The Pennsylvania State University
The Graduate School

IMPROVING AUTONOMOUS SOARING VIA ENERGY STATE ESTIMATION AND EXTREMUM SEEKING CONTROL

A Thesis in
Aerospace Engineering
by
Shawn C. Daugherty

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

December 2013

The thesis of Shawn C. Daugherty was reviewed and approved* by the following:

Jacob W. Langelaan
Associate Professor of Aerospace Engineering
Thesis Advisor

Joseph F. Horn
Associate Professor of Aerospace Engineering

George A. Lesieutre
Professor of Aerospace Engineering
Head of the Department of Aerospace Engineering

*Signatures are on file in the Graduate School.

Abstract

This research is motivated by the significant potential of soaring UAVs to efficiently accomplish both civil and scientific missions by atmospheric energy harvesting. Viable missions include surveillance, communication relay, and environmental monitoring.

This thesis seeks to improve the utility of small, autonomously controlled gliders by extending the range and endurance of these vehicles. This is accomplished through the exploitation of energy from columns of warm, rising air known as thermals. Thermals occur naturally and are utilized by large birds and sailplane pilots to soar for several hours and cover distances of hundreds of kilometers without any source of propulsion. This thesis analyzes limitations imposed by current algorithms and describes solutions in the form of improved energy estimation methods and turn optimization through extremum seeking control.

The thermal centering algorithm, based on Reichmann's method, uses the second derivative of total energy as a feedback term to remain centered around the thermal core. Due to the controller's susceptibility to latency, conventional filtering methods greatly restrict the capabilities of the centering controller. This thesis discusses an alternative estimation method, 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. Significant improvements were observed including: the ability to track a larger range of thermals, rapid thermal centering, and improved disturbance rejection.

The problem of optimal thermal soaring was also addressed. Assuming a Gaussian updraft distribution, any given thermal has an optimal flight radius that can be computed for the specific aircraft. However, determining this flight path has proven to be a difficult problem that has not been adequately addressed. A solution is proposed; climb rate maximization using extremum seeking control with

turn radius as the varying parameter. Simulations demonstrated steady turn rate convergence while driving the climb rate to the optimal value.

A simulation environment based on a commercially available soaring simulator is described, with a low level aircraft controller implemented on an Arduino Mega 2560 single board computer. This environment was used for testing and validation of the aforementioned methods. The utility of the energy estimator and extremum seeking controller is demonstrated in this high fidelity simulation.

Table of Contents

List of Figures

List of Tables

Acknowledgments

Dedication

To my loving family.

Chapter 1

Introduction

This thesis seeks to improve the range and endurance of autonomous soaring vehicles by advancing the way thermal updraft energy is estimated and utilized. This research was motivated by the significant potential that small autonomous soaring aircraft have to efficiently perform various civil and scientific missions. The relatively small payload capacity and lower aerodynamic efficiency of smaller aircraft require efficient use of sensor data and intelligent control methods to be used reliably for these missions. Therefore, by improving the efficiency of atmospheric energy harvesting in thermal soaring conditions, the utility of these vehicles are substantially augmented.

Many different soaring methods currently exist that allow the extraction of energy from the atmosphere. Of these soaring types, thermal soaring is an extremely effective way to gain a significant amount of energy in the form of potential energy. For this reason, this thesis specifically focuses on the problem of autonomous thermal soaring. This problem can be broken down into: accurately centering a thermal, maintaining flight within the thermal, and flying the path that optimizes the rate of energy gain from the thermal.

This problem is exceptionally challenging because atmospheric energy conditions are difficult to estimate due to extensive and unpredictable disturbances. The first and second derivatives of energy (an already noisy signal) are required for thermal centering control. Differentiating this signal amplifies noise in a system that is susceptible to latency. Additionally, thermals tend to drift with the direction of the wind making trajectory planning for optimal soaring a difficult task.

This thesis: (a) describes an energy estimation method that increases thermal centering performance; (b) introduces an adaptive method for optimizing the rate of energy gain in a thermal; and (c) compares the performance of the proposed techniques with current algorithms quantifying the results in a high fidelity simulation.

1.1 Motivation

Small uninhabited 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). Therefore, flight control methods that enhance range and endurance can greatly improve the overall utility of these vehicles.

Large birds such as hawks, eagles, and vultures as well as human sailplane pilots exploit atmospheric energy 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.

(a) Red-tailed hawk. (b) JS-1 Sailplane. (c) RQ-11 Raven UAV.

Figure 1.1. Applications where wind energy harvesting is used to remain aloft [1–3].

Several methods of extracting energy from the wind have been observed and used by glider pilots. These methods can be broken down into three main categories: updraft soaring, gradient soaring, and gust soaring. Soaring capable updrafts are caused by three different phenomenon: uneven solar heating of the ground (thermal soaring); deflection of horizontal winds by a slope (ridge soaring); and long-period oscillations in the atmosphere (wave soaring). Thermal soaring is an extremely common technique used by glider pilots, relying on the formation of rising air masses to gain potential energy. Ridge soaring has shown to be an effective way of extending the range and endurance of flight along ridge lines [5]. Additionally, wave lift has been used for high-altitude soaring and enabled a world record soaring altitude of 50,699 feet [11].

Energy can also be gained using wind gradients and the rapid conversion between potential and kinetic energy; this is known as dynamic soaring. This type of soaring accumulates energy in the form of airspeed and is therefore limited by aerodynamic and structural properties of the sailplane.

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 [10, 12, 13]. Soaring by manned gliders has also been the subject of a large amount of research [14–18].

1.2 Previous and Related Work

Autonomous soaring UAVs were proposed by John Wharington in 1998, with techniques such as reinforcement learning used for flight control [19, 20]. In 2005, Michael Allen published an analysis that showed that the endurance of small UAVs could be significantly increased using thermal updrafts [21]. In 2007, Allen proposed and demonstrated a soaring controller that is based on Reichmann's method of thermal soaring [5, 10]. 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 on the SB-XC, a small-scale sailplane from RnR Products. 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 [10].

In 2008, Daniel Edwards built upon Allen's research by improving the estimate of thermal location [12]. Using the centroid-based position acquired from Allen's method, a grid was created and an assumed curve was fit to each node. The node with the best fit was selected as an improved estimate of the thermal center location. Edwards demonstrated successful autonomous thermal soaring using this method and placed third in the Montague Cross-Country Challenge using a 5 kg SB-XC glider named ALOFT (autonomous locator of thermals) [22].

In 2009, Andersson et al. performed a probability analysis demonstrating the benefits of using a cooperating team of UAVs to increase the probability of finding lift [6]. Additionally, it was shown that thermal centering can be achieved without performing the computationally and memory expensive thermal drift calculations. The stability of this simplified thermal centering controller was also demonstrated [23]. Test results showed that the energy filter contributes a significant amount of latency in the feedback; enough to degrade thermal centering performance (especially in smaller, weaker thermals) [13]. In 2012, Andersson et al. published flight tests showing soaring simulations using the this centering controller. It was mentioned that an altered centering controller was later used to 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.

In summary, a significant amount of research and flight testing of autonomous soaring UAVs has taken place recently. Successful thermal soaring has been demonstrated in both simulations and real flight proving both the utility of these vehicles and the potential for improvement.

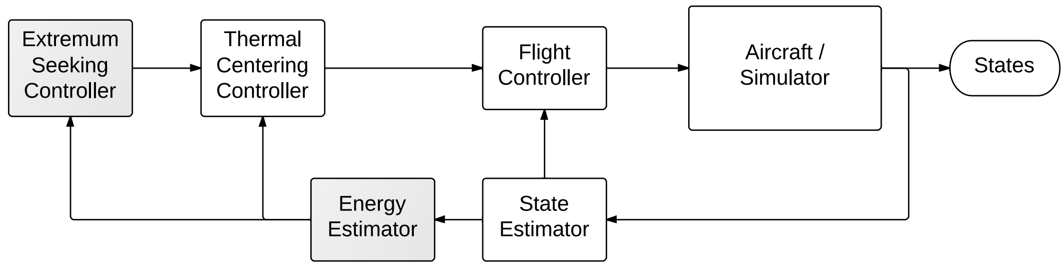

Figure 1.2. Overview of thermal soaring system.

1.3 System Overview

A system to enable optimal thermal soaring is shown in Figure 1.2. The (highlighted) "Energy Estimator" and "Extremum Seeking Controller" are the contributions that this thesis adds to previous autonomous thermal soaring research. The state estimator uses the available sensing package to provide estimates of the vehicle's current state to the flight controller. Information is fed into the estimator from an array of sensors. These sensors include: a pitot-static probe, GPS, and an inertial measurement unit (IMU) including accelerometers, rate gyroscopes, and a magnetometer. These sensors provide raw data to the estimator in the form of: airspeed, altitude, GPS position, orientation, and body-axis accelerations and rotation rates. The state estimator provides an enhanced estimate of the states using sensor fusion and noise filtering methods.

The flight computer implements low-level control laws and commands the deflection of the vehicle's control surfaces. It uses an estimate of the vehicle's state and various reference inputs which can come from a mission computer, a high-level controller, or a human interface device. In this case, reference inputs are received from the thermal centering controller. The flight computer uses a common approach: state feedback enclosed within successive control loops. The details of this controller are further discussed in Chapter 4.

For the thermal soaring application, the need for a thermal centering controller becomes relevant. The thermal centering controller establishes a desired turn rate which maintains flight within the rising airmass. The controller accomplishes this using an estimate of the vehicle's energy states. The energy state estimator provides an estimate of the energy states using vehicle state information such as airspeed and altitude (as discussed in Chapter 3). This thesis introduces an enhanced energy estimation routine as well as an original extremum seeking controller providing more efficient atmospheric energy harvesting through turn rate convergence.

The extremum seeking controller (introduced in Chapter 2) improves the climb rate by adapting the steady-state turn rate (and consequently the turn radius). Past and present energy information is used to seek out the optimal climb rate which is both aircraft and thermal dependent. The aircraft then flies at this optimal steady-state turn rate using the thermal centering and low-level controllers.

These basic functions work together to provide the capabilities necessary for the autonomous soaring flight discussed in this thesis.

1.4 Problem Description

Upon discovery of a thermal, a glider can center on the region of highest lift by maintaining a constant turn around the updraft core. While circling within the rising air mass, energy is extracted in the form of potential energy. However, when a glider performs a banked turn, the lift vector is reoriented from the vertical direction causing the aircraft to fly at a higher lift coefficient to remain aloft. Increasing the lift coefficient moves the aircraft to a less desirable part of the sink polar and causes the sink rate (and the rate of energy loss) to increase. Therefore, determining the circling radius that produces the maximum climb rate is a delicate balance between updraft velocity and aircraft sink rate. In addition, this optimal circling radius is dependent upon the characteristics of the thermal as well as the performance of the individual aircraft that is being used. This implies that in order to maximize the energy gain, the glider must remain centered on the thermal and the energy states of the vehicle must be observed. Summarizing, the thermal soaring problem can be broken down into three basic problems, energy estimation, thermal centering, and climb rate optimization.

1.5 Contributions

Primary contributions of this thesis are as follows:

- **Energy estimator design**

  Increases the range of thermals that can be used for energy gain as well as improved thermal tracking due to reduced latency in energy estimates.

- **Method for optimizing the rate of energy gain**

  Adapts the size of the turn within a thermal to continuously track the radius that provides the best climb rate for a given aircraft.

- **Comparison against current thermal soaring algorithms**

  Simulations were performed verifying the performance of the proposed estimator and controller. Performance gains were quantified by simulating current autonomous soaring methods using our specific aircraft configuration.

1.6 Reader's Guide

The remainder of this thesis is organized as follows. Chapter 2 presents information about the formation and structure of thermals. Additionally, it proposes a control method that actively seeks the steady-state turn radius that provides maximum rate of climb in a thermal. Chapter 3 describes the use of Savitzky-Golay filters for estimating total energy and rates of change of total energy, of which are used for thermal centering. Chapter 4 discusses the design of the low-level autopilot and the various control loops that are used to track high-level reference commands. Additionally, simulation results (obtained using Silent Wings flight simulator) are presented here. Chapter 5 concludes this research by summarizing the contributions of this thesis and discussing recommendations of future work.

Chapter 2

The Thermal Soaring Problem

This chapter defines the soaring problem that was introduced in Chapter 1. It has three purposes: (a) describe thermal formation and identification; (b) outline thermalling techniques used by human glider pilots and how these are applicable to the problem of autonomous soaring; (c) describe the flight trajectory that provides maximum climb rate within a thermal. Section 2.1 describes the lifecycle of various thermal models and indicators that a thermal is present. Section 2.2 describes thermal centering techniques including the one that was chosen for this autonomous soaring application. Section 2.3 presents the kinematic equations for a glider in a steady turn and the derivation of vehicle energy states. Additionally, this section uses these equations to determine the optimal turn conditions required to maximize the climb rate in a thermal. Section 2.4 presents a method for actively tracking optimal turn conditions without the use of thermal mapping techniques.

2.1 Formation and Identification

Thermals are parcels of warm, rising air caused by the uneven heating of the ground and consequently the convective heating of the air directly above these warm areas. Local differences in the conductivity of the Earth's surface and various other factors cause atmospheric instability near the ground. Warm pockets are formed and begin their upward ascent after some triggering event, such as a horizontal wind pushing the parcel of air over a ridge. Because air inside the thermal is warmer than the surrounding air, it continues to rise until it cools and becomes neutrally buoyant.

Thermals are very complex and random formations. However, glider pilots typically conceptualize the structure of thermals using one of two different models, the column model or the bubble model. In the column model, the structure remains attached to the ground for a longer period of time and the column of vertical wind widens (and slows) as altitude increases due to the changing air pressure. The bubble model resembles an air bubble rising to the surface of a boiling pot of water. In the bubble model, the plume detaches from the ground much sooner and upward wind velocity is concentrated inside the bubble. Typically thermals are viewed as circular in cross section, with the vertical component of wind increasing smoothly towards the center of the thermal. Thermal airmasses also tend to drift laterally as they rise due to prevailing or other horizontal winds.

