

Real-Time Unmanned Aerial System (UAS) Based Interference Localization in a GNSS Denied Environment

Adrien Perkins, *Stanford University*
Yu-Hsuan Chen, *Stanford University*
Sherman Lo, *Stanford University*
Chiawei Lee, *USAF Test Pilot School*
J. David Powell, *Stanford University*

BIOGRAPHIES

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 Bachelor of Science in Mechanical Aerospace Engineering from Rutgers University in 2013 and his Master of Science in Aeronautics and Astronautics from Stanford University in 2015.

Yu-Hsuan Chen is a research engineer at the Stanford University GPS Laboratory.

Sherman Lo is a senior research engineer at the Stanford University GPS Laboratory.

Chiawei Lee is an Instructor Flight Test Engineer and Chief of Test Safety at the United States Air Force Test Pilot School. His primary responsibility is in the Flight Test Foundations Branch where he oversees the student and staff Test Management Projects. He received his Bachelor of Science in Aerospace Engineering from the University of California, Los Angeles and Master of Science in Aeronautics and Astronautics from Stanford University.

J. David Powell is a professor emeritus of Aeronautics and Astronautics at Stanford University. He received his Ph.D. in Aeronautics and Astronautics from Stanford University in 1970. He has conducted research on the use of GPS-based attitude determination augmented with inertial sensors, applications of the FAA's WAAS for enhanced pilot displays, flight inspection of aircraft landing systems, and the use of WAAS and new displays to enable closer spacing of parallel runways.

ABSTRACT

An unmanned aerial system (UAS) based system has been developed to be able to rapidly localize the source of GNSS interference to minimize the impact interference would have on critical infrastructure, such as an airport. Previously, the project, known as Jammer Acquisition with GPS Exploration and Reconnaissance or JAGER, has demonstrated capabilities of localizing a signal source in real time by physically rotating the antenna, however this approach proved to be a slow approach to localization [1]. Since that demonstration, the sensor has been upgraded to be electronically steered, the localizing algorithm and path planning has been improved to account for the higher rate of bearing measurements, and a GNSS independent navigation system has been developed. Each of these pieces have been demonstrated individually in both flight tests and simulation, however, until now, the full system has not been operated together in real-time [2]–[4]. This paper presents the real-time implementation of all the core sub-systems making up the JAGER platform working together to localizing a signal source in a GNSS denied environment.

The JAGER system is made up of three sub-systems: the interference localization system, a GNSS denied navigation system, and a path planning system. The interference localization system is comprised of an electronically beam steered antenna making signal strength measurements to create received signal strength (RSS) patterns from which bearing information can be used to estimate the RFI source location. The GNSS denied navigation system is built around using optical flow from a gimbal stabilized FLIR Systems Boson infrared (IR) camera to measure velocity and in turn estimate the position of the vehicle. Finally, the path planning system sends velocity commands to the autopilot to fly a desired trajectory to get the best bearing measurements for estimating the RFI source.

Flight tests were performed at Edwards Air Force Base with a 2W, 5MHz wide, 2.48 GHz signal source being localized by JAGER. The GNSS denied environment was created in software by disabling the use of GNSS by the autopilot for real time guidance and control. The flight tests were conducted at night, in the early morning (pre-sunrise), and in the late morning to stress test the IR based vision system and demonstrate the capabilities in a range of time periods. The location chosen at Edwards Air Force Base also consisted of a large hangar door that presented “urban environment”-like challenges (namely reflections of the signal source). The flight tests comprised of 35 different mission attempts, where a mission consists of starting from a known location and fully autonomously executing the localization of the signal source with no user interaction until the source was determined to be localized. Of the 35 different mission attempts, 16 successfully localized the signal source with an average final 2D localization error of 13m, and the vision system demonstrated a drift rate of 4% and 6.1% during the day and night, respectively.

INTRODUCTION

Three key systems, radio frequency interference (RFI) localization system, the path planner, and the GNSS independent navigation system, must work in concert to allow JAGER to localize GNSS RFI. In addition to these systems, JAGER is equipped with a Pixhawk autopilot, running PX4 that handles the autonomous control of the drone throughout the flight. A high level view of the architecture for the real-time implementation used in this paper is depicted in Figure 1, and the physical implementation is depicted in Figure 2. For this real-time implementation, all three systems are implemented using the Robot Operating System (ROS) framework on an Intel

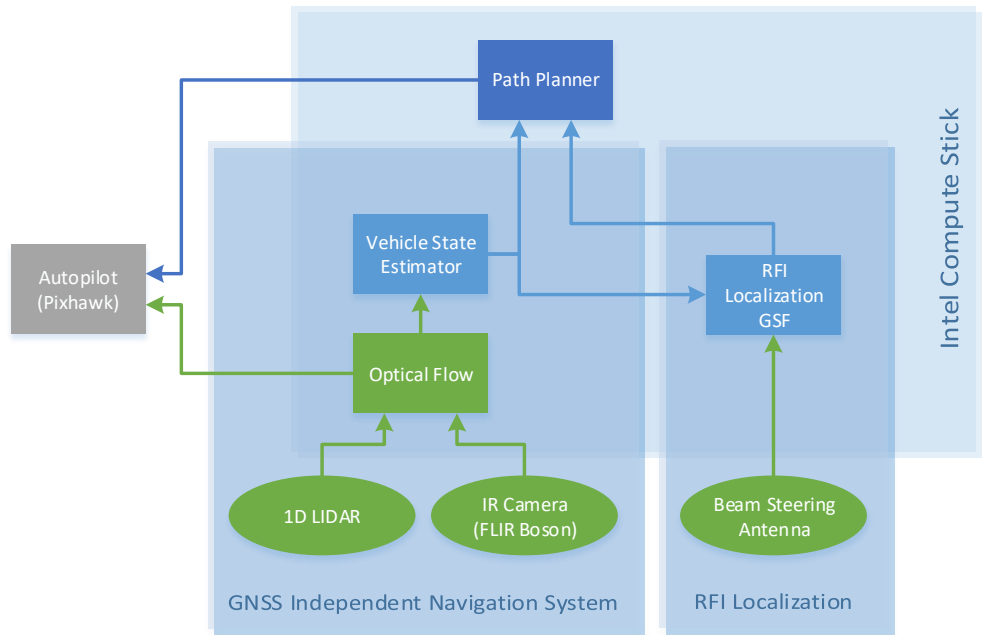


Figure 1: JAGER system overview for the real-time implementation used

Core-i3 Compute Stick. This paper examines the operation and performance of these three systems under real time RFI through a series of localization flight tests. To successfully localize the RFI source in flight, the system had to overcome a few key challenges posed by the environment: first, the localization algorithm and the path planner had to work together, using real-time noisy measurements and estimates, to fly the best trajectory for localization, resulting in the development of a two phase path planner; second, the GNSS independent navigation system had to handle the different quality of features during different times of day to provide an estimate of the vehicle’s location that was usable by the RFI localization system; and lastly, all of these systems had to operate under real-world conditions that presented varying levels of noise in the measurements, resulting in the development of a confidence metric in the measurement to remove noisy measurements from the RFI localization system.

The localization system is comprised of two main pieces: the sensor and the estimator. The sensor used is a 3-element beam steering antenna (BeaSt) that measures the signal strength at a collection of azimuth angles to create an receive signal strength (RSS) pattern [5]. From the RSS pattern, bearing information can be extracted by finding the location of the lobe(s) present in the pattern [4]. The estimator (the RFI Localization GSF block in Figure 1) is described in more detailed in the following section.

