Spline Mapping to Maximize Energy Exploitation of Non-Uniform Thermals

John J. Bird
jjbird@gmail.com

Jack W. Langelaan
jlangelaan@psu.edu

Department of Aerospace Engineering
The Pennsylvania State University
State College, Pennsylvania, USA

Abstract

A method is described for modeling and maximizing the use of thermals by small unmanned aerial vehicles. A spline model is used to map thermals of arbitrary structure with no *a priori* knowledge of their shape. A candidate thermal exploitation method is developed to showcase the capability of this flexible mapping technique. Simulation results show the utility of the proposed approach both for simple Gaussian thermals and non-uniform thermals and compare climb rate for map-based thermalling and more traditional spiral climbing techniques.

Nomenclature

$\mathbf{c}$	basis spline coefficients
g	acceleration due to gravity
$\mathbf{h}$	observation model
k	spline order
n	number of knots supporting basis splines
r	turn radius
w	updraft velocity measurement
$\dot{z}$	sink rate
$C(w)$	contour at level w
K_t	Kalman gain
$N_{\cdot,k+1}$	basis spline of order k
$\mathcal{P}_k$	space of piecewise polynomial functions of order k
$\mathbf{P}$	state covariance
Q	process noise
R	measurement covariance
$\mathcal{S}$	spline function
V_a	airspeed
$\mathbf{X}$	system state
ϕ	bank angle
λ	knot location

Introduction

Thermal soaring has been practiced by pilots of manned sailplanes since the invention of the variometer in the 1920s. Recently, the proliferation of small UAVs has sparked an interest in automated soaring methods. Work by Allen established that substantial gains could be made by exploiting thermals [1]; flight tests by Allen, and later Andersson demonstrated that autonomous aircraft could extend endurance by harvesting energy from thermals [2, 3]. Edwards demonstrated the use of thermalling in cross-country flight by an autonomous aircraft, placing third in competition with piloted RC aircraft [4].

Thermalling controllers to date have generally used variations of Reichmann's method to fly a constant radius circle around the center of a thermal, drawing from techniques used by piloted sailplanes to maximize climb rate [2, 5–7]. Thermal modeling in autonomous aircraft is rudimentary, typically estimating the strength and size of a radially-symmetric thermal with the objective of establishing a nominal turn rate for circling [2, 6]. Edwards does propose a method to estimate the size and orientation of an elliptical thermal [4], but no treatment has been given to an arbitrarily shaped thermal and no controllers developed to exploit such knowledge. Manned sailplane instrumentation is similarly rudimentary. While recent instrumentation presents climb rate along the flight path and sectors of maximum lift while thermalling, the pilot must mentally construct a model of the lift environment [8], as no automatic modeling capability is provided.

This paper presents a method using splines to map the air motion in a thermal without assuming a thermal structure, allowing a more fully descriptive model of a thermal to be constructed. A Kalman filter method is also presented which allows the map to be efficiently constructed by a sailplane in climb and allows the model to remain current with changes in the thermal. Further, a thermal exploitation method is presented which uses contours of constant lift from the thermal map as flight paths for the sailplane. The effectiveness of the contour path thermalling method is established by comparing climb rates with Allen's and Andersson's circling methods.

Thermal Modeling

Tensor Product Splines

Splines can be used to efficiently model complex functions of unknown shape, allowing complex non-linear functions to be described as a piecewise polynomial. Partitioning a function by a number of "knots," a different polynomial is defined on each segment with continuity of the k^{th} derivative at the knots, where k is the order of the spline. If the spline is written as a linear combination of functions, known as basis splines (or B-splines), then it represents a linear mapping and can be used in linear estimation algorithms. In this form the spline is written [9]:

$$\mathcal{S}(x) = \sum_{i=-k}^{n} c_i N_{i,k+1}(x) \in \mathcal{P}_k \tag{1}$$

Presented at the XXXI OSTIV Congress, 8–15 August 2012, Uvalde TX USA

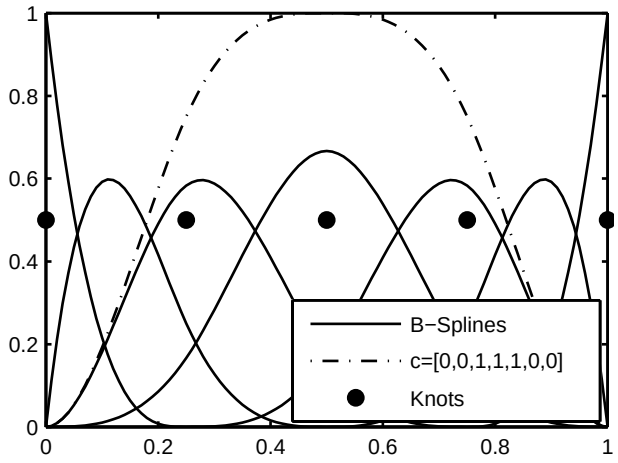

Fig. 1: Basis splines and one example spline with knots [0, 0.25, 0.5, 0.75, 1], order 3.

where N is the value at point x of a set of basis splines of order k defined on the set of knots $\lambda_j, j = 0\ldots, n+k-1$. Calculation of N is accomplished using a triangular scheme, described by Diercx [9]. The modeling of complex functions by linear combinations of basis splines is demonstrated in Fig. 1.