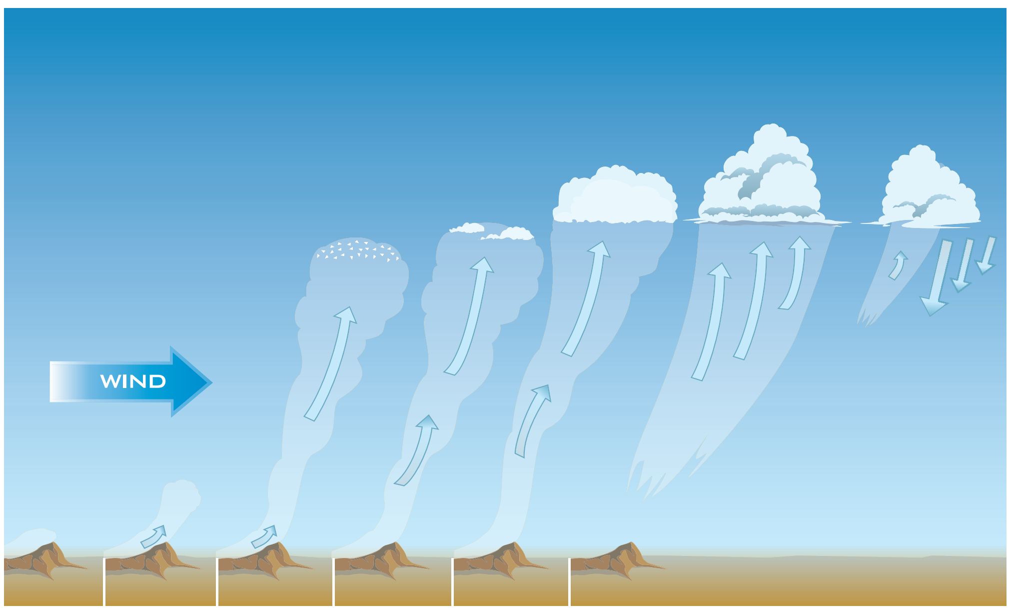

Figure 2.1. Formation of a cumulus cloud from a moisture-rich thermal [4].

When thermals contain significant levels of moisture, they have the potential to form cumulus clouds as shown in Figure 2.1. In fact, newly forming cumulus clouds are oftentimes used by glider pilots to locate thermals. However, the presence of cumulus clouds does not guarantee the existance of thermals. Unfortunately, locating thermals is a difficult task due to the minimal number of visual indicators that exist. This is especially difficult for autonomous systems, where high-level decisions must be made on-the-fly using only the available sensors. Glider pilots have been known to locate thermals using the following indicators:

- cumulus cloud formation
- birds soaring in vertical lift
- other gliders circling in lift
- dark surface regions surrounded by brighter regions

If none of these indicators exist, pilots may simply glide until a thermal is randomly found. Additionally, "house" thermals (thermals that repeatedly occur in the same location) sometimes exist and can be utilized by pilots that have knowledge of their existence. A thermal's existence is made known by a sudden vertical acceleration felt by the pilot or sensed using a variometer. Once a thermal is found, the glider can fly a circular path inside this area of lift to climb to higher altitudes [4, 5].

2.2 Centering Techniques

Thermal soaring, known to some glider pilots as "static soaring", is a commonly used method of extending flight time and endurance. By flying in regions with an upward component of wind, potential energy can be harvested for later use. Because these thermal regions are relatively small, the aircraft must continuously change course to remain in the area that provides a positive climb rate. Quickly and accurately centering the core of the thermal enables improved energy extraction. Glider pilots have become relatively efficient at this circling task and have developed many different techniques to quickly center the core of the thermal. All of the subsequent techniques rely on the assumption that thermals follow one of the conceptual models discussed in Section 2.1. Specifically, the assumption is made that the thermal is strongest at the core and has a relatively circular cross section.

2.2.1 Piggott Technique

Piggott's thermal centering technique relies on the instantaneous vertical acceleration felt by the pilot to originally place the aircraft in a bank around the core [24].

This technique is as follows; while flying a gentle turn and a surge of lift is encountered, immediately steepen your turn. This will move the circle towards the place where the lift was encountered. Piggott's technique requires immediate knowledge to work therefore instrumentation such as a variometer is useless. Pilots that use this technique typically fly "by the seat of their pants," meaning they use the rapid vertical acceleration that is felt by the seat accelerating them upwards and not using instrumentation. However, this method relies on the circumstance that winds at the core are encountered therefore continuous centering corrections are not feasible.

2.2.2 Reichmann Technique

Reichmann's method attempts to follow a steady-state turn radius within the thermal by using changes in climb rate to continuously re-center the core of the thermal [5]. This method consists of the following three rules:

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

However, positive climb rate does not guarantee that the glider is flying in an upward component of wind. By pulling back on the stick, a glider can pull up and climb temporarily by trading kinematic energy for potential energy. This is known as a "stick thermal." Ultimately, to center a thermal using this method we want to use changes in the compensated rate of climb or changes in the rate of total energy gain. This is illustrated in Figure 2.2.

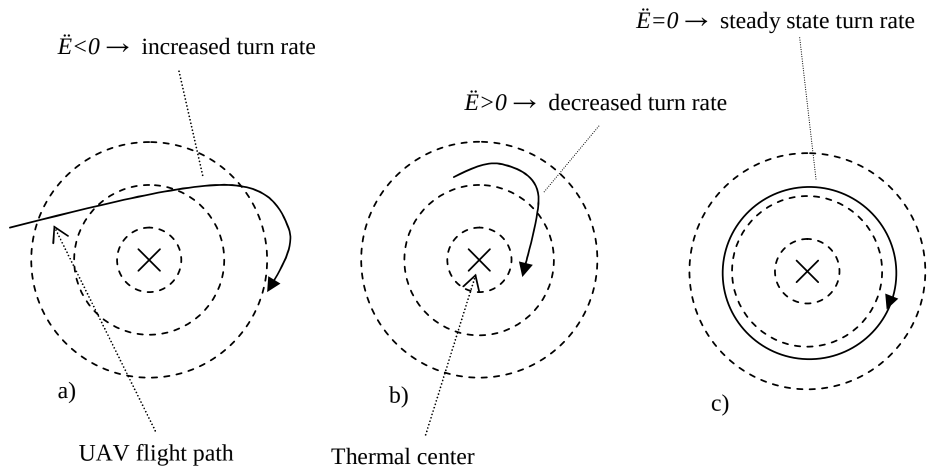

Figure 2.2. Reichmann's method of centering a thermal updraft using total energy acceleration. [5, 6]

This method or slight variations of this method have been used to successfully center thermals using autonomous gliders in the past [10, 12, 13, 22, 25]. For this reason, the autonomous centering controller used in this research is derived from Reichmann's method and is discussed in greater detail in Section 2.2.4.

2.2.3 Mapping

Glider pilots typically rely on TE-compensated variometers to determine the velocity of the upward component of wind as they circle inside a thermal. However, these instruments have extensive lag (on the order of a few seconds) associated with them. The human brain is able to easily correct for this lag and build a mental map of the thermal. Using this mental map, pilots are able to center a thermal by continuously correcting their flight path by adjusting bank angle.

The autonomous application of this approach requires sensor information to construct and update the map. *Spline Mapping to Maximize Energy Exploitation of Non-Uniform Thermals* by Bird and Langelaan showed that a thermal map could be constructed and exploited to maintain climb in a thermal [26]. A Kalman filtering approach was used to estimate the updraft profile from which the optimal trajectory could be calculated. However, this research is based on the assumption that the aircraft's orientation in the wind axis (α, β) is observable. With a GPS sensor providing only a ground relative position, obtaining an accurate estimate of thermal drift due to prevailing winds and localized gusts has proven to be a difficult problem. Additional sensing equipment (such as wind vanes) as well as supplementary computing power are required for this task. Provided the movement of the airmass is observable, the time required to initialize an accurate map hinders the climb rate of the vehicle during this phase. This is apparent in the results presented by Bird and Langelaan as the mapping method outperformed the Reichmann method only when the thermal structure was complex and the formation remained intact for an extended period of time.

2.2.4 Autonomous Centering

Accurately centering a thermal is important because it prevents the thermal from being lost as well as improves the average climb rate of the glider. Under autonomous control, Reichmann's method (Section 2.2.2) has proven to be an effective thermal centering method. In 2007, Allen codified a thermal centering controller that was based on this method and has since been adapted by Edwards as well as Andersson [10,12,13]. Andersson has shown that reliable thermal centering can be achieved using this simplified control law:

$$\dot{\psi}_{cmd} = \dot{\psi}_{ss} - k\ddot{e} \tag{2.1}$$

where $\dot{\psi}_{ss}$ is the desired nominal steady-state turn rate, k is a proportional gain, and $\dot{\psi}_{cmd}$ is the commanded turn rate that is sent to the low-level autopilot. This autonomous centering controller uses the second derivative of energy ($\ddot{e}$) as a feedback term to continually adjust the turn rate allowing the glider to remain centered about the thermal core as illustrated in Figure 2.2. Because this controller is a codified form of the Reichmann method, it is implied that the thermal profile resembles a Gaussian distribution.

The feedback term ($\ddot{e}$) plays a crucial role in the performance of this controller as the performance diminishes with excessive latency. Estimation of $\ddot{e}$ is a difficult problem and is discussed in Chapter 3. Although the structure of thermals can change drastically from one to the next, this algorithm performs relatively well as it continuously turns the aircraft toward areas of lift.

2.3 Steady Turn in a Thermal

This section describes the kinematic equations that govern a glider performing a steady-state turn. These equations are used to determine an expression for turn rate (which is used for thermal centering) and to determine a soaring trajectory that maximizes the rate of energy gain in a thermal.

2.3.1 Kinematics of a Glider in a Steady Turn

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 2.3. 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:

$$\dot{x} = v_a \cos\gamma \cos\psi + w_x \tag{2.2}$$

$$\dot{y} = v_a \cos\gamma \sin\psi + w_y \tag{2.3}$$

$$\dot{z} = v_a \sin\gamma + w_z \tag{2.4}$$

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 [27].

The summation of the forces in the stability axes for a glider performing a steady turn are:

$$L \sin\phi = \frac{m{v_a}^2}{R} \tag{2.5}$$

$$L \cos\phi = mg\cos\gamma \tag{2.6}$$

$$D = mg\sin\gamma \tag{2.7}$$

where R is the radius of the turn, L is the lift force and D is the drag force.

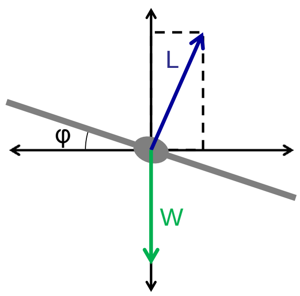

Figure 2.3. Body forces in a steady turn.

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 2.5 can be written as:

$$L \sin \phi = m v_a \dot{\psi} \tag{2.8}$$

Dividing equation Equation 2.8 by Equation 2.6 gives:

$$\tan \phi = \frac{v_a \dot{\psi}}{g \cos \gamma} \tag{2.9}$$

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

$$\dot{\psi} = \frac{g \tan \phi \cos \gamma}{v_a} = \frac{v_a \cos \gamma}{R} \tag{2.10}$$

Sink rate is determined by the aerodynamic properties of the aircraft. Using the standard definition of lift and drag coefficients, Equations 2.6 and 2.7 become:

$$\frac{q S C_L \cos \phi}{mg} = \cos \gamma \tag{2.11}$$

$$\frac{q S C_D}{mg} = \sin \gamma \tag{2.12}$$

Dividing Equation 2.12 by Equation 2.11:

$$\frac{C_D}{C_L \cos \phi} = \tan \gamma \tag{2.13}$$

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

$$C_D = \sum_{i=0}^{n} a_i C_L^i \tag{2.14}$$

Drag polar coefficients for the SB-XC are given in Appendix A. Substituting Equation 2.14 into Equation 2.13:

$$\gamma = \tan^{-1}\left(\frac{\sum_{i=0}^{n} a_i C_L^i}{C_L \cos\phi}\right) \tag{2.15}$$

where C_L is a function of airspeed. Equation 2.15 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.

2.3.2 Total Energy and Rates of Change of Total Energy

This section derives the equations that describe the energy states of the glider, including the total energy and the first and second derivatives of total energy. These energy states are useful for maximizing the climb rate in a thermal as well as accurately centering the thermal.

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

$$E = mgh + \frac{1}{2}mv^2 + E_s \tag{2.16}$$

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 2.16 can be modified to show the vehicle's specific energy using units of height. Specific energy (denoted by e) is total energy divided by weight.

$$e = \frac{E}{mg} = \frac{v^2}{2g} + h + e_s \tag{2.17}$$

In the case of an aerial vehicle, v is the velocity of the aircraft relative to the wind (v_a). Additionally, because we are interested in the energy states of an unpowered glider, stored energy is assumed to be zero. Differentiating Equation 2.17, we acquire an expression for $\dot{e}$:

$$\dot{e} = \frac{v_a \dot{v}_a}{g} + \dot{h} \tag{2.18}$$

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

$$\ddot{e} = \frac{\dot{v}_a^2 + v_a \ddot{v}_a}{g} + \ddot{h} \tag{2.19}$$

As mentioned in Section 2.2.4, this $\ddot{e}$ term is used as feedback for the thermal centering controller. Estimating this value is discussed in Chapter 3.

2.3.3 Optimal Turn Rate

In this application, optimal turn rate is defined as the turn rate that maximizes the rate of energy gain that is achievable given the current environmental conditions. Equation 2.18 in the previous section gives an expression for $\dot{e}$ in terms of v_a, $\dot{v}_a$, and $\dot{h}$. Since $\dot{h} = -\dot{z}$, The vertical kinematics equation (Equation 2.4) can be substituted into Equation 2.18.

$$\dot{e} = \frac{v_a \dot{v}_a}{g} - (v_a \sin\gamma + w_z) \tag{2.20}$$

Assuming a steady turn, $\dot{v}_a = 0$. This reduces $\dot{e}$ to the following:

$$\dot{e} = -v_a \sin\gamma - w_z \tag{2.21}$$

Equation 2.21 shows that the rate of energy gain is the upward wind speed minus the glider's sink rate in still air ($v_a \sin\gamma$ is sink rate). Equation 2.15 shows that γ (and therefore sink rate) is a function of both v_a and ϕ. Rewriting Equation 2.13, we obtain an expression for $\sin\gamma$:

$$\sin\gamma = \frac{C_D \cos\gamma}{C_L \cos\phi} \tag{2.22}$$

In a turn, γ is assumed to be small compared to ϕ. Using this information we can approximate $\cos\gamma = 1$:

$$\sin\gamma = \frac{C_D}{C_L \cos\phi} \tag{2.23}$$

Additionally, the small angle approximation allows the turn rate (from Equation 2.10) to be written independently of γ:

$$\dot{\psi} = \frac{g \tan \phi}{v_a} = \frac{v_a}{R} \tag{2.24}$$

Using the n^{th} order polynomial approximation for the drag polar, Equation 2.23 becomes:

$$\sin \gamma = \frac{\sum_{i=0}^{n} a_i C_L^i}{C_L \cos \phi} \tag{2.25}$$

This equation can now be substituted into Equation 2.21 to solve for sink rate in still air terms of v_a and ϕ:

$$\dot{e} = -v_a \frac{\sum_{i=0}^{n} a_i C_L^i}{C_L \cos \phi} - w_z \tag{2.26}$$

Determination of the optimal climb conditions can be achieved by adding a thermal model. For this analysis, it is assumed that the updraft profile is a Gaussian (or normal) distribution with a set core strength w_0 and radius R_{th}. Vertical wind speed varies with distance from the core as:

$$w_z(R) = -w_0 e^{-\frac{1}{2}\left(\frac{R}{R_{th}}\right)^2} \tag{2.27}$$

where R defines the distance from the thermal center. A thermal profile is shown in Figure 2.4.

