

ABSTRACT

EDWARDS, DANIEL J. Autonomous Soaring: The Montague Cross Country Challenge. (Under the direction of Dr. Larry Silverberg).

A novel method was developed for locating and allowing gliders to stay in thermals (convective updrafts). The method was applied to a 5 kg, glider, called ALOFT (autonomous locator of thermals), that was entered in the 2008 Montague Cross-Country Challenge held on 13-15 June 2008 in Montague, California. In this competition, RC (remote controlled) gliders in the 5 kg class competed on the basis of speed and distance. ALOFT was the first known autonomously soaring aircraft to enter a soaring competition and its entry provided a valuable comparison between the effectiveness of manual soaring and autonomous soaring. ALOFT placed third in the competition in overall points, outperforming manually-flown aircraft in its ability to center and utilize updrafts, especially at higher altitudes and in the presence of wind, to fly more optimal airspeeds, and to fly directly between turn points. The results confirm that autonomous soaring is a bona fide engineering sub-discipline, which is expected to be of interest to engineers who might find this has some utility in the aviation industry.

Autonomous Soaring: The Montague Cross Country Challenge

by
Daniel J. Edwards

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Aerospace Engineering

Raleigh, North Carolina

2010

APPROVED BY:

Dr Larry Silverberg
Committee Chair

Dr Ashok Gopalarathnam

Dr Charles Hall

Dr Mihail Sichitiu

BIOGRAPHY

Dan Edwards started his obsession with airplanes at the age of 12 when his mother saw him eyeing an RC airplane at a model store and said: “Do you want this to be your birthday present?” The seed was sown.

Before turning 14, Dan bought an RC simulator and taught himself to fly. After seeing his first airplane crash as a result of radio interference from a nearby train while under the control of his father, Dan set out to build another. Cheap Fun was born from Stu Richmond plans: a simple gas powered glider, it survived for over 10 years and logged hundreds of flight hours before being rebuilt in 2007.

In college, Dan quickly found a group interested in flying robots, the Aerial Robotics Club. By his sophomore year, Dan was president of the club and oversaw the construction and flight testing of three aircraft. Teaming with a core group of around 6 peers, his fostered tinker sense grew beyond aircraft construction and started leaning toward autonomous system development. Culminating in a 12 ft span, 45 lb behemoth carrying \$15,000 worth of avionics and payload, Dan grew up with the club and still volunteers as a graduate advisor.

One summer as a coop with the Tactical Electronic Warfare Division at the US Naval Research Laboratory, Dan took his hobby-work on Autonomous Soaring and proposed it for a basic research project. In the meantime, Dan received an offer to intern at Dryden Flight Research Center, Rosamond, CA working under Michael Allen on his autonomous soaring project. Upon returning to NRL, the proposed project was picked up as a fully funded 3-year effort. The dissertation that follows is now part of this history.

ACKNOWLEDGMENTS

The author would like to thank Adam Propst and Chris Bovais, as well as the project sponsors: US Naval Research Laboratory, NASA Dryden Flight Research Center, Jim D. Gray and Associates, Dean Gradwell, Wolfgang Schreiner, Tri-M Systems, and Phytex America. Additionally, many thanks to Michael Allen, Brady Baggs, Matt Hazard, Dave Burke, and JB Scoggins for your support, advice, and friendship. Thanks also to Dr. Larry Silverberg, Dr. Charles Hall, Dr. Ashok Gopalarathnam, and Dr. Mihail Sichitiu. Most especially, thanks to my patient wife Mary.

TABLE OF CONTENTS

LIST OF TABLES.....	v
LIST OF FIGURES.....	v
Introduction.....	1
Updraft Identification	2
Centroid Method	3
Evolutionary Search Method	5
Simultaneous Iteration Method.....	5
Guidance Commands.....	8
Deciding to orbit in a thermal or to leave it.....	9
Deciding how to orbit	11
Selecting airspeed	13
Deciding how to search for thermals	13
Flight Testing.....	14
The Montague Cross Country Challenge	20
Final Remarks and Conclusions	25

LIST OF TABLES

Table 1: Competition results.....	25
-----------------------------------	----

LIST OF FIGURES

Figure 1: Speed ring -- altitude versus updraft velocity (netto m/s).....	10
Figure 2: Maintain turn direction approach.	12
Figure 3: Switching turn direction.	12
Figure 4: ALOFT hardware.	15
Figure 5: Sink polar of the SBXC at 5.56 kg.....	16
Figure 6: Evolutionary Search center point updates on an error surface.....	18
Figure 7: Simultaneous Iteration center point updates on an error surface.	19
Figure 8: 2008 Montague teams and their aircraft (photo courtesy Ron McElliot).	21
Figure 9: Course on Day 1, turn point sequence 1-2-3-2-4-5-6-2-1.....	22
Figure 10: ALOFT altitude versus time, Day 2.....	23
Figure 11: ALOFT entering, orbiting, and leaving an updraft.	24

Introduction

The general idea behind soaring flight is to use air motion to offset an aircraft's inherent sink rate. Manned gliders routinely stay aloft for several hours after being towed to an initial altitude, often out-climbing their starting altitude. A novice pilot can be taught to soar with little training. The pilot only needs to learn how to recognize the signs of upward air motion [1] – but can a machine do the same? The challenge lies in transforming in situ learned behaviors into a form that a computer understands. The goal of this work was to demonstrate the feasibility of autonomous soaring and ultimately to bring these results to the attention of engineers who might find some utility for autonomous soaring in the aviation industry and further develop the necessary tools.

Previous work in autonomous soaring suggests its feasibility. Allen performed simulations suggesting endurance improvements of 8 to 14 hours using autonomous soaring techniques [2]. He then developed and demonstrated the Centroid Method to locate, in flight, thermal centers [3]. Wharington proposed a neural-network approach based on reinforcement learning [4]. A recursive method using an unscented Kalman filter was recently developed by Hazard [5], significantly reducing the computational requirements for thermal identification. Several papers have also given optimized strategies for selecting airspeed [6-9] and optimum turn radii [10]. These strategies are directly applicable to autonomous soaring. The number of flight-tests with autonomous soaring is very small; among them the most notable are Allen [3] and Edwards [11].

Section II describes how to determine the location and parameters of a thermal. Allen's Centroid Method [3] is described first. Then improved thermal identification algorithms that

were implemented on ALOFT are described. Section III describes the principle guidance decisions that are necessary in autonomous soaring, namely, how and when to orbit in a thermal and how to search for thermals. Section IV describes the flight testing of ALOFT, which logged 164 flights and 71 hours of flight time. Section V describes ALOFT's performance in the Montague Cross-Country Challenge and the paper ends in section VI with final remarks and conclusions.

Updraft Identification

This section describes how updraft velocity is measured and then how this data is used to estimate the updraft velocity distribution.

The vertical velocity of a local parcel of air is estimated by analyzing the work done on the aircraft by both wind loading and drag. The aircraft is modeled as a point mass acted on by a vertical wind force and a horizontal drag force. The work done on the aircraft is

$$\begin{aligned}
 U_W + U_D &= E - E_0 \\
 E &= \frac{1}{2}mv^2 + mgh
 \end{aligned}
 \tag{1}$$

where U_W is work done by the wind, U_D is energy done by drag, m is mass, v is aircraft speed, and h is altitude. In Eq. (1) $v^2 = v_x^2 + \dot{h}^2$ where v_x is horizontal velocity. Dividing by mg and differentiating with respect to time to give the terms units of velocity, we get

$$\frac{\dot{U}_W}{mg} + \frac{\dot{U}_D}{mg} = \frac{v\dot{v}}{g} + \dot{h}
 \tag{2}$$

In Eq. (2), h and v are measurements from vehicle telemetry and $-\dot{U}_D/mg$ is sink rate due to drag, which is typically a function of velocity and bank angle. The sink rate can be measured off-line by a series of constant speed flight conditions that relate it to the aircraft's (measured) speed [12]. The measurements and calculations in Eq. (2) are used by a netto variometer [6] to estimate the instantaneous updraft velocity term $v_{ud} = \dot{U}_W/mg$. Note that the netto variometer, since it measures the vertical motion of air mass, estimates the local air mass vertical motion better than the more common total energy variometer.

