Safe2Ditch Steer-to-Clear Development and Flight Testing

Bryan J. Petty
Analytical Mechanics Associates
Hampton, Virginia, United States of America
bryan.j.petty@nasa.gov

Patricia C. Glaab
Aeronautical Systems Analysis Branch
NASA Langley Research Center
Hampton, Virginia, United States of America,
patricia.c.glaab@nasa.gov

Louis J. Glaab
Aeronautics Systems Engineering Branch
NASA Langley Research Center
Hampton, Virginia, United States of America
louis.j.glaab@nasa.gov

***Abstract*— This paper describes a series of small unmanned aerial system (sUAS) flights performed at NASA Langley Research Center in April and May of 2019 to test a newly added Steer-to-Clear feature for the Safe2Ditch (S2D) prototype system. S2D is an autonomous crash management system for sUAS. Its function is to detect the onset of an emergency for an autonomous vehicle, and to enable that vehicle in distress to execute safe landings to avoid injuring people on the ground or damaging property. Flight tests were conducted at the City Environment Range for Testing Autonomous Integrated Navigation (CERTAIN) range at NASA Langley.**

Prior testing of S2D focused on rerouting to an alternate ditch site when an occupant was detected in the primary ditch site. For Steer-to-Clear testing, S2D was limited to a single ditch site option to force engagement of the Steer-to-Clear mode. The implementation of Steer-to-Clear for the flight prototype used a simple method to divide the target ditch site into four quadrants. An RC car was driven in circles in one quadrant to simulate an occupant in that ditch site. A simple implementation of Steer-to-Clear was programmed to land in the opposite quadrant to maximize distance to the occupant's quadrant. A successful mission was tallied when this occurred.

Out of nineteen flights, thirteen resulted in successful missions. Data logs from the flight vehicle and the RC car indicated that unsuccessful missions were due to geolocation error between the actual location of the RC car and the derived location of it by the Vision Assisted Landing component of S2D on the flight vehicle. Video data indicated that while the Vision Assisted Landing component reliably identified the location of the ditch site occupant in the image frame, the conversion of the occupant's location to earth coordinates was sometimes adversely impacted by errors in sensor data needed to perform the transformation. Logged sensor data was analyzed to attempt to identify the primary error sources and their impact on the geolocation accuracy.

Three trends were observed in the data evaluation phase. In one trend, errors in geolocation were relatively large at the flight vehicle's cruise altitude, but reduced as the vehicle descended. This was the expected behavior and was attributed to sensor errors of the inertial measurement unit (IMU). The second trend showed distinct sinusoidal error for the entire descent that did not always reduce with altitude. The third trend showed high scatter in the data, which did not correlate well with altitude. Possible sources of observed error and compensation techniques are discussed.

I. INTRODUCTION

Safe2Ditch (S2D) is an autonomous crash management system for sUAS. Its function is to detect the onset of an emergency for an autonomous vehicle, and to enable that vehicle in distress to execute safe landings to avoid injuring people on the ground, damaging property, and lastly preserving the vehicle and payload. S2D was designed with sUAS as a primary target host as a most constraining case for weight and size. However, the system can provide autonomous crash management for any air vehicle with proper integration of an autopilot, avionics, and vehicle diagnostic systems.

S2D concept testing began in 2015 with primary avionics integration. NASA Langley Research Center performed a series of sUAS flights to test the integration of sub-elements of the S2D system with a representative sUAS multi-rotor. Reference [1] describes the interfacing of software in Safe2Ditch, previous testing procedures, hardware used for the testing of the system, and results of the testing procedures from the 2015 concept tests.

Remote vision capability was integrated in 2016 with the help of Brigham Young University (BYU). Reference [2] describes Safe2Ditch communication using Robot Operating System (ROS) message passing, detection of moving ground objects, and algorithm details for identification and geolocation of moving targets.

Prior flight-tests of the S2D prototype tested rerouting to an alternate ditch site if a moving object occupied the primary site. In 2019, S2D development and testing focused on Steer-to-Clear functional behavior, which activates when Safe2Ditch detects occupants in the selected site, but has no alternative ditch site within range. During Steer-to-Clear control, the vehicle continually monitors the activity in the ditch site as it descends. When it reaches a final decision height, also called the commit altitude, the vehicle lands in the spot determined to provide the best clearance of occupants. Simulation testing of the Steer-to-Clear mode was leveraged to create a behavior that avoided site inhabitants while remaining within the flight maneuvering capabilities of the vehicle. Flight-testing of the new Steer-to-Clear mode took place in the first half of 2019.

II. SAFE2DITCH SYSTEM OVERVIEW

The Safe2Ditch system has five main software components, including the Health Monitor, the Ditch Site Selector, the

Navigation and Route (Nav/Route) Optimizer, Vision Assisted Landing, and a central Intelligent Hub, shown in Fig. 1.

A. Health Monitor

The Health Monitor continually examines components of the vehicle, such as battery and motor status, to determine if there is a need for an emergency landing. If the Health Monitor believes the vehicle can no longer complete its initial mission, it signals the other components of Safe2Ditch to trigger the emergency landing protocol. The Health Monitor uses information collected from the vehicle to estimate the time remaining until the vehicle will lose control and crash, also known as "time-to-live". If this calculated time is close to or less than the current time remaining in the mission, the vehicle will execute an emergency landing.

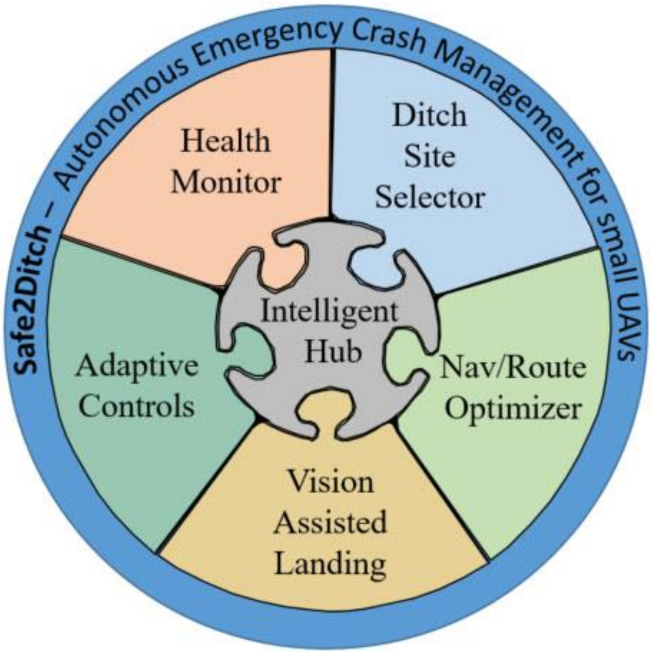

Fig. 1. Safe2Ditch Components

Alternatively, the system is designed to automatically engage in the event that remaining battery percentage received from the battery is less than a user-defined percentage. For range safety compliance, a safety pilot or the operator of the ground control station can also engage an emergency landing, at their discretion. Once the emergency landing procedure starts, the Intelligent Hub informs the other components. The Intelligent Hub provides the time-to-live to the Ditch Site Selector and performance limitations to the Nav/Route Optimizer.