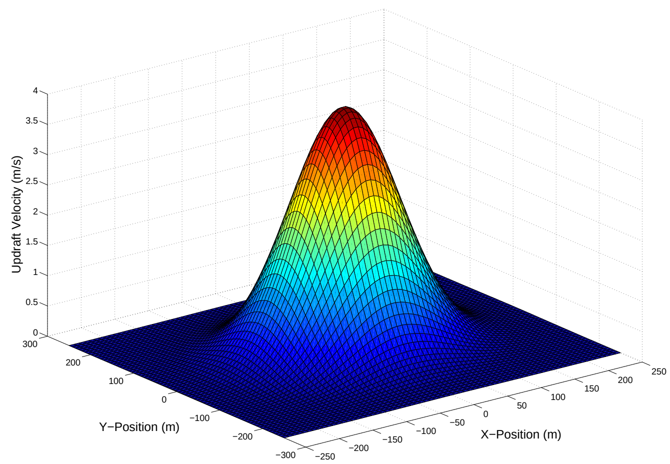

Figure 2.4. Gaussian thermal model with a $R_{th} = 70\ m$ and $w_0 = 4\ m/s$.

Given that the vertical component of wind varies with distance from the thermal core, one can compute the optimal airspeed and bank angle to fly in a given thermal by maximizing:

$$\dot{e} = -v_a \frac{\sum_{i=0}^{n} a_i C_L^i}{C_L \cos\phi} + w_0 e^{-\frac{1}{2}\left(\frac{R}{R_{th}}\right)^2} \tag{2.28}$$

Using Equation 2.24, we can write this equation strictly in terms of v_a and ϕ:

$$\dot{e} = -v_a \frac{\sum_{i=0}^{n} a_i C_L^i}{C_L \cos\phi} + w_0 e^{-\frac{1}{2}\left(\frac{v_a^2}{R_{th} g \tan\phi}\right)^2} \tag{2.29}$$

Equation 2.29 therefore suggests that one can schedule an airspeed and bank angle that maximizes the climb rate for a given thermal. 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 2.5 shows the rate of change of energy for an SB-XC circling in the thermal. Note the strong dependence of $\dot{e}$ on turn radius and comparatively weak dependence on airspeed. Flying at a nominal airspeed of approximately 12.5 m/s gives good performance, and this remains true as thermal radius varies.

The problem now is to determine the steady-state turn rate that maximizes energy gain without perfect knowledge of the thermal. Thermal mapping has been discussed in Section 2.2.3, but thermal mapping methods [12, 26] have many drawbacks. Thermal mapping methods typically:

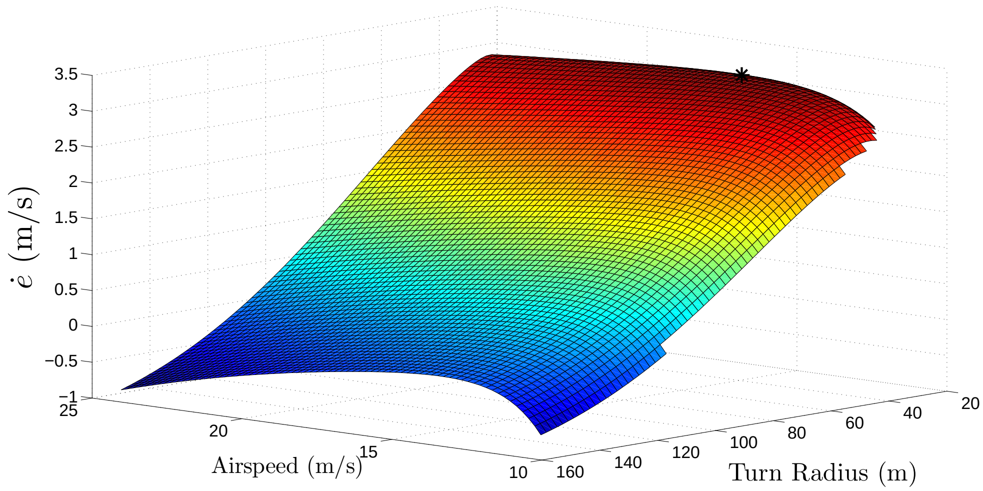

Figure 2.5. Climb rate as a function of turn radius and airspeed for the SB-XC glider centered on the thermal shown in Figure 2.4.

- require significant processing power and have higher memory requirements
- take overhead time to generate a usable map
- require exploration that leads the aircraft out of high lift areas
- do not perform as well in windy conditions
- require an accurate aircraft model

However, an alternative solution to this problem is presented in Section 2.4. This solution actively seeks areas where maximum climb rate can be achieved and doesn't require a map of the thermal.

2.4 Extremum Seeking Control for Steady-State Turn Convergence

This section introduces a control method that increases the rate of energy gain by actively tracking the optimal turn rate in a thermal. Using only a time history of commanded turn rate and corresponding energy rates, the controller continuously decides to tighten or widen the turn seeking out the maximum rate of energy gain.

This is accomplished by perturbing turn radius around its current nominal value and using the recorded energy information to determine the gradient $\frac{\partial \dot{e}}{\partial R}$, the nominal turn radius is then changed in the direction of increasing $\dot{e}$. The repetition of this process causes the steady-state turn radius to converge to the optimal value. A flowchart illustrating this adaptive thermalling algorithm is shown in Figure 2.6.

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 Edwards [22] and Allen [10]. When the algorithm is first triggered, the turn radius is initialized to a relative value that depends on the scale of the aircraft being used. For a full-size glider with a 15 meter wingspan, 80 meters is a typical starting radius. For the small-scale SB-XC glider, an initial turn radius could be as small as 20 meters. This initial radius should be small enough that the thermal won't be lost by reaching a steady-state value that is larger than the thermal. Additionally, a small initial radius allows centering to occur faster because more loops can be completed in the same amount of time. However, thermal centering performance diminishes as radius decreases due to the increasing effects of estimate lag. Therefore, the initial turn radius must be chosen by negotiating a balance between climb performance and stability.

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). Once the specified number of loops are completed, the steady-state turn radius is altered 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 $\bar{\dot{e}}$ values are compared. If an increase is observed, the turn radius is perturbed in the same direction. However, if a decrease is observed, the turn radius is perturbed in the opposite direction. The algorithm follows the logic table presented in Table 2.1 and repeats this process until the thermal dissipates or another task is assigned.

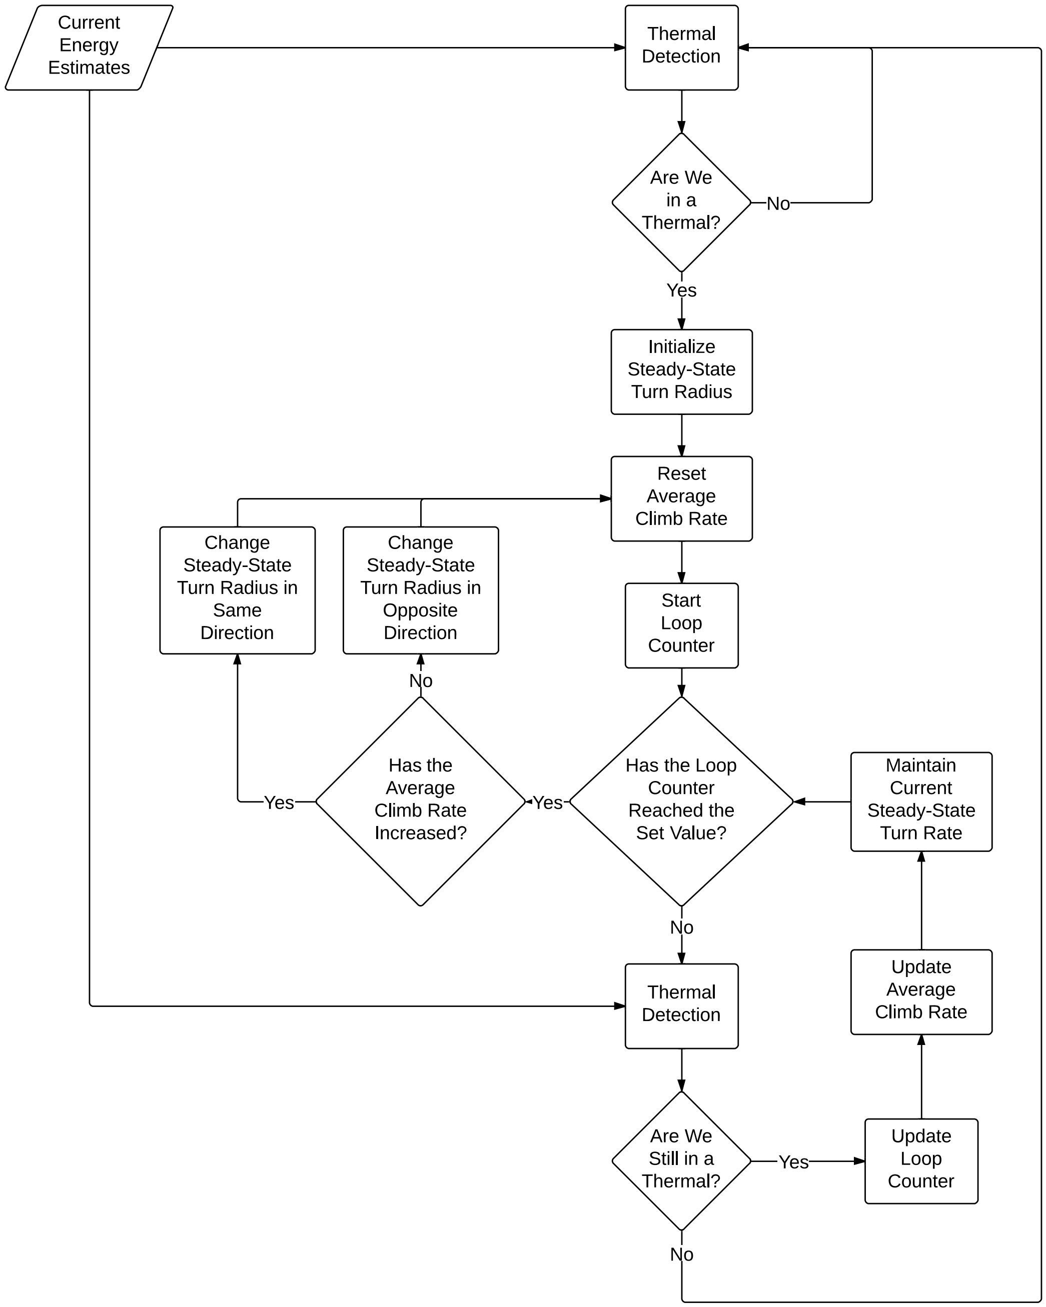

Figure 2.6. Flowchart showing the structure of the turn radius convergence algorithm.

Table 2.1. Logic table for thermal radius convergence.

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

While proofs of stability of extremum seeking control for thermal soaring have not yet been derived, preliminary simulations in Simulink show stable convergence. The following sample simulation shows the results of an SB-XC circling in a Gaussian thermal with $w_0 = 2.56$ and $R_{th} = 75$.

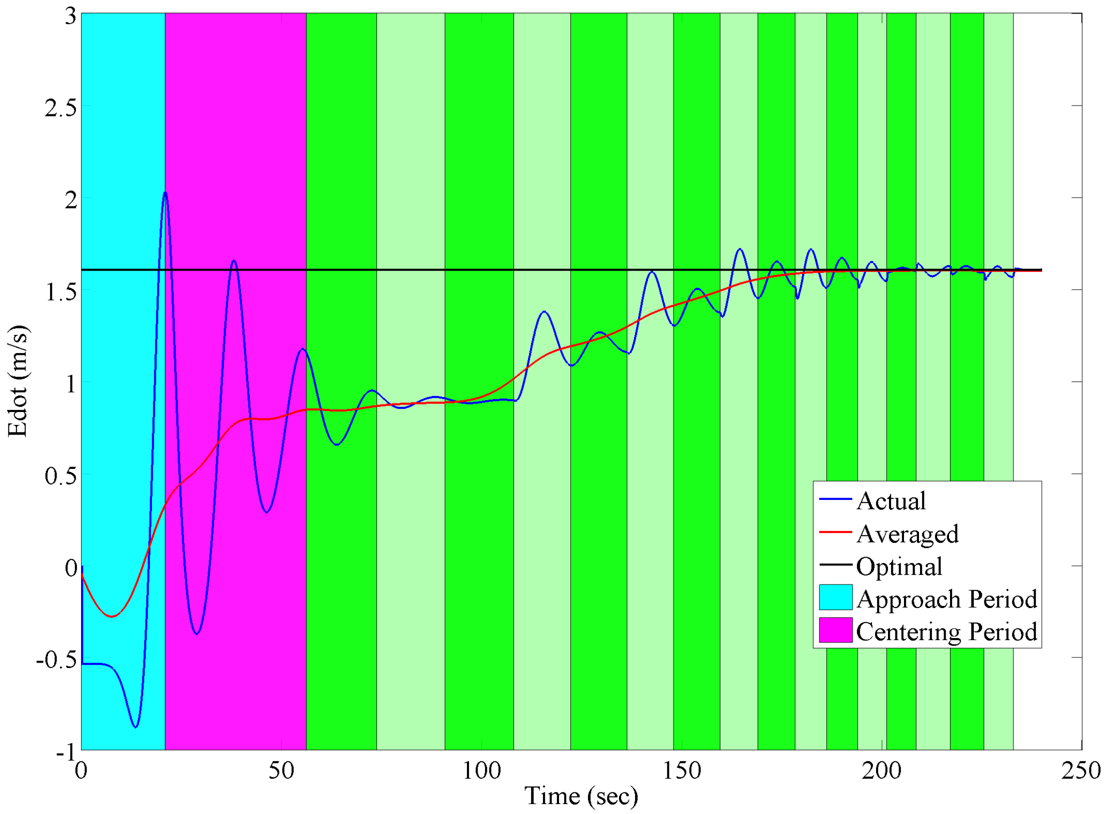

Figure 2.7. Convergence of $\dot{e}$ to the optimal value.

