AIAA Infotech@Aerospace Conference
and
AIAA Unmanned...Unlimited Conference
6 - 9 April 2009, Seattle, Washington

AIAA 2009-2021

Autonomous Guidance and Control of the Flying Fish Ocean Surveillance Platform

Ryan D. Eubank* and Ella M. Atkins†
University of Michigan, Ann Arbor, MI, 48109

Daniel Macy‡
Karem Aircraft, Lake Forest, CA, 92630

The Flying Fish unmanned aerial system is designed to perform persistent surveillance on the open ocean by gathering data as a drifting surface buoy and periodically repositioning via self-initiated flight. The robust stable airframe supports an avionics package with inertial navigation, pressure, and ultrasonic range sensors interfaced with an embedded Linux computer and extended-range modem to a ground station. Flying Fish maintains a preset watch area, defined by center coordinates and a radius, through autonomous flight/drift cycles. This paper describes the multi-mode guidance logic and gain-scheduled control system that enables fully-autonomous operations. The mode-switching logic relies on geometric and temporal rules to select the appropriate flight mode, and intercept-guidance to ensure the aircraft properly flies across its watch circle. The control system relies on traditional decoupled proportional-derivative control laws with a mode-based gain scheduling scheme. Flying Fish has been extensively flight-tested. This paper presents results from early flights to autonomous watch circle maintenance operations.

Nomenclature

u	=	airspeed
z	=	altitude
ϕ	=	roll angle
θ	=	pitch angle
ψ	=	yaw/heading angle
δ_a	=	aileron position
δ_r	=	rudder position
δ_e	=	elevator position
δ_t	=	throttle position

I. Introduction

THE Flying Fish Unmanned Aerial System[1] (Fig. 1) is designed to serve as an unanchored ocean buoy capable of persisting in its deployed region despite currents and wind/wave disturbances. Unlike many UAS of comparable scale, Flying Fish operates without the need for conventional runways or launch apparatus, instead operating out of any body of water with sufficient expanse for takeoff and approach to landing. The ultimate mission of Flying Fish is to persist indefinitely in a region, with a "drop and go" strategy that ultimately can allow a support ship to deploy the vehicle unattended long-term. While most current unmanned aerial systems terminate guidance and control logic shortly after each landing, truly

Figure 1: Flying Fish Landing, Monterey Bay, CA[1]

* Graduate Research Assistant, Department of Aerospace Engineering, 2016 FXB, Student Member.
† Associate Professor, Department of Aerospace Engineering, 3009 FXB, Associate Fellow.
‡ Engineer, 1 Capitol Drive.

Copyright © 2009 by Ryan Eubank and Ella Atkins. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission.

unattended operations will require Flying Fish to execute numerous consecutive flights without operator intervention. To achieve this goal, Flying Fish must not only robustly compute and follow its flight plans but it also must self-manage all decisions through an always-on guidance strategy. With always-on guidance, Flying Fish computes, initiates, and adjusts complete flight sequences as required for its persistent surveillance mission.

The history of manned seaplanes dates to nearly the dawn aviation, but examples of unmanned seaplanes are sparse and self-initiated autonomy remains a largely unexplored UAS research area. Perhaps the most notable precursor to this work is Pisanich and Morris' Autonomous Cargo Amphibious Transport (ACAT) project.[2] ACAT yielded a novel UAS design capable of autonomous takeoff, waypoint flight, and auto-landing to paved runways and smooth water surfaces. The ACAT system was tested only in flat water conditions, with significant development required to operate in the variable ocean environment. Its autonomous flight capabilities were also derived from predefined trajectory following and initiation rather than response to mission requirements. Two notable contemporaries of the Flying Fish project, developed by industry teams, include the Sea-Scout and Gull UAS. Oregon Iron Works' Sea-Scout successfully demonstrated autonomous ocean landing in early 2006,[3] and the Gull UAS of Warrior Ltd. is approaching, for moderate sea conditions, the level and type of autonomy achieved by the ACAT team.[4] Both of these aircraft are of a decidedly different class than Flying Fish being heavy gas-powered systems built for higher structural loads and for carrying heavier sensor payloads such as LIDAR range-finders. The Flying Fish design is focused on self-renewable operations and a careful balance of airframe robustness and weight/efficiency. Our payload will necessarily be lightweight, with current research focused primarily on autonomy and characterization of the ocean environment.

Details of airframe design, and construction, and simulation in an ocean environment are presented elsewhere.[1,5] The avionics system, also presented previously, is summarized in Fig. 2. A Microbotics MIDG-II inertial navigation system (INS) provides 6-DOF filtered state estimate from a set of three-axis gyros, accelerometers, and magnetometers in combination with GPS. The sensor system also includes two pressure transducers for redundant airspeed measurements and a waterproof ultrasonic altimeter for low altitude surface ranging. A Gumstix 400MHz embedded Linux computer provides onboard computational power, working in concert with an Atmel-based Robostix expansion card. The integrated system provides analog-to-digital conversion, multi-port serial communications, 802.11b Wi-Fi, and large-volume data storage via Compact Flash. A Microbotics Servo Switch Controller provides a fail-safe switch between computer and radio-controlled pilot servo commands as well as a means for the computer to log pilot inputs. Long-range communication and ground station telemetry are enabled with a Digi XTend 900MHz radio modem. The embedded Linux computer runs a customized version of the University of Michigan's open-source flight management system,[6] with special-purpose guidance and control software developed for Flying Fish.

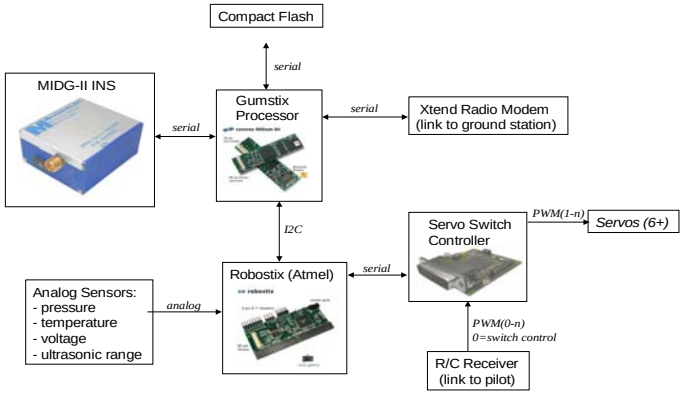

Figure 2: Flying Fish Avionics Block Diagram

This paper will describe the implementation of the Flying Fish guidance and control system starting below with an overview of the decoupled, gain-scheduled, control laws followed by details of the mode-switching goal-based guidance strategy. The University of Michigan first demonstrated fully-autonomous flight during a DARPA-sponsored sea trial in California's Monterey Bay in October of 2007. The aircraft “hopped” repeatedly without operator intervention in order to maintain a prescribed watch circle. Subsequent tests have been conducted locally with a second multi-day operation in Monterey Bay in May 2008. Below we summarize the evolution of flight tests and performance improvements realized over the course of testing to-date. The paper concludes with a summary of first-generation system achievements and a discussion of ongoing research and development activities underway to deploy a truly persistent platform with solar energy harvesting capability in summer of 2009.