B. Ditch Site Selector

The Ditch Site Selector uses time-to-live from the Health Monitor and information about possible ditch sites, which are loaded in from a file. The data in the ditch site file includes latitude, longitude, altitude, radius, reliability of the site being clear, and a unique identity for the landing zone. For each software update cycle, the Ditch Site Selector scores the desirability of each ditch site based on site proximity, size, and reliability. The Ditch Site Selector uses the latitude, longitude and altitude of the ditch site to calculate the distance to the site at any given time. Once the distance is calculated, the Ditch Site Selector computes the estimated flight time using the cruise speed of the vehicle. If this estimated flight time is greater than the time-to-live, the Ditch Site Selector considers it out of range and eliminates that ditch site as a potential solution.

After all sites within range are determined, the Ditch Site Selector computes an overall desirability with a higher score for ditch sites with larger radii. The weight factor of the radius criteria is configurable by the user. The algorithm degrades the desirability of a site when the Vision Assisted Landing component detects occupants within the ditch site. Once the Ditch Site Selector calculates the desirability of all sites, it chooses the one with the highest desirability and informs the other components. The Ditch Site Selector continues to monitor as the vehicle approaches and descends, and updates the target ditch site choice if the desirability ranking changes.

C. Nav/Route Optimizer

The Nav/Route Optimizer receives the location of the selected site from the Ditch Site Selector and creates the flight plan to that site. If available, the Nav/Route Optimizer considers constraints such as current vehicle performance, risk assessment of surrounding areas and no-fly areas when developing a route. Without constraints, it will use a straight path to the top of descent. The Nav/Route Optimizer calculates the top of descent to allow the descent angle to match the camera mount angle to center the ditch site in the camera's field of view for the Vision Assisted Landing. Once at the top of descent, the Nav/Route Optimizer begins to perform Steer-to-Clear calculations. Steer-to-Clear assists in avoiding occupants in the ditch site by adjusting the touchdown point within the selected site. The current version of the Steer-to-Clear algorithm in the prototype system separates the ditch site into quadrants along the north/south and east/west lines of the site. The algorithm then assigns all moving occupants to quadrants based on their latitude, longitude and their calculated bearing to the latitude and longitude of the selected ditch site. After the algorithm assigns all occupants to quadrants, it selects the quadrant farthest from the occupants as the safest target. If no quadrants have occupants, the vehicle lands in the middle of the site. The algorithm continues to execute until the vehicle is below a user defined commit altitude, or until the Ditch Site Selector determines a better site.

D. Vision Assisted Landing

The vision system is composed of two parts, the visual front-end and the Recursive-Random Sample Consensus (Recursive-RANSAC) analysis software [3]. The system receives video coming from the mounted camera and compares features from each image at a set time step. The image differences in each time step are used to generate a perspective transformation through homography. Homography is an image analysis technique that resolves details within an image using vectors, which allows recognition of common features in images captured from different perspectives. The homography and features combine to calculate the velocity magnitude of moving objects within a predefined threshold. Recursive-RANSAC uses the velocity magnitudes of the moving objects to calculate the position, velocity, acceleration and jerk of detected moving ground objects [4]. The Vision Assisted Landing performs the transformation from the camera field of view to the north-east-down axis using sensor inputs from the autopilot Global Positioning System (GPS) position and attitude. The axial transformation requires the angle of the camera relative to the copter and the ground. The Nav/Route Optimizer also uses this angle to calculate the top of descent to descend at an angle matching that of the camera to allow the camera to have full view of the ditch site. Once the axial transformation is complete, the Visual Assisted Landing sends the GPS coordinates of the moving ground objects to the other components of Safe2Ditch.

E. Intelligent Hub

The Intelligent Hub provides the configuration, communication, and timing services needed between components. It also creates the subsystems and interfaces needed to communicate with external parts of the system.

F. Steer-to-Clear Objectives

The governing algorithms to create Steer-to-Clear functionality were tested in simulation to accomplish the primary objective of occupant avoidance. A secondary objective was to minimize stress to the vehicle and its performance. An early version of the algorithm commanded continual steering by the autopilot as the simulated copter dynamically reacted to the moving occupant in the ditch site. As the object moved through the quadrants on the ground, the copter continually moved in response as it descended to the commit altitude. For the flight test algorithm, the dynamic maneuvering before the commit altitude was eliminated because the maneuvering provided no benefit and had the potential to overload the communication network. The revised algorithm saves battery life and communication bandwidth with dynamic *monitoring* (i.e., no maneuvering) prior to the commit altitude, with steering and maneuvering to the final selected quadrant once the vehicle reaches the commit altitude.

In simulation with perfect geolocation, smooth flight of the vehicle without wind or gust disruption, little to no signal latency, and perfect compliance of the camera mounting angle to the design angle, the copter landed in the predicted best quadrant in each test. These simulation cases were very useful for debugging and refining the algorithms. In flight tests with real hardware and weather, however, imperfect signals or resolution constraints degraded the system's performance and success rate. Understanding the impact of signal errors on the overall system performance to inform system requirements for real-world use were a main goal of this round of testing of the Steer-to-Clear functionality.

A remote-controlled (RC) car was driven in the ditch site to create ground motion to simulate an occupant in the site. For each test run, the RC car motion was isolated to a single quadrant (varied from flight to flight) to simplify prediction of the targeted landing quadrant by the copter. When the copter targeted and landed in the quadrant expected to be rated as the best location, this was tallied as a successful mission. Successful mission flights required the Ditch Site Selector to select the correct site, the Vision Assisted Landing to identify and communicate the location of the RC car, and the Nav/Route Optimizer to provide the correct touchdown location for landing. To correctly geolocate the RC car to the quadrant in which it actually drove, the Vision Assisted Landing had to locate the car as an occupant in the ditch site, transform the location of the car in the image frame to the earth frame of reference, and communicate that location to the rest of the system. The transformation of the occupant's location from the image frame to the earth frame required the flight vehicle's GPS position and attitude, and required accurate mounting of the camera to the design mount angle (which shook slightly in response to abrupt vehicle motion). The RC car carried its own GPS receiver and logged its location data. In post-processing, the GPS location measured and logged by the car was compared to the derived latitude and longitude location of the occupant by the Vision Assisted Landing component using the imaged ditch site.

III. Test Setup

A. Test Site

Flight-testing occurred at NASA Langley's City Environment Range for Testing Autonomous Integrated Navigation (CERTAIN) [5]. This test site allows within-line-of-sight operation of unmanned vehicles, up to 120 meters (400 feet) above ground level. Required weather conditions for operations require 4.8 kilometers (3 miles) of visibility, 304.8 meters (1000 feet) cloud ceiling, and winds less than 32.1 kilometers per hour (20 miles per hour).

B. Test Hardware

The Eagle XF8, shown in Fig. 2 and the ACV-2 S-1000, shown in Fig. 3, were the test vehicles for Steer-to-Clear. A Pixhawk Cube 2 autopilot running APM v3.5.4 and 3DR uBlox GPS were used on the copter. An Imaging Development Systems UI-1250ML camera send images to the Vision Assisted Landing component. This camera has a USB 2.0 interface. During the tests, the camera ran at an image resolution of 800x600 at 30 frames per second. The flight computer running Safe2Ditch software was an NVIDIA Jetson TX2 with an ARMv8 64-bit CPU and 8GB of Low-Power Double Data Rate Synchronous Dynamic Random Access Memory. The Jetson TX2 interfaced with the autopilot through an Orbitty carrier board, which provides interfaces via serial ports. A Traxas Xmaxx RC car simulated motion in the ditch site to reduce the risk to test personnel. The RC car was equipped with the same GPS and Pixhawk used on the vehicle to perform assessment of the vehicles ability to geolocate the position of the ground occupants.

