Maximizing Miniature Aerial Vehicles

Obstacle and Terrain Avoidance for MAVs

BY STEPHEN GRIFFITHS, JEFF SAUNDERS, ANDREW CURTIS, BLAKE BARBER, TIM McLAIN, AND RANDY BEARD

Unmanned aerial vehicles (UAVs) are playing increasingly prominent roles in defense programs and strategy around the world. Technology advancements have enabled the development of large UAVs (e.g., *Global Hawk, Predator*) and the creation of smaller, increasingly capable UAVs. The focus of this article is on smaller fixed-wing miniature aerial vehicles (MAVs), which range in size from .25–2 m in wingspan. As recent conflicts have demonstrated, there are numerous military applications for MAVs, including reconnaissance, surveillance, battle damage assessment, and communications relays.

Civil and commercial applications are not as well developed, although potential applications are extremely broad in scope. Possible applications for MAV technology include environmental monitoring (e.g., pollution, weather, and scientific applications), forest fire monitoring, homeland security, border patrol, drug interdiction, aerial surveillance and mapping, traffic monitoring, precision agriculture, disaster relief, ad hoc communications networks, and rural search and rescue. For many of these applications to develop to maturity, the reliability of MAVs will need to increase, their capabilities will need to be extended further, their ease of use will need to be improved, and their cost will have to come down. In addition to these technical and economic challenges, the regulatory challenge of integrating UAVs into the national and international airspace must be overcome.

Critical to the more widespread use of MAVs is making them easy to use by nonpilots, such as scientists, forest firefighters, law enforcement officers, or military ground troops. One key capability for facilitating ease of use is the ability to sense and avoid obstacles, both natural and man made. Many of the applications cited require MAVs to fly at low altitudes in close proximity to structures or terrain. For example, the ability to fly through city canyons and around high-rise buildings is envisioned for future homeland security operations. For MAVs to be effective tools, the challenge of operating in complex environments must be automated, allowing the operator to concentrate on the task at hand.

Performing obstacle and terrain avoidance from a fixed-wing MAV platform is challenging for several reasons. The limited payload and power availability of MAV platforms place significant restrictions on the size, weight, and power requirements of potential sensors. Sensors such as scanning laser detection and ranging (LADAR) and radar are typically too large and heavy for MAVs. Related to limits on sensor payload are those on computing resources. For most MAVs, the primary computational resource is the excess capacity in the autopilot microcontroller. Additional computational capacity can be added, but computers such as PC104-based systems generally exceed the payload capacity of MAVs; smaller microcontrollers are typically used.

Another challenge posed by fixed-wing MAVs is that they move fast: ground speeds are often in the range of 10–20 m/s (22–44 mi/h). Contrary to the computational limits imposed, obstacle avoidance algorithms must execute and act quickly. Unlike ground robots and unmanned rotorcraft, fixed-wing MAVs cannot stop or slow down while avoidance algorithms process sensor information or plan maneuvers. Reactions must be immediate. Obstacle sensing is further complicated by the fact that sensor readings are altered by changes in aircraft attitude, especially the

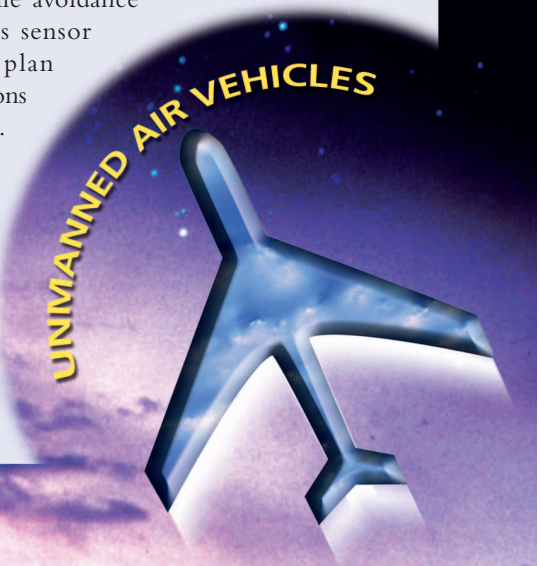

1070-9932/06/$20.00©2006 IEEE

rolling motions that occur during turns. Attitude changes affect not only the pointing direction of the sensor, but also cause motion of fixed objects in the field of view. For avoidance maneuvers to be successful, obstacle and terrain detection must account for the effects of aircraft attitude changes. All of the challenges associated with MAV obstacle and terrain avoidance are compounded by the reality that for MAVs, mistakes are costly or even catastrophic; crashes can result in damage to or loss of the MAV and failure to complete the objectives of the flight.

As evidenced by the recent Defense Advance Research Projects Agency (DARPA) Grand Challenge, capable obstacle avoidance and terrain navigation systems have been developed for ground vehicles. Obstacle avoidance and path planning have been active areas of research for many years, and the associated robotics literature is immense. While providing a guiding influence, most of the proposed methods fail to deal with the sensing and computational challenges imposed by the limited payload capabilities of MAVs.

As autonomous MAVs and feasible obstacle sensors are recent technological developments, the body of experimental research directed specifically toward MAV obstacle and terrain avoidance is small. Related to terrain avoidance is work focused on utilizing vision processing techniques to estimate height above ground. Chahl et al. demonstrated that mimicking the landing behavior of bees, by maintaining constant optic flow during a landing maneuver, could be used to successfully control the descent of a MAV [1]. The development of lightweight sensors for the measurement of optic flow has enabled their use in MAVs [2], [3], [4]. Barrows et al. have demonstrated that these sensors can be used to follow undulations in terrain with low-flying MAVs [5].

This article presents MAV obstacle and terrain avoidance research performed at Brigham Young University (BYU). Our work builds on the notion of utilizing useful but imperfect map information to plan nominal paths through city or mountain terrain. Because maps may be limited in resolution, out of date, or offset in location, MAVs must also utilize sensory information to detect and avoid obstacles unknown to the path planner. In this article, we present research utilizing laser range finder and optic-flow sensors to detect obstacles and terrain. Avoidance algorithms using this sensor information are discussed briefly, and flight test results from our MAVs are presented.

BYU Miniature Aerial Vehicle Platforms

Over the past five years, BYU has been involved in the development of MAV airframes, autopilots, user interfaces, sensors, and control algorithms. This section describes the experimental platform developed specifically for the obstacle avoidance research described in this article.

