

Improving Autonomous Soaring via Energy State Estimation and Extremum Seeking Control

Shawn C. Daugherty*

Jack W. Langelaan†

The Pennsylvania State University, University Park, PA 16802, USA

This paper introduces autonomous soaring methods that enhance the performance of small autonomous gliders in a thermal soaring environment. Thermal centering control is aided by an asymmetric Savitzky-Golay filter that computes estimates of total energy, rate of change of total energy and the second derivative of total energy using polynomial approximations over a moving time window. Climb rate in the thermal is maximized using extremum seeking control with turn radius as the varying parameter. A simulation environment based on a commercially available multiplayer soaring simulator is described, with low level aircraft control implemented on an Arduino Mega single board computer. Higher level control is implemented on a laptop computer that communicates with the Arduino autopilot over a serial link. The utility of the thermal soaring controller is demonstrated in this high fidelity simulation: stable thermal centering and good convergence to a maximum climb rate is observed, with climb performance of the new controllers exceeding previous methods.

I. Introduction

SMALL UNMANNED AERIAL vehicles (UAVs) have become prevalent in military operations (the RQ-11 Raven and Wasp III are two examples) and have significant potential for civil and scientific missions such as environmental monitoring and meteorology. However, range and endurance of small UAVs is limited both by the relatively small payload capacity (which limits the sensing package that can be carried as well as fuel or batteries) and by the lower aerodynamic efficiency typical of smaller aircraft (due to lower operating Reynolds numbers). Flight control methods that enhance range and endurance can therefore greatly improve overall utility of these vehicles.

Large birds such as hawks, eagles and vultures as well as human sailplane pilots exploit energy available in the atmosphere to fly several hours and cover distances of many hundreds of kilometers without flapping wings or the use of engines. Exploiting this atmospheric energy is known as soaring, and it has the potential to significantly change the missions that can be flown with small UAVs.

Autonomous soaring flight has become the focus of a significant amount of research over the past several years, with both simulations and hardware demonstrations showing the utility of soaring for extending endurance of small UAVs.¹⁻³ Soaring by manned gliders has also been the subject of a large amount of research.⁴⁻⁸

This paper focuses on autonomous soaring in thermals (a column of rising, warm air, triggered by uneven solar heating of the ground and instability in the atmosphere). Specifically, it seeks to maximize the rate of energy gain during thermal soaring by: proposing an energy estimation method that increases thermal centering performance and disturbance rejection; introducing an extremum seeking control method for optimizing the rate of energy gain once the thermal is centered.⁹

The remainder of this paper is organized as follows. Section II discusses background and related research. Section III describes the kinematics of soaring flight and outlines the equations governing thermal centering. Section IV describes the use of Savitzky-Golay filters for estimating total energy and rates of change of

*Graduate Student, Department of Aerospace Engineering, Student Member AIAA.

†Associate Professor, Department of Aerospace Engineering, Associate Fellow AIAA.

total energy. Section V describes extremum seeking control as it applies to autonomous thermal soaring. Section VI describes the Silent Wings simulation environment and Section VII presents results of simulations that show the utility of the thermalling controller proposed here. Finally, concluding remarks are presented in Section VIII.

II. Background and Related Work

Autonomous soaring UAVs were proposed by John Wharington in 1998, with techniques such as reinforcement learning used for flight control.^{10,11} In 2005, Michael Allen published an analysis that showed that the endurance of small UAVs could be significantly increased using thermal updrafts.¹² In 2007, Allen proposed and demonstrated a soaring controller that is based on Reichmann's method of thermal soaring.^{1,13} Allen's controller estimated the total energy of the aircraft and the required turn rate ($\dot{\psi}$) to properly center the thermal at each time step. In addition, a queue of total energy estimates were kept and updated to estimate the size, strength, and location of the thermal by iteratively fitting an assumed velocity profile. Using this information, the turn rate was adjusted to correct for thermal drift caused by prevailing winds. Allen successfully performed test flights using the proposed centering controller. According to Allen, the UAV climbed an average of 172 meters (567 feet) while experiencing 23 different thermals. However, the controller's performance diminished in smaller, weaker thermals due to latency in energy rate and energy acceleration estimates.¹

In 2008, Daniel Edwards built upon Allen's research by improving thermal location estimates and demonstrated successful thermal soaring by placing third in the Montague Cross-Country Challenge using a 5 kg SB-XC glider named ALOFT (autonomous locator of thermals).^{2,14}

In 2009, Andersson et al. performed an analysis demonstrating the benefits of using a cooperating team of UAVs to increase the probability of finding lift.¹⁵ It was shown that stable thermal centering can be achieved by using strictly feedback control and omitting thermal drift estimation.¹⁶ In 2012, Andersson et al. published successful flight test results using this centering controller. However, this thermal centering controller suffered from degradation in thermal centering performance (especially in smaller, weaker thermals) due to significant latency in the energy estimates.³ A modified centering controller was described that provided phase lead by feeding back the third derivative of energy. However, a quantification of the improvement provided by this altered controller was not published, it was described as potential future work.

III. Thermal Centering Control

A. Kinematics of a Glider in a Steady Turn

In order to acquire energy from a thermal, the glider must remain inside the region of rising air. This requires the glider to remain in a turn about the thermal core. When an aircraft is in a steady turn, a component of the lift vector is pointed towards the center of the turn. This component of force provides the centripetal acceleration required for the turn. The vertical component of the lift vector counteracts weight, and in constant speed flight this vertical component must equal weight exactly. As a result, the total lift force in a turn is greater than in level flight, thus the aircraft must fly at a higher lift coefficient. A free-body diagram showing this phenomenon is shown in Figure 1. This kinematics approach assumes perfect control of the vehicle; that is the low-level flight controller is able to track heading, bank angle, and airspeed commands with negligible delay. This assumption allows the motion of the vehicle to be represented using a point-mass model. Vehicle kinematics are given by:




Figure 1. Body forces in a steady turn.

$$\dot{x} = v_a \cos \gamma \cos \psi + w_x \quad (1)$$

$$\dot{y} = v_a \cos \gamma \sin \psi + w_y \quad (2)$$

$$\dot{z} = v_a \sin \gamma + w_z \quad (3)$$

where v_a is airspeed, γ is flight-path angle relative to the surrounding airmass, ψ is heading, ϕ is bank angle, and w_x , w_y , w_z represent the velocity of the airmass.¹⁷ The summation of the forces in the stability axes for a glider performing a steady turn are:

$$L \sin \phi = \frac{mv_a^2}{R} \quad (4)$$

$$L \cos \phi = mg \cos \gamma \quad (5)$$

$$D = mg \sin \gamma \quad (6)$$

where R is the radius of the turn, L is the lift force and D is the drag force. To determine an equation that relates turn rate ($\dot{\psi}$) to turn radius (R), let's focus on the forces acting in the radial direction. Assuming a steady turn with no sideslip, Equation 4 can be written as:

$$L \sin \phi = mv_a \dot{\psi} \quad (7)$$

Dividing equation Equation 7 by Equation 5 gives:

$$\tan \phi = \frac{v_a \dot{\psi}}{g \cos \gamma} \quad (8)$$

This expression can be rearranged, solving for $\dot{\psi}$:

$$\dot{\psi} = \frac{g \tan \phi \cos \gamma}{v_a} = \frac{v_a \cos \gamma}{R} \quad (9)$$

Sink rate is determined by the aerodynamic properties of the aircraft. Using the standard definition of lift and drag coefficients, Equations 5 and 6 become:

$$\frac{qSC_L \cos \phi}{mg} = \cos \gamma \quad (10)$$

$$\frac{qSC_D}{mg} = \sin \gamma \quad (11)$$

Dividing Equation 11 by Equation 10:

$$\frac{C_D}{C_L \cos \phi} = \tan \gamma \quad (12)$$

Relating C_D as a function of C_L , the drag polar can be approximated using a polynomial of order n :

$$C_D = \sum_{i=0}^n a_i C_L^i \quad (13)$$

Drag polar coefficients for the SB-XC are given in the Appendix. Substituting Equation 13 into Equation 12:

$$\gamma = \tan^{-1} \left(\frac{\sum_{i=0}^n a_i C_L^i}{C_L \cos \phi} \right) \quad (14)$$

where C_L is a function of airspeed. Equation 14 gives an expression for flight path angle as a function of airspeed and bank angle. This implies that given a vehicle drag polar, $\dot{\psi}$ can be written is a function of only v_a and ϕ . Now that we have an expression that relates $\dot{\psi}$, v_a , and ϕ , turn rate can be controlled by scheduling v_a and ϕ commands.