Fig. 2. ACV-2 S-1000

Fig. 3. Eagle XF8

C. Test Procedure

A geofence feature of the autopilot insures that operations do not violate test site boundaries. The geofence allowed for vehicle operations up to 120 meters (400 feet) altitude and up to 152.4 meters (500 feet) laterally from the home position. Before a test occurred, the test team performed a preflight safety checklist on the vehicle. A mission within the geofence was loaded into the vehicle's autopilot. The mission included an auto-takeoff, four guidance waypoints, and a loop back to the initial waypoint. The vehicle flew the mission until the remaining battery percentage dropped below the configured minimum (or until the safety pilot used a transmitter switch) to engage Safe2Ditch. Once Safe2Ditch engaged, the vehicle began flying to the top of descent and the RC car began driving in a quadrant within the ditch site. When an alternate ditch site is available within range, Safe2Ditch reroutes if the vision interface detects an occupant in the current landing zone. For testing of Steer-to-Clear, the team provided only one ditch site to Safe2Ditch to force landing in the occupied site. Fig. 4 shows the layout of the ditch site provided and Fig. 5 shows the general descent route.

Once the copter reached the top of descent, it began its descent to the ditch site. When the vehicle reached an altitude below the commit altitude, the vehicle choose where to land based on the most recent Steer-to-Clear calculations. The RC car exited the landing zone to avoid potential collision. The decision altitude of 20 meters was selected because it was determined to be the lowest altitude at which the camera can image the entire 9.144-meter (30-foot) ditch site, shown in Fig. 3. The copter initiated an autoland once it reached the calculated landing latitude and longitude.

The team retrieved the copter and unloaded the log files to a storage device. The RC car stored the log file of each run on its own storage device, which the team retrieved at the end of the testing. After securing the log files, the team reset and reconfigured the copter and the car for another test.

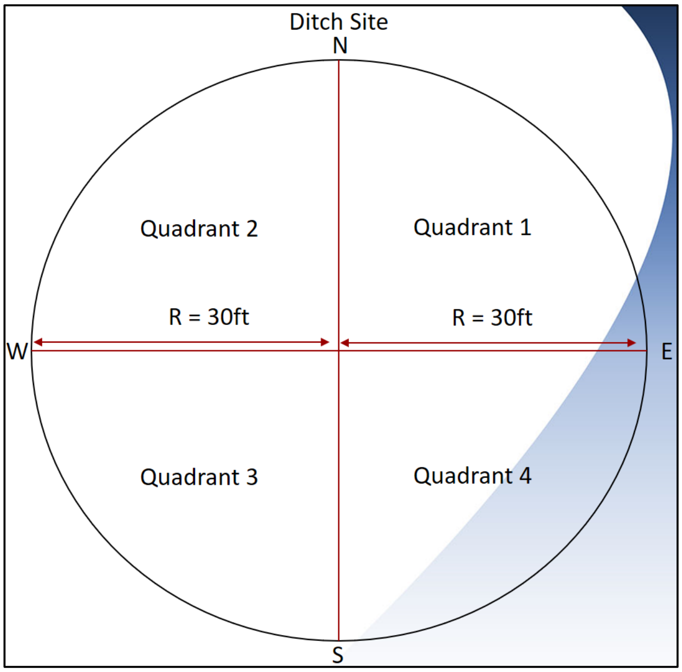

Fig. 4. Layout of Ditch Site used in test.

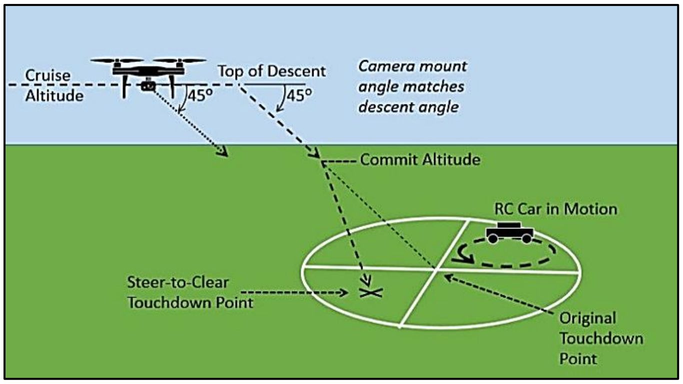

Fig. 5. Basic route to the ditch site, with touchdown spot adjustment by Steer-to-Clear.

D. Log Files Produced

Safe2Ditch produces two logs of interest to determine if the system is functioning properly – the ROS bag file and the Safe2Ditch event log file. The bag file contains messages passed between the components of Safe2Ditch during a run. Examples of the messages from the bag file used for analysis of Steer-to-Clear include video feed of the camera, latitude and longitude of the targets seen, and GPS location of the copter. The Safe2Ditch event log contains important actions that occur during a run of Safe2Ditch, such as the selected ditch site, the determination of the number of occupants within the ditch site, and the distance to the current Safe2Ditch guidance waypoint. The Safe2Ditch event log also informs the user of which quadrant Safe2Ditch chose and how many occupants the Vision Assisted Landing identified in each quadrant. This information is helpful when examining the performance of Steer-to-Clear. The RC car and the copter produce a telemetry log of the GPS module on board, which contain the GPS position and a GPS time stamp used to synchronize the copter and car data logs in post-processing.

IV. RESULTS

Testing for Steer-to-Clear occurred over multiple days in March and April of 2019. Overall, nineteen tests of Steer-to-Clear yielded usable data. Out of the nineteen tests, thirteen resulted in successful quadrant selection by Steer-to-Clear, which is approximately a 68% success rate. Flights in March used the ACV-2 S1000 and flights in April used the Eagle XF8. The Eagle hosted four out of the nineteen tests with three successful flights, yielding a 75% success rate. The S1000 hosted the other fifteen flights and had ten successful flights, yielding a success rate of 66%.

Before the flight test evaluation, the distance error between the RC car's actual location and its derived location computed by the Vision Assisted Landing component ("LaDE" for latitude and "LoDE" for longitude in the plot legends) was expected to be largest at cruise altitude. This distance error was expected to decrease as the vehicle descended. While this was generally the case, the amount of error sometimes varied more than expected between runs. Test results showed three general trends: convergence, high sinusoidal error, and high scatter error. Two examples of each trend were chosen and examined in this paper. Five graphs are presented for each of the chosen flights– the distance error between the RC car's logged location and the Vision Assisted Landing component's calculated location of the RC car, the actual quadrant of the RC car versus the Steer-to-Clear derived quadrant, and the roll, heading and pitch of the vehicle during descent. A negative one value in the quadrant graph indicates that the RC car was determined to be outside of the range of the ditch site and not within a quadrant.

The RC car drove in continuous circles within its assigned quadrant for each run. This resulted in repeating sinusoids for latitude and longitude location.

A. Expected Convergence Results

The test results shown in this subsection behaved as predicted for the location error between the derived and actual RC car locations converging with descent. The low distance error allowed the Steer-to-Clear function to determine the correct occupied quadrant.

1) Test Run Six on March 13th, 2019

Fig. 6 shows the distance between the RC car and the calculated GPS position of the occupant to be large at the beginning of descent, but stabilizing between 0 and 2 meters once the flight vehicle altitude reached 45 meters. The error in distance below 45 meters altitude is predominantly from GPS location error on the car and flight vehicle. Fig. 7 shows the quadrant selected by Steer-to-Clear on the descent phase. Larger location errors resulted in the wrong quadrant assessment above 47 meters, but as the distance error decreases, Steer-to-Clear selected and maintained the correct quadrant until the commit altitude of 20 meters.