Updraft parameters can be identified recursively or in batch. The batch method is preferred because batch data can be smoothed with no filter time-lag, as explained later. Care must still be taken to avoid a batch that is so long that it encompasses multiple updrafts. Recursive methods are generally more computationally efficient [5] and may be beneficial in low-power applications of autonomous soaring techniques; they were not considered in this research.

Centroid Method

Consider first the method by Allen [3] in which a thermal is modeled in a wind-corrected coordinate system that moves with the free stream air velocity. The updraft velocity distribution is assumed to be normally distributed and to be circular in cross-section, written

$$v_{ud}(x, y) = W e^{-\left(\frac{D}{R}\right)^2} \quad D^2 = (x - x_c)^2 + (y - y_c)^2 \quad (3)$$

where W is the vertical velocity of a thermal's center, R is characteristic updraft radius, and D is the distance between any point (x, y) and the thermal's center (x_c, y_c) . The interest lies in

determining W , R , x_C , and y_C from the measurement data $v_k = v_{ud}(x_k, y_k)$ ($k = 1, 2, \dots, N$). For rudimentary guidance, only the center point (x_C, y_C) is needed, so an algorithm that determines the center point is of primary interest. The data can be measured over a time period in which the aircraft completes about two revolutions. With this amount of data, the center point can be approximated coarsely by averaging the positions with a weighting that is equal to the squares of the updraft velocities. Thus,

$$x_C = \frac{\sum_{k=1}^N x_k v_k^2}{\sum_{k=1}^N v_k^2}, \quad y_C = \frac{\sum_{k=1}^N y_k v_k^2}{\sum_{k=1}^N v_k^2} \quad (4)$$

Although the center averaging approximation is coarse, Eq. (4) is insensitive to the characteristically large variations in the updraft velocities within a thermal [14]. Indeed, Eq. (4) is the solution to weighted minimization of $J = \sum_{k=1}^N v_k [(x_k - x_c)^2 + (y_k - y_c)^2]$, in which the data is weighted by the square of the measured updraft velocities. Equation (4) does not account directly for a thermal's lateral drift (which is different than the free stream velocity), so the center point needs to be updated in time. The thermal's drift velocity can be approximated by taking a difference between thermal center estimates over consecutive batches of data, although this can lead to an instability. The instability is avoided by setting drift velocity to average wind velocity.

Evolutionary Search Method

The method by Allen provides a good first estimate of the thermal's center point. Once found, the optimal W and R in Eq. (3) can be found and the estimate of the thermal's center point can be improved in an evolutionary manner as described below. Sampling the natural log of Eq. (3) produces the set of linear algebraic equations:

$$\mathbf{Ax} = \mathbf{b} \quad \text{in which} \quad \mathbf{A} = \begin{bmatrix} 1 D_1^2 \\ 1 D_2^2 \\ \vdots \\ 1 D_N^2 \end{bmatrix} \quad \mathbf{x} = \begin{pmatrix} \ln(W) \\ -R^{-2} \end{pmatrix} \quad \mathbf{b} = \begin{pmatrix} \ln(v_1) \\ \ln(v_2) \\ \vdots \\ \ln(v_N) \end{pmatrix} \quad (5)$$

The least squares fit $\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$ determines W and R . The least square error $e = (\mathbf{Ax} - \mathbf{b})^T (\mathbf{Ax} - \mathbf{b})$ is a function of the center point so different center points can be tested to incrementally reduce the least square error. Improved estimates of the center point are found using an evolutionary search as follows: Candidate center points that lie on a circle of radius r around the initially estimated center point are selected. Next, the candidate center point that produces the smallest error is selected as the updated center point. Finally, this step is repeated in an iteration using circles of decreasing radii r until convergence is reached. This method is described in more detail in [11].

Simultaneous Iteration Method

The evolutionary search method sampled the natural log of Eq. (3) in order to turn the data into a set of linear algebraic equations in terms of unknown parameters, minimizing the modified parameters given in Eq. (5) rather than the actual updraft parameters used in Eq. (3). The Simultaneous Iteration Method developed below avoids the need of taking a natural

log and can estimate simultaneously a thermal's center point and its other parameters. Furthermore, the shape of the thermal can be relaxed from a thermal that is static in the wind-corrected coordinate system and that has a circular cross-section to one whose cross-section is elliptical. The updraft velocity distribution is expressed in the more general form

$$v_{ud} = e^{f(x,y)}, \quad f(x,y) = a_0 + a_1(x - a_4)^2 + a_2(y - a_5)^2 + a_3(x - a_4)(y - a_5) \quad (6)$$

where a_0, a_1, \dots, a_5 are unknown parameters. The parameter a_0 is associated with the magnitude of the vertical velocity of the thermal's center, a_1 and a_2 are associated with the radii of the elliptical cross-section in the x and y directions, respectively, a_3 is associated with the elliptical cross-section's axis of rotation, $a_4 = x_C$ and $a_5 = y_C$. Next, assume that the parameters in the previous batch of data were $a_0^0, a_1^0, \dots, a_5^0$, in which case the interest lies in finding the changes $\Delta a_0, \Delta a_1, \dots, \Delta a_5$ in the parameters. From Eq. (6), the change in the updraft velocity distribution depends on the changes in the parameters by

$$\Delta v_{ud} = \sum_{r=0}^5 \frac{\partial v_{ud}}{\partial a_r} \Delta a_r = \sum_{r=0}^5 \frac{\partial f}{\partial a_r} e^f \Delta a_r \quad (7a)$$

in which

$$\begin{aligned} \frac{\partial f}{\partial a_0} &= 1, & \frac{\partial f}{\partial a_1} &= (x - a_4^0)^2, & \frac{\partial f}{\partial a_2} &= (y - a_5^0)^2, & \frac{\partial f}{\partial a_3} &= (x - a_4^0)(y - a_5^0) \\ \frac{\partial f}{\partial a_4} &= -2a_1^0(x - a_4^0) - a_3^0(y - a_5^0), & \frac{\partial f}{\partial a_5} &= -2a_2^0(y - a_5^0) - a_3^0(x - a_4^0) \end{aligned} \quad (7b)$$

where $\Delta v_{ud} = v_{ud} - v_{ud}^0$ in which v_{ud}^0 denotes the estimate of v_{ud} calculated from Eq. (6) using the parameters taken from the previous batch of data. Equation (7) is evaluated at the times t_k ($k = 1, 2, \dots, N$) to yields the set of linear algebraic equations:

$$\mathbf{Ax} = \mathbf{b}, \quad \mathbf{A} = \mathbf{EA}_0, \quad \mathbf{x} = \begin{pmatrix} \Delta a_1 \\ \Delta a_2 \\ \vdots \\ \Delta a_5 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \Delta v_{u1} \\ \Delta v_{u2} \\ \vdots \\ \Delta v_{uN} \end{pmatrix} \quad (8a)$$

in which

$$\mathbf{E} = \text{diag}(e^{f_1} \quad e^{f_2} \quad \dots \quad e^{f_N}), \quad \mathbf{A}_0 = \begin{bmatrix} \frac{\partial f_1}{\partial a_0} & \frac{\partial f_1}{\partial a_1} & \dots & \frac{\partial f_1}{\partial a_5} \\ \frac{\partial f_2}{\partial a_0} & \frac{\partial f_2}{\partial a_1} & \dots & \frac{\partial f_2}{\partial a_5} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_N}{\partial a_0} & \frac{\partial f_N}{\partial a_1} & \dots & \frac{\partial f_N}{\partial a_5} \end{bmatrix}_0 \quad (8b)$$

where f_k denotes the function f evaluated at x_k and y_k . Equation (8) is an over-determined system of linear algebraic equations whose least square error needs to be minimized. However, for real-time stability, it is desirable to prevent the elliptical cross-section of the updraft velocity distribution from becoming too elliptical and to penalize large changes in the parameters during an iteration. Therefore, we can minimize a performance functional that considers both of these objectives. The general form of the performance functional is

