JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. , No. ,

Adaptive Glide Slope Control for Parafoil and Payload Aircraft

Michael Ward* and Mark Costello†
Georgia Institute of Technology, Atlanta, Georgia 30332

DOI: 10.2514/1.59260

The invention of the steerable, gliding, ram-air parafoil enabled the possibility of precision, autonomous aerial payload delivery. Research and development work on guided airdrop systems has focused primarily on evolutionary improvements to the guidance algorithm, although the navigation and control algorithms have changed little since the initial autonomous systems were developed. Recent work has demonstrated the potential for dramatic improvements in landing accuracy through the incorporation of canopy incidence variation to achieve glide slope control for parafoils. The current work presents the development of a control law to implement glide slope control on an autonomous airdrop system. Autonomous landings with the control law in both simulation and flight test demonstrate an improvement in landing accuracy by a factor of two, though the improvement can be even greater in especially windy conditions. Finally, the ability to perform in-flight system identification to adapt internal control mappings to match flight data and provide dramatic improvements in landing accuracy when there is a significant discrepancy between the assumed and actual flight characteristics is demonstrated in both simulation and flight test.

Nomenclature

b	=	canopy span
c	=	chord
C_L, C_D	=	lift and drag coefficients
C_l, C_m, C_n	=	roll, pitch, and yaw moment coefficients
GS	=	glide slope
h	=	altitude above ground level
$\mathbf{I}_B, \mathbf{J}_B, \mathbf{K}_B$	=	unit vectors defining body fixed reference frame
$\mathbf{I}_C$, $\mathbf{J}_C$, $\mathbf{K}_C$	=	unit vectors defining canopy fixed reference frame
p, q, r	=	body frame roll, pitch and yaw rates
$\tilde{u}$, $\tilde{v}$, $\tilde{w}$	=	components of velocity of canopy with respect to air
V_0	=	forward component of airspeed
V_G	=	ground speed
V_w	=	wind speed
x, y, z	=	inertial position in north, east, down and frames
$\dot{z}$	=	descent rate
α	=	angle of attack
β	=	sideslip
δ	=	control input
$\bar{\delta}$	=	normalized control input
δ_A	=	asymmetric brake deflection
δ_B	=	symmetric brake deflection
δ_L	=	left brake deflection
δ_R	=	right brake deflection
δ_I	=	incidence angle control input
Γ	=	incidence angle
χ	=	azimuth angle
ψ	=	heading angle

I. Introduction

THE use of parachutes to deliver cargo provides a unique capability for the rapid deployment of very large payloads to remote and inaccessible locations. Traditional airdrop systems are based on round, unguided parachutes which slow the descent rate of a payload by producing a large amount of drag. An alternative form of airdrop to drag-based, round parachutes is a form of steerable, gliding parachute invented by Domina Jalbert in the 1960s known as a parafoil. The potential application of the parafoil for autonomous guided cargo delivery was recognized immediately, and the first autonomous flight tests were performed in 1966 by Knapp and Barton [1]. Early autonomous systems made use of radio beacons for guidance [1–5], but the advent of GPS finally provided the practical, reliable, and accurate position feedback required to enable precision autonomous airdrop. In the early 1990s, NASA and the United States Army started programs to develop guided parafoil aircraft, and the guidance, navigation, and control (GNC) software for both programs was developed at the Charles Stark Draper Laboratory, Inc. [6–9]. The basic flight profile established during these programs is the standard for all current guided parafoil algorithms [10–25]. There are three basic phases to a guided airdrop flight: 1) go to the target, 2) loiter, and 3) execute a landing maneuver. Current autonomous airdrop GNC algorithms typically take advantage of the loiter phase to perform in-flight wind estimation. This wind estimate is then used to plan the final landing maneuver. There is normally sufficient time during loiter to obtain a very accurate wind estimate, so unless there is some type of malfunction, all landing errors can be traced to some source of uncertainty during the final landing approach. The primary sources of uncertainty for an airdrop system are deviations from the assumed wind profile and deviations from the assumed flight-dynamic model. The current work seeks to improve parafoil accuracy by addressing both of these sources of uncertainty by handling deviations from the assumed wind with improved longitudinal control authority and deviations from the assumed flight-dynamic model with adaptive control strategies.

The control mechanisms used on autonomous parafoil systems have not changed since the development of the first systems in the 1960s. The trailing edge of the canopy is deflected downward asymmetrically to turn and symmetrically to change speed or to flare while landing. Symmetric deflection of the trailing edge brakes produces an increase in both lift and drag; this provides a reduction in speed but little change in glide angle until stall. The limited longitudinal control of current airdrop systems makes it difficult to address the possibility of deviations from the assumed wind during landing approach. Yakimenko et al. showed that deviations in the wind below an altitude of 100 m can shift the landing point of their system by over 100 m off target [25]. The average accuracy of current guided parafoils is somewhere between 75 and 100 m [26], and the accuracy numbers for various systems in the same weight class are often very similar because these systems have all converged to essentially the same GNC strategies. A point of diminishing returns is being reached with the evolutionary updates to the current approach.

Received 8 June 2012; revision received 24 October 2012; accepted for publication 6 November 2012; published online XX epubMonth XXXX.

*Graduate Research Assistant, Guggenheim School of Aerospace Engineering, Member AIAA.
†Professor, Guggenheim School of Aerospace Engineering, Woodruff School of Mechanical Engineering, Associate Fellow AIAA.

The most direct way of introducing the potential for large gains in accuracy potential is to increase control authority through the addition of new control mechanisms. Slegers et al. demonstrated the use of variable rigging geometry to obtain glide slope control of a parafoil and payload system [27]. This novel control mechanism was explored in depth with a flight-test program examining the coupled effect of incidence angle and trailing edge brake deflection on the flight characteristics of canopies of two different aspect ratios by Ward et al. [28]. Based on this understanding of incidence angle as a control mechanism, the current work develops a glide slope control algorithm to choose optimal incidence angle and brake deflection inputs in response to variations in the wind. This is the first time a control algorithm has been presented to control glide slope over ground using incidence angle and brake deflection.

Model uncertainty can be addressed by identifying the needed dynamic and control characteristics in-flight using a system identification algorithm. The key quantities to estimate for an autonomous airdrop system are the airspeed, descent rate, and the lateral control response. Recently, the need for in-flight estimation of these characteristics has been recognized and efforts to address this are beginning to appear in the literature. Jann [20] and Carter et al. [23] developed GNC algorithms with the ability to estimate airspeed in-flight. Calise and Preston [29] developed an adaptive control law for lateral control of a parafoil. A thorough examination was performed of the benefits of estimating the system characteristics in-flight with different levels of sensor error and turbulence was performed by Ward et al. [30]. In the current work, a robust methodology is developed to adapt the internal model of the flight-dynamic characteristics to match these in-flight estimates.

The paper begins with a summary of the flight-dynamic simulation used to develop and test the autonomous precision landing algorithm. The simulation is validated against flight test data. The basic guidance, navigation and control algorithm used for autonomous landings is then presented. The basic GNC algorithm is then expanded with the development of a completely new and original glide slope control algorithm which makes use of both symmetric brake and incidence angle to control glide slope during final approach. This control algorithm is tested extensively in simulation and compared against the basic landing algorithm. The effectiveness of the glide slope control method is also compared to the basic algorithm in an extensive flight-test program performed with atmospheric conditions varying from complete calm up to conditions where the mean wind speed exceeded the air speed of the flight-test vehicle and the updrafts were strong enough to produce significant climb rates. To handle uncertainty in the internal model of the system's control response, a robust and efficient methodology is developed for adapting the internal flight-dynamic model to in-flight estimates of the key flight characteristics. Extensive simulation results are presented demonstrating the ability of this method to identify and compensate for large amounts of model uncertainty in a wide variety of simulated atmospheric conditions. Finally, flight test results demonstrate the effectiveness and robustness of this adaptive control technique in actual flight tests in windy and turbulent conditions.

II. Simulation

Modern engineering of airdrop systems leans heavily on flight-dynamic modeling and simulation to predict a multitude of drop events virtually, so that GNC software can be developed and tested in a cost efficient manner. There is a large body of work on different methods of parafoil dynamic modeling. The simplest approach is to model the entire parafoil and payload aircraft as a rigid body and decouple the longitudinal and lateral dynamics to obtain a reduce order model [31–33]. Although these reduced order models can provide insight into a particular aspect of parafoil flight dynamics, it is necessary to model the full set of rigid body states in order to provide a realistic environment to test lateral and longitudinal guidance and control algorithms [34,35]. In reality, the relative motion between the payload and canopy can be significant, which calls for additional degrees of freedom beyond the rigid body motion

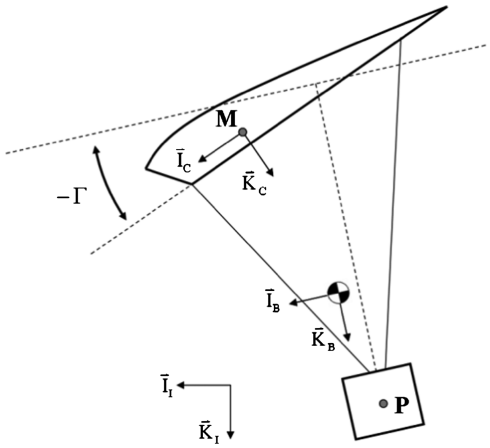

Fig. 1 Parafoil and payload schematic.

[36–38]. Gorman and Slegers presented a comparison of the dynamics of six, seven, eight and nine degree of freedom (DOF) parafoil simulations [39]. They concluded that a seven DOF simulation is required if any sensor information related to payload orientation is used for control, however, a six DOF model is sufficient if the autonomous GNC algorithm uses only position and velocity feedback. The only sensors used in the current work are a GPS receiver and barometric altimeter, so the minimal order model which captures the relevant dynamics is a six DOF model.

Figure 1 shows a schematic of a parafoil and payload system. With the exception of movable parafoil brakes, the parafoil canopy is considered to be a fixed shape. The canopy is allowed to rotate with respect to the system through the incidence angle Γ. The combined system of the parafoil canopy and the payload are modeled with a six DOF, including three inertial position components of the total system mass center as well as the three Euler orientation angles. The canopy aerodynamic forces and moments are computed about the canopy aerodynamic center (point *M* in Fig. 1). The transformation from the body frame (frame *B* in Fig. 1) to the canopy reference frame (frame *C* in Fig. 1) is defined by a single axis rotation in pitch by the canopy incidence angle.

The equations of motion for this six DOF parafoil and payload representation have been derived previously [27,30,40], but the aerodynamic model used to represent the canopy is altered slightly for the current work. The aerodynamic forces and moments on the canopy are computed at the aerodynamic center in the reference frame attached to the canopy. The velocity components of the canopy aerodynamic center, with respect to the air, are expressed in the canopy reference frame and are denoted with tildes, $\tilde{u}$,$\tilde{v}$, $\tilde{w}$, so that the airspeed, angle of attack, and angle of sideslip are given by:

$$\tilde{V} = \sqrt{\tilde{u}^2 + \tilde{v}^2 + \tilde{w}^2}, \qquad \alpha = \tan^{-1}\tilde{w}/\tilde{u}, \qquad \beta = \sin^{-1}\tilde{v}/\tilde{V} \tag{1}$$

The aerodynamic forces on the canopy in the canopy reference frame are computed as functions of these aerodynamic quantities

$$\begin{Bmatrix} F_{X,\text{canopy}} \\ F_{Y,\text{canopy}} \\ F_{Z,\text{canopy}} \end{Bmatrix} = \frac{1}{2}\rho\tilde{V}^2 S[T_a]\begin{Bmatrix} -C_D \\ C_{Y\beta}\beta \\ -C_L \end{Bmatrix} \tag{2}$$

where T_a is a y axis rotation by the angle of attack. The angle of attack used to compute lift and drag forces is offset by a term proportional to symmetric brake deflection:

$$\alpha' = \alpha + \alpha_{\delta B}\delta B \tag{3}$$

The lift and drag coefficients depend on this modified angle of attack and symmetric brake deflection:

Fig. 2 Self powered system in flight.

$$C_D = C_{D0} + C_{D0\delta B}\delta B + (C_{D\alpha 2} + C_{D\alpha 2\delta B}\delta B)\alpha'^2 \tag{4}$$

$$C_L = C_{L0} + C_{L0\delta B}\delta B + (C_{L\alpha} + C_{L\alpha\delta B}\delta B)\alpha' + C_{L\alpha 3}\alpha'^3 \tag{5}$$

The aerodynamic moments are given by Eq. (6):

$$\begin{Bmatrix} M_{X,\text{canopy}} \\ M_{Y,\text{canopy}} \\ M_{Z,\text{canopy}} \end{Bmatrix} = \frac{1}{2}\rho\tilde{V}^2 S \begin{Bmatrix} bC_l \\ cC_m \\ bC_n \end{Bmatrix} \tag{6}$$

The roll and pitch moment coefficients are dependent on the canopy angular rates:

$$C_l = \frac{b}{2\tilde{V}}(C_{lp}p + C_{lr}r) \tag{7}$$

$$C_m = \frac{c}{2\tilde{V}}C_{mq}q \tag{8}$$

The yaw moment contains additional terms to model the effect of asymmetric brake deflection. The sensitivity of the yaw moment to differential brake input is modified by the canopy incidence angle and the canopy angle of attack:

$$C_n = \frac{b}{2\tilde{V}}(C_{np}p + C_{nr}r) + C_{n\beta}\beta + (C_{nA} + C_{nAI}\delta_i + C_{nA\alpha}\alpha')\delta_A \tag{9}$$

The control actuators are modeled by filtering and rate limiting the control commands:

$$\dot{\delta}_{\text{actual}} = (\delta_{\text{command}} - \delta_{\text{actual}})/\tau_\delta \quad |\dot{\delta}_{\text{actual}}| \le \dot{\delta}_{\max} \tag{10}$$

The simulation model was generated to match flight test results for the parafoil and payload aircraft used for autonomous flight testing. The test vehicle is shown in-flight in Fig. 2. The test vehicle uses high-torque servos with extended arms to actuate the trailing edge brakes. The rear suspension line groups are attached directly to the payload. The front suspension line groups are attached to a winch servo which is used to provide incidence angle control. An electric motor allows the system to climb under its own power to a simulated release altitude, at which point the power to the motor is cut and the system glides back down. Only the gliding portion of the flight is modeled in simulation, so the ability of the payload to create thrust is not part of the simulation model. The test vehicle uses an airborne guidance unit (AGU) consisting of a GPS sensor, barometric altimeter, and flight computer.

