An Overview of MAV Research at Brigham Young University

Timothy W. McLain
Dept. of Mechanical Engineering
Brigham Young University
Provo, UT 84602 USA

mclain@byu.edu

Randal W. Beard
Dept. of Electrical and Computer Engineering
Brigham Young University
Provo, UT 84602 USA

beard@ee.byu.edu

D. Blake Barber
Lincoln Laboratories
Massachusetts Institute of Technology
Lexington, MA 02420 USA

d.blake.barber@gmail.com

Nathan B. Knoebel
Dept. of Mechanical Engineering
Brigham Young University
Provo, UT 84602 USA

nathan.knoebel@gmail.com

ABSTRACT

This paper summarizes research efforts at Brigham Young University related to the control of miniature aerial vehicles (MAVs). Recent results in the areas of vector field path following, precision landing and target prosecution, target localization, obstacle detection and avoidance, tailsitter aircraft control, and cooperative control are presented.

1.0 INTRODUCTION

Researchers at Brigham Young University (BYU) have been involved in the study of unmanned aerial vehicles (UAVs) and miniature aerial vehicles (MAVs) since the late 1990's. BYU's efforts have spanned the range from basic research to technology transfer and commercialization. In this paper, we briefly describe a number of BYU-developed technologies with the potential for military and commercial application.

The small size of MAVs has several implications on their performance capabilities. First and foremost, they have limited payload capacity (size, weight, and power) and are therefore unable to carry significant computational resources or sensors of the highest accuracy and capability. Second, their small size and relatively low flight speeds make them susceptible to degraded performance caused by high winds and wind gusts. These challenges, imposed by the small size of MAVs, must be overcome for MAVs to be utilized successfully. Future advancements in miniaturization and performance of sensors and computers will enable increased success. However, efforts must also be made to utilize existing sensor and computer capabilities in novel and innovative ways to enhance the utility of MAV systems in the immediate future. The following subsections describe recent developments at BYU towards this objective.

2.0 VECTOR FIELD PATH FOLLOWING

For MAVs, such as those of primary interest in this work, wind disturbances, dynamic characteristics, and the

McLain, T.W.; Beard, R.W.; Barber, D.B.; Knoebel, N.B. (2007) An Overview of MAV Research at Brigham Young University. In *Platform Innovations and System Integration for Unmanned Air, Land and Sea Vehicles (AVT-SCI Joint Symposium)* (pp. 27-1 – 27-20). Meeting Proceedings RTO-MP-AVT-146, Paper 27. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int/abstracts.asp.

quality of sensing and control all limit the achievable tracking precision. For MAVs wind speeds are commonly 20 to 60 percent of the desired airspeed. Effective path tracking strategies must overcome the effect of this ever present disturbance.

Implicit in the notion of trajectory tracking is that the vehicle is commanded to be in a particular location at a particular time and that this location typically varies in time, thus causing the vehicle to move in the desired fashion. With fixed-wing MAVs, the desired position is constantly moving and the approach of tracking a moving point can result in significant problems for MAVs if disturbances, such as those due to wind, are not accounted for properly. For example, if the MAV is flying into a strong wind (relative to its commanded ground speed), the progression of the trajectory point must be slowed accordingly. Given that wind disturbances vary and are often not easily predicted, trajectory tracking can be very challenging in anything other than calm conditions. Rather than pursuing the trajectory tracking approach, we have pursued path following where the objective is to be *on the path* rather than at a certain point at a particular time. With path following, the time dependence of the problem is removed.

We implement path following through the construction of a vector field surrounding the path to be followed [1, 2]. The vectors of the field provide course commands to guide the MAV toward the desired path. Figure 2.1 shows examples of vector fields for straight-line and circular-orbit paths. Rather than computing a vector field of course commands, course commands are computed as the MAV flies based on its current location relative to the desired path.

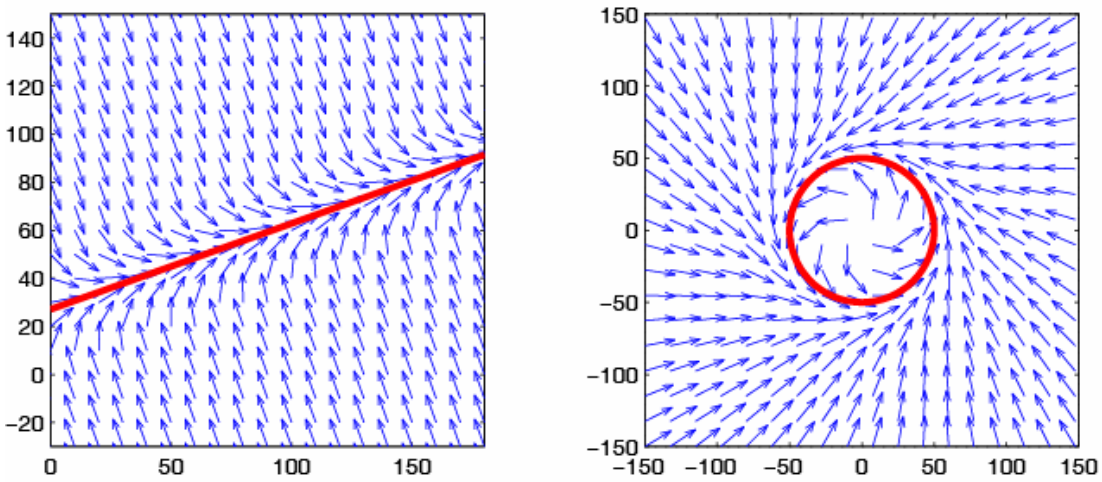

Figure 2.1. Vector fields for straight-line and circular paths.

Critical to the success of the method is the utilization of measurements of course χ and groundspeed V_g for feedback instead of heading ψ and airspeed V_a. The relationship between course/groundspeed and heading/airspeed is depicting in Figure 2.2. Assuming that the altitude and airspeed of the MAV are held constant (or nearly so), a simple kinematic model of the navigational dynamics of the MAV is

$$\begin{aligned} \dot{x} &= V_a \cos\psi + W_x \\ \dot{y} &= V_a \sin\psi + W_y \end{aligned} \tag{1}$$

where W_x and W_y are the x and y components of the wind velocity. Drawing on the relationships between course, heading, airspeed, groundspeed, and windspeed, the model in (1) can be reformulated as

$$\begin{aligned} \dot{x} &= V_g \cos \chi \\ \dot{y} &= V_g \sin \chi . \end{aligned} \tag{2}$$

The key distinction is that the equations of motion are expressed in terms of groundspeed and course and are independent of the wind velocity. We have shown that by using ground-referenced measurements (i.e., course and groundspeed instead of heading and airspeed) in conjunction with the vector field approach to control the path of the vehicle, wind-disturbance rejection can be improved significantly, which is vitally important for small, low-speed MAVs.

