^\circ} R_n(\alpha) R_t(\alpha + \gamma) \quad (1)$$

The cross-correlation c is computed for every possible shift γ. The shift yielding the highest cross-correlation coefficient is taken to be the bearing to the jammer.

The measured pattern must be normalized because patterns of RSSI are similar across various distances, but are scaled differently due to effects of range and transmitter power. This normalization is carried out by subtracting the mean RSSI value from the pattern and dividing by the standard deviation.

To get our "truth" pattern R_t, we sampled RSSI at every 10-degrees at distances ranging from 10 to 40m, normalized the resulting patterns, and took the mean of these normalized patterns.

Finally the max3 method is an improvement on the max method where the bearing is the mean of the bearing of the two crossings of 3dB below the maximum RSSI value for the pattern, as depicted in Figure 4.

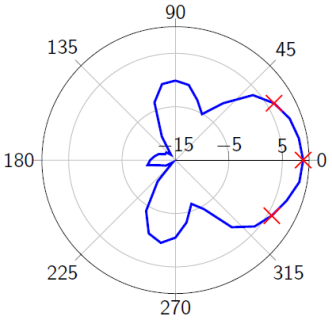

Figure 4: Depiction of Max3 method.

Flight Area

Test flights were performed at the Joint Interagency Field Experimentation (JIFX) event hosted by the Naval Postgraduate School. For most of the measurements the vehicle was flown at an altitude of 100ft AGL with a handful of measurements made near the signal source at an altitude of 50ft AGL. When the localization algorithm was tested, a 9 by 9 grid (each cell 11m on a side) was used as the world with the signal source located in the top right cell and the vehicle starting in the center cell, 62m from the signal source.

RESULTS

During flight tests with the JAGER vehicle, 88 different experimental gain patterns were created and bearing calculations were made with each of the three previously described methods. The configuration geometry and performance of the WiFly module resulted in splitting the gain patterns into 3 different classes of distance from the signal source: near, far and ideal. The performance of each of the bearing methods is described within this classification.

The POMDP based localization algorithm was successfully executed to locate the signal source. Leveraging the performance results of the cross-correlation and max3 methods, the model for the POMDP was updated and produced a significantly different flight profile. In addition to the POMDP based localization algorithm, a baseline algorithm was also used in order to demonstrate the advantages of the POMDP based algorithm. This section presents the results obtained from these flight tests.

Effects of Distance

Throughout the experiment, measurements were made at distances from the signal source ranging from directly overhead to almost 350m away. Figure 5 shows a view of all the locations in which measurements were taken during flight tests, with the signal source being in the center of the main grouping. Each of the markers represents one measurement and the color of the marker represents roughly the maximum RSSI value measured at each location. Green points has an RSSI value greater than -65dBm and the color turns through yellow and orange to dark red where the RSSI value was less than -80dBm. As expected, as the vehicle traveled further from the signal source, the maximum RSSI value measured dropped as can be clearly seen in the series of measurements moving in a straight line aware from the signal source. Of interest is the fact that while there is a good range of green and yellow markers, the markers turn orange and red as the vehicle approaches the signal source. This behavior is due to the fact that in the region near the signal source, the signal is no longer captured by the main lobe (due to the geometry of the antenna and gain pattern), which result in much poorer measurements near the signal source.

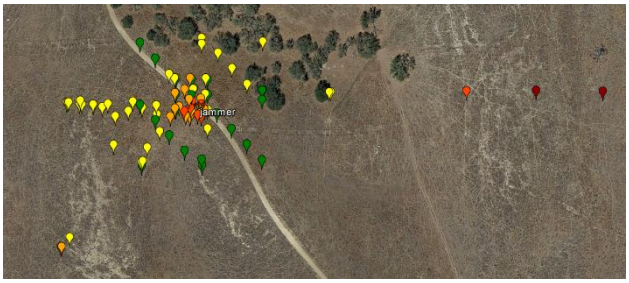

Figure 5: Location of all experimental gain pattern measurements colored by signal strength

Measurement Classification

Due to the effects of distance on the measurements and the configuration of the antenna on the vehicle, all the measurements were split into three different classifications: near, ideal, and far.