$$J = (\mathbf{Ax} - \mathbf{b})^T \mathbf{W}_E (\mathbf{Ax} - \mathbf{b}) + (\mathbf{x}_P + \mathbf{x} - \mathbf{x}_D)^T \mathbf{W}_D (\mathbf{x}_P + \mathbf{x} - \mathbf{x}_D) \quad (9)$$

in which $\mathbf{x}_P = [a_0^0 \ a_1^0 \ \dots \ a_5^0]^T$ denotes the parameters from a previous batch of data and $\mathbf{x}_D = [a_{0D} \ a_{1D} \ \dots \ a_{5D}]^T$ denotes a desired set of parameters. The first term in Eq. (9) is a least square error and the second term is weighted measure of the difference between the unknown parameters $\mathbf{x}_P + \mathbf{x}$ and the desired parameters \mathbf{x}_D . \mathbf{W}_E is an $N \times N$ weighting matrix associated with the error and \mathbf{W}_D is a 6×6 weighting matrix associated with the desired set of parameters. Equation (9) is a function of the changes \mathbf{x} in the parameters. Minimizing it with respect to \mathbf{x} yields the closed-form solution

$$\mathbf{x} = (\mathbf{A}^T \mathbf{W}_E \mathbf{A} + \mathbf{W}_D)^{-1} (\mathbf{A}^T \mathbf{W}_E \mathbf{b} + \mathbf{W}_D (\mathbf{x}_D - \mathbf{x}_P)) \quad (10)$$

Equation (10) reduces to the weighted least squares solution $\mathbf{x} = (\mathbf{A}^T \mathbf{W}_E \mathbf{A})^{-1} \mathbf{A}^T \mathbf{W}_E \mathbf{b}$ when $\mathbf{W}_D = 0$ and it reduces to $\mathbf{x} = \mathbf{x}_D - \mathbf{x}_P$ when $\mathbf{W}_E = 0$. Desirable results are achieved by tuning the weights. Toward this end, it is best to first scale the parameters \mathbf{x} and \mathbf{b} so their magnitudes are of the same order. This produces weights that are scaled and can then be compared with each another more readily. The diagonal elements of the weighting matrices place emphasis on minimizing the corresponding parameters. The weights can also minimize

the thermal's eccentricity. (Notice that $W(a_1 - a_2)^2 = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}^T \begin{bmatrix} W & -W \\ -W & W \end{bmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}$ minimizes eccentricity.)

Guidance Commands

The control strategy for autonomous soaring divides into 4 steps: state measurement, updraft identification, guidance commands, and flight control. Following standard practices,

the state measurements are obtained in this first step from a GPS/INS system, such as that obtained from autopilot telemetry. Filtering of signals is normally performed in this step. The second step, which is unique to autonomous soaring, was described in the previous section. The third step, guidance commands, represents a series of decisions: when to soar, how to soar, how to search, and how fast to fly in and between thermals. This section discusses the considerations that enter into the decisions making up the guidance commands. The fourth step, flight control, responds to the guidance commands and is performed by an autopilot.

Deciding to orbit in a thermal or to leave it

The decision is made to enter a thermal soaring mode (to latch) when confidence is reached that the aircraft is in a sufficiently strong thermal. . This occurs when the average updraft velocity \bar{v}_{ud} exceeds a threshold. The average updraft velocity can be determined from updraft velocity data taken over a window on the order of 10 seconds. Figure 1 gives the threshold strength (updraft velocity) for latching at a given altitude. A typical threshold is

$$\bar{v}_{ud} > 0.6 \text{ m/s at } 600 \text{ m} \quad (11)$$

The decision is made to leave a thermal when any one of a number of escape conditions is met. The thermal can become increasingly weak; the average climb rate can go to zero; the aircraft altitude can reach a maximum height determined by visual sight limits; or the updraft velocity can be less than a speed ring threshold. The speed ring threshold is a go/no-go speed condition that maximizes cross-country speed over a course (MacCready speed ring). Typical threshold criteria to unlatch are

$$\dot{h}_{av} < 0, \quad h > h_{LOS}, \quad \bar{v}_{ud} < 1.4 \text{ m/s at } 600\text{m} \quad (12)$$

where h_{LOS} is an altitude limit set by the pilot's limit of visual sight. These algorithms do not require an estimate of the updraft velocity distribution. However, these conditions can be improved upon by accounting for a thermal's stability in both size and updraft variability, which would be feasible when the updraft velocity distribution presented in the previous section is being estimated with a high level of confidence. The threshold for staying latched is higher than the threshold for the initial latch so the aircraft will make investigative turns even into small thermals, but will leave if they turn out to be weak.

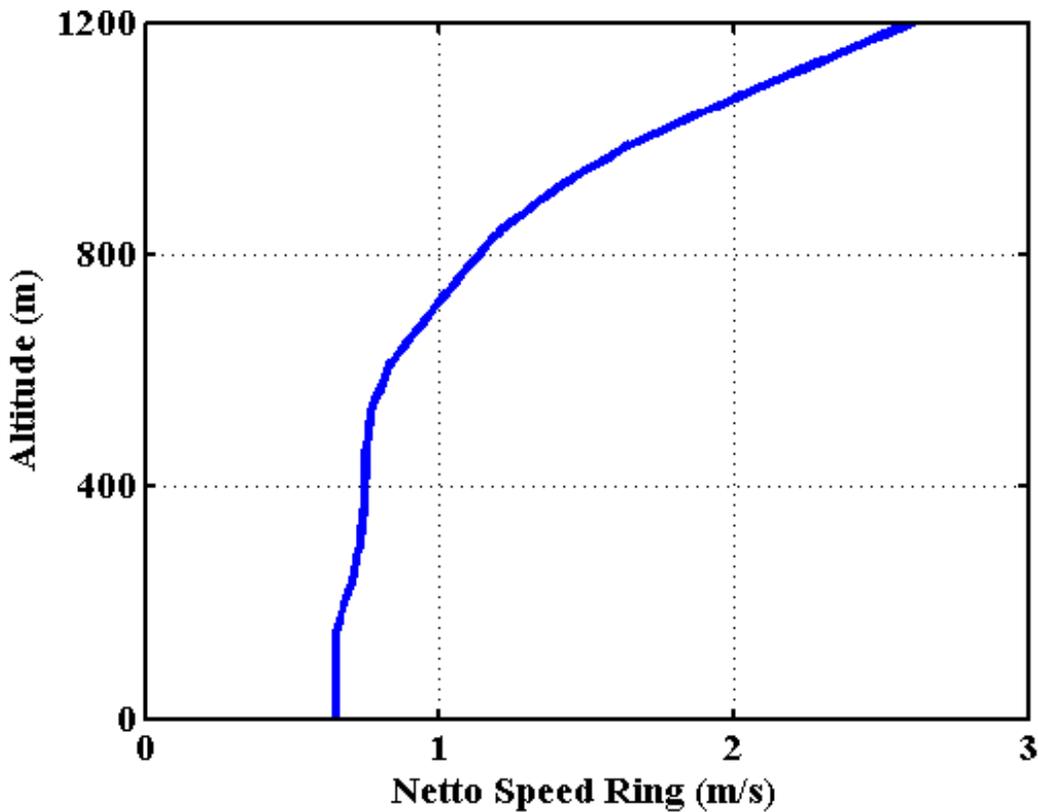


Figure 1 Speed ring – altitude versus updraft velocity (netto m/s).

Deciding how to orbit

Once the decision is made to orbit in a thermal, several subsequent decisions are made regarding how to accomplish this task, like determining orbit radius and orbit direction. It is sufficient to turn as sharply as is necessary to stay in a thermal, recognizing that wider turns are more efficient [10]. It is also known that thermals tend to get wider at higher altitudes [13]. Therefore a simple lookup table can be used to determine the radius in which the aircraft should turn. Greater optimality would exploit data provided by the updraft distribution.

The initial decision to orbit clockwise or counter-clockwise can be difficult because the observability of the center of the updraft is poor from a straight line of measurements [5]. This paper suggests adopting the approach of maintaining turn direction; without loss of generality, always turn left (counterclockwise). The subsequent orbit center corrections look similar to manned centering methods (See Fig. 2) which shift the orbit center rather than roll across wings level for an opposite turn direction. This approach avoids the figure eight correction (See Fig. 3) which incorporates a tight turn in its first loop that is outside of the thermal and therefore results in a significant loss of altitude. The never changing turn direction approach is effective even when turning away from the thermal initially (as in Fig. 2) because the additional data provided during the turn improves the estimate of the thermal's center and prevents the aircraft from "losing" the thermal. This approach differs from manual piloting practices because the manual pilot does not use a wind-corrected coordinate system.