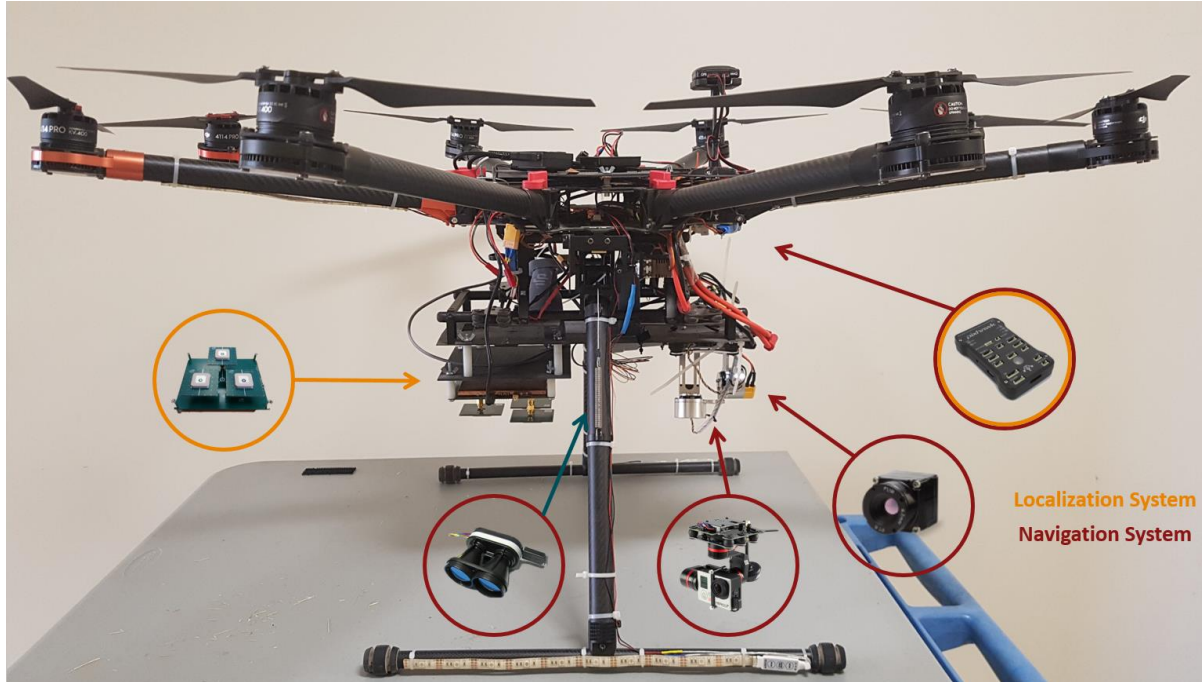


Figure 2: JAGER with the various sensors highlighted

The GNSS independent navigation system is comprised of a mechanically stabilized FLIR Systems Boson 640 infrared (IR) camera with a 4.9mm lens providing a 95 degree horizontal field of view, optical flow processing, a Lightware SF 11/C one dimensional (1D) light detection and ranging (LIDAR) and vehicle velocity estimator. Nominally the camera is capable of a frame rate of 60Hz, however, due to computation limits on the Intel Core-i3 Compute Stick, for these flight tests the camera operated at 30Hz. The series of images is used to provide optical flow measurements using OpenCV's 'goodFeaturesToTrack' for feature detection of up to 250 features per frame and OpenCV's implementation of an iterative Lucas-Kanade method with pyramids to calculate the optical flow [2], [6], [7]. The optical flow measured from the camera is converted to velocity using the known height above the ground, as measured by the 1D LIDAR. This velocity measurement is used by an extended Kalman filter (EKF) that estimates the 2D position and 2D velocity of JAGER throughout the flight [2] using only the take-off position to initialize. This system has been previously discussed and has been demonstrated to perform with a drift rate of 0.4% of the distance traveled [2]. A caveat with these prior results is that JAGER flew a predetermined path using GNSS so it was able to execute very smooth maneuvers and maintain the desired speed. This will not be true for an operational system flying in RFI and in this paper, the autopilot onboard only uses the information from the optical flow, hence no GNSS, for navigation and control. This results in a degraded ability to maintain a target velocity and reduces the overall performance of the navigation system (i.e. increase the drift rate). In the Flight Test Results section, those effects are presented and discussed in more detail.

The path planner onboard uses the current state estimate for the signal source location from the localization system and the current estimate of JAGER's position from the navigation system to compute the desired velocity vector to fly a circular trajectory around the signal source. The following section provides a more detailed description of the path planner. The desired velocity vector is the command that is sent to the Pixhawk autopilot which handles the actual control of the drone itself. For these flight tests, the Pixhawk autopilot was configured to have no access to GNSS position and velocity information for any of its command and control capabilities. This effectively creates a GNSS denied environment to the autopilot (the inflight or real time operations of JAGER). Instead, the autopilot was configured to use the optical flow information directly and was therefore closing all of its internal control loops on the optical flow information.

For the flight tests described in this paper, the goal was to repeatedly run localization missions to evaluate and demonstrate the capability of the localization, path planning and navigation systems to work in concert to successfully localize a signal source. Relatively benign and challenging conditions were tested. The paper is organized as follows. The first section details the localization and path planning algorithms needed to understand the mission trajectory. The Test Setup section provides a more detailed description of the physical implementation of the sensors and algorithms onboard JAGER, as well as describes the test environment and signal source used throughout the flight tests. Next, the Flight Test Results section discussed several missions

in detail and provides results for all the missions flown throughout test campaign. Finally, the paper provides some closing conclusions on the test campaign and addresses areas of remaining work on the JAGER system.

LOCALIZATION AND PATH PLANNING

Localization of the RFI takes measured bearings to derive an estimated jammer location and confidence which then feeds the path planner. BeaSt provides a high rate of bearing measurement (configured for 3Hz for these flights) enabling the use of existing bearing-only localization algorithms [5]. This algorithm takes multiple measurements while handling various measurement errors and uncertainties to provide an estimate of RFI source location and confidence in the localization. The path planning algorithm takes this information to determine how to move quickly to better localize the source. This section briefly discusses the localization algorithm used throughout this test campaign and special considerations to the existing path planning strategy needed to enable real-time operations [2].

Localization Estimator

In a perfect environment, at each time step, there would only be as many bearing measurements as RFI source, and those bearing measurements would be in the direction of their respective source, with some Gaussian noise. However, due to regulatory and size reasons, BeaSt only has 3 elements, which result in RSS patterns with large side lobes considered to be clutter – unwanted echoes in the system and measurement which can cause erroneous or additional bearing measurements. For a single source, a typical RSS pattern has two or three lobes, as shown in the theoretical patterns in Figure 3a and Figure 3b, meaning that a typical measurement can result in up to three possible bearings depending on the relative strength of the lobes, of which two would be considered due to clutter.

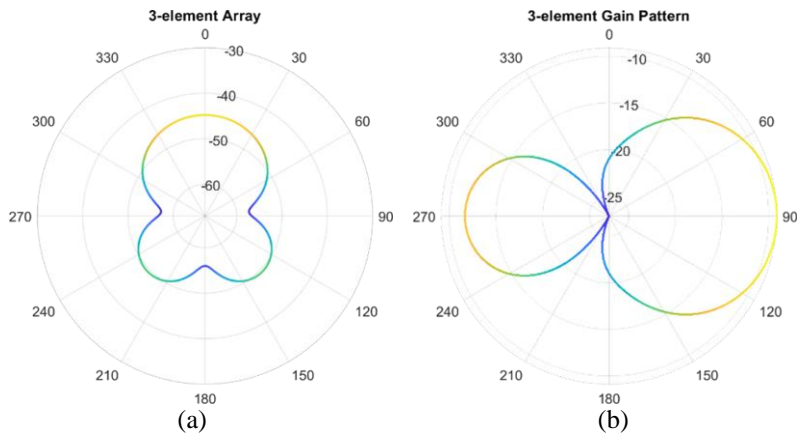


Figure 3: example RSS patterns with (a) one main lobe and two side lobes and (b) one main lobe and one side lobe

In addition to clutter caused by BeaSt’s geometry, the environment can also corrupt the RSS pattern. For example, multipath can cause interference in the RSS values making up the RSS pattern, resulting in the corrupt pattern shown in Figure 4. Therefore, a confidence metric was implemented that is computed based on how well the RSS pattern matches an expected, nominal, RSS pattern for a single source (the patterns in Figure 3). This confidence metric is used to ignore bearing measurements from an RSS pattern with low confidence.