Near

Near measurements are measurements made where the signal source is within the cone underneath the vehicle where the main lobe of the antenna no longer reaches as depicted in Figure 6. More specifically this is when the distance between the vehicle and the signal source is less than ~0.58H, where H is the height of the vehicle above the ground. When the signal source was near the vehicle, we did not obtain null measurements, but rather obtained gain patterns such as the one shown in Figure 7. These gain patterns do not resemble the ideal gain pattern of the antenna due to the noise in the measurements from the signal source not being picked up by the main lobe, making it challenging for any of the bearing calculation methods to successfully determine the bearing.

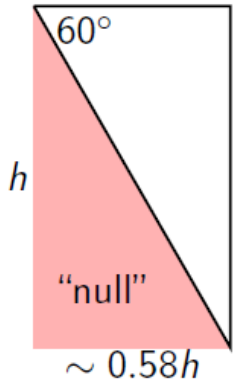

Figure 6: Depiction of vertical antenna pattern below the vehicle

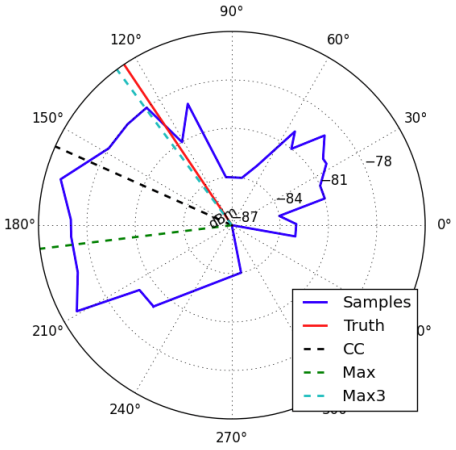

Figure 7: Gain pattern at 7m from signal source (Near)

Far

Far measurements are any measurement further than 200m from the signal source. This limit was determined by examining the measurements made at a range of distances and observing when the majority of the RSSI measurements were very close to the advertised sensitivity of the WiFly module. It also happened to be that at these distances, the resulting gain patterns no longer had enough measurements to clearly resemble the ideal gain pattern of the antenna. Figure 8 shows a gain pattern from 250m away with a true bearing of 267-degrees and demonstrates the partial pattern that is measured. When using cross-correlation, this partial measurement can be mistaken for a side lobe instead of the main lobe which causes the erroneous bearing estimate of 182-degrees shown in Figure 8. The max and max3 methods are still able to calculate reasonable bearings of 276 and 285-degrees, respectively, from these measurements due to the fact that for the most part only the main lobe is able to receive the signal from this distance, so the few measurements are in the correct direction.

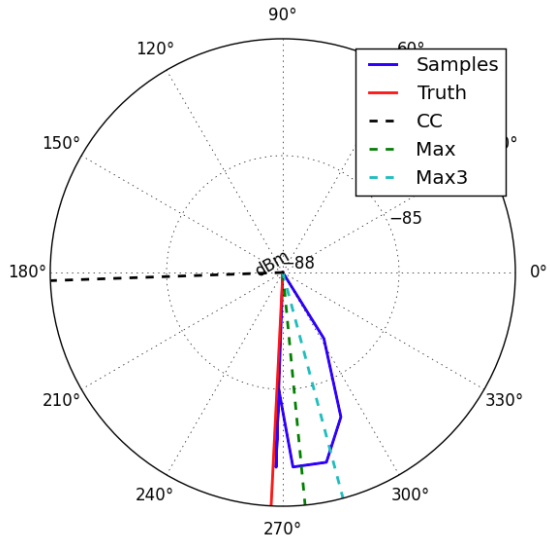

Figure 8: Gain pattern at 250m from signal source (Far)

Ideal

This finally leaves the ideal category, which is any measurement made between near and far. We have called it ideal as it is when the signal source was at an ideal distance from the vehicle to be able to both be within the main lobe of the antenna and within reasonable range of the WiFly's sensitivity. In the ideal range, the gain patterns produced resemble the true pattern of the antenna, as shown in Figure 9.

