

# Spline Mapping to Maximize Energy Exploitation of Non-Uniform Thermals

John J. Bird, Jack W. Langelaan  
Aerospace Engineering, The Pennsylvania State University

This paper describes a method for modeling and maximizing the use of thermals by small unmanned aerial vehicles. The method uses inertial and air-data sensors on board the aircraft to map a thermal and determine a path to fly within the thermal to maximize the energy harvested by the aircraft. Simulation results show modeling capability in simple Gaussian and non-uniform thermals and compare climb rate for map-based thermalling and simple spiral climbing techniques.

## Introduction

Thermal Soaring has been practiced by pilots of manned sailplanes since the invention of the variometer in the 1920's. 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;<sup>1</sup> flight tests by Allen, and later Andersson demonstrated that autonomous aircraft could extend endurance by harvesting energy from thermals.<sup>2,3</sup> Edwards demonstrated the use of thermalling in cross-country flight by an autonomous aircraft, placing third in competition with piloted RC aircraft.<sup>4</sup>

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.<sup>2,5-7</sup> 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.<sup>2,6</sup> Edwards does propose a method to estimate the size and orientation of an elliptical thermal,<sup>4</sup> 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,<sup>8</sup> as no automatic modeling capability is provided.

This paper presents a method using splines to model 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-filtering method is also presented which allows the model 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 is presented which uses contours of constant lift from the thermal model 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:<sup>9</sup>

$$s(x) = \sum_{i=-k}^g c_i N_{i,k+1}(x) \quad (1)$$

where  $N$  is the value at point  $x$  of a set of basis splines of order  $k$  defined on the set of knots  $\lambda_j, j = 0, \dots, g + 1$ . Calculation of  $N$  is accomplished using a triangular scheme, described by Diercx.<sup>9</sup>

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:<sup>9</sup>

$$s_{R_{i,j}} \in \mathcal{P}_k \otimes \mathcal{P}_l \quad (2)$$

If  $\mathcal{P}_k(x)$  and  $\mathcal{P}_l(y)$  are written as basis splines as in equation 1, then the resulting bivariate spline can be written as the tensor product of the two spline functions, which can be written:<sup>9</sup>

$$s(x, y) = \sum_{i=-k}^g \sum_{j=-l}^h c_{i,j} N_{i,k+1}(x) M_{j,l+1}(y) \quad (3)$$

With knots  $\lambda_i$  in  $x$  and  $\mu_j$  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<sup>10</sup> ensures that the final spline function is linear, thus  $c_{i,j}$  can be estimated using any linear estimation procedure.

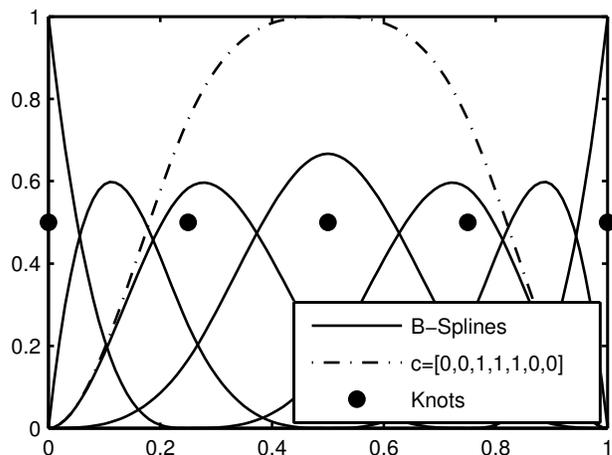


Fig. 1 Basis splines and one example spline with knots [0, 0.25, 0.5, 0.75, 1], order 3

### Thermal Modeling with Splines

The tensor product definition of a bivariate spline makes the thermal modeling 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,

$$\mathbf{h}(x) = N_{i,k+1}(x) \otimes M_{j,l+1}(y) \quad (4)$$

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 model with time:

$$\begin{aligned} \hat{\mathbf{c}}_{t|t-1} &= \hat{\mathbf{c}}_{t-1|t-1} \\ \hat{\mathbf{P}}_{t|t-1} &= \hat{\mathbf{P}}_{t-1|t-1} + \mathbf{Q}_t \end{aligned} \quad (5)$$

