Combined Lateral and Longitudinal Control of Parafoils Using Upper-Surface Canopy Spoilers

E. Scheuermann,* M. Ward,† M. R. Cacan,‡ and M. Costello§

Georgia Institute of Technology, Atlanta, Georgia 30332-0150

Precision placement of guided airdrop systems necessarily requires some mechanism enabling effective directional control of the vehicle. Often this mechanism is realized through asymmetric deflection of the parafoil canopy trailing-edge brakes. In contrast to conventional trailing-edge deflection used primarily for lateral steering, upper-surface bleed air spoilers have been shown to be extremely effective for both lateral and longitudinal (i.e., glide slope) control of parafoil and payload systems. Bleed air spoilers operate by opening and closing several spanwise slits in the upper surface of the parafoil canopy, thus creating a virtual spoiler from the stream of expelled ram air. The work reported here considers the autonomous landing performance of a small-scale parafoil and payload system using upper-surface bleed air spoilers exclusively for both lateral steering and glide slope control. Landing accuracy statistics computed from a series of Monte Carlo simulations in a variety of atmospheric conditions and experimental flight tests were found to be in good agreement. Median miss distances for the combined lateral and longitudinal control logic are on the order of 13 m, indicating an improvement in landing accuracy of nearly 50% over similar systems employing only lateral steering control.

Nomenclature

F	=	vehicle turn rate mapping
$\mathrm{GS_{min}}, \mathrm{GS_{max}}$	=	minimum and maximum system glide slope
GS_c	=	commanded system glide slope
H	=	system glide slope mapping
h	=	current altitude, m
$\boldsymbol{I}_I, \boldsymbol{J}_I$	=	inertial reference frame axes along the north and east directions
$\boldsymbol{I}_\mathrm{wf}, \boldsymbol{J}_\mathrm{wf}$	=	wind-fixed reference frame axes along the downwind and crosswind directions
L	=	instantaneous distance from target, m
R	=	turn radius, m
$\hat{V}_0$	=	estimated horizontal projection of vehicle airspeed, m/s
$\hat{V}_{W,x}, \hat{V}_{W,y}$	=	estimated wind velocity components along inertial north and east directions, m/s
x, y	=	inertial position components of system mass center, m
x_T, y_T	=	target coordinates, m
$x_\mathrm{wf}, y_\mathrm{wf}$	=	wind-fixed components of system mass center, m
$\dot{z}$	=	system sink rate, m/s
α	=	aerodynamic angle of attack, rad
β	=	sideslip angle, rad
$\delta a, \delta s$	=	asymmetric, symmetric spoiler deflection
$\delta l, \delta r$	=	left, right spoiler deflection
χ_0	=	system total velocity azimuthal angle, rad
ψ	=	system heading angle, rad
ψ_c	=	commanded system heading angle, rad
ψ_W	=	wind direction, rad

*Graduate Research Assistant, Woodruff School of Mechanical Engineering (Corresponding Author).

†Research Engineer, Guggenheim School of Aerospace Engineering.

‡Graduate Research Assistant, Woodruff School of Mechanical Engineering.

§Professor, Guggenheim School of Aerospace Engineering, Woodruff School of Mechanical Engineering.

I. Introduction

AIRDROP systems offer the unique ability to deliver large payloads to undeveloped or otherwise inaccessible locations. In contrast with unguided ballistic parachutes, autonomously guided systems using steerable, ram air parafoil canopies have been shown to dramatically improve payload delivery accuracy and precision. These systems are released from a carrier aircraft at altitude and then guided to some desired landing location, often using deflection of the canopy trailing-edge brakes as the only available mechanism for directional control of the vehicle. In addition to the limited number of available control channels, the problem is further complicated when considering the effects of variable atmospheric conditions including changes in the assumed wind speed and direction used to plan the system approach trajectory.

Current autonomous parafoil and payload systems are controlled by asymmetric deflection of the canopy trailing-edge left and right brakes, providing an effective means for lateral–directional control. In contrast with asymmetric brake deflection, symmetric brake deflection predominantly causes a reduction in forward flight speed, with small changes in system glide slope until stall [1]. Although these systems have demonstrated substantial improvement in landing accuracy over similarly sized unguided systems, their limited number of available control channels and the effectiveness and independence of each control channel makes them highly susceptible to atmospheric gusts and other unknown surface conditions. In recent years, several researchers have demonstrated that adding longitudinal or glide slope control is a very effective means for reducing impact point errors and can greatly improve landing accuracy [2–5]. Several mechanisms capable of effective glide slope control have been identified, including symmetric brake deflection for airspeed control [2], in-flight adjustment of the canopy incidence angle [3,4,6], and actuation of upper-surface bleed air spoilers [1,7]. This latter mechanism is the subject of the current work and consists of several spanwise slits in the upper surface of the parafoil canopy that, when opened, create a virtual spoiler by releasing pressurized air from within the canopy cell. Much like conventional aircraft spoilers, opening of these spanwise slits creates a disturbance in the airflow over the parafoil wing, resulting in a localized perturbation of the associated aerodynamic forces. By varying the location of these slits along the upper surface of the canopy, significant changes in the lateral and longitudinal dynamics of the vehicle are achieved [1].

This current work documents the autonomous landing performance of a small-scale parafoil and payload system using upper-surface canopy spoilers for both lateral directional (steering) and longitudinal (i.e., glide slope) control. The following sections detail

the onboard guidance, navigation, and control (GNC) algorithm used for all flight experiments documented within this paper, including specific details regarding the newly developed spoiler control logic. Next, the test vehicle and associated onboard electronics are described, followed by the results from several system identification flights aimed at quantifying the effect of symmetric spoiler actuation on steady-state vehicle flight characteristics. Last, autonomous landing performance is investigated in both simulation and flight experiments, where nearly a 50% reduction in median miss distance was noted when using the combined control logic.

II.　Test Vehicle Description

To investigate the autonomous landing performance of upper-surface canopy spoilers, a small parafoil and payload system was developed, specifically modified to eliminate use of conventional trailing-edge brakes in favor of upper-surface canopy spoilers. The parafoil and payload system shown in Fig. 1 weighs approximately 2.3 kg and uses a 0.96 m² (61.0 cm chord, 157.5 cm span) rectangular planform canopy constructed from zero-porosity ripstop nylon. The system was designed to be hand-launched and features an electric motor and propeller, enabling sustained climbing flight. Spoiler actuation is achieved by two servoactuators connected to control lines attached directly to the leading edge of the upper-surface canopy slit. A third servoactuator is used to vary the length of the leading-edge risers, thus changing the canopy incidence between climbing and gliding flight. Additionally, an onboard autopilot equipped with a single GPS receiver and barometric altimeter is used for controlling the vehicle during execution of automated scripts or fully autonomous flight.