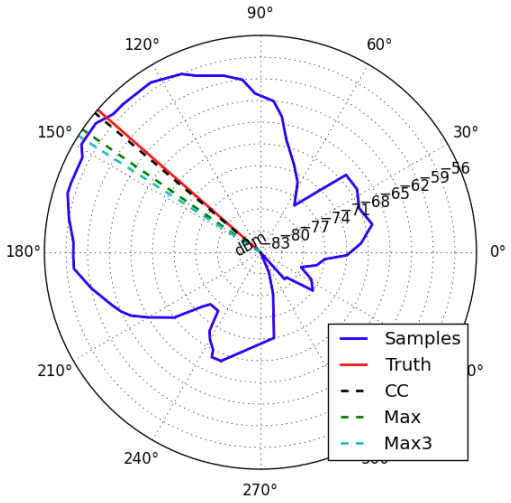

Figure 9: Gain pattern at 26.6m from signal source (Ideal)

In the flight tests performed, the majority of the measurements taken were in the ideal range, as shown in Table 1. Only a couple of measurements were made in the far range so no detailed analysis will be presented for measurements in the far range.

Table 1: Number of gain patterns created in each classification

Near	19
Ideal	67
Far	2
Total	**88**

Performance of Bearing Methods

In this section we highlight the performance of each of the three bearing calculation methods on all the measurements. An overview of the standard deviation of all the results can be seen in Table 2. Overall Max3 outperformed the other two methods. Near the signal source, all three methods are not able to perform well, with all three having very high standard deviations as shown in Table 2. The noise in the measurements near the signal source made each of these methods unreliable. At ideal distances, max3 and cross-correlation performs similarly while max is a little worse.

The simplicity of max and max3 result in these methods outperforming cross-correlation at far distances. This is seen in Figure 10 which shows error as a function of distance for each of the methods.

The proper characterization of the antenna is vital to the performance of the POMDP localization algorithm. Using the results presented with the max3 and cross-correlation methods, the POMDP model can be updated to better reflect the measurements in order to improve the flight profile for localization.

Table 2: Standard deviation [deg] of calculated bearing for each method

	CC	**Max**	**Max3**
Overall	27.9	25.4	22.1
Ideal	13.7	17.9	13.2
Near	45.1	44.1	39.4

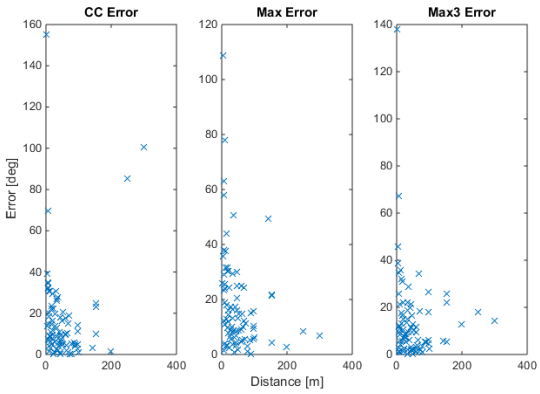

Figure 10: Error as a function of distance for each of the bearing calculation techniques for all measurements

Max

The max method is the simplest method used to calculate the bearing to the signal from a given set of measurements. While overall it does slightly outperform the cross-correlation method, this is due to its apparent superiority far from the signal source in these test flights. While there are too few measurements far from the signal source to warrant detailed analysis, the simplicity of the max method along with the fact that as the antenna gets further from the signal source only the main lobe will be able to pick up the signal, gives it an advantage when far from the signal source.

The simplicity of the max method is also the reason for the poor performance in calculating the bearing. This method can too easily pick a wrong estimate if there is a spike in what should be a smooth main lobe as depicted in Figure 15. These spikes cause a large spread in the errors in calculating bearing seen in Figure 11.