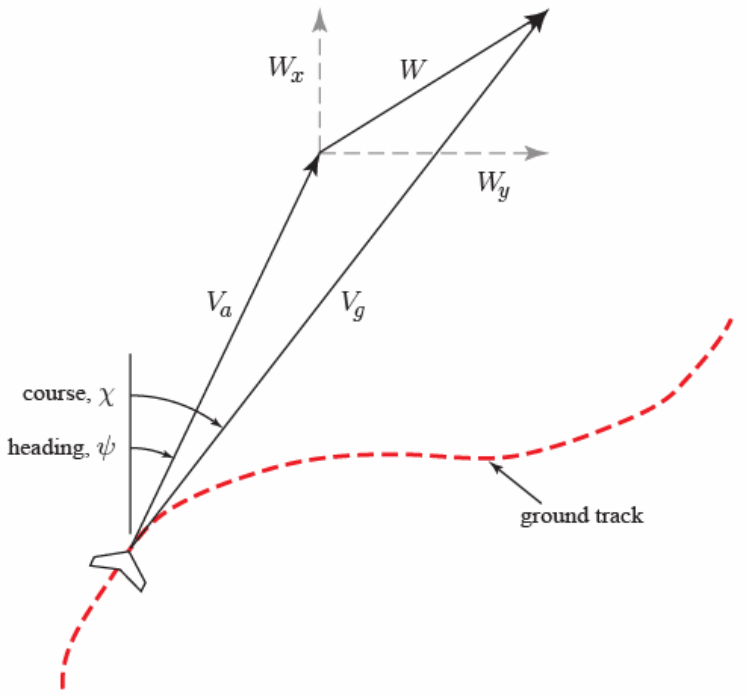

Figure 2.2. Relationship between airspeed, windspeed, and groundspeed.

Experimental results validate the potential value of the approach for MAVs flying in windy conditions. Figure 2.3 shows path following results for a MAV with a 1.3 m wingspan flying a path through city streets. While the flight paths are real, the 30 m wide streets are from a virtual city environment. Winds were about 30 percent of the commanded airspeed. The average tracking error was 3.4 m. These results demonstrate that MAVs can follow paths with precision provided that the effects of wind are dealt with effectively.

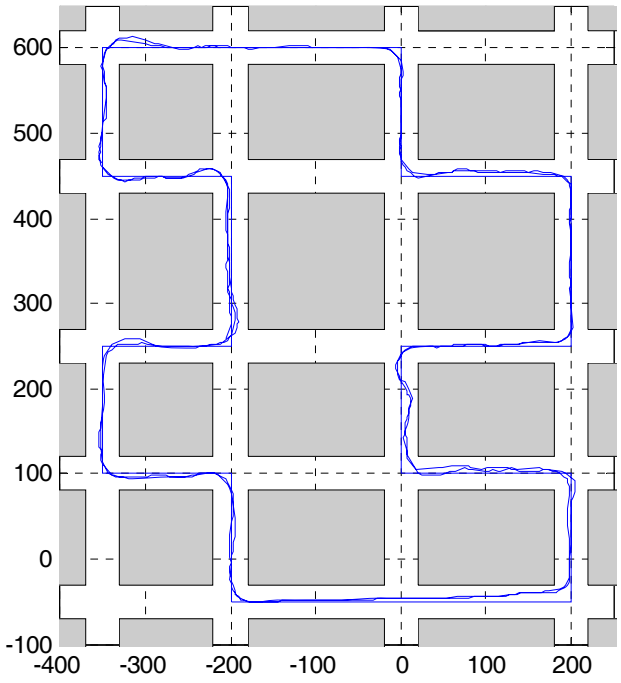

Figure 2.3. Urban canyon flight test results.

3.0 PRECISION LANDING – TARGET PROSECUTION

The ability to land a MAV accurately on a designated landing point is useful for several applications including:

- safe recovery of a MAV in hostile conditions
- supply delivery in rescue operations
- precision strike on a target with minimal collateral damage.

Several challenges make precision landing/strike difficult for small MAVs. The disturbance imposed by wind is always an issue with MAVs, and for the present, sensor errors such as those common in the measurement of attitude and altitude cause difficulties. For precision landing, the most critical information is to know the location of the MAV relative to the desired landing location. Developing reliable methods for measuring this information is central to the success of a precision landing approach.

We have developed precision landing solutions for two situations: one in which the location of the desired landing point is know exactly, and another for the situation where the location of the landing point is known only approximately.

3.1 Precision Landing – Known Landing Location

In the first situation, we use GPS measurements of the MAV location (x and y) to determine a straight-line path to the landing target. Instead of a GPS measurement of altitude, we measure the altitude above ground

level (AGL) to determine the appropriate glideslope to the desired landing point and to control the altitude of the MAV along this glideslope. We have measured the AGL altitude using three methods. We have made a flat-earth assumption and measured barometric altitude as an estimate of altitude AGL. We have used a small laser range finder to measure altitude AGL directly. This worked well with the major disadvantage being the size (32 × 78 × 84 mm) and weight (170 g) of laser range finder. As an alternative to the bulky laser range finder, we have developed a small, lightweight optic-flow sensor with the capability of computing range and utilized it to measure altitude AGL [3].

The optic-flow sensor, shown in Figure 3.1, is constructed by attaching a lens to an Agilent ADNS-2610 optical mouse sensor. The ADNS-2610 has a small form factor, measuring only 10 mm by 12.5 mm and runs at 1500 frames per second. It requires a light intensity of at least 80 mW/m^2 at a wavelength of 639 nm, or 100 mW/m^2 at a wavelength of 875 nm. The ADNS-2610 measures the flow of features across an 18 by 18 pixel CMOS imager. It outputs two values representing the total optic flow across the sensor's field of view in both the *x* and *y* directions. With an appropriate lens and the optic-flow sensor pointing toward the ground, the sensor measures the magnitude of the flow of objects on the ground relative to the MAV. Combining this information with groundspeed measurements from GPS allows the altitude AGL to be determined.

Figure 3.1. Optic-flow sensor.

Precision landing results using altitude AGL estimates were obtained using barometric pressure measurements. Results for 26 landing attempts are shown in Figure 3.2. The mean landing error was 7.6 m while the standard deviation of the landing error was 5.4 m. The accuracy of the results was enhanced by the flatness of the landing site. Errors were largest in the direction of the glideslope path. The combination of a shallow glideslope angle and errors in estimates of the AGL altitude caused overshoot or undershoot of the target.

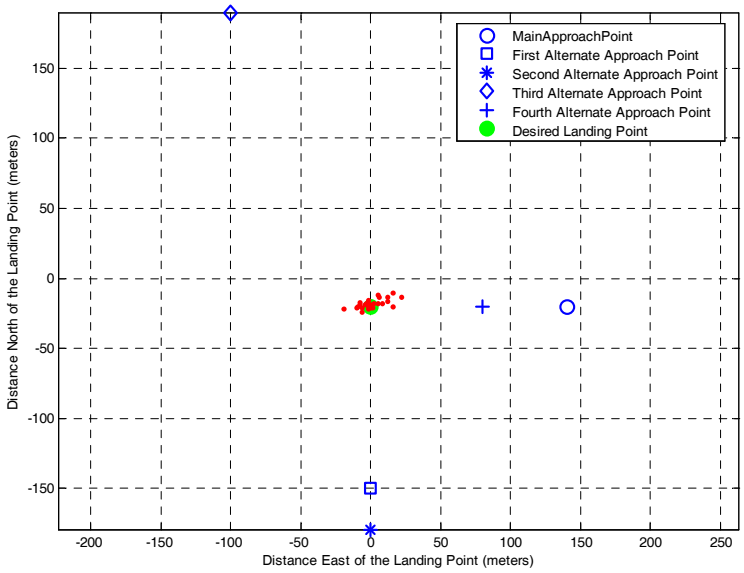

Figure 3.2. Autolanding results using barometric pressure altimeter.