Figure 2.7 shows the convergence of $\dot{e}$ to the maximum value that can be achieved with this aircraft/thermal combination. The aircraft initially centered the thermal with a 40 meter turn radius ($\phi \approx 27.5$ deg). This centering period is represented in Figure 2.7 by the magenta colored region. Once the thermal was centered, the steady state turn conditions were determined by the extremum seeking controller. Convergence of the corresponding bank angles can be seen in Figure 2.8.

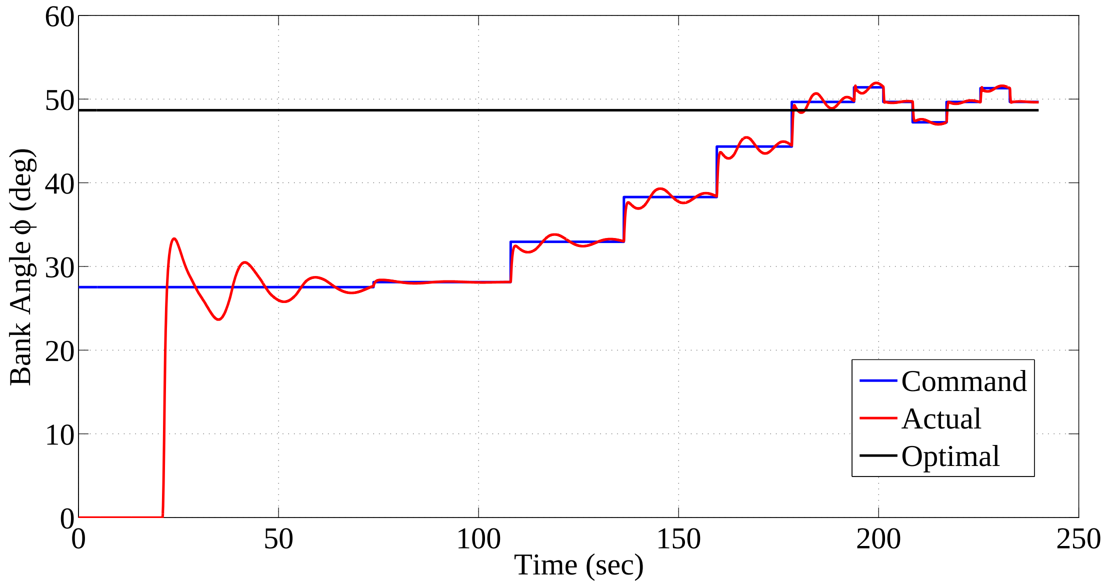

Figure 2.8. Bank angle converging toward the optimal value.

Notice how the bank angle does not remain exactly on the optimal value. As mentioned in Section 2.1, thermals have the tendency to expand as they rise; this slight perturbation around the optimal value provides robustness by allowing the controller to track the maximum climb rate throughout the lifetime of the thermal. Also, this version of the controller perturbs the turn radius proportionally to the magnitude of the scaled gradient with the intent of providing faster and more accurate convergence. This is represented by:

$$|\Delta R_{k+1}| = K_{sc} \left| \frac{\Delta \dot{e}_k}{\Delta R_k} \right| \tag{2.30}$$

where K_{sc} is a proportional gain. The flight path for this simulation is shown in Figure 2.9.

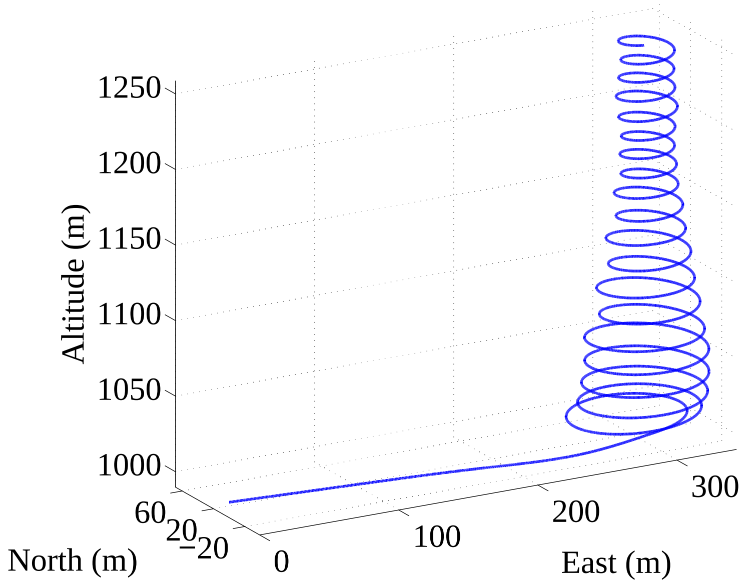

Figure 2.9. Flight path converging toward the optimal turn radius.

These preliminary simulations showed promising results. Further testing of the extremum seeking controller was done in the Silent Wings soaring simulator with a more accurate aircraft model as well as more realistic atmospheric conditions. These results are presented in Section 4.4.2.

Chapter 3

Energy State Estimation

This chapter defines the problem of energy state estimation. It has three purposes: (a) analyze on-board sensors capable of providing state information that is useful for energy estimation; (b) explore current energy estimation schemes; (c) propose a differentiation and filtering method to provide improved energy state estimates.

This chapter describes the acquisition of the $\ddot{e}$ estimate that is used as feedback in the thermal centering controller (Chapter 2). Section 3.1 discusses the capabilities of sensors that can be carried on-board the aircraft and their potential contribution to an accurate energy state estimate. Section 3.2 analyzes traditional numerical differentiation and filtering methods. It then introduces the Savitzky-Golay filter as a basis for performing successive numerical differentiations and describes an adapted version of the estimator that reduces phase lag thus improving centering performance.

3.1 Onboard Sensors Used for Estimation

Since there is no way of measuring $\ddot{e}$ directly, an alternative method must be used to acquire this value. One method is to use a TE-compensated variometer to measure $\dot{e}$ and differentiate the incoming signal in real-time. Another method is to measure v_a and h and differentiate to obtain values of v_a, $\dot{v}_a$, $\ddot{v}_a$, and $\ddot{h}$ that appear in Equation 2.19. However, discrete signal differentiation is a difficult task, especially if significant sensor noise is present. According to Allen and Lin, "Latency in energy rate and energy acceleration estimates was found to be the primary cause of reduced controller performance in weak thermals [10].” Therefore, in order to accomplish robust and precise centering performance, we must choose the method of obtaining $\ddot{e}$ based on the performance of available sensors.

Small autonomous UAVs have the ability to carry a variety of sensing instruments. Table 3.1 lists and discusses sensors that are typically found on autonomous and manned gliders.

Table 3.1. Sensors Capable for Use on an Autonomous Soaring Aircraft

Sensor	**Symbol**	**Measurement**
Accelerometer	a_x, a_y, a_z	Accelerations in the inertial frame
Gyroscope	$\omega_x, \omega_y, \omega_z$	Rotation rates in the inertial frame
Magnetometer	ψ	Aircraft heading
GPS	lat, lon, h	Position relative to the Earth’s surface
Altimeter	h	Altitude
Pitot Probe	v_a	Indicated airspeed (IAS)
Variometer	$\dot{h}$	Climb rate
TE-Variometer	$\dot{e}$	Rate of energy gain/loss

TE-compensated variometers provide relatively accurate $\dot{e}$ measurements; these measurements can be differentiated numerically to obtain an estimate of $\ddot{e}$. However, TE-variometers introduce a relatively large amount of measurement lag (typically 2 seconds or more). Human glider pilots are able to compensate for this lag by creating a mental map of the thermal. However, autonomous thermal mapping adds complexity to the system and requires excessive amounts of processing power and memory to run in real-time. Meeting these computational requirements can increase costs, restrict the already limited payload of small UAVs, and create additional problems by altering mass, wing loading, inertia, and other aircraft properties.

Unlike the TE-variometer, the measurement of v_a and h adds a negligible amount of lag to the system. Airspeed can be measured using a pitot probe and height can be measured using an altimeter or a GPS unit. However, due to measurement noise, acquiring an accurate representation of the first and second derivatives of the airspeed and height measurements is difficult. Additionally, wind dynamics are very unpredictable and measurements can fluctuate significantly on short time scales. Conventional filtering methods have been used to reduce measurement noise so that reasonably smooth estimates of $\dot{e}$ and $\ddot{e}$ can be obtained, but this introduces significant phase lag that degrades centering performance (especially in smaller thermals).

3.2 Differentiation and Filtering Methods

Various differentiation and filtering methods were explored for estimating $\dot{v}_a$, $\ddot{v}_a$, and $\ddot{h}$. To assess the effectiveness of different filtering and estimation methods, an analytic function that is qualitatively representative of the rate of energy change of a thermalling glider was defined:

$$y(t) = (t/10)\sin(2t) + 3t + 10 \tag{3.1}$$

This function was sampled at a rate of 31 Hz (matching the rate of the low-level autopilot discussed in Section 4.2) and corrupted with zero-mean Gaussian random noise with $\sigma = 0.05$ to represent sensor noise. The function of Equation 3.1 and its time derivatives

$$y'(t) = \left(\frac{t}{5}\right)\cos(2t) + \left(\frac{1}{10}\right)\sin(2t) + 3 \tag{3.2}$$

$$y''(t) = \frac{2}{5}\left(\cos(2t) - t\sin(2t)\right) \tag{3.3}$$

are used as reference signals to compare filtering and estimation methods.

3.2.1 Finite Difference Methods

In this section, estimating derivatives using a conventional finite difference method is explored. Various order central difference methods were tested. The 3, 5, and 9 point central difference formulas are represented as:

$$f'(t) = \frac{f(t+\Delta t) - f(t-\Delta t)}{2\Delta t} \tag{3.4}$$

$$f'(t) = \frac{f(t-2\Delta t) - 8f(t-\Delta t) + 8f(t+\Delta t) + f(t+2\Delta t)}{12\Delta t} \tag{3.5}$$

$$f'(t) = \frac{\begin{array}{l} f(t-4\Delta t) - 32f(t-3\Delta t) + 168f(t-2\Delta t) - 672f(t-\Delta t) + \cdots \\ \qquad 672f(t+\Delta t) - 168f(t+2\Delta t) + 32f(t+3\Delta t) - 3f(t+4\Delta t) \end{array}}{840\Delta t} \tag{3.6}$$

respectively. These equations were used to estimate the first and second derivatives of a noisy signal ($\sigma = 0.05$). These results can be seen in Figure 3.1.

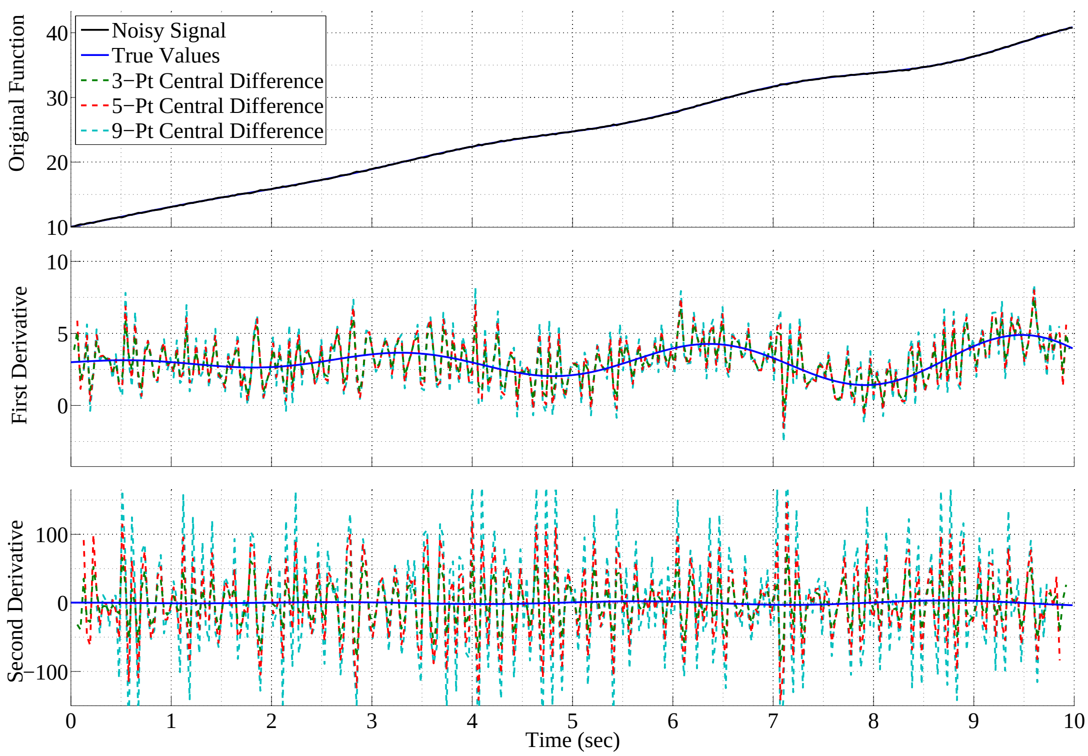

Figure 3.1. Estimating the first and second derivatives of a noisy signal ($\sigma = 0.05$) using various central difference formulas.

It can be seen that by the addition of even a small disturbance to the original signal, the estimate of the second derivative is nearly unrecognizable. It can also be seen that the least noisy estimate is given by the 3-point central difference. Increased noise presence is observed as higher order central difference formulas are used. This may seem counterintuitive but this occurrence is explained in Section 3.2.2. Additionally, the central difference formulas are non-causal (they require future measurements to compute current derivative estimates). This adds estimation delay which is equal to one-half of the window size where M is the number of measurements in the current window:

$$t_{lag} = \frac{(M-1)\Delta t}{2} \tag{3.7}$$

Due to the extreme amount of noise in the derivative estimates, straightforward numerical differentiation is not a feasible differentiation method. Pre-filtering the signal can reduce measurement noise before numerical differentiation methods are used. Because the measurement noise lives at a much higher frequency than the energy dynamics, a low-pass filter is a viable candidate. Using various first order low-pass filters combined with the 3-point central difference method, derivative estimates of the noisy signal were computed. These results can be seen in Figure 3.2. Note that measurement noise added to this input signal is double the amount used in Figure 3.1.