Modification from conventional trailing-edge brakes to upper-surface spoiler control required a series of spanwise slits carefully cut into the upper-surface fabric of the canopy. Each slit is actuated by pulling down from a single control line attached at the center of the leading-edge side of the upper-surface opening that runs down through the bottom surface of the canopy to the payload. As the upper-surface slit opens, a stream of high-pressure ram air from within the canopy cell is expelled, creating a virtual aerodynamic spoiler. When not actuated, the spanwise tension and internal pressure from within the canopy cell is sufficient to force the upper-surface spoiler closed, preventing any further airflow to the outside. Additionally, a small piece of fabric or sealing flap is added to the actuated edge of the slit to help prevent air leakage when the spoiler is closed. A cross-sectional view of the upper-surface canopy spoiler mechanism within a single cell is shown in Fig. 2. Note that the upper-
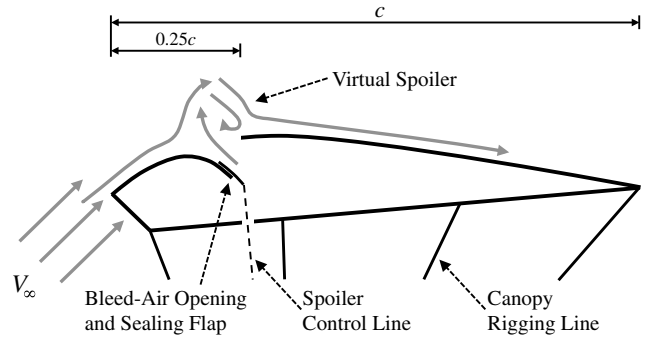

Fig. 2　Cross-section view of upper-surface canopy spoiler control mechanism.

surface bleed air opening is located at approximately 0.25c back from the leading edge, where c represents the mean airfoil chord. This configuration is similar to that tested by Gavrilovski et al. in [1].

In total, eight cells (four left, four right) within the test vehicle canopy were modified to include upper-surface spoiler openings. Figure 3 provides a front-view illustration of the test vehicle canopy, where each shaded gray region denotes the relative location of those cells with spoiler openings. Four cells within the canopy center section were left unmodified, whereas two groups of four cells on either side of the center section were configured for spoiler control. This arrangement was selected to balance the expected lateral and longitudinal control authority of the vehicle, noting that spoiler openings nearest the canopy center contribute primarily to changes in vehicle airspeed and descent rate, whereas openings nearest the canopy wingtips contribute mostly to changes in turn rate. In practice, the upper-surface spoilers are actuated as two distinct groups as opposed to eight individual openings, enabling two-channel control of the left and right spoiler openings using two servoactuators.

III.　System Characterization

One prerequisite to the onboard GNC algorithm is accurate knowledge of the parafoil and payload system lateral and longitudinal response to various asymmetric and symmetric spoiler openings. Accordingly, a series of test flights were performed to measure the vehicle steady-state response (i.e., airspeed, descent rate, and turn rate) as a function of different combinations of spoiler opening. During each flight, several periods of constant control input were held for 20–30 s or longer, allowing the system to fly in a near-circular path for at least one complete rotation while GPS and barometric altitude data were continuously logged within the autopilot memory. Circular paths were used to facilitate accurate
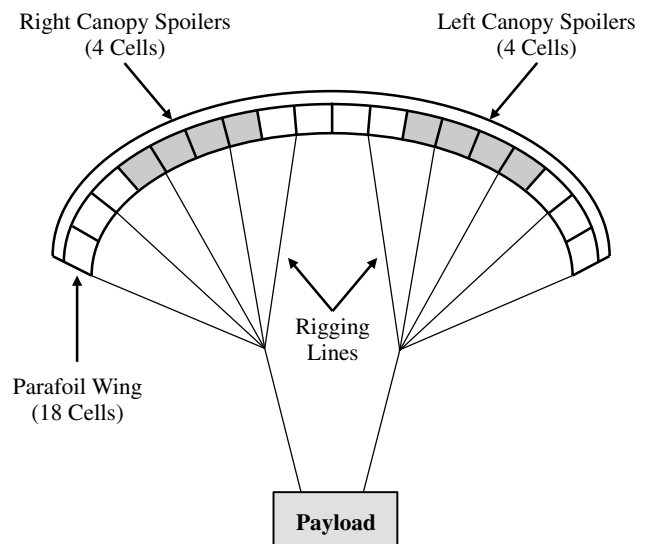

Fig. 1　Test vehicle during gliding flight.

Fig. 3　Upper-surface canopy spoiler control mechanism (front view).

Fig. 4 Measured vehicle turn rate vs normalized asymmetric spoiler deflection δa.

Fig. 6 Measured vehicle descent rate vs normalized symmetric spoiler deflection δs.
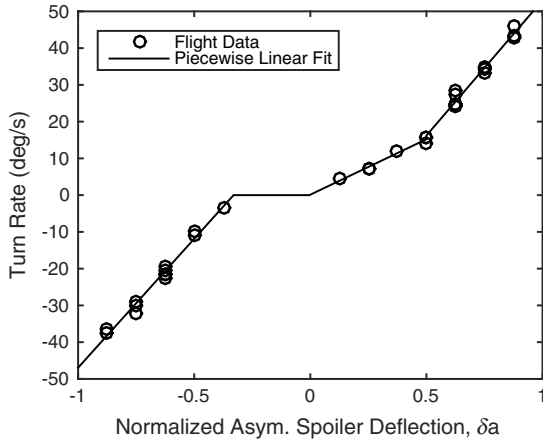

Fig. 7 Estimated vehicle glide slope vs normalized symmetric spoiler deflection δs.

wind estimation and subsequent compensation. All flight data recorded during each period of constant control input were later postprocessed and time-averaged to obtain a series of points representative of the observed steady-state flight characteristics.

Figure 4 shows the measured steady-state turn rate as a function of normalized asymmetric spoiler deflection, denoted δa. Asymmetric spoiler deflection is defined as simply the difference between the left and right spoiler openings, $\delta a = \delta r - \delta l$, where δr and δl are the individual right and left spoiler deflections, respectively, and δl, $\delta r \in [0, 1]$. For the test vehicle, it was observed that turn rate response is both nonlinear and asymmetric depending on turn direction (right turns are defined as positive values of turn rate). In the case where the vehicle is turning left, the measured turn rate response is nearly linear, although a significant deadband was observed for nearly the first 50% of normalized asymmetric spoiler deflection. Right turns did not exhibit any deadband behavior in response to asymmetric spoiler deflection, yet the measured turn rate response was not consistent across the entire range of travel.

By inverting the vehicle turn rate response shown in Fig. 4, a piecewise linear mapping is created, where $\dot\psi$ represents the vehicle turn rate in degrees per second. The mapping function $F(\dot\psi)$ shown in Eq. (1) will be later used within Sec. IV. Note that asymmetry and deadband behavior in the turn rate response are likely the result of improper control line or rigging line lengths. Additionally, asymmetries in the construction of the canopy itself can also result in asymmetric turn rate response:

$$\delta a = F(\dot\psi) = \begin{cases} 0.012\dot\psi - 0.430 & \text{if } \dot\psi < 0 \\ 0.033\dot\psi & \text{if } 0 \leq \dot\psi \leq 15 \text{ deg}/\text{s} \\ 0.013\dot\psi + 0.500 & \text{if } \dot\psi > 15 \text{ deg}/\text{s} \end{cases} \quad (1)$$
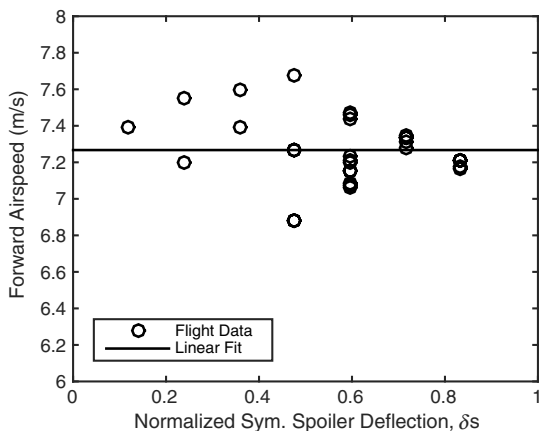