Part of the distance error may be attributable to the orientation of the vehicle as it descended, since geolocation of the occupant is dependent on the camera having a 45 degree angle to the ditch site. Fig. 8, Fig. 9 and Fig. 10 show the roll, heading and pitch of the vehicle through descent, respectively. The roll of the vehicle ranged between -5 and 0 degrees, but generally remained between 0 and -4 degrees. The heading ranging only between -47 and -49 degrees and stayed relatively constant through the descent. The pitch of the vehicle ranged from 0 to 5 degrees below 30 meters, and between 1 and 4 degrees above 30 meters.

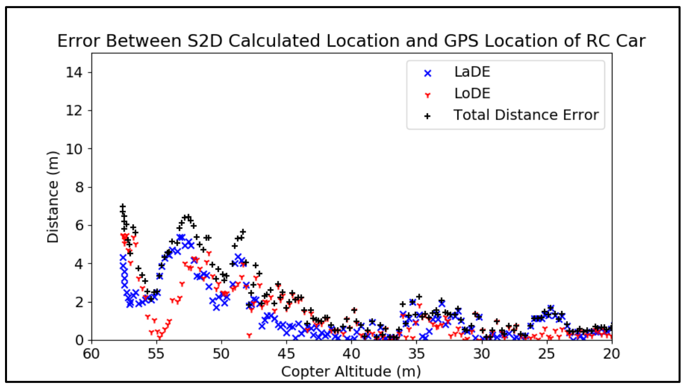

Fig. 6. Distance error between actual and derived RC car location.

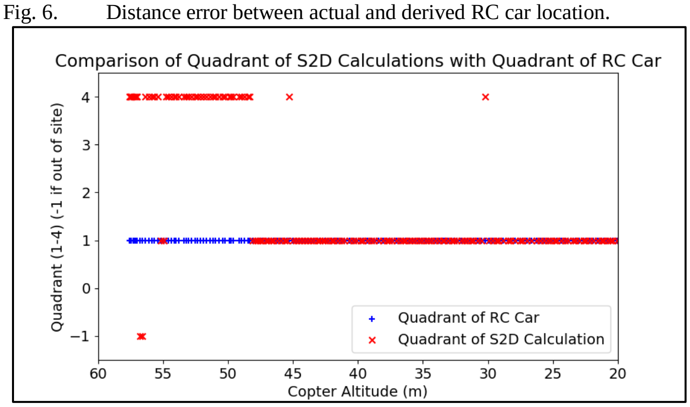

Fig. 7. S2D calculated quadrant versus actual quadrant of the RC car.

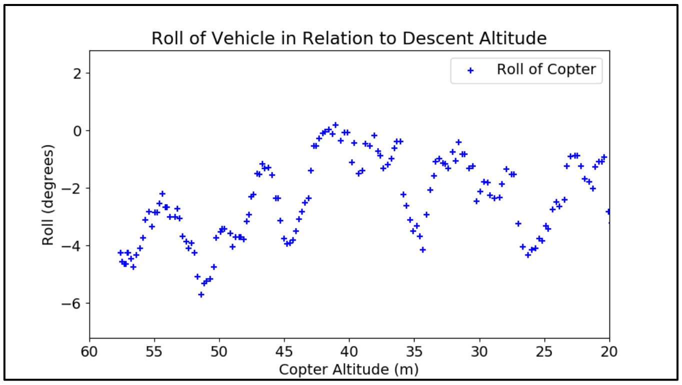

Fig. 8. Roll attitude of the vehicle as it descends to the ditch site.

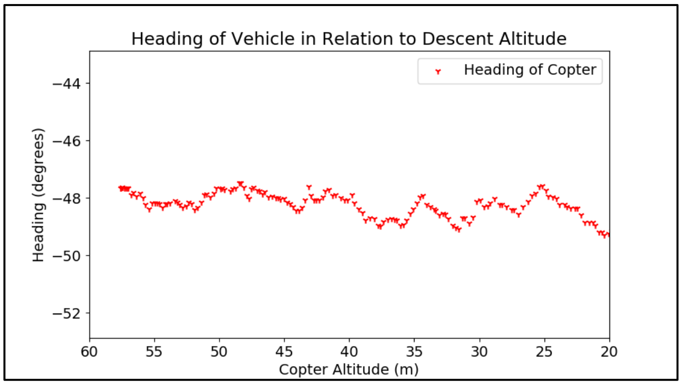

Fig. 9. Heading attitude of the vehicle as it descends to the ditch site.

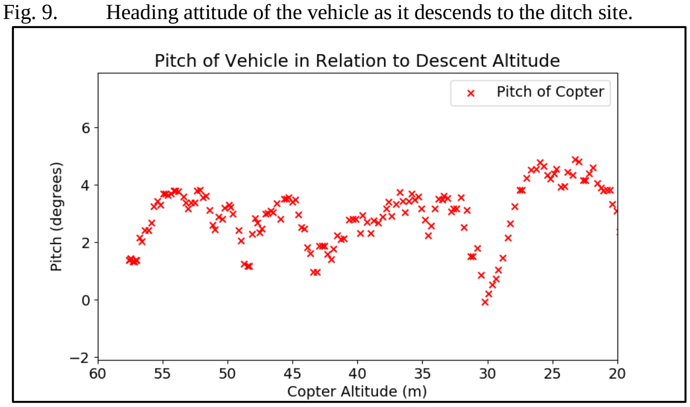

Fig. 10. Pitch attitude of the vehicle as it descends to the ditch site.

2) *Test Run Five on March 29th, 2019*

Like the previous test, this test had a large distance error at the start of descent, which decreased significantly later in the descent, as shown in Fig. 11. The distance error dropped to around 6 meters at just over 55 meters altitude and slowly decreased through the descent. Fig. 12 shows that Steer-to-Clear selected the correct quadrant when the distance error fell to 6 meters at just above 55 meters altitude. Fig. 13 shows the roll on the descent, which ranged from -2 to 3 degrees. Fig. 14 shows heading during this test. Fig. 15 shows the pitch of the vehicle during descent, which primarily ranged between 0 and 3 degrees.

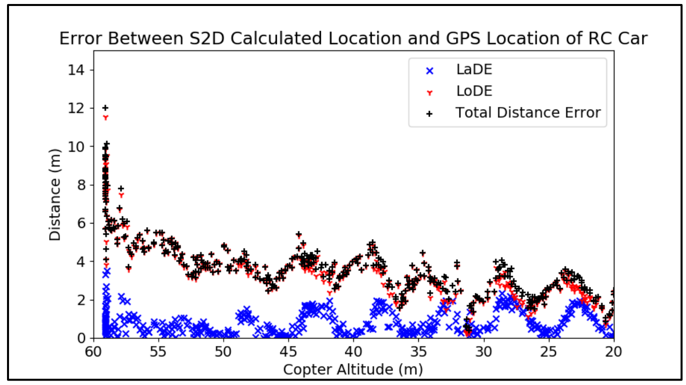

Fig. 11. Distance error between actual and derived RC car location.