Airframe

Figure 1 shows the airframe used for obstacle avoidance experiments. The airframe has a 1.5-m wingspan and was constructed with an expanded polypropylene (EPP) foam core covered with Kevlar. This design was selected for its durability, usable payload, ease of component installation, and flight characteristics. The airframe can carry a 0.4-kg payload and can remain in flight for over 45 min at a time. The collision avoidance sensors that are embedded in the airframe include three optic-flow sensors, one laser ranger, and two electro-optical cameras as shown in Figure 2. Additional payload includes the Kestrel autopilot, batteries, a 1-W, 900-MHz radio modem, a 12-channel global positioning system (GPS) receiver, and a video transmitter.

Kestrel Autopilot

The collision avoidance algorithms described in this article were implemented on Procerus Technologies' Kestrel Autopilot version 2.2 [6]. The autopilot is equipped with a Rabbit 3400 29-MHz processor, three-axis rate gyros, three-axis accelerometers, absolute and differential pressure sensors, and a variety of interface ports. The autopilot

Figure 1. Airframe used for collision avoidance experiments.

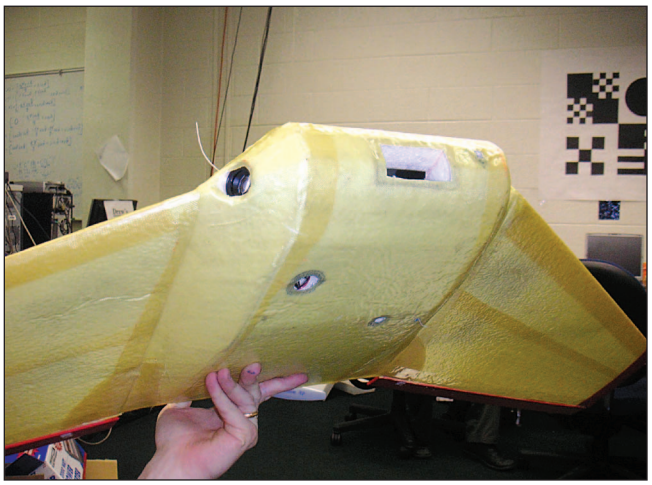

Figure 2. Sensors used for collision avoidance. The round hole on the right and the large hole on the belly are the optic-flow sensors. The square hole in the center is the laser range, and the other two round holes are for electro-optical cameras.

Given a nominal waypoint path, it is essential for the MAV to have the ability to track the path with precision.

measures 3.8 cm × 5.1 cm × 1.9 cm and weighs 18 g. The autopilot also serves as a data acquisition device and is able to log 175 kB of user-selectable telemetry at rates up to 60 Hz. The optic-flow sensors and the laser ranger used in this article are connected directly to the autopilot and the collision avoidance algorithms are executed on board the Rabbit processor.

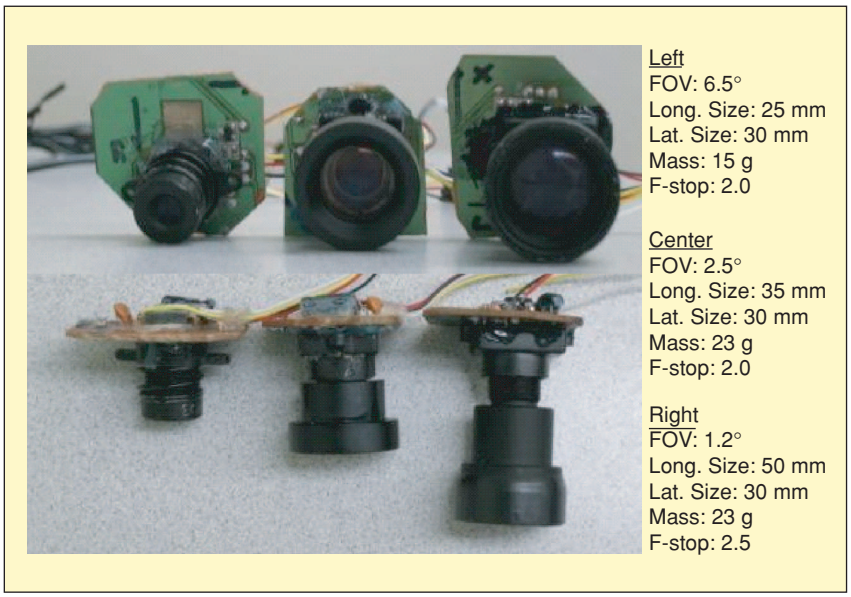

Figure 3. Optic-flow sensors with three different lens configurations: 1.2°, 2.5°, and 6.5° field of view. The optic-flow sensors are constructed by attaching a lens to an optical mouse chip.

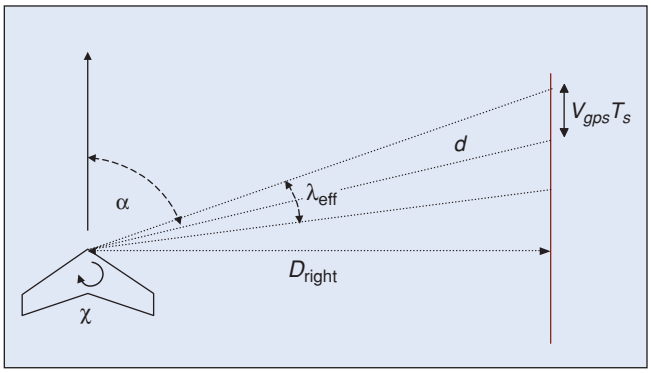

Figure 4. The optic-flow sensor is used to compute the distance to an obstacle based on the distance traveled between samples ($V_{gps}T_s$) and the effective field of view λ.

Optic-Flow Sensors

The MAV is equipped with three optic-flow sensors. Two of the optic-flow sensors are forward looking but swept back from the nose by $\alpha = 60°$. The third optic-flow sensor points down to determine the height above ground. The optic-flow sensors, shown in Figure 3, are 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 1,500 frames/s. 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 × 18 pixel complementary metal-oxide semiconductor (CMOS) imager. It outputs two values, δp_x and δp_y, representing the total optic flow across the sensor's field of view in both the x and y directions. The flow data in the camera y direction correspond to lateral motion of the MAV and are ignored.