B. Autonomous Thermal Centering

The autonomous thermal centering controller that was developed by Allen¹ and altered by Andersson³ has been shown to sufficiently center and maintain climb without the use of thermal mapping. This controller is based on Reichmann's thermal centering method.¹³ Reichmann's centering method consists of the following three rules (as illustrated in Figure 2):

1. As climb improves, flatten the circle (approximately 15–20 degree bank angle)
2. As climb deteriorates, steepen the circle (approximately 50 degree bank angle)

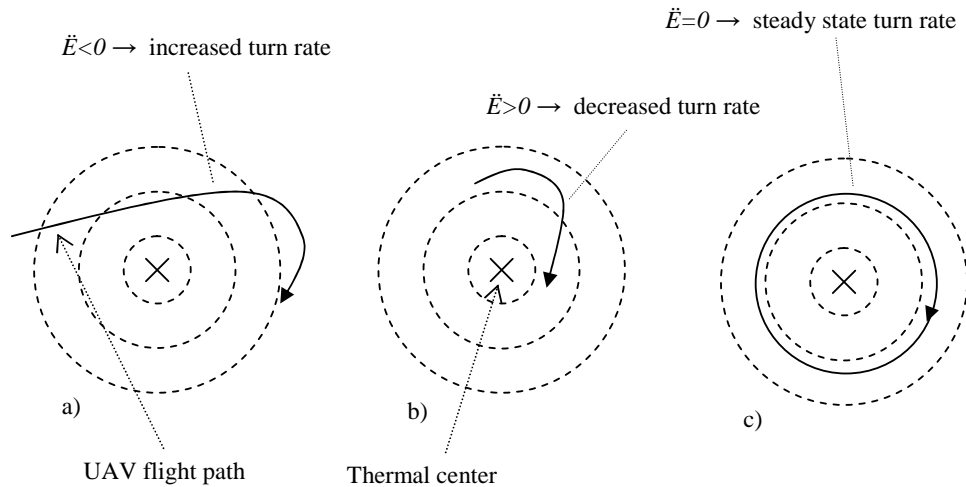


Figure 2. Reichmann's method of centering a thermal updraft using total energy acceleration.^{13, 15}

3. If climb remains constant, keep constant bank (approximately 25–30 degree bank angle)

Using the second derivative of energy (\ddot{e}) as a feedback term, this method was written as a control law:

$$\dot{\psi}_{cmd} = \dot{\psi}_{ss} - k\ddot{e} \quad (15)$$

where $\dot{\psi}_{ss}$ is the desired nominal steady-state turn rate and k is a gain.¹⁵ This controller centers the aircraft's trajectory around the core of the thermal by adjusting the turn rate to follow a contour of constant \dot{e} . This control law has the built-in assumption that the thermal profile somewhat resembles a Gaussian, circular in shape with the highest updraft velocity at the core.

Choosing the steady-state turn rate ($\dot{\psi}_{ss}$) is equivalent to choosing the steady-state circling radius (R_{ss}) when the glider is holding a constant airspeed. The equation for steady-state turn rate (assuming a small flight path angle) is

$$\dot{\psi}_{ss} = \frac{v_a}{R_{ss}} \quad (16)$$

The optimal steady-state turn rate is dependent on the sink polar of the aircraft and the characteristics of the thermal. An important question now is to determine the values of v_a and R_{ss} that maximize energy gain in the thermal: this is discussed in Section V.

Climb performance is also affected by how quickly and accurately the thermal is centered. Using the centering controller presented in Equation 15, energy "acceleration" must be determined accurately without introducing excessive phase lag into the system.

IV. Estimating Total Energy and its Rates of Change

The total energy of an aircraft can be represented by the sum of its potential, kinetic, and stored energies:

$$E = mgh + \frac{1}{2}mv^2 + E_s \quad (17)$$

In this analysis, we are primarily concerned with the aircraft's climb rate and the height of the aircraft above the ground. For this reason, Equation 17 can be modified to show the vehicle's specific energy using units of height. Specific energy (e) is total energy divided by weight.

$$e = h + \frac{v^2}{2g} + e_s \quad (18)$$

In the case of a glider, we are assuming no propulsion source and v is the velocity of the aircraft relative to the wind (v_a). To optimize the amount of energy harvested from a thermal, the rate of energy gain (\dot{e})

must be maximized. Differentiating Equation 18, \dot{e} can be written as follows:

$$\dot{e} = \frac{v_a \dot{v}_a}{g} + \dot{h} \quad (19)$$

Note that \dot{v}_a is the rate of change of airspeed, and this cannot be measured directly using accelerometers (which provide a measure of acceleration, i.e. the rate of change of inertial speed). Differentiating once more gives an expression for \ddot{e} :

$$\ddot{e} = \frac{\dot{v}_a^2 + v_a \ddot{v}_a}{g} + \ddot{h} \quad (20)$$

This \ddot{e} term is used for the thermal centering controller as discussed in Section III. Obtaining an measurement of \ddot{e} is tricky, however: numerically differentiating total energy results in a *very* noisy signal; filtering the the numerical derivative to remove enough noise that the signal is recognizably smooth results in unacceptably large lag in the signal. A filter that can compute a smooth estimate of total energy and its derivatives with acceptable lag is therefore required.

A. Savitzky-Golay Filter

Adaptive filters can use a high order estimate and still provide sufficient noise rejection. An example of an adaptive convolution filter is the Savitzky-Golay (SG) filter.¹⁸ This filter performs a polynomial fit on the measurements within a moving window of time (Figure 3), and it is adaptive in the sense that the computed polynomial fit to the measured values is continuously changing with time.

The SG filter (along with other adaptive filters) require that the window contains an overdetermined system, meaning the number of points in the window is greater than the order of the polynomial plus one. This allows the polynomial coefficients to be determined using the Gauss-Markov theorem. This theorem states that the best linear unbiased estimate (BLUE) of the polynomial coefficients is given by the ordinary least squares estimate.¹⁹ Given a window of time T containing M measurements, and a desired polynomial of order N of the following form:

$$p(x) = a_0 + a_1x + \dots + a_{N-1}x^{N-1} + a_Nx^N \quad (21)$$

where $1 < N + 1 < M$, the minimum variance estimate of the polynomial coefficients

$$\hat{\mathbf{a}} = \begin{bmatrix} a_0 & a_1 & \dots & a_{N-1} & a_N \end{bmatrix}^T \quad (22)$$

can be found using the linear least squares solution. The linear least squares estimate of $\hat{\mathbf{a}}$ is




Figure 3. Approximating numerical derivatives using a Savitzky-Golay filter.

$$\hat{\mathbf{a}} = \mathbf{K}\mathbf{y} \quad (23)$$

where \mathbf{y} is the measurement vector and \mathbf{K} is

$$\mathbf{K} = (\mathbf{D}^T \mathbf{D})^{-1} \mathbf{D}^T \quad (24)$$

Because we are estimating the polynomial coefficients, \mathbf{D} is the Vandermonde matrix:

$$\mathbf{D} = \begin{bmatrix} 1 & x_1 & x_1^2 & \dots & x_1^N \\ 1 & x_2 & x_2^2 & \dots & x_2^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_M & x_M^2 & \dots & x_M^N \end{bmatrix} \quad (25)$$

Using M measurements taken every time step (Δt), the size of the window is:

$$T = (m - 1)\Delta t \quad (26)$$

This time parameter is used to normalize the window, preventing any scaling problems that can arise from fitting higher order polynomials. The equivalent Vandermonde matrix ($\bar{\mathbf{D}}$) becomes

$$\bar{\mathbf{D}} = \begin{bmatrix} 1 & \frac{t_0-t_0}{T} & \left(\frac{t_0-t_0}{T}\right)^2 & \dots & \left(\frac{t_0-t_0}{T}\right)^N \\ 1 & \frac{t_1-t_0}{T} & \left(\frac{t_1-t_0}{T}\right)^2 & \dots & \left(\frac{t_1-t_0}{T}\right)^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \frac{t_M-t_0}{T} & \left(\frac{t_M-t_0}{T}\right)^2 & \dots & \left(\frac{t_M-t_0}{T}\right)^N \end{bmatrix} \quad (27)$$

where t_m , $m = 0 \dots M$ is the time each corresponding measurement was taken and t_0 is the oldest timestamp in the current window. Using this Vandermonde matrix normalizes the window on the interval of $[0, 1]$. The coefficients of this normalized polynomial fit (denoted by $\bar{\mathbf{a}}$) can now be estimated using

$$\bar{\mathbf{a}} = \bar{\mathbf{K}}\mathbf{y} \quad (28)$$

where \mathbf{y} is the vector of measurements in the current window.

Once this normalized polynomial fit is acquired, the polynomial is evaluated at the center point of the current window. In addition, the first and second derivatives at this point can be approximated and evaluated by analytically differentiating the polynomial fit. Once this point is evaluated, the window is shifted to incorporate the new measurement and the oldest point is eliminated.