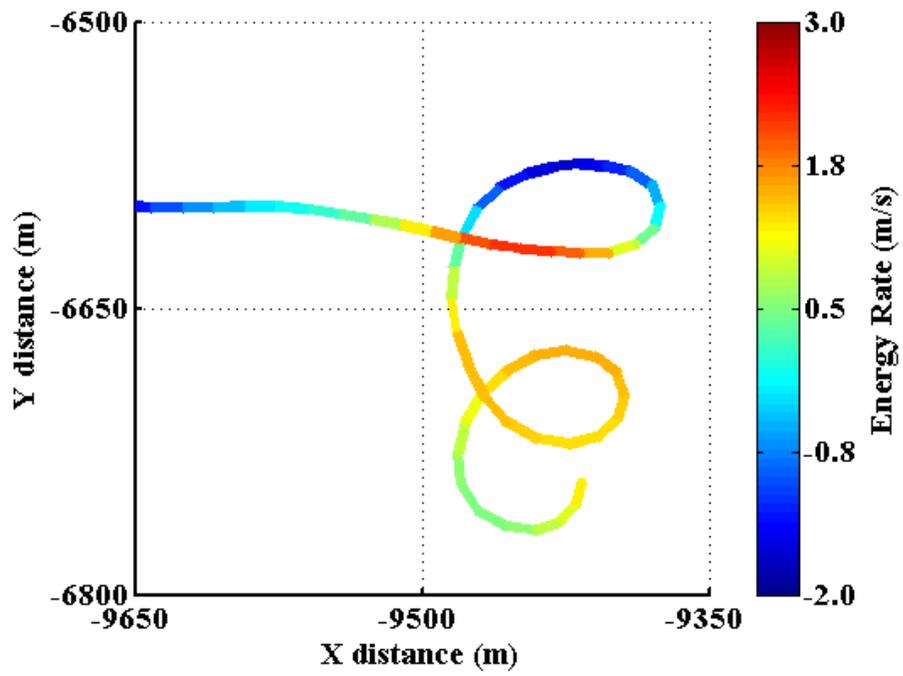


Figure 2 Maintain turn direction approach.

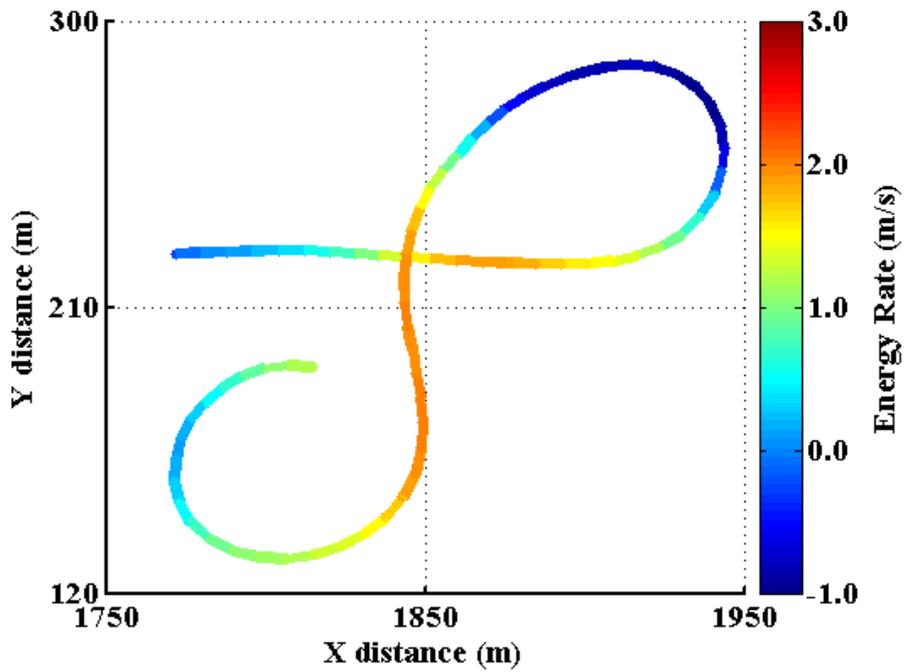


Figure 3 Switching turn direction.

Selecting airspeed

Energy consumption increases significantly with airspeed, making it important in autonomous soaring to fly at a proper speed. During inter-thermal cruise, pilots decrease airspeed in rising air and increase it in sinking air in a well known dolphin-soaring mode [9]. This decision can instead be made on the basis of gust soaring [8]. This paper adopts the speed ring approach, which optimizes commanded speed-to-fly based on anticipated conditions (See Eq. 13). A speed ring curve specified in terms of altitude balances the behavior of being conservative at low altitudes and aggressive at high altitudes [7] and was empirically tuned (See Fig. 1). The commanded airspeed is

$$v = v_h + \sqrt{\frac{1}{a}(c + v_{ud} - v_{sr})} \quad (13)$$

in which the headwind v_h is determined from the wind velocity estimate and aircraft heading, a and c are constants determined from the sink polar [12], v_{ud} is measured by the netto variometer, and v_{sr} is the anticipated next updraft strength [6]. In this research, 1.3 times the minimum sink speed is used during orbits to avoid approaching stall.

Deciding how to search for thermals

The decision is made to search for a new thermal once leaving the last thermal. The method of search can be influenced by terrain features [14], expert knowledge of the area, and real-time atmospheric measurements. In absence of this information, the shortest path from A to B is a straight line. The straight line path is decidedly not an optimal path. However, when the density of thermals along a path is sufficiently high, the straight-line

approach is effective. Other works assuming well-known updraft fields and optimum transit times [15] are not yet applicable without a sensor that can measure the vertical wind fields ahead of the aircraft; such a sensor is suggested by Charrett [16].

Flight Testing

ALOFT is an SBXC airframe produced by RnR Composites, a highly popular RC glider used in cross-country soaring. ALOFT has a flight weight of 5.56 kg, a wingspan of 4.3 m, a best L/D of 24:1 at 10.9 m/s, and a minimum sink rate of 0.43 m/s at 10.0 m/s [12]. Figure 4 shows ALOFT's hardware for autonomous soaring, specifically, a Cloud Cap Technologies Piccolo II autopilot, a JR 2.4 GHz RC receiver, an RxMux safety switch, a Skymelody variometer, and an 87 W-hr lithium-ion battery. The onboard telemetry is down-linked to a ground control station (GCS) laptop computer where the soaring calculations are performed and the resulting guidance commands are then up-linked to the autopilot. Note that the Skymelody variometer is used for operator situational awareness only and not used by the soaring algorithms.