The concept of a spline can be generalized to more dimensions through use of the tensor product spline. In the tensor product spline, knot intervals are defined along each coordinate direction and the domain is then divided into cells defined by the Cartesian product of the knot intervals. In two dimensions this forms a rectangular mesh, with x and y knot intervals. The spline may be represented on each rectangle by the product of two polynomials, one along each coordinate direction [9]:

$$\mathcal{S} \in \mathcal{P}_k \otimes \mathcal{P}_l$$

If $\mathcal{P}_k(x)$ and $\mathcal{P}_l(y)$ are written as basis splines as in Eq. 1, then the resulting bivariate spline is the tensor product of the two spline functions, which can be written [9]:

$$\mathcal{S}(x,y) = \sum_{i=-k}^{n} \sum_{j=-l}^{m} c_{i,j} N_{i,k+1}(x) M_{j,l+1}(y)$$

With knots λ_i in x and μ_j in y fixing N and M, the values $c_{i,j}$ will define the shape of the spline function. The bilinearity of the tensor product [10] ensures that the final spline function is linear, thus $c_{i,j}$ can be estimated using any linear estimation procedure.

Thermal Modeling with Splines

The tensor product definition of a bivariate spline makes the thermal mapping process conceptually very simple: knots are defined on an interval bounding the region containing the thermal, and the order of the model is specified. The coefficients defining the spline's shape can then be estimated from measured updraft velocity to approximate the observed shape of the thermal. In order to keep the updates simple while allowing the thermal to change with time, a Kalman filter is used to estimate the shape of the thermal. The states of the Kalman filter are taken to be the spline coefficients, $c_{i,j}$ and the observation model is the tensor product of the two spline bases,

$$\mathbf{h}(x) = N_{i,k+1}(x) \otimes M_{j,l+1}(y) \tag{2}$$

The process noise is chosen to represent the expected change in the thermal parameters with time. With no state transition, the prediction step simply represents the increase in the uncertainty of the thermal map with time:

$$\begin{aligned} \hat{\mathbf{c}}_{t|t-1} &= \hat{\mathbf{c}}_{t-1|t-1} \\ \hat{\mathbf{P}}_{t|t-1} &= \hat{\mathbf{P}}_{t|t-1} + Q_t \end{aligned}$$

With the observation model defined as in Eq. 2 and the measurement noise chosen to represent the error in the measurement of vertical air motion, the Kalman filter update step proceeds:

$$\begin{aligned} K_t &= \hat{\mathbf{P}}_{t|t-1}\mathbf{h}^T(\mathbf{h}\hat{\mathbf{P}}_{t|t-1}\mathbf{h}^T + R_t)^{-1} \\ \hat{\mathbf{X}}_t &= \hat{\mathbf{c}}_{t|t-1} + K_t(w - \mathbf{h}\hat{\mathbf{c}}_{t|t-1}) \\ \hat{\mathbf{P}}_{t|t} &= (I - K_t\mathbf{h})\hat{\mathbf{P}}_{t|t-1} \end{aligned}$$

This filter allows the map to be rapidly updated, requiring storage of only the coefficient array c and its covariance matrix. The memory requirements are fairly modest, $\mathbf{c}$ has only $n = (k+g+l+h-2)$ elements, where k and l are the orders of the splines in the x and y directions, g and h are the number of knots in the x and y directions respectively. The update step does not even require matrix inversion, as $(\mathbf{h}\hat{\mathbf{P}}_{t|t-1}\hat{\mathbf{h}}^T + R_t)$ reduces to a scalar value.

The use of a Kalman filter means that the vertical wind velocity component, w can be incorporated as a raw measurement from a variometer with no prior filtering (i.e. a vario time constant of 0). The ability to directly incorporate this noisy measurement allows the elimination of filters which cause significant lag in most variometers, with the disadvantage that a good estimation of the vertical wind component requires a number of samples near a point to converge. In this paper a direct measurement of the wind field using the method described by Langelaan [11] is assumed, however measurements from a netto variometer could be used instead.

Path Planning

With a more complete understanding of the lift environment surrounding the sailplane comes the need for a method to leverage this information in harvesting energy from the thermal. This section proposes a path planning scheme which uses contours of the thermal map as candidate paths for the sailplane. A technique is also presented to balance exploration of the lift environment with exploitation of known areas of strong lift.