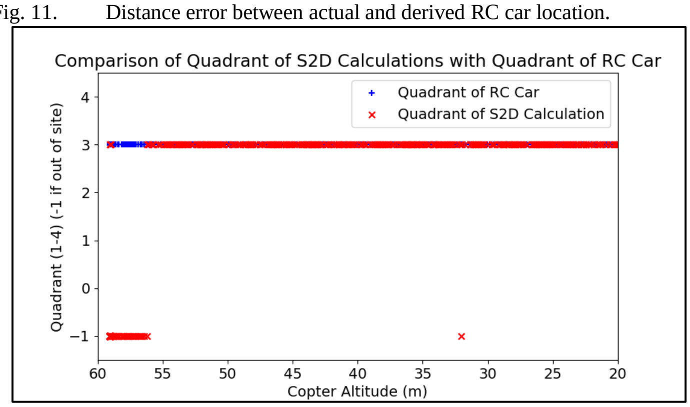

Fig. 12. S2D calculated quadrant versus actual quadrant of the RC car.

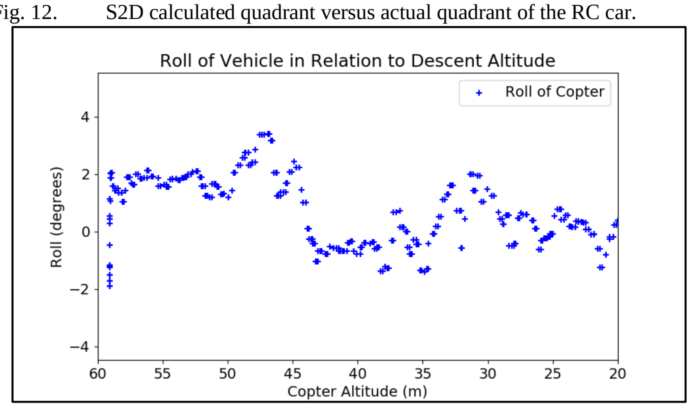

Fig. 13. Roll attitude of the vehicle as it descends to the ditch site.

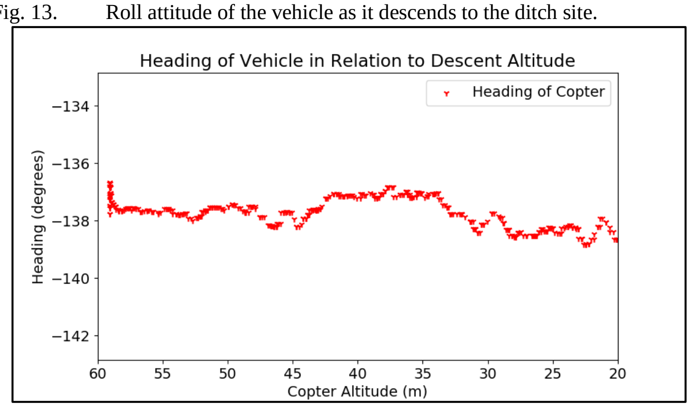

Fig. 14. Heading attitude of the vehicle as it descends to the ditch site.

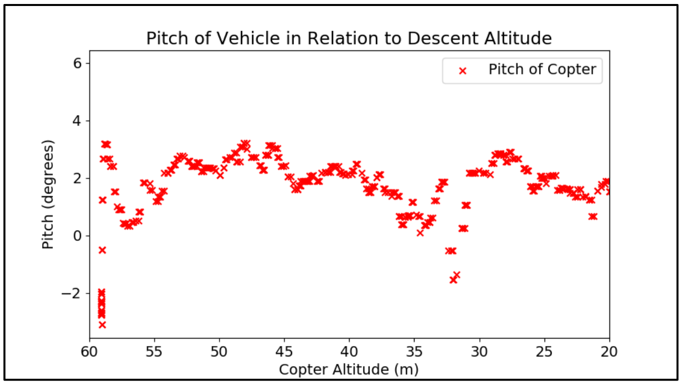

Fig. 15. Pitch attitude of the vehicle as it descends to the ditch site.

B. High Sinusoidal Error Results

High sinusoidal error results are those where the distance error maintained a large and distinctly sinusoidal pattern throughout the descent. One possibility for this pattern is that the error is due to latency between the derived and actual location of the car and is sinusoidal because the car was driven in a circle. Alternatively, the approach heading of the vehicle to the ditch site in conjunction with attitude error or latency in the attitude transmission might exacerbate a sensor error in either latitude or longitude. Because of this persistent error, the quadrant assigned to the occupant by Steer-to-Clear often did not match the quadrant that the RC car actually occupied.

1) Test Run Two on March 13th, 2019

Fig. 16 shows the distance error in the system started out high and decayed sinusoidally during descent. Fig. 17 shows the error in quadrant placement due to this large distance error and sinusoidal pattern. While this test produced larger error than that in the previous section, the error did decrease over the duration of the descent. Fig. 18 shows the roll of the vehicle during descent ranged from negative 2.5 to positive 1.5 degrees. Fig. 19 shows the heading of the vehicle during the descent ranged from -126 to -124. Fig. 20 shows the pitch of the vehicle ranged just between approximately 1 and 4 degrees. These attitude variations were smaller than the attitude variations in the previous section, though the total distance error was greater.

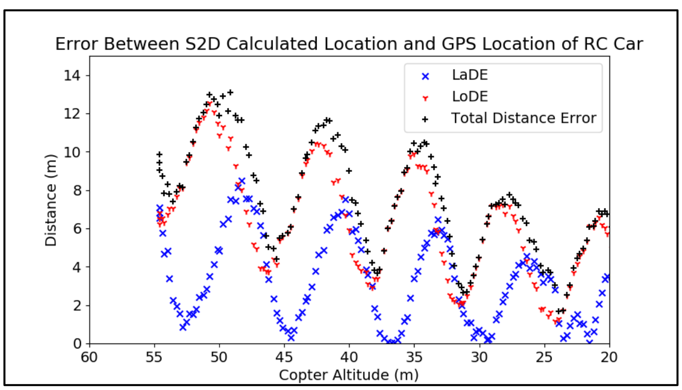

Fig. 16. Distance error between actual and derived RC car location.

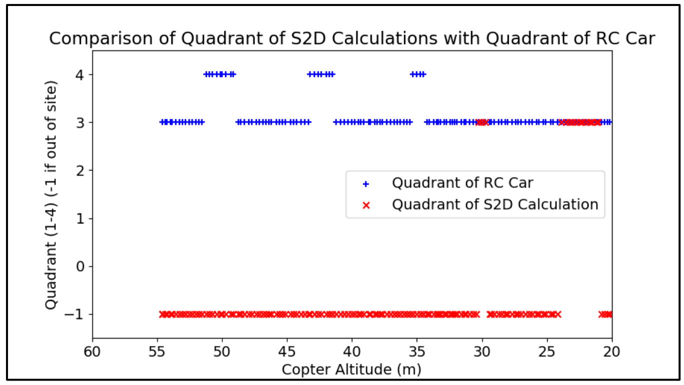

Fig. 17. S2D calculated quadrant versus actual quadrant of the RC car.