With the observation model defined as in equation 4 and the measurement noise chosen to represent the error in the measurement of vertical air motion, the Kalman filter update step proceeds:

$$\begin{aligned} \mathbf{K}_t &= \hat{\mathbf{P}}_{t|t-1} \mathbf{h}^T (\mathbf{h} \hat{\mathbf{P}}_{t|t-1} \mathbf{h}^T + R_t)^{-1} \\ \hat{\mathbf{X}}_t &= \hat{\mathbf{c}}_{t|t-1} + \mathbf{K}_t (w - \mathbf{h} \hat{\mathbf{c}}_{t|t-1}) \\ \hat{\mathbf{P}}_{t|t} &= (\mathbf{I} - \mathbf{K}_t \mathbf{h}) \hat{\mathbf{P}}_{t|t-1} \end{aligned} \quad (6)$$

This filter allows the model to be rapidly updated, requiring storage of only the coefficient array  $\mathbf{c}$  and its covariance matrix. The memory requirement 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  $(\mathbf{h} \hat{\mathbf{P}}_{t|t-1} \mathbf{h}^T + R_t)$  reduces to a single value.

The use of a Kalman filter means that the vertical wind velocity component,  $w$  does not have to be filtered prior to use. 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.

### Path Planning

With the more complete model 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 model 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 model, a level set can be taken of the model, describing a closed path around the estimated thermal structure which has a constant vertical wind speed. In order to optimize the aircraft climb rate, a cost function can be defined to be the mean climb rate achieved during one orbit of a contour at level  $w$  (parameterized in polar coordinates by  $\theta$ ):

$$J(w) = \oint_{C(w)} (w - z(\theta)) d\theta \quad (7)$$

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

$$\begin{aligned} \dot{z} &= \dot{z}(\phi, V_a) \\ \phi &= \tan^{-1} \left( \frac{V_a^2}{R(\theta)g} \right) \end{aligned} \quad (8)$$

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. The controller implemented in this investigation is developed from the guidance method presented by Park et al.<sup>11</sup> They present a controller 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.

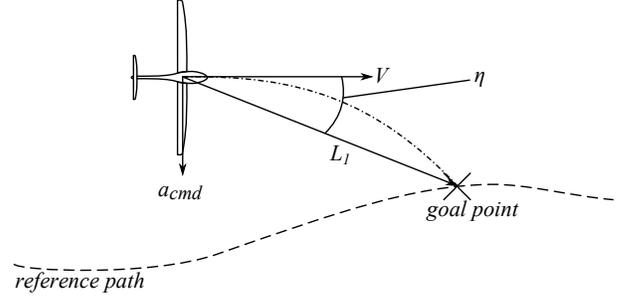


Fig. 2 Illustration of Nonlinear guidance law from Park, et al

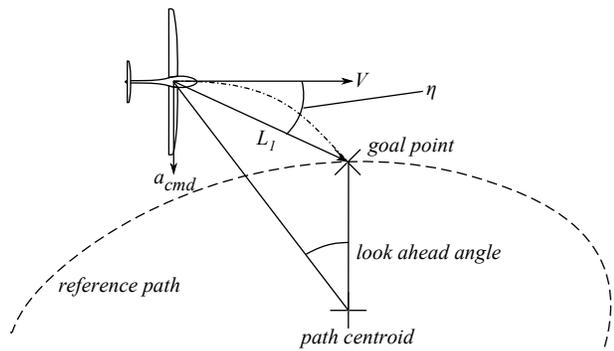
The goal point progresses along the reference path to always be at distance  $L_1$  away from the aircraft. The lateral acceleration  $a_{cmd}$  is then given by:

$$a_{cmd} = 2 \frac{V_a^2}{L_1} \sin(\eta) \quad (9)$$

This guidance law gives good convergence and excellent tracking when compared with PID controllers,<sup>11</sup> 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.
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 figure 3.



**Fig. 3 Park's Nonlinear Guidance Law Modified for Circular Trajectories**

The lateral acceleration command is then calculated as in equation 9. 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, however paths of such complexity are unlikely to be encountered even in non-uniform thermals.

### 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 windfield model 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.<sup>12</sup>

### Path Optimization with *a priori* Windfield Knowledge

To assess the potential of the contour planning method independently 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 listed below:

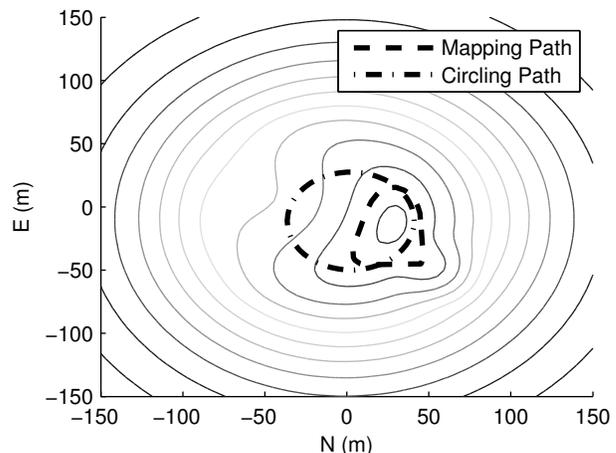
Parameter	Min	Max	Mean	$\sigma$
Thermal Center (m)	-10	10	0	3
Thermal Strength (m/s)	-1.33	5.33	2	1
Number Of Cores	0	5	-	-
Core Strength (m/s)	-0.333	1.333	0.5	0.25
Core Radius (m)	0	97	30	20
Core Center (m, N or E)	-53	53	0	27

**Table 1 Thermal Initialization Parameters for Simulation**

To evaluate the climb rate achieved on a given path, the difference between aircraft sink rate 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 based 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.<sup>2,5,6</sup> In 54% of cases the mapping approach showed better performance than the baseline.

Mean Climb Rate for Mapping 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%

**Table 2 Results for Path Planning on an *a priori* known wind field**



**Fig. 4 Flight Paths for Circling and Mapping Methods Given *a priori* Knowledge of the Windfield**

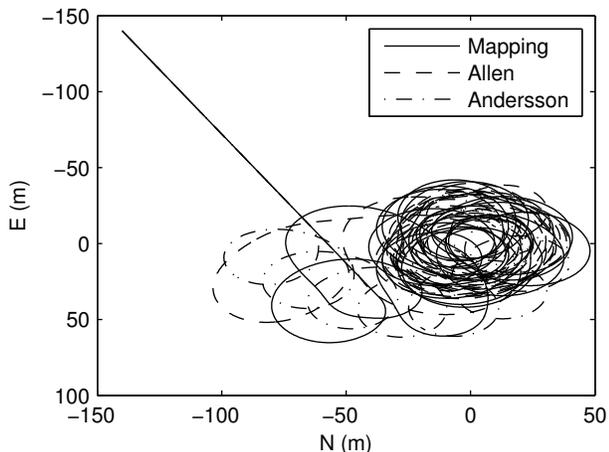
### Thermalling in an Unknown Windfield

A second simulation is used to evaluate the stability of the spline 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<sup>2</sup> and the second uses Andersson's controller.<sup>5</sup> For this simulation, thermals are modeled as Konovalov single and four cell thermals of random strength and size.<sup>13</sup> 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 Konovalov type 1 (single cell Gaussian) thermal.

Climb Rate for the Full 240 Second Simulation			
	Mapping	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			
	Mapping	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

**Table 3 Results for Thermal Exploitation of a Konovalov Type 1 Thermal**

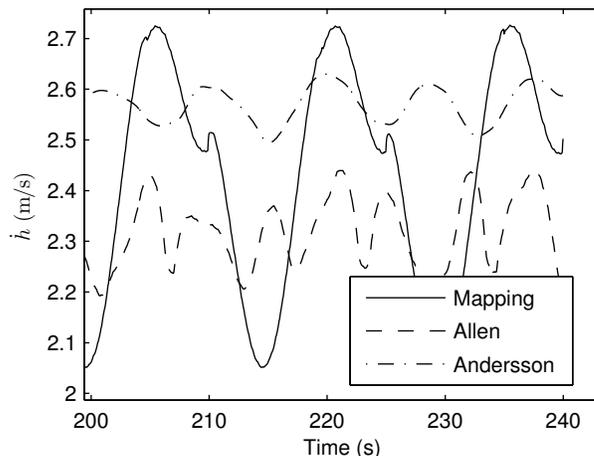
In examining the bulk simulation results it is immediately apparent that the planning method converges to the thermal much more robustly than Allen’s method. If the aircraft only grazes the thermal, often Allen’s method will turn the wrong way or fail to turn in time to intercept the thermal and flies 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, there is a clear advantage in stability of convergence for the mapping controller. In comparing Andersson’s controller and the 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 the thermalling controller from operating until the aircraft had entered the thermal.



**Fig. 5 Flight Paths for Three Thermalling Techniques During a Four Minute Simulation of an Encounter with a Konovalov Type 1 Thermal  $C_0 = 3.2m/s$ ,  $R = 114.46m$**

If the examination is restricted only to the cases where the aircraft successfully intercepted the thermal (approximately half of the total runs), the 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 figure 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.