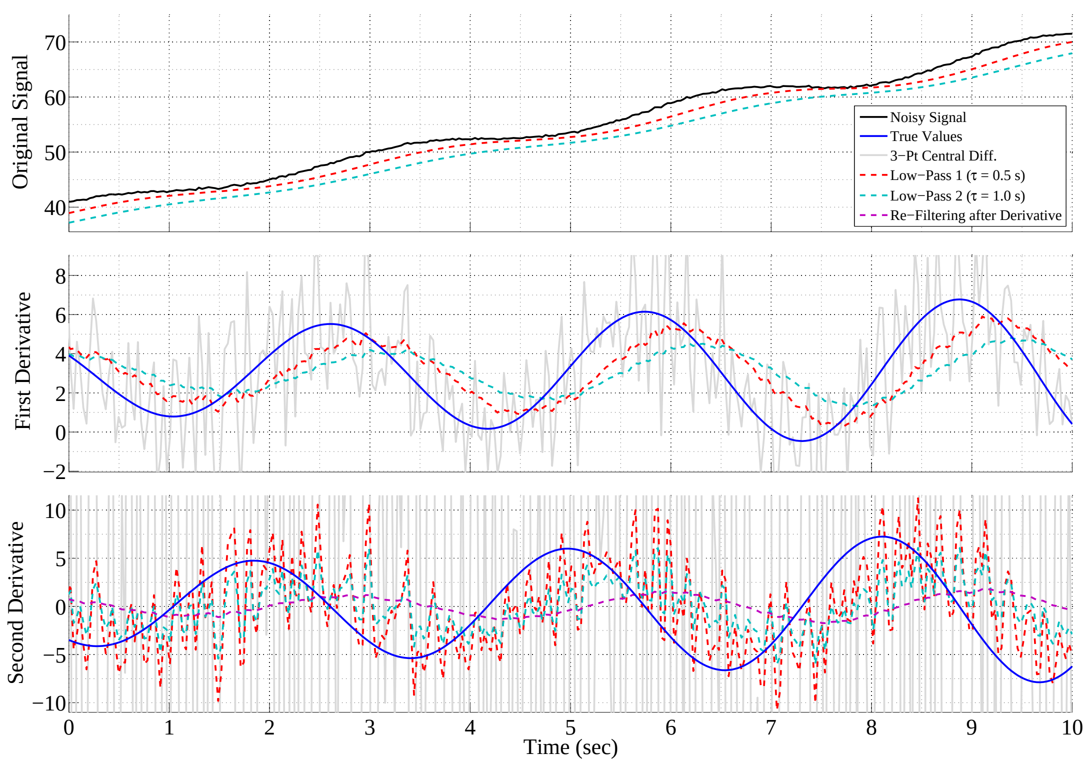

Figure 3.2. Estimating the first and second derivatives of a noisy signal ($\sigma = 0.1$) using a low-pass filter and 3-point central difference formula.

The smoothest estimate of the second derivative was produced after filtering and numerically differentiating twice. However, this estimate introduced over a second of lag between the time the measurement was taken and the time the estimate was available. Also, the amplitude of the estimates are significantly less than the actual value. It can be observed that the smoother the estimate, the higher the amplitude error. Figure 3.2 also shows that a trade off exists between lag time and estimation error. With increased noise, this latency and estimate error grows significantly. The energy filters used by Allen (shown in Section 4.4.1) severely limited the centering performance due to lag on the order of $\sim$3 seconds [6].

3.2.2 Savitzky-Golay Filter

Adaptive filters can use a higher order estimate and still provide much better noise rejection. An example of an adaptive convolution filter is the Savitzky-Golay (SG) filter [28]. This filter performs a polynomial fit on the measurements within a moving window of time (Figure 3.3), and it is adaptive in the sense that the computed polynomial fit to the measured values is continuously changing with time.

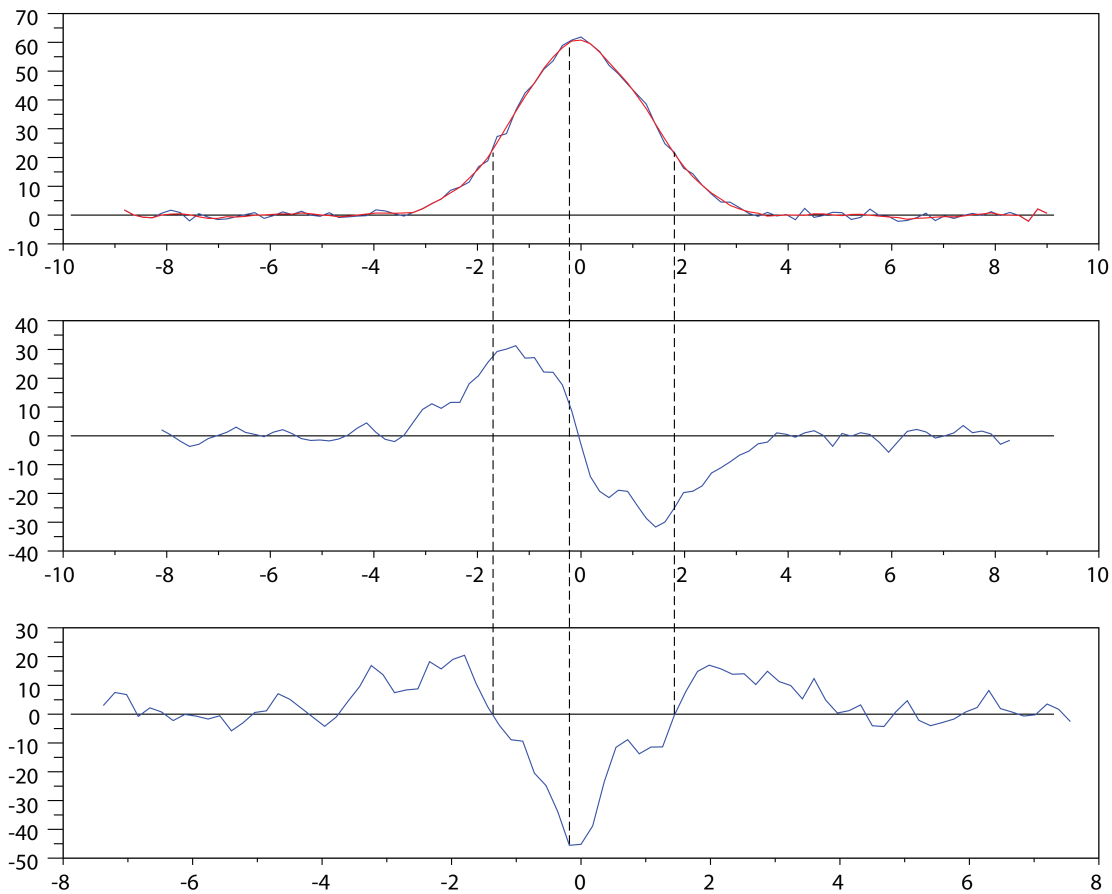

Figure 3.3. Approximating numerical derivatives using a Savitzky-Golay filter [7].

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 [29]. Given a window of time, T, containing M measurements, and a desired polynomial of order N of the following form:

$$p(x) = a_0 + a_1 x + \cdots + a_{N-1} x^{N-1} + a_N x^N \tag{3.8}$$

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

$$\hat{\mathbf{a}} = \begin{bmatrix} a_0 & a_1 & \cdots & a_{N-1} & a_N \end{bmatrix}^T \tag{3.9}$$

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

$$\hat{\mathbf{a}} = \mathbf{K}\mathbf{y} \tag{3.10}$$

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

$$\mathbf{K} = (\mathbf{D}^T\mathbf{D})^{-1}\mathbf{D}^T \tag{3.11}$$

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

$$\mathbf{D} = \begin{bmatrix} 1 & x_1 & x_1^2 & \cdots & x_1^N \\ 1 & x_2 & x_2^2 & \cdots & x_2^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_M & x_M^2 & \cdots & x_M^N \end{bmatrix} \tag{3.12}$$

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

$$T = (m-1)\Delta t \tag{3.13}$$

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

$$\bar{\mathbf{D}} = \begin{bmatrix} 1 & \frac{t_0-t_0}{T} & (\frac{t_0-t_0}{T})^2 & \cdots & (\frac{t_0-t_0}{T})^N \\ 1 & \frac{t_1-t_0}{T} & (\frac{t_1-t_0}{T})^2 & \cdots & (\frac{t_1-t_0}{T})^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \frac{t_M-t_0}{T} & (\frac{t_M-t_0}{T})^2 & \cdots & (\frac{t_M-t_0}{T})^N \end{bmatrix} \tag{3.14}$$

where $t_m,\ m = 0 \ldots M$ 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

$$\bar{\mathbf{a}} = \bar{\mathbf{K}}\mathbf{y} \tag{3.15}$$

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:

$$\mathbf{D} = \begin{bmatrix} 1 & t - \Delta t \\ 1 & t + \Delta t \end{bmatrix} \tag{3.16}$$

Substituting into Equation 3.11, $\mathbf{K}$ becomes:

$$\mathbf{K} = \frac{1}{2} \begin{bmatrix} \frac{t}{\Delta t} + 1 & -\frac{t}{\Delta t} + 1 \\ -\frac{1}{\Delta t} & \frac{1}{\Delta t} \end{bmatrix} \tag{3.17}$$

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

$$\mathbf{y} = \begin{bmatrix} f(t - \Delta t) \\ f(t + \Delta t) \end{bmatrix} \tag{3.18}$$

Substituting into Equation 3.15 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

$$f'(t) = \frac{f(t + \Delta t) - f(t - \Delta t)}{2\Delta t} \tag{3.19}$$

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.

This derivation exemplifies the reason why the SG filter outperforms the central difference solution. Because the central difference solution fits a polynomial curve with same number of polynomial coefficients as measurements in the chosen measurement window, the size of the measurement window and the order of the polynomial fit are directly related. Therefore the central difference simply precomputes the weighting coefficients for the fully determined system. However, the SG filter presents the ability to choose the order of the polynomial fit independent of the window size, as long as the window size is greater than the number of coefficients in the polynomial. Tuning these values alter the amount of smoothing that is performed on the signal.

The ability of the SG filter to accurately estimate polynomial coefficients is briefly demonstrated in Figure 3.4. The analytic function given in Equation 3.1 was used as a comparison. The polynomial coefficients generated by the Savitzy-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 3.4 shows that smooth derivatives can be obtained, while numerical derivatives (shown in gray) are enormously affected by measurement noise.

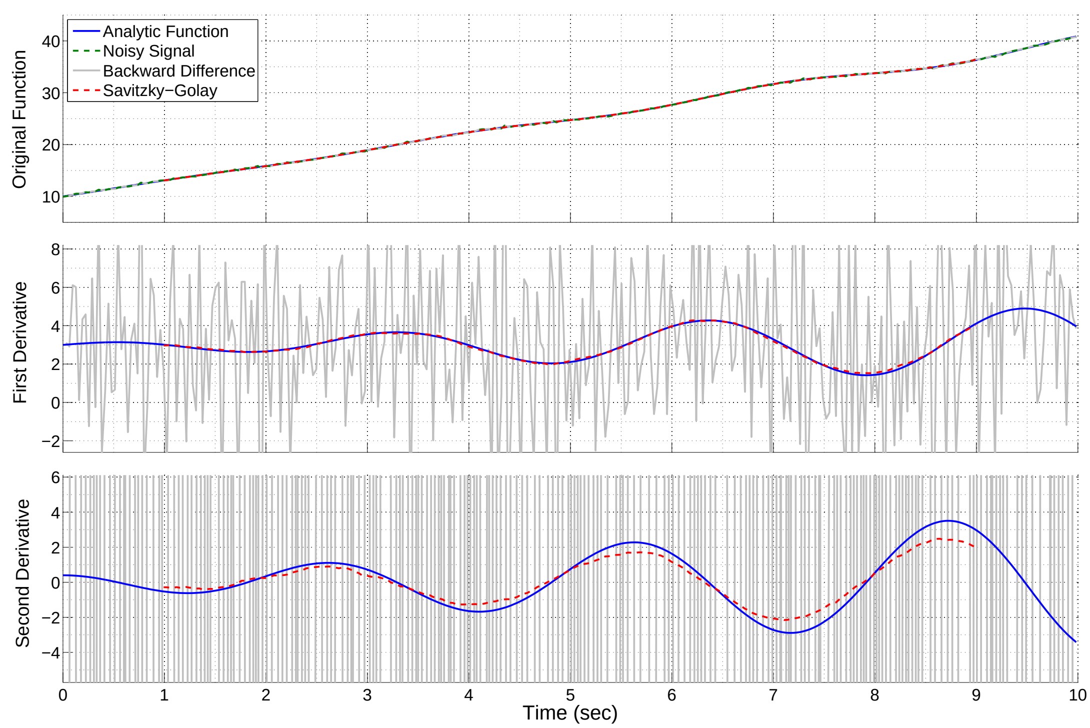

Figure 3.4. Estimating the first and second derivatives of a noisy signal ($\sigma = 0.1$) using a symmetric Savitzky-Golar filter.

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 is proposed; rather than computing estimates at the center of the time window, estimates are computed at the end of the time window.

3.2.3 Backward Savitzky-Golay Filter

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,

$$e(t_m) = a_0 + a_1 \left(\frac{t_m - t_0}{T}\right) + \ldots + a_N \left(\frac{t_m - t_0}{T}\right)^N + \mathcal{N}(0, \sigma^2) \tag{3.20}$$

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 $m = 1 \ldots M$ 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:

$$\hat{e}(t_m) = \sum_{n=0}^{N} \hat{a}_n \left(\frac{t_m - t_0}{T}\right)^n \tag{3.21}$$

$$\dot{\hat{e}}(t_m) = \sum_{n=1}^{N} n\hat{a}_n \left(\frac{t_m - t_0}{T}\right)^{(n-1)} \tag{3.22}$$

$$\ddot{\hat{e}}(t_m) = \sum_{n=2}^{N} n(n-1)\hat{a}_n \left(\frac{t_m - t_0}{T}\right)^{(n-2)} \tag{3.23}$$

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 $\left(\frac{t_m - t_0}{T}\right) = 1$. The energy estimates reduce down to

$$\hat{e}(t_m) = \sum_{n=0}^{N} \hat{a}_n \tag{3.24}$$

$$\dot{\hat{e}}(t_m) = \sum_{n=1}^{N} n\hat{a}_n \tag{3.25}$$

$$\ddot{\hat{e}}(t_m) = \sum_{n=2}^{N} n(n-1)\hat{a}_n \tag{3.26}$$

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 3.5, which uses the analytic function (given in Equation 3.1) 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.