Contour Selection

As the aircraft constructs and updates the thermal map, a level set can be taken which describes a closed path around the estimated thermal structure that has a constant vertical wind speed. In order to optimize the aircraft climb rate, a cost function is defined to be the mean climb rate achieved during one orbit of a contour at level w (parameterized in polar coordinates by θ):

$$J(w) = \oint_{C(w)} (w - \dot{z}(\theta)) d\theta$$

Making the assumption that the aircraft is in steady-state turning flight as it traverses the path, the sink rate $\dot{z}(\theta)$ can be related to the

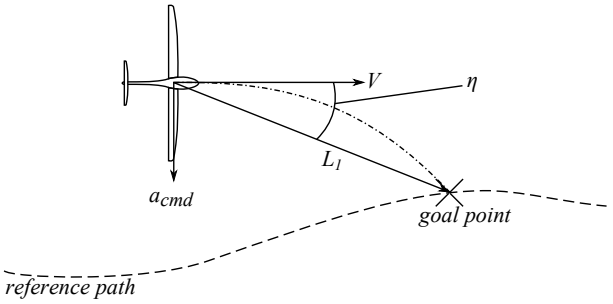

Fig. 2: Illustration of Park's nonlinear guidance law.

flight path curvature through the sailplane polar at a given bank angle, allowing the cost function to be evaluated relatively easily.

$$\begin{aligned} \dot{z} &= \dot{z}(\phi, V_a) \\ \phi &= \tan^{-1}\left(\frac{{V_a}^2}{r(\theta)g}\right) \end{aligned}$$

The term $r(\theta)$ is computed by locally fitting a circular arc at each point on the trajectory. With a cost function defined, the optimal path can be selected through the use of an optimization function to minimize the cost (coordinates here are defined positive down so a negative climb rate indicates an altitude gain).

Path Control

With a path defined, a controller is needed to keep the aircraft following the desired contour. The controller used here is a high level controller, developed under the assumption that lower level control (roll angle, airspeed control, etc) is already provided for on the UAV platform. To prevent unrealistic aircraft motions, the roll rate was restricted to $225^\circ/\text{s}$ (high, but not completely unreasonable for a small UAV platform), and the roll rate tended to remain below $45^\circ/\text{s}$. The controller implemented in this investigation is developed from the guidance method presented by Park [12], which generates a lateral acceleration command from the bearing to a reference point located on the desired path at a fixed distance from the vehicle as illustrated in Fig. 2.

The goal point progresses along the reference path so that it is always located at a distance L_1 from the aircraft. The lateral acceleration a_{cmd} is then given by:

$$a_{cmd} = 2\frac{{V_a}^2}{L_1} sin(\eta) \tag{3}$$

This guidance law gives good convergence and excellent tracking when compared with PID controllers [12], but presents several problems in this application. First, it cannot be guaranteed that there will be a point on the path that is distance L_1 away, especially when the contour is recalculated. Second, a closed path is more easily parameterized in polar coordinates. For these reasons, Park's guidance law is modified to use a constant look-ahead angle instead of distance. Use of the modified controller proceeds as:

1. The desired contour and aircraft position are shifted to put the path centroid at $(0,0)$. The path and aircraft position are then transformed to polar coordinates.

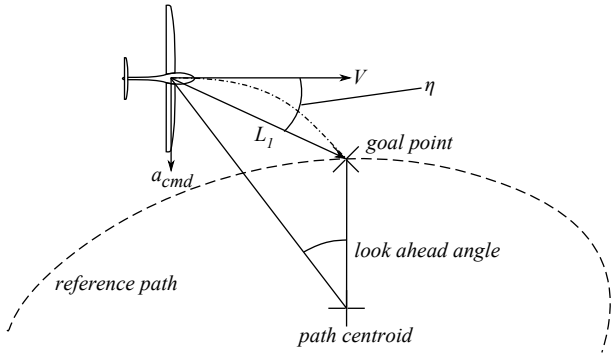

Fig. 3: Park's nonlinear guidance law modified for circular trajectories.

2. The goal point is selected to lie on the desired contour at a look-ahead angle of 15 deg.

3. L_1 is calculated as the distance from the aircraft position to the goal point.

This modified process is pictured in Fig. 3.

The lateral acceleration command is then computed as in Eq. 3. This guidance method gives accurate tracking and rapid convergence for paths that are not too complicated, but can fail for paths with overly skewed dimensions or for paths that loop back on themselves. The thermal models investigated in this paper did not present such problems. However, if real thermals prove to be sufficiently complex, this path following method may need to be revisited.

Windfield Exploration

Mapping the windfield to improve climb rate suffers from the quandary inherent in simultaneous mapping and exploitation of any resource - insufficient mapping of the wind field potentially leaves an area unexplored which could improve climb rate, but a thorough exploration takes time which degrades average climb rate. In an attempt to balance these competing objectives, a dither is applied to the aircraft goal location's radial distance from the path centroid. In this investigation a sinusoidally varying dither is applied with amplitude of 20 meters and period of 15 seconds. A dither amplitude based on the local uncertainty in the thermal map may deliver higher performance, but a fixed dither is used here for simplicity. This dither allows the aircraft to explore a region close to the current trajectory.