Table 3: Mean and standard deviation for max bearing calculation

	Mean [deg]	**Std. Dev. [deg]**
Overall	4.7	25.4
Ideal	4.3	17.9
Near	5.5	44.1

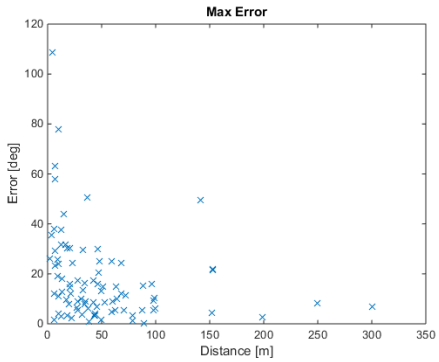

Figure 11: Bearing error as a function of distance for max technique

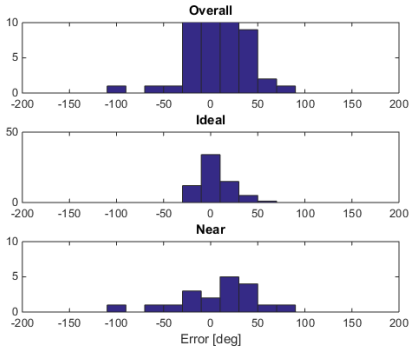

Figure 12: Distribution of errors for max technique for different distance classifications

Cross-Correlation

Cross-correlation is the most complex of the methods used and in the ideal range is one of the best performing methods (on par with the max3 method).

The overall performance of the cross-correlation suffered from the poor performance near and far from the router. Since this method requires a known "truth" pattern, when the experimental measurements don't yield enough results to create a full pattern, the cross-correlation can mistakenly identify the partial pattern for a side lobe instead of a main lobe as was seen in Figure 8.

In the ideal range it greatly outperforms the max method as expected. When looking at the bearing error shown in Figure 13, it can be seen that the errors are much more tightly grouped near zero than those seen in Figure 11 for the max method. The outliers for the far measurements can also be clearly seen in this figure, but this is again due to not enough RSSI measurements to create enough of a pattern to cross correlate against a "truth" pattern created at a much closer distance.

The increase in performance from near to ideal can clearly be seen in Figure 14, where the ideal category has a tight distribution, while the near measurements are very spread out in error.

Table 4: Mean and standard deviation for cross-correlation bearing calculation

	Mean [deg]	Std. Dev. [deg]
Overall	-0.1	27.9
Ideal	0.25	13.7
Near	8.5	45.1

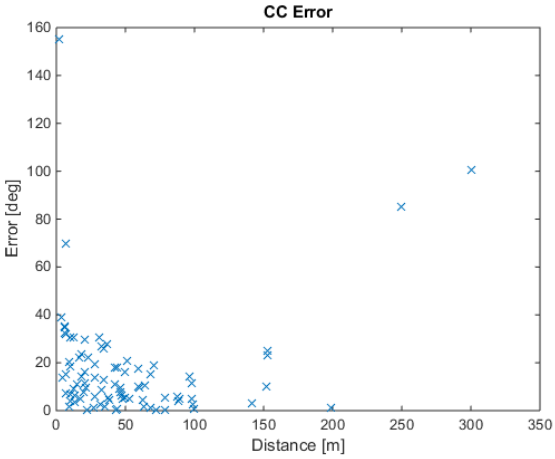

Figure 13: Bearing error as a function of distance for cross-correlation technique

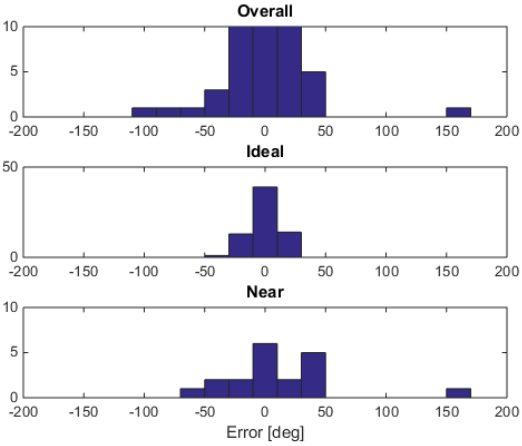

Figure 14: Distribution of errors for cross-correlation technique for different distance classifications

Max3

Max3 is the strongest of the three bearing calculation methods tested; overall it performed the best and max3 has the advantage of simplicity over cross-correlation. It is able to perform on par with cross-correlation in the ideal range and Figure 16 shows a similarly close grouping of errors as was seen in Figure 13 for cross-correlation. Looking at the distribution of the errors paints the same picture with max3's distribution (Figure 17) is very similar to cross-correlation's distribution (Figure 14).