If the number of points contained in the window (M) is equal to the order of the polynomial fit plus one ($N + 1$) the system is no longer overdetermined and application of the SG filter simply becomes the central difference solution of order (N). To demonstrate, consider a first order polynomial ($N = 1, M = 2$), and \mathbf{D} becomes:

$$\mathbf{D} = \begin{bmatrix} 1 & t - \Delta t \\ 1 & t + \Delta t \end{bmatrix} \quad (29)$$

Substituting into Equation 24, \mathbf{K} becomes:

$$\mathbf{K} = \frac{1}{2} \begin{bmatrix} \frac{t}{\Delta t} + 1 & -\frac{t}{\Delta t} + 1 \\ -\frac{1}{\Delta t} & \frac{1}{\Delta t} \end{bmatrix} \quad (30)$$

Given two measurements, the measurement matrix for a central difference is

$$\mathbf{y} = \begin{bmatrix} f(t - \Delta t) \\ f(t + \Delta t) \end{bmatrix} \quad (31)$$

Substituting into Equation 28 yields the coefficients of the polynomial fit. The second coefficient is the numerical derivative (or slope) of the curve at time t . This estimated derivative evaluates to

$$f'(t) = \frac{f(t + \Delta t) - f(t - \Delta t)}{2\Delta t} \quad (32)$$

which is exactly equal to the first order central difference solution. Using an SG filter (where $M > N + 1$ by definition), noise can be further reduced by essentially acting as a low-pass filter.

Performance of the SG filter in estimating polynomial coefficients is briefly demonstrated in Figure 4. An analytic function

$$y(t) = (t/10) \sin(2t) + 3t + 10 \quad (33)$$

is evaluated at increments of $\Delta t = 0.02$ and corrupted by zero-mean Gaussian random noise with $\sigma = 0.1$. The polynomial coefficients generated by the Savitzky-Golay filter were used to calculate the derivative at each time step (using a third order polynomial and a window size of two seconds). Figure 4 shows that smooth derivatives can be obtained (and, on the scale of the plot, the first derivative obtained using the SG filter is lies on top of the analytic function's derivative), while numerical derivatives (shown in gray) are enormously affected by measurement noise.

Note, however, that the symmetric SG filter is non-causal (it requires future measurements to compute current polynomial coefficients). Equivalently, it provides estimates that are delayed by half the time window. To reduce the effect of delay, an asymmetric Savitzky-Golay filter can be used: rather than computing estimates at the center of the time window, estimates are computed at the end of the time window.




Figure 4. Estimating first and second derivatives of a noisy signal ($\sigma = 0.1$) using a backward difference filter and a symmetric Savitzky-Golay filter.

B. Savitzky-Golay Filter for Total Energy Estimation

Good thermal centering performance requires that \dot{e} and \ddot{e} are available with minimal (ideally zero) delay. Given a polynomial order N and a window size T , the least-squares estimate of the polynomial fit to total energy variation over the window is computed. At time t_m ,

$$e(t_m) = a_0 + a_1 \left(\frac{t_m - t_0}{T} \right) + \dots + a_N \left(\frac{t_m - t_0}{T} \right)^N + \mathcal{N}(0, \sigma^2) \quad (34)$$

where t_0 is time at the beginning of the window and $\mathcal{N}(0, \sigma^2)$ denotes zero-mean Gaussian random measurement noise. Given $m = 1 \dots M$ measurements, an estimate $\hat{\mathbf{a}}$ for the polynomial coefficients is computed (as in the preceding section), and finally total energy (and its derivatives) are computed at the desired time:

$$\hat{e}(t_m) = \sum_{n=0}^N \hat{a}_n \left(\frac{t_m - t_0}{T} \right)^n \quad (35)$$

$$\dot{\hat{e}}(t_m) = \sum_{n=1}^N n \hat{a}_n \left(\frac{t_m - t_0}{T} \right)^{(n-1)} \quad (36)$$

$$\ddot{\hat{e}}(t_m) = \sum_{n=2}^N n(n-1) \hat{a}_n \left(\frac{t_m - t_0}{T} \right)^{(n-2)} \quad (37)$$

To minimize delay, total energy and its derivatives are computed at the end of the time window (i.e.

when $t_m = t_0 + T$, so $(\frac{t_m - t_0}{T}) = 1$. The energy estimates reduce down to

$$\hat{e}(t_m) = \sum_{n=0}^N \hat{a}_n \quad (38)$$

$$\dot{\hat{e}}(t_m) = \sum_{n=1}^N n \hat{a}_n \quad (39)$$

$$\ddot{\hat{e}}(t_m) = \sum_{n=2}^N n(n-1) \hat{a}_n \quad (40)$$

essentially creating an asymmetric version of the SG filter. However, as one can expect, the asymmetric filter's estimates are not as smooth as those produced by the symmetric filter. This is illustrated in Figure 5, which uses the analytic function (given in Equation 33) corrupted with zero-mean Gaussian random noise ($\sigma = 0.1$) to test the performance of both the symmetric SG filter and the asymmetric (backward) SG filter.




Figure 5. Estimating first and second derivatives of a noisy signal ($\sigma = 0.1$) using the backwards SG filter.

In Figure 5, the estimates that were computed using the symmetric filter are delayed by half the time window (1 second) to show when the estimates would actually be available to a controller. There is a clear reduction in latency at the cost of overshoot and a noisier estimate.

When the measurements change slope quickly, the polynomial fit often overshoots this change, causing a delay in the derivative estimates. By altering the window size and fit order, we can tune the filter to perform well by adequately rejecting noise with minimal overshoot and reasonable latency. Cases were run for several different combinations of window size and fit order for the function given in Equation 33. The average error for each of the cases is shown Figure 6.

This figure demonstrates that fitting small time windows (containing few measurements) with a high order polynomial yields a great deal of error. Intuitively, this makes sense because higher order polynomials can have many more turning points than lower order polynomials, causing the fit to track the measurement




Figure 6. Average error for various polynomial orders and window sizes.

noise as well as the base signal. This is the equivalent of setting the filter's crossover frequency too high. Furthermore, fitting a large window with a polynomial of low order will shift the crossover frequency of the filter the opposite direction to a point where the underlying data is being significantly affected. The best performance exists at some balancing point between window size and fit order (and may be limited by available computing hardware). The larger the window, the higher the fit order should be. The combination that yields the best performance is dependent upon the dynamics of the aircraft and the quality of the sensors. For testing purposes, the backwards SG filter was implemented on a low-level autopilot. This autopilot was used in a feedback loop to remotely fly the SB-XC in Silent Wings flight simulator (www.silentwings.no). Andersson's centering controller, with the addition of the turn radius convergence algorithm as discussed in Section V, was used to heuristically tune the energy estimator.

V. Maximizing Rate of Energy Gain Using Turn Radius Convergence

Once the thermal is centered using the energy estimation techniques discussed in Section IV, the energy estimates at various turn radii can also be used to continually improve the turn radius and the associated steady-state turn rate of the aircraft.

The climb rate in a thermal is the combination of the sink rate of the aircraft in still air and the thermal updraft velocity. The sink rate in still air can be determined from the aircraft's sink polar and adapted for various bank angles. Assuming the thermal updraft velocity follows a Gaussian distribution with a set core strength w_0 and radius R_{th} , vertical wind speed varies with distance from the core as

$$w_{th}(x, y) = w_0 e^{-\left(\frac{(x-x_0)^2 + (y-y_0)^2}{2R_{th}^2}\right)} \quad (41)$$

where x_0 , y_0 define the location of the thermal core. A profile is shown in Figure 7(a).

The net climb rate in the thermal is the sum of the glider's sink rate (at its airspeed and bank angle) and the vertical wind speed. Flying in a tight circle close to the thermal core will give the largest vertical wind speed, but sink rate will be higher. Intuition therefore suggests that once a desired turn radius has been selected, one should choose an airspeed and bank angle that minimizes the sink rate for that turn radius. In practice, however, once the airspeed is "close enough" to the optimal value, sink rate changes very little with changes in airspeed for a given turn radius. Figure 8 shows the rate of change of energy for an SB-XC circling in a thermal of radius 70 m with core strength 4 m/s (properties of the SB-XC are given in the Appendix).