Simulation Results

In order to evaluate the benefit of modeling and path planning in thermals, several simulations were run for both planning and circling thermal exploitation techniques. Two types of simulation were run — a simulation to compare optimal climb rate achieved by planning and circling paths given *a priori* knowledge of the entire wind field, as well as a kinematic simulation of an aircraft flying in thermals with no prior windfield knowledge. In all simulations a perfect inner loop controller is assumed to test only the effectiveness of the outer loop guidance method. Measurements of vertical air motion were corrupted with zero mean Gaussian noise with standard deviation of 0.5 m/s in order to simulate the noise in the aircraft's sensors [11].

Table 1: Thermal initialization parameters for simulation.

Parameter	Min	Max	Mean	σ
Thermal Center (m)	-10	10	0	3
Thermal Strength (m/s)	-1.33	5.33	2	1
Number Of Cores	0	5	-	-
Core Strength (m/s)	-0.333	1.333	0.5	0.25
Core Radius (m)	0	97	30	20
Core Center (m, N or E)	-53	53	0	27

Table 2: Results for thermalling in an *a priori* known wind field.

Mean Climb Rate for Contour Planning Glider:	2.21 m/s
Mean Climb Rate for Circling Glider:	2.08 m/s
Minimum Improvement in Climb Rate:	-44%
Maximum Improvement in Climb Rate:	66%
Mean Improvement in Climb Rate:	6.2%
5th Percentile Improvement:	-12.8%
95th Percentile Improvement:	34.4%

Path Optimization with *a priori* Windfield Knowledge

To assess the potential of the contour planning method independent of the thermal model quality, a simulation was developed to plan both circular and contour paths on an *a priori* known windfield. The simulation compares contour paths with paths generated by optimizing a circular path centered at the thermal centroid, calculated using Allen's lift-weighted centroid method. Due to the complexity of the interaction between wind field structure and climb rate, a Monte Carlo approach was taken where each run was seeded with a random thermal composed of a Gaussian thermal with half-sine "cores" superimposed to form a more complex wind field. The thermal parameters and their ranges are given in Table 1.

To evaluate the climb rate achieved on a given path, the difference between aircraft sink rate (adjusted for load factor) and thermal rise rate is integrated around the path to determine mean climb rate. MATLAB's nonlinear optimization tools were then used to find the path maximizing climb rate. For the contour planning path, the contour level and aircraft speed are used as optimization targets. The circling method used circle radius and aircraft speed as independent variables. The simulation uses aerodynamic characteristics for an RnR Products SBXC sailplane, a 4.5 m span radio controlled sailplane commonly used in autonomous soaring experiments [2, 5, 6]. In 54% of cases the mapping approach showed better performance than the baseline. Results are summarized in Table 2 and flight paths for a typical case are shown in Fig. 4.

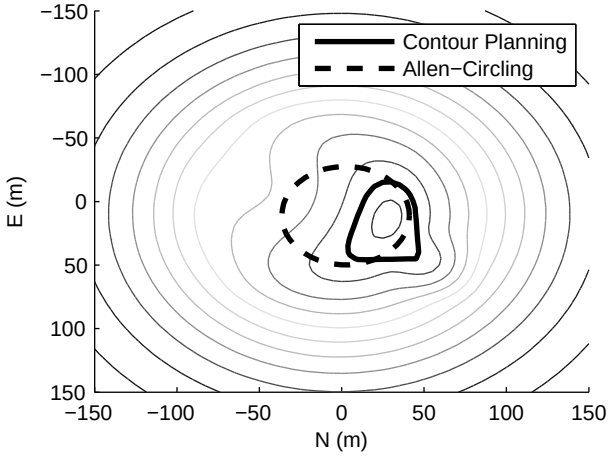

Fig. 4: Flight paths for circling and contour planning methods given *a priori* knowledge of the windfield.

Thermalling in an Unknown Windfield

A second simulation is used to evaluate the stability of the spline mapping and contour planning method as it explores and exploits a thermal. For comparison two other gliders are also simulated to compare the climb rate and flight paths for the different methods. One of the other gliders circles using Allen's method [2] and the second uses Andersson's controller [5]. For this simulation, thermals are modeled using Gedeon's single and four cell thermals of random strength and size [13]. The gliders are started at the same location at one corner of a box surrounding the thermal, with an initial heading into the box at a random angle between 0 and 90 degrees. The simulation is then run for four minutes to give the aircraft time to find and center the thermal. Figure 5 illustrates the flight paths flown by the three gliders during one such thermal encounter with a type 1 (single cell Gaussian) thermal and Table 3 summarizes the performance of the three methods.