Autolanding results were also obtained using optic-flow measurements to estimate altitude AGL. Results from 27 autolanding attempts are shown in Figure 3.3. The mean landing error for these attempts was 4.3 m while the standard deviation of the landing error was 2.2 m. This represents an improvement of 43 percent in the mean error and 59 percent drop in the standard deviation of the error over results obtained using barometric pressure measurements alone. In addition, direct measurement of altitude AGL provides robustness to varying terrain and altitude differences between the take-off and landing locations.

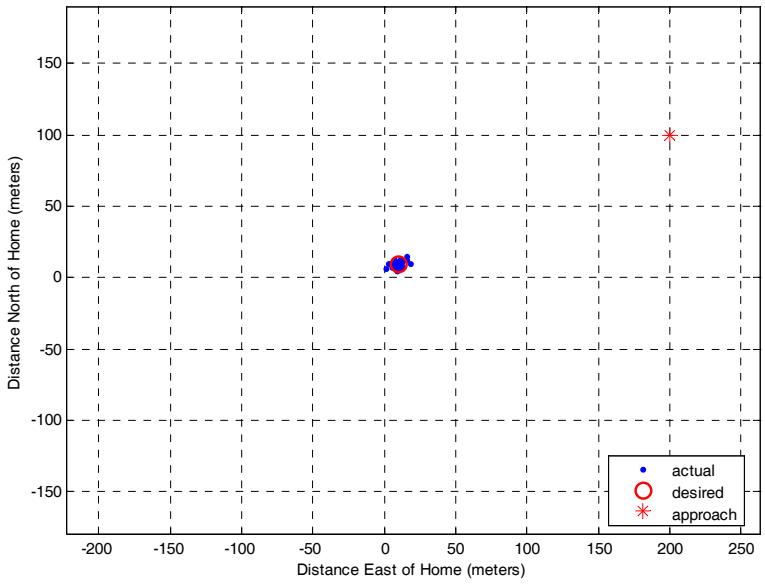

Figure 3.3. Autolanding results using optic-flow sensor.

3.2 Precision Landing – Uncertain Landing Location

Some landing scenarios require the MAV to land at a precise location that may be unknown to the operator prior to the time when landing is necessary. In such situations it is convenient for the operator to select an appropriate landing location in the real-time video imagery acquired by the MAV. From the user-selected landing point in the image, a linear glideslope to the target can be determined. By tracking the landing point in consecutive video frames, the glideslope can be continually modified to overcome errors due to wind, sensor misalignment, or attitude estimation inaccuracies. Following this glideslope will guide the MAV to the desired landing location. We have implemented vision-based landing approaches with vision processing carried out either on the ground station computer or on board the MAV [4]. In this paper, we will focus on our results using onboard vision processing.

Onboard vision processing is enabled by the Helios board shown in Figure 3.4, which was developed at BYU [5, 6]. The Helios board utilizes a Xilinx Virtex-4 FX FPGA and an embedded 450 MHz PowerPC CPU to perform vision processing calculations. It has 8 MB of RAM and 64 MB of SDRAM. It measures 2.5 × 2.5 × 0.5 inches and weighs 2 oz. It typically requires 1 to 3 W of power during operation. For the autolanding experiments reported in this paper, the desired landing location was marked by a square red tarp with sides 2 m in length. The Helios board performed color segmentation and center of mass calculations to distinguish and track the designated landing point. The pixel locations of the landing point center of mass were utilized to provide feedback to the autolanding guidance algorithms.

Figure 3.4. Helios FPGA board for onboard vision processing.

The autolanding guidance algorithm is straightforward conceptually. The approach is to provide flight path angle and course commands to the MAV that align its motion vector with the ray extending from the MAV to the desired landing point. This ray is determined from the MAV attitude and the pixel location of the target in the image. Under perfect conditions (i.e. no wind, perfect sensors), the objective would be to guide the MAV to keep the landing point in the center of the camera image. However, with wind disturbances, sensor misalignments, and sensor biases common to MAVs, additional measures must be taken to ensure accurate landing results. One approach that we have pursued successfully involves estimation of wind disturbances and sensor biases, followed by corresponding corrections to course and glideslope commands to account for these disturbance and bias errors.

15 vision-based autolanding attempts were made on two separate days. On these attempts, the MAV landed on the target on 13 of the 15 attempts. The two attempts that missed the target both landed within 5 m of the

desired landing point. Figure 3.5 shows the results from a particularly challenging attempt where the MAV had a tailwind of 80 percent of the commanded airspeed. Due to the high groundspeed of the MAV only 15 vision samples were available to correct the course of the MAV as it flew toward the target. Once the target was acquired by the vision system, the MAV made a hard right turn to land at the designated point. Figure 3.6 shows the imagery obtained by the vision system for another landing attempt at three different ranges. These images demonstrate the sensitivity of the color segmentation algorithms and their ability to clearly distinguish the target during the descent to the landing target. Current research is being directed towards implementation of feature tracking algorithms on the Helios board to allow landing targets to be of arbitrary color and geometry.

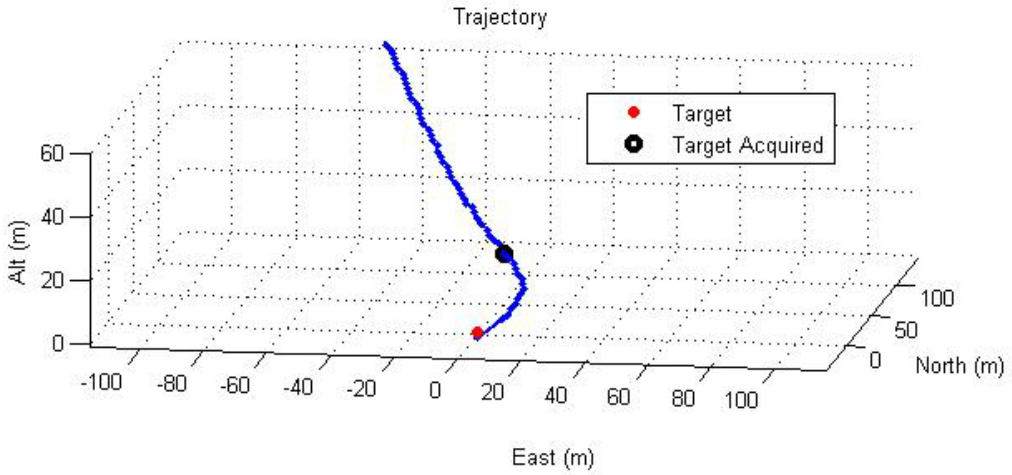

Figure 3.5. Autolanding results with vision. Tailwind 80% of groundspeed.