Figure 4 indicates how distance is computed using the optic-flow sensor. The optical mouse chip outputs an optic-flow displacement $(\delta p_x, \delta p_y)^T$ at its internal sample rate (1,500 Hz). Since the collision-avoidance loop is executed at $T_s = 20$ Hz, the total optical displacement is integrated over T_s to produce $(\Delta p_x, \Delta p_y)$. The distance to the object D is related to the measured distance d by the expression

$$D = d \cos\phi \sin\alpha,$$

where ϕ is the roll angle of the MAV. From geometry, the measured distance to the object is given by

$$d = \frac{V_{\text{gps}} T_s}{\tan(\frac{\lambda_{\text{eff}}}{2})},$$

where λ_{eff} is the effective field of view. The effective field of view is given by

$$\lambda_{\text{eff}} = \lambda_{\text{cam}} \frac{\Delta p_x}{P_x} - \dot{\chi} T_s,$$

where λ_{cam} is the field of view of the camera, P_x is the size of the pixel array along the direction of motion, and $\dot{\chi}$ is the yaw rate with respect to the ground. Using similar reasoning for left-looking and down-looking optic-flow sensors we can derive the following expression:

$$D_{\text{right}} = \frac{V_{\text{gps}} T_s}{\tan\left(\frac{\lambda_{\text{cam}} D_{\text{right}} p_x}{2 P_x} - \frac{\dot{\chi} T_s}{2}\right)} \cos\phi \sin\alpha$$

$$D_{\text{left}} = \frac{V_{\text{gps}} T_s}{\tan\left(\frac{\lambda_{\text{cam}} D_{\text{left}} p_x}{2 P_x} + \frac{\dot{\chi} T_s}{2}\right)} \cos\phi \sin\alpha$$

$$D_{\text{down}} = \frac{V_{\text{gps}} T_s}{\tan\left(\frac{\lambda_{\text{cam}} D_{\text{down}} p_x}{2 P_x} - \frac{\dot{\theta} T_s}{2}\right)} \cos\theta \cos\phi.$$

Laser Ranger

For the experiments discussed in this article, we used the Opti-Logic RS400 Laser range finder. The range finder has a range of 400 m with an update rate of 3 Hz. It weighs 170 g and consumes 1.8 W of power. Figure 2 shows the laser ranger mounted in the airframe. It is important to note that the RS400 is not a scanning laser range finder. Scanning laser range finders are currently too heavy and consume too much power for MAV applications. The RS400 returns a single distance measurement and must be steered by maneuvering the airframe.

Path Planning and Following

The first step in our approach for navigating through complex environments is to plan a nominal path based on known information about the environment, which is usually in the form of a street map or topographic map. The MAV must be able to accurately follow the nominal path to avoid known obstacles. This section discusses the methods for planning and following the nominal path. Subsequent sections will discuss reactive, sensor-based obstacle avoidance strategies for obstacles unknown during the planning process.

Planning the Nominal Path

When planning paths through complex environments, the computational requirements for finding an optimal path can be significant and unrealistic for near-real-time execution [7]. Because of this, recent research has focused on randomized techniques to quickly find acceptable, though not necessarily optimal, paths [8], [9]. Path planning for MAVs is also difficult because of the dynamic constraints of flight. Many common path planning algorithms are inadequate for fixed-wing MAV systems because they do not handle turn-radius limitations and airspeed constraints effectively.

One randomized method that addresses these limitations is the rapidly exploring random tree (RRT) algorithm [7], [10]. RRTs use a dynamic model of the system to build a tree of traversable paths. The search space is quickly explored by applying control inputs to states already in the tree. Working with the precise control inputs ensures that the dynamic constraints are not violated; however, it also results in an open-loop solution. This would be adequate if we had a perfect model of the system and no disturbances, but this method is not satisfactory for an actual MAV because of model inaccuracies and disturbances such as wind.

Similar to Frazzoli et al. [11], we have extended some of the concepts of RRTs to plan paths in the output space. Through this work, we have developed a useful a priori path planner for the MAVs [12]. Our modified RRT algorithm searches the output states instead of the inputs and produces a list of waypoints to track. This is sufficient if we can bound the error of the controlled MAV from the waypoint path. For a given waypoint path, we can determine the expected trajectory of the MAV [13] and ensure that only traversable paths are built into the search tree. Branches in the tree are checked to ensure that they pass tests on turn radius and climb rate and are collision-free. Figure 5 depicts the growth of an RRT path through a simulated urban environment.

For our preliminary flight tests, we considered a simple scenario: a single unknown obstacle placed directly in the flight path.

Vector Field Path Following

Given a nominal waypoint path, it is essential for the MAV to have the ability to track the path with precision. MAVs must track these paths despite dynamic limitations, imprecise sensors and controls, and wind disturbances, which are often 20–60% of airspeed [14]. Trajectory tracking, which requires the MAV to be at a specific location at a specific time, is difficult in such wind conditions. As an alternative, we have developed a path following approach where the focus is simply to be on the path, instead of at a specific point that evolves in time. Similar research in [15] describes a maneuvering method focused on converging to the path then matching a desired speed along the path. Our path following method is based on the creation of course vector fields that direct the MAV onto the desired path.

The vector field method produces a field of desired course commands that drive the MAV toward the current path segment. At any point in space, the desired course can be easily calculated. This desired course is used to command heading and roll control loops to guide the MAV onto the desired path.

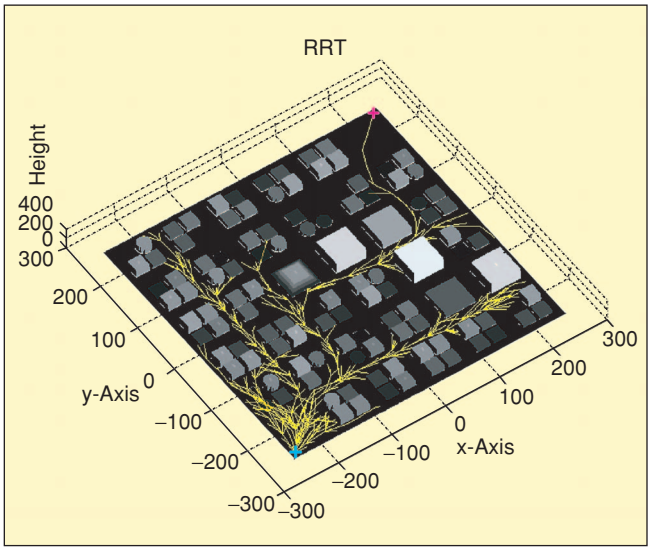

Figure 5. This figure shows the growth of an RRT path tree through a simulated urban enviroment. The algorithm is terminated once a feasible path to the destination (red X) is found.