The mass and geometry parameters for the payload and canopy are given in Table 1. The canopy is a rectangular planform, airdrop-style parafoil with an aspect ratio of 2.35.

The aerodynamic parameters used to model the parafoil canopy are given in Table 2.

The time constants and rate limits for the brake deflection actuators are 0.1 s and 25 cm/s, respectively. The time constant and rate limit for the incidence angle actuation are 0.5 s and 4 deg/s, respectively.

Simulation results are compared to flight test for a step input in incidence angle. The flight test results are derived from the onboard GPS receiver and barometric altimeter. The control input is shown in Fig. 3 and consists of pulling the nose of the canopy down 6 cm for a duration of 12 s. The AGU records only control commands, so the actual control deflections are estimated using the actuation models described above. To put this in context, the canopy collapses when the incidence angle is lowered beyond 12 deg. The airspeed and descent rate response of the simulation model are compared to flight test data in Fig. 3.

An especially important feature to capture in the simulation model is the difference in turn rate response created as the canopy incidence angle is changed. Figure. 4 shows the steady state turn rate vs differential brake deflection of the simulation model compared to simulation data at three different incidence angle settings.

Table 1 Mass and geometry parameters

Parameter	Value	Units
Total mass	2.2	kg
Span (b)	1.88	m
Chord (c)	0.8	m
Wing area	1.5	m^2
I_{XX}	1.68	$kg\,m^2$
I_{YY}	0.80	$kg\,m^2$
I_{ZZ}	0.32	$kg\,m^2$
I_{XZ}	0.09	$kg\,m^2$
A	0.05	kg
B	0.35	kg
C	1.85	kg
P	0.07	$kg\,m^2$
Q	0.06	$kg\,m^2$
R	0.05	$kg\,m^2$

Table 2 Aerodynamic parameters

Parameter	Value	Parameter	Value
C_{L0}	0.24	C_{mq}	−2.5
$C_{L\alpha}$	2.14	$C_{Y\beta}$	1.0
$C_{L\alpha 3}$	−1.53	$C_{n\beta}$	0.1
C_{D0}	0.12	C_{lp}	−0.02
$C_{D\alpha 2}$	0.33	C_{np}	0.0
$C_{L0\delta B}$	0.0	C_{lr}	0.0
$C_{L\alpha\delta B}$	0.39	C_{pr}	1.85
$C_{D0\delta B}$	0.043	$C_{n\alpha}$	−0.04-0.04
$C_{D\alpha 2\delta B}$	2.06	$C_{n\delta\alpha}$	−0.01
$\alpha_{0\delta B}$	0.11	$C_{n\delta I}$	-0.017−0.017

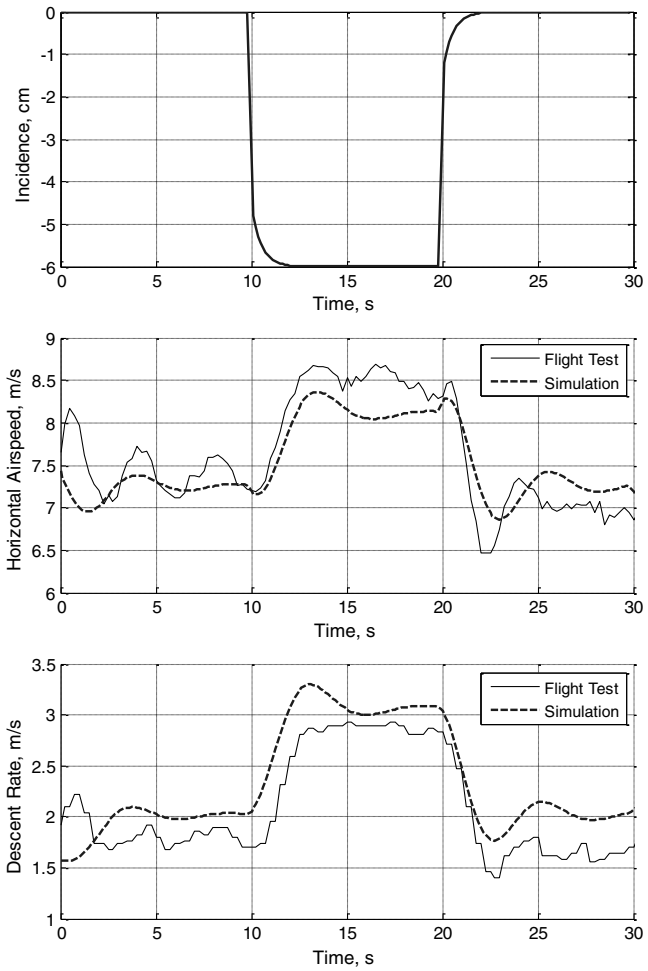

Fig. 3 Airspeed and descent rate response to incidence step.

III. GNC Algorithm for Autonomous Landing

This section describes the basic GNC algorithm used for autonomous landings. The basic algorithm shares many features with state-of-the-art guided parafoils discussed in the literature [10–25].

A. Guidance

The guidance algorithm splits the flight into four main phases: initialization, loiter, final approach, and landing. A simple terminal guidance algorithm for autonomous airdrop systems is an implementation of the T approach [20].

This is an approach commonly used by human sky divers and paragliders. The system loiters just downwind of the target by performing a series of figure eight turns. The turns are always made into the wind. This allows the system to enter the final approach trajectory quickly in case conditions change rapidly. Figure 5 shows a simulated flight trajectory starting near the end of the loitering phase.

1. *Initialization*

The purpose of this phase is to provide reasonable initial state estimates of wind and airspeed to the navigation algorithm. The simplest way to obtain these estimates is to hold a constant differential brake deflection long enough for the system to fly at least one complete circle. The length of the initialization phase and initial control setting are input as parameters to the GNC algorithm. This phase is entirely open loop; homing to the target area is accomplished using the loitering guidance logic.

2. *Loiter*

The loiter phase begins immediately after the initialization phase and consists of figure eight turns performed just downwind of the target. This is implemented by assigning homing targets at a specified distance downwind and a specified distance perpendicular to a line drawn straight downwind from the target. When the system reaches a specified radius from the target, the target is switched to the other side of the downwind line. The initial turn when the target is switched is always into the wind, resulting in a figure eight pattern. In very windy conditions it is desirable to prevent the system from flying too far downwind. This is accomplished by tilting the loiter targets into the wind so that the figure eight pattern begins close to the target and slowly drifts downwind to obtain the desired offset for final approach.

During loiter, the altitude required to reach the target from the current location is computed constantly. (Note: heavily filtered state estimates of airspeed, descent rate, and wind components are used for guidance planning and are denoted with a subscript F, while instantaneous state estimates are denoted with a subscript k. Also, descent rate is constrained to a reasonable minimum value for all guidance calculations.) The time remaining in the flight is:

$$T = h_k / \dot{z}_F \tag{11}$$

The effective distance from the target accounting for the wind is:

$$d = \sqrt{(x_k + V_{WX,F}T)^2 + (y_k + V_{WY,F}T)^2} \tag{12}$$

The altitude required to cover this distance is determined from the current estimate of the glide ratio. Some additional altitude is required to make the turn to the target, which is assumed to occur at a constant nominal turn rate. The sum of the altitude required to reach the target is the sum of the altitude used to turn to the target and the altitude used to glide to the target:

$$T_{\text{turn}} = |\chi_0 - \text{atan}\,2(-y_k, -x_k)|\dot{\chi}_{\text{nom}} \tag{13}$$

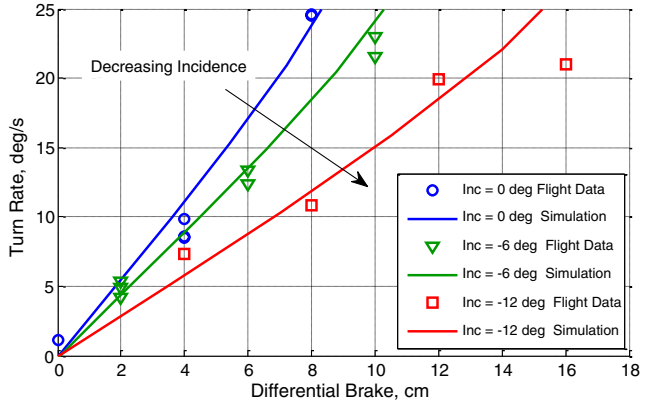

Fig. 4 Turn rate vs differential brake and incidence angle.

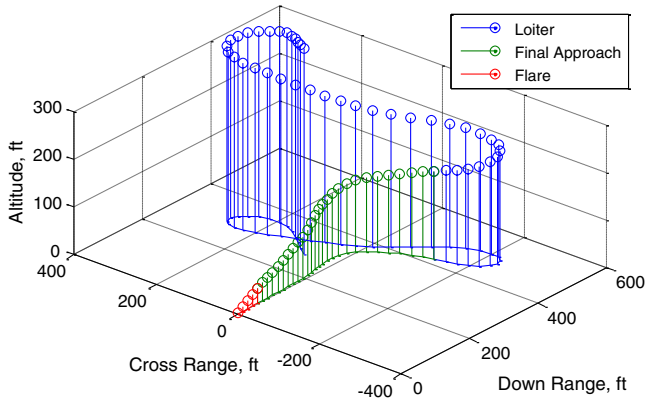

Fig. 5 Typical landing trajectory.

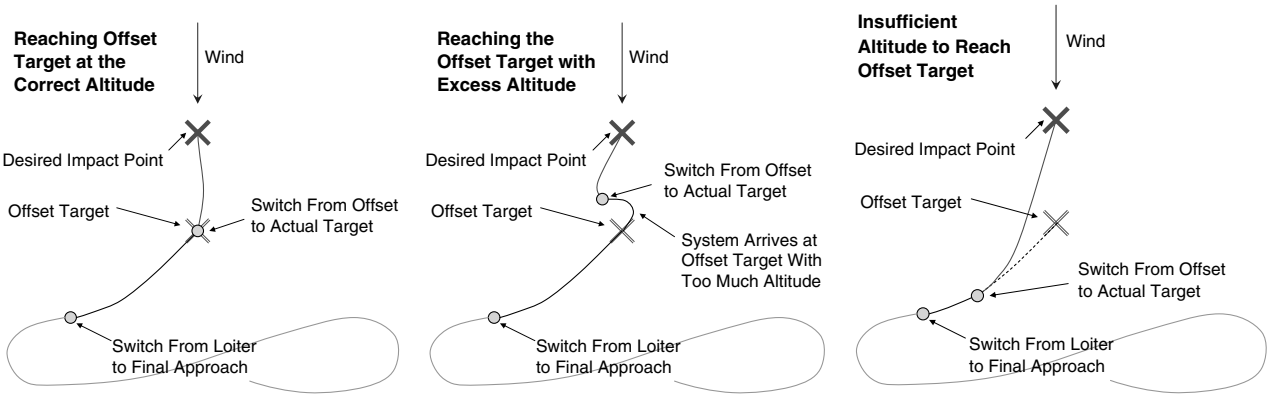

Fig. 6 Two stage final approach examples.

$$h_{\text{req}} = d\frac{\dot{z}_F}{V_0} + T_{\text{turn}}\dot{z}_F \tag{14}$$

The altitude margin is defined as the difference between the current altitude and the altitude required to reach the target. When the altitude margin falls below a specified value, the guidance algorithm switches from the loiter phase to the approach phase.