Note the strong dependence of \dot{e} on turn radius and comparatively weak dependence on airspeed. Flying at a nominal airspeed of approximately 13 m/s gives good performance, and this remains true as thermal radius varies. The problem now is to determine the thermalling radius that maximizes energy gain in an

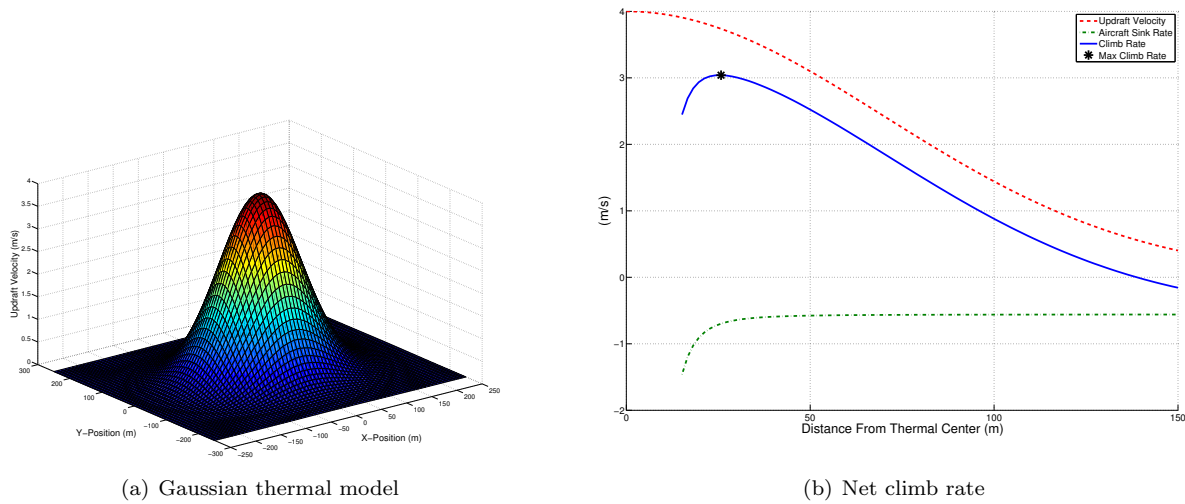


Figure 7. Left: Gaussian thermal model with a $R_{th} = 70$ m and $w_0 = 4$ m/s; Right: Climb rate as a function of distance from the thermal center for SB-XC. The thermal pictured on the left was used. The maximum climb rate for this aircraft/thermal combination is 3.04 m/s and the optimal turn radius is 25.8 m.




Figure 8. Sink rate as a function of turn radius and airspeed for the SB-XC glider.

unknown thermal. Various approaches to mapping the thermal (so that an optimal radius can be computed) have been proposed,^{2,20} but these methods can become computationally intensive.

Extremum seeking control is a non-model based approach where system parameters are continuously perturbed to compute a gradient. This gradient is used to autonomously seek an extremum and track the extremum.²¹ It has been used in many applications including flow control²² and UAV motion planning.²³ This approach has several advantages for thermal soaring:

1. An estimate of the thermal profile is not required.
2. The aircraft's sink polar is not required.
3. The algorithm works for multiple aircraft without the need for additional tuning.
4. The algorithm is simple to implement; computation and memory requirements are minimal.
5. The algorithm adapts to changing thermal conditions.

To maximize energy gain during thermal soaring, the gradient $\frac{\partial \dot{e}}{\partial R}$ is computed by perturbing turn radius around its current nominal value; the nominal turn radius is then changed in the direction of increasing \dot{e} . A flowchart of the adaptive thermalling controller is shown in Figure 9. Using this algorithm the aircraft flies at or near the optimal radius and has the capability to adapt as the thermal profile changes.

The algorithm is triggered when the decision to take the thermal is made. Methods for deciding when to lock on and when to leave the thermal are discussed by Edwards¹⁴ and Allen.¹ When the algorithm is first triggered, the turn radius is initialized to a relatively small value. This value depends on the scale of the aircraft being used. For a full-size glider with a 15 meter wingspan, 80 meters is a viable starting radius. For the small-scale SB-XC glider, an initial turn radius could be as small as 20 meters. This prevents the aircraft from "losing" smaller thermals by reaching a steady-state radius that is larger than the thermal. In addition, it allows centering to occur faster because more loops can be completed in the same amount of time.

Once the thermal is centered at this initial radius, the algorithm computes the average rate of energy gain over a loop (or sequence of loops). When the specified number of loops are completed, the steady-state turn radius is increased for the next sequence of loops. Once again, the average \dot{e} is calculated recursively while flying at this steady-state radius. The current and previous \bar{e} values are compared. If an increase is observed, the turn radius is once again increased. However, if a decrease is observed, the turn radius is perturbed in the opposite direction. The algorithm follows the logic table presented in Table 1 and repeats this process until the thermal dissipates or another task is assigned.

Table 1. Logic table for thermal radius convergence.

Condition	Action
$\bar{e}_k \geq \bar{e}_{k-1} \ \& \ R_{ss,k} > R_{ss,k-1}$	Increase R_{ss}
$\bar{e}_k \geq \bar{e}_{k-1} \ \& \ R_{ss,k} < R_{ss,k-1}$	Decrease R_{ss}
$\bar{e}_k < \bar{e}_{k-1} \ \& \ R_{ss,k} > R_{ss,k-1}$	Decrease R_{ss}
$\bar{e}_k < \bar{e}_{k-1} \ \& \ R_{ss,k} < R_{ss,k-1}$	Increase R_{ss}

While proofs of stability of extremum seeking control for thermal soaring have not yet been derived, performance tests using a high-fidelity soaring simulator shows that the controller indeed converges to the maximum climb rate and that this rate is maintained. Further, stable thermal centering still occurs. The simulator and results are described in the next two sections.