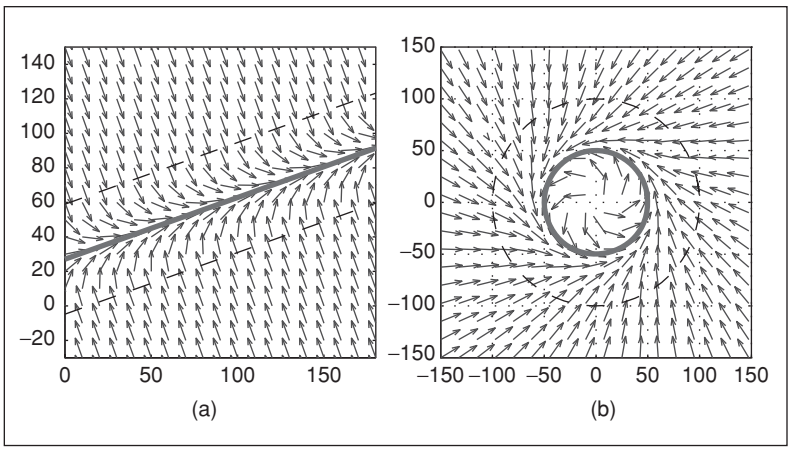

Figure 6. *Path following in wind is accomplished by creating a vector field of desired course commands based on the lateral deviation from the path: (a) shows a possible vector field for a straight-line waypoint path segment and (b) shows a possible vector field for orbit following.*

The vector field method uses only the current path segment to find the desired course, avoiding possible singularities and sinks resulting from sums of vectors. Many paths planned for MAVs can be approximated by combinations of straight line segments and circular arcs [16]. Figure 6 shows examples of vector fields for linear and circular paths.

To account for wind, we use the course and groundspeed instead of heading and airspeed to control the MAV. Groundtrack motion is the vector sum of the MAV motion relative to the surrounding air mass and the motion of the air mass relative to the ground. Since course direction includes the effects of wind, control based on course is much more effective at rejecting wind disturbances. In implementing the vector field approach, course measurements from GPS are compared with the desired course from the vector field to determine the appropriate control inputs to keep the MAV on the path.

For a given path, the vector field is divided into a transition region and an outer region. This is similar in some respects to the belt zone technique developed by Loizou et al. [17] Outside the transition region, the vector field drives the MAV toward the transition region along a constant course. Once inside, the vector field changes linearly from the entry course direction to the desired course along the path. The effect is to smoothly drive the MAV to follow the path, with larger effort as the error from the path increases. In [14] it is shown that for any initial condition, the MAV will enter the transition region in finite time then converge to the desired course asymptotically.

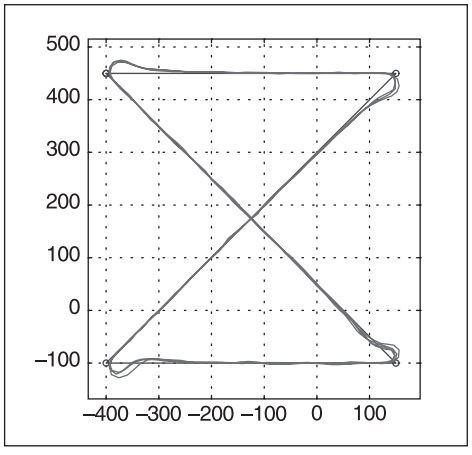

Figure 7. *This figure shows telemetry data for four consecutive traversals of a waypoint path. Wind speeds during the flight were 20% of the MAV airspeed. Note the repeatability of the trajectories even in significant wind.*

Flight tests have demonstrated the effectiveness of the vector field path following method, even in windy conditions. Figure 7 demonstrates path following for straight line segments with acute angles. Wind speeds were approximately 20% of the airspeed during these tests. The vector field method has been shown to be effective in tracking paths of lines and orbits with wind speeds of up to 50% of the airspeed of the MAV.

Reactive Obstacle and Terrain Avoidance

Despite having an effective a priori path planner, we cannot guarantee that the flight path will be free of obstacles. Our path planner assumes a perfect model of the terrain, but this assumption is not realistic. If an urban terrain model is missing a newly constructed building or a large antenna or tree, a path leading to a collision could result. Our canyon models are based on 10 m United States Geological Survey (USGS) data. They are fairly accurate but cannot represent small obstacles like trees and power lines. In addition, the GPS sensor used on the MAV has a constant bias that can be as large as 10 m. Path planners can produce a nominal path prior to flight, but the MAV must also have the ability to sense and reactively avoid unanticipated obstacles and terrain in real time.

The following sections present reactive planners for producing deviations from a nominal path to enable obstacle and terrain avoidance. The first section presents a method for sensing and avoiding obstacles directly in the flight path and shows results for reactive avoidance of a building. The second section presents an approach for staying centered between obstacles as might be required for flying through a corridor. Flight test results are presented that demonstrate autonomous navigation of a winding canyon.

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 platform. The speed with which avoidance decisions must be made and carried out also causes difficulties. For obstacle avoidance in urban environments, we have developed a heuristic algorithm that utilizes a laser ranger to detect and avoid obstacles. The laser ranger points directly out the front of the MAV and returns range data for objects directly in front of the MAV with a 3-Hz update. For our preliminary flight tests, we considered a simple scenario: a single unknown obstacle placed directly in the flight path.

Algorithm

Consider the scenario shown in Figure 8, where obstacle avoidance is required. The MAV has a forward ground velocity V and a minimum turn radius R and is assumed to be tracking the given waypoint path at the time the obstacle is detected by the laser, which has a look-ahead distance L. Figure 8(a) shows the instant when the obstacle is detected by the laser ranger. The basic idea is to construct an internal map of obstacles detected by the laser and to modify the waypoint path to maneuver around the obstacles in the internal map. We will refer to the internal representation of obstacles as *map obstacles*. When the laser detects the location of an obstacle, we are unsure about the size and height of the obstacle. We propose representing map obstacles as cylinders with radius R equal to the minimum turn radius of the MAV and height equal to the current altitude of the MAV. As shown in Figure 8(b), there are two alternate waypoint paths that maneuver around the map obstacle. The endpoints of the waypoint paths are selected so that the new waypoint paths are tangent to the obstacles in the internal map. As shown in Figure 8(a), the new waypoints are located a distance $dR/\sqrt{d^2 - R^2}$ from the original waypoint path, where d is the turn-away distance from the obstacle. If both waypoint paths are collision free, then the algorithm randomly selects between the two paths as shown in Figure 8(c). Since the map obstacle may be smaller than the actual obstacle, the laser may again detect the obstacle as it maneuvers on the modified path. If that is the case, a new map obstacle is added to the internal map as shown in Figure 8(d). This process is repeated until the MAV maneuvers around the obstacle as shown in Figure 8(e) and (f).