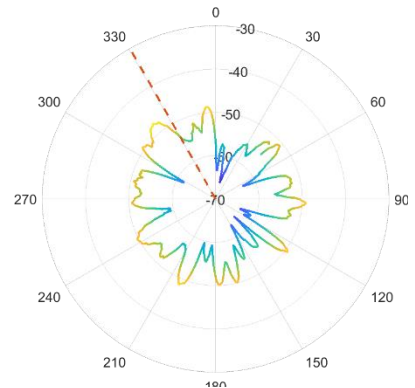


Figure 4: example pattern corrupted by noise

As a result, a bearing only localization algorithm capable of handling cluttered measurements and corrupt patterns was implemented. At its core, the localization estimator uses a Gaussian Sum Filter (GSF) that enables handling the range ambiguity of bearing only measurements [8], [9]. Range ambiguity exists because the transmitted RFI signal power is not known. So, from any individual measurement, we cannot tell if the measurement is due to a faraway strong signal or a weak nearby signal or something in between. A GSF is used to be able to represent the “cone” of uncertainty that results from a bearing only measurement, depicted in Figure 5. In the example shown in Figure 5, the range uncertainty is handled with the creation of multiple weighted individual estimates, with three shown in the figure, which sum together to create the GSF. The initial distribution of the estimates is made to span the specified minimum and maximum range and is created using the method in [9]. The minimum and maximum is a user-settable parameter for the filter and may be chosen based on factors such as sensitivity and anticipated jammer power levels. For these flight tests, the minimum and maximum range was set to

10m and 500m, respectively, though these values do not greatly affect the performance of the estimated, provided that the true range lies within the bounds (i.e. a conservative value of 2km for a maximum range could be set if desired, but 500m was chosen based on the capabilities of the 3-element beam steering antenna with the interference source power used in these flight trials).

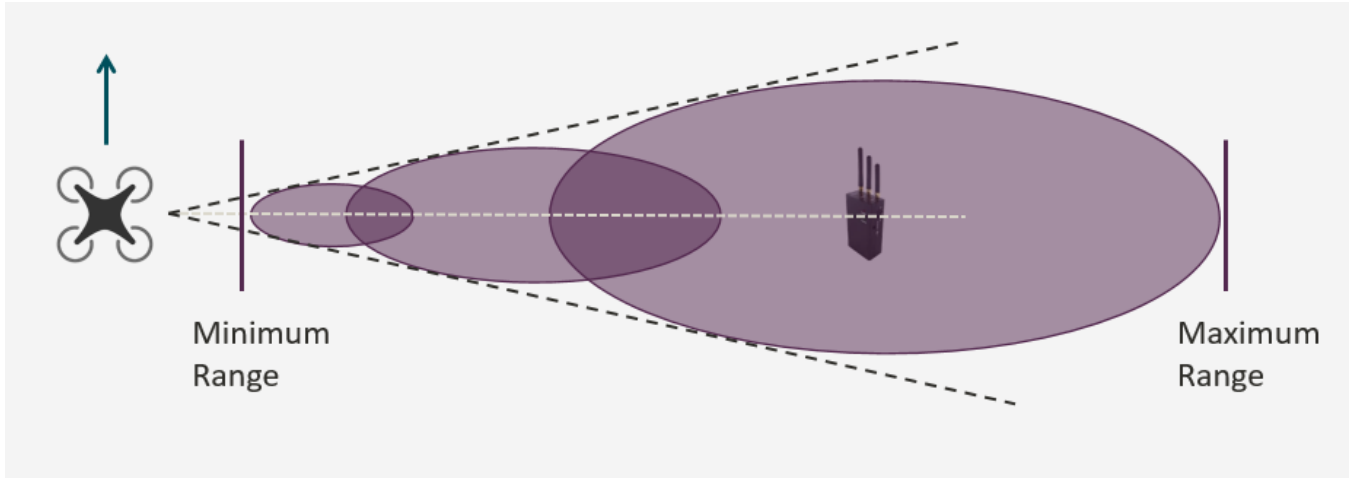


Figure 5: depiction of the initialization of a Gaussian Sum Filter to estimate the location of a signal source given a bearing measurement and a defined minimum and maximum range

To handle the cluttered measurements, instead of using a single GSF, the algorithm employs a bank of GSF [8]. At each time step, each of the bearing measurements update the GSF with the highest association or initializes a new estimate if none of the associations are above a given threshold. Additionally, estimates that have not been updated recently are pruned periodically to keep the bank of filters small. Each of the filters in the bank are weighted based on likelihood that the filter represents the estimate of the true source which is needed by the path planner, described in the following sub-section.

Path Planning

For bearing-only localization, it is important to have geometrically diverse measurements which means the sensor must “out-maneuver” the source, for example by encircling the source. For a static source, this means that JAGER needs to move in a direction that is not directly towards or away from the source location. Previous work has demonstrated that a very effective path planning strategy is to execute either a circle or an inward spiral around the signal source, such as the circle trajectory flown in Figure 6 [2]. There are merits to the different approach angles, however for short ranges (e.g. < 500m) from the signal source, the difference between the circle and inward spiral trajectory is minimal, therefore, for this demonstration, a circle approach is taken [2].

Flying a circle path requires roughly knowing the location of the RFI source to determine the direction of travel. Theoretically the bearing measurement could be used to determine the initial direction of travel. In practice, since the measured RSS pattern can have up to 3 different bearing measurements, the path planner is unable to use the bearing measurements to determine the direction of travel. Therefore, the path planner is configured to use the bearing to the current estimate of the source location as a guide for flying the circular trajectory. However, given that when the missions starts – due to the clutter in the bearing

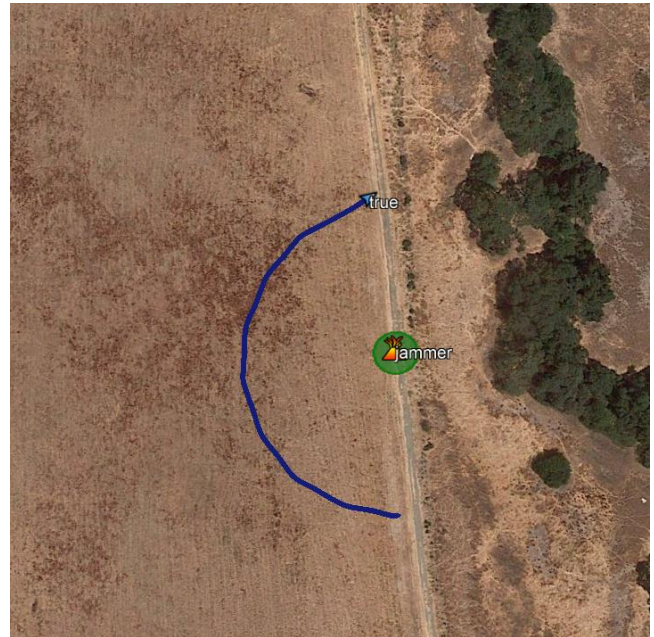


Figure 6: example circle trajectory flown around the RFI source resulting in very accurate estimate of the signal source location (green circle)

measurements – there is an initialization period before there is an estimate for the source location, the path planner was modified to support two phases of flight: an initialization phase and a controlled localization phase, as outlined in Figure 7. During the initialization phase, JAGER is configured to fly a straight line in a direction that is set by the user at takeoff. Once the localization algorithm has initialized a high likelihood estimate for the source location, then the path planner switches to a controlled localization phase. During this controlled localization phase, JAGER is configured to fly a circle around the mean of the current estimate for the source location (where a circle is simply defined to be 90 degrees from the bearing to the estimate—this is an open loop circle, there is no range information being used). This two phase approach is apparent in the flight paths shown in the Flight Test Results section.

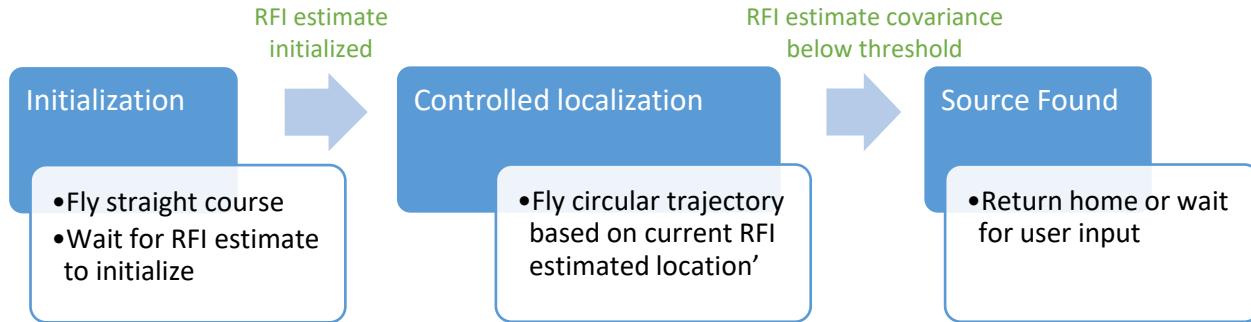


Figure 7: overview of the path planner stages and actions within each stage

TEST SETUP

This section describes the environment at used during our flight tests at Edwards Air Force Base, provides details on the signal source used, and explains the three times of day selected for all the flight tests.