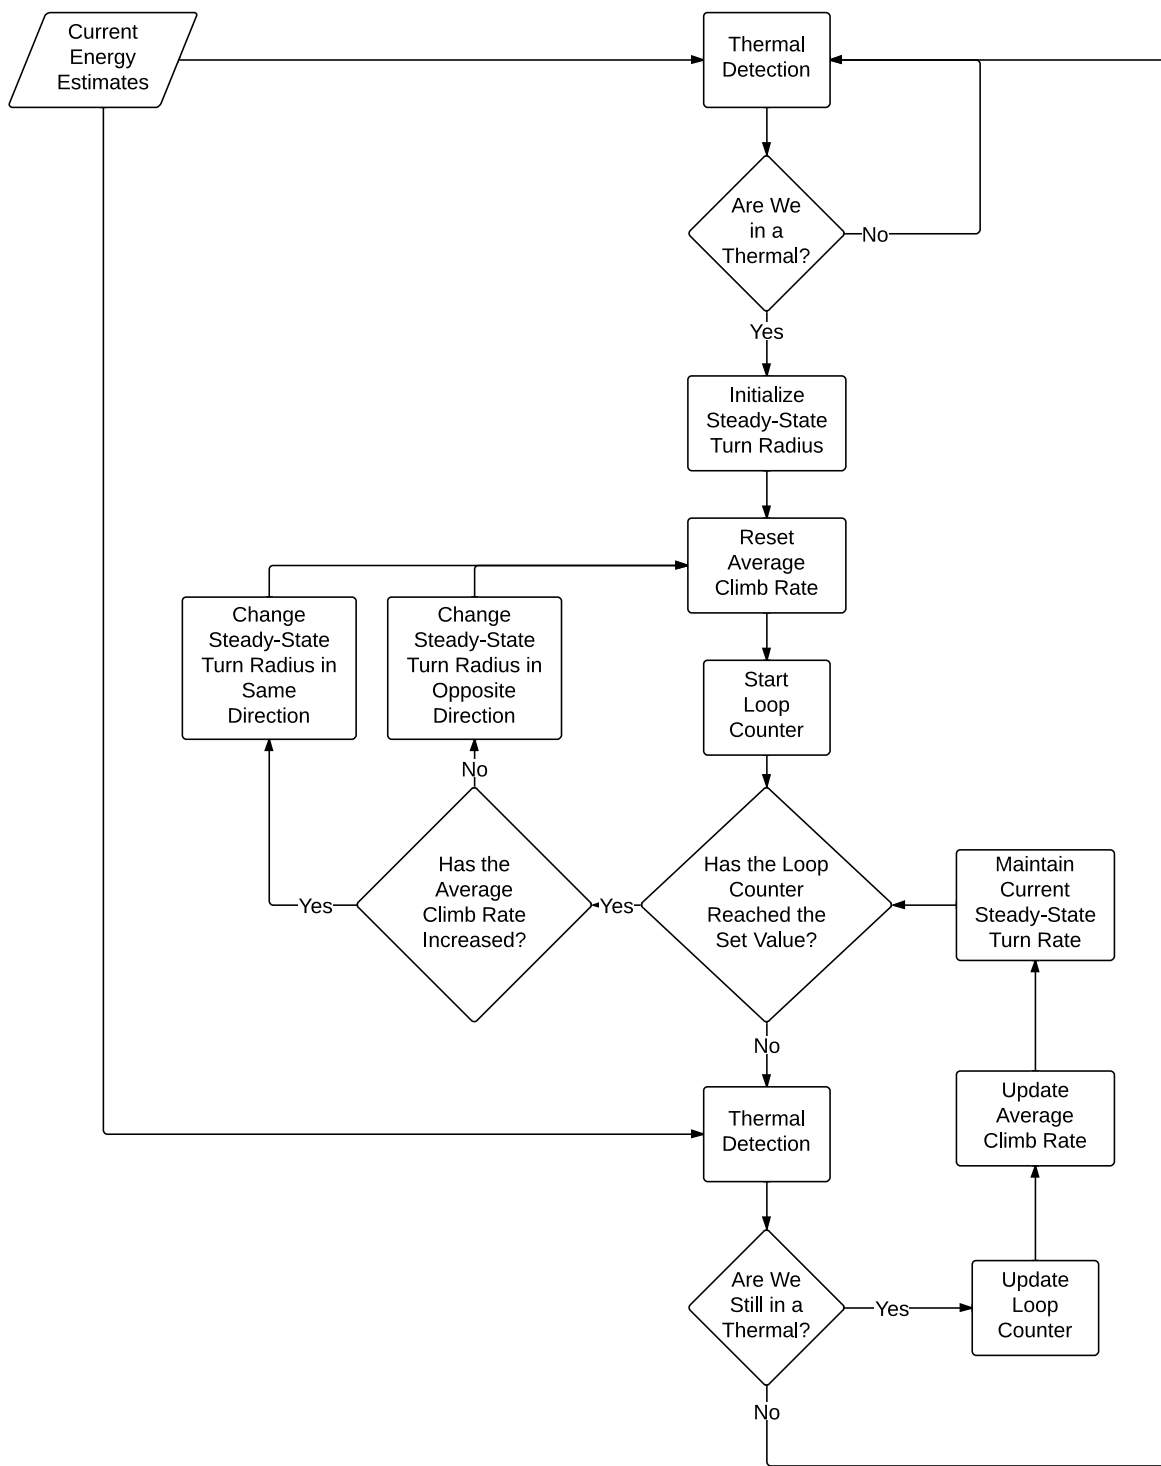


Figure 9. Flowchart showing the structure of the turn radius convergence algorithm.

VI. Simulation

The Silent Wings flight simulator (www.silentwings.no) was used as the primary testing environment for the methods discussed throughout this paper. Silent Wings is a commercially available soaring simulator and multiplayer game that is committed to replicating the soaring experience as accurately as possible. Silent Wings is used by soaring pilots around the world for training purposes, to practice soaring, or to compete against other pilots.

This simulator generates extremely realistic atmospheric conditions, including accurate pressure and temperature distributions and dynamic thermal formations that follow the full thermal cycle from release to decay. Wind characteristics such as average wind speed, wind direction, turbulence levels, and wind shear can be adjusted as well as thermal properties such as average thermal size, average core strength, and the number of thermals present in the current map. The aerodynamics of the preloaded aircraft are accurately modeled by simulating the airflow over each section of the aircraft model. However, the simulator is robust in the sense that it allows you to create and fly add-on aircraft using the native panel model or a linearized aircraft model. Silent Wings also simulates the load factor and various stresses on the aircraft, causing damage to specific aircraft regions if the limits are exceeded. The gliders also use simulated instruments such as a TE-compensated variometer and a GPS allowing measurements to be generated in real-time. These measurements, along with the current state of the aircraft, can be retrieved in real-time. Unlike Condor (www.condorsoaring.com), another comparable soaring simulator, Silent Wings provides the capability to control aircraft remotely through a UDP connection. This enables the simulation of flights where aircraft are controlled autonomously using state feedback. Figure 10 shows a glider thermal soaring under autonomous control in Silent Wings. In addition, Silent Wings provides the ability to join or set up your own multiplayer server where you can soar with other aircraft. This permits the performance of different controllers to be compared by flying in the exact same environmental conditions as well as the testing of autonomous versus manned flight.

Using the states being sent from Silent Wings in real-time, a low-level autopilot system was designed to remotely control the aircraft in simulation. The autopilot is powered by an Arduino Mega 2560 microcontroller and communicates with Silent Wings (at a rate of 30-50 Hz) using an external Ethernet shield. The low-level autopilot software was designed in-house and uses successive loop closure to control the desired longitudinal and lateral/directional aircraft states. It has the ability to control airspeed, pitch attitude, bank angle, heading, turn rate etc. Matlab/Simulink was used as a mission computer to send high-level commands to the autopilot system through a serial connection. Additionally, state information was forwarded from the autopilot to the mission computer at a rate of 1-2 Hz. Figure 11 shows a graphical representation of the testing and simulation network that was used.

Initially, an aerodynamic model of the SB-XC was created for this aircraft using XFLR5 (www.xflr5.com). XFLR5 is a software tool that was created specifically for aerodynamic analysis of aircraft flying at low Reynolds numbers. This program uses XFOIL, a subsonic airfoil analysis tool developed at the Massachusetts Institute of Technology (web.mit.edu/drela/Public/web/xfoil) to estimate the 2-dimensional drag polar, maximum lift coefficient, and other aerodynamic properties. XFLR5 utilizes this 2-D airfoil data to build a 3-D aerodynamic model using lifting line theory, vortex lattice method, and a 3-D panel method. This aerodynamic model was then implemented (as a nonlinear panel model) in Silent Wings soaring simulator. This model was fine tuned using flight test data published by Edwards⁷ and adapted for our aircraft weight. This SB-XC model was used to test the performance of the soaring algorithms in a realistic environment.



Figure 10. Glider autonomously climbing in a thermal. Simulation was performed using an SB-XC model in Silent Wings flight simulator.

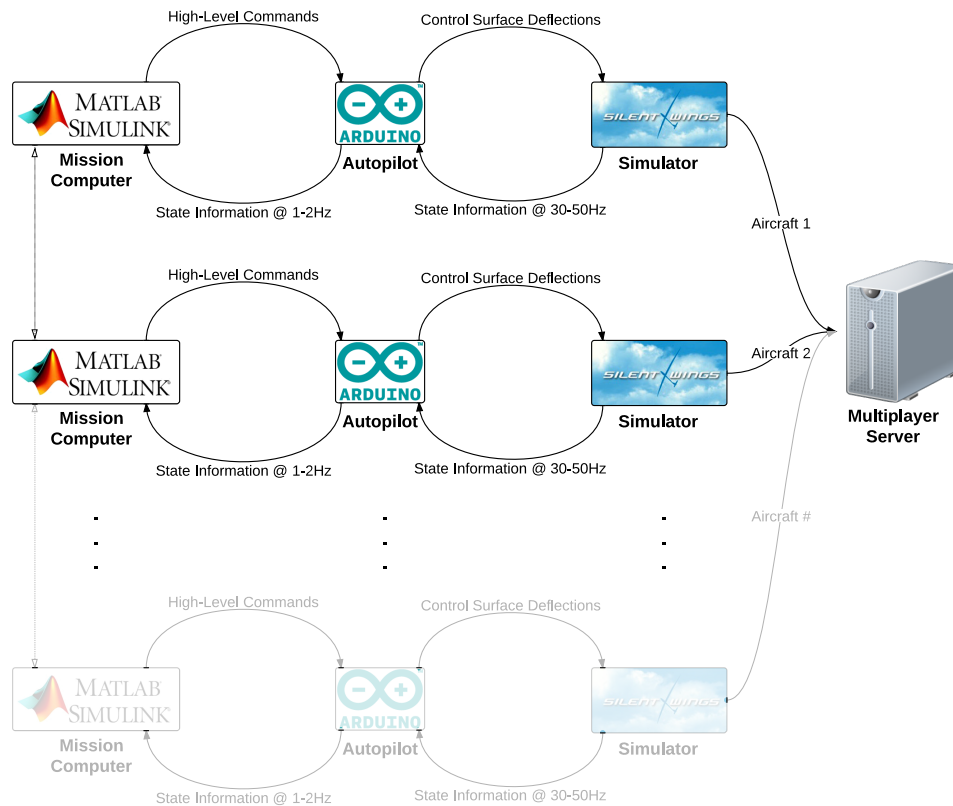


Figure 11. Complete testing and simulation architecture. This system can be used to compare an unbounded number of aircraft and control systems in the same environment by setting up a multiplayer server.

VII. Results

The results of various thermal soaring simulations are reported in this section. The purpose of these simulations is to analyze the proposed improvements and their ability to extend the endurance of small-scale soaring UAVs via thermal soaring. A model of the SB-XC small-scale glider was used exclusively for testing in Silent Wings soaring simulator. The simulations discussed in this section investigate: thermal centering performance based on energy estimation methods; the utility of the extremum seeking controller.

A. Centering Performance

The performance of the thermal centering controller (Equation 15 in Section III) is directly dependent on the quality and timeliness of the \ddot{e} estimates. For this reason, the merit of the state estimator can be determined by examining thermal centering performance. Centering performance includes the aircraft's ability to center a thermal quickly, maintain climb in a large range of thermal sizes and strengths, and reject disturbances such as turbulence and wind shear.

An energy estimator that has been successfully used for thermal soaring with the SB-XC can be seen in Figure 12. This estimator was designed by Allen and Lin¹ and uses a series of filters to reject noise and differentiate the energy signal.