If we assume zero wind, then the two-dimensional (2-D) navigation for the MAV is given by

$$\dot{n} = V \cos\chi$$
$$\dot{e} = V \sin\chi$$
$$\dot{\chi} = \frac{g}{V} \tan\phi,$$

where g is the gravitational constant, and ϕ is the roll angle of the MAV. On most MAVs, the roll angle is limited between $-\bar{\phi} \le \phi \le \bar{\phi}$. We will assume that the roll dynamics of the MAV are sufficiently fast to assume near instantaneous transitions between $\pm\bar{\phi}$. Therefore, the minimum turn radius is given by $R = V^2/g \tan\bar{\phi}$.

We would like to establish a minimum turn-away distance D so that we are guaranteed to avoid collision with a single rectangular obstacle. The first step is to determine the bounds on the forward and lateral motion of the MAV when it transitions from one waypoint path to the next.

Claim: After the insertion of a map obstacle, the MAV requires at most a forward distance of $(2/\sqrt{3})R$ *and a lateral distance of* $\sqrt{2/3}R$ *to transition onto the new waypoint path while avoiding the map obstacle.*

Assuming the ability to roll instantaneously between $\pm\bar{\phi}$, the motion of the MAV during the transition can be constrained to lie on circles of radius R. As shown in [13], the path length of the transition increases monotonically with the angle between the old and new waypoint paths. Therefore, the forward and lateral distances are maximized when the angular separation is maximized, which occurs when instantaneous motion of the MAV follows a circle of radius R that just touches the map obstacle, as shown in Figure 9(b). The claim follows directly from standard geometrical arguments. Note that the maximum angular separation is therefore given by $\theta = \tan^{-1}(1/\sqrt{2}) \approx 36^\circ$.

The algorithms we have developed enable the MAV to center itself within a corridor or canyon or to fly near walls with a specified offset.

Claim: Avoidance of a collision with a flat wall is guaranteed if the the turn-away distance D satisfies

$$D > \left(\frac{8 + 2\sqrt{6}}{2\sqrt{3}}\right) R. \tag{1}$$

Consider the worst-case scenario, shown in Figure 9(c), of a MAV that is initially traveling perpendicularly to a flat wall. The MAV detects an obstacle and inserts a waypoint at maximum angle $\tan^{-1}(1/\sqrt{2})$. After aligning its heading with the waypoint path, the wall is again detected, a map obstacle is inserted, and a new waypoint with maximum angle $\tan^{-1}(1/\sqrt{2})$ is planned. This scenario will repeat itself at most three times since $3\tan^{-1}(1/\sqrt{2}) > \pi/2$. Therefore, the maximum forward direction is bounded by

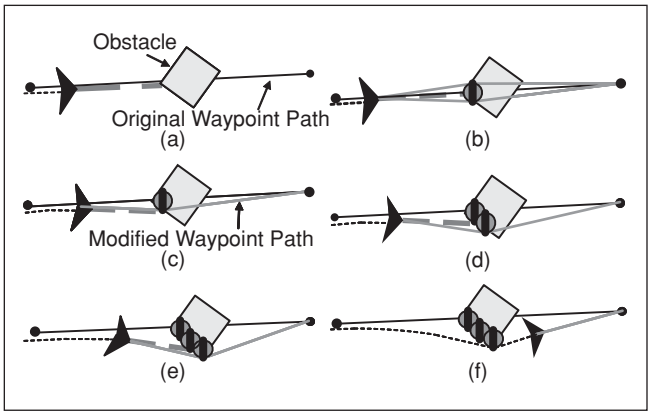

Figure 8. *Obstacle avoidance algorithm: (a) the laser detects the obstacle; (b) a map obstacle of radius R is inserted into the map, and two candidate waypoint paths are constructed; (c) a modified waypoint path is randomly selected; (d) the obstacle is again detected by the laser and another map obstacle is constructed; and (e)–(f) the process repeats until the MAV is able to maneuver around the obstacle.*