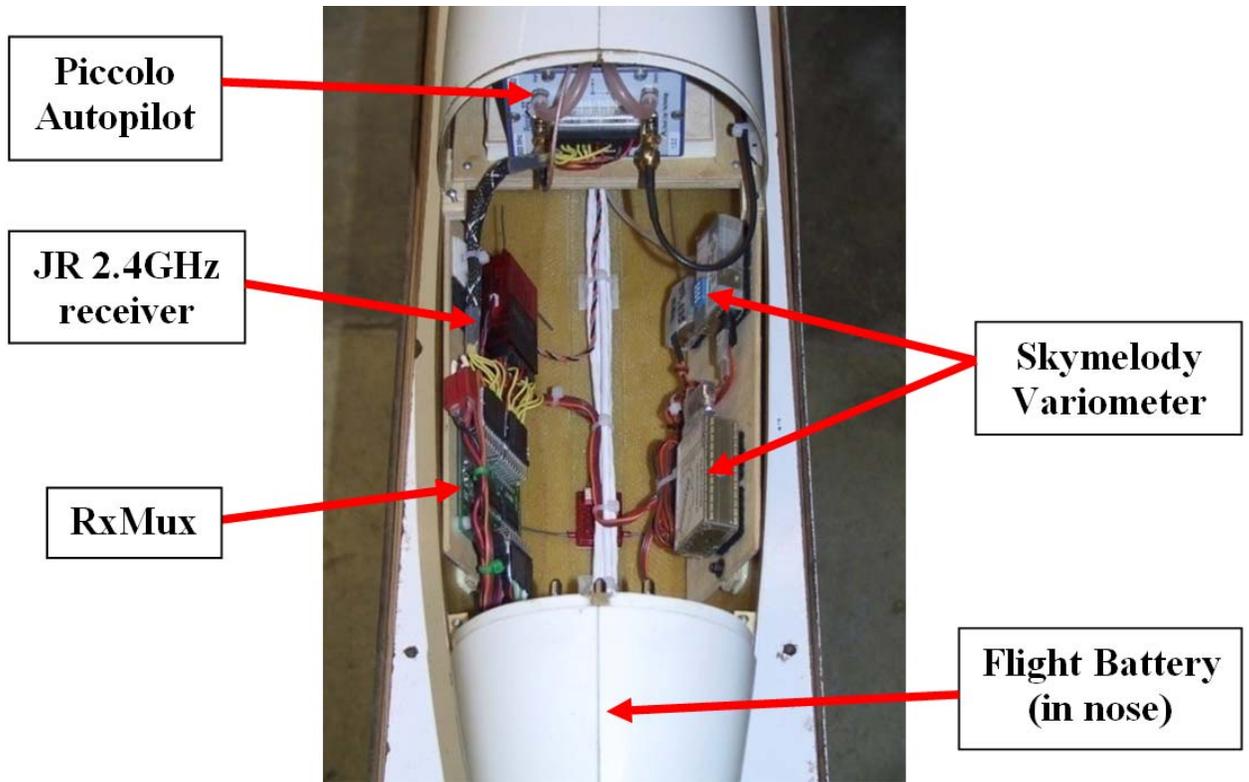


Figure 4 ALOFT hardware.

The autonomous soaring calculations represent a high-level hierarchy that is kept independent of the flight controls. The development of an independent hierarchy for autonomous soaring decoupled the testing, troubleshooting, and validation of the autonomous soaring algorithms and allowed a COTS (commercial off the shelf) autopilot to be leveraged as inner-loop control. This significantly simplified algorithm integration and is suggested as preferred over the coupled approach in [8]. Also, note that the off-board execution of the autonomous soaring calculations was not actively constrained by data link throughput because the guidance commands in autonomous soaring are performed at a low frequency on

the order of 2 Hz. In the RC soaring community, it is well known that the guidance should be performed with “a light touch...let the plane read the air and tell you what it needs” [1].

The sink polar of the ALOFT was determined in several constant airspeed flight tests and fit to a quadratic function of relative speed [6]. The sink polar for ALOFT was determined to be

$$\dot{h} = -0.0232v^2 + 0.4634v - 2.759 \quad (14)$$

(See Fig. 5). As shown, a minimum sink rate of 0.48 m/s was obtained when the aircraft’s horizontal velocity was 11.2 m/s [12]. Similar sink polars can be determined including flap deflections to improve the perceived aircraft performance, but were not used in this implementation.

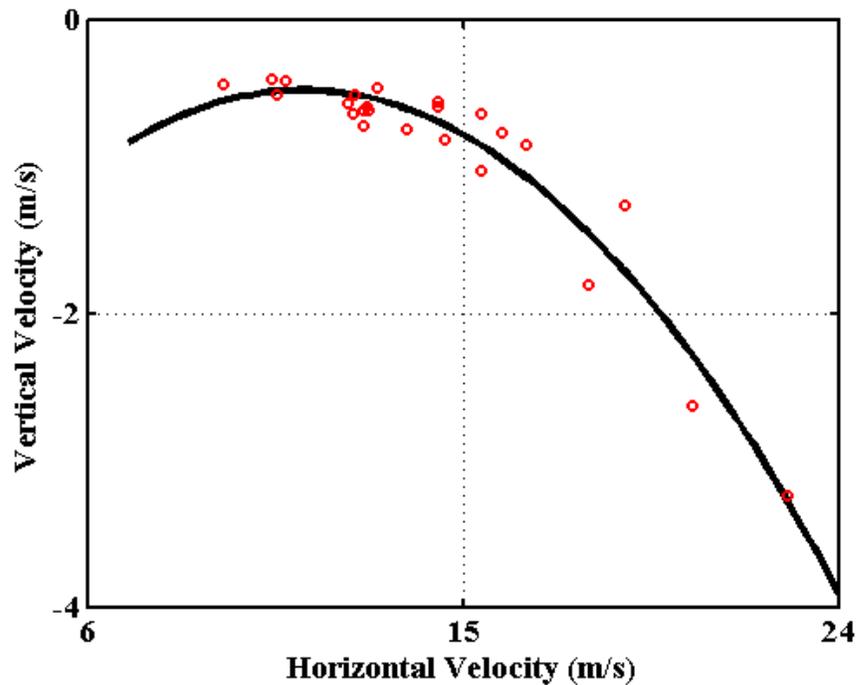


Figure 5 Sink polar of the SBXC at 5.56kg.

The period of the batch data used in Eq. (5) was set to 45 seconds at 2 Hz as a compromise. On a windy day (> 10 m/s) 45 second old data is obsolete before it can be used whereas on a calm day the amount of data collected in the batch is sufficient. The batch data consists of measurement times t_k , aircraft locations x_k and y_k , altitude h_k , airspeed v_k , and West and South wind speed components v_W and v_S . The batch data was updated in a first-in first-out manner.

The computed netto variometer signal (Eq. 2b) was noisy due to differentiation of the autopilot's digitized velocity data and therefore needed to be filtered. To avoid the phase shift that is normally found in real-time filtering, a unique no-lag, real-time filter was developed. The lag was eliminated by realizing that the use of the first and last 5 seconds of data is less important than having a smooth batch. See [11] for more discussion about the filtering method employed.

The autopilot's wind estimate was filtered by using the average last 5 minutes of wind data; when inside the thermal the wind estimation approximated the thermal's drift velocity and when between thermals approximated the ambient wind velocity.

The different methods of identifying the thermal's center point are compared in Fig. 6 and 7 using flight data. The initial center point was determined by the Centroid Method. Notice that the data surrounds the initial center point; this is a requirement of the Centroid Method. When the aircraft is not orbiting around the thermal's center, the error in the Centroid Method will be large, albeit tending in the correct direction.

Figure 6 shows the updating of the center point estimated by the Evolutionary Search Method. The updated center points move down the error surface from the Centroid Method's center point, reaching the outside of the flight path. In the Montague flights, three fixed steps of size 25m, 15m, and 7m were used to prevent extreme center estimates when flying in irregular thermals.

Figure 7 shows the progression of the Simultaneous Iteration Method overlaid on the same error surface as in Fig. 6. The initial move is toward the global minimum, but subsequent steps show pulling toward the desired center (set to the centroid location), balancing the desire to minimize error and fit to some desired parameters. In this example, $\mathbf{W}_E = 1000 \mathbf{I}$ and $\mathbf{W}_D = \mathbf{I}$, where \mathbf{I} is the identity matrix, and the desired parameters are set to the values obtained by the Centroid Method.