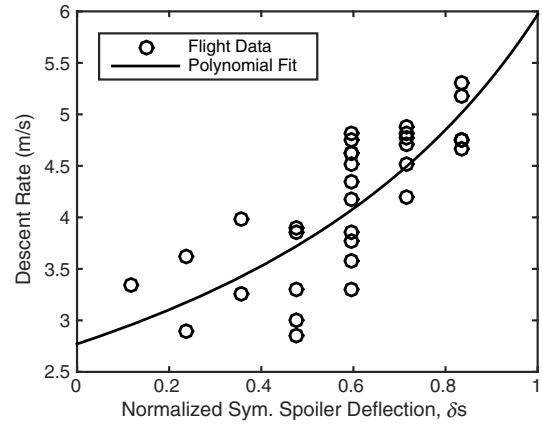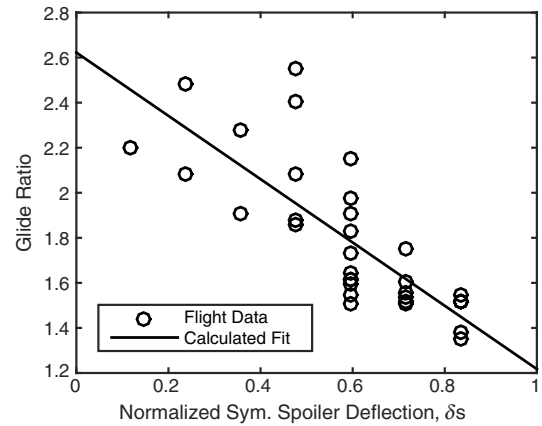**Fig. 5 Estimated vehicle airspeed vs normalized symmetric spoiler deflection δs.**

In addition to vehicle turn rate, additional steady-state characteristics including airspeed, descent rate, and glide ratio were also computed as a function of normalized symmetric spoiler deflection, denoted δs. Here, symmetric spoiler deflection is defined as the average opening between the left and right spoiler, $\delta s = 0.5(\delta l + \delta r)$. These results are shown in Figs. 5–7. Estimated forward airspeed values shown in Fig. 5 were relatively constant for all values of symmetric spoiler deflection, resulting in an average value of 7.3 m/s. Forward airspeed refers only to the horizontal projection of vehicle airspeed. In terms of measured descent rate shown in Fig. 6, significant changes with increasing symmetric spoiler deflection were noted, resulting in values ranging from approximately 3.0 m/s with no control input to nearly 5.5 m/s at near-full symmetric deflection.

Figure 7 presents the range of feasible glide ratios observed in flight tests as a function of symmetric spoiler input. Glide ratio or glide slope is simply the ratio of forward airspeed over descent rate. This result is extremely important and demonstrates that the effective longitudinal control authority of the upper-surface canopy spoiler mechanism is capable of reducing the test vehicle glide slope from approximately 2.6 to less than 1.5 in a nearly linear manner. As a result, a simple linear mapping is created by inverting the vehicle glide slope response, where GS denotes the vehicle glide slope. This mapping function $H(\text{GS})$ shown in Eq. (2) will be later used within Sec. IV:

$$\delta s = H(\text{GS}) = -1.4\text{GS} + 2.6 \quad (2)$$

Steady-state lift and drag coefficients were also estimated according to the procedure described by Ward et al. in [8] and are shown in Figs. 8 and 9 for three different levels of normalized symmetric spoiler deflection. The resulting curves were formed from
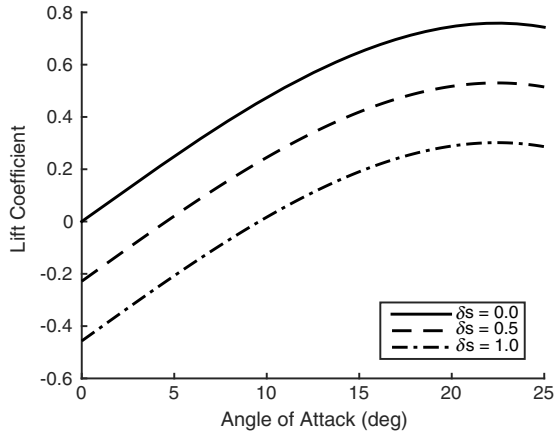
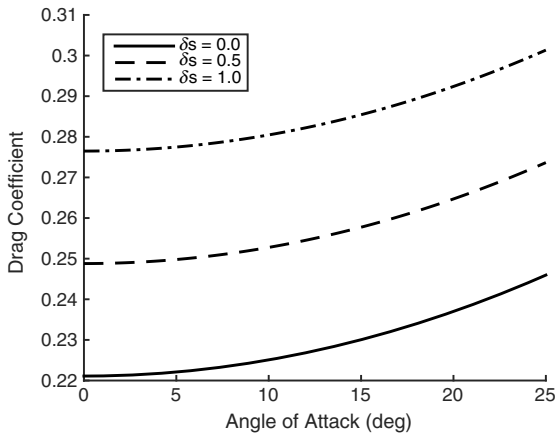**Fig. 8 Vehicle lift coefficient vs angle of attack.**

Fig. 9 Vehicle drag coefficient vs angle of attack.

simple polynomial expressions to match the coefficients estimated from flight data. Note that angle of attack was estimated as the difference between the total flight path angle and the nominal incidence angle of the canopy assuming a pitch angle of zero. Nominal incidence angle of the test vehicle canopy during glide was fixed at −10 deg. These models were used in constructing an accurate aerodynamic model of the vehicle flight characteristics for purposes of trajectory simulation and landing accuracy prediction (refer Sec. V). Details of the six-degree-of-freedom parafoil and payload system equations of motion can be found in [8,9]. Validation of the test vehicle dynamic model with flight data is omitted here for brevity. However, previous work has shown the model to accurately represent the observed flight characteristics of similarly sized parafoil and payload system [8,9].

The important point to take away from the resulting curves shown in Figs. 8 and 9 is the relative effect of increasing symmetric spoiler deflection. Here, the effect varies linearly with each coefficient and simply shifts each curve up or down. Increasing symmetric spoiler deflection is characterized by decreased lift with increased drag. These results are consistent with the expected effects inherent to aerodynamic spoilers for conventional fixed-wing aircraft. Additionally, the effect of symmetric spoiler deflection is somewhat opposite that of increasing symmetric trailing-edge brake deflection. In the latter case, both lift and drag are increased with increasing symmetric brake deflection [9].

IV. Guidance, Navigation, and Control Algorithm for Autonomous Landing

The guidance, navigation, and control (GNC) algorithm is responsible for path planning, estimating relevant system states and atmospheric winds, and specifying the proper control inputs to follow

the desired trajectory. As reference for the reader, a complete description of the GNC algorithm is presented. Where appropriate, details specific to the use of upper-surface canopy spoilers are provided.

A. Guidance

The guidance algorithm is tasked with planning an appropriate path to reach the target landing area, given relevant state and atmospheric wind estimates. It consists of four separate phases (initialization, loiter, final approach, and terminal guidance), where each phase is briefly described in the following sections. In the current work, all path planning is performed within a wind-fixed or wind-based reference frame (denoted with subscript wf). Originally introduced by Goodrick et al. [10] and later by Jann [11], the wind-fixed frame is a moving reference frame that offsets the current vehicle position by some distance equal to the expected wind drift the vehicle will experience during the remaining portion of the flight. Additionally, the I_{wf} axis is rotated to align with the wind direction. Use of a wind-fixed frame is advantageous in that it allows all guidance and path-planning calculations to be performed as if the system were flying through completely still air. Figure 10 provides a visual illustration of the wind-fixed reference frame.