**Fig. 6 Climb Rate During the Final 40 seconds of Simulation in Konovalov Type 1 Thermal  $C_0 = 3.2m/s$ ,  $R = 114.46m$**

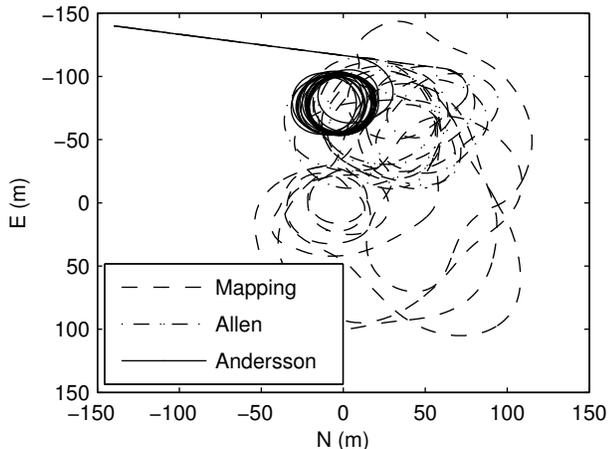
The periodic notch that can be seen in the planning method climb rate occurs at the replanning intervals. When a new contour is determined the aircraft is often some distance away from the contour and begins an aggressive maneuver to intercept the proper trajectory, temporarily increasing its sink rate.

Climb Rate for the Full 240 Second Simulation			
	Mapping	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			
	Mapping	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

**Table 4 Results for Thermal Exploitation of a Konovalov Type 2 Thermal**

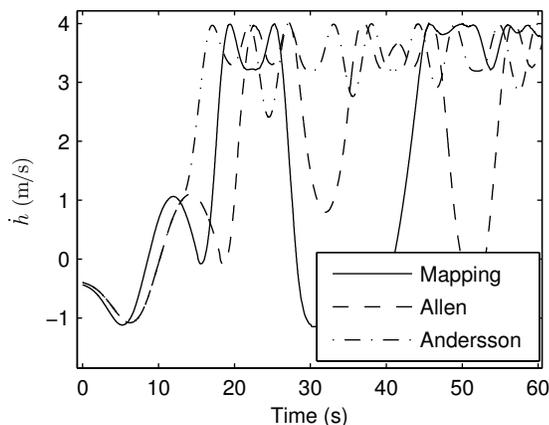
Both circling techniques were also tested for the Konovalov type 2 four cell thermal.<sup>13</sup> Again, the convergence advantage of the planning method can be seen with about twice the convergence rate of Allen’s method. Examining only the converged thermals, the two methods have nearly identical mean climb rates over the course of a four minute simulation. Examining 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 Figure 7. With no clear maximal point in the thermal, the mapping glider traverses an irregular trajectory as it explores 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. Chasing these irregularities naturally leads the mapping 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.



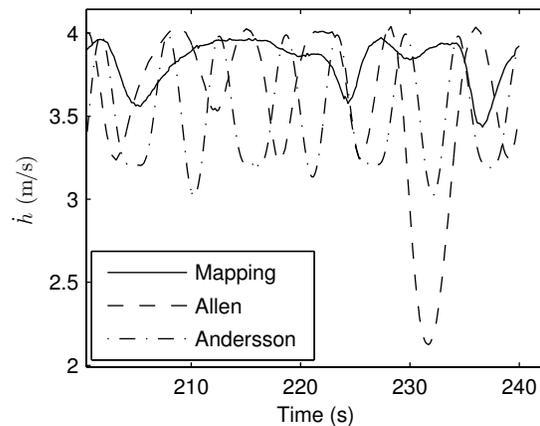
**Fig. 7 Flight Paths for Three Thermalling Techniques During a Four Minute Simulation of an Encounter with a Konovalov Type 2 Thermal  $C_0 = 4.4m/s$ ,  $R = 42.93m$**

Comparing the mapping glider and the Andersson-circling glider, the mean climb rate is similar for the two techniques. The mapping 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 achieved by the mapping 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, but 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.