In examining the bulk simulation results it is immediately apparent that the planning method converges to the thermal much more robustly than Allen's method. If the aircraft only grazes the thermal, often Allen's method will turn the wrong way or fail to turn in time to intercept the thermal and flies away from the area of lift. It should be noted that in both cases the aircraft controllers are only trying to thermal — there is no thresholding or logic for a thermal/cruise decision in the simulations. With the addition of such logic some of this advantage may be negated as the cases grazing the thermal in cruise may not trigger an attempt at thermalling for either method. Even so, ther-

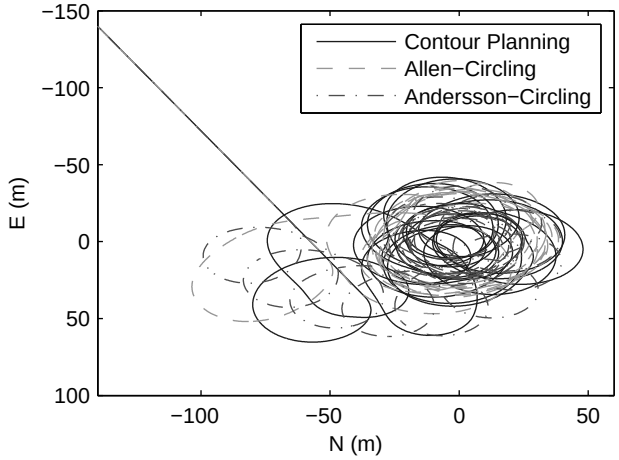

Fig. 5: Flight paths for three thermalling techniques during a four minute simulation of an encounter with a type 1 thermal, $\mathbf{C_0 = 3.2}$ m/s, $\mathbf{R = 114.46}$ m.

Table 3: Results for exploitation of type 1 thermals, averaged over 113 simulations

Climb Rate for the Full 240 Second Simulation			
	Contour Planning	Allen	Andersson
Min (m/s)	0.19	0.23	0.16
Max (m/s)	3.12	2.87	3.26
Mean (m/s)	1.47	1.42	1.46
Mean Climb Rate in the final 30 Seconds of Simulation			
	Contour Planning	Allen	Andersson
Min (m/s)	0.28	0.26	0.25
Max (m/s)	3.47	3.47	3.65
Mean (m/s)	1.64	1.60	1.74

mal mapping conveys a clear advantage in stability of convergence. In comparing Andersson's controller and the contour planning method, it is seen that Andersson's controller has very robust convergence characteristics within the thermal itself, but is very sensitive to thermal/cruise logic. If the aircraft is not definitely in the thermal when the controller begins operation then the controller will converge slowly to the center, if it converges at all. In order to ensure that the controller had a chance of succeeding it was necessary to add a logic switch to prevent Andersson's controller from operating until the aircraft had entered the thermal.

If the examination is restricted only to the cases where all aircraft successfully intercepted the thermal, the contour planning glider out climbed the Allen-circling glider by an average of 3.5%. Some of the climb advantage can be attributed to the reduced time required to center a thermal (under one turn in some situations), but, as can be seen in Fig. 6, the final climb rate is also superior, the mean climb rate in the final 30 seconds of simulation was 2.5% better for the planning glider than for the circling one. Comparing the mapping and Andersson-circling gliders, the total climb achieved is nearly identical, with the planning circling glider achieving a total climb less than 1% better on average. Comparison of the final climb rates indicates that the planning glider has an advantage in more rapid centering: despite a lower mean climb rate, the Andersson-circling glider achieved a climb rate in the final 30 seconds of simulation 5.7% better on average than did the planning glider. With a simple, Gaussian type thermal, this is to be expected as this thermal model plays to the strengths of the Andersson-circling technique. Both the planning glider and Andersson-circling glider have some room to improve climb rate in the simple Gaussian thermal. The gains used for Andersson's controller could be tuned more finely than those used here, allowing more rapid convergence. For the planning glider, a more sophisticated dithering algorithm would improve the final climb rate as the simple dithering algorithm takes the aircraft into non-optimal areas even after the thermal model has converged.

The periodic notch that can be seen in the planning method climb rate occurs at the replanning intervals. When a new contour is determined the aircraft is often some distance away from the contour and begins an aggressive maneuver to intercept the proper trajectory, temporarily increasing its sink rate.

The thermalling techniques were also tested for the type 2, four cell thermal [13]. Again, the mapping glider converges to the thermal much more consistently than the other gliders. Examining only the cases where all methods converged, the methods have nearly identical mean climb rates over the course of a four minute simulation. Examining

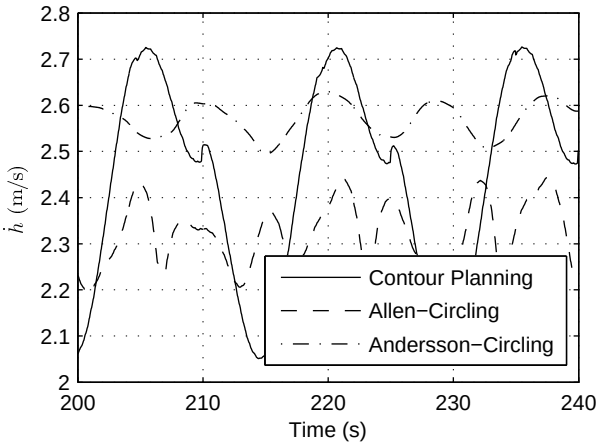

Fig. 6: Climb rate during the final 40 seconds of simulation in type 1 thermal, $\mathbf{C_0 = 3.2}$ m/s, $\mathbf{R = 114.46}$ m.