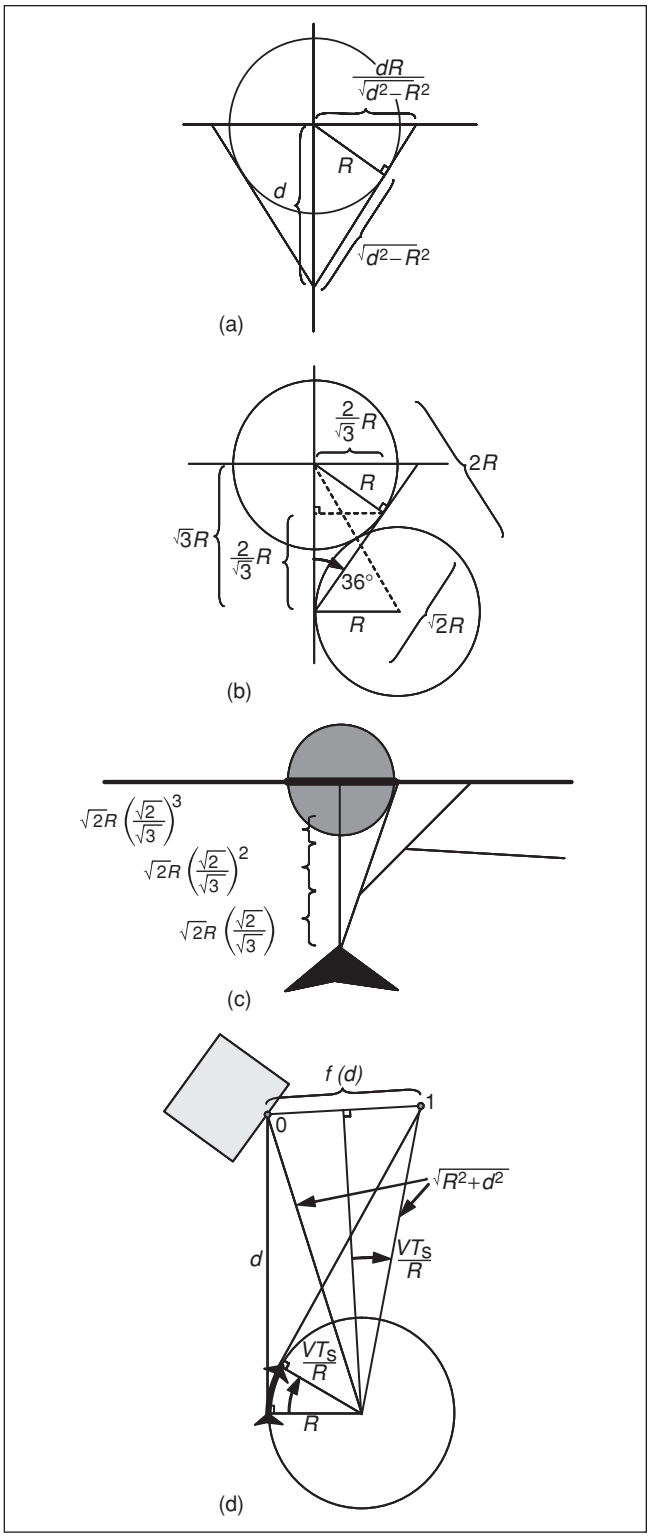

Figure 9. *(a) The waypoint path is constructed so that it is perpendicular to the map obstacle. The radius R ensures collision free passage around the map obstacle. (b) The maximum heading change in waypoint paths is when the MAV must make a full bank to maneuver around the obstacle. (c) An approximation of the minimum distance required to avoid a straight wall if the laser is only sampled when the MAV is on the waypoint path. (d) The geometry used to calculate the distance between two consecutive laser updates.*

$$\sqrt{2}R\left(\left(\sqrt{\frac{2}{3}}\right)^1+\left(\sqrt{\frac{2}{3}}\right)^2+\left(\sqrt{\frac{2}{3}}\right)^3\right)=\left(\sqrt{\frac{8+2\sqrt{6}}{2\sqrt{3}}}\right)R.$$

We note that the algorithm described above requires that the laser detect points on the obstacle that are outside of the map obstacles as soon as they become visible. Is this feasible given the update rate of the laser? Let T_s be the time between laser updates.

Claim: The maximum distance between laser updates at a range of $d \leq L$ is given by

$$f(d) = 2\sqrt{R^2 + d^2}\sin\left(\frac{VT_s}{2R}\right).$$

Assuming the vehicle is turning at its maximum rate, the change in heading between updates is VT_s/R. Utilizing the geometry depicted in Figure 9(d), the calculation of $f(d)$ is straightforward. To ensure overlap of map obstacles between samples, we require that $f(D) < R$, which implies that

$$T_s < \frac{2R}{V}\sin^{-1}\left(\frac{R}{2\sqrt{R^2+D^2}}\right).$$

For our airframes, typical values are $V = 13$ m/s, $R = 25$ m, which implies from (1) that $D = 93$ m and $T_s < 0.5$ s. The laser ranger sample period of 0.33 s satisfies this constraint, thus ensuring that map obstacles overlap between samples.

Results

For initial testing of the reactive avoidance algorithm, we chose to deal with a single obstacle only. It was important that the obstacle be tall enough to allow the MAV to fly at a safe altitude. Flying at an altitude of 40 m also prevented the laser ranger from detecting points on the ground that might be mistakenly interpreted as obstacles and allowed for losses of altitude that can occur during aggressive maneuvers.

For our flight tests, we used the tallest building on the BYU campus (the Kimball Tower), which is 50 m high and 35 m square. The surrounding buildings are only about 20 m in height. The MAV was directed to fly at 40 m altitude from the south side of the building to the north along a waypoint path that passed directly through the building. No information about the location or the dimensions of the building was provided to the MAV. A GPS telemetry plot of the results is shown in Figure 10.

As the MAV approached the building, the laser ranger detected the building and calculated its position. When the MAV came within 93 m of the building, the reactive planner generated a path around the building, and the MAV began to track the path. Notice that as the MAV began to pass the building, it turned towards the original waypoint path and detected the building a second time. This caused the MAV to execute a second avoidance maneuver before rejoining the original waypoint path. The MAV successfully

avoided the building without human intervention. Figure 11 shows images of the MAV and its camera view as it executed the avoidance maneuver.

Remote Environment Terrain Avoidance

As small MAVs become more reliable and maneuverable, their missions will involve navigating through complex terrain, such as mountainous canyons and urban environments. In this section, we focus on terrain avoidance for flying in corridors and canyons. The algorithms we have developed enable the MAV to center itself within a corridor or canyon or to fly near walls with a specified offset. The algorithms utilize optic-flow sensors like those shown in Figure 3. To validate our algorithms, canyon navigation flight experiments were carried out in a mountain canyon.

Canyon Navigation Algorithm

The first step in navigating through a canyon or urban corridor is to select a suitable path through the terrain. This can be done using the RRT algorithm discussed earlier, or the operator can utilize maps to define waypoints for the MAV to follow. Preplanned paths will rarely be perfect, and some paths could lead the MAV near or even into uncharted obstacles. Reasons for this include inaccurate or

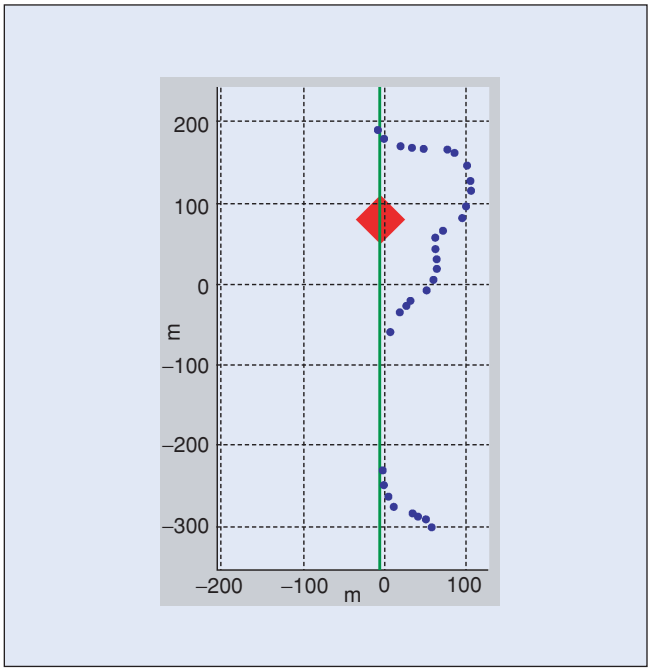

Figure 10. Flight results for collision avoidance using a laser ranger. The green line indicates the planned waypoint path, and the dotted line indicates the GPS track of the MAV.