3. Approach

A two stage final approach is used where the system first tracks an offset target on its way to the desired impact point. This offset target is placed downwind of the desired impact point and the altitude of this offset target is set to lie just above the nominal glide path to the actual impact point. While homing to the offset target, the system computes the altitude margin for reaching the desired impact point. When the altitude margin for reaching the impact point reaches zero, the system begins homing to the actual target. Figure 6 shows some example approach trajectories demonstrating the logic of the two stage approach.

If the system reaches the offset target with the correct amount of altitude, it then flies to the actual target and lands into the wind. If the system reaches the offset target with excess altitude, it loiters over the offset target until the excess altitude is burned off. Finally, if the system runs out of altitude margin on the way to the offset target, the offset target is abandoned and the system flies straight at the desired impact point. This allows the approach trajectory to adapt to changes in the wind during final approach.

4. Landing

The goal of the landing maneuver is to minimize the kinetic energy of the system just prior to impact. This is accomplished by first releasing the trailing edge brakes to zero deflection and setting the incidence angle to the maximum nose-up setting (this is the setting corresponding to minimum airspeed with zero brakes), then applying full symmetric brake at the maximum rate (flare). The initiation of the brake release and the flare occur at specified altitudes.

B. Navigation

The term navigation is used here to refer to the task of estimating the parafoil aircraft states from the available sensor data. Sensor data is assumed to be a three-dimensional position and the three components of inertial velocity obtained from a single, commercial GPS receiver and a barometric altimeter. The barometric altimeter, when well-calibrated, provides an improved estimate of altitude and descent rate compared to the GPS receiver.

An extended Kalman filter [41] observer produces an estimate of the wind vector and system heading angle from GPS data by solving the vector diagram in Fig. 7. The ground track velocity vector is measured with GPS. The resolution of this ground track vector into the airspeed V_0 and wind vector V_w is not unique. When performing system identification offline, a series of vector diagrams at a variety of heading angles are solved simultaneously to obtain unique airspeed and wind estimates.

To ensure a stable, real-time wind vector estimate, the wind vector is estimated gradually over a series of measurements and the airspeed is not estimated as a state in the observer. A model to compute airspeed as a function of the control inputs is either loaded into the flight computer beforehand or estimated in-flight with the process described in the next section. Note that solution of the vector diagram does not directly yield an estimate of the parafoil heading angle ψ, but rather the azimuth angle χ_0. The two are related by the sideslip angle β, which is normally small for parafoil and payload aircraft.

Two implementations of extended Kalman filter observers for wind and heading estimation of airdrop systems are given in [30], the first using only GPS measurements and the second using both GPS and a heading sensor. The navigation algorithm used throughout the current work makes use of the GPS-only observer derived in [30]. The process noise variance for the wind estimates are set to $0.05\ (\text{m/s})^2$, the process noise variance for the heading rate is set to $0.005\ (\text{rad/s})^2$, and the measurement noise variance for the velocity measurements is set to $4\ (\text{m/s})^2$.

C. Control

Lateral control is provided by a model predictive controller tracking a commanded heading angle. The controller uses an internal model of the turn rate dynamics to determine an optimal set of control inputs given a set of heading commands. The goal of the model predictive controller is to determine a vector of control inputs that will minimize the error between the predicted output vector and a commanded output vector over a finite time horizon with a minimum amount of control effort. The controller assumes the turn rate response is related to the differential brake deflection by the first order linear model specified by Eqs. (15–17), where τ is the turn rate time constant, s is the control sensitivity, and Δt is the update interval for the controller:

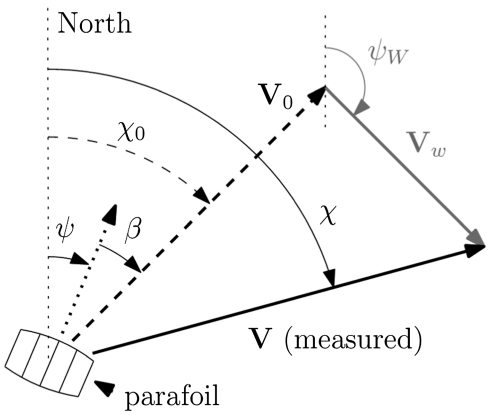

Fig. 7 Decomposing measured velocity vector.

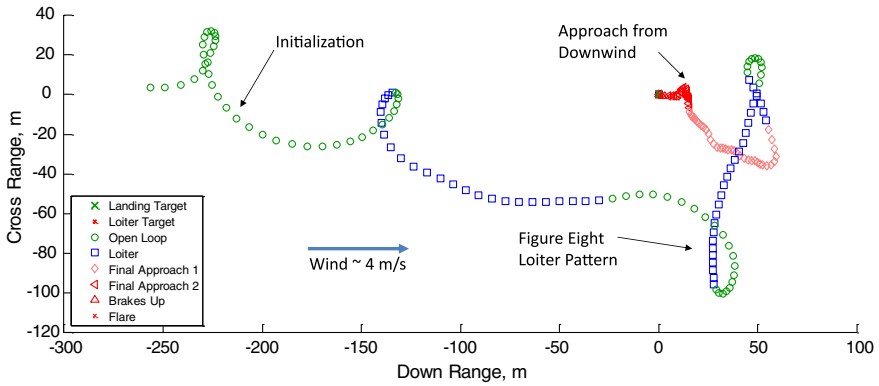

Fig. 8 Example simulated autonomous flight.

$$x_{k+1} = Ax_k + Bu_k, \qquad y_k = Cx_k \tag{15}$$

$$x_k = \begin{Bmatrix} \chi_k \\ \dot{\chi}_k \end{Bmatrix}, \qquad u_k = \delta_A, \qquad y_k = \chi_k \tag{16}$$

$$A = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 - \Delta t/\tau \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ s\Delta t/\tau \end{bmatrix}, \qquad C = [1 \quad 0] \tag{17}$$

The A, B, and C matrices describing the internal model are used to generate a commanded brake differential δ_a, given a current state estimate from navigation x_k, and a vector of heading commands Y_c, from guidance

$$\delta_{A,C} = k_1(Y_c - K_{CA}x_k) \tag{18}$$

where the terms k_I and KCA refer standard forms of the model predictive controller gain matrices [27,42].

Parafoil and payload aircraft typically exhibit a random turn bias so that zero differential brake input usually results in a nonzero turn rate. The internal model used by the controller is used to estimate this bias over time with a heavily damped filter:

$$\delta_{A,\text{BIAS}} = \delta_{A,\text{BIAS}} + K_I(\dot{\chi}_C - \dot{\chi})/s \tag{19}$$

This bias estimate is subtracted from the differential brake commands from the model predictive controller:

$$\delta_A = \delta_{A,C} + \delta_{A,\text{BIAS}} \tag{20}$$

D. Example Simulated Flight

A simulated flight trajectory from an altitude of 250 m is shown below. The average wind speed is 4 m/s from the north and the standard deviation of the vertical wind for the Dryden turbulence model was set to 0.6 m/s.

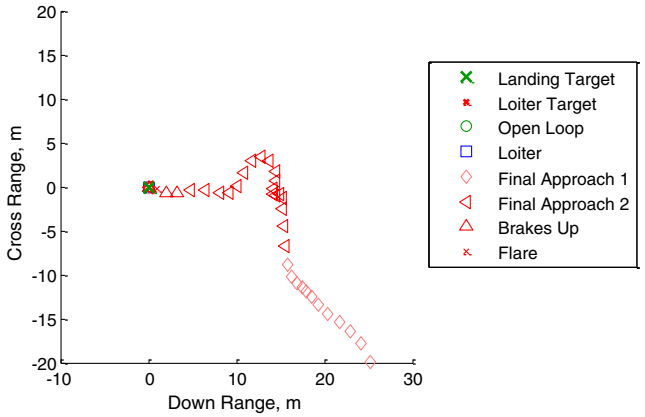

Fig. 9 Example simulated autonomous flight, zoomed in on landing.

IV. Glide Slope Controller

Longitudinal control is obtained with a nonlinear proportional control strategy based on the coupled use of canopy incidence angle and trailing brake deflection to track a commanded glide slope over ground. This is an entirely new concept for guided airdrop control and is one of the primary contributions of this paper.

A. Controller Development

The glide slope control strategy is based on a nonlinear proportional control law which can be divided into two parts, 1) generating a commanded glide slope based on the current state of the parafoil aircraft and the environment, and 2) determining the correct control input to achieve the commanded glide slope.

In discussing the formulation of the glide slope command logic, it is helpful to consider a quantity called the glide path to target

$$\text{GS}_T = d/z \tag{21}$$

where d is the distance downwind from the target and z is the current altitude above the target. The glide slope control strategy is to make a straight line final approach to the target from directly downwind. While on final approach, if the glide slope over ground of the parafoil and payload aircraft is equal to the glide path to target, the aircraft will intersect the target. If the glide path to target is steeper than the minimum glide slope over ground of the parafoil, the system will fly past the target, and if the glide path to target is shallower than the maximum glide slope over ground of the parafoil, the system will land short of the target. In this way, the minimum and maximum glide slopes over ground define the boundaries of the region from which the parafoil and payload aircraft will be able to reach the target. To maximize the ability of the system to reject any disturbances during final approach, the system should seek to maintain a nominal glide path in the center of this region. A normalized glide path error is defined as follows:

$$e_{\text{GS}} = \frac{\text{GS}_{\text{nom}} - \text{GS}_T}{(\text{GS}_{\text{max}} - \text{GS}_{\text{min}})/2} \tag{22}$$

When the glide path error is 0, the system is on an intercept course with the target on the nominal glide path, when the glide path error is 1, the system will hit the target with the controls set for minimum glide slope over ground, and when the glide path error is −1 the system will hit the target with the controls set for maximum glide slope over ground. To minimize control inputs near the nominal glide path, the glide slope commands are made proportional to the square of the glide path error. Finally, an additional parameter e_{SAT} is used define the magnitude of glide path error at which the controls saturate. The parameter e_{SAT} is set to a value less than one so that the controls will saturate prior to the system reaching the minimum and maximum glide slope boundaries. The resulting glide slope command logic

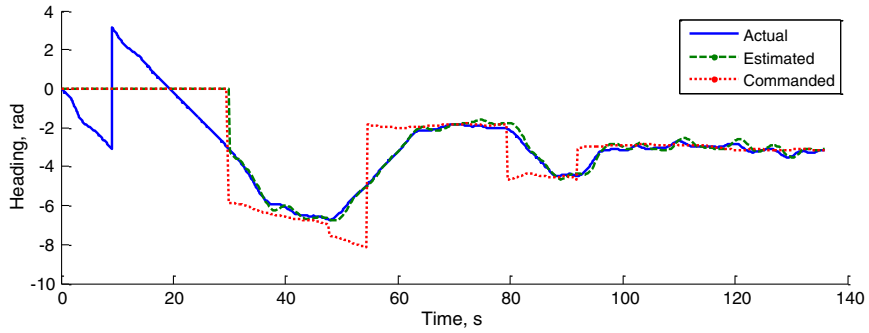

Fig. 10 Simulated heading angle estimation and tracking.