The Environment

The flight tests were performed at Edwards Air Force Base at the location shown in Figure 8. This environment was chosen because it possesses many of the potential features seen at an airport, specifically a taxiway and some off-taxiway vegetation, which is an example deployment location for a system such as JAGER. Furthermore, the environment also presented some urban environment type challenges, namely a large metal hangar door that reflected the RFI source signal. This environment challenged both the localization scheme and the navigation scheme, enabling both the demonstration of the capabilities of JAGER and the understanding of the limitations of the system as it is designed. Localization is challenged by the variety of the environment (e.g. the hangar door) and navigation is challenged by the potentially feature poor environment of an airport.



Figure 8: satellite view of the test environment, with the large red box depicted the rough flight region, the yellow and orange boxes marking regions of interference, and the pins marking starting locations for the missions flown

The test area, depicted in dark red in Figure 8, was about 270m long and 175m wide. Within the test area, there were roughly three different starting areas. Starting area A

is about 80m north of the source, starting area B is about 100m east of the source, and starting area C is about 90m south of the source.

Within this test area, the hangar provided two regions of different types of interference. Figure 8 highlights two different regions in orange and yellow. In the orange region, the reflected signal appeared stronger than the main signal from the source, resulting in a distorted RSS pattern with the lobe pointing towards the hangar door rather than towards the source itself, as seen in the pattern in Figure 9 (the dashed line represents the true bearing to the source and the solid cone represents the measured bearing). For comparison, Figure 11 shows an RSS pattern measured outside of these regions of interferences and has a lobe aligned with the true bearing to the source. In the yellow region, the interference resulted in the RSS patterns not having any distinct main lobe pointing in any direction, such as the pattern in Figure 10. The effect and handling of these regions are discussed in more detail in the Flight Test Results section.

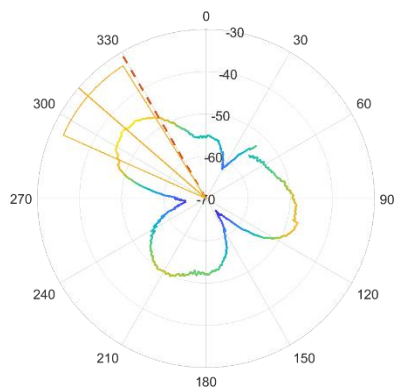


Figure 9: example RSS pattern distorted by strong reflected signal, resulting in lobe pointing towards hangar door rather than true bearing (dashed line)

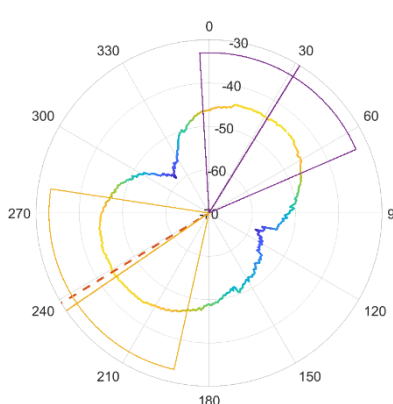


Figure 10: example RSS pattern from a region with no interference with a measured bearing aimed towards the true source location

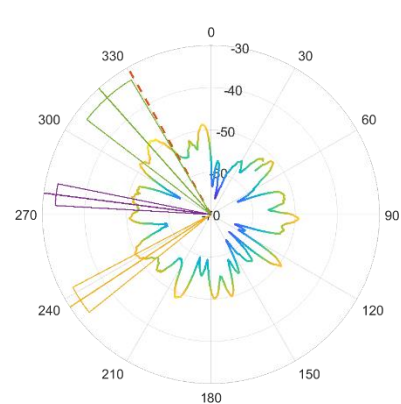


Figure 11: example RSS pattern corrupted by interference resulting in no predominant lobe and incorrect bearing measurements

The Time of Day

For these tests, flights were performed in three different times of day, each providing a very different environment and challenge for the GNSS independent navigation system: late in the night (at least an hour after sunset), early in the morning (about an hour before sunrise), and late in the morning. The late night flights provided an environment to demonstrate the capabilities of an IR camera over a vision camera as it is a time where a vision camera would not be able to work successfully. The difference in image definition is apparent in Figure 12 which shows an image from vision camera (left) and an IR camera (right) taken around 11 pm local time.

Even though the IR camera is capable of operating at night, daytime flights were flown to assess the difference in performance between day and night conditions, as daytime conditions do generally offer better features, even in IR. Figure 13 shows IR images from above the side of the taxiway taken during the day (left, a) and at night (right, b). It can be seen that there are more details in the day image (a) than in the night image (b). For both of these images, the camera is in a “white hot” configuration, meaning that the white regions of the image are the hottest and the black regions of the image are the coolest. It can be seen in these images that this same terrain changes in what is relatively the warmest and coolest region of the image from day to night. This transition happens fairly quickly in the day to night transition, but happens fairly slowly in the night to day transition, which motivated the inclusion of the early morning time period for the flight tests.

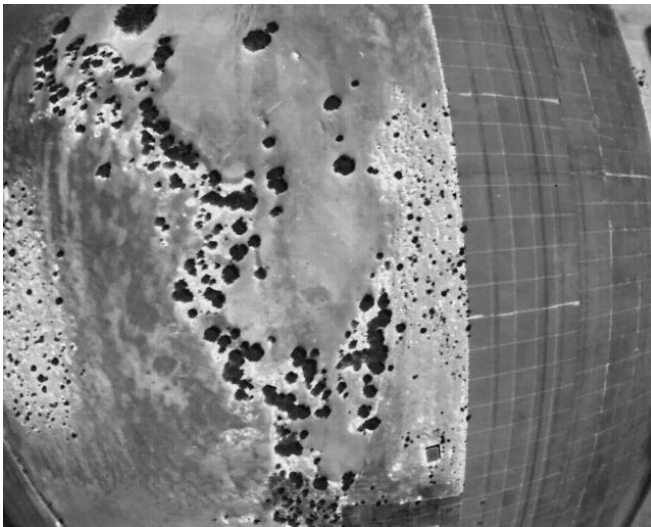


(a)

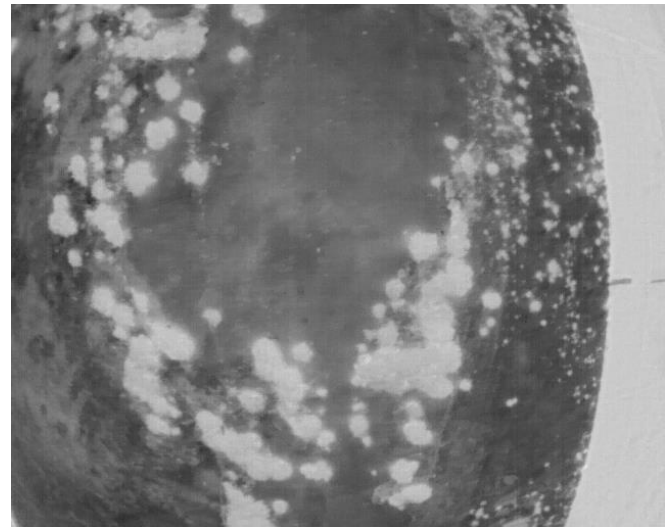


(b)

Figure 12: comparison of regular vision (a) and IR (b) camera performance at 11pm



(a)



(b)

Figure 13: comparison of IR image during the day (a) and at night (b)

Signal Source

The RFI source used for these flight trials was in the 2.4 GHz band, as it was not possible to broadcast interference in the GNSS bands due to legal reasons. Furthermore, BeaSt, the beam steering antenna, is designed and tuned for the 2.4 GHz spectrum to enable practice flights on Stanford campus with the same 2.4 GHz based RFI source. Note that the entire system design is agnostic to frequency (with the exception of the physical layout of the antenna element on BeaSt and the antennas themselves) so the resulting behavior for the localization mission, be it tuned for 2.4GHz or 1.575GHz, is the same.

The signal source used was a commercial off-the-shelf video transmitter connected to a 2 W power amplifier with a 3dB omnidirectional dipole antenna. The transmitter was centered on 2.48 GHz and had a bandwidth of about 5 MHz. This 5 MHz bandwidth source is different from the more typical chirp style GNSS jammer that can be found online [10], but lab tests comparing the video transmitter to a generated chirp source resulted in the RSS patterns as measured by BeaSt to be nearly identical between the chirp style source and the commercial off-the-shelf video transmitter used in these flight trials.