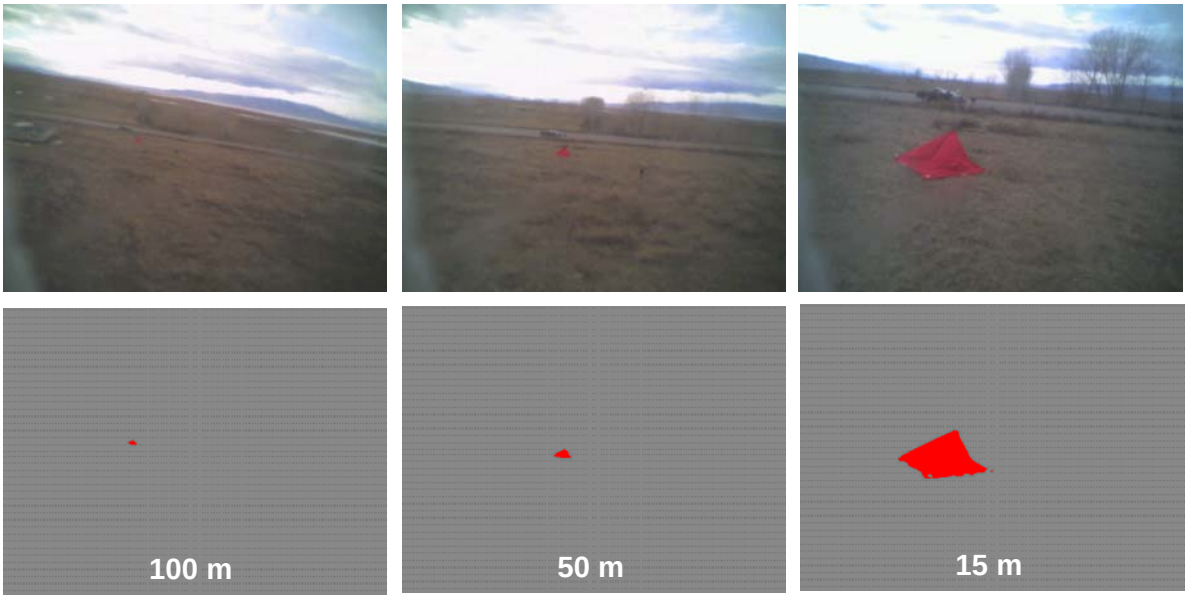

Figure 3.6. Images from onboard camera and output of color segmentation algorithm.

4.0 TARGET LOCALIZATION

The military value of target tracking and localization from a small MAV platform is obvious and significant. This importance of this capability has been demonstrated with large platforms, such as the Predator. Extending this capability to smaller, more widely available MAV platforms is attractive, but presents challenges. The geometry associated with target localization is depicted in Figure 4.1. In our implementation of target localization, the following information is assumed to be available: pixel location of the target in the image frame, position of the MAV, attitude of the MAV, gimbal pointing angles, and camera calibration parameters. From this information, the target localization algorithm calculates the three-dimensional position of the target in the world coordinate frame [7].

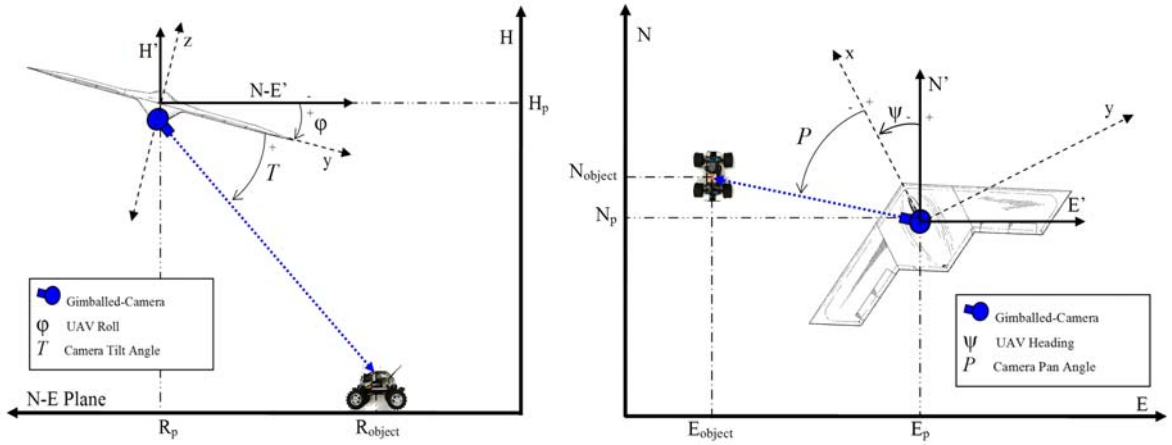

Figure 4.1. Target localization geometry.

The most significant challenges associated with accurate target localization stem from the lack of precise attitude estimates for the MAV platform. Pitch and roll are difficult to estimate with a high degree of accuracy with the MEMS sensors typically used for MAV control. Measurements of heading are not available – heading is often approximated by the course of the MAV which is estimated from successive GPS measurements. In high-wind conditions (relative to the desired airspeed), the course is often a poor approximation of the heading. These measurement difficulties are further complicated by alignment errors between the airframe, the autopilot, and the camera gimbal.

At BYU, we have developed strategies for target localization that mitigate these ever-present sensor errors including recursive least squares (RLS) filtering, bias estimation, flight path selection, and wind estimation. Using these methods we have been able to localize fixed targets to within 2 to 3 m of precision consistently. Figure 4.2 shows the results of a localization experiment. The plot on the left shows the MAV location, the instantaneous estimates of the target location, and the estimate of the target location resulting from the RLS filtering of the instantaneous estimates. This particular plot illustrates the benefits of the RLS filter. Even though the instantaneous estimate errors are about 30 m, the error in the RLS filter estimate is only 6 m. The errors in the localization estimates exhibit a circular pattern about the target. This is due primarily to sensor biases and misalignment. By using a bias estimation approach, this deterministic error can be removed from the instantaneous estimates. The improvement in localization brought about by estimating and correcting biases in the measurements is shown in the right plot of Figure 4.2. With bias correction, instantaneous errors are reduced to 5 to 10 m and the localization error of the filtered estimate is less than 3 m.