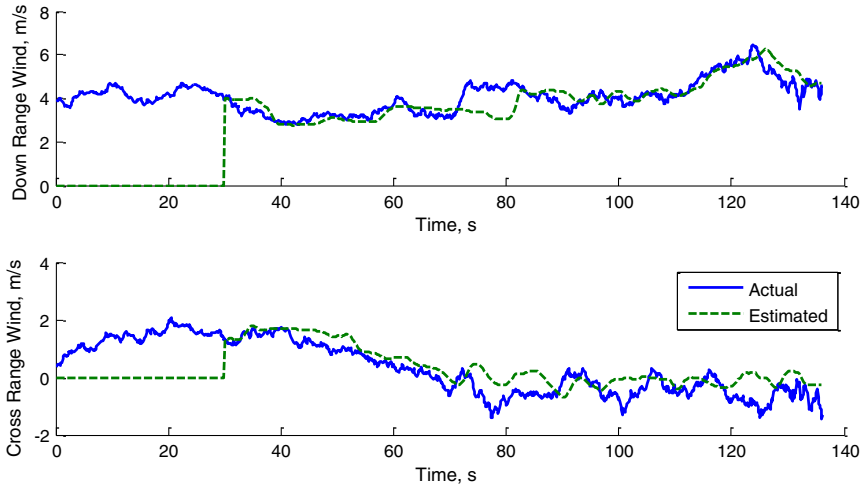

Fig. 11 Wind estimation in autonomous simulation.

given in Eqs. (21–24) provides a simple method for choosing a commanded glide slope GS_C based on the current glide path to target and the range of glide slope over ground that the parafoil can achieve in the current atmospheric conditions:

$$\tilde{e}_{GS} = \begin{cases} \min(e_{GS}/e_{SAT}, -1), & e_{GS} < 0 \\ \max(e_{GS}/e_{SAT}, 1), & e_{GS} > 0 \end{cases} \tag{23}$$

$$GS_C = \begin{cases} GS_{nom}(1 - \tilde{e}_{GS}^2) + \tilde{e}_{GS}^2 GS_{max}, & e_{GS} < 0 \\ GS_{nom}(1 - \tilde{e}_{GS}^2) + \tilde{e}_{GS}^2 GS_{min}, & e_{GS} > 0 \end{cases} \tag{24}$$

An example of the glide slope commands generated with this method is shown in Fig. 12. For this scenario, the minimum glide slope over ground is set at 1, the maximum is set at 3, and the normalized error at which the controls saturate, e_{SAT}, is set at 0.5. The plot shows how the commanded glide slope is generated to bring the system smoothly onto an intercept course with the target on the nominal glide path. If the system is outside the boundaries set by the e_{SAT} parameter, the maximum or minimum glide over ground is commanded to bring the system back toward the nominal glide path, and if the system is outside the minimum and maximum glide slope boundaries it will not be able to reach the target.

The second part of the glide slope control algorithm is the selection of appropriate control inputs to achieve the commanded glide slope. This is done by inverting the mapping of incidence angle and brake deflection to glide slope over ground. For convenience and computational efficiency, mappings of incidence angle and brake deflection to horizontal airspeed and descent rate are stored as polynomial functions of normalized brake and incidence angle inputs:

$$V_A = V_{A,0} + \partial V_I \tilde{\delta}_I + (\partial V_B + \partial V_{BI} \tilde{\delta}_I + \partial V_{BI2} \tilde{\delta}_I^2) \tilde{\delta}_B \tag{25}$$

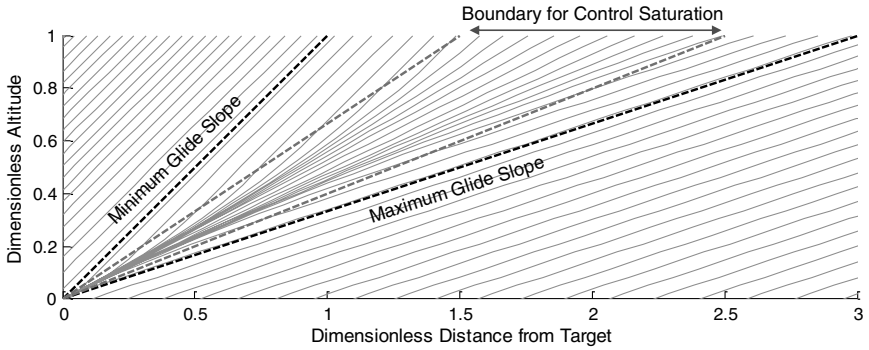

Fig. 12 Visualization of commanded glide slope logic.

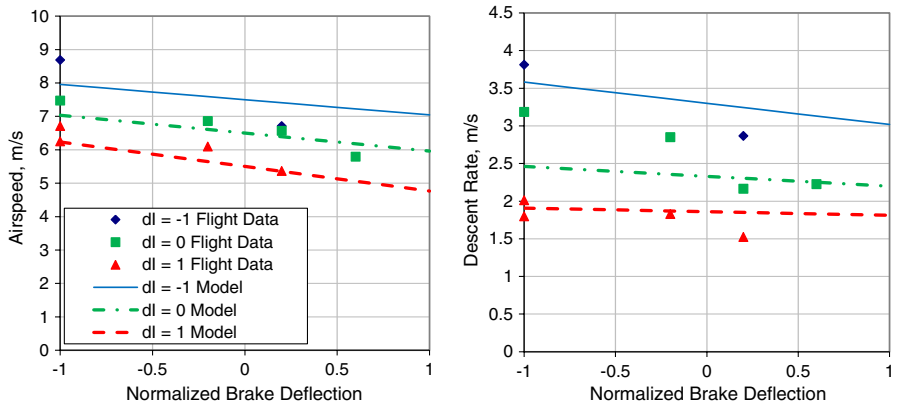

Fig. 13 Airspeed and descent rate models compared to flight test data.

$$\dot{z} = \dot{z}_0 + \partial\dot{z}_I\tilde{\delta}_I + \partial\dot{z}_{I2}\tilde{\delta}_I^2 + (\partial\dot{z}_B + \partial\dot{z}_{BI}\tilde{\delta}_I + \partial\dot{z}_{BI2}\tilde{\delta}_I^2)\tilde{\delta}_B \quad (26)$$

Brake deflection and incidence angle are normalized to span the range −1 to 1. In practice, the limits of incidence angle are actually a function of the level of brake deflection. The limit on incidence angle is assumed to be linear function of brake input:

$$\tilde{\delta}_B = 2(\delta_B - \delta_{B,\text{MIN}})/(\delta_{B,\text{MAX}} - \delta_{B,\text{MIN}}) - 1 \quad (27)$$

$$\delta_{I,\text{MIN}} = \delta_{I\text{MIN}0} + \delta_{I\text{MIN}B}\delta_B \quad (28)$$

$$\tilde{\delta}_I = 2(\delta_I - \delta_{I,\text{MIN}})/(\delta_{I,\text{MAX}} - \delta_{I,\text{MIN}}) - 1 \quad (29)$$

The speed over ground is determined by subtracting the wind speed from the forward airspeed estimate, and the glide slope over ground is determined as the ratio of the speed over ground to the descent rate:

$$\text{GS} = (V_A - V_W)/\dot{z} \quad (30)$$

Sample mappings were created to fit the flight test data vehicle used for autonomous landings. The models are compared to flight test data in Fig. 13, and the coefficients for the polynomial models are listed in Table 3.

Contours of constant glide slope over ground for these example mappings are shown vs incidence angle and brake deflection for four different wind levels in Fig. 14. These plots demonstrate how very different the effect of the control inputs on glide slope over ground can be in different wind conditions. As stated above, the effect of incidence angle on glide slope over ground can be completely opposite in different wind conditions. There is a complex interaction between incidence angle and trailing edge brake in determining the glide slope over ground in different wind conditions.

It is also clear from the plots in Fig. 14 that a wide range of control inputs can produce the same glide slope over ground. The approach taken here is to restrict the controls to lie on a line drawn on the glide slope mapping from the maximum glide point to the minimum glide point. This ensures that the full range of glide over ground is achieved, and the problem of inverting the nonlinear mapping to obtain the controls to achieve a given glide slope command is reduced to a line-search problem. The attraction of this approach is the simplicity of implementation and minimal computation time required.

Table 3 Airspeed and descent rate model parameters

Parameter	Value, m/s	Parameter	Value, m/s
V_0	6.5	$\dot{z}_0$	2.33
∂V_I	−1	$\partial\dot{z}_I$	−0.72
∂V_B	−0.536	$\partial\dot{z}_{I2}$	0.25
∂V_{BI}	−0.138	$\partial\dot{z}_B$	−0.132
∂V_{BI2}	−0.059	$\partial\dot{z}_{BI}$	0.117
		$\partial\dot{z}_{BI2}$	−0.033

These lines of optimal control inputs are plotted on top of the glide slope contours in Fig. 14. The line-search problem is solved with successive three point quadratic approximations. In fact, because the minimum and maximum glide slope configurations always lie on the boundaries of the mapping, these quantities can be determined using the same line search algorithm. Note from Fig. 14 that the control combination for minimum glide slope is minimum incidence angle and minimum brake deflection in no wind, but the control combination for minimum glide slope changes to maximum incidence angle and maximum brake deflection when the wind increases to 2 m/s. This could result in large oscillations in control input if the estimated wind speed is near the boundary where this switch occurs. To handle this, a hysteresis effect is imposed. If the wind estimate crosses this boundary, the control combinations corresponding to minimum glide slope are not updated immediately. The wind estimate must remain on one side of the boundary for more than 5 s prior to the control combination corresponding to minimum glide slope is updated.

Also note from Fig. 14 that the minimum glide slope in a 6 m/s wind is actually negative. A negative glide slope implies that the wind speed is greater than the airspeed, so the system will actually be traveling backward over ground when flying into the wind. The nominal airspeed of the system used here is 6.5 m/s. This can be increased to 9 m/s with the zero brake deflection and the incidence angle set to the maximum nose-down position, and the airspeed can be decreased to 5 m/s by applying full brake with the incidence angle in the maximum nose-up position. This allows the system to transition from positive to negative glide slopes in wind speeds exceeding 5 m/s. The GNC algorithms make no assumptions about the relative magnitude of the airspeed and wind speed, so no modifications are required to handle the case when the wind speed exceeds the airspeed.

B. Simulation of Controller on Final Approach

Simulation results of an example final approach from an altitude of 100 m are shown in Fig. 15. Two cases are shown; the first case uses the glide slope controller described above, and the second case has the controls fixed. The minimum, maximum, and nominal glide slope lines shown in Fig. 15 were determined based on the average winds during the final approach. The wind profiles used for the example approach are shown in Fig. 16. The average wind is 5 m/s and the standard deviation of the vertical wind component used for the Dryden turbulence model is 0.8 m/s. Note that this is a different wind profile from the one used for the autonomous landing simulation shown in Fig. 8. Initially, the system is nearly on the nominal glide path with the controls centered. Once the altitude drops below approximately 40 m, the head wind weakens and, to make matters worse, a positive vertical wind component picks up. As shown by the controls-fixed flight path, this combination of changes in the wind would normally cause the system to overshoot by nearly 40 m. As shown in Fig. 17, the glide slope controller reacts to this change in the

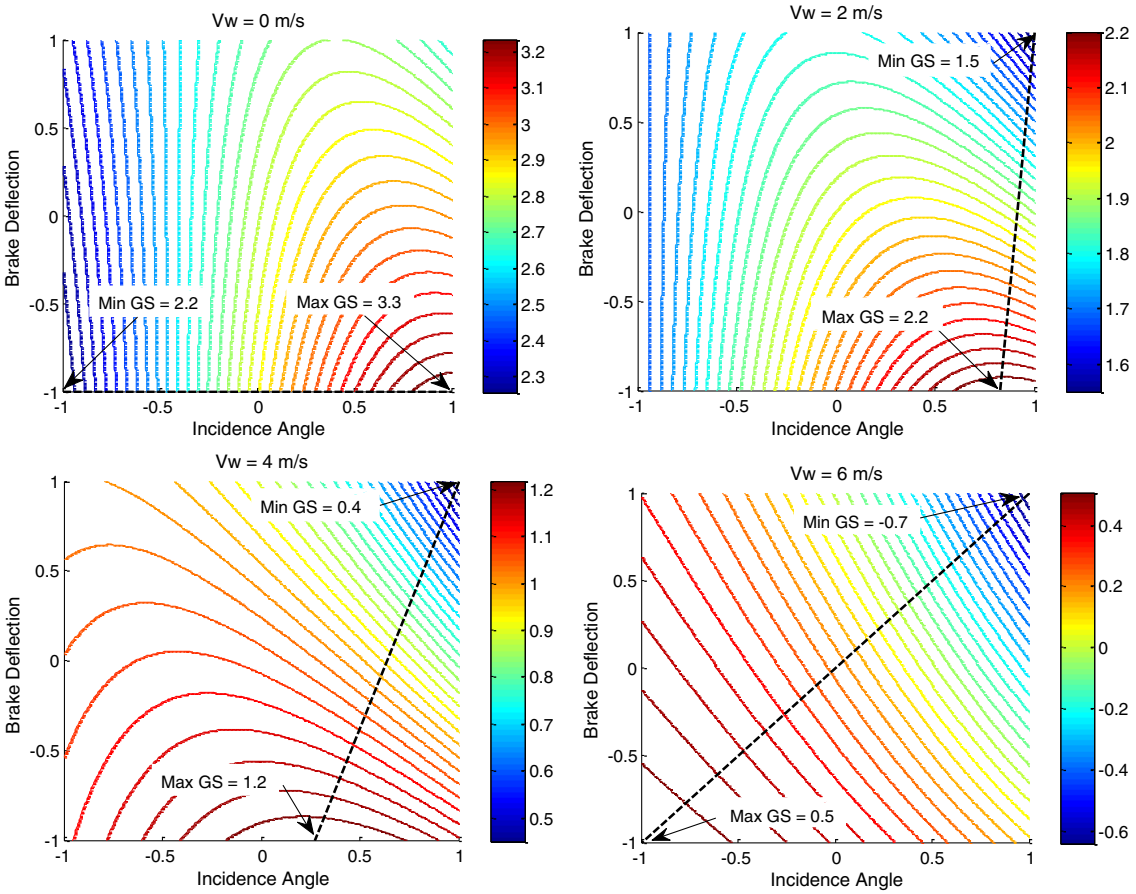

Fig. 14 Range of optimal control inputs at increasing wind levels.

wind by quickly applying a large amount of trailing edge brake and increasing the incidence angle to the maximum setting. This causes a large reduction in forward flight speed and a significant reduction in the glide slope over ground, allowing the system to stay on the nominal glide path to the target.