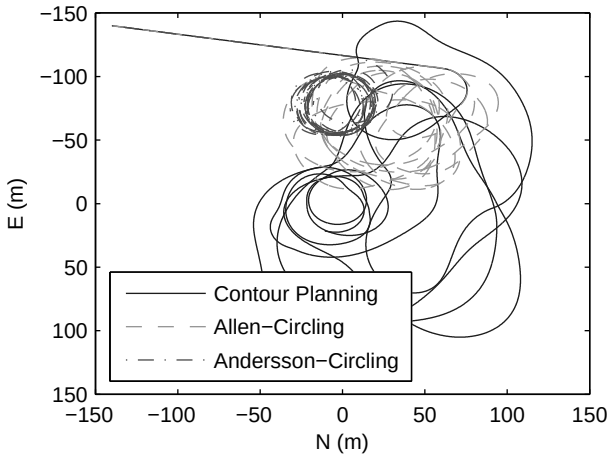

Fig. 7: Flight paths for three thermalling techniques during a four minute simulation of an encounter with a type 2 thermal, $\mathbf{C_0 = 4.4}$ m/s, $\mathbf{R = 42.93}$ m.

the final 30 seconds of climb shows that the steady-state climb rate is superior for the Allen-circling glider, with a steady state climb rate averaging 3% better than the mapping technique. The reason for the discrepancy between mean and steady-state climb rates for the four cell thermals becomes apparent when examining the flight paths in Fig. 7. With no clear maximal point in the thermal, the contour planning glider traverses an irregular trajectory as it maps the thermal. Unlike the simple Gaussian thermal which is rapidly mapped and has a clear and easily determined structure, the complexity in the type 2 thermals occasionally leads to phantom peaks in the model such as that visible in Fig. 8. Chasing these irregularities naturally leads the planning glider to fully explore the thermal and limits the uncertainty in the model, but also degrades the mean climb rate. Table 4 presents the differences in climb for several simulations using the type 2 thermal structure.

Comparing the contour planning glider and the Andersson-circling glider, the mean climb rate is similar for the two techniques. The planning glider achieves a mean climb rate 1.4% better than the Andersson-circling glider, and in the final 30 seconds of simulation the climb rate

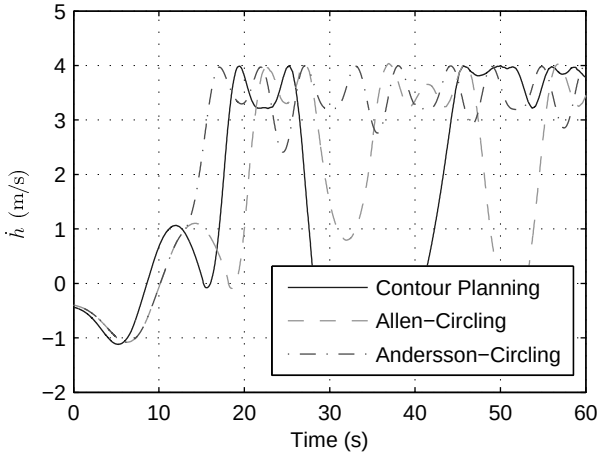

Fig. 8: Climb rate during the first 60 seconds of simulation in Konovalov type 2 thermal, $\mathbf{C_0 = 4.4}$ m/s, $\mathbf{R = 42.93}$ m.

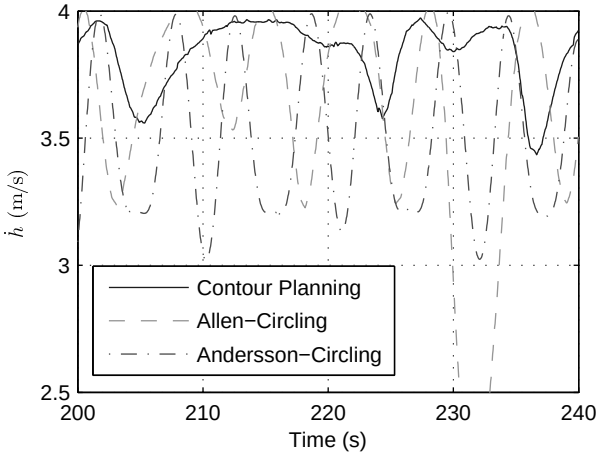

Fig. 9: Climb rate during the final 40 seconds of simulation in Konovalov type 2 thermal, $\mathbf{C_0 = 4.4}$ m/s, $\mathbf{R = 42.93}$ m.