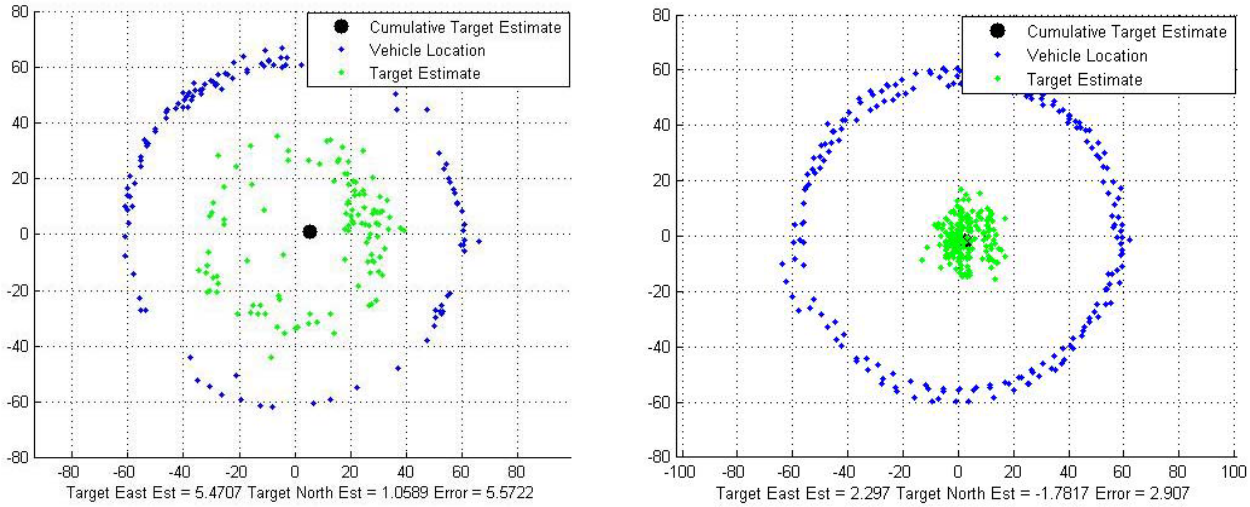

Figure 4.2. Target localization flight test results. The left figure shows results without bias estimation, while the right figure shows results with bias estimation.

The importance of wind estimation in obtaining accurate target location estimates is highlighted by the results shown in Figure 4.3. The results shown in Figure 4.2 were obtained on a day with calm winds, while the results shown in Figure 4.3 were obtained from flights in extremely high winds (above 10 m/s). The irregular flight pattern shown was caused by the high winds. The wind estimation method we have developed estimates the mean wind speed and direction and does not account for gusts. However, localization results from flight tests indicate that the gusts are normally distributed about the mean wind speed, and that the effects of gusts are removed by the RLS filter. The resulting localization estimate for the high-wind test was within 2 m of the known location of the target.

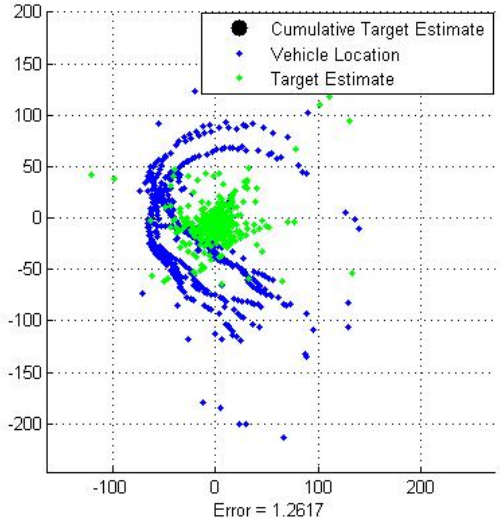

Figure 4.3. Target localization results in high-wind conditions.

5.0 OBSTACLE DETECTION AND AVOIDANCE

The ability to operate in constrained environments is of great importance for MAVs utilized in military applications. Whether in city corridors defined by streets and buildings, or canyons formed by mountainous terrain, MAVs must be able to fly amidst natural and man-made obstacles. We have demonstrated obstacle avoidance with a fixed-wing MAV using two different sensing technologies: optic-flow sensors (discussed in Section 3) and a laser range finder. Optic-flow sensors were applied to the navigation of a mountain canyon corridor [8], while the laser range finder was utilized to avoid a building obstructing the flight path of a MAV [9].

5.1 Canyon Navigation

Given a terrain map of a canyon or mountainous region, safe paths can be generated using path planning algorithms, or manually by defining waypoints. Flights through congested areas using only preplanned paths based on maps are susceptible to collisions with unanticipated obstacles. Reasons for this include inaccurate or biased terrain data, GPS bias error, and the existence of new unmapped obstacles. In addition to careful preplanning, it is important for MAVs to have reactive capabilities to sense and avoid potential hazards.

In this research, the MAV follows a given preplanned path using the vector field approaches described in Section 2. At each time step along the path, the MAV measures its lateral distance to objects on the left and right using the optic-flow range sensors and computes an offset δ from its current path that will place the MAV in the center of the corridor. This is depicted in Figure 5.1 below.

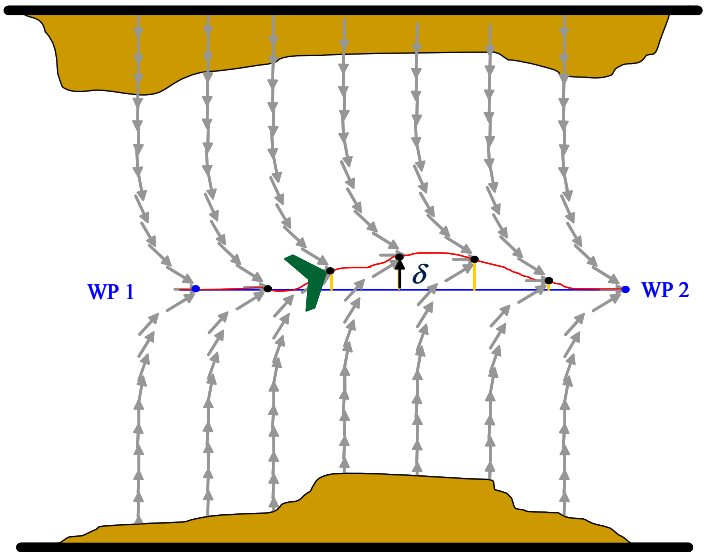

Figure 5.1. Lateral range measurements used to offset the preplanned path and center the MAV in the corridor.

Goshen Canyon in central Utah was chosen as the site to test our canyon navigation approach. Photographs of the flight tests taken by observers and the onboard camera are shown in Figure 5.2. In the first flight through the canyon, a path was planned that followed the road down the middle of the canyon. The MAV navigated the canyon with only minor adjustments to its path. For the second flight, the planned path was intentionally biased into the east canyon wall. The subsequent successful flight verified the effectiveness of the optic-flow-

based navigation algorithm in that it corrected the planned path toward the center of the canyon, enabling the MAV to avoid the canyon walls.

Figure 5.2. Optic-flow-based navigation of Goshen Canyon, Utah. The main frame shows the MAV as seen from a position on the canyon wall. The inset frame shows the canyon as seen from the MAV.

5.2 Reactive Obstacle Avoidance

Reactive obstacle avoidance from a MAV platform is challenging because of the size and weight limitations for sensing and computation hardware imposed by the limited payload of the MAV. The speed with which avoidance decisions must be made and carried out also causes difficulties. As a first step towards obstacle avoidance in urban environments, we have created a reactive avoidance algorithm that utilizes a laser range finder to detect and avoid obstacles. The laser range finder points directly out the front of the MAV and returns range data for objects in front of the MAV at a 3 Hz update.

The scenario of interest is depicted in Figure 5.3 below, where an obstacle obstructs a preplanned waypoint path. Once the laser range finder detects the obstacle, the reactive avoidance algorithm must plan a path that avoids the obstacle and then rejoins the nominal path. Figure 5.3 (a) shows the instant when the obstacle is detected by the laser range finder. The basic idea is to construct an internal map of the obstacles detected by the laser and to modify the waypoint path to maneuver around the obstacles in the internal map. When an obstacle is detected by the laser, a virtual obstacle is inserted into the map at the location detected by the laser. As shown in Figure 5.3 (b) there are two alternate paths for maneuvering around the virtual obstacle. If both paths are collision free, one path is selected arbitrarily, as shown in Figure 5.3 (c). Since the virtual obstacle is smaller than the actual obstacle, the laser again detects the obstacle as it flies the modified waypoint path. A new virtual obstacle is inserted at the appropriate location (as shown in Figure 5.3 (d)) and the process is repeated until the MAV maneuvers around the obstacle, as shown in Figures 5.3 (e) and (f).