The video transmitter was set up on a tripod, depicted in Figure 14, that also contained a GNSS receive to survey the location of where the tripod was placed yielding a very accurate location of the source for the analysis of the results and performance of the system.

FLIGHT TEST RESULTS

Throughout the test campaign 35 missions were performed from the three different starting locations previously described. Of the 35 missions flown, 19 resulted in successfully localizing the source, 15 flights initialized to the wrong location, resulting in an inability to properly localize the source, and 1 flight never initialized, also resulting in an inability to localize the source. A more specific breakdown of the number of successful localizations and failed initializations for each starting point is shown in Figure 15 and recapped in Table 1. Of the 35 missions, 18 were flown at night, 10 were flown in the early morning, and 7 were flown in the late morning. For each time of day, the flights were fairly evenly distributed across the starting locations.

This paper examines a few example missions in detail to highlight the performance of the systems facing different challenges: two successful and one unsuccessful to illustrate how the system performed. Additionally, some high-level statistics and information on all the flights are presented to give an overall view of the flight tests. It is worth noting that all the successful missions played out very similarly and therefore the two successful missions discussed in more detail here are representative of the rest of the successful flights. Additionally, the missions that failed also played out very similarly due to the failure condition being the same across the missions.



Figure 14: mobile signal source platform used in flight tests



Figure 15: overview of the mission success from the three different starting locations

Table 1: mission success overview, broken down by starting location

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

Mission 1

The first mission discussed starts from location B in a counter-clockwise direction at 10:45pm. Most of the successful missions followed the same main elements as this flight as presented here. All the phases of the flight are laid out in Figure 16. The first step, shown in Figure 16a, is the end of the predefined straight line for the initialization of the estimate, in this case a heading of 80 degrees was chosen. Once an estimate reached a high enough confidence level to be initialized as the likely source estimate, the result of the flight plays out as an arc around the source until the covariance of the estimate drops below a given threshold, as shown in Figure 16b, and the filter considers the source as found. Since we had a real-time display of the estimate, we let JAGER continue to fly to both collect more data and reduce the covariance size further, to the final state, shown in Figure 16c.

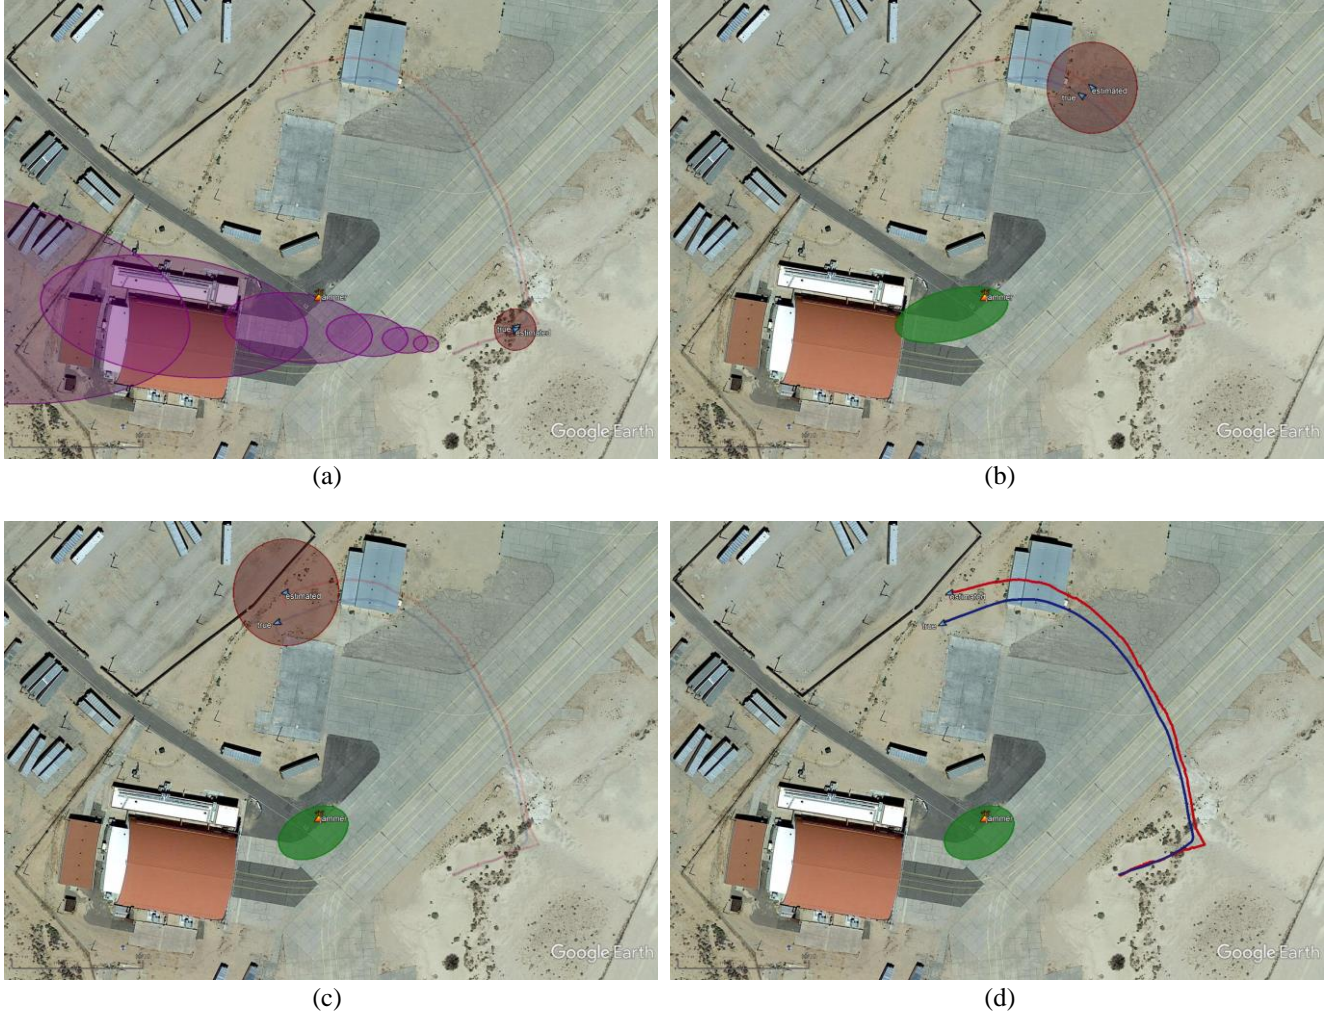


Figure 16: snapshots of different states of a mission flown at 10:45pm from starting location B with the source estimate in purple and green, the true flight path in blue, and the estimated flight path in red.

Looking at the full highlighted path in Figure 16d, the expected circular trajectory does not have a constant radius of curvature. This change in radius of curvature is a result of the path planner algorithm using the mean of the estimate for the center of the circle. Therefore, in the early stages of the controlled localization, when the GSF is still comprised of many evenly weighted estimates (e.g. the state in Figure 16a), the estimated position of the source is at a point much further out. As more measurements come in and the mean of the estimate moves closer to the true source location, the radius of curvature decreases, resulting in the tighter circle seen in the later portion of the flight.

Figure 17 shows the covariance of the estimate and the true position error in the estimate for the source location throughout the flight. Note that the covariance and error is plotted against the angle traveled around the source rather than time as plotting against angle traveled normalizes the results of different starting distances from the source and removes some distortion that results from JAGER not flying at a perfectly constant velocity. It can be seen that once JAGER has flown enough around the source, the covariance of the estimate quickly drops. The threshold used to consider the source as found is when the 95% confidence ellipse is less than 1000 m², which in this case is triggered at about 30 degrees traveled around the source. In simulation, the more typical result is closer to 100 degrees around the source, and as is discussed later, this mission is atypically quick in localizing the source. However, all missions follow a similar trend with the covariance, and position error, quickly dropping after enough of an angle around the source has been traveled.