The benefit over the cross-correlation method of not requiring a known "truth" pattern in order to be able to make calculations allows max3 to perform well when the number of measurements is very small and the gain pattern is mostly incomplete. However, just like cross-correlation and max, max3 has difficulty making accurate bearing calculations when close to the router.

The advantage max3 has over the simple max is well illustrated in Figure 15. While the gain pattern looks very promising, there is a spike along the otherwise mostly smooth main lobe at 116-degrees. This spike is off from the true 92-degree bearing which results in the max method estimating an incorrect bearing. By taking the mean of the bearing of the two crossing points 3dB below the max (marked in blue x's), effects from spikes like the ones depicted are reduced allowing for a much better estimate of 93-degrees.

Max3 is able to still perform well at distance because the measurements being made are still good measurements, there is just only a couple of them. This means that the RSSI measurements we do get are in the reasonably correct bearings.

Again near the signal source, max3 suffers from similar problems as max and cross-correlation due to the noisy measurements, though not as badly as the other two methods. Of interest is the near distribution for max3 (Figure 17) compared to the near distribution for cross-correlation (Figure 14) and max (Figure 12). While the distribution can be described as flat for cross-correlation and max, the distribution for max3 is not as flat and begins to form a peak, showing that despite the noise, max3 appears to be able to make slightly more reasonable estimates, albeit with a much larger mean than in the ideal range (Table 5). Through this characterization, the performance seen from max3 can be used to update the POMDP model to improve the localization algorithm, as described in section *(Effect on Localization Algorithm)*.