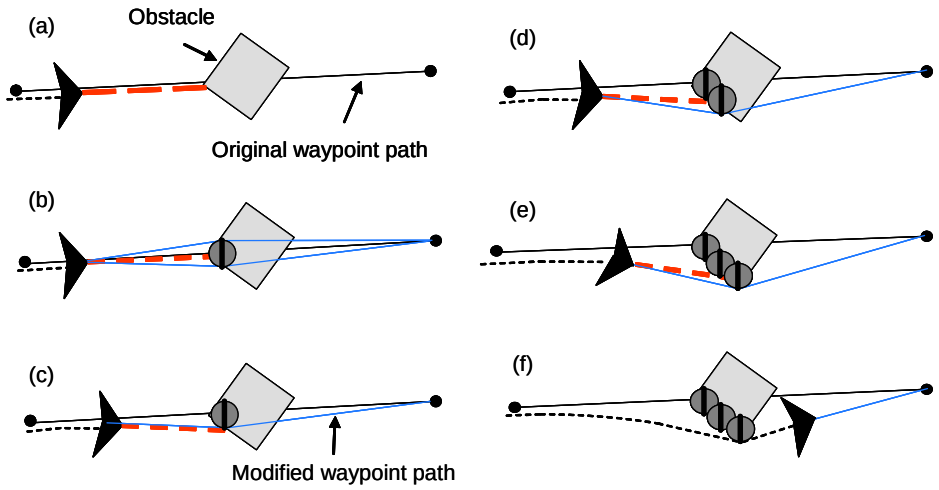

Figure 5.3. Reactive obstacle avoidance algorithm.

For our initial obstacle avoidance flight tests we chose to deal with a single obstacle only. As depicted in Figure 5.3, we planned a straight-line path through an obstacle to demonstrate the effectiveness of the reactive obstacle avoidance approach. The obstacle was chosen to be the tallest building on the BYU campus which is approximately 50 m high. The surrounding buildings were only about 20 m high, which allowed the MAV to fly unobstructed with the exception of the single obstacle in its path. The MAV was directed to fly at an altitude of 40 m along a path passing through the building. No information about the location or the dimensions of the building was provided to the MAV. A plot of the MAV's position obtained from GPS is shown in Figure 5.4.

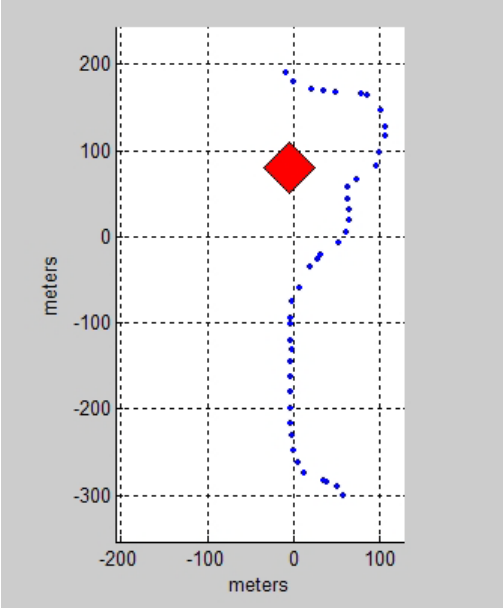

Figure 5.4. Reactive obstacle avoidance flight test results.

As the MAV approached the building, the laser range finder detected the building and returned a range measurement to the building. In response to this obstacle detection, the reactive planner generated an avoidance path around the building. Once the MAV began its turn, it no longer received laser returns from the obstacle. As it started to pass the building, the MAV attempted to rejoin the original waypoint path and detected the building a second time. Based on this detection, the MAV executed a second avoidance maneuver before rejoining the original waypoint path. The MAV was able to avoid the obstacle in its path without any intervention from an operator. Figure 5.5 shows an image of the camera view from the MAV as it initiated the avoidance maneuver. Current research is extending this work to more complex environments.

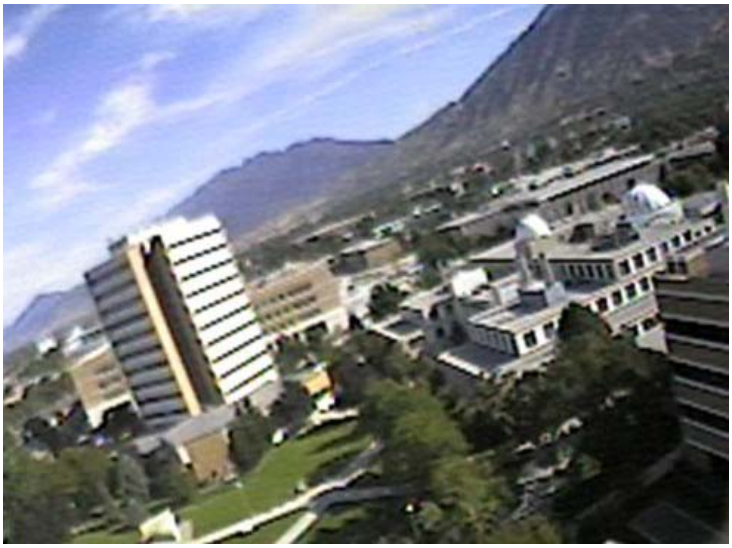

Figure 5.5. In-flight image of building obstacle during avoidance maneuver.

6.0 TAILSITTER AIRCRAFT

VTOL (Vertical Take-Off and Land) MAVs have inherent advantages due to their capability to hover. Such vehicles can persistently image a target in flight, fly in confined areas, and take off and land in tightly restricted regions. These capabilities greatly enhance the autonomy of the MAV, reducing the level of human interaction in recovery and deployment, and also enabling perch and stare functions. Tailsitter MAVs are fixed-wing aircraft with VTOL capabilities that combine the advantages of VTOL with the benefit of efficient flight. BYU's tailsitter MAV is shown in flight in Figure 6.1. The tailsitter MAV, with its many advantages, also poses estimation and control challenges.

Traditionally aircraft attitude is represented by a set of 3-2-1 body-referenced Euler angles (φ, θ, ϕ). This formulation contains a singularity when θ approaches $\pi/2$. A tailsitter by nature hovers vertically at this singularity, and thus is incompatible with this attitude representation. Also elevator control based on error in θ and rudder control based on error in ϕ, which are traditionally used for aircraft attitude tracking, degrade as ϕ and θ become larger than $\pi/4$ respectively. For fixed-wing MAVs, it is common to estimate aircraft heading using course measurements from GPS. In hover, however, the course flown is independent of aircraft heading and an additional measurement must be provided to correctly estimate the heading of the MAV.

Figure 6.1. BYU Tailsitter in flight.