Figure 11. In-flight image of the Kimball Tower on the BYU campus during the collision avoidance maneuver.

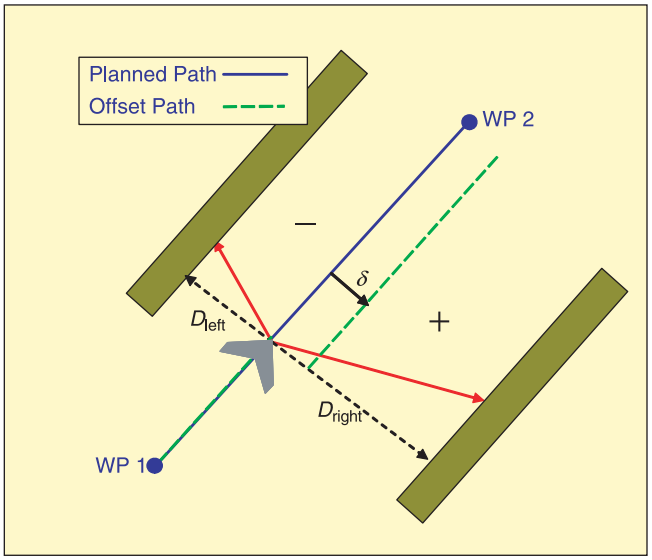

Figure 12. Using the measurements from the optic-flow sensors, the planned path (solid blue) is shifted by δ to create a new desired path (dashed green) that is centered between the canyon walls.

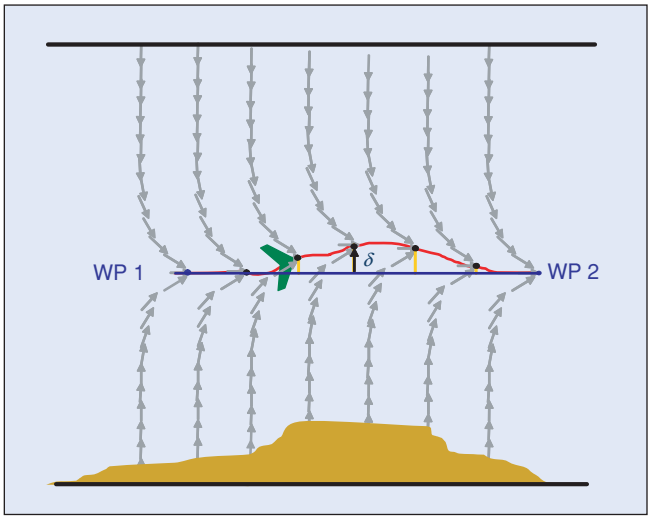

Figure 13. The adjusted path (red) is offset from the preplanned path (blue) by the calculated offset (δ) at each time step to center the desired path between the canyon walls, thus shifting the vector field along with it.

To broaden the range of applications for MAVs, methods to enable operation in environments of increased complexity must be developed.

biased terrain data, GPS error, and the existence of obstacles that have been added since the terrain was mapped. Therefore, it is important that the MAV be able to make adjustments to its path to center itself between walls and other potential hazards.

In our approach, the MAV follows its preplanned path using the vector field following method. At each time step along the path, the MAV computes its lateral distance from objects to the left and right using the optic-flow ranging sensors. Using this information, the MAV computes an offset δ from its planned path

$$\delta = \frac{1}{2}(D_{\text{right}} - D_{\text{left}}), \tag{2}$$

where D_{left} and D_{right} are distances to walls on the left and right measured by the optic-flow sensors. Shifting the desired path by this offset centers the desired path between the detected walls as shown in Figure 12. As Figure 13 illustrates, shifting the desired path also shifts the vector field accordingly. To improve the performance of this method, the optic ranging sensors are pointed forward at a 30° angle. This reduces lag caused by filtering the sensor readings and allows the MAV to detect obstacles ahead of its current position.

Figure 14. This figure shows the MAV as it enters Goshen Canyon. The inset is an image from the camera on board the MAV.

Figure 15. Results from the second flight through Goshen Canyon. Flight test results show the planned path (green) and the actual path (blue). The planned path was intentionally biased to the east, forcing the MAV to offset from its planned path of centering itself through the canyon.

Flight Test Results

Goshen Canyon in central Utah was chosen as a flight test site. This canyon was selected for its steep winding canyon walls, which reach over 75 m in height, as well as its proximity to BYU and low utilization. Flight tests through Goshen Canyon were conducted using the fixed-wing MAV discussed previously. Photographs of the flight tests taken by observers and the onboard camera are shown in Figure 14. In the first flight through the canyon, the planned path was selected to follow the road. 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 to verify that the navigation algorithms would correct the planned path toward the center of the canyon, enabling the MAV to avoid the canyon walls.

Figure 15 shows results from the second flight, demonstrating that the MAV biased its desired path up to 10 m to the right to avoid the canyon walls. If the MAV had not biased its path, it would have crashed into the east canyon wall.

Summary

Miniature aerial vehicles have demonstrated their potential in numerous applications. Even so, they are currently limited to operations in open air space, far away from obstacles and terrain. To broaden the range of applications for MAVs, methods to enable operation in environments of increased complexity must be developed. In this article, we presented two strategies for obstacle and terrain avoidance that provide a means for avoiding obstacles in the flight path and for staying centered in a winding corridor. Flight tests have validated the feasibility of these approaches and demonstrated promise for further refinement.

Acknowledgments

This work was funded through Air Force Office of Scientific Research (AFOSR) awards FA9550-04-1-0209 and FA9550-04-C-0032.

Keywords
Miniature aerial vehicle, obstacle avoidance, terrain navigation, autonomous flight.