As seen in Fig. 10, Δx and Δy represent the expected drift of the vehicle due to horizontal winds along north and east directions, respectively. These values are computed as the integral of the wind profile and vehicle descent rate from ground level to the current vehicle altitude h, according to Eqs. (3) and (4), where $V_{W,x}(z)$ and $V_{W,y}(z)$ represent the north and east wind velocity components with changing altitude, and $\dot{z}(z)$ is the vehicle sink rate with changing altitude. Orientation of the wind-fixed frame is determined according to the estimated wind direction as shown in Eq. (5). In practice, the two argument arctangent function atan2($\cdot$) commonly found in many computer programming languages may be used to properly resolve the wind direction, where the resulting angle spans the range $(-\pi, \pi]$:

$$\Delta x = \int_0^h \frac{V_{W,x}(z)}{\dot{z}(z)}\,\mathrm{d}z \approx \frac{h}{\dot{z}}\hat{V}_{W,x} \tag{3}$$

$$\Delta y = \int_0^h \frac{V_{W,y}(z)}{\dot{z}(z)}\,\mathrm{d}z \approx \frac{h}{\dot{z}}\hat{V}_{W,y} \tag{4}$$

$$\psi_W = \tan^{-1}\frac{\hat{V}_{W,y}}{\hat{V}_{W,x}} \tag{5}$$

Determination of the exact wind profile is difficult without the use of some external measurement source (e.g., weather balloons, dropsondes, Lidar, meteorological forecasting, etc.) and a real-time data uplink. Sink rate can also vary substantially during flight depending on the specific control mechanism used and the bank angle of the vehicle while turning. An approximation to the integrals shown in Eqs. (3) and (4) is the use of a constant descent rate for the vehicle, which is typical for most airdrop systems [11]. Additionally, the wind profile is reduced to two scalar components, denoted $\hat{V}_{W,x}$ and $\hat{V}_{W,y}$, representative of the estimated wind components at the current vehicle altitude using the onboard navigation algorithm (refer to Sec. IV.B for details on wind estimation). These estimates are
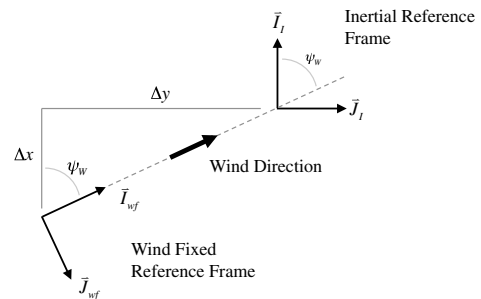

Fig. 10 Wind-fixed reference frame.

continually updated throughout vehicle descent yet provide no indication of the winds below the current altitude of the vehicle. As expected, perfect knowledge of the expected vehicle drift is not possible if the winds near ground level vary significantly from those at altitude. However, as the onboard wind estimates change during flight, the position and orientation of the wind-fixed frame adjusts accordingly.

1. Initialization

Autonomous airdrop systems typically execute some type of homing maneuver immediately following deployment from the carrier aircraft, where the system attempts to fly either directly above the target location or to some predefined loiter area [12–15]. Various system parameters, including sink rate and turn rate bias, can also be estimated and properly compensated during homing flight [12, 13]. In the current work, the goal of the initialization phase is to provide reasonable initial state estimates of wind and vehicle airspeed to the navigation algorithm. This is accomplished by holding a constant asymmetric spoiler input, allowing the system to enter into an open-loop steady turn completing at least one full rotation. The duration of the initialization phase and initial turn rate are dependent on the specific vehicle and designed as input parameters to the GNC algorithm. Typical values include a total initialization time of 35 s and a turn of approximately 15 deg/s. Turn direction is arbitrarily selected.

Using the resulting GPS velocity measurements recorded throughout the entire open-loop turn, the task of estimating the horizontal wind components and vehicle airspeed can be cast as a simple linear regression problem. Two necessary assumptions are made, where each wind component and the forward airspeed of the vehicle are constant throughout the entire open-loop turn. Here, forward airspeed indicates the horizontal projection of the total vehicle airspeed. Further details on wind and airspeed estimation can be found in [8,9] and are omitted here for brevity.

2. Loiter

The loiter or energy management phase of the descent begins immediately following initialization and consists of a series of maneuvers intended to systematically reduce the excess altitude of the system before landing. In previous work, various loiter patterns have been explored, including consecutive spiral or S-shaped turns [12,16], figure eights [3,13,14], and Dubins paths [17]. The loiter strategy employed here is similar to that found in [11], where an alternating sequence of trajectories are generated between two fixed waypoints. Each waypoint is located on either end of a T-shaped pattern whose geometry remains stationary with respect to the wind-fixed frame. Use of a T-shaped pattern ensures that each loiter waypoint is located downwind of the intended target to facilitate an upwind landing. Recall that the axes of the wind-fixed frame are always rotated to coincided with the estimated wind direction. A graphical illustration of the T-shaped loiter pattern is shown in Fig. 11.

Figure 11 also presents an example trajectory between subsequent loiter waypoints. Here, Dubins paths [18] are used for trajectory planning, consisting of two constant radius arcs joined by a single straight line segment. Arc radius, denoted R, is dependent on both the forward airspeed and maximum turn rate of the vehicle and is left as an input parameter to the GNC algorithm. As an illustrative example, the arc radius in Fig. 11 was chosen as 25 m. Locations of the loiter waypoints are also parameterized according to arc radius, where each loiter target was located a distance $5R$ downwind of the intended target and offset a distance of $4R$ in the crosswind direction. Aside from the fixed arc radius, additional constraints for each Dubins path include fixed initial and final vehicle positions and tangent directions. Also, the direction of the first turn is set to equal that of the second turn from the previously computed path, except for the case immediately following initialization where the direction of the first turn is arbitrary. The second turn is then chosen according to the minimum-distance path satisfying all constraints. Although each turn is often opposite in direction, resulting in a continuous figure-eight
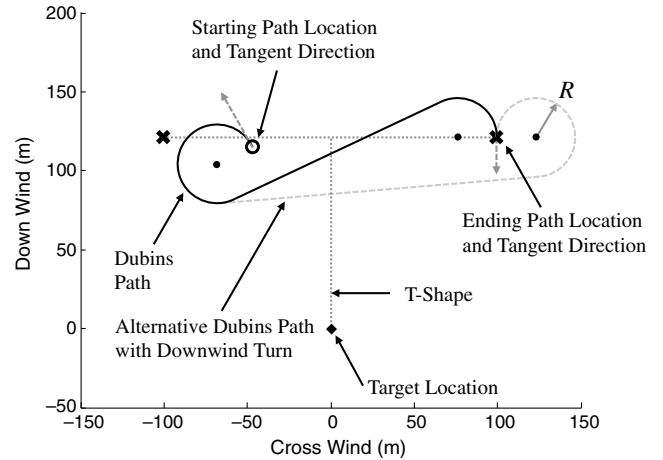

Fig. 11 T-shaped loiter pattern relative to the wind-fixed frame.

pattern between each waypoint, this is not strictly enforced such that trajectories with initial and final turns in the same direction are also admissible.

During loiter, altitude required to reach the target from the current vehicle position is continuously computed each update cycle using Eqs. (6) and (7). The instantaneous distance to the target, denoted L, is defined by the arc length with radius R required to turn from the current heading to point directly at the target and the straight line segment between the end of this turn and the target. Figure 12 provides a graphical depiction of the instantaneous distance from the target:

$$L = |\Delta\psi|R + \sqrt{(x_1 - x_T)^2 + (y_1 - y_T)^2} \qquad (6)$$

$$h_{req} = L \frac{\dot{z}}{\hat{V}_0} \qquad (7)$$

In Eq. (6), $\Delta\psi$ represents the change in vehicle heading necessary to point directly at the target, (x_1, y_1) are the end point coordinates of the circular arc, and (x_T, y_T) are the target coordinates. This distance is then converted into the required height using Eq. (7), where $\dot{z}$ and $\hat{V}_0$ represent the measured descent rate and estimated horizontal projection of vehicle airspeed, respectively. Once the current vehicle altitude drops below h_{req}, the guidance algorithm switches from loiter to final approach.