II. Independent Gain-Scheduled Control Laws

The Flying Fish airframe has a conventional aerodynamic configuration with a twin vertical tail design. The

structure is strong but relatively heavy to maximize resilience to the harsh ocean environment. The addition of the twin pontoon flotation system provides hydrodynamic stability and minimizes drag during takeoff, but adds appreciable drag during flight. During initial flight tests, the vehicle proved to be naturally stabile and displayed relatively docile response to small-magnitude control surface inputs. The control surfaces sizing is generous, however, enabling appreciable aerodynamic forces to be applied if desired. Since the vehicle need not perform aerobatic maneuvers and is inherently stable the decision was made to implement a decoupled lateral/longitudinal controller to minimize controller complexity. Decoupling the dynamics allows independent control development and execution for straight flight and gentle turning behaviors. Examples in nearly every modern aircraft control textbook as well as in academic literature[7] provide extensive precedent for this treatment. The decoupled control model was extended to the takeoff dynamics through the additional of open-loop behaviors and rule-based guidance logic, as discussed below. For our analysis, it was assumed that the digital control system operated at a sufficiently high frequency to allow continuous time dynamics and control equations to approximate system performance. Notably, the high frequency INS filter operates independent of the control system so that even for slower control updates the sensor feedback suffers minimal aliasing and provides the control system with undistorted and timely state reports.

Under these assumptions a set of proportional-derivative (PD) control laws were written for the decoupled continuous dynamics and applied to the full flight regime using a gain-scheduling scheme. The controllers were tuned for each of several flight modes discussed under guidance below. In this section the individual longitudinal and lateral control laws will be described as will the open-loop pitch controller for takeoff.

A. Longitudinal Control Law

The longitudinal control law is shown in Figure 3. Inputs include the desired pitch angle θ_d and pitch rate θ_d and the output is elevator position δ_e. The desired pitch rate is set to zero to damp the rate of control reaction. The desired pitch angle is determined in real-time by the guidance system to satisfy the requirements of each flight mode, with an airspeed reference for climbing and descending flight and an altitude reference for cruise flight. Note that this paper treats this scheme as a guidance law driving a PD controller, but it is structurally identical to an inner-loop/outer-loop control scheme analogous to that applied by Kingston and Beard[7] for small-scale UAS. The distinction is largely semantic, but from an organizational perspective it is easier to classify the mode-specific reference generation loops as intercept-guidance laws and incorporate them into the mode-switching strategy. Throttle position, δ_t, is also handled by the guidance logic and is typically set to a predetermined value for each guidance mode.

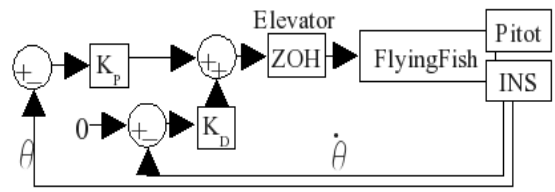

Figure 3: Longitudinal Control Loop

B. Lateral Control Law

Flying Fish lateral dynamics involve slightly more complicated interactions. The vehicle's large twin rudder surfaces have the potential to generate non-negligible adverse roll moments. However, comparably large aileron surfaces provide sufficient control authority to compensate the adverse roll and still meet roll performance goals. These conditions led to a further decoupling of the lateral control into separate yaw and roll loops. The high roll-authority controller can achieve and hold a desired bank angle regardless of rudder action. Further, separating roll control from the yaw/heading system reduces the possibility of trajectory tracking errors causing divergent roll behavior. The lateral control laws, having the same form as the longitudinal control law, are shown in Figure 4. The roll control loop takes desired roll angle ϕ_d and desired roll rate ϕ_d and generates aileron positions δ_a. The inputs to the yaw control law are the desired heading angle ψ_d and desired yaw rate ψ_d while the output is rudder position δ_r. Again the guidance system provides the real-time reference trajectories via an intercept-guidance system and the reference rates are set to zero to damp control response speed.

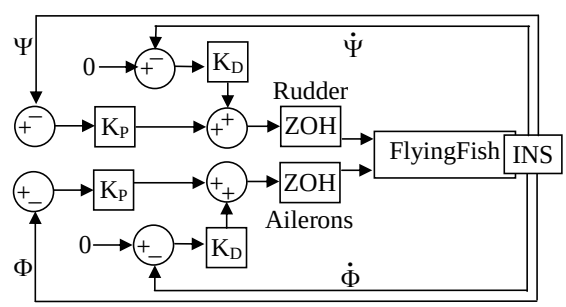

Figure 4: Lateral Control Loop

C. Open Loop Control

A major challenge for flight control of an ocean-based UAS is the complex hydrodynamic interactions that dominate takeoff. Flying from the ocean's surface the vehicle will traverse a set of superimposed wave forms with

amplitudes and periods that potentially differ over orders of magnitude. The largest slowest waves, referred to as swell, have amplitudes on the order of a several meters and periods approaching 10-30 seconds. More difficult to handle, from a dynamics perspective, are the moderate and small amplitude waves, or chop, which range in amplitude from a few centimeters to nearly a meter with very short periods and fast dynamics. In the case of Flying Fish the lateral dynamics are actually slower and have reduced control authority on the water because the rudder effectiveness is greatly diminished below flight speeds and the low and wide stance of the airframe tends to opposes rolling moment. Difficulties arise, however, in the longitudinal dynamics as hydroplaning to takeoff speeds subjects the aircraft to fast impulsive pitch disturbances that are difficult to characterize and compensate. As a result the longitudinal control system is unable to guarantee pitch stability when ramping up to takeoff speeds, at least with our simple decoupled PD control strategy. Pitch control was therefore driven with open-loop commands for the takeoff roll. The Flying Fish team captured and characterized the inputs of a human test pilot for a variety of takeoff conditions. The vehicle was then trimmed and balanced to transition to and maintain a stable hydroplane with only throttle input. An open-loop script of elevator commands was developed to handle the required pitch inputs for takeoff based on time, airspeed, and pitch thresholds, as described below.

III. Goal-Driven Guidance and Mode Switching

The unusual mission of Flying Fish created an interesting guidance challenge for the development team. In order to simplify the guidance task to the extent possible we decomposed operations into a mode-based algorithm. An overview of guidance modes are presented below followed by specifics for each mode.