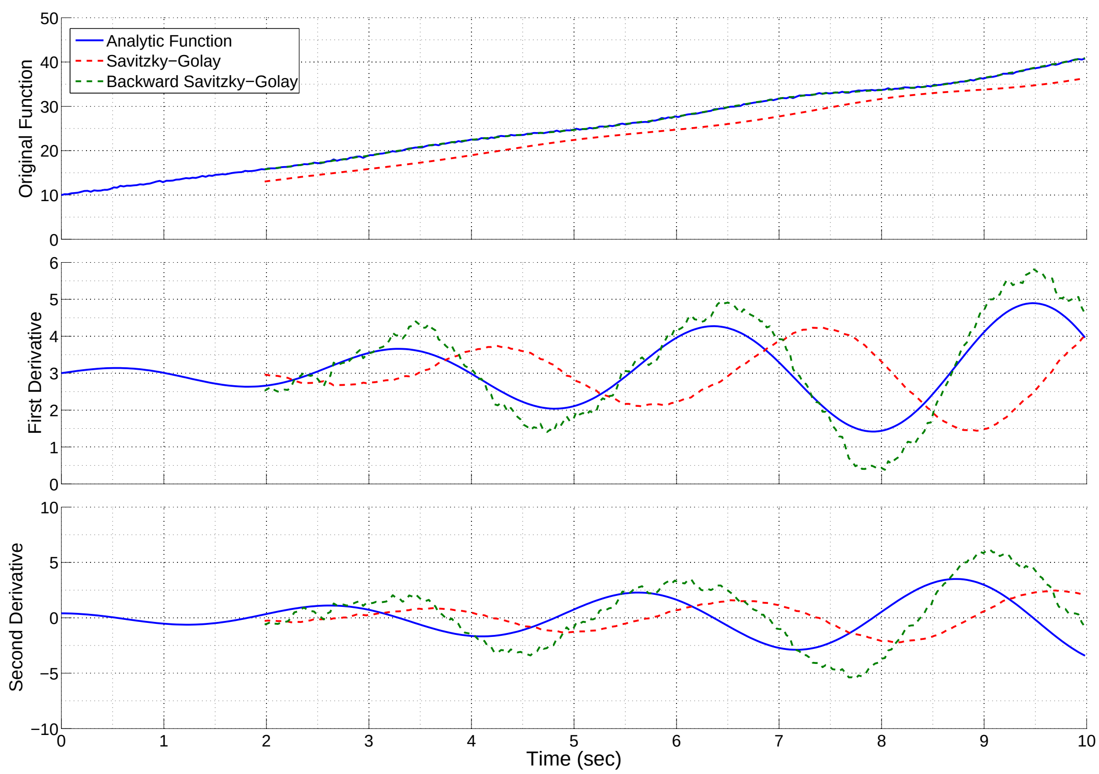

Figure 3.5. Estimating the first and second derivatives of a noisy signal ($\sigma = 0.1$) using a delayed Savitzky-Golay filter and a backward SG filter.

In Figure 3.5, 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

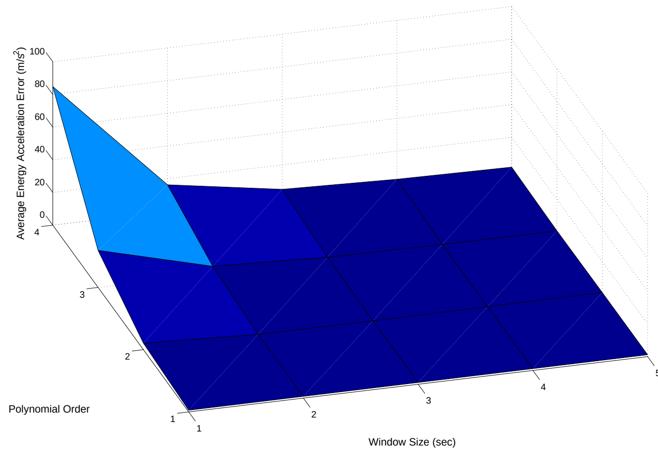

Figure 3.6. Average error for various polynomial orders and window sizes.

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 3.1. The average error for each of the cases is shown Figure 3.6.

This figure demonstrates that fitting a small time window (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 noise as well as the base signal. This is the equivalent of setting a conventional 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 siginficantly 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 backward SG filter was implemented on a low-level autopilot. This autopilot was used in a feedback loop to remotely fly full-size gliders in Silent Wings flight simulator (www.silentwings.no). The thermal centering controller (presented in Section 2.2.4) was used to heuristically tune the energy estimator.

Chapter 4

Modeling, Control, and Simulation

This chapter describes the simulation environment that was used and presents the results from various thermal soaring simulations. Specifically, it has four purposes: (a) validate the aircraft model that was used in simulations; (b) construct the simulation environment including physical hardware that was used; (c) implement a low-level flight controller; (d) analyze the performance of the thermal soaring algorithms presented in this thesis.

As the SB-XC was the selected aircraft for this implementation, a build-up of the aerodynamic model and validation of this model is introduced in Section 4.1. Section 4.2 outlines the simulator structure and provides details about the hardware and software that was used for each of the different systems. Section 4.3 describes the control laws used to implement low-level aircraft control. Section 4.4 concludes the chapter with a performance analysis of the proposed thermal soaring techniques.

4.1 Aircraft Modeling and Validation

Simulations using the SB-XC model were performed in the Silent Wings simulation environment. To ensure accurate vehicle performance, an SB-XC aerodynamic model was created using XFLR5 (www.xflr5.com). XFLR5 is a software package designed specifically for the analysis of aircraft flying at low Reynolds numbers. By utilizing XFoil's airfoil analysis tool combined with 3D wing analysis capabilities based on Lifiting Line Theory, Vortex Lattice Method, and 3D Panel Methods, an aerodynamic model was created for the SB-XC. The interface for this software is shown in Figure 4.1 where the results of a flow path analysis can be seen.

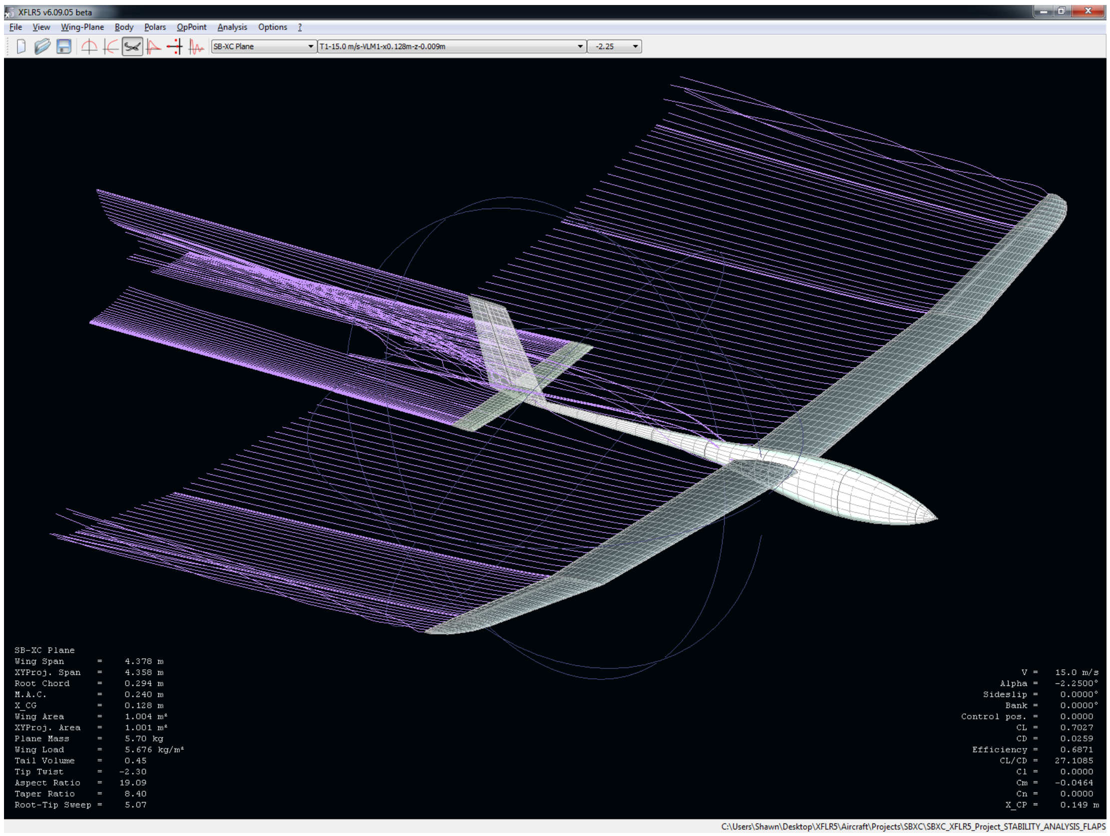

Figure 4.1. Interface showing the SB-XC model in XFLR5

An estimate of the stability derivatives provided a basis for modal frequency and damping ratio estimates. Using this XFLR5 model as a starting point, a model of the SB-XC was developed in Silent Wings.

The Silent Wings simulator accepts two types of aerodynamic models, a panel model and a linear model. The panel model is much more sophisticated in that it breaks the aircraft down into panels and simulates the flow over each panel. By the addition of the various force contributions due to each panel, the aircraft behaves in a realistic manner for a wide range of flight conditions. The linear model uses various stability derivatives that provide realistic dynamics near a specific operating condition. Originally, a linear SB-XC model was created but a panel model was later developed to ensure greater accuracy and ease the constraints of the operating condition.

Due to unforeseen complications with autonomous soaring regulations and time constraints, actual flight testing was unable to be performed to verify this model. However, because the SB-XC has been a popular vehicle for the autonomous soaring community, the performance of this vehicle has been the focus of past analyses. In 2007 (updated in 2008), Dan Edwards published a performance analysis of the SB-XC using previously recorded data and data that he collected specifically for this purpose [8]. A sink polar was generated for the aircraft as shown in Figure 4.2.

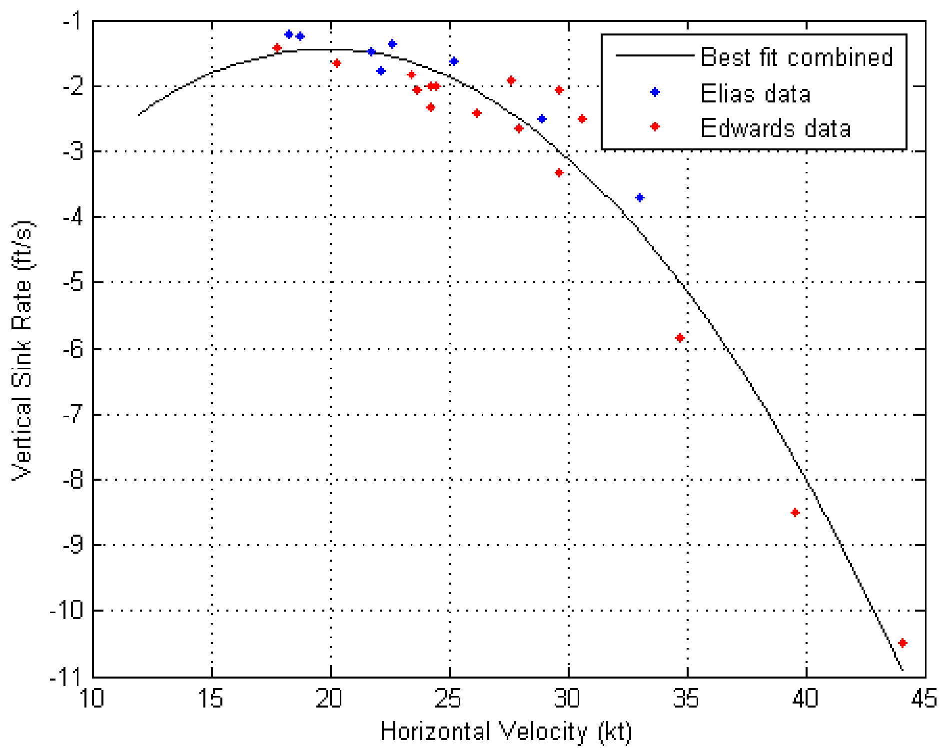

Figure 4.2. Sink polar generated using flight test data [8]

As suggested by Reichmann [5], a second order polynomial fit was used to represent the sink polar. Edwards calculated the coefficients of this polynomial which can be seen in Equation 4.1 (using units of kts).

$$V_{vert} = -0.0095\, V_{horiz}^2 + 0.3782\, V_{horiz} - 4.6072 \tag{4.1}$$

This sink polar was adapted to account for the mass difference of our specific

SB-XC flight system. Our flight-ready SB-XC was measured at 5.7 kg; this increased mass is likely due to the on-board energy estimation computer as well as other contributing factors. The additional mass changes the wing loading which results in a down and rightward shift in the sink polar. As discussed by Edwards [8], this shift is proportional to a scale factor that is determined by the ratio of the reference mass to our vehicle mass as shown in Equation 4.2.

$$m_{rat} = \frac{m_{ref}}{m} \tag{4.2}$$

Scaling both the vertical and horizontal wind speeds yield the corresponding velocities.

$$V_{vert} = \frac{V_{vert,ref}}{m_{rat}} \tag{4.3}$$

$$V_{horiz} = \frac{V_{horiz,ref}}{m_{rat}} \tag{4.4}$$

Our scaled sink polar (Equation 4.5) was compared to the original. This can be seen in Figure 4.3.

$$V_{vert} = -0.0162\, V_{horiz}^2 + 0.3782\, V_{horiz} - 2.7018 \quad (m/s) \tag{4.5}$$

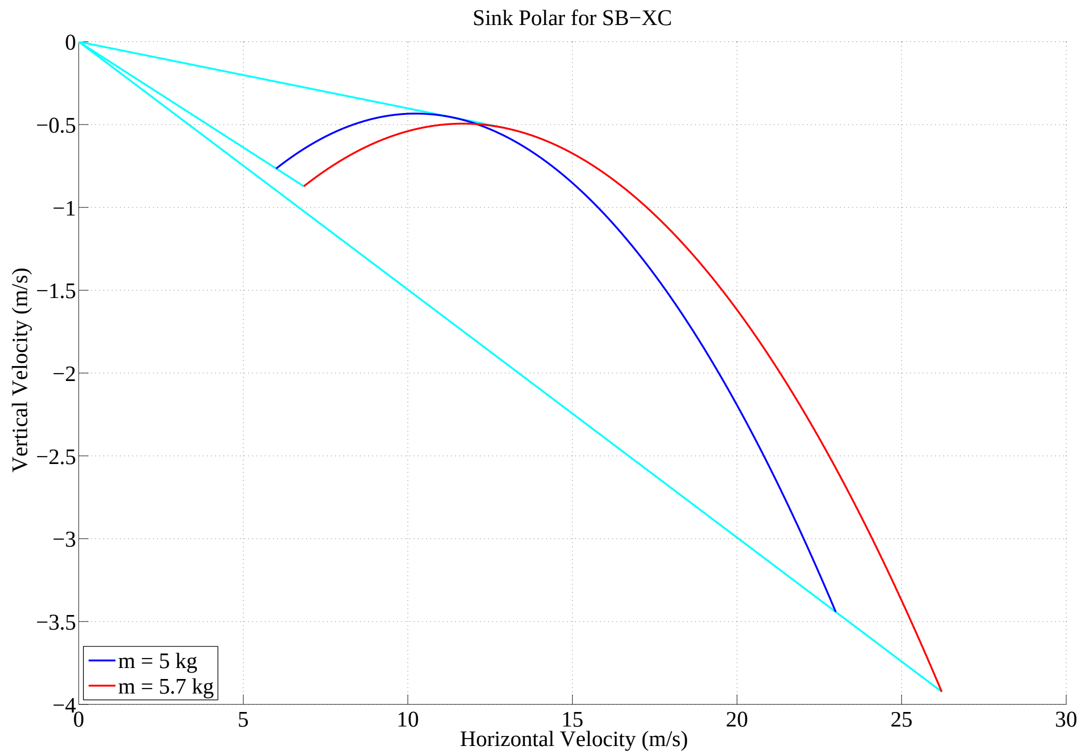

Figure 4.3. Shifted sink polar due to mass difference

Using this sink polar as a benchmark, the SB-XC drag model was fine tuned to match this curve (Figure 4.4).