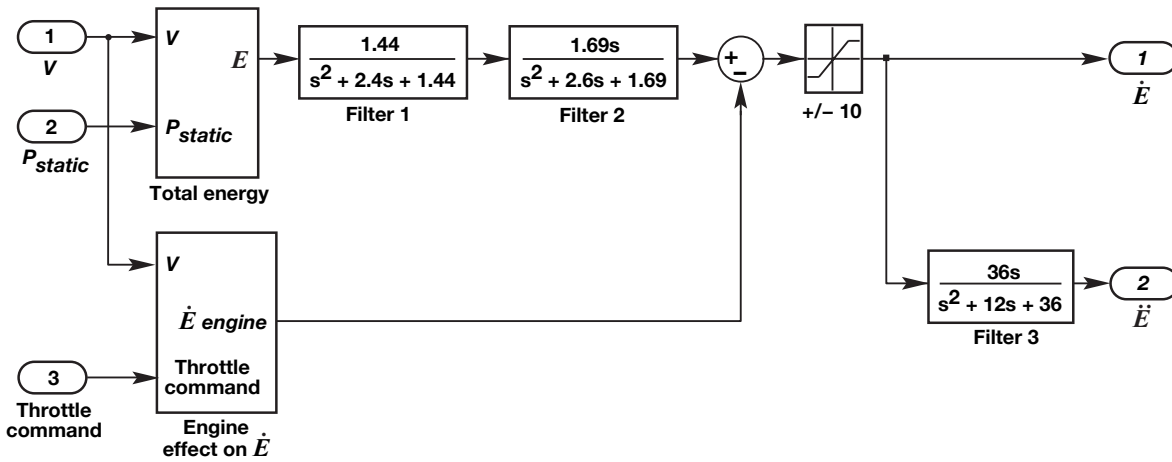


Figure 12. Energy estimator designed by Allen and Lin¹

Andersson et al. discuss the possibility of augmenting the thermal centering controller to accept an \ddot{E} term in an attempt to reduce latency.²⁴ However, this was discussed as potential future work and an analysis has yet to be completed showing the efficacy of this controller. It is known that other energy filters have been successfully used for autonomous thermal centering, unfortunately these filter values remain unpublished. Therefore, for the sake of comparison, the energy estimator shown in Figure 12 was used as a baseline for the following thermal soaring simulations.

The following analysis was performed to examine the estimator’s ability to enable thermal centering in a wide range of thermals.⁹ The thermals were generated by Silent Wings and resemble a Laplacian of Gaussian surface. An example thermal profile can be seen in Figure 13.




Figure 13. Thermal generated by Silent Wings ($R_{\text{therm}} = 80\text{m}$, $w_{\text{core}} = 4\text{m/s}$).

Thermals ranged from a minimum radius of 55 m to a maximum radius of 105 m and core strengths ranging from 2 to 6 m/s. Figure 14 shows the initial behavior of the SB-XC as it encounters these different thermals. The aircraft was flown directly into the core of the thermal; this was to analyze the robustness of the system by starting at the most difficult location for the estimator/controller to center a thermal.




Figure 14. Initial centering response of the SB-XC model to various thermals.

Results for each of these simulations are summarized in Tables 2, 3, and 4.

Table 2. SB-XC model flying in thermal ($R_{th} = 55$ m).

Core Str (m/s)	Filter	Locked On?	Avg ROC (m/s)
2 m/s	BSG	Yes	1.23
	Allen	No	N/A
4 m/s	BSG	Yes	3.27
	Allen	No	N/A
6 m/s	BSG	Yes	5.30
	Allen	No	N/A

Table 3. SB-XC model flying in thermal ($R_{th} = 80$ m).


Core Str (m/s)	Filter	Locked On?	Avg ROC (m/s)
2 m/s	BSG	Yes	1.37
	Allen	Yes	1.35
4 m/s	BSG	Yes	3.36
	Allen	Yes	3.32
6 m/s	BSG	Yes	5.24
	Allen	Yes	5.20

Table 4. SB-XC model flying in thermal ($R_{th} = 105$ m).

Core Str (m/s)	Filter	Locked On?	Avg ROC (m/s)
2 m/s	BSG	Yes	1.33
	Allen	Yes	1.34
4 m/s	BSG	Yes	3.41
	Allen	Yes	3.40
6 m/s	BSG	Yes	5.27
	Allen	Yes	5.24

The results show that both filtering methods were able to steadily climb in the 80 and 105 m thermals. However, only the model using the BSG filter was able to lock on to the smallest (55 m radius) thermals. Figure 14 also shows that the BSG filtering method allowed for much quicker and smoother centering; typically centering the thermal within the first $1\frac{1}{2}$ rotations. The reduced latency in the BSG filter allows accurate centering of small thermals. Note that simulations where the Allen filter was able to lock onto a thermal, the average climb rate was similar to that of the BSG filter (with the BSG filter averaging a climb rate about 0.6% greater).

The next simulation tests the estimator/controller's ability to reject disturbances (which may arise from wind shear, turbulence, and many other factors). This simulation imposes disturbances in the form of thermal drift and change of wind direction. A series of shear layers were defined, which causes the thermal to change its drift direction rapidly. The lateral directions of the disturbances are illustrated in Figure 15. Note that this scenario was designed specifically to be severe, and it is unlikely that one would encounter a sequence of shear layers of this magnitude. The thermal tracking performance was tested using the backward Savitzky-Golay (BSG) filter and Allen's differentiation filters. Each filter was subjected to three different sets of cross-winds of varying strength. These simulations were performed with 1.5, 3, and 5 m/s lateral wind strengths. The thermal used for this simulation maintains a profile similar to the one shown in Figure 13 and an average radius and strength of 80 meters and 4 m/s respectively. The path of the aircraft through the thermals can be seen in Figure 16 and the corresponding altitude plot is shown in Figure 17.

**Figure 15. Cross-wind altitude locations.**




Figure 16. Path of the SB-XC in a thermal subjected to rapidly-changing cross-winds.

Table 5. Results of SB-XC subjected to rapidly-changing cross-winds

Wind	Filter	Finished	Avg ROC (m/s)	Vertical Gain (m)
1.5 m/s	BSG	Yes	3.34	1058
	Allen	No	3.05	872
3.0 m/s	BSG	Yes	3.24	1052
	Allen	No	2.23	895
5.0 m/s	BSG	No	3.19	310
	Allen	No	3.10	261




Figure 17. Altitude of SB-XC in a thermal subjected to various cross-winds.

The proposed BSG energy filter was able to provide better average climb rate and thermal tracking than Allen's energy filter. The BSG estimator allowed the thermal centering controller to maintain steady climb through the entire life of the thermal in both the 1.5 and 3 m/s cases. In these cases, Allen's filter enabled climbing through multiple wind shear layers but both lost tracking around 950 m AGL. The 5 m/s cross-wind proved to be too much for both cases, as they lost tracking almost immediately upon experiencing such large disturbances. The inability of Allen's filter to track the thermals is likely due to excessive estimate lag.

B. Extremum Seeking Control

This section shows the results of performance testing of the extremum seeking controller (discussed in Section V). This high-level controller uses the rate of energy gain gradient ($\frac{\Delta \dot{e}}{\Delta R}$) to choose steady-state turn commands that ultimately converge to the optimal value. Results are compared with constant-radius thermalling controllers; in all cases the backward Savitzky-Golay energy filter and the centering controller presented in Section III were used. Simulations were performed in Silent Wings soaring simulator.

Simulations of flights in thermal sizes ranging from an 80 to 200 m radius were conducted. The performance of the extremum seeking controller was compared to the fixed turn radius values of 25 and 35 m. These fixed values were chosen because they represent a typical thermalling radius that provides near optimal climb rate for the SB-XC. Additionally, the extremum seeking controller was initialized with a turn radius of 35 m; this radius is large enough that phase lag in energy estimates won't cause the thermal to be lost but small enough that it should provide good lift during the initialization phase with almost any usable thermal. Figure 18 shows the altitude, rate of energy gain, and commanded turn radius of the SB-XC while flying in a medium-sized thermal ($w_{core} = 4$ m/s, $R_{th} = 120$ m). Table 6 shows the corresponding average climb rate values.