A. Multi-Mode Guidance Overview

Flying Fish satisfies its persistent surveillance mission through alternating periods of passive drift and repositioning flight. Despite uncertainty, especially of drift behavior, the mission remains inherently cyclic and the robust satisfaction of a single cycle is extensible to continuous operation. A single cycle can be decomposed into the following segments: drift, acceleration, takeoff, climb, cruise, descent, landing, drift. In each mode a desired trajectory can be defined (Fig. 5), along with switching strategy between modes. From a guidance and control perspective each segment of the flight-drift operation can be implemented using straightforward techniques whereas the complex behavior resulting from the sequencing of these steps might otherwise be very difficult to automate robustly.

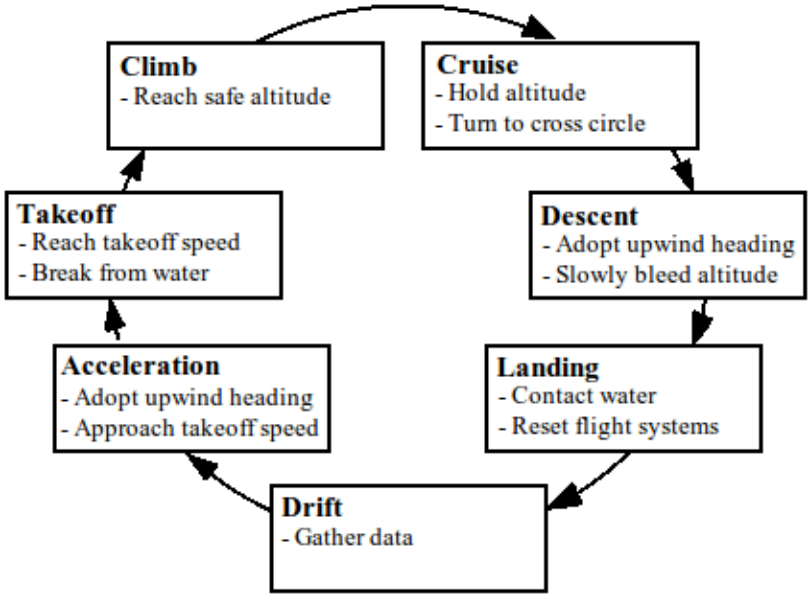

Figure 5: Mode-Based Guidance System.

The decoupled PD control laws discussed above were tuned independently for each flight mode based on piloted flight data and assessments of each autonomous flight operation. The guidance logic uses the mode-specific gain sets and trim states to reconfigure the control system when switching between flight modes. Each mode also has a set of state bounds that are used by the guidance logic as mode-transition points, as a means of diagnosing off-nominal/out-of-range behavior, and to set mode-appropriate constraints on reference variables. Finally, each guidance mode has a minimum operation time that prevents disturbances from triggering premature transitions.

B. Guidance Implementation Details

To begin a deployment the watch region (Fig. 6) must first be specified. A watch circle can be set manually and requires specification of the watch circle radius and GPS coordinates of the center of the circle. To facilitate more rapid deployment an automated watch-circle-setting routine was developed for Flying Fish that sets a watch region centered on the aircraft's initial GPS position and then gathers data to develop a baseline environmental

characterization. The initial environmental data include minimum, maximum, and average wind speed and wind heading as well as initial wave statistics from the INS. Once the watch circle is set the vehicle is considered "on station" and begins its first sequence of guidance modes as specified below.

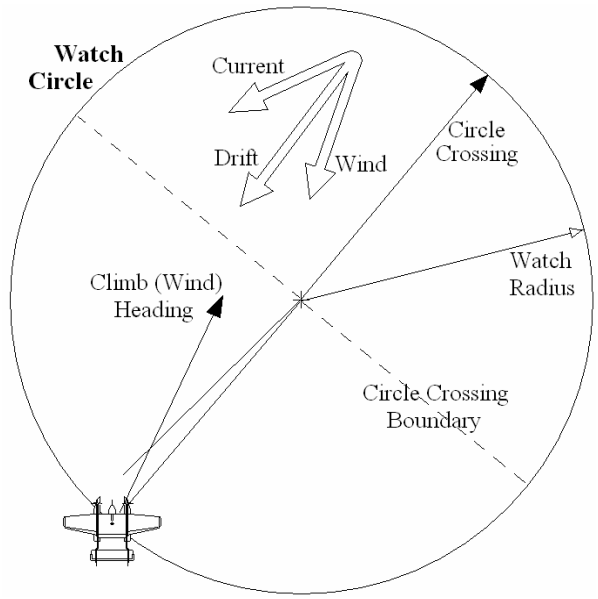

Figure 6: Watch Circle Schematic

1. Drift

This is the only mode in which no control action is exerted. The system commands the control surfaces to predefined trims for the duration of the drift. The system could also shut down servos in this mode to conserve power. In drift mode the guidance system monitors range and direction to the center of the watch circle. When the range exceeds the watch circle radius the guidance system initiates flight with the takeoff logic shown in Figure 7.

2. Acceleration

To take off the vehicle must drive up into a hydroplane (i.e., accelerate onto the pontoon step) and achieve aerodynamic takeoff speed. The open loop longitudinal trim is constant moderate up elevator to counter nose-down moment from hydrodynamic drag. Lateral controllers are given wings-level roll and wind-heading yaw references. The guidance mode transitions to takeoff time when both a minimum operating time has elapsed and a minimum airspeed is achieved. The minimum time is particularly critical to reduce the likelihood of a wind gust triggering premature transition to takeoff.

3. Takeoff

At liftoff speed the autopilot initiates a scripted two-part open-loop elevator command. An initial hard elevator pulse is required to pitch up, breaking the aircraft free from the water. The elevator must then be quickly relaxed as the required break-away elevator will quickly rotate the aircraft to stall. With the elevator near neutral the aircraft gains speed at a slow climb, then entering the climb mode.

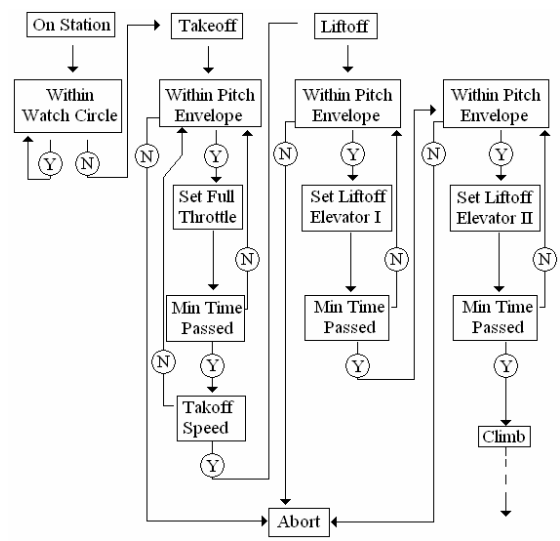

Figure 7: Takeoff Logic Tree

4. Climb

Once the aircraft rotates into a climb without stalling or violating the pitch envelope, the climb guidance mode becomes active. The guidance law adopted for climb generates a desired pitch angle for the longitudinal controller using the reference pitch and airspeed for the climb mode, by:

$$\theta_d = \theta_{ref} - G_p(u_{ref} - u) \qquad (1)$$

Climb continues for a preset time. The final altitude, set to the reference cruise altitude, is based on GPS altitude achieved after this preset time. Note that this altitude is appreciable and based on GPS readings since out of the ultrasonic altimeter's maximum range of 7 meters.