V. In-flight System Identification

The strategy used for implementing adaptive control in the present work is based on in-flight system identification. Periodically during the flight, an open-loop maneuver is performed, flight characteristics are identified from the observed response during the maneuver, the quality of the identified characteristics is evaluated, and if estimates meet the quality criteria, the control laws are updated with the new information. This is quite different from the typical adaptive control law where the control gains are updated continuously. The primary benefit of the approach proposed here is the ability to robustly extract accurate information about the flight characteristics using a minimal amount of low quality sensor information in spite of very large and frequent disturbances from atmospheric turbulence.

System identification relies on state estimates that are degraded by atmospheric turbulence and sensor errors. The traditional scenario when system identification is performed on the ground takes advantage of the ability to estimate parameters over multiple maneuvers and multiple flights so that the effects of sensor errors and turbulence are averaged out. A study of the benefits of in-flight system identification [30] showed that it is always beneficial to

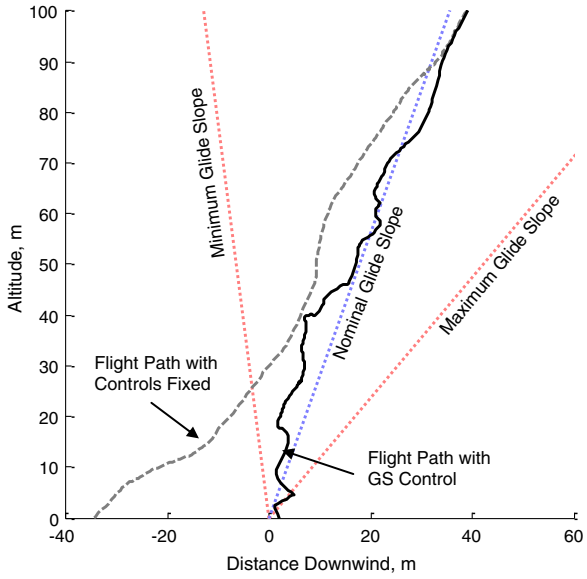

Fig. 15 Simulated final approach trajectories with and without glide slope control.

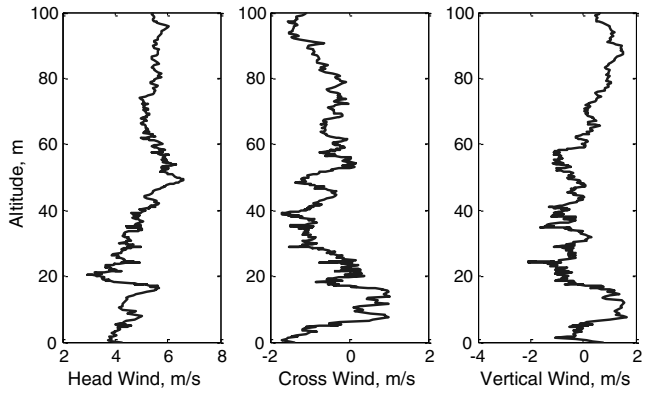

Fig. 16 Wind profiles for example approach simulation.

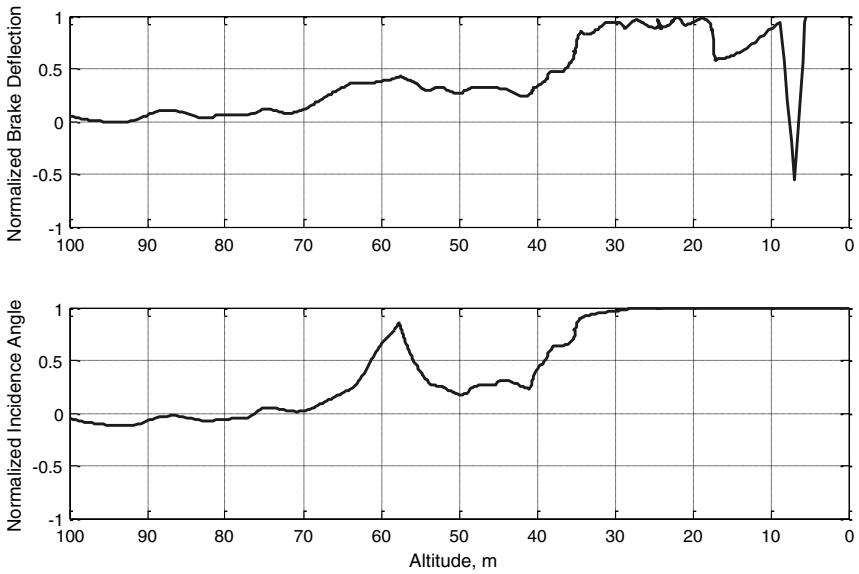

Fig. 17 Control inputs during example approach simulation.

estimate steady state quantities in-flight with reasonable levels of model uncertainty, sensor noise, and wind because the estimation is performed by averaging over a series of measurements. On the other hand, transient characteristics, estimated in-flight over small windows of data during maneuvers are much more sensitive to sensor noise and turbulence, and the results show that it is not always beneficial to estimate these transient characteristics in-flight. In other words, for transient characteristics, the degradation in the quality of in-flight estimates from sensor noise and turbulence is comparable to the degradation in the quality of fixed estimates from model uncertainty. It is important to note that this result is dependent on the nature of the air vehicle. Airdrop systems typically have very benign flight dynamics, so precise control of the transient response is not required. For a vehicle with very lightly damped or unstable flight dynamics, precise transient control and, hence, accurate knowledge of the transient characteristics would be critical.

In-flight system identification is integrated into the loiter phase of the autonomous flight. The loiter phase consists of a series of figure eight turns created by tracking alternating loiter targets. Each time a loiter target is reached, a turn toward the next target is initiated. Reliable estimates of steady state characteristics are obtained by holding a gentle turn with constant control deflection, so this turning phase after a loiter target is reached is a convenient time to estimate steady state flight characteristics. Normally, the system turns until it is facing the next target and then begins homing to it, but while estimating quantities in-flight, the constant control segment is held until a criterion for a valid airspeed estimate (described below) is met. Pairs of left and right turning segments are completed at various incidence angle and brake levels to build an internal mapping of the controls to airspeed, descent rate, and turn rate. A long loiter phase allows a greater opportunity to update the internal model. Conversely, if there is little or no time for loiter, then the internal model is not updated in-flight.

The first step in the system identification process is wind estimation. The same method used to estimate wind and airspeed developed specifically for offline system identification for parafoil and payload aircraft [40] is used for the in-flight system identification of the steady state flight characteristics. An upper bound on the airspeed estimation error for a segment of circling flight with constant control deflection was derived as the following

$$\frac{\sigma(\hat{V}_0)}{\sigma(V_G)} < \frac{1}{\sin{(\delta/4)^2}}, \qquad \delta = \min(|\Delta\psi|, 2\pi) \tag{31}$$

where $\sigma(\hat{V}_0)$ is the upper bound on the standard deviation of the airspeed estimation error, $\sigma(V_G)$ is the standard deviation of the measured ground speed when flying in a straight line (note that this term is meant to account for both disturbances from turbulence and measurement noise), and $|\Delta\psi|$ is the net change in heading angle spanned by the flight segment. The right hand side of Eq. (28) provides a convenient metric for evaluating the quality of an airspeed estimate obtained in-flight. If the result is one, it indicates the system has completed at least one complete circle. This should always be true for the initialization phase of the autonomous flight. When performing a series of estimates during the loiter phase of the flight, the result from Eq. (28) is evaluated at each time step while recording data for a constant control segment. When the value falls below a threshold, it indicates that a valid airspeed estimate has been obtained, the constant control segment is ended, and the system returns to the normal loiter routine.

The remaining flight characteristics of interest are the descent rate and the heading rate. The descent rate can simply be averaged over the constant control segment and the standard deviation of the descent rate can be used as a metric to evaluate the quality of the estimate from a particular flight segment. One method for estimating the heading rate is to use the wind vector obtained from the linear regression to go back and solve the vector diagram in Fig. 7 for every measurement of the constant control segment. This provides a series of heading angles which can then be differentiated to obtain heading rate. An alternative way to obtain heading rate is based on convenient approximation derived by Calise and Preston [43],

$$\dot{\psi} \approx \dot{\chi} V_G / V_0 \tag{32}$$

which allows heading rate to be determined from the ground speed, course rate (which can be obtained by differentiating the GPS ground course), and airspeed. Because the airspeed is a constant during the flight segment, the numerator in the heading angle approximation can be obtained at each measurement update and the airspeed can be divided out at the end of the flight segment. This approximation allows an efficient computation of the average heading rate over the constant control segment. Although this approximation can produce significant errors for an instantaneous estimate, the error when using this approximation to obtain an average heading rate over a large number of measurements with significant heading variation is quite small. In simulation, heading rates obtained with this approximation agreed with the heading rates calculated with the exact method described above to within 1%. The quality of a heading rate estimate is evaluated based on the standard deviation in heading rate over the flight segment.

VI. Performance of Adaptive Glide Slope Control

All autonomous landing results are based around the flight-test vehicle shown in Fig. 2. The vehicle is a hand-launched, all electric powered parafoil. The use of a self-powered payload allows very rapid flight testing without the need for a drop plane and parachute packing in between flights. The canopy used is a miniature airdrop style canopy which shares many design features with the canopies used on full-scale systems. The typical flight profile begins with a rapid climb under power using manual control. At this point, power to the motor is shut off and the system enters gliding flight. Control of the payload is handed off to the autonomous GNC algorithm running on the onboard pilot, and the system is allowed to fly itself all the way to the ground. The flight log is downloaded from the autopilot after each flight. Results from a total of 51 autonomous flights are presented.

A. Performance of Glide Slope Control in Simulation

Predicted landing accuracy for the baseline case of autonomous landing using only lateral control is shown in Fig. 18 for different wind speeds and levels of turbulence. The turbulence level is the standard deviation of the vertical wind component used in the Dryden gust model. This plot was generated by running 50 autonomous landing simulations with average wind speeds of 0, 2, 4, 5, and 6 m/s at turbulence levels of 0.0, 0.2, 0.4, 0.6, 0.8, and 1.0 m/s for a total of 1500 simulated landings. In general, increasing the level of turbulence increases the average miss distance. This is the expected trend for any autonomous airdrop landing algorithm. The plot shows that when using only lateral control, the landing accuracy becomes extremely sensitive to turbulence at high wind speeds. This is because the wind speed is approaching the nominal forward airspeed of 6 m/s for this particular system. A strong wind gust when the average wind is near the forward airspeed can push the system downwind of the target. If the gust persists, there is no way for the system to make progress upwind toward the target. Another interesting trend in this plot is the increase in average miss distance when the wind speed is near 2 m/s in turbulent conditions. The reason for this increase has to do with changes in the wind direction. In strong winds, even strong wind gusts will not produce dramatic changes in the wind direction. However, in light winds, even a small wind gust can result in a large change in the wind direction. The result is that in light, gusty conditions, it is not uncommon for the system to set up on final approach into the current wind direction, only to have the wind shift dramatically so that the system is landing across the wind or even downwind. As the mean wind speed increases, this occurs less frequently, and the system is almost always facing into the wind during the entire approach.

Figure 19 shows the simulated landing accuracy for autonomous landings with the glide slope controller. The same set of mean wind and turbulence combinations described above were run for a total of 1500 simulated landings. The average miss distance is significantly lower in most combinations of wind and turbulence. In zero wind with no turbulence, the average miss distances are essentially the same because there are no disturbances for the glide slope controller to correct. Although increasing turbulence does tend to increase the average miss distance, increasing the average wind speed actually tends to decrease the average miss distance in all but the most

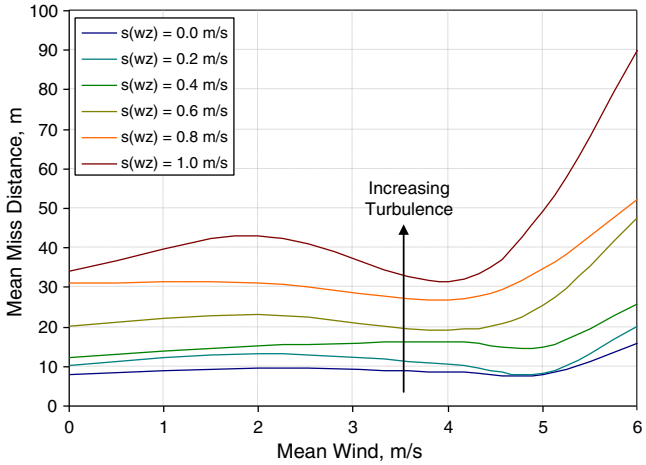

Fig. 18 Predicted landing accuracy vs wind speed using only lateral control.