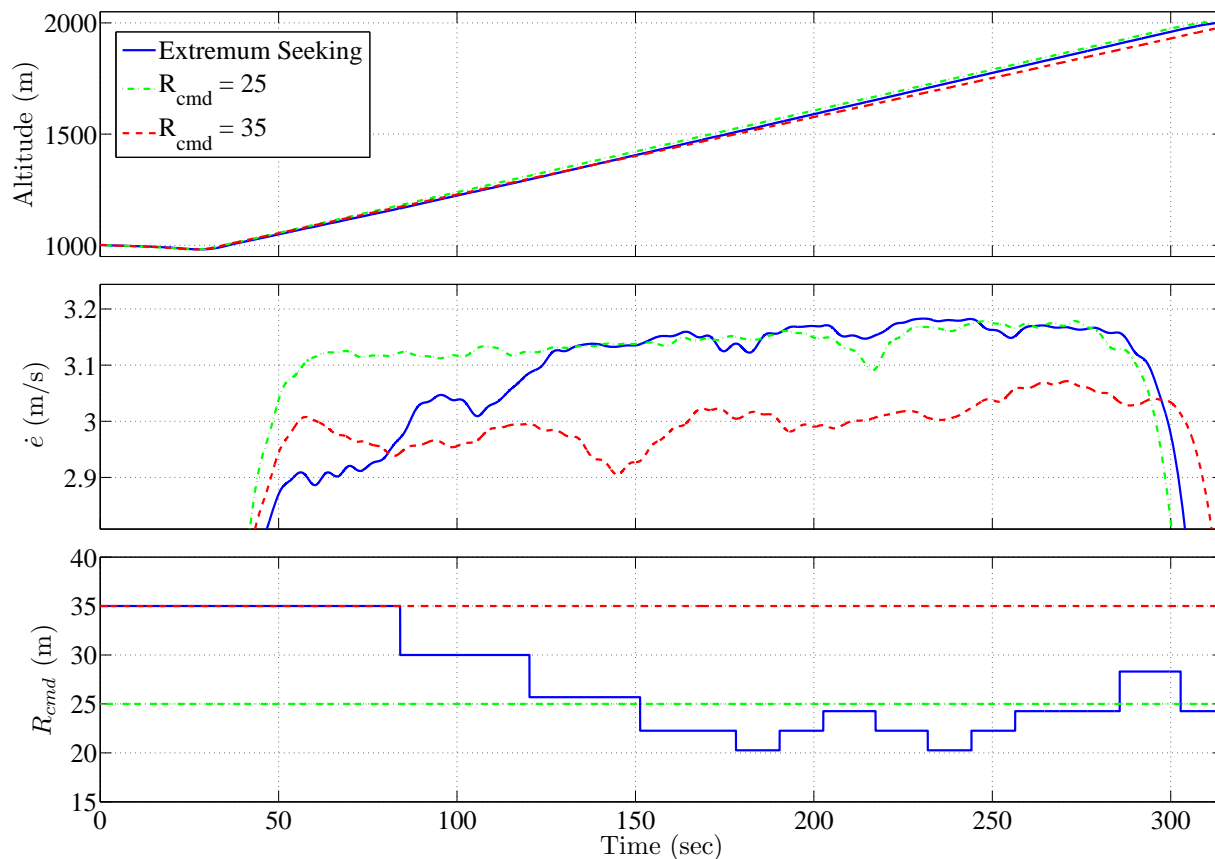


Figure 18. Extremum seeking control in a thermal ($w_{core} = 4$ m/s, $R_{th} = 120$ m).

Table 6. Average climb rates ($w_{core} = 4$ m/s, $R_{th} = 120$ m).

Radius Cmd (m)	Avg ROC (m/s)
Extremum Seeking	3.63
R = 25	3.67
R = 35	3.50

This simulation shows good convergence for the extremum seeking controller. The rate of change of energy starts low (as with the 35 m case) but the controller consistently tightens the turn until it converges to a near-optimal radius that is likely just under 25 m; the exact value is unknown due to unknowns in the simulated thermal environment. This is demonstrated by the fact that \dot{e} slightly surpasses the 25 m case as the commanded radius was tightened below 25 meters. Although the 25 m case had a slightly better average rate-of-climb than the extremum seeking controller, the extremum seeking controller achieved a higher peak energy than both the 25 and 35 m flight paths. Despite the extremum seeking controller starting with an initial turn radius of 35 m, after the ($\sim 1,000$ m) climb it outperformed the 35 m case by almost 4%. Figure 19 and Table 7 show results for flight within a wide thermal with a slightly weaker core ($w_{core} = 3$ m/s, $R_{th} = 200$ m).




Figure 19. Extremum seeking control in a thermal ($w_{core} = 3 \text{ m/s}$, $R_{th} = 200 \text{ m}$).

Table 7. Average climb rates ($w_{core} = 3 \text{ m/s}$, $R_{th} = 200 \text{ m}$).

Radius Cmd (m)	Avg ROC (m/s)
Extremum Seeking	2.72
R = 25	2.65
R = 35	2.68

This simulation also shows good convergence of the controller as it converged rapidly to a value around 28 m and maintained a higher climb rate throughout the duration of climb. Additionally, the extremum seeking controller provided a more uniform \dot{e} than the other controllers. Performance gains (in the form of climb rate) over the 25 and 35 m cases were 2.5% and 1.5% respectively with the 35 m case outperforming the 25 m case (as opposed to the previous simulation). The next set of simulations exposed the controller to a narrow but strong thermal ($w_{core} = 5 \text{ m/s}$, $R_{th} = 80 \text{ m}$). These results are shown in Figure 20 and Table 8.

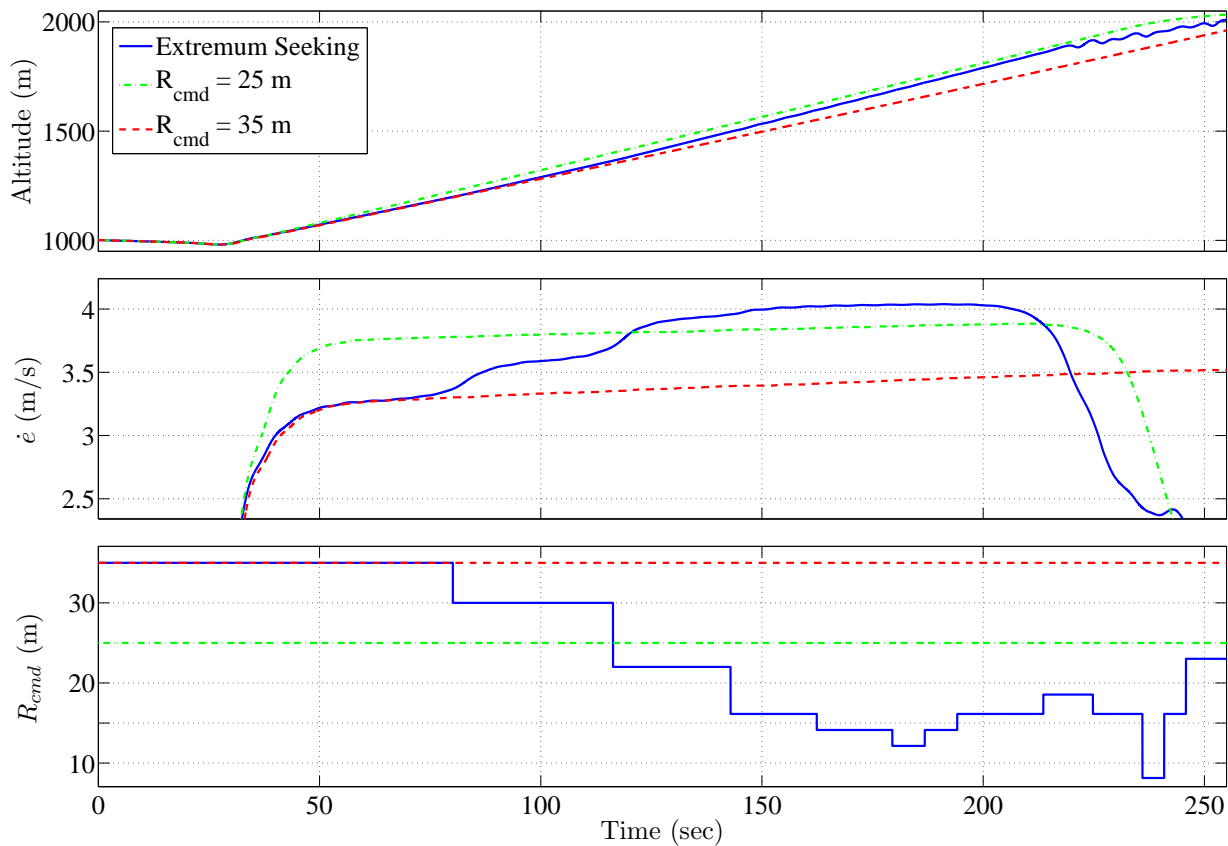


Figure 20. Extremum seeking control in a thermal ($w_{core} = 5$ m/s, $R_{th} = 80$ m).

Table 8. Average climb rates ($w_{core} = 5$ m/s, $R_{th} = 80$ m).