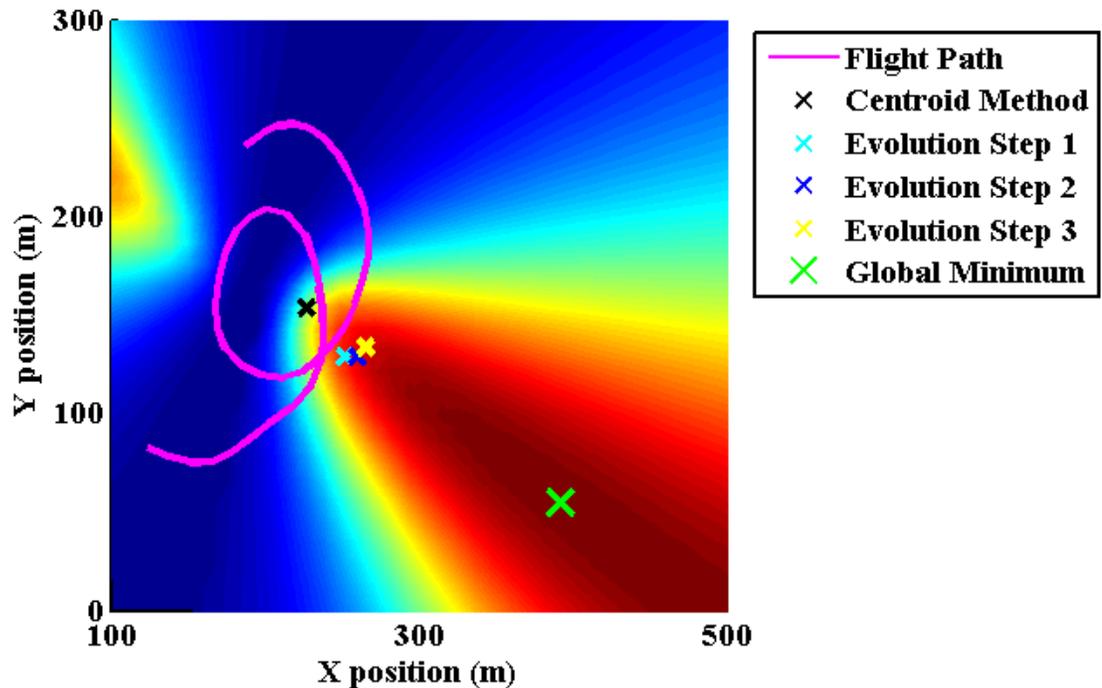


Figure 6 Evolutionary Search center point updates on an error surface.

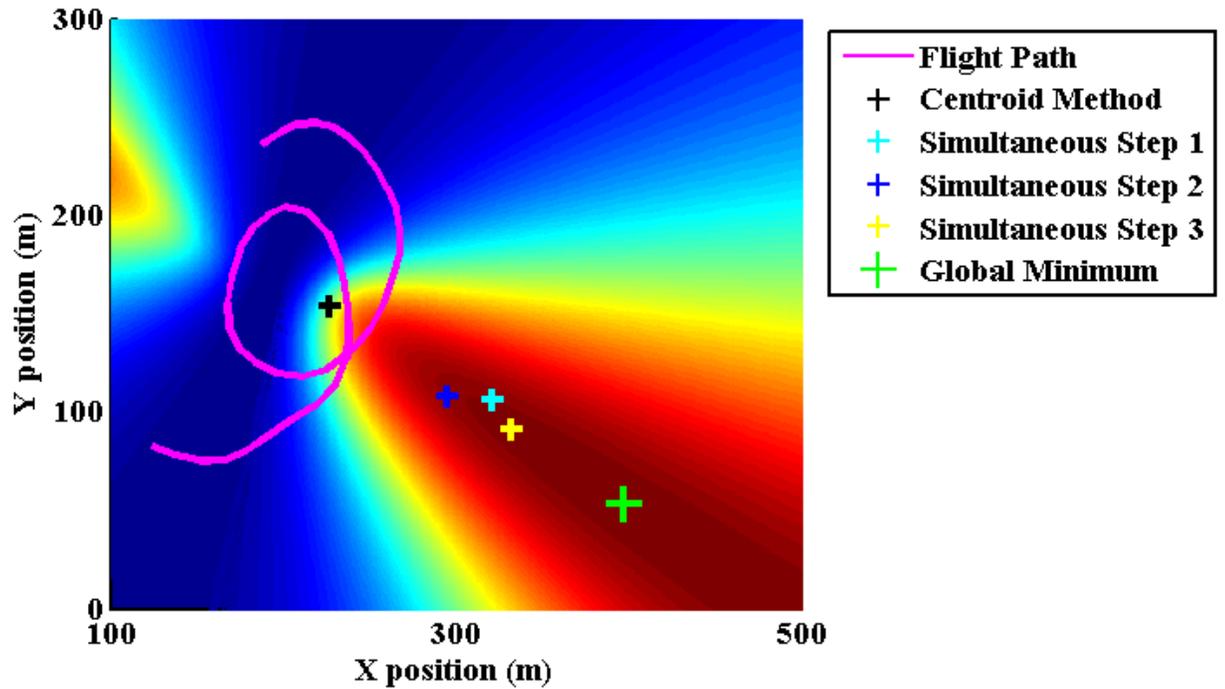


Figure 7 Simultaneous Iteration center point updates on an error surface.

The Evolutionary Search Method was used in the Montague Cross-Country Challenge (discussed in more detail in the next section). A high level of variability is present in the updraft velocity measurements, especially when the updrafts are ill-formed (typical of windy days). The Evolutionary Search Method is particularly attractive under these conditions because it starts from the stable Centroid Method center point and is only allowed to move slightly from it. The Simultaneous Iteration Method using $\mathbf{W}_E = \mathbf{I}$ and $\mathbf{W}_D = \mathbf{0}$ showed marked improvements in centering ability when the updrafts are well-formed, but these results are sensitive to initial conditions and updraft velocity measurement noise. The Evolutionary Search Method was implemented in the Montague competition instead of the Simultaneous Iteration Method because of greater familiarity with the former at that time.

Since its construction in early 2003 and as of July 2009, ALOFT logged 164 flights for a total air time of 70.7 hrs, of which 45.5 hrs were flown by the autopilot and 19.8 hours were in the auto-soar mode (actively soaring in thermals). The maximum continuous autopilot flight was for 5.28 hours and the maximum altitude reached was 5500 feet AGL (limited by h_{LOS}).

The Montague Cross Country Challenge

The Montague Cross Country Challenge arguably brings together the best RC cross-country (XC) sailplane teams in the US for three days of competitive soaring. Most of the teams used the SBXC airframe, like ALOFT. The majority of the teams, including the ALOFT team, down-linked an audible variometer signal for sensing the effects of updrafts. The noteworthy differences between the RC teams and the ALOFT team were: a) the high experience level of the RC teams compared to the modest experience level of the ALOFT team, since this was the ALOFT team's first entry in this competition and the first known entry of an autonomous soaring aircraft in any soaring competition, and b) the autonomous soaring algorithms available in ALOFT, not available to the RC teams. The competing teams and their aircraft are shown in Fig. 8.



Figure 8 2008 Montague teams and their aircraft (photo courtesy Ron McElliot).

The teams flew a different course each of the three days, receiving a briefing of the day's weather and task at 9AM and being released to an open winch at 11AM. Time cards were turned in each day by 5PM. Day 1 was a speed task over a 2 hour minimum period and along a specified 28.08 mile course, and an open course thereafter (See Fig. 8). Day 2 was an open distance task along a 72 mile specified course. Day 3 was a speed task, but only along a specified 15.92 mile course.