Table 4: Results for exploitation of type 2 thermals, averaged over 127 simulations.

Climb Rate for the Full 240 Second Simulation			
	Contour Planning	Allen	Andersson
Min (m/s)	-0.02	0.06	-0.47
Max (m/s)	4.0	3.46	3.82
Mean (m/s)	1.47	1.46	1.45
Mean Climb Rate in the final 30 Seconds of Simulation			
	Contour Planning	Allen	Andersson
Min (m/s)	-0.02	0.08	-0.44
Max (m/s)	4.16	3.88	3.92
Mean (m/s)	1.62	1.67	1.55

achieved by the planning glider is 4.5% better than the Andersson-circling glider. The flight path trace bears this out - the Andersson-circling glider immediately starts turning in the edge of the thermal, achieving an initial climb rate advantage. Once the planning glider has sufficiently mapped the thermal it can catch up by flying a path in a more consistent portion of the thermal, seen in the smaller variation in climb rate depicted in Fig. 9.

The climb rates achieved by the Andersson and Allen techniques in the two thermals illustrate the sensitivity these two techniques have to assumptions built into their algorithms about thermal structure. Using the parameters specified by the authors of these controllers [2, 5], the two controllers exhibit "preferred" thermal sizes. As specified, the Andersson controller prefers a small thermal, flying tight circles which gives it good performance in the type 1 thermals with a clear and narrow core. The Allen controller prefers a larger thermal, making it better suited to centering the wide core of the type 2 thermals, where the Andersson controller ends up stuck on the edge and achieves a lower climb rate. The planning controller runs a course in between, delivering consistent performance in several thermal structures, though not achieving maximum climb rate in either.

Thermal Modeling

Mapping a thermal while soaring is useful for more than just the contour controller presented, it could enable other controller types or the tuning of existing controllers. While quantitative evaluation of such a model is problematic, an example model constructed during one simulation run is presented below. Evolution of the thermal map for one of the type 2 thermals is illustrated in Fig. 10. Qualitatively the figure shows the algorithm presented is capable of mapping even complex thermal structures.

The broken outer ring observed in the map is the result of the aircraft not flying in that region. Since the spline model is purely descriptive, windfield features will not be modeled for areas where the aircraft did not gather data.

Conclusion

A method has been presented for mapping non-uniform thermals by aircraft in soaring flight. A path-planning method using the thermal map to maximize exploitation of thermals has also been presented.

Simulations using several thermal structures show that thermal maps can be constructed by a sailplane in climb given measurements reasonably available on board the aircraft. The utility of the map is established through the performance of the contour-following controller which achieves mean climb rates similar to existing thermalling controllers, and exhibits resiliency to differing thermal structure and size. Further improvements in the thermal map quality can be made through the implementation of algorithms to automatically place the spline knots. In addition to the thermalling controller presented in this paper, the utility of thermal mapping could be extended to an adaptive thermal size for Andersson or Allen's controllers, or displaying a better picture of the lift environment to the pilot of a manned sailplane.

While the contour following controller presented here shows promise, further examination should be made of the cost incurred by frequent control surface action needed in order to follow the more complex paths generated by the controller.

Acknowledgments

Portions of the work presented here were funded by the National Science Foundation under Grant Number IIS-0746655.

References

[1] Allen, M., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *AIAA 2005-1025*, 2005.