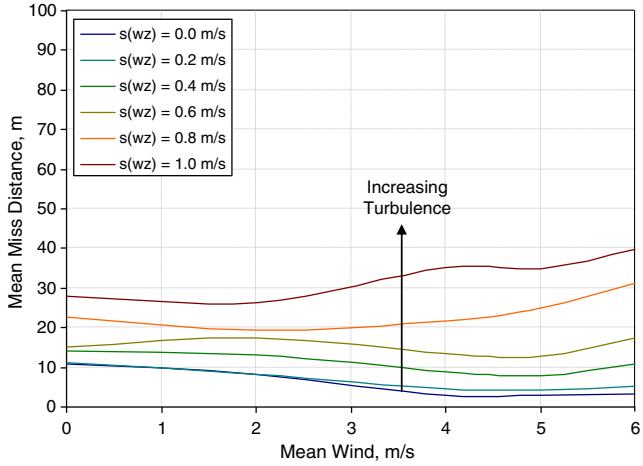

Fig. 19 Predicted landing accuracy vs wind speed using glide slope control.

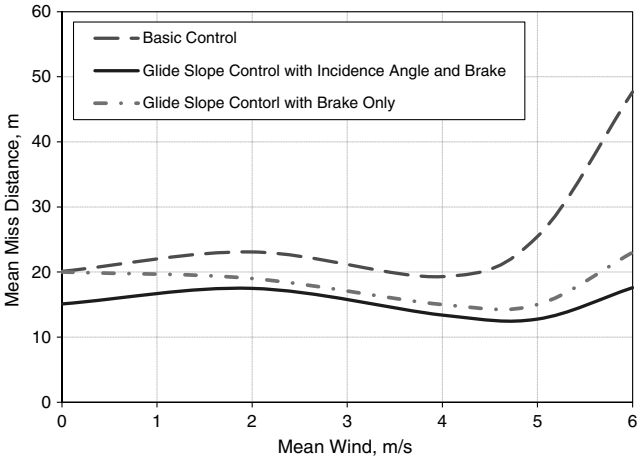

Fig. 20 Comparing predicted landing accuracy with basic control and with glide slope control.

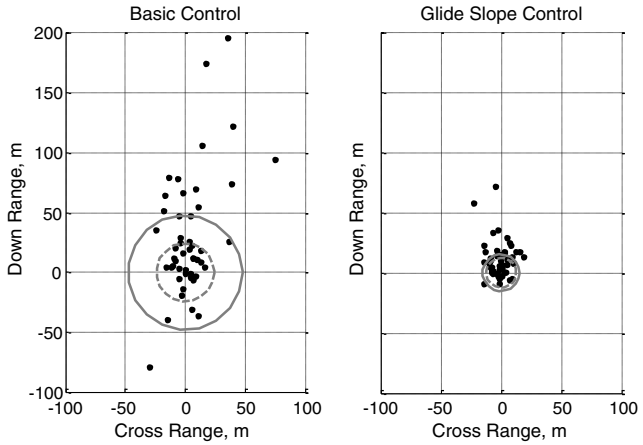

Fig. 21 Comparing simulated landing accuracy of basic and glide slope control algorithms with average wind of 6 m/s.

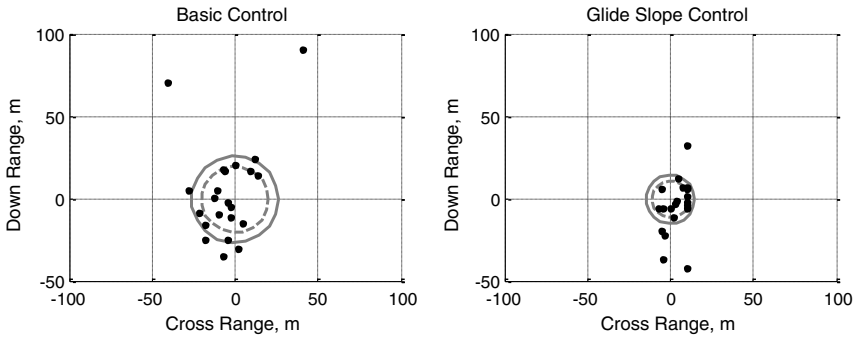

Fig. 22 Landing dispersions from autonomous flight tests.

turbulent cases. This is because the control authority of the glide slope controller is increased in windy conditions due primarily to the increasing effect of trailing edge brakes on glide slope over ground as the wind speed increases.

Figure 20 shows the landing accuracy using the basic control algorithm compared to landing with glide slope control at different average wind levels. Two cases with glide slope control are shown, the first using incidence angle and brake deflection, the second using only brake deflection. The turbulence level for this plot was set at 0.6 m/s. In light winds the addition of brake deflection provides no improvement, although glide slope control with incidence angle and brake deflection provides a 30% improvement in landing accuracy.

As the average wind increases to 5 m/s, the glide slope controller with incidence angle and brake improves landing accuracy by a factor of two, and as the wind increase to 6 m/s, the glide slope controller with incidence angle and brake improves accuracy by a factor of three. In moderate wind conditions, the glide slope controller using only brake deflection provides nearly the same performance as the combined incidence angle and brake controller. This is a result of the inverse relationship between aerodynamic glide ratio and airspeed when using incidence angle. In very high winds, the addition of incidence angle begins to again provide some additional benefit compared to the brake only control due to the ability of incidence angle to provide large increases in airspeed. The important thing to take away from these plots is that the performance of any autonomous landing algorithm can be heavily dependent on the average wind and especially the level of turbulence. In calm conditions there is little improvement in landing accuracy when using glide slope control because there is essentially nothing for the controller to do. It is in windy conditions when the addition of glide control begins to produce dramatic effects due largely to the wide range of forward airspeed that can be achieved by varying incidence angle in concert with the trailing edge brakes.

Figure 21 compares the simulated landing dispersion of the basic and glide slope control algorithms with a turbulence level of 0.6 m/s and an average wind speed of 6 m/s. Positive downrange indicates downwind, so the wind is blowing from the bottom of these plots and the system is approaching the target from positive downrange. Two circles are drawn around the target, the solid circle has a radius equal to the average miss distance and the dashed circle has a radius equal to the median miss distance. The mean and median miss distances for the basic control algorithm are very different due to the very large miss distances when the system is blown downwind. By comparison, the glide slope control algorithm is able to prevent exceedingly large miss distances in the same condition due to the ability to change airspeed to counter the strong, varying winds.

B. Performance of Glide Slope Control in Flight Test

The basic control algorithm and the glide slope control algorithm were tested in a series of autonomous flight tests. A total of 22 flights were performed with the basic control algorithm and 20 flights were performed with the glide slope control algorithm. The control algorithm was switched between basic control and glide slope control between each flight to ensure that the testing for the two control algorithms took place in similar weather conditions. The landing

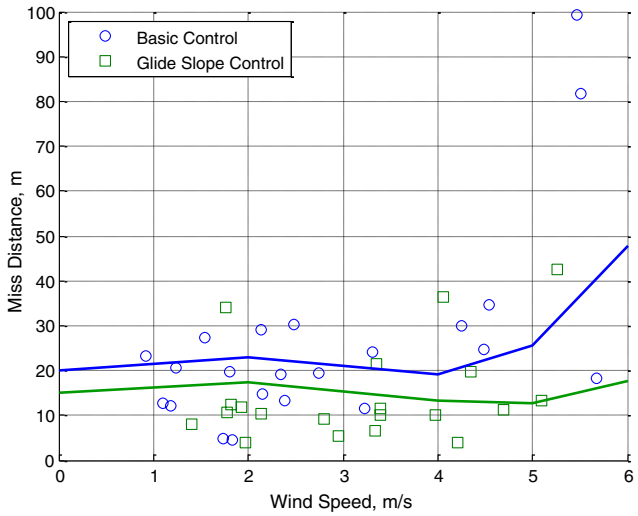

Fig. 23 Comparing flight test results of landing accuracy vs wind speed with and without glide slope control.

dispersions for each control algorithm are shown in Fig. 22. Note: all reported impact points are determined from the onboard GPS position reading at the instant the system touches down. The miss distances are plotted vs average wind speed in Fig. 23 with the average miss distances from simulation plotted as solid lines. The median miss distance, often referred to in the airdrop community as the circular error probable, and mean miss distance predictions from simulation are compared with statistics computed from the flight test data in Table 4. The flight test results agree quite well with simulation. The landings with glide slope control are generally closer to the target. Furthermore, at high wind speeds there is a rapid increase in miss distance with the basic control algorithm, although no such increase occurs with the glide slope control algorithm.

The excellent agreement between the accuracy predictions from simulation and the actual landing accuracy observed in flight test means that a high degree of confidence can be placed in trends in landing accuracy predicted with the simulation model. This allows a number of interesting trade studies to be run very quickly which would otherwise entail weeks of flight testing.

C. Performance of Adaptive Glide Slope Control in Simulation and Flight Test

One of the primary difficulties of in-flight estimation of parafoil flight characteristics is the sensitivity to turbulence. A simulation study was performed to determine the influence of the turbulence

Table 4 Simulated and actual landing accuracy statistics

	Simulation		Flight test	
	CEP, m	Mean miss, m	CEP, m	Mean miss, m
Basic control	19.1 s	27.2	20.1	26.2
Glide slope	12.3	15.5	10.9	14.7

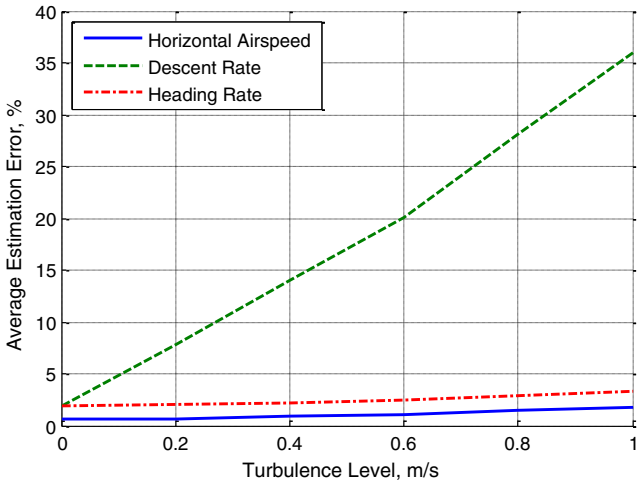

Fig. 24 In-flight estimation error vs turbulence level.

level on the error in the in-flight estimates of three flight characteristics of interest, horizontal airspeed, descent rate, and heading rate. Fifty cases were run with an average wind of 3 m/s and turbulence levels of 0, 0.2, 0.4, 0.6, 0.8, and 1.0 m/s, for a total of 300 cases. The airspeed, descent rate, and heading rate values were compared to the known values from the simulation model. The average errors in each estimated quantity are plotted against the turbulence level in Fig. 24. It is clear that reliable estimation of descent rate is extremely difficult to obtain in-flight. Horizontal airspeed and heading rate, on the other hand, can be estimated very reliably in-flight even in very turbulent conditions. The reason for this is that the horizontal winds can be exposed by maintaining a gentle turn, allowing the airspeed to be distinguished from the wind vector, and heading rate can only by perturbed temporarily as the wind is changing in magnitude. In contrast, large slowly varying disturbances to descent rate can be caused by the vertical wind component, and there is no way to distinguish the wind component from the aerodynamic descent rate when only a measurement of the absolute descent rate is available.

Although it is not possible to separate the vertical wind component from the aerodynamic descent rate, it is possible to observe the presence of a significant vertical wind component. During the constant control segments used to produce the airspeed, descent rate, and heading rate estimates, any variation in descent rate from the mean value can only result from a vertical wind component. This means that while the average descent rate over a constant control segment contains contributions from both the vertical wind and the aerodynamic descent rate, the standard deviation of the descent rate over the segment contains only contributions from the vertical wind. A strong correlation was found between the average in-flight estimation error of descent rate and the confidence interval in descent rate over the segment, where the confidence interval is approximated as twice the standard deviation over the square root of the number measurement sample. The strong correlation between the two quantities means that a reliable prediction of the quality of a descent rate measurement obtained in-flight can be obtained based on the computed standard deviation. For the remaining simulation results, a threshold of 0.155 m/s is set for the confidence interval on standard deviation, which corresponds to an average estimation error of 10%. If the confidence interval on the descent rate during a constant control segment exceeds this threshold, the descent rate data obtained for this segment is discarded.