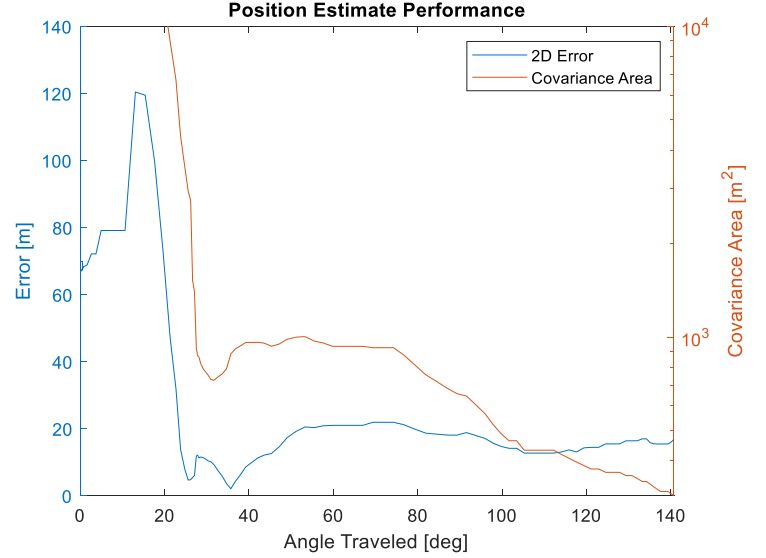


Figure 17: performance of the estimate of the source localization for mission 1

On the navigation front, the error in the estimated location of JAGER compared to the GNSS position of JAGER is plotted in Figure 18. Additionally, the velocity, as measured by the optical flow system and as estimated by the EKF used compared to the GNSS velocity measurement, is plotted in Figure 19. The EKF for the state of the vehicle assumes a constant velocity, which is not the case with the path planning strategy having to take the turn from the predefined heading to the circle trajectory, results in a sharp jump in position error. Overall the error growth over time, for this flight, was 6% of the distance traveled, which was more typical for the flights at night.

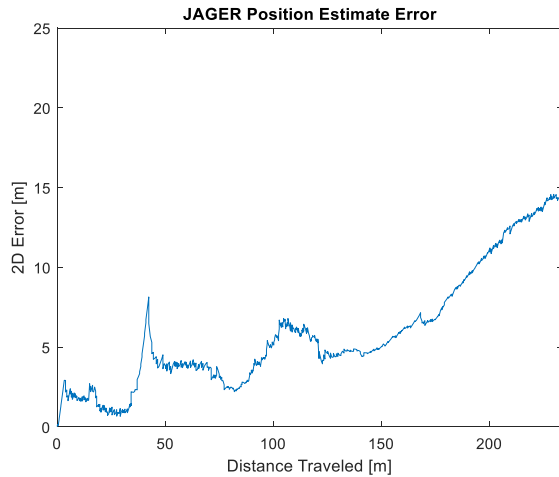


Figure 18: position error of the estimate of JAGER's position as a function of distance traveled in mission 1

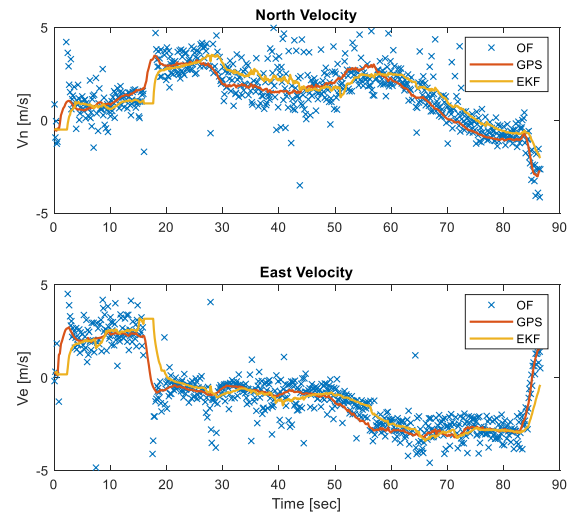


Figure 19: measured optical flow velocity plotted with the GPS (red) and estimated (yellow) velocity throughout the duration of the flight

Mission 2

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

As flown, before including the confidence information, the flight progressed as depicted in Figure 20. In Figure 20a, the initialize estimate of the filter is in the proper direction towards the source. As JAGER flies through the region of interference that caused distorted patterns aimed towards the hangar door (e.g. the pattern in Figure 9), the estimate begins to converge to the hangar door, depicted in Figure 20b, instead of the source itself. Continuing through the region of interference that results in antenna patterns with low confidence (no discernable lobes or only cluttered bearing measurements), the estimate can be seen in Figure 20c to continue to converge towards the hangar door. Finally, by the time JAGER reached a region of less information, the estimate slowly started to pull away from the hangar door, but the final result, shown in Figure 20d, does not capture the jammer within its error ellipse.

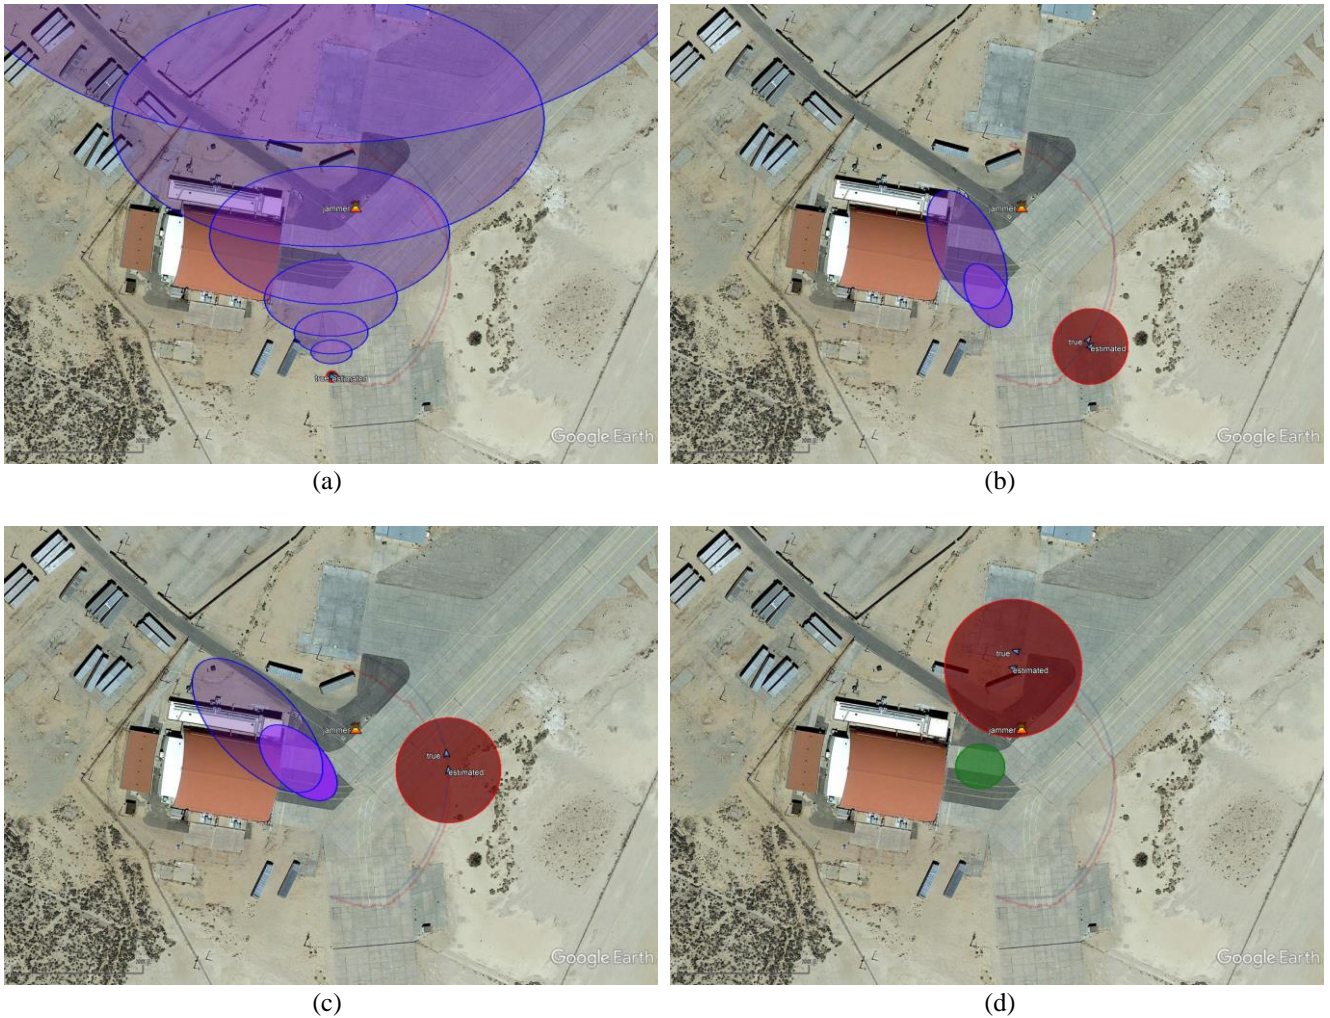


Figure 20: snapshots of mission 2 as flown without use of confidence metric on RSS pattern measurements with the RFI estimate in the violet ellipses, the vehicle estimate as the red ellipse and line, and the true path in blue

After experiencing these effects from the hangar door, the localization estimator was adjusted to exclude “low confidence” RSS patterns, such as the pattern in Figure 10, where there is no discernable lobe. Rerunning the exact same mission with the exact same measurements (in post processing), resulted in a better estimate throughout the flight, depicted in Figure 21. Due to measurements being excluded early in the flight due to interference, the covariance of the estimate remains large, visually seen in Figure 21a and Figure 21b, until JAGER reaches the region of less interference. As a result, the final estimate, shown in Figure 21c, converges more accurately on the true location of the signal source.