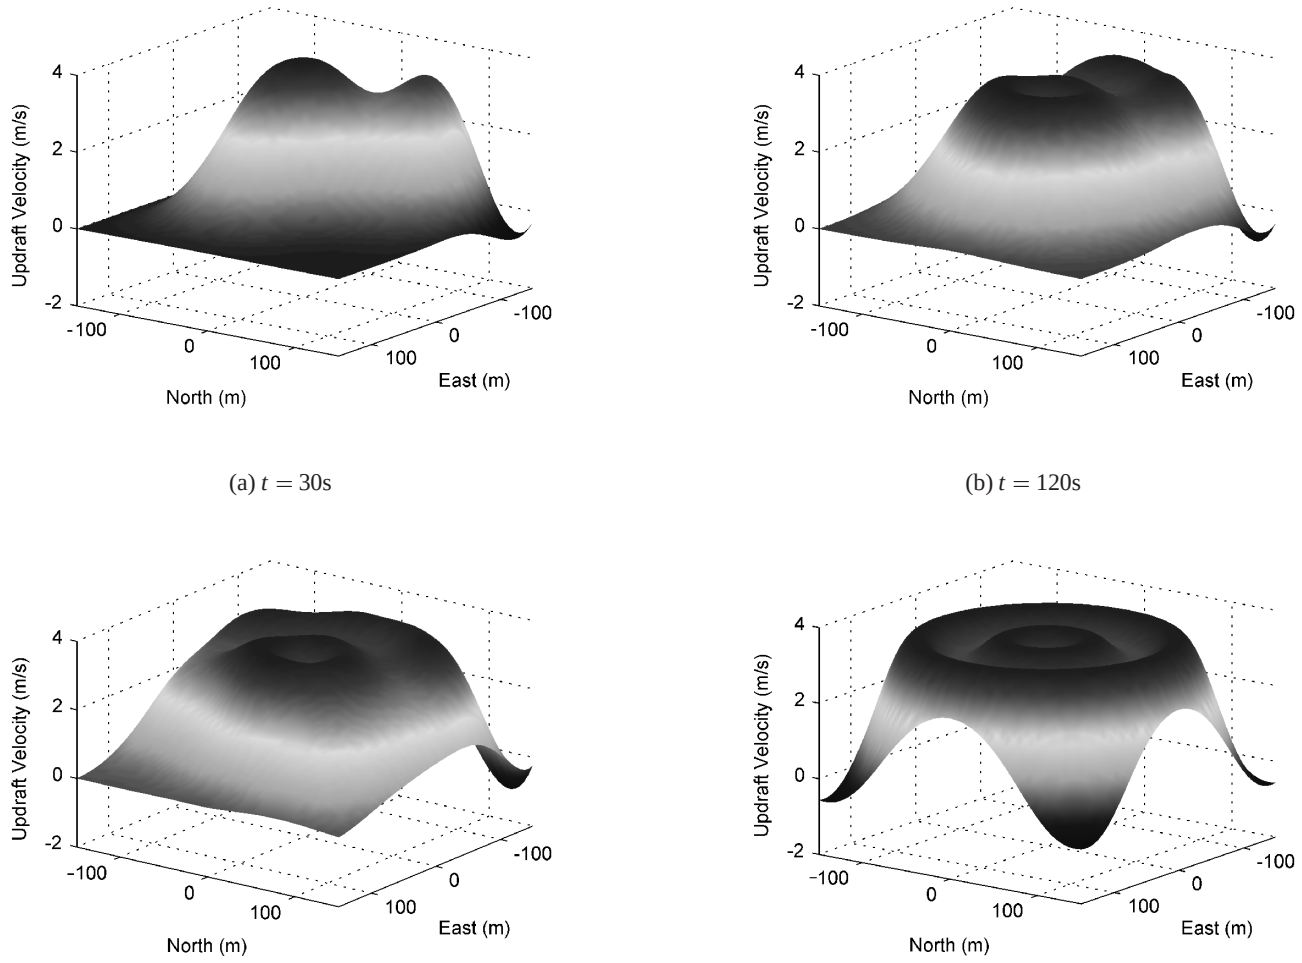

(a) $t = 30$s

(b) $t = 120$s

(c) $t = 240$s (end of simulation)

(d) True Structure

Fig. 10: Evolution of the thermal map for a type 2 thermal, $\mathbf{C_0 = 4.4}$ m/s, $\mathbf{R = 42.93}$ m. Breaks in the outer ring are in regions not sampled by the sailplane.

[2] Allen, M. and Lin, V., "Guidance and Control of an Autonomous Soaring UAV," Tech. Rep. TM-2007-214611, NASA Dryden Flight Research Center, April 2007.

[3] Andersson, K., Kaminer, I., and Jones, K., "Autonomous Soaring; Flight Test Results of a Thermal Centering Controller," *AIAA 2010-8034*, 2010.

[4] Edwards, D. and Silverberg, L., "Autonomous Soaring: The Montague Cross-Country Challenge," *Journal of Aircraft*, Vol. 47, 2010.

[5] Andersson, K. and Kaminer, I., "On Stability of a Thermal Centering Controller," *AIAA 2009-2043*, 2009.

[6] Edwards, D., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA 2008-7244*, 2008.

[7] Reichmann, H., *Cross Country Soaring*, Soaring Society of America, 1993.

[8] XCSoar, *XCSoar 6.3 User Manual*, 6th ed., July 2012.

[9] Diercx, P., *Curve and Surface Fitting with Splines*, Oxford University Press, 1993.

[10] Boor, C. D., *A Practical Guide to Splines*, Springer-Verlag, 1978.

[11] Langelaan, J., Alley, N., and Neidhoefer, J., "Wind Field Estimation for Small Unmanned Aerial Vehicles," *AIAA 2010-8177*, 2010.

[12] Park, S., Deyst, J., and How, J., "Performance and Lyapunov Stability of a Nonlinear Path-Following Guidance Method," *Journal of Guidance, Control, and Dynamics*, Vol. 30, 2007.

[13] Gedeon, J., "The Influence of Sailplane Performance and Thermal Strength on Optimal Dolphin-Flight Transition Piloting Techniques," *XV OSTIV Congress*, 1976.