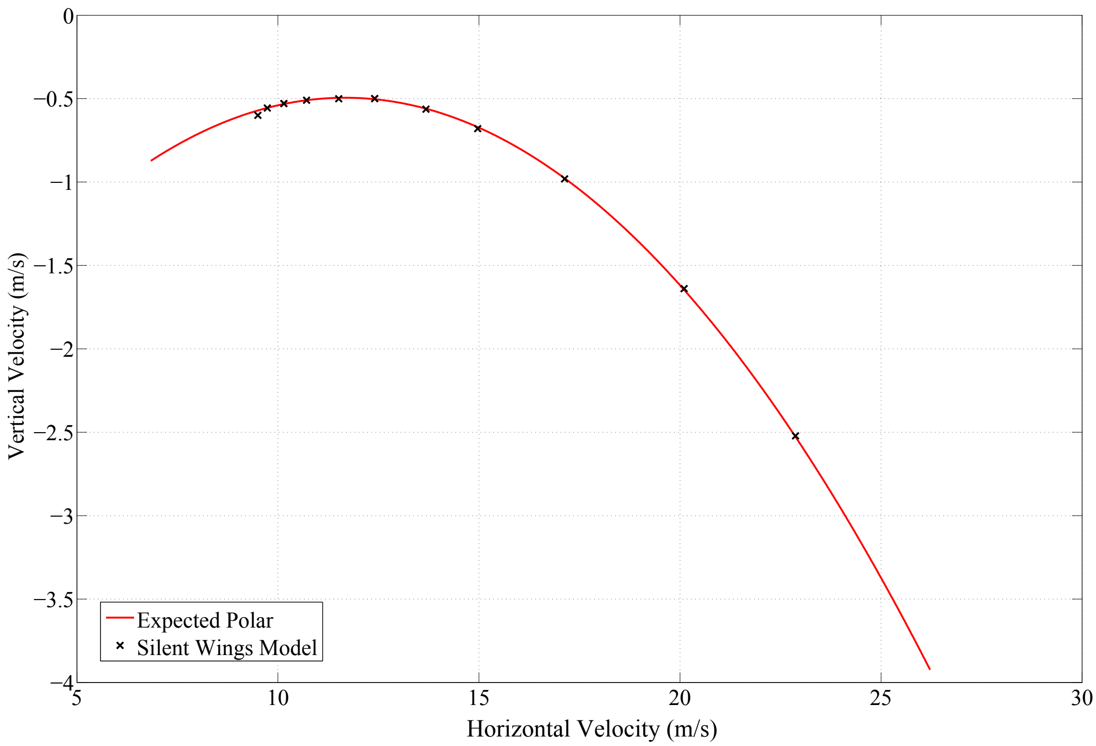

Figure 4.4. Comparison of model performance against the expected sink polar

This tuned aircraft model allows the SB-XC to be flown in simulation with confidence that the performance characteristics are similar to that of the actual aircraft.

4.2 The Simulation Environment

Silent Wings soaring simulator (www.silentwings.no) was used as the primary testing environment for the methods discussed throughout this thesis. Silent Wings is a commercially available soaring simulator and multiplayer game that is committed to replicating the soaring experience as accurately as possible. It 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 the user 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. Many of the gliders come with simulated instruments such as a TE-compensated variometer and a GPS, allowing measurements to be generated in real-time.

Figure 4.5. Glider autonomously climbing in a thermal. Simulation was performed using an SB-XC model in Silent Wings flight simulator.

A UDP interface facilitates the retrieval of these real-time measurements, along with the current state of the aircraft, by an outside source. Unlike Condor (http://www.condorsoaring.com), another comparable soaring simulator, Silent Wings provides the capability to control aircraft remotely through a similar UDP interface. This enables the simulation of flights where aircraft are controlled autonomously using state feedback. Figure 4.5 shows a glider while 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 real-time data from Silent Wings, a low-level autopilot system was designed to remotely control the aircraft in simulation. This autopilot is powered by an Arduino Mega 2560 microcontroller and communicates with Silent Wings (at a specified 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. This autopilot facilitated the testing of the energy estimator, thermal centering controller, and the extremum seeking controller in the Silent Wings flight simulator.

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 so that high-level decisions could be made such as when to switch to soaring mode, etc. Figure 4.6 shows a graphical representation of the testing and simulation network that was used.

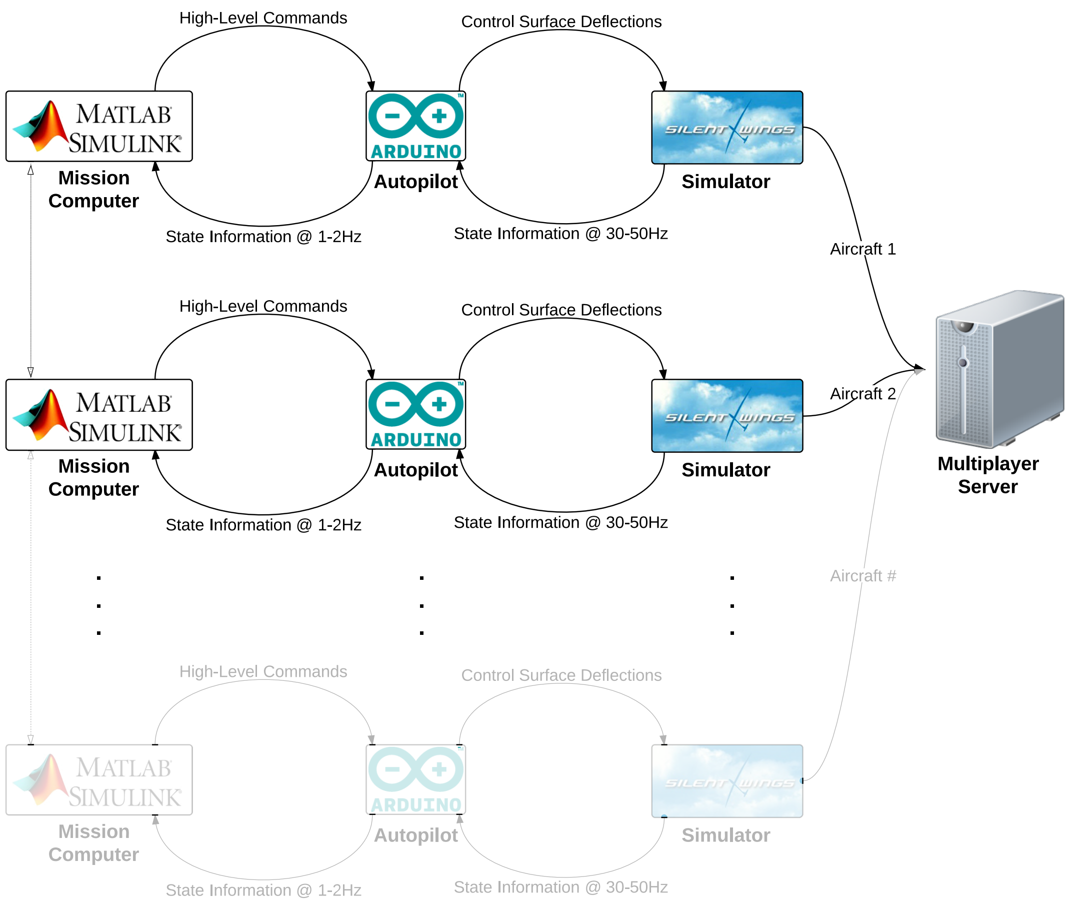

Figure 4.6. Complete testing and simulation architecture. This system can be used to test and compare an unbounded number of aircraft and corresponding control systems in the same environment by setting up a multiplayer server.

Although real flight testing was unable to be completed, the infrastructure for real autonomous flight using the SB-XC was set in place. Low-level autonomous control is handled by the commercially available Piccolo SL autopilot (by Cloud Cap Technology). The Piccolo SL (shown in Figure 4.7) has the ability to track the following commands for fixed wing aircraft:

- Indicated Airspeed
- Altitude
- Bank Angle

- Flap Angles
- Heading
- Vertical Rate
- Pitch

in addition to a navigational mode that can use the aforementioned command loops to follow a flight plan. Various Piccolo autopilots have been used successfully for this specific application [10, 13, 22].

In order to support the estimation and control algorithms discussed throughout this document, an onboard computer system communicates directly with the Piccolo. This payload receives state information (around 25 Hz) and returns bank angle commands to the Piccolo when thermal soaring mode is engaged. Both the Piccolo autopilot and the thermal soaring computer are controlled and monitored from the ground station computer via a 900 MHz frequency hopping spread spectrum (FHSS) radio system. Additionally, a radio transmitter connected to the ground station maintains the ability for a pilot takeover in case of an emergency. This system is illustrated in Figure 4.8.

Figure 4.7. Piccolo SL autopilot by Cloud Cap Technology

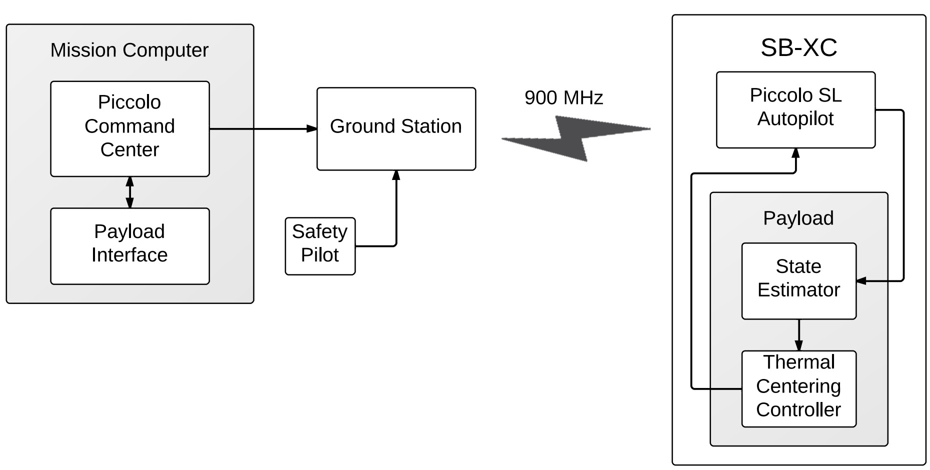

Figure 4.8. Structure of the flight testing system.

The Piccolo Command Center and Payload Interface GUIs can be seen in Figures 4.9 and 4.10 respectively.

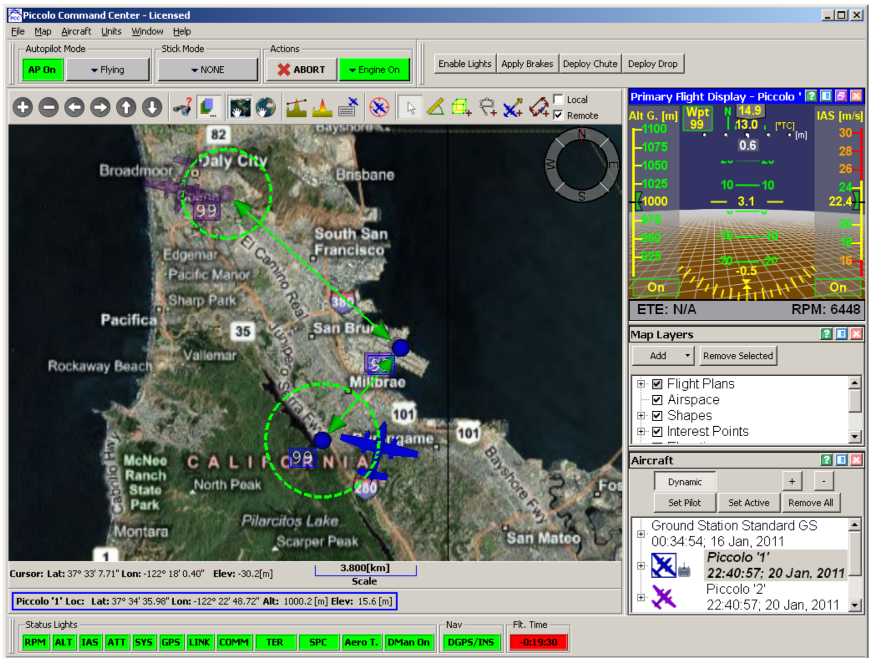

Figure 4.9. A screenshot of Piccolo Command Center. This software is used to monitor and manage the Piccolo autopilot.