References
[1] J. Chahl, M. Srinivasan, and S. Zhang, "Landing strategies in honeybees and applications to uninhabited airborne vehicles," *Int. J. Robot. Res.*, vol. 23, no. 2, pp. 101–110, 2004.

[2] G. Barrows and C. Neely, "Mixed-mode VLSI optic flow sensors for in-flight control of a micro air vehicle," in *Proc. SPIE,* San Diego, Aug. 2000, pp. 52–63.

[3] F. Ruffier and N. Franceschini, "Visually guided micro-aerial vehicle: automatic take off, terrain following, landing and wind reaction," in *Proc. 2004 IEEE Int. Conf. Robotics Automation*, New Orleans, 2004, pp. 2339–2346.

[4] J.-C. Zufferey and D. Floreano, "Toward 30-gram autonomous indoor aircraft: Vision-based obstacle avoidance and altitude control," in *Proc. 2005 IEEE Int. Conf. Robotics Automation*, Barcelona, Apr. 2005, pp. 2594–2599.

[5] G.L. Barrows, J.S. Chahl, and M.V. Srinivasan, "Biomimetic visual sensing and flight control," *Aeronaut. J., London: Royal Aeronaut. Soc.*, vol. 107, no. 1069, pp. 159–168, 2003.

[6] Procerus Technologies [Online]. Available: http://procerusuav.com.

[7] S.M. LaValle and J.J. Kuffner, "Randomized kinodynamic planning," *Int. J. Robot. Res.*, vol. 20, no. 5, pp. 378–400, May 2001.

[8] L.E. Kavraki, P. Švestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Automat.*, vol. 12, no. 4, pp. 66–80, 1996.

[9] N.M. Amato and Y. Wu, "A randomized roadmap method for path and manipulation planning," in *Proc. IEEE Int. Conf. Robotics Automation*, Minneapolis, MN, 1996, pp. 113–120.

[10] S.M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," Computer Science Dept., Iowa State Univ., Tech. Rep. 98-11, Oct. 1998.

[11] E. Frazzoli, M.A. Dahleh, and E. Feron, "Real-time motion planning for agile autonomous vehicles," *AIAA J. Guid., Contr. Dynam.*, vol. 25, no. 1, pp. 116–129, 2002.

[12] J.B. Saunders, B. Call, A. Curtis, R.W. Beard, and T.W. McLain, "Static and dynamic obstacle avoidance in miniature air vehicles," in *Proc. Infotech@Aerospace Conf.*, Sept. 2005, AIAA-2005-6950.

[13] E.P. Anderson, "Constrained extremal trajectories and unmanned air vehicle trajectory generation," Master's thesis, Brigham Young Univ., Provo, Utah, Apr. 2002.

[14] D.R. Nelson, "Cooperative control of miniature air vehicles," Master's thesis, Brigham Young Univ., Provo, Utah, Dec. 2005.

[15] R. Skjetne, T. Fossen, and P. Kokotović, "Robust output maneuvering for a class of nonlinear systems," *Automatica*, vol. 40, pp. 373–383, 2004.

[16] E.P. Anderson, R.W. Beard, and T.W. McLain, "Real time dynamic trajectory smoothing for uninhabited aerial vehicles," *IEEE Trans. Contr. Syst. Technol.*, vol. 13, no. 3, pp. 471–477, May 2005.

[17] S. Loizou, H. Tanner, V. Kumar, and K. Kyriakopoulos, "Closed loop motion planning and control for mobile robots in uncertain environments," in *Proc. 42nd IEEE Conf. Decision Control*, 2003, pp. 2926–2931.

Stephen Griffiths received his M.S. degree in mechanical engineering degree from Brigham Young University in April 2006. His master's research investigated terrain navigation strategies for small unmanned aerial vehicles. He is currently employed in the intelligent and autonomous control systems group at Scientific Systems Company, Inc., Woburn, Massachusetts.

Jeff Saunders is a Ph.D. student in the Electrical and Computer Engineering Department at Brigham Young University (BYU). He received a B.S. in electrical engineering from BYU in 2004. In summer 2006, he spent time as a visiting researcher at the Air Force Research Laboratory Air Vehicles Directorate. He currently pursues research in the area of obstacle and collision avoidance for unmanned aerial vehicles.

Andrew Curtis is a graduate student at Brigham Young University. He received his B.S. in electrical engineering in April 2006 and is now pursuing an M.S. in electrical engineering. His current research focuses on path planning and trajectory generation for miniature unmanned aerial vehicles.

Blake Barber is a graduate student in the mechanical engineering department at Brigham Young University (BYU). He received a B.S. degree in mechanical engineering from BYU in 2005. He has been actively involved in researching and testing new flight control algorithms and sensor platforms for miniature unmanned aerial vehicles. His research includes the use of optic-flow sensors for range measurement, autonomous landing, path following, and target localization.

Tim McLain is an associate professor in the mechanical engineering department at Brigham Young University (BYU). He received an M.S. degree in mechanical engineering from BYU in 1987 and a Ph.D. in mechanical engineering from Stanford University in 1995. He has been actively involved in the control of air and underwater vehicles and robotic systems for the past 18 years. During the summers of 1999 and 2000, he was a visiting scientist at the Air Force Research Laboratory. Since that time, he has pursued research involving the modeling and control of miniature unmanned aerial vehicles (UAVs), real-time trajectory generation for UAVs, and cooperative control of UAV teams. He is a Senior Member of the IEEE and AIAA and is a member of the AIAA Unmanned Systems Program Committee.

Randy Beard is currently an associate professor in the Department of Electrical and Computer Engineering at Brigham Young University (BYU). He received his Ph.D. in electrical, computer, and systems engineering from Rensselaer Polytechnic Institute in 1995. He was a research fellow at the Jet Propulsion Laboratory, California Institute of Technology in 1997 and 1998. His research interests are in nonlinear control, multiple vehicle coordination, and autonomy for unmanned air vehicles (UAVs) and mobile robots. He is the cofounder of the BYU MAGICC Lab that conducts research in miniature UAVs and cooperative control. He is a Senior Member of the IEEE and an associate editor of *IEEE Control Systems Magazine*.

Address for Correspondence: Tim McLain, Department of Mechanical Engineering, Brigham Young University, 435 S CTB, Provo, UT 84602 USA. Phone +1 801 422 6537. E-mail: mclain@byu.edu.