Table 5: Mean and standard deviation for max3 bearing calculation

	Mean [deg]	**Std. Dev. [deg]**
Overall	4.5	22.1
Ideal	1.2	13.2
Near	15.0	39.4

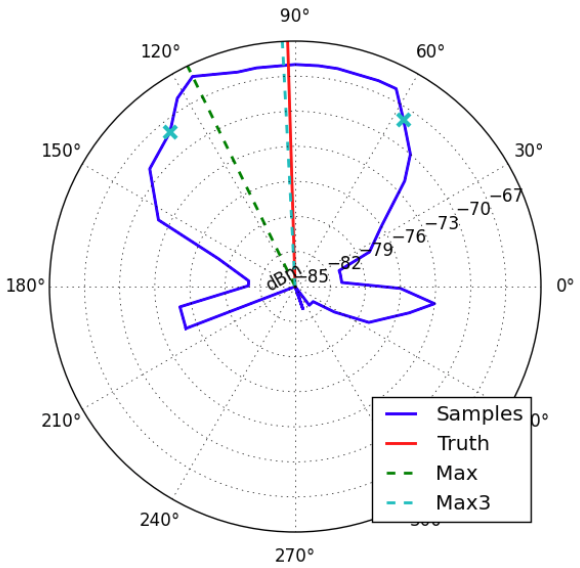

Figure 15: Gain pattern explaining benefits of Max3 method over max

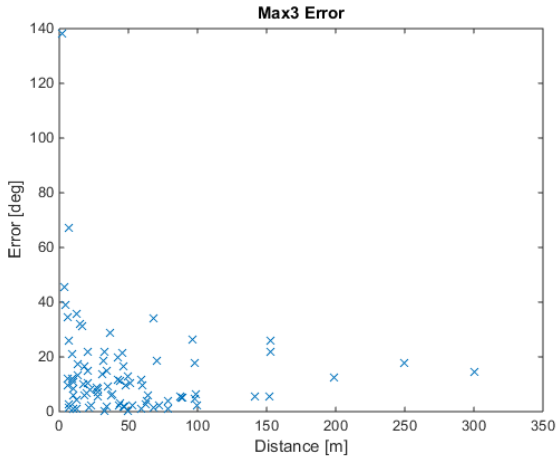

Figure 16: Bearing error as a function of distance for max3 technique

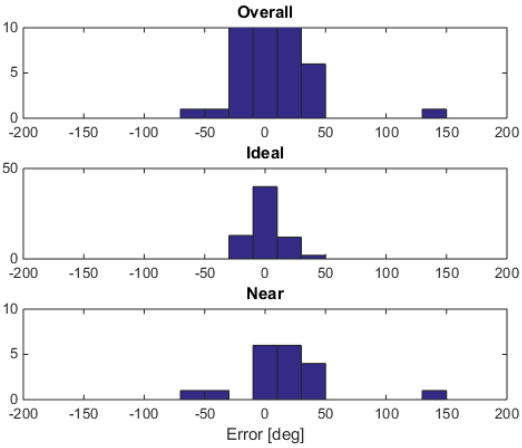

Figure 17: Distribution of errors for max3 technique for different distance classifications

Baseline Localization

One of the goals of our flight tests was to determine the feasibility of the POMDP approach and begin to understand the performance of the POMDP method. A simple baseline method is used for comparison. The baseline method used in this test was a variable step greedy algorithm that moved in the direction of the calculated bearing (using the max method) with a variable step size. The step size used was determined with equation (2) where δ is the tolerance in bearing similarity between measurements (set to 20-degrees for this test), α is the step increase factor (set to 1.5 for these tests), s is the step size and b is the calculated bearing.

$$s_{i+1} = \begin{cases} \alpha * s_i & if\ |b_i - b_{i-1}| < \delta \\ s_0 & \end{cases} \quad (2)$$

Using this baseline method, JAGER was able to move towards the location of the signal source, and with the assistance of a user monitoring the behavior, was able to locate the signal source. The flight path of the vehicle for this test can be seen in Figure 18. It can be seen that the vehicle requires a couple similar measurements before moving quickly towards the signal. Once the vehicle approached the signal source, the jump in calculated bearing resulted in the step size to drop back down to the initial step size.

With a user in the loop with this baseline method, a good estimate of the location can be determined by watching the behavior of the vehicle. Looking at Figure 18, it can be seen that the vehicle kept crossing its path near one location, which can be determined to be an estimate of the location of the signal source. However it is worth noting that the baseline method does take 4 steps to get in the region of the signal source and then another 4 or 5 steps for the user to be confident that the vehicle is in the vicinity of the signal source.

Figure 18: Flight path of the variable step size, greedy localization algorithm. White lines are true bearing from measurement locations, red lines are cross-correlation bearings and black lines are max bearings.

POMDP Localization

With a baseline determined, the POMDP approach was executed from the same starting location and used the simple max bearing method for determining bearing from each location. This localization took a mere 2 steps and 3 measurements to be able to locate the signal source. Figure 19 shows the state updates as the vehicle made subsequent measurements. It can be seen that after the first measurement is made at the starting location, the vehicle is able to immediately narrow down the location of the signal source to a small region within the grid.

It can also be seen that unlike the simple method of moving slowly in the direction of the max bearing, the POMDP method can make large changes in order to get to the next best location to make a measurement.

When running this algorithm, we had an assumption that when the vehicle is in the same cell as the signal source, a null measurement would be made. This is the reason for the vehicle moving to the dark red cells at each of its steps in this execution of the localization. Unfortunately near and over the signal source resulted in noisy measurements and that noise resulted in location of the signal source being off by one cell.

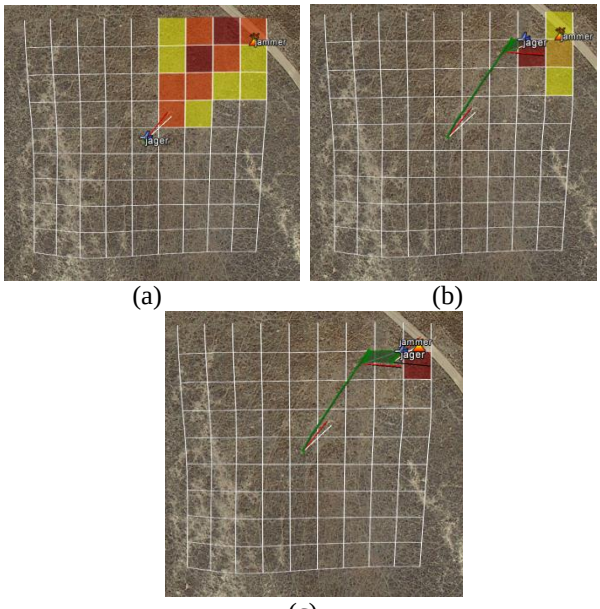

Figure 19: POMDP signal source belief state at each step. Darker the red in the cell, the more likely the signal source is in that location.