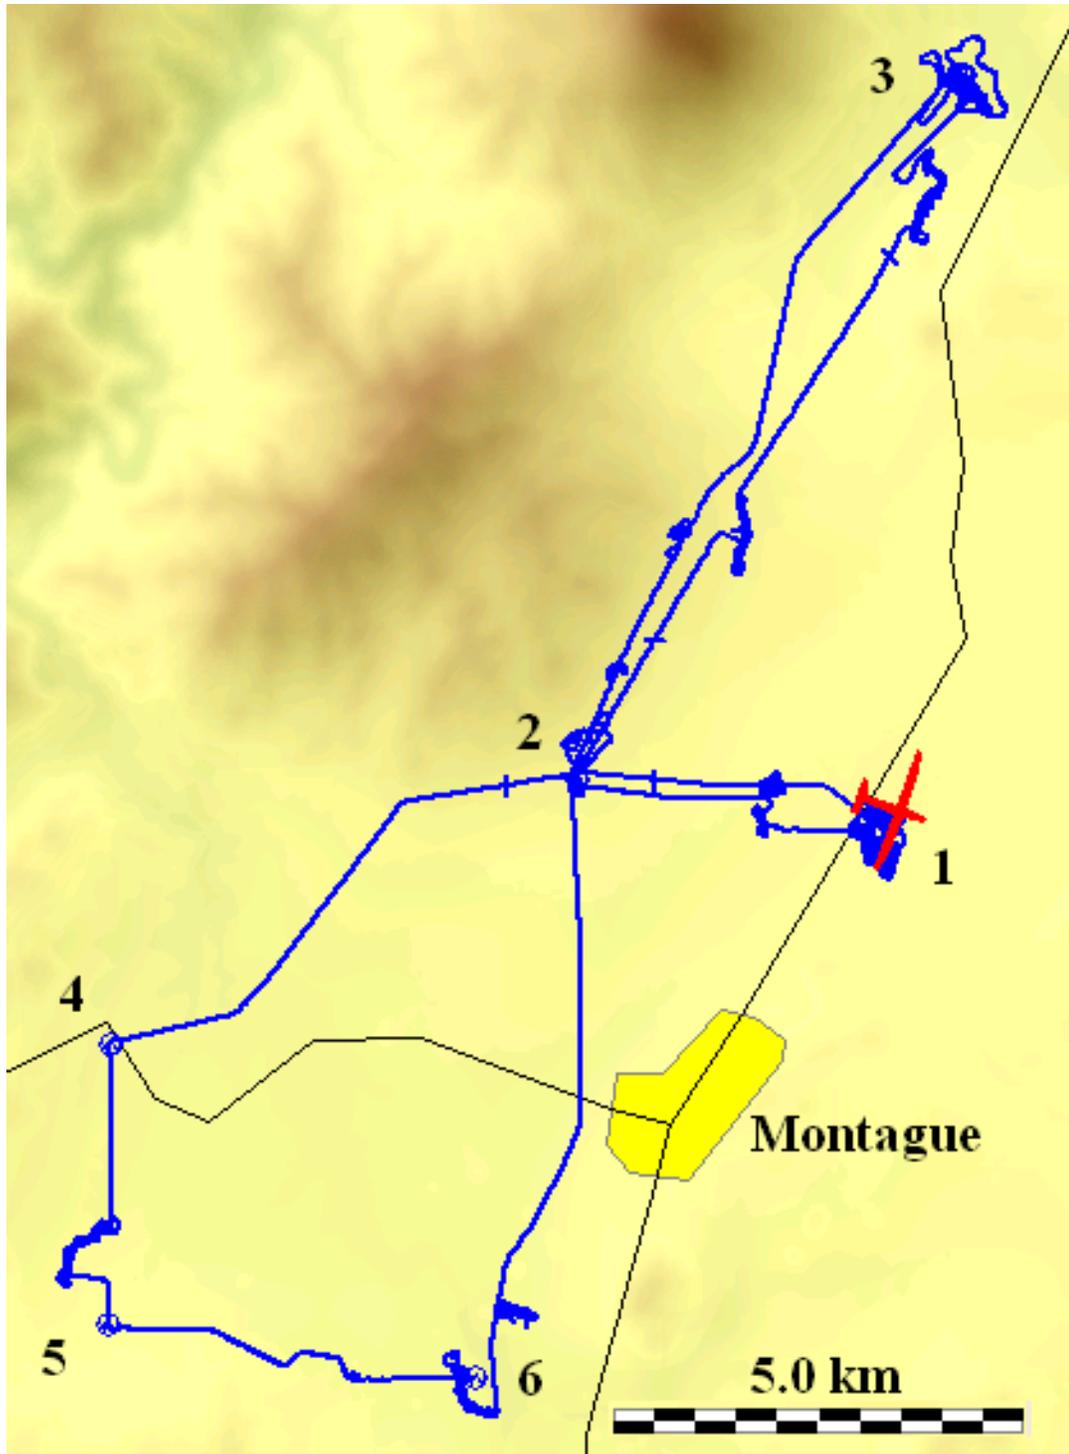


Figure 9 Course on Day 1, turn point sequence 1-2-3-2-4-5-6-2-1.

Figure 9 also shows the flight path taken between waypoints, including areas where the vehicle stopped to orbit in thermals. The ALOFT algorithms were optimized more for speed tasks, but were adjusted on-the-fly to be more conservative to optimize distance by modifying the speed ring altitude curve. By doing so, the aircraft stayed in the weaker updrafts over higher altitudes. The ALOFT team improved its speed ring lookup table and thermal latching limits at the end of each day, resulting in corresponding performance improvements.

The altitude profile for Day 2 is shown in Figure 10, demonstrating how altitude was traded along the course and regained using updrafts. Active soaring is shown in red, representing about 50% of the flight time for the weak day.

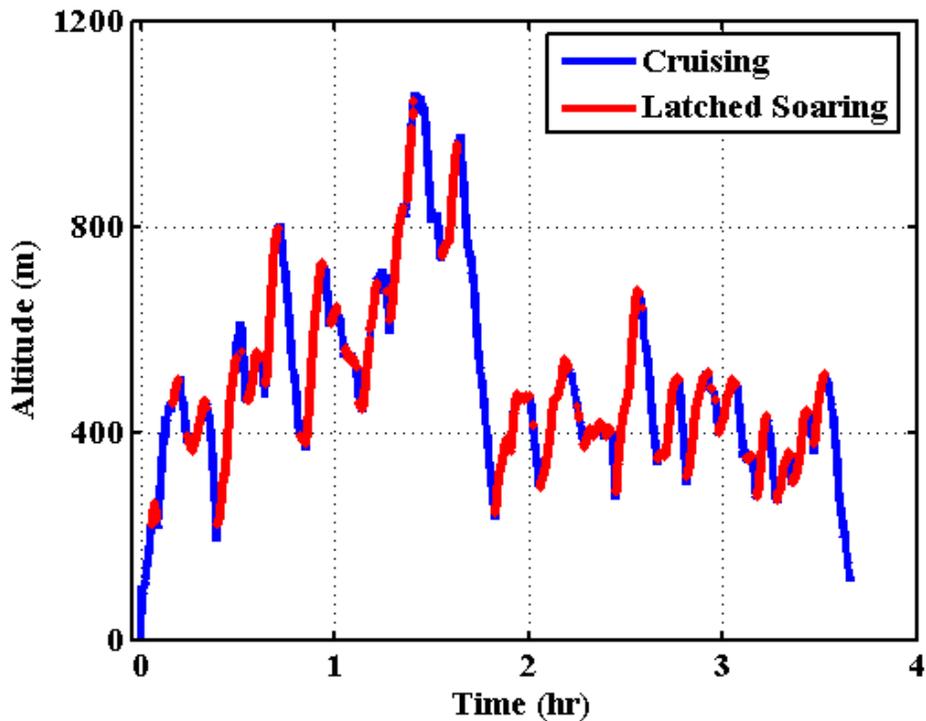


Figure 10 ALOFT altitude versus time, Day 2.

An example of the soaring behavior is shown in Figure 11. ALOFT is flying in a straight line using dolphin soaring to optimize its speed given the location's vertical wind speed. Then it enters an updraft and takes an investigative turn into the updraft. Entering the latched mode, it continues to update the updraft center and follow the shifting updraft core, which carries the aircraft up quickly initially and tapers off with altitude. At this point, the speed-ring altitude curve unlatches the soaring mode and ALOFT continues along its path. This climb netted 365 m in 3.2 minutes.