An additional challenge in controlling the flight of a tailsitter comes from its nonlinear dynamics that are difficult to model. The tailsitter flies in, and transitions between, two drastically different flight modes: hover and level flight. Tailsitter aerodynamics, particularly in stall conditions, are complicated and difficult to model. Critical states, such as angle of attack and sideslip angle, are difficult to measure or estimate on MAVs because of payload limitations.

To overcome the challenges associated with tailsitter control, quaternion estimation and adaptive quaternion control methods have been developed [10]. These methods utilize a magnetometer to estimate the heading of MAV in hover. Quaternion attitude representations lack the singularities inherent with the Euler angle formulation. Further, by implementing quaternion-based control approaches utilizing model reference adaptation, the system can quickly learn and compensate for changes in flight conditions and modeling errors. Quaternion-based navigational control algorithms for hover position hold, level flight waypoint tracking, and transitions between these two modes have also been developed.

The model reference adaptive quaternion controller has been derived using Lyapunov and least-squares theory applied to simplified tailsitter dynamics. Simulation results for the controller are shown in Figure 6.2 for a flight that included a takeoff to hover waypoints, a transition to level-flight waypoint tracking, a transition to hover waypoints, and then a hover land. Figure 6.2 shows the reference model and the commanded and actual quaternion elements, while Figure 6.3 shows the adaptive terms (both truth and adapted values) for the pitch-axis dynamics. The transitions between hover and level flight occurred at 17 and 63 seconds. Despite large changes in system parameters, the least squares algorithm estimates parameters well and the controller tracks the reference model as desired.

For hover position hold, successive loop closure is used to hold altitude, and a desired quaternion is generated for north-east coordinate position tracking. The altitude controller gives throttle commands based upon control loops closed around propeller-wash airspeed and altitude. The north-east controller develops a desired quaternion that tilts the aircraft from vertical for translation in the desired direction. Flight test results from this controller with the quaternion estimation described earlier have been obtained and are shown in

Figure 6.4. In this test, the MAV was commanded to move from its take-off point to 30 m north and 30 m east. Current research is being directed towards hardware implementation of the adaptive control algorithms, improving transitions between level flight and hover, and improving the pinpoint landing capabilities of the MAV.

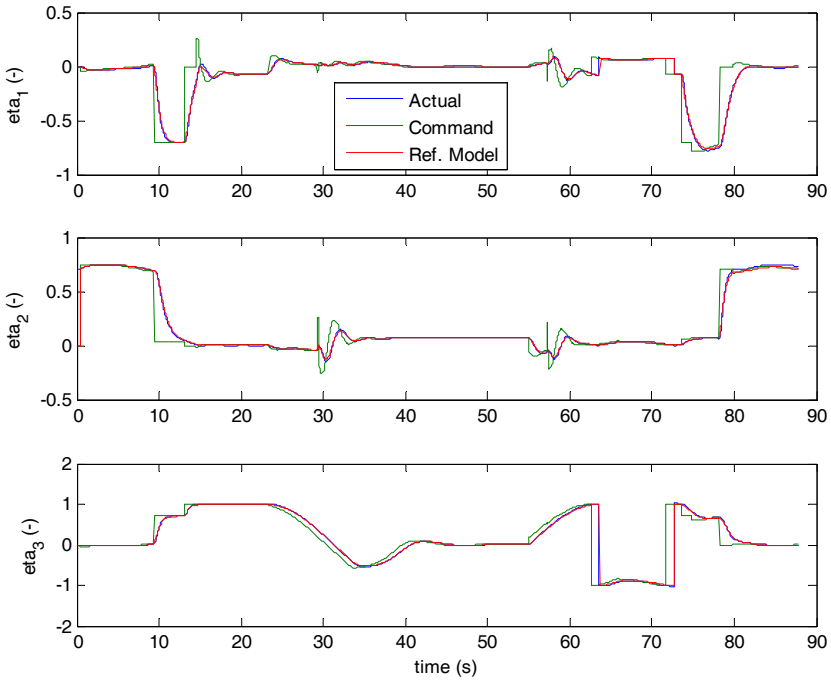

Figure 6.2. Quaternion attitude tracking in simulation of hover, level, and hover flight.

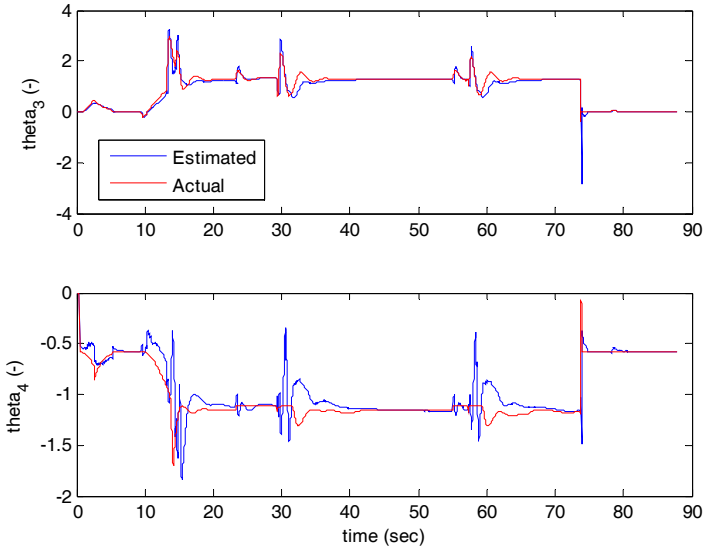

Figrue 6.3. Estimated parameter tracking in simulation of hover, level, and hover flight.

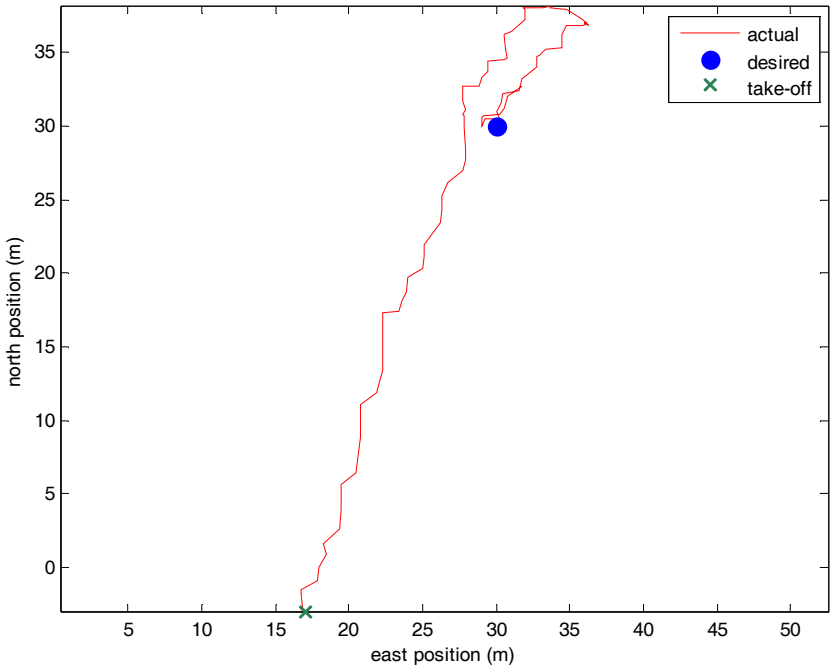

Figure 6.4 Hover position tracking in hardware.