3. Final Approach

A desirable feature of any autonomous airdrop system is the ability to land facing directly into the wind. This provides for the greatest
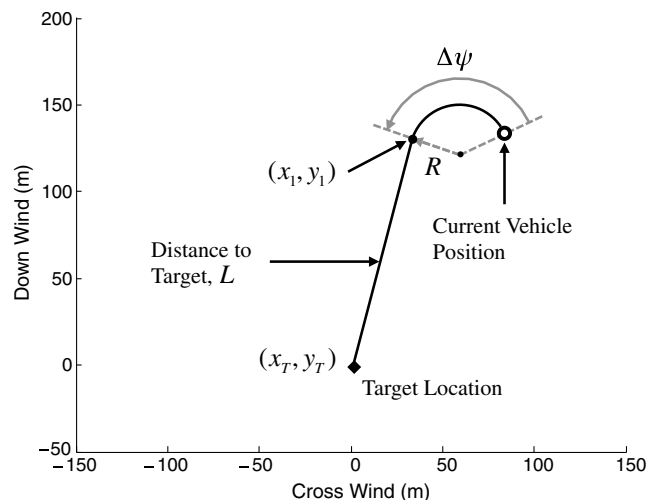

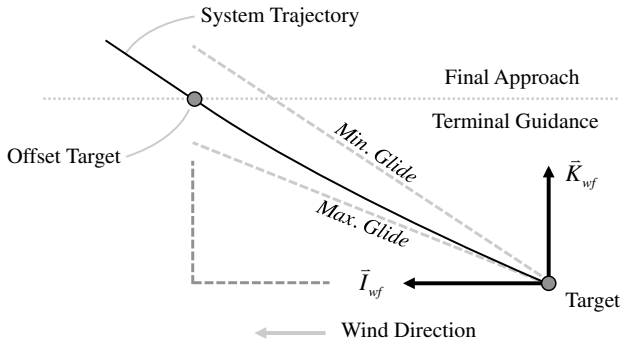Fig. 12 Instantaneous distance to target.

Fig. 13 Terminal guidance strategy.

reduction in ground speed just before touchdown and helps maximize payload survivability during impact. Similar to that of Jann [11], a two-stage final approach technique is employed where the parafoil and payload system initially begins tracking to an offset target following loiter before transitioning to the actual desired impact point. This technique facilitates landing of the vehicle pointing into the estimated wind vector by forcing the system to fly along the vertical stem of the T-shaped pattern. The offset target is located directly downwind of the desired impact point at an altitude centered within the feasible glide range of the vehicle. This ensures that sufficient longitudinal control margin exists to compensate for errors in the approach trajectory caused by changes in the estimated winds and other outside disturbances. Figure 13 provides an illustration of the longitudinal control strategy employed during final approach and terminal guidance.

Often the decision to leave loiter and begin final approach is responsible for highly accurate or highly inaccurate landings. Accordingly, the ability of the vehicle to actively control its glide ratio during approach provides a significant advantage for reducing impact point errors when compared to systems with lateral steering only. In the current work, upper-surface canopy spoilers are used for longitudinal control only during the final two phases of the guidance algorithm, namely final approach and terminal guidance. In the case where the required glide slope necessary to intersect the intended target is less than the nominal system glide slope (i.e., the system is likely to overshoot the intended target), the upper-surface canopy spoilers are opened symmetrically, thus temporarily increasing the vehicle sink rate and allowing the system to "drop" down onto the correct glide path. In the opposite case where the system is likely to land short of the intended target given its current glide path, the upper-surface spoilers are simply closed to lower the vehicle sink rate, thus temporarily increasing the resulting glide path.

4. Terminal Guidance

The guidance logic transitions from final approach to terminal guidance at the instant the altitude of the vehicle drops below the height of the offset target. Recall that the offset target is located directly downwind of the desired impact point at an altitude slightly higher than the nominal glide of the vehicle. The idea behind this technique is that if the vehicle reaches the offset target at the correct altitude, the remaining portion of the descent will focus primarily on traversing the stem of the T-shaped pattern while maintaining the proper heading and glide path necessary to intersect the target. In some ways, this technique is analogous to that of an instrument landing system for fixed-wing aircraft. If terminal guidance is entered either above the minimum glide path or below the maximum glide path, the vehicle will inevitably overshoot or land short of the intended target, respectively. In this case, the guidance logic will simply saturate the symmetric spoiler control while simultaneously keeping the vehicle pointed either at the target or directly upwind of the target in an effort to minimize miss distance.

Just before ground impact, the vehicle executes a landing maneuver intended to minimize the kinetic energy of the system. In systems using conventional trailing-edge brakes for control, this is typically accomplished by applying full symmetric brake just before

touchdown. However, for the small-scale test vehicle using upper-surface canopy spoilers with no trailing-edge control, a third actuator is used to raise the canopy trim angle just before impact by lengthening the leading-edge risers. This creates a flaring effect similar to that resulting from full trailing-edge brake deflection. For larger autonomous systems, variable canopy incidence angle is not a typical feature, and a separate mechanism is required for kinetic energy reduction during landing.

B. Navigation

Following the open-loop initialization maneuver, the navigation algorithm is tasked with continually updating estimates of the atmospheric wind vector and current vehicle heading and heading rate using available sensory information. These measurements include position and velocity information from GPS and barometric pressure for altitude and altitude rate. This estimation process is accomplished using a discrete extended Kalman filter observer [19] based on the solution of the vector diagram shown in Fig. 14. A complete description of navigation algorithm is omitted for brevity. Additional details of the estimation process can be found in [8,9].

It is important to note that decomposition of the measured vehicle ground-speed vector, denoted V_{GPS}, into the horizontal projections of vehicle airspeed V_0 and atmospheric wind V_W is not unique. Solution of the vector diagram in Fig. 14 requires the assumption that airspeed is constant throughout the entire flight to formulate a more tractable problem. Additionally, solution of the vector diagram does not directly yield an estimate of the parafoil heading angle ψ, but rather the azimuthal angle χ_0. However, the sideslip angle β is typically small for parafoil and payload aircraft, in which case the azimuthal angle is assumed to be equal to the actual vehicle heading angle.

C. Control

The parafoil and payload system is controlled by asymmetric and symmetric opening of the upper-surface canopy spoilers. For all constant-radius segments of the planned vehicle path, the asymmetric spoiler deflection is essentially held constant (i.e., open-loop) until the required change in vehicle heading angle is achieved. Selection of the required asymmetric deflection value is based on the known vehicle turn rate response and the corresponding turn rate, given the path radius. All straight segments of the planned path use a nonlinear proportional–integral controller, where the commanded heading angle, denoted ψ_c, is chosen to point directly at some desired waypoint from the current vehicle location. Waypoints may include the end point of any straight path segment during loiter and the offset target or actual target during final approach and terminal guidance. A nonlinear proportional component is used to reduce control effort when the heading error is small. By comparing the commanded heading angle with the estimated vehicle heading from navigation $\hat{\psi}_{nav}$, vehicle turn rate commands, denoted $\dot{\psi}_c$, are calculated using Eqs. (8) and (9):
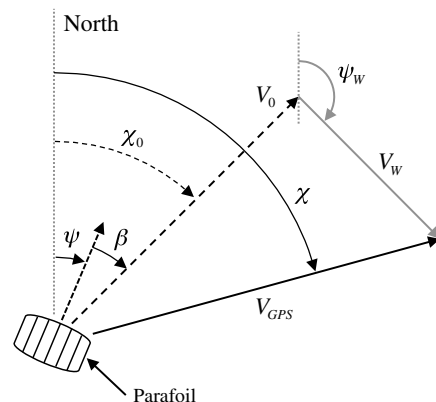