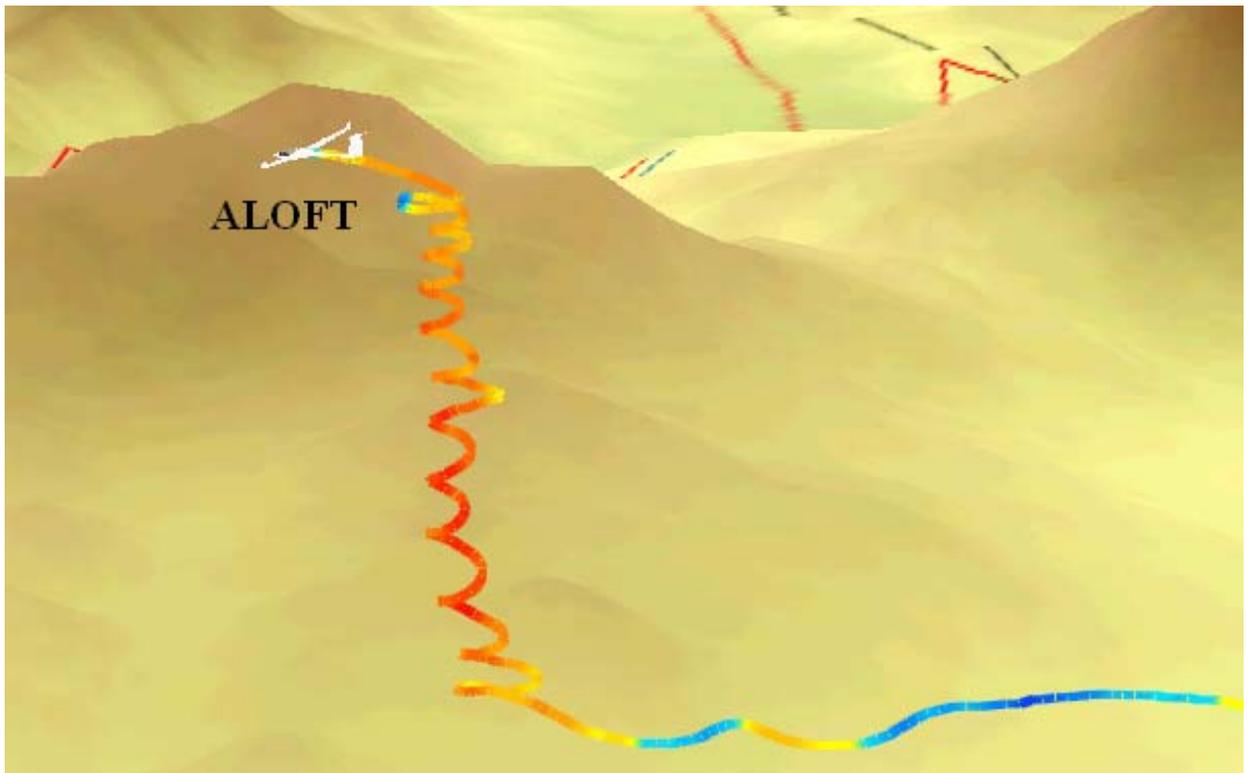


Figure 11 ALOFT entering, orbiting, and leaving an updraft.

The competition results are given in Table 1. Typical of RC soaring competitions, man-on-man scoring normalized the scores each day based on the top performer, who received

1000 points for the day. Over the 3 day period, ALOFT accumulated 2283 points compared to the overall leader who accumulated 2901 points. ALOFT placed third in the competition in overall points.

Table 1 Competition results

Day	Winner	ALOFT	Placement	Points
1	17.26 mph	12.50 mph	6th	724
2	53.76 miles	39.44 miles	5th	733
3	59.45 min	72 min	4th	825
Overall	2901 points		3rd	2283

The winning pilot of the competition, John Elias, and a 10 year XC veteran with 150 hours of competition experience commented on autonomous soaring: “It is my impression that ALOFT flies with greater precision, controls speed more accurately, maintains a more direct route to turn-points, and thermals as well as or better than the typical manually flown glider. However, when conditions are confusing and weak, I think an experienced pilot can make use of many terrain factors and visual indicators that ALOFT simply cannot.”

Final Remarks and Conclusions

This paper developed a method for locating and actively guiding gliders to stay in thermals. Starting with a first estimate by Allen, an Evolutionary Search Method and a Simultaneous Iteration Method were developed; both were verified using flight data to be able to locate thermals and identify the parameters that describe a thermal’s velocity

distribution. The Evolutionary Search Method was then implemented in the ALOFT vehicle for competing at Montague.

The ALOFT team, having no experience with the Montague course, did not take advantage of expert knowledge of the terrain. Weather limited the updrafts to weaker and lower altitudes than typical test conditions, forcing ALOFT to stay within the same altitude limits as an RC aircraft. Despite these drawbacks, ALOFT out-climbed other teams in the same thermals because ALOFT had better knowledge of the center of the updrafts and could fly more precisely in them. Also, ALOFT better optimized speed-to-fly for longer glides and flew directly between turn points, resulting in more efficient inter-thermal cruise. Placing third in overall points against similarly classed aircraft flown by experienced RC pilots establishes autonomous soaring as a bona fide engineering sub-discipline, which is expected to be of interest to engineers who might find this has some utility in the aviation industry.

The short-term question arises how autonomous soaring can further be improved. The following lists some of the areas where progress would be helpful:

- 1) Path planning: The search algorithms in autonomous soaring constrain an aircraft's admissible paths. The adaptation of UAV path planning algorithms using autonomous soaring is an open problem.
- 2) In-flight knowledge of updraft velocity distributions: The search algorithms will be greatly improved with in-flight remote sensing of updraft velocity distributions in the neighborhood of the aircraft.
- 3) Powered flight: The development of guidance algorithms in the presence of powered flight requires additional study, namely in accurate inclusion of the new energy term.

Postscript: Since the Montague Cross-Country Challenge in 2008, ALOFT surpassed the 2005 FAI F3 goal-and-return world distance record of 39.1 km. On October 5, 2008, it unofficially set a new record on its first attempt, travelling 48.6 km (97.2 km round-trip) in a flight time of 3.5 hrs. During the flight, ALOFT was under full autopilot control for approximately 99% of the time.

REFERENCES

- [1] Thornburg, D., *Old Buzzard's Soaring Book*, Pony Express, New Mexico 1993.
- [2] Allen, M., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *AIAA Guidance, Navigation, and Control Conference*, AIAA 2005-1025, AIAA, Reno, Nevada, 2005.
- [3] Allen, M. "Guidance and Control of an Autonomous Soaring UAV," NASA TM-214611, 2007.
- [4] Wharington, J., "Autonomous Control of Soaring Aircraft by Reinforcement Learning," Ph.D. Dissertation, Aerospace Engineering Dept., Royal Melbourne Institute of Technology, Melbourne, Australia, 1998.
- [5] Hazard, M., "Unscented Kalman Filter for Thermal Parameter Identification," *AIAA Region II Student Conference*, AIAA, Washington, DC, 2009.
- [6] Reichmann, H., *Cross Country Soaring*, 7th ed., Soaring Society of America, Inc, New Mexico, 1993.
- [7] Cochrane, J., "MacCready Theory with Uncertain Lift and Limited Altitude," *Technical Soaring*, Vol. 23, No. 3, 1999, pp. 88.
- [8] Langelaan, J., "Gust Energy Extraction for Mini- and Micro- Uninhabited Aerial Vehicles," *AIAA Aerospace Sciences Meeting and Exhibit*, AIAA 2008-0223, AIAA, Reno, Nevada, 2008.
- [9] Metzger, D. and Hendrick, J., "Optimal Flight Paths for Soaring Flight," *Journal of Aircraft*, Vol. 12, No. ii, 1975, pp. 867, 871.
- [10] Thomas, F., *Fundamentals of Sailplane Design*, College Park Press, Maryland, 2005.
- [11] Edwards, D., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA Guidance, Navigation, and Control Conference*, AIAA 2008-7244, AIAA, Honolulu, Hawaii, 2008.
- [12] Edwards, D., "Performance Testing of RNR's SBXC Using a Piccolo Autopilot," *Technical Papers Database* [online database], URL: http://www.xcsoaring.com/techPicts/Edwards_performance_test.pdf [cited 5 January 2010].
- [13] Lenschow, D., "The role of thermals in the convective boundary layer," *Boundary-Layer Meteorology*, Vol. 19, No. 4, 1980, pp. 509 - 532.
- [14] Andersson, K. "Extending Endurance for Small UAVs by Predicting and Searching for Thermal Updrafts," Mechanical and Astronautical Engineering Dept., Naval Postgraduate School, Monterey, California, 2009 (unpublished).
- [15] Kahveci, N., Ioannou, P., and Mirmirani, D., "Optimal Static Soaring of UAVs Using Vehicle Routing with Time Windows," *AIAA Aerospace Sciences Meeting and Exhibit*, AIAA-2007-158, AIAA, Washington, DC, 2007.
- [16] Cherrett, G., "Derivation of requirements for a LIDAR based thermal detection sensor for autonomous soaring flight," M.S. Thesis, School of Engineering., Cranfield University., Bedfordshire, United Kingdom, 2008.