Radius Cmd (m)	Avg ROC (m/s)
Extremum Seeking	4.73
R = 25	4.88
R = 35	4.33

This simulation also shows rapid convergence of the turn radius. Additionally, the maximum \dot{e} for the extremum seeking controller is noticeably higher than the other two cases. However, since the thermal is narrow with a strong core, the optimal turn radius is very small. Despite hard limits on the bank angle command, the tight turn commands combined with disturbances in the thermal cause the controller to become slightly unstable. This instability is likely caused by: attempting to fly near stall conditions in a turbulent environment; pushing the limits of the phase margin of the centering controller; and frequent changes in the commanded turn rate (since this frequency is dependent on loop size). Because all of these factors are a result of attempting to fly an extremely tight turn, this problem could be corrected by performing a stability analysis and adjusting the limits (and gains) accordingly. Although this stability analysis has not yet been completed, it is a good candidate for future work. Despite the instability towards the end of the simulation, the controller provided an average climb rate of 4.73 m/s. This showed 9% improvement over the 35 m case but was 3% slower on average than the 25 m case.

In all three of the previous scenarios, although convergence was observed from an initial turn radius of 35 m it appears that using 25 m as an initial radius for the extremum seeking controller would have provided better initial climb rates. For this reason, the first case was flown using a 25 m initial radius. This is shown in Figure 21 and Table 9.




Figure 21. Extremum seeking control in a thermal ($w_{\text{core}} = 4 \text{ m/s}$, $R_{\text{th}} = 120 \text{ m}$).

Table 9. Average climb rates ($w_{\text{core}} = 4 \text{ m/s}$, $R_{\text{th}} = 120 \text{ m}$).

Radius Cmd (m)	Avg ROC (m/s)
Extremum Seeking ($R_0 = 35$)	3.63
Extremum Seeking ($R_0 = 25$)	3.76
$R = 25$	3.67
$R = 35$	3.50

It can be seen that the extremum seeking controller is able to significantly outperform all of the other cases. Starting at the same steady-state turn radius of 25 m, the extremum seeking controller was able to improve the climb rate by almost 4% over the case where the radius remained constant. This result solidifies the utility of the extremum seeking controller via improved climb rate.

VIII. Conclusion

This paper has introduced a method for optimizing climb rate during thermal soaring through extremum seeking control. This is a model-free approach (neither a model of the thermal nor a model of the aircraft is required) that can adapt to changing thermal conditions in real-time.

Thermal centering is aided using estimates of total energy, rate of change of total energy and second derivative of total energy using an asymmetric (backward) Savitzky-Golay filter. This filter computes a polynomial fit to the measurements of total energy over a moving time window, and then computes derivatives of the estimated total energy from the computed polynomial. This approach significantly reduces the phase lag associated with other filtering approaches, improving performance of the thermal centering algorithm.

A simulation environment based on a commercially available soaring simulator was briefly presented. This environment allows hardware in the loop simulations: here low-level flight controls were implemented on an Arduino Mega single board computer communicating with the simulator over UDP; higher level control was implemented on a laptop communicating with the autopilot over a serial link.

Successful simulations were conducted using the small-scale SB-XC glider to autonomously track and improve climb rate in thermals. Stable thermal centering, in addition to improvements caused by the extremum seeking control algorithm, were observed from these simulations once again demonstrating the feasibility of using small-scale autonomous sailplanes to complete long endurance missions.

References

- ¹Allen, M. J. and Lin, V., "Guidance and Control of an Autonomous Soaring Vehicle with Flight Test Results," *AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada, January 2007.
- ²Edwards, D. J., "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," *AIAA Guidance, Navigation and Control Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia, August 2008.
- ³Andersson, K., Kammer, I., and Jones, K. D., "Autonomous Soaring: Flight Test Results of a Thermal Centering Controller," *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2010-8034, American Institute of Aeronautics and Astronautics, Toronto, Canada, August 2010.
- ⁴Arho, R., "Optimal Dolphin Soaring as a Variational Problem," *OSTIV Publication XIII*, Organisation Scientifique et Technique Internationale du Vol à Voile, 1974.
- ⁵de Jong, J. L., "The Convex Combination Approach: A Geometric Approach to the Optimization of Sailplane Trajectories," *OSTIV Publication XVI*, Organisation Scientifique et Technique Internationale du Vol à Voile, 1981, pp. 182–201.
- ⁶Cochrane, J. H., "MacCready Theory with Uncertain Lift and Limited Altitude," *Technical Soaring*, Vol. 23, No. 3, July 1999, pp. 88–96.
- ⁷MacCready Jr., P. B., "Optimum Airspeed Selector," *Soaring*, January-February 1958, pp. 10–11.
- ⁸Sander, G. and Litt, F. X., "On Global Optimal Sailplane Flight Strategy," *NASA Conference Publication 2085, Science and Technology of Low Speed and Motorless Flight*, edited by P. W. Hanson, NASA, Hampton, Virginia, March 1979, pp. 355–376.
- ⁹Daugherty, S., *Improving Autonomous Soaring via Energy State Estimation and Extremum Seeking Control*, Master's thesis, The Pennsylvania State University, December 2013.
- ¹⁰Wharington, J., *Autonomous Control of Soaring Aircraft by Reinforcement Learning*, Ph.D. thesis, Royal Melbourne Institute of Technology, Melbourne, Australia, November 1998.
- ¹¹Wharington, J. and Herszberg, I., "Optimal Semi-Dynamic Soaring," Tech. rep., Royal Melbourne Institute of Technology, October 1998.
- ¹²Allen, M. J., "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," *43rd AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada, January 2005.
- ¹³Reichmann, H., *Cross-Country Soaring*, Thomson Publications, Santa Monica, California, 1978.
- ¹⁴Edwards, D. J. and Silverberg, L. M., "Autonomous Soaring: The Montaque Cross-Country Challenge," *Journal of Aircraft*, Vol. 47, No. 5, September-October 2010, pp. 1763–1769.
- ¹⁵Andersson, K., Kammer, I., Jones, K. D., Dobrokhodov, V., and Lee, D.-J., "Cooperative UAVs Using Thermal Lift to Extend Endurance," *AIAA Infotech@Aerospace Conference*, AIAA Paper 2009-2043, American Institute of Aeronautics and Astronautics, Seattle, Washington, April 6-9 2009.
- ¹⁶Andersson, K. and Kammer, I., "On Stability of a Thermal Centering Controller," *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2009-6114, American Institute of Aeronautics and Astronautics, Chicago, Illinois, August 2009.
- ¹⁷Chakrabarty, A. and Langelaan, J. W., "Energy-Based Long-Range Path Planning for Soaring-Capable Unmanned Aerial Vehicles," *Journal of Guidance, Control, and Dynamics*, Vol. 34, July-August 2011.
- ¹⁸Savitzky, A. and Golay, M. J. E., "Smoothing and Differentiation of Data by Simplified Least Squares Procedures," *Analytical Chemistry*, Vol. 36, 1964, pp. 1627–1639.
- ¹⁹Simon, D. J., *Optimal State Estimation*, John Wiley & Sons, Inc., Hoboken, New Jersey, 1st ed., 2006.
- ²⁰Bird, J. and Langelaan, J., "Spline Mapping to Maximize Energy Exploitation of Non-Uniform Thermals," *Proceedings of the XXXI OSTIV Congress*, Organization Scientifique et Technique du Vol a Voile, August 2012.
- ²¹Ariyur, K. B. and Krstic, M., *Real-Time Optimization by Extremum Seeking Control*, John Wiley & Sons, Inc., Hoboken, New Jersey, 2003.
- ²²Becker, R., King, R., Petz, R., and Nitsche, W., "Adaptive Closed-loop Separation Control on a High-Lift Configuration using Extremum Seeking," *AIAA Journal*, Vol. 45, No. 6, 2007, pp. 1382–1392.
- ²³Dixon, C. and Frew, E., "Controlling the Mobility of Network Nodes using Decentralized Extremum Seeking," *Proceedings of the 45th IEEE Conference on Decision and Control*, Institute of Electrical and Electronics Engineers, San Diego, California, December 13-15, 2006 2006.
- ²⁴Andersson, K., Kammer, I., Dobrokhodov, V., and Cichella, V., "Thermal Centering Control for Autonomous Soaring: Stability Analysis and Flight Test Results," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 3, May-June 2012, pp. 963–974.

Appendix: Vehicle Properties

Note that a third order polynomial is used to relate C_D to C_L : this provided a better fit to the computed data over the full speed range.

Table 10. Parameters for SB-XC glider.

variable	value	description
m	5.7 kg	mass
S	0.996 m ²	wing area
$f(C_L)$	$0.0166C_L^3 + 0.0535C_L^2 - 0.0437C_L + 0.0276$	
$v_{a,min}$	10 m/s	
$v_{a,max}$	35 m/s	
η_p	0.80	efficiency of the propeller
η_m	0.90	efficiency of motor
η_{esc}	0.95	efficiency of speed controller
a,b,c	-0.0162, 0.3782, -2.7018	sink rate polar fit in $\frac{m/s}{v_a}$