Six flights were conducted with the adaptive glide slope control algorithm with an average miss distance of 15 m. Conditions were

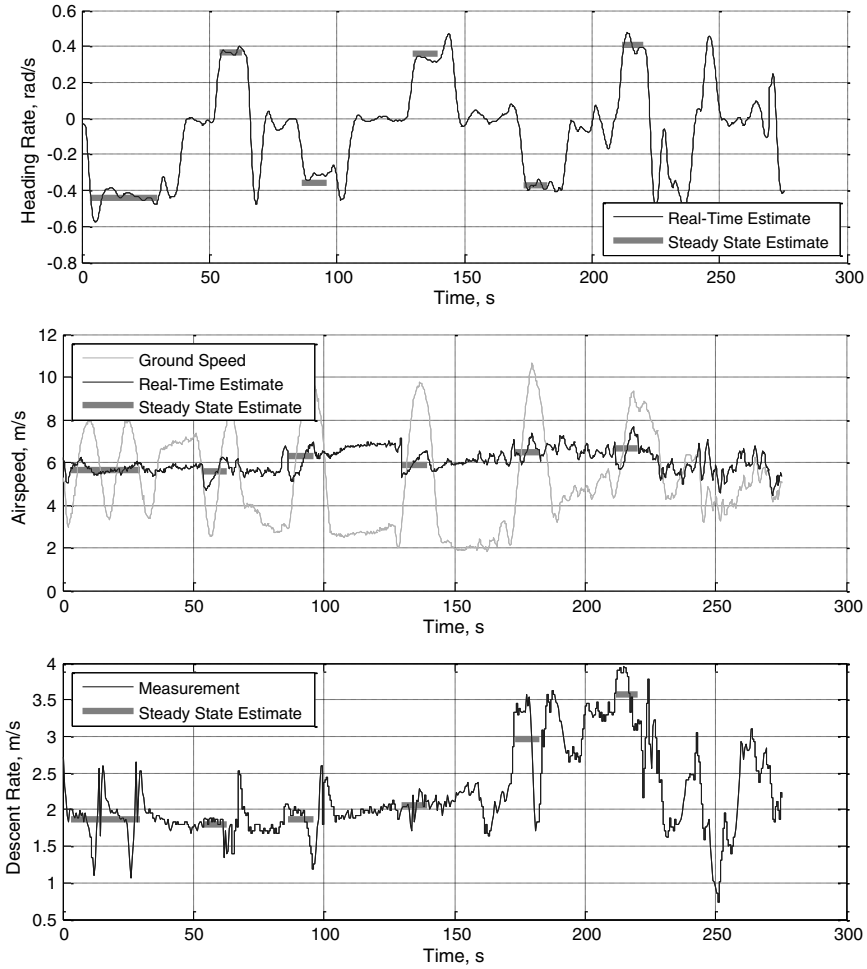

Fig. 25 In-flight heading rate, airspeed, and descent rate estimates.

Table 5 In-flight airspeed, descent, and turn rate estimates

Incidence, deg	Brake deflection, cm	Airspeed, m/s	Descent rate, m/s	Turn rate, rad/s	Airspeed confidence	Descent rate confidence, m/s	Turn rate confidence, rad/s
0	8	5.64	1.87	−0.44	1	0.061	0.013
0	8	5.60	1.79	0.37	1.5	0.049	0.024
0	0	6.29	1.87	−0.36	1.5	0.067	0.024
0	0	5.91	2.06	0.36	1.5	0.043	0.023
−6	8	6.48	2.97	−0.37	1.5	0.189	0.026
−6	8	6.67	3.58	0.41	1.5	0.104	0.031

significantly more turbulent on this testing day compared to the flights shown in Fig. 22. Three flights performed with the nonadaptive glide slope control algorithm on the same, turbulent day as the adaptive algorithm had miss distances of 36, 34, and 8 m. The concern with using an adaptive algorithm is that in turbulent conditions, the in-flight estimates may be very poor, and if the adaptive algorithm is not sufficiently robust, landing results could actually be worse with the adaptive algorithm. Although these six adaptive flights are insufficient to compute statistically relevant landing accuracy numbers, it is clear from the excellent landing accuracy obtained in turbulent conditions that the adaptive control algorithm works properly in real flight conditions.

The heading rate, airspeed, and descent rate estimates from one of the adaptive flights are shown in Fig. 25. The steady state estimates obtained from the in-flight system identification algorithm for each segment of constant control are plotted on top of the real-time estimates from the navigation algorithm. The large variation in descent rate during the constant control segments apparent in Fig. 25 are evidence of the turbulent atmospheric conditions. There are six segments in total, each consisting of a left and right turn at a specified combination of incidence angle brake deflection. The steady state estimates for each segment are given in Table 5. The airspeed confidence is computed based on Eq. (28), while the descent rate and turn rate confidence values are twice the standard deviation over the square root of the number of measurements.

The airspeed confidence level is 1 for the initialization segment because this segment covers more than a complete circle, and the confidence level 1.5 for remaining segments because this was the threshold used to define the minimum length of each segment. The descent rate confidence interval threshold was set to 0.65 m/s for the flight tests, so the descent rate estimates for segments three, five and six are considered invalid.

Figure 26 shows how the internal models of turn rate, airspeed and descent rate are updated from the initial guesses to match the estimates obtained in-flight. The updated models demonstrate the successful adaptation of the internal flight-dynamic model to match in-flight observations in actual, turbulent flight conditions.

To explore the benefits of in-flight system identification a series of autonomous landing simulations were performed with different levels of model uncertainty and different levels of adaptation. All of the flights included the six segment in-flight system identification and mapping update procedure used for the flight tests described above. To model uncertainty in a reliable and easily quantifiable way, perturbations are made to the internal model used by the autonomous landing algorithm rather than the parafoil simulation. Perturbations were made to the internal airspeed, descent rate, turn rate, and turn bias models. The perturbations to airspeed, descent rate, and turn rate are made as scalers, while the turn bias perturbations are absolute values. The forms of the perturbations are given in Eq. (30), where the perturbations are labeled p, the subscript INT indicates the baseline internal model, and the subscript GNC indicates the perturbed value used by the GNC algorithm:

$$V_A = p_{0,V}V_0 + p_{I,V}\partial V_I\tilde{\delta}_I + (p_{B,V}\partial V_B + \partial V_{BI}\tilde{\delta}_I + \partial V_{BI2}\tilde{\delta}_I^2)\tilde{\delta}_B \tag{33}$$

$$\dot{z} = p_{0,Z}\dot{z}_0 + p_{I,Z}\partial\dot{z}_I\tilde{\delta}_I + \partial\dot{z}_{I2}\tilde{\delta}_I^2 + (p_{B,Z}\partial\dot{z}_B + \partial\dot{z}_{BI}\tilde{\delta}_I + \partial\dot{z}_{BI2}\tilde{\delta}_I^2)\tilde{\delta}_B \tag{34}$$

$$\frac{\partial\dot{\chi}}{\partial\delta_{A\,\mathrm{GNC}}} = \frac{\partial\dot{\chi}}{\partial\delta_{A\,\mathrm{INT}}}(p_{0,TR} + p_{I,TR}\tilde{\delta}_I + p_{B,TR}\tilde{\delta}_B) \tag{35}$$

$$\delta_{A,\mathrm{BIAS,GNC}} = p_{0,\mathrm{DAB}} + p_{I,\mathrm{DAB}}\tilde{\delta}_I + p_{B,\mathrm{DAB}}\tilde{\delta}_B) \tag{36}$$

The ranges of the perturbations are given in Table 6. The range for the bias term is normalized by the maximum differential brake deflection, so all the ranges are nondimensional. The ranges were chosen to represent a typical level of uncertainty when flying a parafoil and payload aircraft for the first time. A model uncertainty of zero corresponds to no perturbation of the internal model, while a model uncertainty of one indicates that the majority of the perturbations will be at the maximum values. This is done by selecting the perturbations randomly from bounded normal distribution scaled according to achieve the desired uncertainty level.

Fifty simulated landings were performed at model uncertainty levels of 0, 0.25, 0.5, 0.75 and 1.0. The mean wind was set to 3 m/s and the turbulence level set to 0.5 m/s for all cases. The set of simulations was repeated for a basic case using glide slope control with no adaptation of the internal model. In the second case, the internal model of turn rate and control bias was adjusted to match in-flight measurements. In the third case, the internal turn rate and horizontal airspeed models were adjusted to match flight data. Finally, for the fourth case, all of the internal models were adjusted based on flight test data. The average landing accuracy vs model uncertainty for each case is shown in Fig. 29.

For the first case with no adaptation of the in-flight model, the average miss distance is increased dramatically as model uncertainty

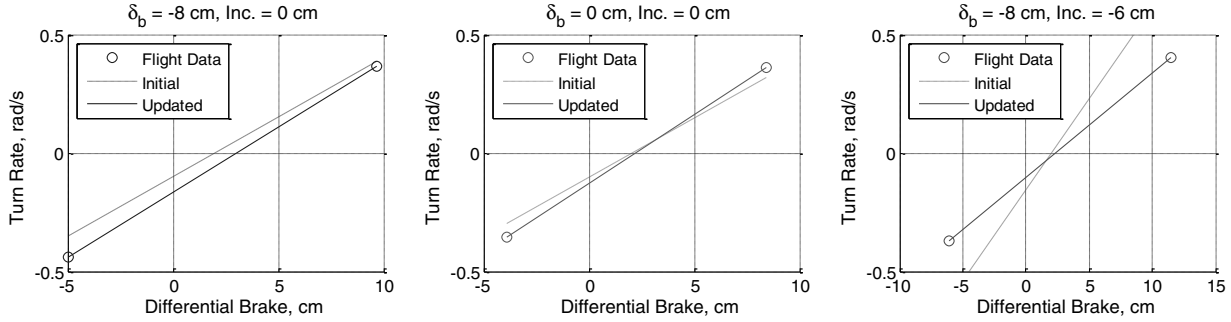

Fig. 26 Initial and updated turn rate model compared to flight data.

Table 6 Ranges of model perturbations

	V	$\dot{z}$	TR	δa_{BIAS}
p_0	(0.5,1.5)	(0.5,1.5)	(0.5,1.5)	(−.3,.3)
p_B	(0,2)	(0,2)	(−0.5, 0.5)	(−.3,.3)
p_I	(0,2)	(0,2)	(−0.5, 0.5)	(−.3,.3)

is increased. The landing accuracy with no model uncertainty is 13 m, while the average landing accuracy with the model uncertainty at the maximum level is 50 m, an increase of a factor of four. This is to be expected because the internal model is diverging from the actual flight characteristics of the system. For the second case, with the turn rate and control bias models adjusted to match in-flight estimates, the landing accuracy is somewhat improved. The average landing accuracy at the maximum level of model uncertainty in this case is reduced to 40 m, a 20% improvement over the case with no adaptation. When both turn rate characteristics and airspeed model are adapted to match in-flight estimates, the landing error is reduced dramatically. For model errors up to 25% there is essentially no degradation in landing accuracy, and with model errors at maximum the landing error has only risen to 20 m. For the final case, where the internal turn rate, airspeed, and descent rate models are all set to match flight test data, there is no improvement in landing accuracy compared to the case where descent rate is not estimated in-flight. In fact, with low levels of model uncertainty, the landing accuracy is actually slightly degraded when the descent rate is estimated in-flight. This indicates that descent rate estimates corrupted by vertical wind were still used to update the internal model, despite the threshold placed on standard deviation. This means that the descent rate of a parafoil and payload system can only be reliably estimated in very calm conditions.

To summarize, in-flight estimation of the turn rate and airspeed characteristics of a parafoil and payload aircraft can be reliably estimated in-flight with only GPS and a barometric altimeter during the descent to target. The use of these in-flight estimates to update the internal turn rate and airspeed models used for GNC calculations can dramatically increase the landing accuracy in situations where there is a significant amount of uncertainty in the aircraft flight characteristics.

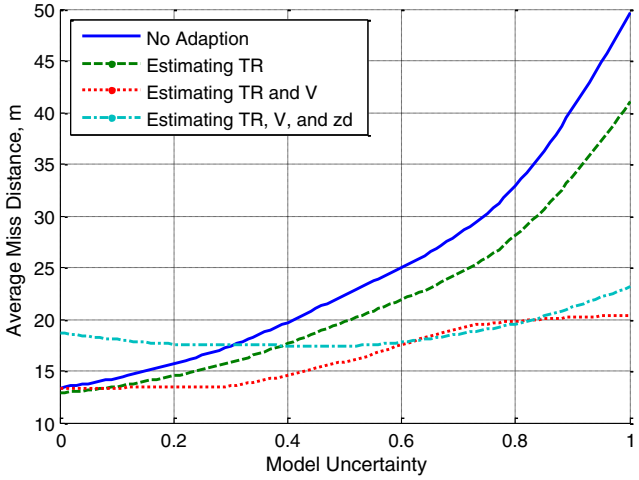

Fig. 29 Simulated landing accuracy vs model uncertainty with increasing levels of in-flight system identification.