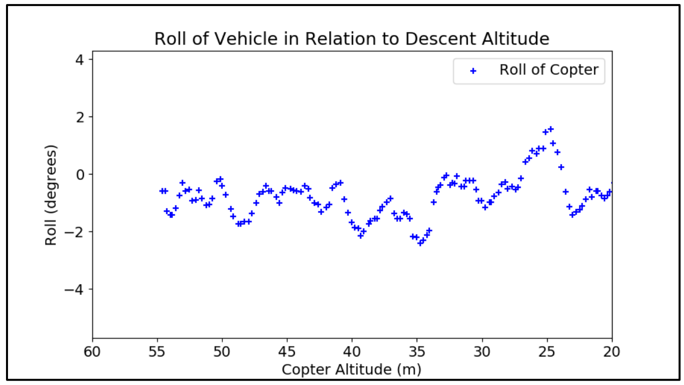

Fig. 18. Roll of the vehicle as it descends to the ditch site.

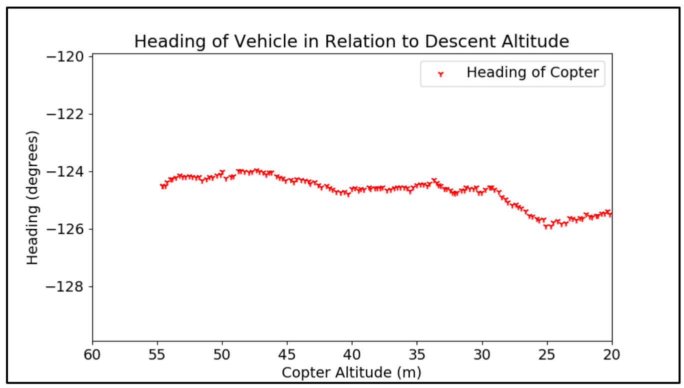

Fig. 19. Heading of the vehicle as it descends to the ditch site.

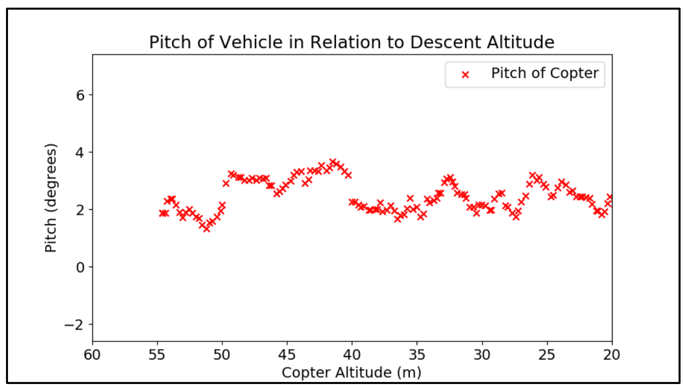

Fig. 20. Pitch attitude of the vehicle as it descends to the ditch site.

2) Test Run Seven on March 13th, 2019

Fig. 21 shows that the distance error in this test followed the same sinusoidal trend as the last test, but the error started lower and decreased more significantly with copter altitude. Fig. 22 demonstrates that a decreasing error as the copter descended allowed Steer-to-Clear to correctly locate the occupant at some points on its path. Fig. 23, Fig. 24, Fig 25 show the roll, heading, and pitch of the vehicle during descent, respectively.

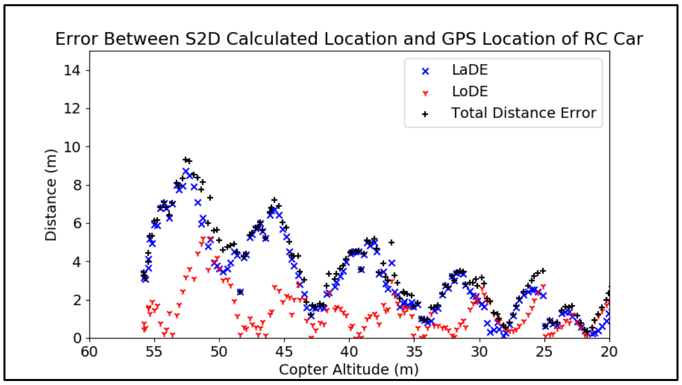

Fig. 21. Distance error between derived and actual RC car location.

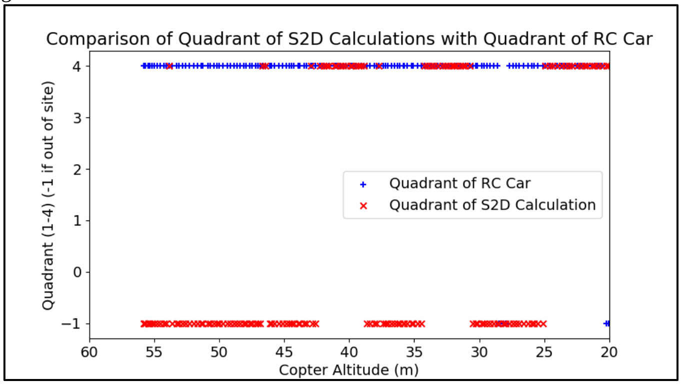

Fig. 22. S2D calculated quadrant versus actual quadrant of the RC car.

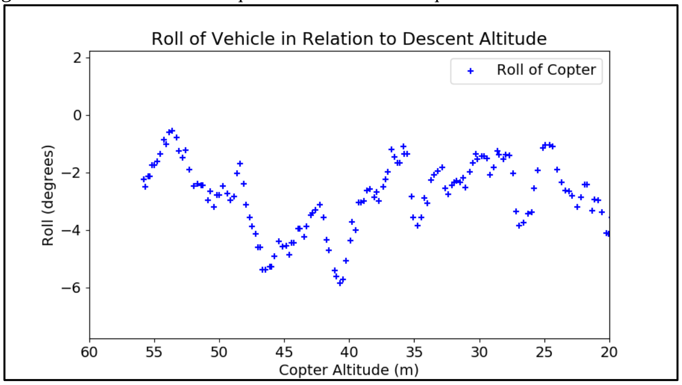

Fig. 23. Roll of the vehicle as it descends to the ditch site.

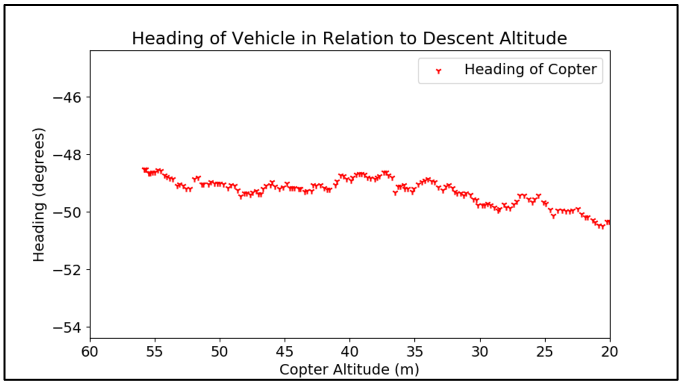

Fig. 24. Heading of the vehicle as it descends to the ditch site.

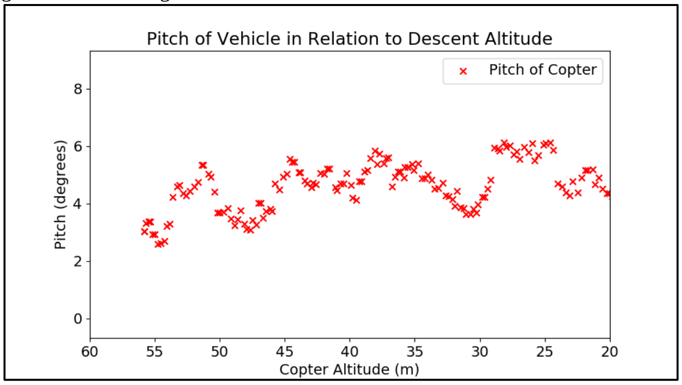

Fig. 25. Pitch attitude of the vehicle as it descends to the ditch site.