Fig. 14 Parafoil ground track velocity decomposition.

$$\Delta\psi_{\text{ratio}} = \frac{\psi_c - \hat{\psi}_{\text{nav}}}{\Delta\psi_{\text{max}}} \qquad (8)$$

$$\dot{\psi}_c = \begin{cases} \dot{\psi}_{\text{max}} & \text{if } \Delta\psi_{\text{ratio}} \geq 1 \\ -\dot{\psi}_{\text{max}} & \text{if } \Delta\psi_{\text{ratio}} \leq -1 \\ \dot{\psi}_{\text{max}} \Delta\psi_{\text{ratio}} \sqrt{|\Delta\psi_{\text{ratio}}|} & \text{else} \end{cases} \qquad (9)$$

Here, $\Delta\psi_{\text{max}}$ is the maximum difference between the commanded and current vehicle heading after which saturation occurs, and $\dot{\psi}_{\text{max}}$ is the maximum allowable turn rate in either direction. From this commanded turn rate, the required asymmetric spoiler deflection is determined according to Eq. (10), where F is the known vehicle turn rate mapping determined through prior system characterization. Last, δa_{bias} is the integral component computed from the difference between the commanded and estimated vehicle heading rates multiplied by the integral gain K_I. The summation over the index i shown in Eq. (11) represents each GNC update in which the vehicle is operating under closed-loop tracking control. The integral component is also restricted to the interval $[-1, 1]$:

$$\delta a = F(\dot{\psi}_c) + \delta a_{\text{bias}}, \qquad \delta a \in [-1, 1] \qquad (10)$$

$$\delta a_{\text{bias}} = K_I \sum_i [\dot{\psi}_c(i) - \hat{\dot{\psi}}_{\text{nav}}(i)], \qquad i = 1, \ldots, N \qquad (11)$$

In addition to lateral steering control, longitudinal control is performed using a simple proportional control law. As previously mentioned, longitudinal control is only active during final approach and terminal guidance phases in which the system is commanded to remain on its nominal glide path to intersect the desired target. Recall that the nominal glide path, denoted GS_{nom}, is that centered with the feasible glide range of the vehicle. Accordingly, the following equations are used to compute commanded values of glide slope, denoted GS_c, where GS_c is confined within the interval $[GS_{\text{min}}, GS_{\text{max}}]$:

$$GS_c = \begin{cases} GS_{\text{min}} & \text{if } \frac{L}{h} \leq GS_{\text{min}} \\ GS_{\text{max}} & \text{if } \frac{L}{h} \geq GS_{\text{max}} \\ GS_{\text{nom}} + K_{\text{GS}}\left(\frac{L}{h} - GS_{\text{nom}}\right) & \text{else} \end{cases} \qquad (12)$$

In Eq. (12), GS_{max} and GS_{min} correspond to the maximum and minimum values of system glide slope, L is the instantaneous distance from the target, h is the current vehicle altitude, and K_{GS} is the glide slope proportional gain. However, it must be noted that asymmetric spoiler deflection used for steering control also creates some level of symmetric deflection, defined as the absolute value of half the calculated asymmetric spoiler deflection, $\delta s = |\delta a|/2$. The amount of additional symmetric spoiler needed to achieve the required glide path is simply the difference between the required level of symmetric deflection according to the known glide slope mapping, denoted as the function H, and the symmetric deflection resulting from asymmetric turning commands. This result is shown in Eq. (13):

$$\delta s = H(GS_c) - \frac{|\delta a|}{2}, \qquad \delta s \in [0, 1] \qquad (13)$$

Individual left and right spoiler deflections, denoted δl and δr, respectively, are then computed using Eq. (14), where δl and δr are restricted within the interval $[0, 1]$. In cases where the individual deflections exceed their maximum value, the control inputs are first set to meet the required asymmetric deflection necessary for lateral steering and then, if possible, opened further to increase symmetric spoiler input until saturation occurs. This creates an inherent prioritization of lateral steering over longitudinal control when the heading angle error is large:

$$\delta l = \begin{cases} -\delta a + \delta s & \delta a < 0 \\ \delta s & \delta a \geq 0 \end{cases}, \qquad \delta r = \begin{cases} \delta a + \delta s & \delta a > 0 \\ \delta s & \delta a \leq 0 \end{cases} \qquad (14)$$

V. Autonomous Landing Performance

The following sections detail the autonomous landing performance of the test vehicle in both simulation and experimental flight tests. First, a single autonomous trajectory is considered to illustrate the overall machinery of the GNC algorithm and combined lateral and longitudinal control logic using upper-surface canopy spoilers. Next, simulation and experimental landing accuracy statistics are presented, including a comparison with an identical system using upper-surface canopy spoilers exclusively for lateral directional control.

A. Single Trajectory Analysis

In this example flight, the test vehicle was manually piloted to an altitude of approximately 400 m. Upon reaching the desired altitude, the electric motor was stopped and the onboard flight computer activated to begin autonomous descent and landing. Figure 15 shows the atmospheric wind magnitude and direction estimated by the onboard navigation algorithm throughout the entire flight. Note that accurate wind estimates are only valid following completion of the open-loop initialization routine at approximately 280 m altitude. Initially, the wind aloft is blowing due west at approximately 2.5 m/s before transitioning to blow north at just over 1.0 m/s near ground level. Actual winds are unknown throughout the entire flight.

Using the estimated wind values, Fig. 16 presents the measured vehicle trajectory in both the wind-fixed and inertial reference frames where the wind-fixed trajectory begins immediately following initialization. In this case, the vehicle enters autonomous flight located approximately 100 m north and 100 m west of the target (origin). Following nearly two complete turns during initialization, the system recognizes that it has flown too far downwind and begins homing toward the first tracking target. Upon reaching this first tracking target, the vehicle turns left 180 deg and begins homing toward the second tracking target, before making the transition from loiter to final approach. During final approach, the vehicle abandons the second tracking target and begins flying toward the offset target located 50 m directly downwind of the desired impact point. Accordingly, as the vehicle nears the offset target, it transitions to the actual desired impact point and maintains this heading until landing. Although ground winds are blowing toward the north, as shown in Fig. 15, the vehicle maintains an approach trajectory mostly from the northwest. This decision highlights the fact that the onboard GNC is attempting to execute the initially planned approach trajectory (i.e., winds assumed from the east, given the relatively low wind
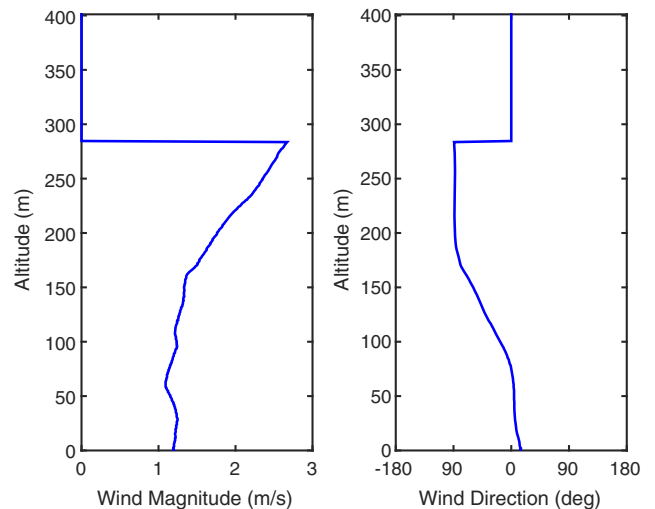