VII. Conclusions

Although all of the simulation and flight test results presented focus on small systems constructed solely for research purposes, all of the findings presented here are applicable to full scale airdrop systems. The guidance, navigation and control algorithms are all based solely on position and velocity feedback already available on all modern guided airdrop systems. Although incidence angle control is not a typical feature of current systems, the promising results presented here suggest that it is worth considering this extra control channel for full-size systems in the future. The extra servo motor required for incidence angle control channel should not be a significant addition in cost or complexity to standard full-size systems. Even without incidence angle control, the in-flight system identification procedure would be simple to implement on current systems and the glide slope control algorithm presented here was shown to be very effective in providing glide slope control in windy conditions using only symmetric trailing edge brake deflection. These modifications could provide a significant improvement in the

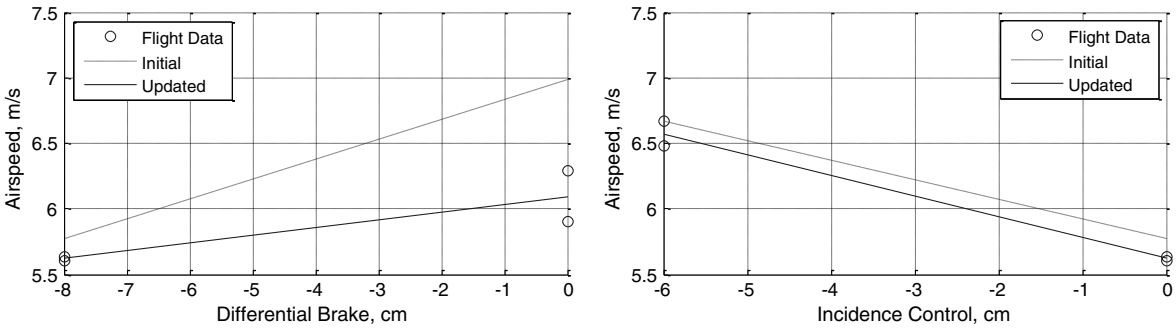

Fig. 27 Initial and updated airspeed model compared to flight data.

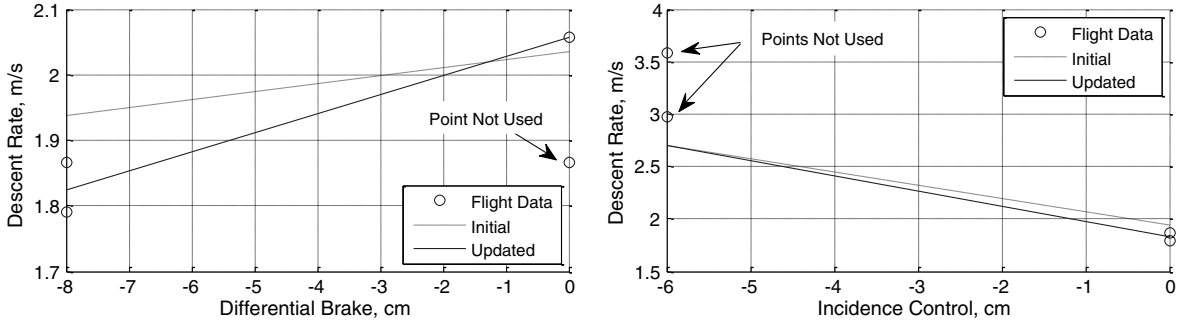

Fig. 28 Initial and updated descent rate model compared to flight data.

performance of current guided systems with only relatively simple modifications to flight software.

References

[1] Knapp, C., and Barton, W., "Controlled Recovery of Payloads at Large Glide Distances, Using the Para-Foil," *Journal of Aircraft*, Vol. 5, No. 2, 1968, pp. 112–118.

[2] Goodrick, T., "Estimation of Wind Effect on Gliding Parachute Cargo Systems Using Computer Simulation," *AIAA Aerodynamic Decelerator Systems Conference*, 14–16 Sept. 1970.

[3] Goodrick, T., Pearson, A., and Murphy, A., "Analysis of Various Automatic Homing Techniques for Gliding Airdrop Systems with Comparative Performance in Adverse Winds," *4th AIAA Aerodynamic Decelerator Conference*, Palm Springs, CA, 1973.

[4] Goodrick, T., "Hardware Options for Gliding Airdrop Guidance Systems," *6th AIAA Aerodynamic Decelerator and Balloon Technology Conference*, , Houston, TX, 1979, pp. 345–350.

[5] Mayer, R., "Terminal Descent Controlled Vehicle Recovery," AIAA Paper 84-0801, 1984.

[6] Murray, J., Sim, A., Neufeld, D., Rennich, P., Norris, S., and Hughes, W., "Further Development and Flight Test of an Autonomous Precision Landing System Using a Parafoil," NASA TM-4599, July 1994.

[7] Sim, A. G., Murray, J. E., Neufeld, D. C., and Reed, R. D., "Development and Flight Test of a Deployable Precision Landing System for Spacecraft Recovery," NASA TM-4525, 1993.

[8] Sim, A., Murray, J., Neufeld, D., and Reed, R., "Development and Flight Testing of a Deployable Precision Landing System," *Journal of Aircraft*, Vol. 31, No. 5, 1994, pp. 1101–1108. doi:10.2514/3.46617

[9] Hattis, P., et al., "Precision Guided Airdrop System Flight Test Results," *Proceedings of the 14th AIAA Aero-dynamic Decelerator Systems Technology Conference*, San Francisco, CA, 3–5 June 1997.

[10] Hogue, J. R., and Jex, H. R., "Applying Parachute Canopy Control and Guidance Methodology to Advanced Precision Airborne Delivery Systems," *Proceedings of the 13th AIAA Aerodynamic Decelerator Systems Technology Conference*, Clearwater Beach, FL, 15–18 May 1995.

[11] Smith, B. D., "Steering Algorithms for GPS Guidance of RAM-AIR Parachutes," *Proceedings the 8th International Technical Meeting of the Institute of Navigation*, Palm Springs, CA, 12–15 Sept. 1995.

[12] Anon., "Development and Demonstration Test of a Ram-Air Parafoil Precision Guided Airdrop System, Vol. 3, Simulation and Flight Test Results," Charles Stark Draper Laboratory, Inc., Final Rept. CSDL-R-2752, Oct. 1996.

[13] Anon., "Development and Demonstration Test of a Ram-Air Parafoil Precision Guided Airdrop System, Addendum: Flight Test Results and Software Updates after Flight 19," Charles Stark Draper Laboratory, Inc., Final Rept. CSDL-R-2771, Dec. 1996.

[14] Hattis, P., et al., "Precision Guided Airdrop System Flight Test Results," *Proceedings of the 14th AIAA Aero-dynamic Decelerator Systems Technology Conference*, San Francisco, CA, 3–5 June 1997.

[15] Gockel, W., "Concept Studies of an Autonomous GNC System for Gliding Parachute," *Proceedings of the 14th AIAA Aerodynamic Decelerator Systems Technology Conference*, San Francisco, CA, 3–5 June 1997.

[16] Soppa, U., Strauch, H., Goerig, L., Belmont, J. P., and Cantinaud, O., "GNC Concept for Automated Landing of a Large Parafoil," *Proceedings of the 14th AIAA Aerodynamic Decelerator Systems Technology Conference*, San Francisco, CA, 3–5 June 1997.

[17] Kaminer, I., and Yakimenko, O., "Development of Control Algorithm for the Autonomous Gliding Delivery System," AIAA Paper 2003-2116, May 2003.

[18] Carter, D., George, S., Hattis, P., Singh, L., and Tavan, S., "Autonomous Guidance, Navigation, and Control of Large Parafoils," AIAA Paper 2005-1643, May 2005.

[19] Jann, T., "Advanced Features for Autonomous Parafoil Guidance, Navigation and Control," AIAA Paper 2005-1642, May 2005.

[20] Carter, D., George, S., Hattis, P., McConley, M., Rasmussen, S., Singh, L., and Tavan, S., "Autonomous Large Parafoil Guidance, Navigation, and Control System Design Status," AIAA Paper 2007-2514, May 2007.

[21] Calise, A., and Preston, D., "Swarming/Flocking and Collision Avoidance for Mass Airdrop of Autonomous Guided Parafoils," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 4, 2008, pp. 1123–1132. doi:10.2514/1.28586

[22] Carter, D., Singh, L., Wholey, L., Rasmussen, S., Barrows, T., George, S., McConley, M., Gibson, C., Tavan, S., and Bagdonovich, B., "Band-Limited Guidance and Control of Large Parafoils," AIAA Paper 2009-2981, May 2009.

[23] Rademacher, B., Lu, P., Strahan, A., and Cerimele, C., "In-flight Trajectory Planning and Guidance for Autonomous Parafoils," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 6, 2009, pp. 1697–1712. doi:10.2514/1.44862

[24] Slegers, N., and Yakimenko, O., "Optimal Control for Terminal Guidance of Autonomous Parafoils," AIAA Paper 2009-2958, 2009.

[25] Yakimenko, O., Slegers, N., and Tiaden, R., "Development and Testing of the Miniature Aerial Delivery System Snowflake," AIAA Paper 2009-2980, 2009.

[26] Tavan, S., "Status and Context of High Altitude Precision Aerial Delivery Systems," AIAA Paper 2006-6793, 2006.

[27] Slegers, N., Beyers, E., and Costello, M., "Use of Variable Incidence Angle for Glide Slope Control of Autonomous Parafoils," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 3, 2008, pp. 585–596.

[28] Ward, M., Gavrilovski, A., and Costello, M., "Flight Test Results for Glide Slope Control of Parafoil Canopies of Various Aspect Ratios," *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, 2011.

[29] Calise, A., and Preston, D., "Design of a Stability Augmentation System for Airdrop of Autonomous Guided Parafoils," AIAA Paper 2006-6776, May 2006.

[30] Ward, M., Slegers, N., and Costello, M., "On the Benefits of In-Flight System Identification for Autonomous Airdrop Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 5, 2010, pp. 1313–1326. doi:10.2514/1.49030

[31] Goodrick, T. F., "Theoretical Study of the Longitudinal Stability of High-Performance Gliding Airdrop Systems," AIAA Paper 1975-1394, 1975.

[32] Crimi, P., "Lateral Stability of Gliding Parachutes," *Journal of Guidance, Control, and Dynamics*, Vol. 13, No. 6, 1990. doi:10.2514/3.20579

[33] Iosilevskii, G., "Center of Gravity and Minimal Lift Coefficient Limits of a Gliding Parachute," *Journal of Aircraft*, Vol. 32 No. 6, 1995, pp. 1297–1302. doi:10.2514/3.46878

[34] Goodrick, T. F., "Simulation Studies of the Flight Dynamics of Gliding Parachute Systems," AIAA Paper 1979-0417, 1979.

[35] Goodrick, T. F., "Comparison of Simulation and Experimental Data for a Gliding Parachute in Dynamic Flight," AIAA Paper 1981-1924, 1981.

[36] Slegers, N., and Costello, M., "Aspects of Control for Parafoil and Payload System," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 6, 2003, pp. 898–905. doi:10.2514/2.6933

[37] Mooij, E., Wijnands, Q. G. J., and Schat, B., "9 DoF Parafoil/Payload Simulator Development and Validation," *AIAA Modeling and Simulation Technologies Conference*, Austin, TX, Aug. 2003.

[38] Prakash, O., and Ananthkrishnan, N., "Modeling and Simulation of 9-DOF Parafoil-Payload System Flight Dynamics," *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, Keystone, CO, 21–24 Aug. 2006.

[39] Gorman, C., and Slegers, N., "Comparison and Analysis of Multi-Body Parafoil Models with Varying Degrees of Freedom," *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Dublin, 23–26 May 2011.

[40] Ward, M., Slegers, N., and Costello, M., "Specialized System Identification for Parafoil and Payload Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 2, 2012, pp. 588–597. doi:10.2514/1.53364

[41] Simon, D., *Optimal State Estimation*, Wiley, Hoboken, NJ, 2006, pp. 123–131, 407–410.

[42] Ward, M., Montalvo, C., and Costello, M., "Performance Characteristics of an Autonomous Airdrop System in Realistic Wind Environments," *AIAA Atmospheric Flight Mechanics Conference*, Toronto, 2–5 Aug. 2010.

[43] Calise, A., and Preston, D., "Approximate Correction of Guidance Commands for Winds," *20th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, Seattle, WA, 4–7 May 2009.