7.0 COOPERATIVE CONTROL

Military missions are typically carried out by teams and the success of the mission often hinges on how well the members of these teams coordinate their activities with one another. At BYU, we have developed a cooperative control framework to enable team-optimal coordination among autonomous or semi-autonomous vehicle teams. The framework requires the definition of coordination variables and coordination functions which encode how information is shared by members of the team [11, 12].

This cooperative control framework has been applied to a variety of problems, perhaps most successfully to problems involving the coordination of timing of actions by members of the team, such as might be required for a simultaneous strike mission. Several flight experiments have been conducted to demonstrate this capability. Figure 7.1 shows telemetry from one of the flight tests. In this test, UAVs were initially searching an area of interest independent of one another. Upon receiving a command from the ground station operator, the UAVs were tasked to arrive at a previously unknown destination simultaneously along defined flight path angles. The UAVs were at different distances from the target requiring them to coordinate their paths and flight speeds. Using this cooperative control framework, the three-UAV teams were able to cooperatively plan paths to arrive simultaneously (within fractions of a second) on a consistent basis. Figure 7.2 is a photograph showing the results of a simultaneous arrival flight experiment.

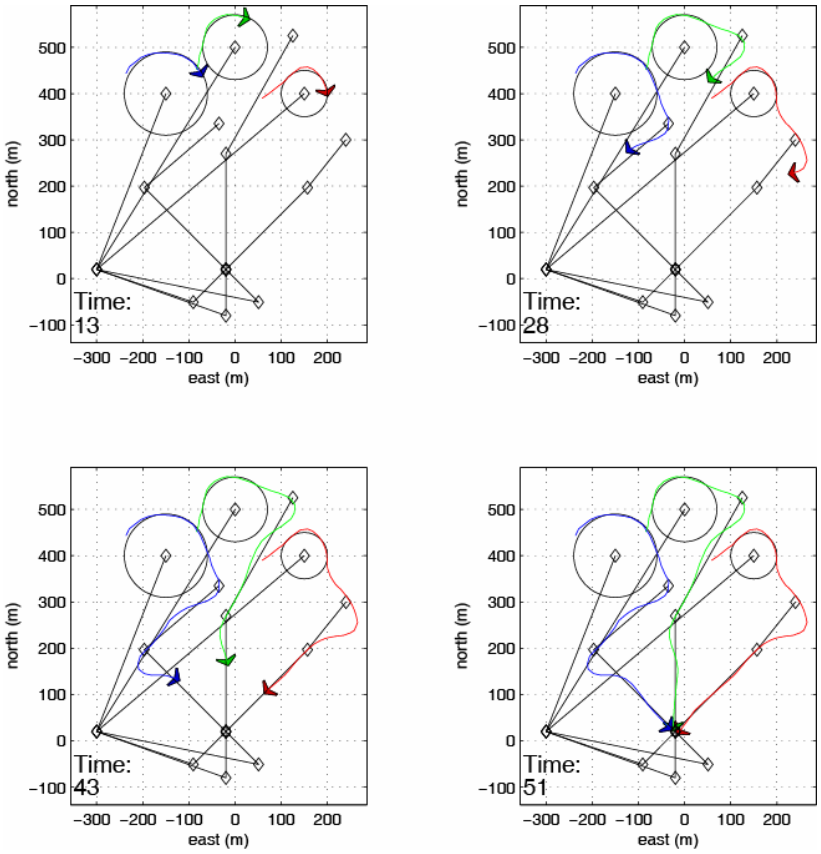

Figure 7.1. Telemetry data from simultaneous arrival experiment.

Figure 7.2. Photo illustrating coordinated simultaneous arrival.

ACKNOWLEDGMENTS

This work was supported by AFOSR award numbers FA9550-0401-0209 and FA9550-04-C-0032, AFRL/MN award number FA8651-05-1-0006, and the Utah State Centers of Excellence Program.

REFERENCES

[1] Nelson, D., Barber, B., McLain, T., and Beard, R. Vector Field Path Following for Small Unmanned Aerial Vehicles, *Proceedings of the American Control Conference,* pp. 5788-5794, June 2006, Minneapolis, Minnesota.

[2] Nelson, D., Barber, B., McLain, T., and Beard, R. Vector Field Path Following for Miniature Air Vehicles, *IEEE Transactions on Robotics*. Accepted for publication.

[3] Barber, D., Griffiths, S., McLain, T., and Beard, R. Autonomous Landing of Miniature Aerial Vehicles, *Proceedings of the AIAA Infotech@Aerospace Conference*, AIAA-2005-6949, September 2005, Washington, DC.

[4] Barber, B., McLain, T., Taylor, C., and Beard, R. Vision-based Landing of Fixed-wing Miniature Air Vehicles, *Proceedings of the AIAA Infotech@Aerospace Conference*, May 2007, Rohnert Park, CA.

[5] Edwards, B., Fife, W., and Archibald, J. A Design Approach For Small Vision-Based Autonomous Vehicles, *Proceedings of the SPIE International Conference on Intelligent Robots and Computer Vision,* Vol. 6384, October 2006.

[6] Fife, W. and Archibald, J. Reconfigurable On-board Vision Processing for Small Autonomous Vehicles, *EURASIP Journal on Embedded Systems,* Article ID 80141, 2007.

[7] Barber, D., Redding, J., McLain, T., Beard, R., and Taylor, C. Vision-based Target Geo-location Using a Fixed-wing Miniature Air Vehicle, *Journal of Intelligent and Robotic Systems,* vol. 47, no. 4, pp. 361-382, December 2006.

[8] Griffiths, S., Saunders, J., Curtis, A., Barber, B., McLain, T., and Beard, R. Obstacle and Terrain Avoidance for Miniature Unmanned Aerial Vehicles, *IEEE Robotics and Automation Magazine, Special Issue on Unmanned Aerial Vehicles: Enabling Technologies & Roadmap for Autonomy,* vol. 13, no. 3, pp. 34-43, September 2006.

[9] Saunders, J., Call B., Curtis, A., Beard, R., and McLain, T. Static and Dynamic Obstacle Avoidance for Miniature Air Vehicles, *Proceedings of the AIAA Infotech@Aerospace Conference*, AIAA-2005-6950, September 2005, Washington, DC.

[10] Knoebel, N., Osborne, S., Snyder, D., McLain, T., Beard, R., and Eldredge, A. Preliminary Modeling, Control, and Trajectory Design for Miniature Autonomous Tailsitters, *Proceedings of the AIAA Guidance, Navigation, and Control Conference,* AIAA-2006-6713, August 2006, Keystone, Colorado.

[11] McLain, T. and Beard, R. Coordination Variables, Coordination Functions, and Cooperative Timing Missions, *AIAA Journal of Guidance, Control, and Dynamics,* vol. 28, no. 1, pp. 150-161, January-February 2005.

[12] Beard, R., McLain, T., Kingston, D. Nelson, D., and Johanson, D. Decentralized Cooperative Aerial Surveillance Using Fixed-Wing Miniature UAVs, *Proceedings of the IEEE,* vol. 94, no. 7, pp. 1306-1323, July 2006.