Fig. 15 Estimated wind speed and direction during example autonomous flight.
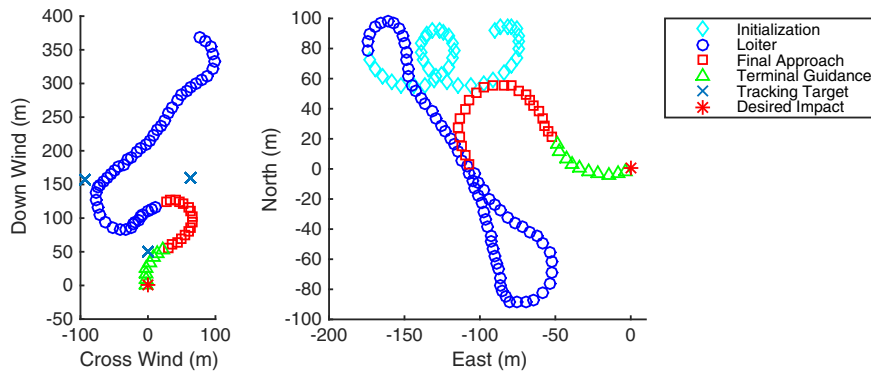
Fig. 16   Example autonomous flight trajectory shown in wind-fixed reference frame (left) and inertial reference frame (right).

magnitudes near ground). Final miss distance for this flight was 3.2 m.

In Fig. 17, vehicle altitude and normalized control inputs are presented as a function of instantaneous distance from the target where the dashed lines represent the maximum and minimum glide path of the vehicle from the target (i.e., glide ratios of approximately 2.5 and 1.5, respectively). Note that the vehicle is initially too high as it approaches the target. As a result, the onboard control logic recognizes this situation and opens the upper-surface canopy spoilers via maximum symmetric control input. Consequently, the vehicle glide ratio is temporarily reduced as it simply drops onto the correct glide path necessary to intersect the desired impact point. Lastly, Fig. 18 compares the actual and commanded vehicle heading angles versus time for the example autonomous flight.

B.   Simulation Results

Using the preceding results from the detailed system characterization discussed in Sec. III, a rigid-body, six-degree-of-freedom simulation model was developed in which the aerodynamic parameters were tuned to match the steady-state flight characteristics of the test vehicle. Recall that lift and drag coefficients are dependent on both the canopy angle of attack and the symmetric deflection of the upper-surface spoilers, as shown in Figs. 8 and 9. Using this computer model, complete autonomous landings were simulated

from altitude to ground impact in a variety of realistic atmospheric conditions, allowing for both refinement of the combined lateral and longitudinal control logic and evaluation of the expected landing performance. Predictions for the relative improvement in landing accuracy over the conventional lateral control only case were also computed.

For each control strategy, a series of Monte Carlo simulations were performed in which 500 autonomous landings were simulated, where the mean wind speed was allowed to vary uniformly from 0 to 5 m/s, and the turbulence level (defined as the standard deviation of the vertical gust component in the Dryden turbulence model) was varied uniformly from 0 to 0.75 m/s. Figure 19 presents the simulated landing accuracy dispersion for each case, where impact point errors have been rotated into a downwind and crosswind frame based on the actual ground winds at the time of landing. Note that the solid and dashed black lines represent the 50 and 90% circular error probable (CEP) regions, respectively. CEP is defined as the minimum radius centered at the target that encompasses either 50 or 90% of all recorded impact points. Landing accuracy statistics indicate a 43% reduction in 50% CEP when using the combined lateral and longitudinal control (21.9 m lateral only, 12.5 m combined lateral and longitudinal control). Similarly, a 38% reduction in 90% CEP was also computed (43.4 m lateral only, 26.8 m combined lateral and longitudinal control).

C.   Flight-Test Results

In an effort to validate the predicted landing accuracy using the parafoil and payload simulation model, a total of 70 autonomous landings were recorded during flight testing in Eloy, Arizona, under a variety of wind conditions (7 m/s or less). For each landing, the test vehicle described in Sec. II was hand-launched and manually piloted during powered climb to an altitude of approximately 400 m. Upon reaching the desired altitude, the electric motor was stopped, and the
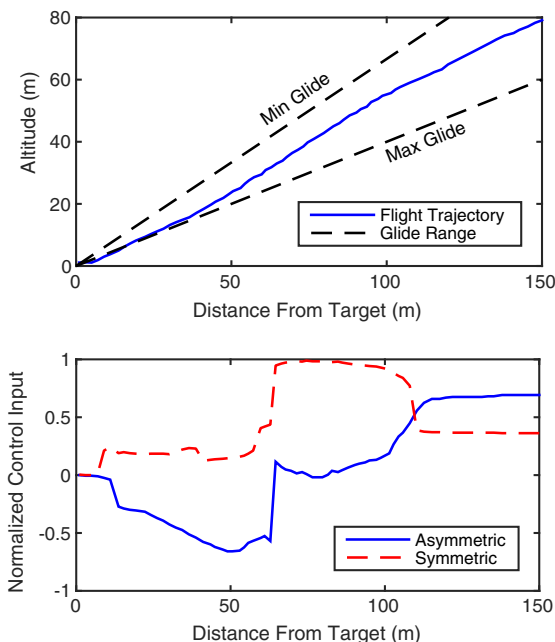

Fig. 17   Altitude (top) and normalized control input (bottom) vs instantaneous distance from the target for example autonomous flight trajectory.
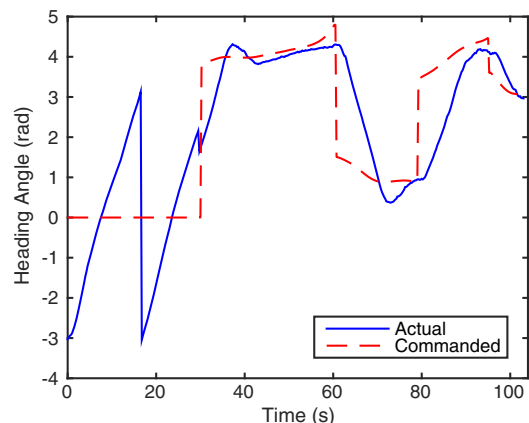

Fig. 18   Actual and commanded vehicle heading angle vs time for example autonomous flight trajectory.
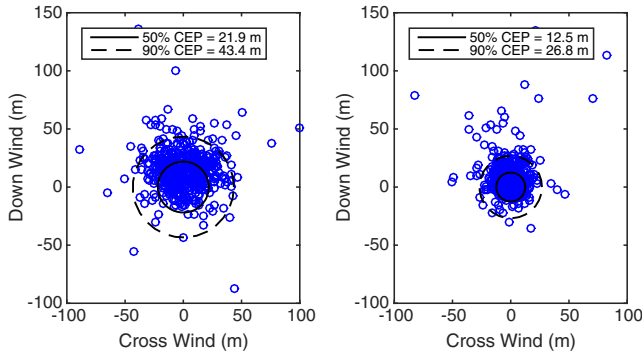
Fig. 19 Simulated landing dispersion for lateral only control (left) and combined lateral and longitudinal control (right) using upper-surface canopy spoilers.
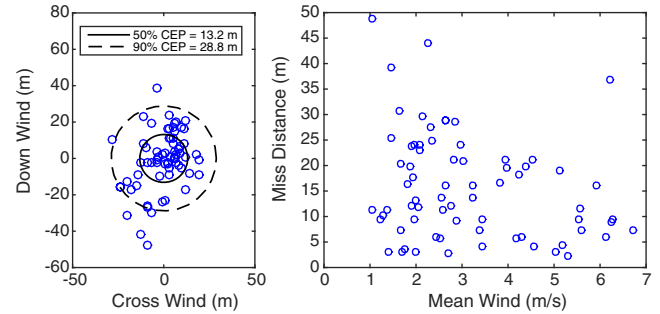

Fig. 20 Experimental landing dispersion (left) and calculated miss distance vs mean wind speed (right) for 70 autonomous flight tests using upper-surface canopy spoiler for combined lateral and longitudinal control.

onboard flight computer was activated to begin its autonomous guided descent and landing. Although the exact starting point for each autonomous flight varied substantially, all flights were initiated sufficiently upwind of the desired impact point to ensure that sufficient reserve exists for the vehicle to reach the target area. Miss distances were calculated based on the GPS location of the intended target and the coordinates of the parafoil and payload system immediately following impact.