C. High Scatter Error Results

High scatter error results were distance error results that did not follow the sinusoidal patterns and were larger than expected. The two test cases presented show sporadic selection of quadrants by Steer-to-Clear during descent.

1) Test Run One on April 4th, 2019

Fig. 26 shows the distance error started out large, peaked at more than 14 meters, and decreased as the vehicle completed the descent. Fig. 27 shows that Steer-to-Clear did not correctly resolve to the correct quadrant until the distance error decreased. Fig. 28 shows the roll of the copter had abrupt changes compared to previous flights. A contributing factor to this could have been the transition in vehicles from the S-1000 to the Eagle. Fig. 29 shows heading trended similarly to the previous tests. Fig. 30 shows the pitch range for this test.

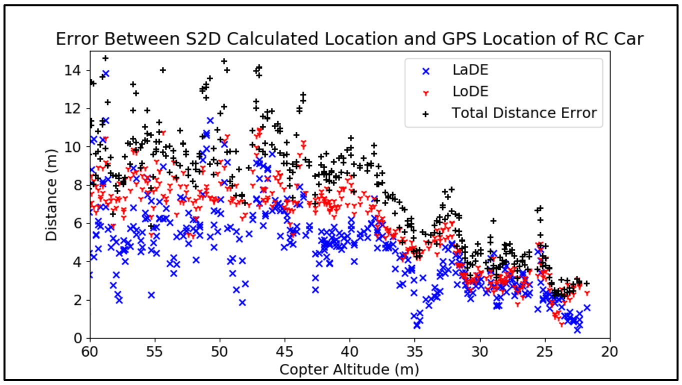

Fig. 26. Distance error between actual and derived RC car location.

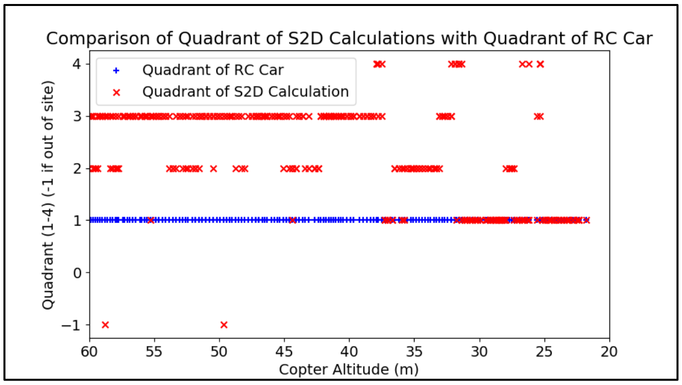

Fig. 27. S2D calculated quadrant versus actual quadrant of the RC car.

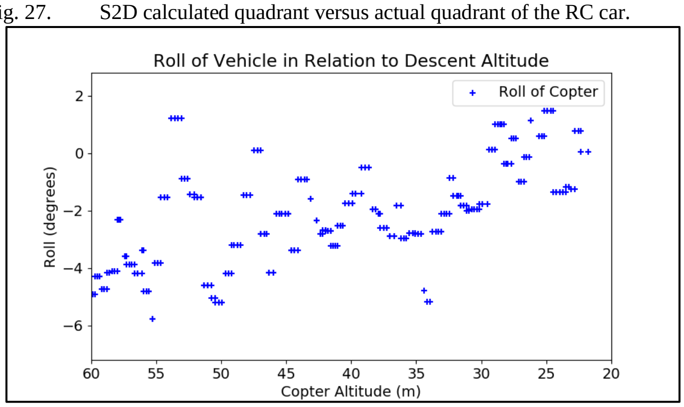

Fig. 28. Roll attitude as the vehicle descends to the ditch site.

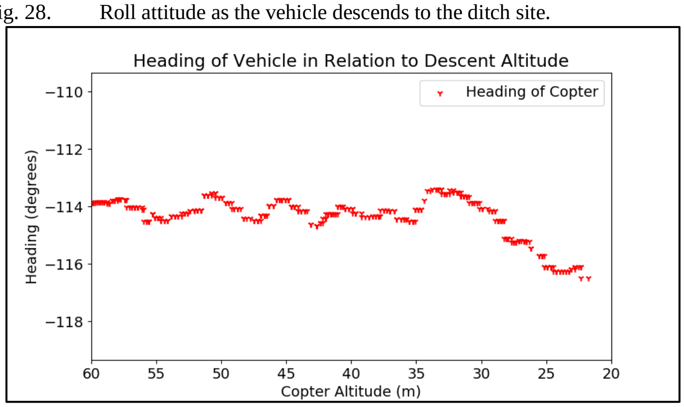

Fig. 29. Heading attitude as the vehicle descends to the ditch site.

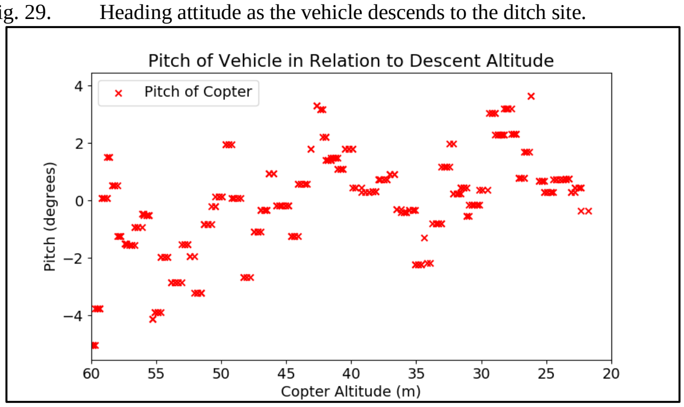

Fig. 30. Pitch attitude of the vehicle as it descends to the ditch site.

2) Test Run Five on April 4th, 2019

Fig. 31 shows high distance error that followed the error in longitude. Despite the large errors, the trend of distance error decreased as the descent continued. Fig. 32 shows that Steer-to-Clear occasionally picked the correct quadrant, but believed the occupant to be outside the ditch site most of the time due to the distance error. Fig. 33 shows a similar scatter in roll attitude as was in Test Run One on April 4th. Fig. 34 shows that the heading varied more than the heading from Test Run One on April 4th, 2019 but magnitude had less variance. Fig. 35 shows the pitch, which values range between -5 and 2 degrees.

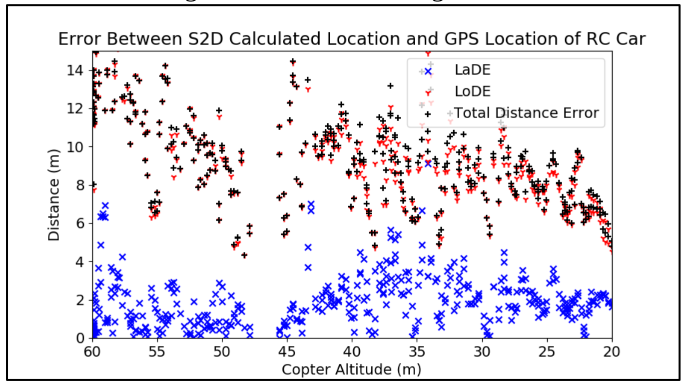

Fig. 31. Distance error between actual and derived RC car location.