5. Cruise

At the termination of the climb the guidance logic switches into the cruise mode. The cruise altitude-intercept guidance law referenced to level flight gives the desired pitch angle:

$$\theta_d = \theta_{ref} - G_p(z_{ref} - z) + G_d(dz) \qquad (2)$$

To keep the vehicle inside the watch area cruise guidance attempts to cross through the center of the watch circle. On the downwind side of the watch circle center, guidance sets the desired heading to match the direction to the center of the circle. Without commanding a bank angle this guidance strategy results in wings-level flight with the heading-to-yaw control (Fig. 4) performing rudder-based turns. A 2.0deg dead-band is used around the desired heading to avoid overshoot oscillations and a 1.5m radius dead-zone surrounding the watch-center helps avoid hard saturated turns while passing close to the watch-center. Once the aircraft reaches the upwind side of the watch circle

center, it resumes a heading determined to be upwind during its last drift cycle. Note that with small watch circle radii (e.g., 250 m or less) the aircraft begins a gradual descent when crossing to the upwind side of the watch circle center.

6. Descent

Once the vehicle either crosses into the upwind side of the watch circle, or the aircraft reaches a certain distance from the edge of the upwind side of the watch circle (for large watch circle regions), the guidance law transitions to a slow descent towards the far upwind edge of the watch region. The descent mode relies on the climb guidance laws (Eq. (1)) with a -1.0 degree reference pitch and reference airspeed of 1.2 times stall speed. The vehicle conducts a powered descent, with 30% throttle reliably enabling a slow and gentle approach to landing. For descent the desired heading is set back to the original wind-heading to maximize the chance of non-slipping flight, without a sideslip sensor.

7. Landing

When the vehicle is within 7 meters of the surface, the ultrasonic altimeter provides data by which the aircraft can reliably predict contact with the water. We maintain the gentle descent until 0.5m above the water surface, at which time the aircraft is flared and throttle reduced to zero. Upon landing, the guidance system deactivates the flight control system, sets the drift trim state, and begins monitoring range and heading to the watch-center for another drift cycle.

C. Improving Guidance Accuracy

In our initial sea trials, the above mode sequence was executed reliably, but the data indicated significant oscillatory behavior when switching between certain flight modes and the vehicle displayed sluggish turning performance.

Oscillatory mode excitations of the phugoid and short-period longitudinal dynamics resulted from impulsive transitions between modes with different intercept-guidance-laws, specifically climb/cruise and cruise/descent. A trajectory smoothing scheme was implemented to address the fact that this behavior increased tracking error, required additional corrective control, and risked destabilizing the control system. Since the reference variables change between modes the issue could not be addressed by interpolation of the guidance references and instead the output of the guidance laws was matched. This was accomplished by computing the guidance laws of the target mode prior to transition, temporarily suspending intercept-guidance, and interpolating between the last desired state and the target state of the next mode.

Turning performance suffered in the original guidance scheme because the control laws were driven for wings-level slipping turns. This strategy was originally employed for simplicity and safety, avoiding the risk of oscillation and instability in the high control authority roll dynamics. However, the lateral forces that can be achieved by yaw/thrust turning are generally much less than the lateral forces achieved with banked/lift turning for aircraft without very high thrust-to-weight ratios or thrust-vectoring capabilities. As a result circle crossings were inaccurate and tended to accumulated error in the direction of the prevailing wind. To improve turning performance another intercept-guidance law, with 10deg bank-angle saturation, was formulated for a watch-center heading reference:

$$\phi_{d,\,ref} = \mathrm{sgn}(\psi_{ref} - \psi)*\min(10.0\mathrm{deg}, |(1/6)(\psi_{ref} - \psi)|) \tag{3}$$

This guidance law drives the desired bank-angle to achieve turning flight over the center of the watch-circle via a discrete, rate-limiting update given by:

$$\phi_{d,\,0} = 0, \quad i = 1, 2, 3, ... \tag{4}$$

$$\phi_{d,\,i} = \phi_{d,i-1} - (0.5)\mathrm{sgn}(\phi_{d,i-1} - \phi_{d,ref}) \tag{5}$$

$$\phi_{d} = \phi_{d,i} \tag{6}$$

The resulting intercept-law in combination with cruise-mode dead-band rules proved to be accurate and effective. Additionally, the original rudder-based turning strategy remains active approximating coordinated turns for higher efficiency turning.

IV. Test Results

A. Preliminary Flight Testing

Initial airworthiness testing, primary piloted flight experiments, and the controller tuning strategy were documented previously.[1,5] It is appropriated here to briefly review these initial control tuning procedures and to provide the previously unpublished data from that period, showing the piloted flight data that guided the takeoff controller design and an example of the earliest autonomous flights of Flying Fish.

All of the early Flying Fish testing was conducted in local freshwater lakes (Fig. 8). To begin the process of implementing the guidance and control system, trim states and pilot inputs were extracted for different flight regime from a series of piloted test flights. With sufficient baseline data, the aircraft was flown to altitude by a human pilot for a series of controller tuning tests. Once at altitude a controller would be activated and the pilot would surrender command to the computer system. In this manner a series of high altitude tests progressed from controlled level flight through climbing, descending, and turning flight capabilities. Eventually entire hop sequences were linked together, at altitude, by the mode switching guidance system. Stability and robustness were favored for the Flying Fish mission over controller speed and accuracy, and as a result controllers were tuned very conservatively, weighting reduced overshoot and stability over other performance metrics. This decision becomes more obvious when considering a slow-descent ocean landing; the potential landing area is immense and a small heading or lateral position error would be negligible whereas a strong roll correction, pitch overshoot, or any instability near the water would be catastrophic.