Figure 20 shows the resulting experimental landing dispersion using combined lateral and longitudinal control, where values for the 50 and 90% CEP regions were calculated to be 13.2 and 28.8 m, respectively. Figure 20 also presents a comparison of landing error with mean wind speed providing a relative indication of the atmospheric conditions encountered during each flight. Mean wind speeds were computed offline by averaging the magnitude of the estimated wind vector at each GNC update cycle immediately following initialization through vehicle touchdown. Additionally, Table 1 provides a tabulated summary of both the simulation and experimental flight-test performance, including comparison with lateral only and combined lateral and longitudinal control cases.

Based on the simulation and experimental landing accuracy results, it is clear that upper-surface canopy spoilers are a viable control mechanism for both lateral and longitudinal control during autonomous flight. Simulation and experimental 50% CEP values are very consistent at 12.5 and 13.2 m, respectively, indicating excellent model agreement with the physical test vehicle. Also, mean wind speeds shown in Fig. 20 ranged from very calm conditions with almost no wind to relatively high wind speeds approaching the maximum flight speed of the vehicle (~7 m/s). In either case, the autonomous system demonstrates the capability to leverage the added control authority of the combined lateral and longitudinal control logic to compensate for such disturbances. Although Fig. 20 indicates a slight increase in miss distance at lower mean wind speeds, this result is primarily due to light and variable wind conditions where the wind direction is constantly changing. These conditions represent some of the more challenging atmospheric conditions for any autonomous airdrop system.

VI. Conclusions

From the simulation results and 70 autonomous landings in flight tests, it has been shown that upper-surface canopy spoilers are a viable mechanism for lateral and longitudinal control of parafoil and payload aircraft. Using a specialized guidance, navigation, and control (GNC) algorithm combined with longitudinal control via symmetric deflection of the upper-surface spoilers, a substantial reduction in miss distance by nearly a factor of 2 was noted in both high-fidelity computer simulation and experimental flight tests using a small-scale test vehicle. Median miss distance for all 70 autonomous landings using upper-surface spoilers for both lateral and longitudinal control was reported as 13.2 m, a 49% reduction when using upper-surface spoilers for lateral steering only.

Advantages of the upper-surface canopy spoiler control mechanism include its ability to create significant changes in lateral and longitudinal vehicle dynamics without the need for complex rigging changes or additional actuators. For most parafoil aircraft, the control lines and actuators necessary for trailing-edge brake deflection can simply be repurposed to control the opening of each upper-surface spoiler slit, enabling a powerful mechanism for improved control authority and landing accuracy. In the current work, a small-scale parafoil and payload test vehicle is used to demonstrate the resulting control authority and relative landing accuracy improvement when using upper-surface canopy spoilers. However, the principles governing operation of the upper-surface spoiler control mechanism are directly applicable to larger, full-scale systems. In fact, lower actuation forces are needed to deform small portions of the upper canopy surface when compared to that of conventional trailing-edge brakes, offering the possibility for significant savings in actuator size, weight, and cost for larger autonomous systems.

Table 1 Simulation and experimental landing accuracy statistics

Metric description	50% CEP, m	90% CEP, m	Reduction in 50% CEP, %
Simulation: lateral only (500 landings)	21.9	43.4	— —
Simulation: combined lateral/ longitudinal (500 landings)	12.5	26.8	43
Flight tests: lateral only (31 landings)	25.6	51.0	— —
Flight tests: combined lateral/ longitudinal (70 landings)	13.2	28.8	49

References

[1] Gavrilovski, A., Ward, M., and Costello, M., "Parafoil Control Authority with Upper-Surface Spoilers," *Journal of Aircraft*, Vol. 49, No. 5, 2012, pp. 1391–1397.
doi:10.2514/1.C031685

[2] Bergeron, K., Fejzic, A., and Tavan, S., "AccuGlide 100: Precision Airdrop Guidance and Control via Glide Slope Control," *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2011-2530, May 2011.
doi:10.2514/6.2011-2530

[3] Ward, M., and Costello, M., "Adaptive Glide Slope Control for Parafoil and Payload Aircraft," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 4, 2013, pp. 1019–1034.
doi:10.2514/1.59260

[4] Slegers, N., Beyer, 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.
doi:10.2514/1.32099

[5] Moore, J. E., "Drogue Assisted Variable Glide Slope Control," *AIAA Aerodynamic Decelerator Systems (ADS) Conference*, AIAA Paper 2013-1380, March 2013.
doi:10.2514/6.2013-1380

[6] Deazley, F., and Hilton, S., "Ram Air Parachute Variable Trim Platform," U.S. Patent 13/429,939.

[7] Bergeron, K., Ward, M., and Costello, M., "Aerodynamic Effects of Parafoil Upper Surface Bleed Air Actuation," *AIAA Atmospheric Flight Mechanics Conference*, AIAA Paper 2012-4737, Aug. 2012.
doi:10.2514/6.2012-4737

[8] Ward, M., Costello, M., and Slegers, N., "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

[9] Ward, M., "Adaptive Glide Slope Control for Parafoil and Payload Aircraft," Ph.D. Dissertation, Georgia Inst. of Technology, Atlanta, 2012.

[10] Goodrick, T. F., Pearson, A., and Murphy, A. L., Jr., "Analysis of Various Automatic Homing Techniques for Gliding Airdop Systems with Comparative Performance in Adverse Winds," *AIAA 4th Aerodynamic Deceleration Systems Conference*, AIAA Paper 1973-0462, May 1973.

[11] Jann, T., "Advanced Featured for Autonomous Parafoil Guidance, Navigation, and Control," *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2005-1642, May 2005.
doi:10.2514/6.2005-1642

[12] Calise, A. J., 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

[13] Carter, D., George, S., Hattis, P., Singh, L., and Tavan, S., "Autonomous Guidance, Navigation, and Control of Large Parafoils," *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2005-1643, May 2005.
doi:10.2514/6.2005-1643

[14] 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," *19th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2007-2514, May 2007.
doi:10.2514/6.2007-2514

[15] 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," *20th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2009-2981, May 2009.
doi:10.2514/6.2009-2981

[16] Kaminer, I. I., and Yakimenko, O. A., "Development of Control Algorithm for the Autonomous Gliding Delivery System," *17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, AIAA Paper 2003-2116, May 2003.
doi:10.2514/6.2003-2116

[17] 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

[18] Dubins, L. E., "On Curves of Minimal Length with a Constraint on Average Curvature, and with Prescribed Initial and Terminal Positions and Tangents," *American Journal of Mathematics*, Vol. 79, No. 3, 1979, pp. 497–516.
doi:10.2307/2372560

[19] Simon, D., "Optimal State Estimation: Kalman, H_∞, and Nonlinear Approaches," Wiley, Hoboken, NJ, 2006, pp. 123–145.