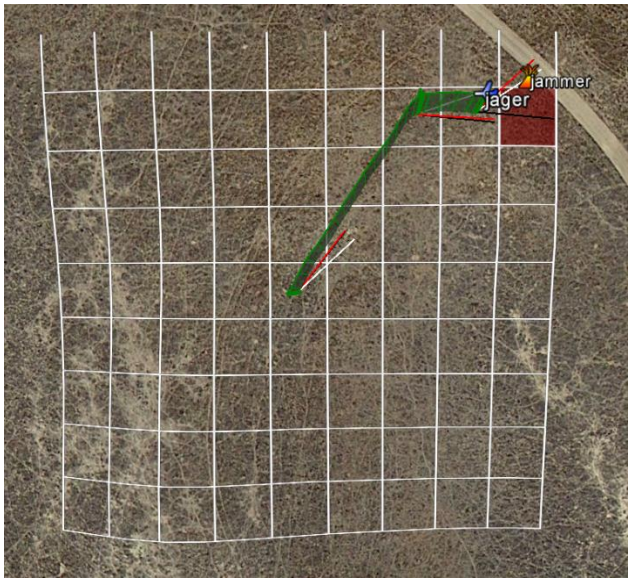

Figure 20: Flight path of the POMDP localization flight with an overlay of the final grid state

Effect on Localization Algorithm

The experiments in this paper were performed to obtain a better observation model for the localization algorithm. Previously, the model assumed 10-degree noise except when the vehicle was in the same cell as the jammer; the model assumed no bearing measurement would be obtained there. These assumptions were used in the experimental trajectory shown in Figure 20 and affected the selected trajectory. The vehicle always moved towards regions with high probability of containing the jammer (the dark red cells). Because we assumed that rotation would only yield a null measurement when over the jammer, receiving a null observation after rotating would convince the vehicle that the jammer was in its current cell. For this reason, the vehicle moves to regions with high probability of containing the jammer; it hoped to receive this high-information measurement and solve the problem with a single rotation.

Experimental results have shown that measurement noise increases greatly close to the jammer. Our new model assumes 40-degree noise if the jammer is in any of the adjacent grid cells when the vehicle rotates, and 13-degree noise if the jammer is farther away. If the vehicle rotates in the same cell containing the jammer, it no longer receives a null measurement. Instead, it can receive any measurement with uniform probability.

Generating a policy with this new model leads to different trajectories. A simulated rerun of the experimental trajectory from Figure 20 is shown in Figure 21. The vehicle avoids the darker cells, which indicate higher probability of containing the jammer. Proximity to the jammer increases measurement noise, reducing the information provided by a measurement. This reduced information requires more rotations to be performed before localizing the jammer, so the vehicle chooses to rotate in cells it believes are farther away from the jammer.

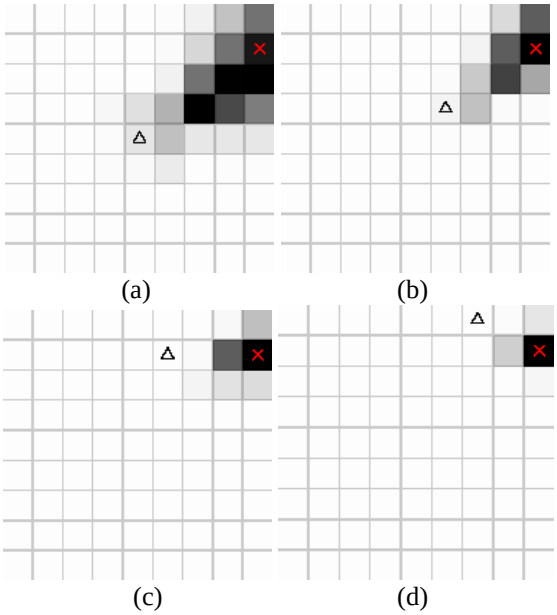

Figure 21: Simulation steps of POMDP with updated model.