Figure 8: Flying Fish, Lake Testing

Once the airborne flight modes were tuned for stability and the high-altitude experimentation possibilities were exhausted, testing moved to the water's surface for taxi and takeoff control development. As outlined in the control section, hydrodynamic interactions during takeoff result in large variable high-speed pitch disturbances (Fig. 9-10) that would be difficult to stabilize without significant additional modeling and a new control strategy. Analysis of a number of flights revealed that the pilot inputs for a well-trimmed takeoff were very steady between the initiation of the "takeoff roll" (throttle up) and the elevator pulse that breaks the vehicle free from the water (Fig. 11). Altitude and airspeed profiles for this flight are presented in figures 12 and 13. This data also demonstrates a right-roll at takeoff (Fig. 9-10), caused by an early stall tendency in the right wing, necessitating closed-loop lateral control for takeoff. Note that control inputs were left in microseconds of PWM signal to provide a common scaling for the

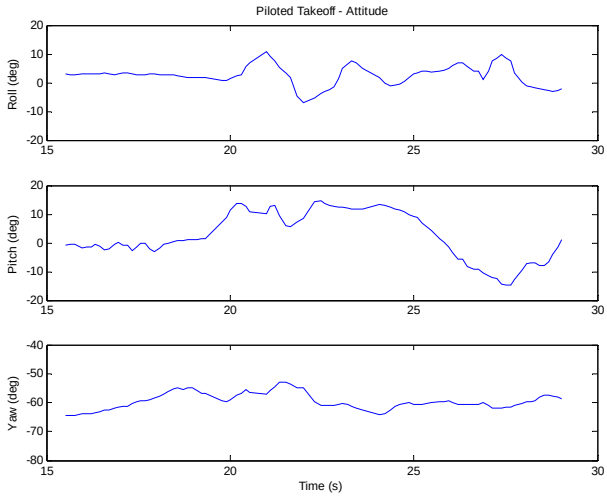

Figure 9: Piloted Flight Data – Attitude

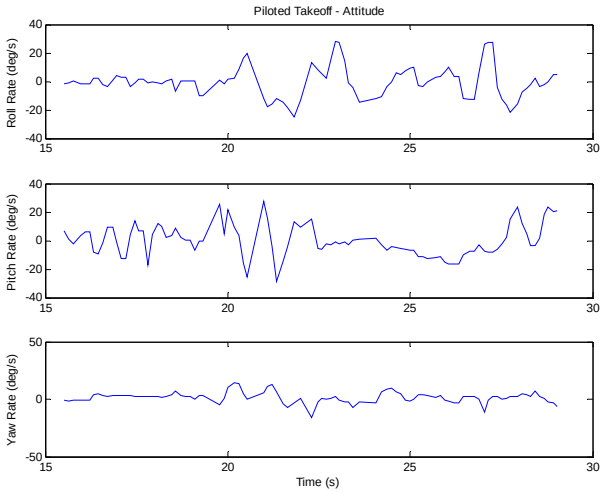

Figure 10: Piloted Flight Data – Rates

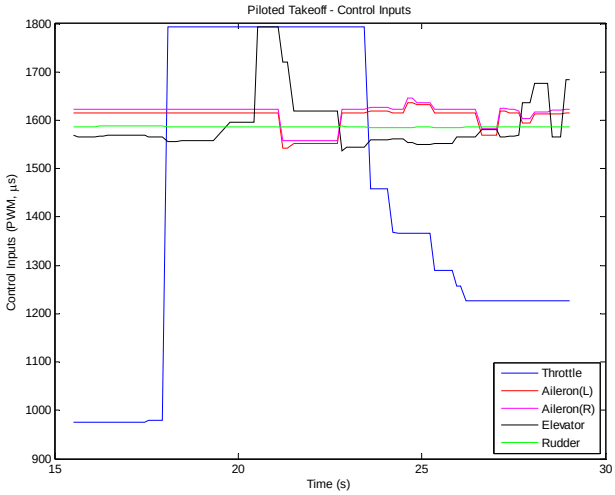

Figure 11: Piloted Flight Data – Control

different control ranges.

With tuned controllers for each flight mode in place the guidance logic was refined, advancing slowly to autonomous hops over the lake. The ground-track, altitude, and airspeed profiles of one of Flying Fish's earliest autonomous hops are presented in figures 14-16. Like the piloted flight data it is based on, the successful takeoff control inputs are steady from the activation of full throttle until the takeoff elevator pulse (Fig. 17) and subsequently relaxed elevator entering the climb stabilizing into the climb mode. The vehicle climbs quickly until initiating a short cruise segment at 136s, climbing slightly during cruise due to the effects of the mode switching. Finally the guidance switches to descent at 139s, climbing rapidly in the first few seconds to reach the desired descent airspeed and the slowly approaching the water. Obvious again are the pitch disturbances during takeoff acceleration and the right-roll tendency (Figs. 18-19).

An important issue uncovered during autonomous flight testing on the lake was that water spray during taxi, takeoff, and landing can easily reach a pitot probe almost anywhere on a vehicle of this scale. Blockages from spray can result in guidance failure (takeoff logic unable to measure liftoff speed) or divergence (airspeed remains constantly low/high, resulting in increasing magnitude pitch references). To-date hydrophobic coatings have been tried with limited success and a compressed-air clearing system has been prototyped but time constraints have precluded comprehensive testing.