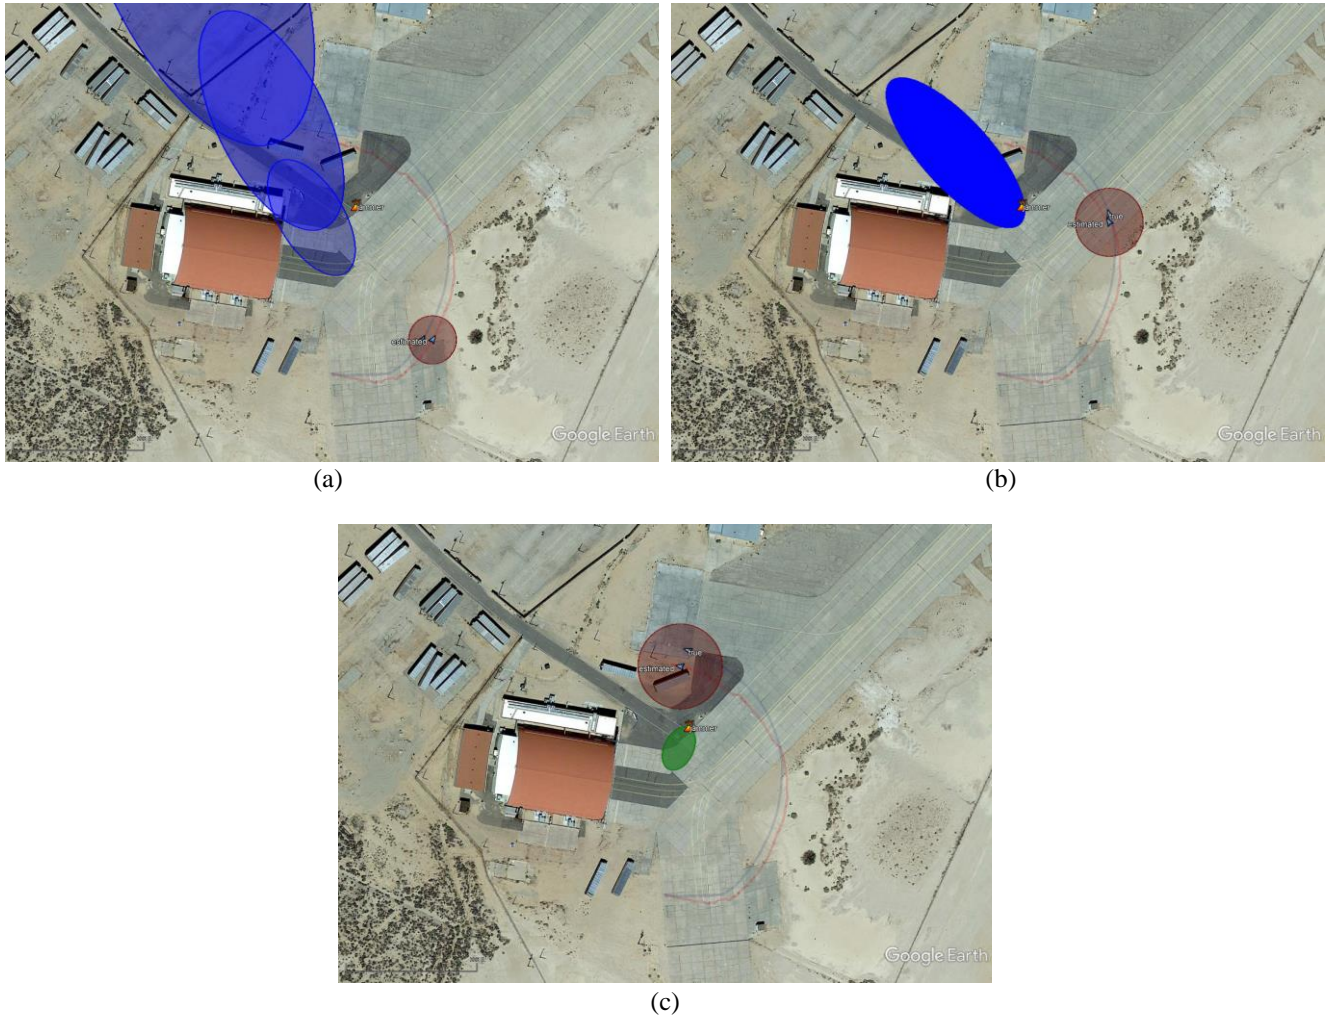


Figure 21: snapshots of mission 2 with the use of the confidence metric to remove corrupt RSS patterns with the source estimate in blue and green, the true flight path in blue, and the estimated flight path in red.

Comparing the covariance and the 2D error of the estimate for both of these runs, in Figure 23 for the original mission flight, and Figure 22 for the rerun removing “low confidence” measurements, it can be seen that the covariance stays a lot larger for the rerun than it does for the original flight, resulting in the estimate to converge more accurately to the source.

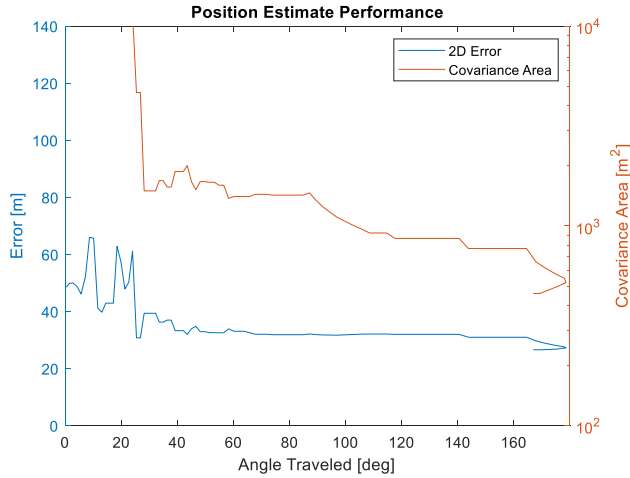


Figure 23: performance of source localization for mission 2 without use of confidence metric

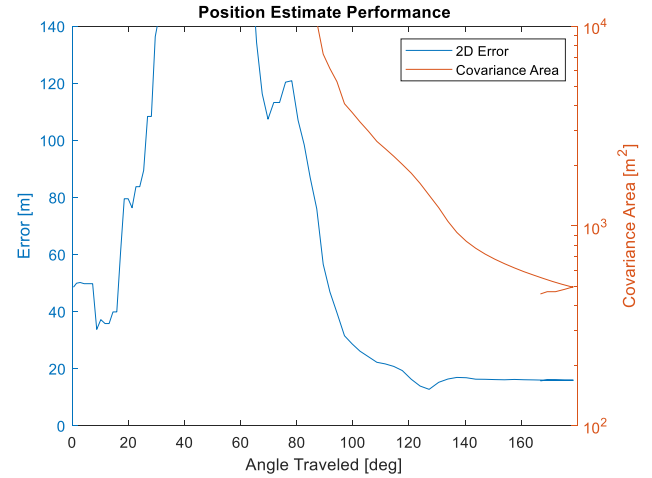


Figure 22: performance of source localization for mission 2 with the use of the confidence metric to exclude measurements

Wrong Initialization

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

Figure 25 and Figure 24 show some example flight paths of wrong initializations. These two examples illustrate the two main types of wrong initialization. Figure 25 demonstrates a wrong initialization that converges almost immediately to a location very close to JAGER. This type of failure was commonly seen when starting in one of the regions of higher levels of interference from the hangar door. Figure 24 demonstrates a wrong initialization that has the estimate from the back lobe as the highest likelihood estimate for the signal source. This type of failure also occurred when starting in location A, with no interference, as it is a result of the cluttered measurements from BeaSt. Both of these examples illustrate the sensitive nature of the initialization scheme implemented for this test with the cluttered measurements that result from using the 3-element phased array antenna. A more robust initialization scheme and other solutions should be possible (e.g. 7 element phased array