CONCLUSION

This paper presents the development of the localization component of a UAV to locate the source of a GPS jamming signal. For the scenarios tested, modeling the localization as a POMDP is a viable solution that is able to locate a static signal source in very few steps. It is faster and has greater confidence than a simple, greedy search baseline solution.

Through extensive test flight using a single directional antenna and rotation-based measurements, three different bearing methods have been analyzed. All three methods suffered when near the signal source due to antenna reception pattern which resulted in very noisy measurements. Of the three, max3 and cross-correlation faired the best in the ideal distance from the signal source. Max3 was able to outperform cross-correlation when the UAV was far the from signal source due to the limitations of cross-correlation requiring a "truth" pattern for correlation. However, cross-correlation can also provide a useful correlation coefficient that can be used in the future to merge several bearing calculation methods.

The characterization of antenna bearing performance is a vital component to the localization process. The characterization affects the optimal behavior determined by POMDP. When we changed our initial assumptions about measurement performance near the jammer to one better informed by our tests, the actions determined POMDP resulted in a significantly different profile.

ACKNOWLEDGMENTS

The authors would like to gratefully acknowledge the Naval Postgraduate School for providing an unmatched space to be able to perform test flights of the JAGER system at the Joint Interagency Field Experimentation events.

The authors would also like to thank the Stanford Center for Position Navigation and Time (SCPNT) and its members for supporting this work.

REFERENCES

[1] G. Gibbons, "FCC Fines Operator of GPS Jammer That Affected Newark Airport GBAS," Inside GNSS, 30 August 2013. [Online]. Available: http://www.insidegnss.com/node/3676. [Accessed 24 September 2015].

[2] J. Spicer, A. Perkins, L. Dressel, M. James, Y.-H. Chen, S. Lo, D. S. De Lorenzo and P. Enge, "Jammer Hunting with a UAV," *GPS World*, pp. 30-38, May 2015.

[3] S. Venkateswaran, J. T. Isaacs, K. Fregene, R. Ratmansky, B. M. Sadler, J. P. Hespanha and U. Madhow, "RF Source-Seeking by a Micro Aerial Vehicle using Rotation-Based Angle of Arrival Estimates," in *American Control Conference (ACC), IEEE*, 2013.

[4] M. Geyer and R. Frazier, "FAA GPS RFI Localization Algorithm," in *Proceedings of the 12th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Nashville, TN., 1999.

[5] E. M. Geyer and B. M. Winer, "Airborne GPS RFI Localization Algorithms," in *Proceedings of the 10th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Kansas City, MO, 1997.

[6] L. P. Kaelbling, M. L. Littman and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99-134, 1998.

[7] O. Madani, S. Hanks and A. Condon, "On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision problems," in *AAAI*, 1999.

[8] L. K. Dressel and M. J. Kochenderfer, "Signal source localization using partially observable markov decision processes," in *AIAA Infotech*, Kissimmee, FL, 2015.

[9] M. Hauskrecht, "Value-function approximations for partially observable Markov decision processes," *Journal of Artificial Intelligence Research*, vol. 13, pp. 33-94, 2000.

[10] J. Pineau, G. Gordon and S. Thrun, "Point-based value iteration: An anytime algorithm for POMDPs," *IJCAI*, vol. 3, pp. 1025-1032, 2003.

[11] H. Kurniawati, D. Hsu and W. S. Lee, "SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces," in *Robotics: Science and Systems*, Zurich, Switzerland., 2008.

[12] Pixhawk, "PX4 autopilot project," [Online]. Available: https://pixhawk.org/.

[13] L-com, *HyperLink Wireless 2.4 GHz 9dBi Radome Enclosed Wireless LAN Yagi Antenna Datasheet.*

[14] J. Graefenstein, A. Albert, P. Biber and A. Schilling, "Wireless Node Localization based on RSSI using a Rotating Antenna on a Mobile Robot," in *PROCEEDINGS OF THE 6th WORKSHOP ON POSITIONING, NAVIGATION AND COMMUNICATION*, 2009.