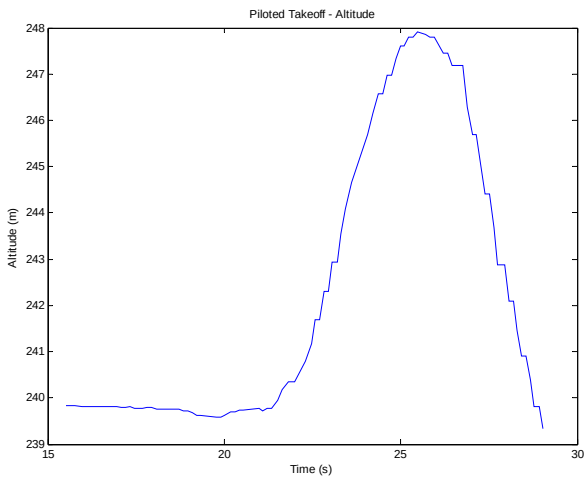

Figure 12: Piloted Flight Data – Altitude

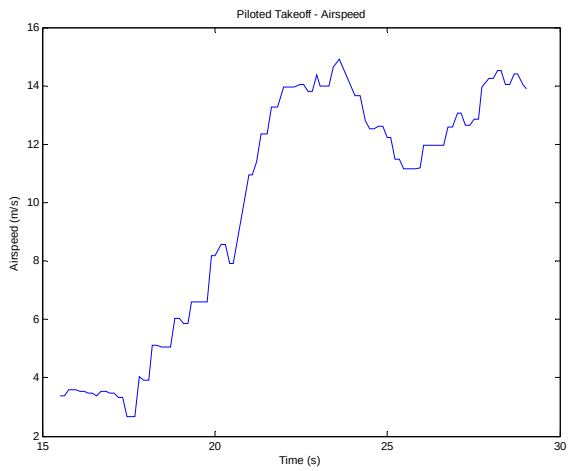

Figure 13: Piloted Flight Data – Airspeed

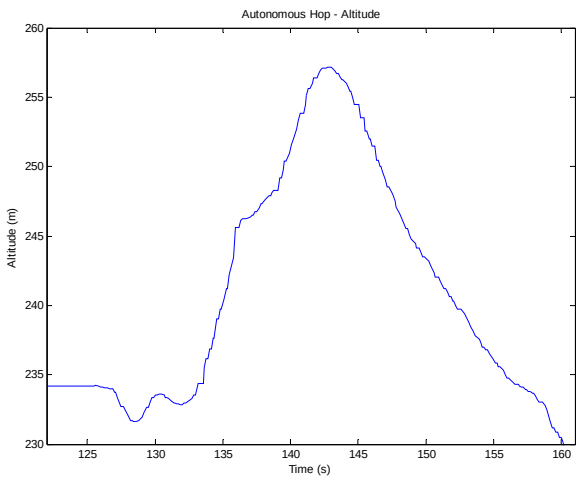

Figure 14: Autonomous Flight – Altitude

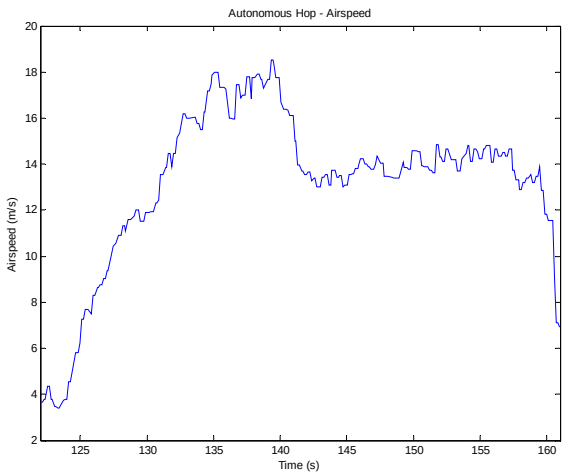

Figure 15: Autonomous Flight– Airspeed

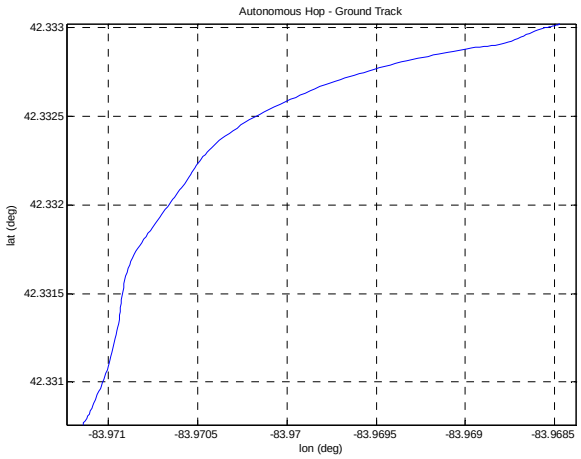

Figure 16: Autonomous Flight – GPS

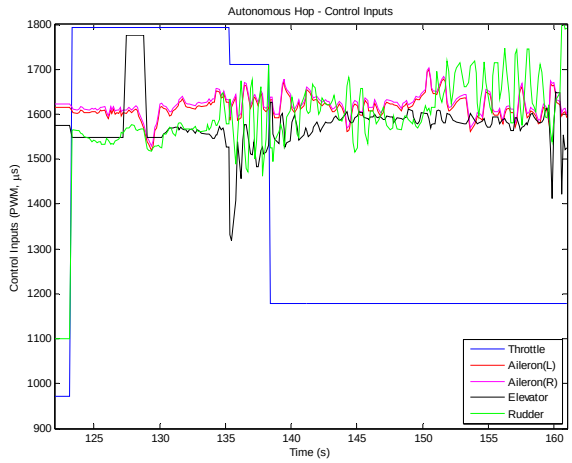

Figure 17: Autonomous Flight – Control

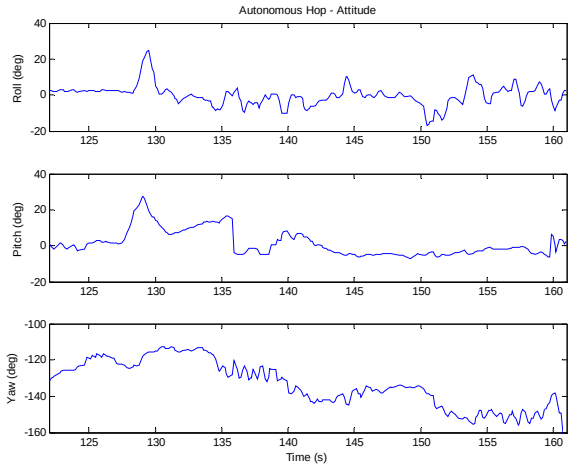

Figure 18: Autonomous Flight– Attitude

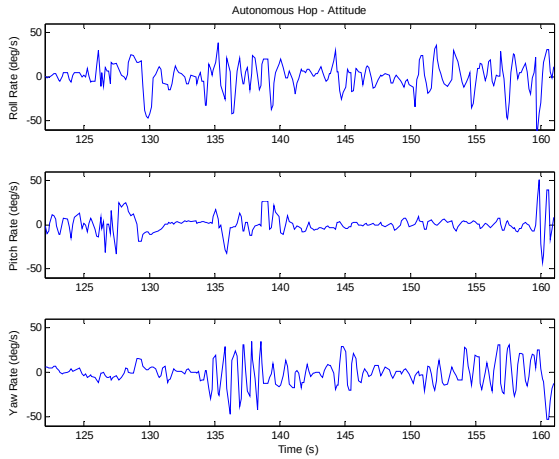

Figure 19: Autonomous Flight – Rates

B. Sea Trials, October 2007

In late October of 2007 the Flying Fish vehicle was deployed for the first time in an ocean environment (Fig. 20). During the first day of the sea-trial piloted flight tests were conducted to demonstrate the airframe's flight capabilities. Testing progressed to minor controller tuning and guidance timing refinement for the new operating environment. Once the vehicle was tuned and trimmed for its target environment autonomous flight tests were conducted (Fig. 21-23). In Figure 22 the effects of the oscillatory excitation from mode-switching can be observed in the large altitude variations immediately after entering cruise at the 383s mark. Following the first autonomy tests Flying Fish was deployed for a watch-circle maintenance test (Fig. 24). The turning controller shortcomings outlined at the end of the guidance section can be seen very clearly in the accumulating error of the GPS ground-track. Extended duration watch-maintenance tests were not feasible with the proof-of-concept vehicle as it lacked the solar-recharge capabilities of the full-scale vehicle that is currently under development. By the end of the 2-day sea trial Flying Fish had flown 27 times, of which 22 of the flights were fully autonomous, in seas ranging from 1-2m with 2-11s periods and 7-8m/s winds.