**Fig. 8 Climb Rate During the First 60 seconds of Simulation in Konovalov Type 2 Thermal  $C_0 = 4.4m/s$ ,  $R = 42.93m$**

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,<sup>2,5</sup> 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

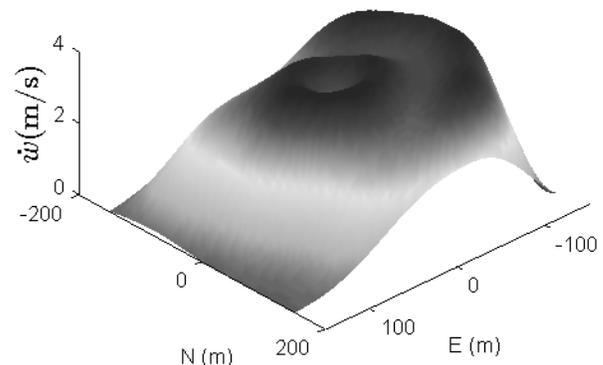


**Fig. 9 Climb Rate During the Final 40 seconds of Simulation in Konovalov Type 2 Thermal  $C_0 = 4.4m/s$ ,  $R = 42.93m$**

controller ends up stuck on the edge and achieves a lower climb rate. The mapping controller runs a course in between, delivering consistent performance in several thermal structures, though not achieving maximum climb rate in either.

### Thermal Modeling

Modelling a thermal while soaring is useful for more than just the contour controller presented earlier, 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. The estimated thermal model for one of the Konovalov type 2 thermals is illustrated in figure 10, with the true structure in figure 11. Qualitatively the figure shows the algorithm presented is capable of modeling even complex thermal structures.



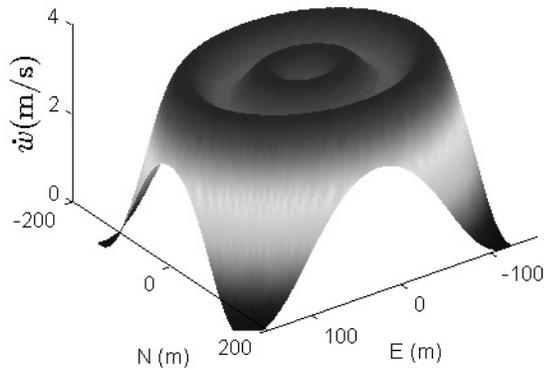
**Fig. 10 Modeled Structure for a Konovalov Type 2 Thermal,  $C_0 = 3.53m/s$ ,  $R = 61m$**

The broken outer ring observed in the model 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 modeling non-uniform thermals by aircraft in soaring flight. A path-planning method using the thermal model to maximize exploitation of thermals has also been presented.

Simulations using several thermal structures show that models of the thermal can be constructed by a sailplane in climb given measurements reasonably available on board the aircraft. The utility of the model is established through



**Fig. 11 True Structure for a Konovalov Type 2 Thermal,**  
 $C_0 = 3.53\text{m/s}$ ,  $R = 61\text{m}$

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 model quality can be made through the implementation of algorithms to automatically place knots in the spline model. In addition to the thermalling controller presented in this paper, the utility of thermal modeling 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.

## References

- <sup>1</sup>Allen, M., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *AIAA 2005-1025*, 2005.
- <sup>2</sup>Michael Allen, V. L., "Guidance and Control of an Autonomous Soaring UAV," Tech. Rep. TM-2007-214611, NASA Dryden Flight Research Center, April 2007.
- <sup>3</sup>Klas Andersson, Isaac Kaminer, K. J., "Autonomous Soaring; Flight Test Results of a Thermal Centering Controller," *AIAA 2010-8034*, 2010.
- <sup>4</sup>Daniel Edwards, L. S., "Autonomous Soaring: The Montague Cross-Country Challenge," *Journal of Aircraft*, Vol. 47, 2010.
- <sup>5</sup>Klas Andersson, I. K., "On Stability of a Thermal Centering Controller," *AIAA 2009-2043*, 2009.
- <sup>6</sup>Edwards, D., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA 2008-7244*, 2008.
- <sup>7</sup>Reichmann, H., *Cross Country Soaring*, Soaring Society of America, 1993.
- <sup>8</sup>XCSoar, *XCSoar 6.3 User Manual*, 6th ed., July 2012.
- <sup>9</sup>Diercx, P., *Curve and Surface Fitting with Splines*, Oxford University Press, 1993.
- <sup>10</sup>Boor, C. D., *A Practical Guide to Splines*, Springer-Verlag, 1978.
- <sup>11</sup>Sanghyuk Park, John Deyst, J. H., "Performance and Lyapunov Stability of a Nonlinear Path-Following Guidance Method," *Journal of Guidance, Control, and Dynamics*, Vol. 30, 2007.
- <sup>12</sup>Jack Langelaan, Nicholas Alley, J. N., "Wind Field Estimation for Small Unmanned Aerial Vehicles," *AIAA 2010-8177*, 2010.
- <sup>13</sup>Gedeon, J., "The Influence of Sailplane Performance and Thermal Strength on Optimal Dolphin-Flight Transition Piloting Techniques," *XV OSTIV Congress*, 1976.