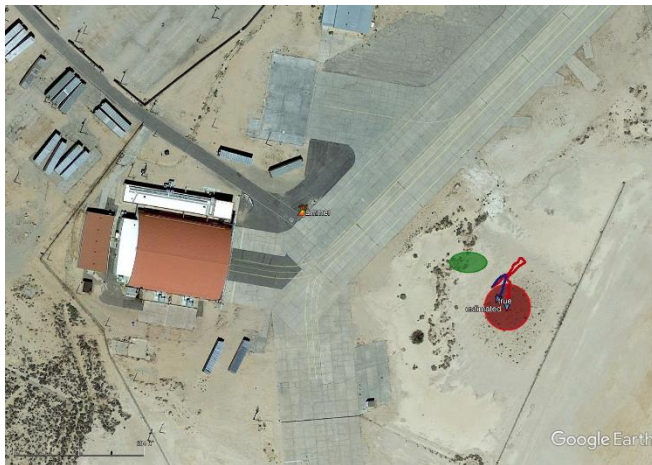


Figure 25: example failed initialization to a high confidence location very close to the starting location

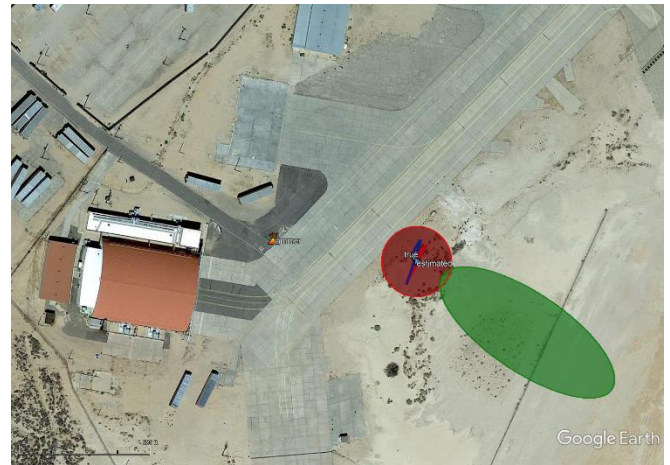


Figure 24: exampled failed initialization on an estimate resulting from a side lobe

antenna).

Both of the failed initialization examples started from location B, however, as seen in Table 1 overview of the performance by starting location, location C was the most challenging starting location. In starting location C, many of the missions failed due to an almost immediate convergence to a location close to the takeoff location (the failure depicted in Figure 25) due to the strong interference from the hangar door. In starting location C, most of the measurements through the initialization phase of the flight were a combination of distorted patterns (see Figure 9) and corrupted patterns (see Figure 10) that caused the initialization to fail.

Flight Test Overview

Table 2 and Table 3 look at the localization and navigation results, respectively, more broadly for all the missions flown. The success of the source localization was not greatly affected by the time of day (i.e. the performance of the navigation algorithm), so the overview results shown in Table 2 are for all the combined missions. While Table 1 shows that the ability to localize the source was much more strongly correlated with the starting location and the quality of the initial set of RSS patterns (and therefore bearing observations), further highlighting the importance of the initialization phase of the estimator, Table 2 does not break down the results by starting location, but rather only includes the successful mission, because, when a mission was successful, there was not a significant difference in the mission statistics (i.e. how the mission played out).

Table 2: overview of the source localization performance for all successful missions

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

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

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

Table 3: overview of the performance of the navigation system for all missions

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

CONCLUSION

A real-time implementation of a fully autonomous, drone based, signal localization system (JAGER) with a GNSS independent navigation system was demonstrated to work in a GNSS unavailable environment at Edwards Air Force Base. The environment chosen was able to both demonstrate the performance of the system in both an open environment (location A) and in more strenuous urban-like environments (locations B and C). As depicted in Figure 15 and Table 1, not every mission resulted in a successful localization due to a combination of the challenges posed by the urban environment of part of the flight area and the cluttered measurements that result from BeaSt being only a 3-element array. These missions with failed initializations provide very important lessons learned on the operating considerations that need to be made with a system like this due to the environment in which it is being deployed. Testing the system in a real-world and real-time scenario, where the path planning is based on the current estimate, shows how having only 3-elements in the array, while feasible, has some drawbacks due to the cluttered measurements. Increasing the array size to a 7-element array would greatly reduce the amount of clutter and may be capable of reducing the number of wrong initializations, resulting in more successful missions.

On the navigation front, the IR camera based system demonstrated to be a feasible system for localization of sources during the day, at night, and in the early morning. The drift rate of the system was low enough to work for these localization missions where the travel time to circle around the source enough to result in a good estimate of the source was small due to the short distance between the starting location and the source location, however it may not be low enough for longer duration missions that would result from signal sources that are further away. Using a constant velocity assumption for the two stage path planning approach adds a large amount of error when transitioning from a straight line to the circle portion of the flight and should be addressed in future development to better account for the expected motion of the vehicle throughout the flight.

Finally, in flying the system in a real world environment enabled a better understanding of the considerations required when operating the JAGER system. With the initialization scheme being so critical to the success of the mission, as configured, it is important to account for the environment around JAGER at the beginning of the flight as it has a much greater impact on the overall performance than any other portion of the flight. Overall, these flight tests demonstrated JAGER's capability of using a 3-element phased array antenna and an IR camera based optical flow system to fully autonomously, in real-time, localize a signal source without the use of GNSS onboard.

ACKNOWLEDGEMENTS

The authors would like to thank Kazuma Gunning for assisting in the flight test campaign. Flying at Edwards Air Force Based enabled the ability to demonstrate JAGER's capabilities at night and in a unique environment, and for that the authors thank Edwards Air Force Base, the 412th Test Wing, and special thanks to Lt. Colonel David Aparicio, the Emerging Technologies Combined Test Force (ETCTF) Commander and Mr. David Freeman (ETCTF) Program Manager for the use of their airspace, their workspace and their time.

The views expressed herein are those of the authors only and are not to be construed as official or any other person or organization.

This paper is Approved for Public Release through the 412th Test Wing, Edwards Air Force Base
Distribution A: 412TW-PA-19478

REFERENCES

- [1] A. Perkins, L. Dressel, S. Lo, T. Reid, K. Gunning, and P. Enge, "Demonstration of UAV-Based GPS jammer localization during a live interference exercise," in *29th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2016*, 2016, vol. 5.
- [2] A. Perkins, Y.-H. Chen, S. Lo, and P. Enge, "Vision Based UAS Navigation for RFI Localization," in *31st International Technical Meeting of The Satellite Division of the Institute of Navigation, ION GNSS+ 2018*, 2018, pp. 2726–2736.
- [3] A. Perkins, Y. Chen, and S. Lo, "Unmanned Aerial System for RFI Localization and GPS Denied Navigation," in *International Flight Inspection Symposium (IFIS)*, 2018.
- [4] A. Perkins, S. Lo, and J. D. Powell, "Measurement Characterization for Localizing Multiple RFI Sources Simultaneously from a UAS," *Proc. ION 2019 Pacific PNT Meet.*, pp. 262–273, 2019.
- [5] A. Perkins, Y.-H. Chen, W. Lee, S. Lo, and P. Enge, "Development of a three-element beam steering antenna for bearing determination onboard a UAV capable of GNSS RFI localization," in *30th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2017*, 2017, vol. 4.
- [6] G. Bradski, "The OpenCV Library," *Dr. Dobb's J. Softw. Tools*, 2000.
- [7] J. Shi and others, "Good features to track," in *1994 Proceedings of IEEE conference on computer vision and pattern recognition*, 1994, pp. 593–600.
- [8] D. Mušicki, "Bearings only single-sensor target tracking using Gaussian mixtures," *Automatica*, vol. 45, no. 9, pp. 2088–2092, Sep. 2009.
- [9] T. Lemaire, S. Lacroix, and J. Sola, "A practical 3D bearing-only SLAM algorithm," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 2449–2454.
- [10] R. H. Mitch *et al.*, "Know your enemy: Signal characteristics of civil GPS jammers," *GPS World*, vol. 23, no. 1, pp. 64–72, 2012.