Figure 20: Flying Fish, Monterey w/Dolphins[1]

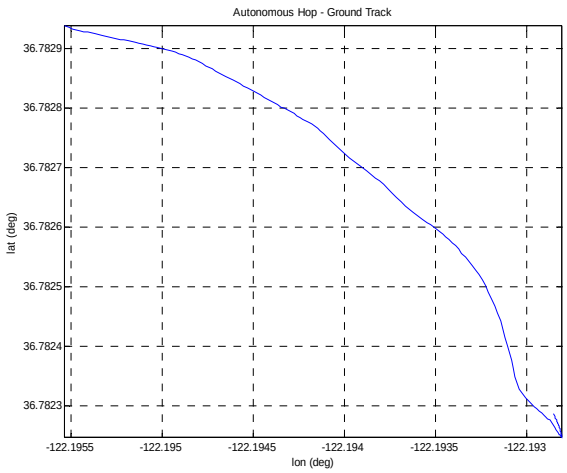

Figure 21: Autonomous Flight – GPS

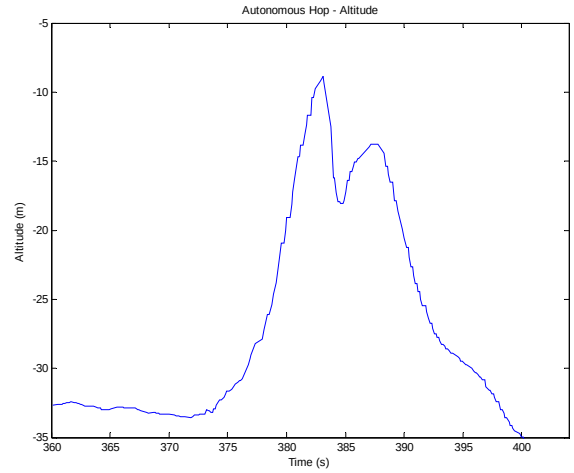

Figure 22: Autonomous Flight– Altitude

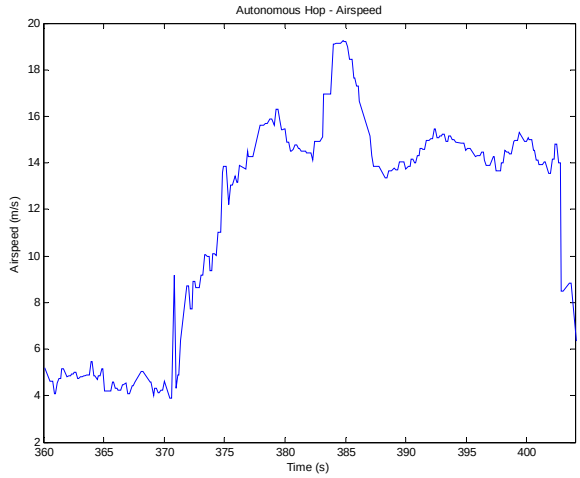

Figure 23: Autonomous Flight – Airspeed

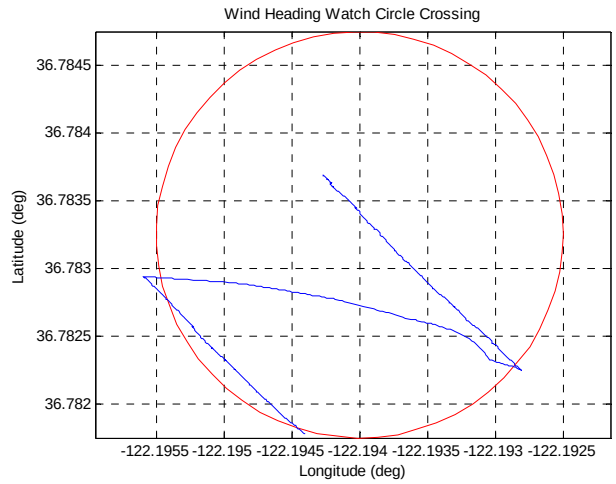

Figure 24: Wind-Heading Watch Crossing

C. Sea Trials, May 2008

The Flying Fish team had a second opportunity to test in Monterey Bay in late May 2008. Unlike the previous sea trials, the May deployment was strictly focused on testing and development. Goals for the three and half ocean testing days included updating the software system for increased performance, refining guidance strategies for more robust watch-maintenance, solving the oscillatory mode excitation problem, achieving increased data rates in order to characterize the ocean surface during drift, and integrating the ultrasonic altimeter to the guidance and control software. Despite issues with pitot probes, an aging airframe, and a near-catastrophic crash after a fail-safe "flight grounding" behavior was inadvertently activated by the loss of a wireless link, the team was able to implemented and test most of the updates. During the first round of test flights software updates yielded increased data rates and execution frequency and a solution to the mode-switch oscillatory mode excitation problem was successfully tested. Figure 25 shows the response of the attitude variables before and after the implementation of the guidance smoothing strategy outlined in guidance section above.

With the guidance-smoothing solution and first round of software updates in place the team turned its attention to the ultrasonic altimeter. The primary challenge for the ultrasonic altimeter was that in saturation the sensor's responses may momentarily mimic unsaturated measurements. Unlike a sensor with a hard saturation limit, at its maximum range the ultrasonic sensor may receive delayed reflections from out-of-range objects and mistakenly interpret these responses as the imaging of an in-range obstacle. As a result the ultrasonic sensor data needs to be constantly assayed for reliability. This was accomplished through a two part filtering strategy, wherein the data was first passed through a moving-window median filter to eliminate most isolated errors and then through a confidence

filter that compared the altimeter response and response gradient to the vehicle state and guidance mode. The result was a binary judgment of the fitness of the altimeter and a smoothed altimeter response signal (Fig. 26). A high-confidence judgment for an altitude at or below 1.0m during decent triggers a pitch-up flare maneuver.