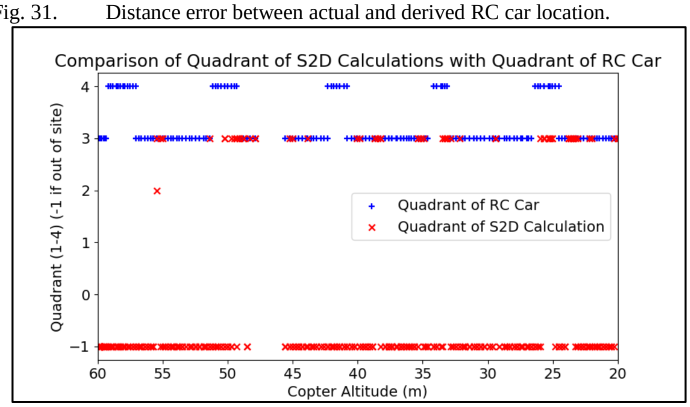

Fig. 32. S2D calculated quadrant versus actual quadrant of the RC car.

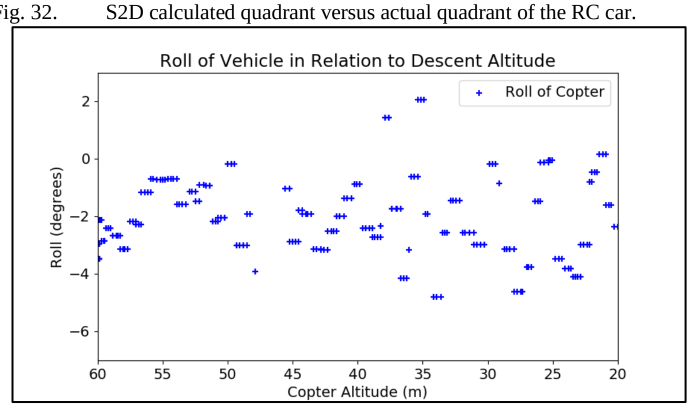

Fig. 33. Roll attitude as the vehicle descends to the ditch site.

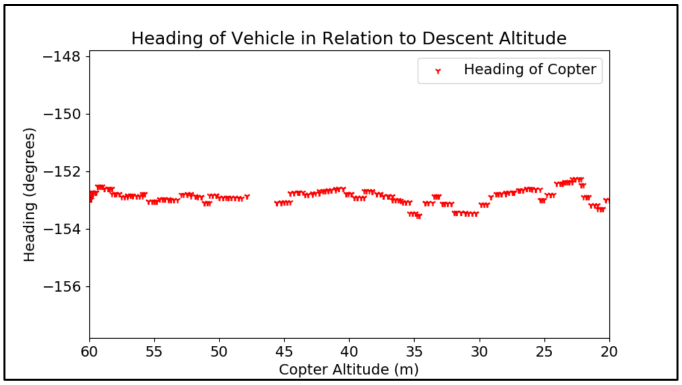

Fig. 34. Heading attitude as the vehicle descends to the ditch site.

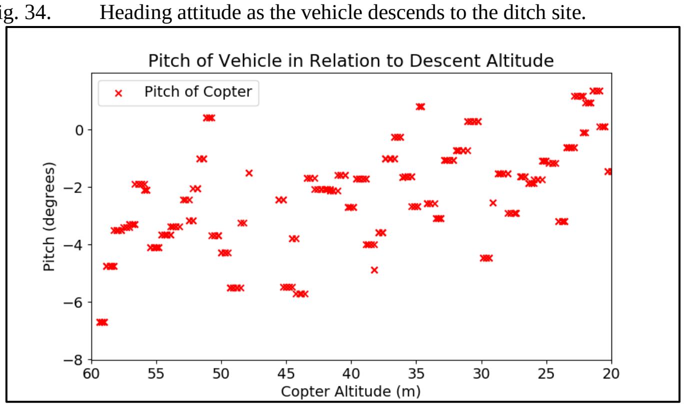

Fig. 35. Pitch attitude of the vehicle as it descends to the ditch site.

V. Discussion

The Steer-to-Clear function worked well for the majority of flights when the Vision Assisted Landing component was able to accurately geolocate the RC car in the ditch site. Accurate geolocation of the RC car is a two-step process within the vision system to first identify an object in the camera frame to be moving, and then to resolve that location in the image to latitude and longitude positions in earth coordinates using sensor inputs from the vehicle's GPS and IMU. The Recursive-RANSAC remote vision system accurately identified the RC car in the ditch site for every test flight. However, errors or latencies in the GPS or IMU signals contributing to the transformation from the image frame to the earth reference frame prevented accurate geolocation of the RC car for some flights. The specific source of the error is still under investigation, and may differ between flight tests. Roll, heading and pitch variation did not always correlate to geolocation error.

As a commercialized product, the geolocation accuracy could become a driver of the recommended minimum size of ditch sites for S2D. Given the 9.144 meter (30 feet) radius of the ditch sites in these tests, the geolocation errors should be less than 2 to 3 meters for the system to select the correct quadrant. Larger ditch sites could have the ability to compensate for geolocation inaccuracies, but might drive the commit altitude (when the vehicle commits to the touchdown spot) higher if the camera field of view limits full coverage of the ditch site. The worst location error observed for these tests at the commit altitude of 20 meters was approximately 8 meters of ground distance. This would be adequate resolution for a ditch site with a radius of approximately four times that error distance – around 32 meters (104.987 feet).

Alternatively, a more accurate IMU on the vehicle could reduce attitude-related error, and a faster flight computer processor could reduce errors associated with signal latency. Brigham Young University (creator of Recursive-RANSAC) is also working on methods to decrease geolocation error using GPS-denied navigation technology. Early experiments suggest a possible order of magnitude decrease in error, which would provide significant improvement for accurate geolocation at much higher altitudes to benefit a system like S2D.

Acknowledgment

The authors would like to thank the Unmanned Traffic Management (UTM) Sub-Project of the Air Traffic Management-eXploration (ATM-X) Project for funding of this work.

The authors would also like to thank Dan Healy and Scott Sims, who supported this effort as CERTAIN Range Safety Officers. They would also like to thank David Bradley and Robert McSwain for their support as Safety Pilots.

References

[1] Louis J Glaab, Patricia C. Glaab, Chester V. Dolph, Robert G. McSwain, Parker L. Lusk, Randall W. Beard, Bryan J. Petty "Integration flight testing of the Safe2Ditch autonomous crash management system for small unmanned aerial systems", NASA/TM-2018-220110, December 2018.

[2] Parker C. Lusk, Patricia C. Glaab, Louis J Glaab, and Randal W. Beard, "Safe2Ditch: emergency landing for small unmanned aircraft systems", Brigham Young University, June 2018.

[3] Niedfeldt, P. C.; Ingersoll, K.; Beard, R. W.; Comparison and Analysis of Recursive-RANSAC for Multiple Target Tracking. IEEE Transactions on Aerospace and Electronic Systems 52(1), 461-476 (2017). DOI 10.1109/TAES.2017.2650818.

[4] Lusk, P. C.; Beard, R. W.: Visual Multiple Target Tracking From a Descending Aerial Platform. In: American Control Conference (ACC), pp. 5088-5093. Milwaukee, WI, USA (2018). DOI 10.23919/ACC.2018.8431915.

[5] J. Brown, "CERTAIN: City Environment Range Testing for Autonomous Integrated Navigation", NASA Langley Research Center, 20160009254, May 2016.