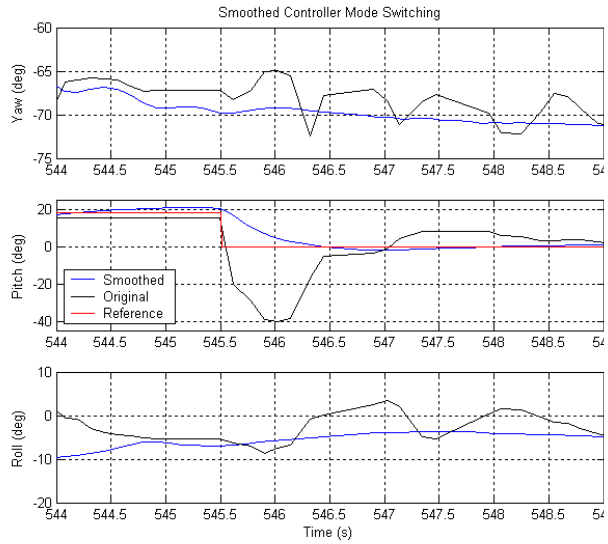

Figure 25: Smoothed Pitch Response

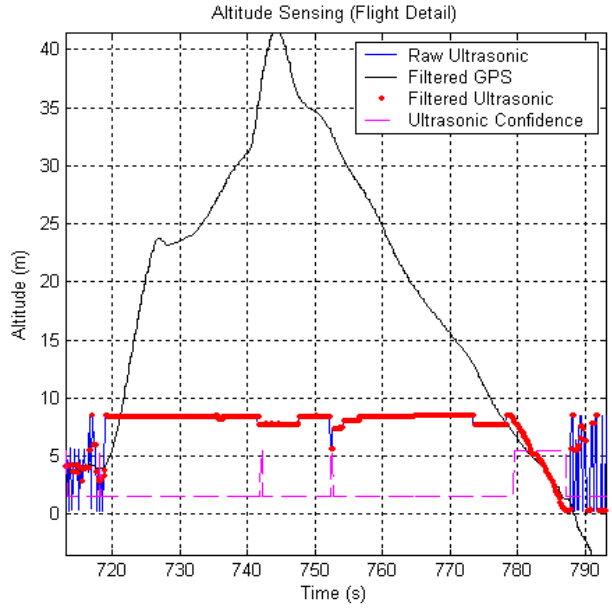

Figure 26: Filtered Ultrasonic Altimeter Response

D. Final Results

Approaching the end of the life-cycle of the proof-of-concept vehicle a final set of tests was organized in late 2008 to address the remaining deficiency in the persistent surveillance system: the need for higher accuracy watch-circle crossing. After identifying a suitably sized test site for watch-circle maintenance experimentation the vehicle was deployed for a series of watch-crossing tests. The reduction in test personnel due to the construction of the full-scale Flying Fish and the late-season Michigan weather made testing very difficult. In fact, the Flying Fish test team observed a novel "inverted" wing icing phenomenon during the last day of testing as the spray raised by the hydroplaning aircraft built up a thick layer of ice over the leading edge and lower surface of the wing and horizontal stabilizer. Despite these difficulties the parameters of the banking intercept-guidance law were adjusted over several flights and a fully-updated guidance system flew directly over the watch-center in the last test of the season (Fig. 27).

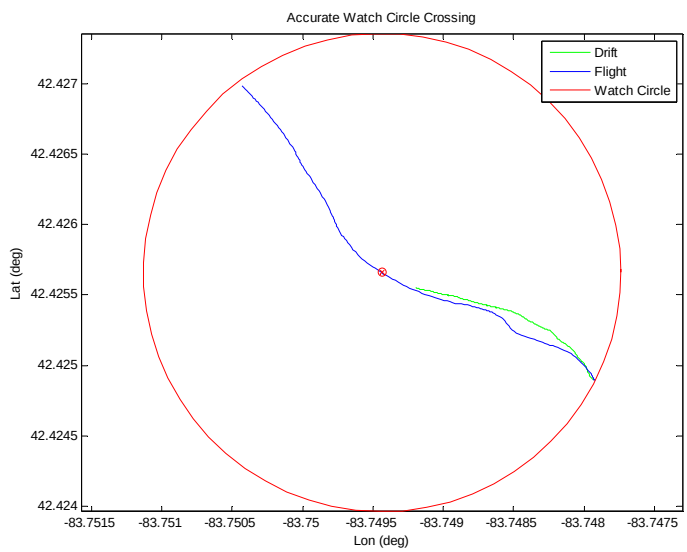

Figure 27: Accurate Watch-Circle Maintenance

V. Conclusion

This paper has presented the gain-scheduled decoupled control laws, multi-mode always-on guidance, and flight test progression of the Flying Fish persistent ocean surveillance UAS. Although unattended flight vehicle operation is challenging, the constrained mission and decomposition of the cyclic mission modes effectively segmented the problem into a number of tractable challenges to which well-understood control and guidance strategies were applied. The result was a first-of-its-kind seaplane UAS with the ability to self-initiate autonomous flight "hops" to

maintain a designated open-water watch circle region.

This work has built a foundation for automation of the next-generation solar-recharging full-scale Flying Fish UAS to be first flown in May 2009. The full-scale system presents a number of new challenges for the guidance and control systems. Foremost are the needs for a balanced energy budget and the ability to survive unsupervised for multiple days in the open ocean. The system will inherit much of the existing guidance and control logic but will also require new algorithms to better incorporate environmental and energy data. Ultimately, additional sensors will be needed to ensure the vehicle is capable of fully unattended/unobserved operation. For example, as the dolphin visitors in Figure 11 illustrate, Flying Fish currently has no notion or means of observing and responding to obstacles or other ocean based systems. It must ultimately at least be capable of identifying ships and boats, as well as being aware of regions of water versus land if operating close to shore.

Acknowledgments

The author thanks faculty, staff, and students of the University of Michigan's Marine Hydrodynamics Laboratory and of the Department of Aerospace Engineering for their collaboration and support. We also thank Moss Landing Marine Laboratories and the Monterey Bay Aquarium Research Institute (MBARI) for supporting our flight tests in Monterey Bay. This work was funded in part under contract N00421-06-C-0094 from the Defense Advanced Research Projects Agency (DARPA). This document has been approved for public release, distribution unlimited.

References

[1] Meadows, G., Atkins, E., Washabaugh, P., Meadows, L., Bernal, L., Gilchrist, B., Smith, D., VanSumeren, H., Macy, D., Eubank, R., Smith, B., Brown, J., "The Flying Fish Persistent Ocean Surveillance Platform," *AIAA Unmanned Unlimited Conference*, AIAA, Seattle, WA, April 2009.

[2] Pisanich, G., Morris, S., "Fielding an Amphibious UAV: Development, Results, and Lessons Learned," *NASA Ames Research Center Publication*, http://ti.arc.nasa.gov/m/pub/407h/0407%20(Pisanich).pdf, NASA, 2001.

[3] "Sea-Scout UAV Information." http://www.oregoniron.com/sea-scout.htm, Oregon Iron Works, Accessed: 13 Apr. 2008.

[4] "Gull UAV Information." http://www.centaurseaplane.com/gull/index.htm, Warrior Ltd., Accessed: 13 Apr. 2008.

[5] Macy, D., Eubank, R., Atkins, E., Bernal, L., Washabaugh, P., Meadows, G., Wild, N., Smith, D., Van Sumeren, H., "Flying Fish: A Persistent Ocean Surveillance Buoy with Autonomous Aerial Repositioning," *AUVSI Conference*, San Diego, CA, June 2008.

[6] Atkins, E., Eubank, R., Klesh, A., "A Reconfigurable Flight Management System for Small-Scale Unmanned Air Systems", *AIAA Infotech at Aerospace Conf.*, AIAA, Seattle, WA, April 2009.

[7] Kingston, D., Beard, R., McLain, T., Larson, M., Ren, W., "Autonomous Vehicle Technologies for Small Fixed Wing UAVs," *Proc. 2nd Unmanned Unlimited Conf.*, AIAA, San Diego, CA, Sept. 2003.