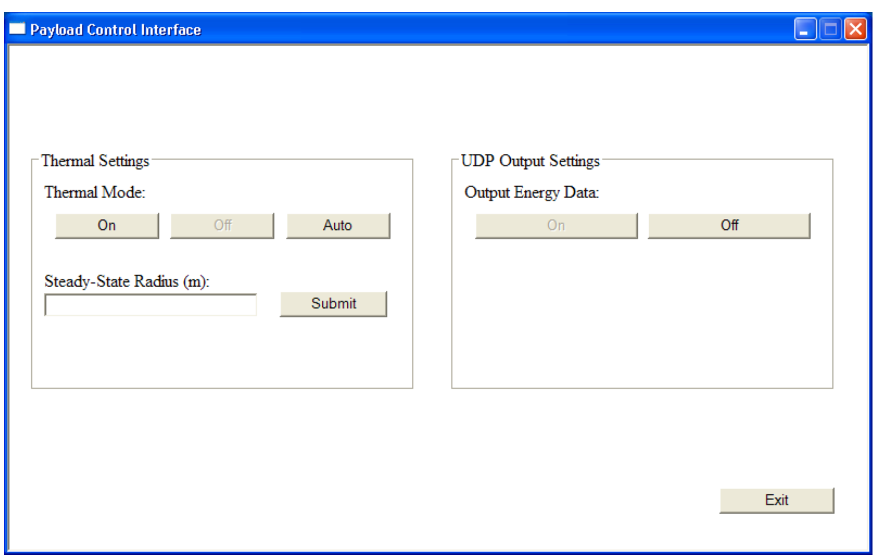

Figure 4.10. Graphical user interface created for interaction with the payload from the ground station computer.

4.3 Low-Level Control

The purpose of the low-level aircraft controller is to drive various vehicle states to a desired value. Using commanded values and an estimate of the vehicle states, the low-level autopilot sends commands to the actuators to force the states to match their corresponding desired values. This state feedback occurs at every time step, controlling the motion of the aircraft. A single-input single-output (SISO) feedback system is shown in Figure 4.11.

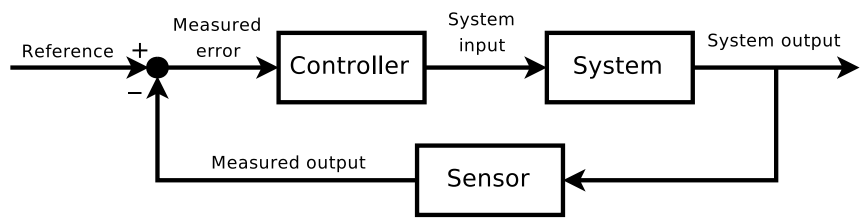

Figure 4.11. Feedback control of a SISO system [9].

The difference between the state estimate and the commanded value is fed through a compensator. The purpose of the compensator is to amplify the error signal and potentially filter an undesired part of the signal such as high frequency sensor noise. This signal is then used as the input to the plant system; in our case this is the aircraft. The aircraft reacts to these inputs and sensors once again measure the state of the system. This process is repeated at every time step driving the measured error to zero.

The control loops in our specific autopilot implementation use PID compensators. PID compensators apply tunable gains to the proportional, integral, and derivative values of the error signal. This approach attempts to eliminate past (steady-state) error, current error, and future (overshoot) error. The robustness of PID compensators make them an ideal choice for our autopilot application. Using a model of the aircraft, the PID gains can be analytically determined based on desired response characteristics. However, for an aircraft model oftentimes the transfer functions are very complex, especially for control inputs to higher-level states such as airspeed. This requires a significantly more complicated compensator to meet the response requirements. This problem is simplified using successive loop closure.

Successive loop closure consists of closing several feedback loops around the open loop plant. This enables the design of multiple, much simpler compensators due to the dissection of the aircraft dynamics into first and second order transfer functions [30]. Since the longitudinal dynamics of an aircraft are decoupled from the lateral/directional modes, each decoupled mode can be controlled using a set of loop closures. The succession of loop closures for the longitudinal and lateral/directional dynamics can be seen in Figure 4.12.

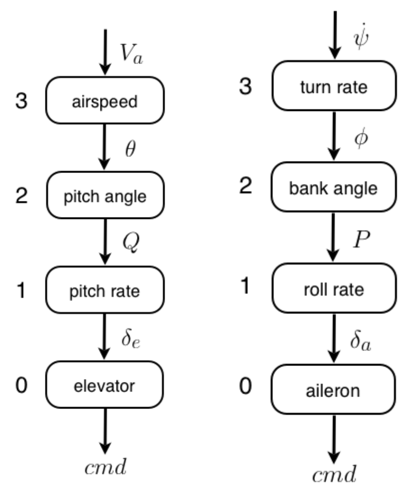

(a) Longitudinal Control loops.

(b) Lateral/ Directional Control Loops.

Figure 4.12. Sequences of control loop closures used for low-level aircraft control.

4.4 Simulation 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.

4.4.1 Centering Performance

The performance of the thermal centering controller (Equation 2.1 in Section 2.2.4) 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.

The nearly exclusive energy estimator that has been used for thermal soaring with the SB-XC can be seen in Figure 4.13. This estimator was designed by Allen and Lin [10] and uses a series of filters to reject noise and differentiate the energy signal.

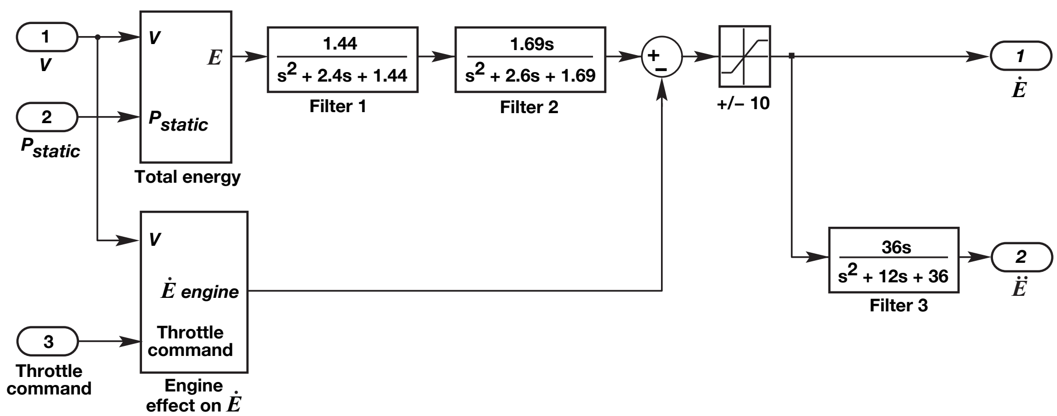

Figure 4.13. Energy estimator designed by Allen and Lin [10]

In a paper by Andersson published in 2012, there is mention of adapting the thermal centering controller to accept an $\dddot{E}$ term in an attempt to reduce latency [25]. However, this was discussed as potential future work and an analysis has yet to be completed showing the efficacy of this controller. For the sake of comparison, the energy estimator shown in Figure 4.13 was used as a benchmark for the following thermal soaring simulations.

The following analysis was completed to examine the estimator's ability to enable thermal centering in a wide range of thermals. The thermals were generated by Silent Wings and resemble a Laplacian of Gaussian surface. An example thermal profile can be seen in Figure 4.14.

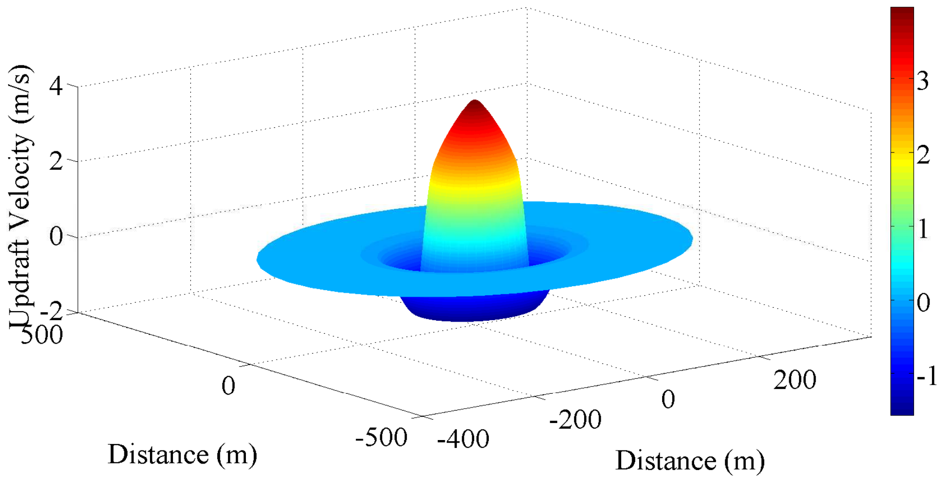

Figure 4.14. Thermal generated by Silent Wings ($R_{therm} = 80m$, $w_{core} = 4m/s$).

Thermals ranged from 55 to 105 m radius and core strengths from 2 to 6 m/s. Figure 4.15 shows the initial behavior of the SB-XC as it encounters 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.

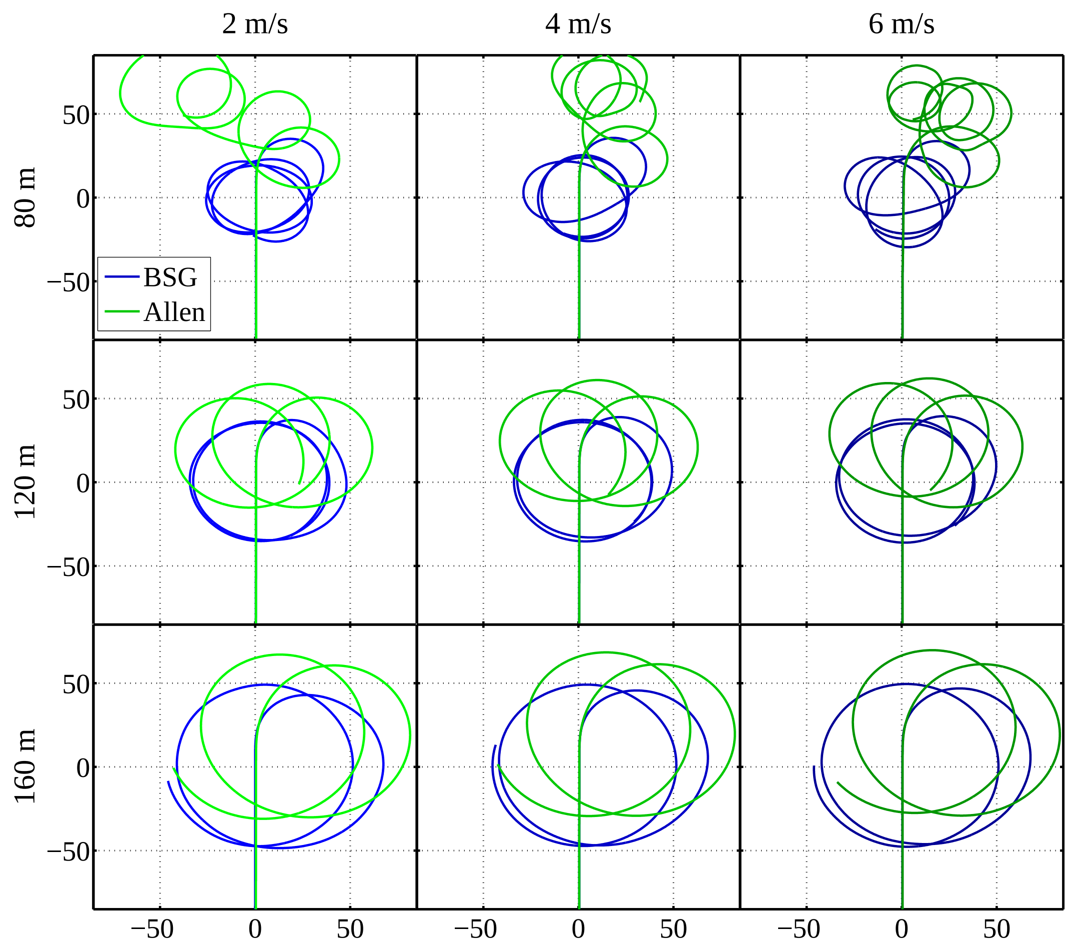

Figure 4.15. Initial centering response of the SB-XC model to various thermals.

The corresponding data for each of these simulations can be found in Tables 4.1, 4.2, and 4.3.

Table 4.1. 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 4.2. 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 4.3. 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 55 m radius thermals. Figure 4.15 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 degradation in centering capability seems to be caused by latency in the energy estimates thus causing the controller to react out of phase. Allen's energy filter is discussed in the AIAA publication *Cooperating UAVs Using Thermal Lift to Extend Endurance* where it is stated that "with a delay of around 2.5 to 3 s. the centering still functions well for bigger and stronger updrafts, but the performance will degrade when the size and strength of the thermal decreases." However, with the BSG filtering method the latency seems much more manageable, allowing the aircraft to successfully track smaller thermals. When the aircraft was able to successfully center a thermal, the average climb rate was nearly the same with the BSG averaging only 0.6% increase in climb rate over Allen's energy filter.

The next simulation tests the estimator/controller's ability to reject disturbances. The motion of wind in the atmosphere is extremely random; this unpredictability must be filtered by the energy estimator. This simulation imposes disturbances in the form of thermal drift and change of wind direction.

The lateral directions of the disturbances are illustrated in Figure 4.16. The rapid change of wind direction exemplifies possible wind shear situations while simulating thermal drift between the nodes. These disturbances alter the thermal direction and the aircraft's flight path through the airmass. Although it is extremely unlikely that this scenario would occur naturally, it magnifies the estimator's response to disturbances. 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 a Laplacian of Gaussian 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 4.17 and the corresponding altitude plot is shown in Figure 4.18.

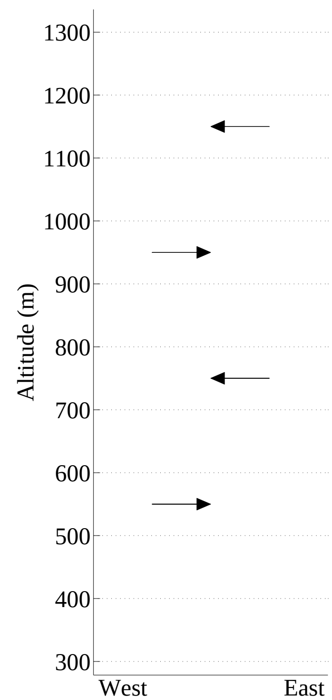

Figure 4.16. Cross-wind altitude locations.

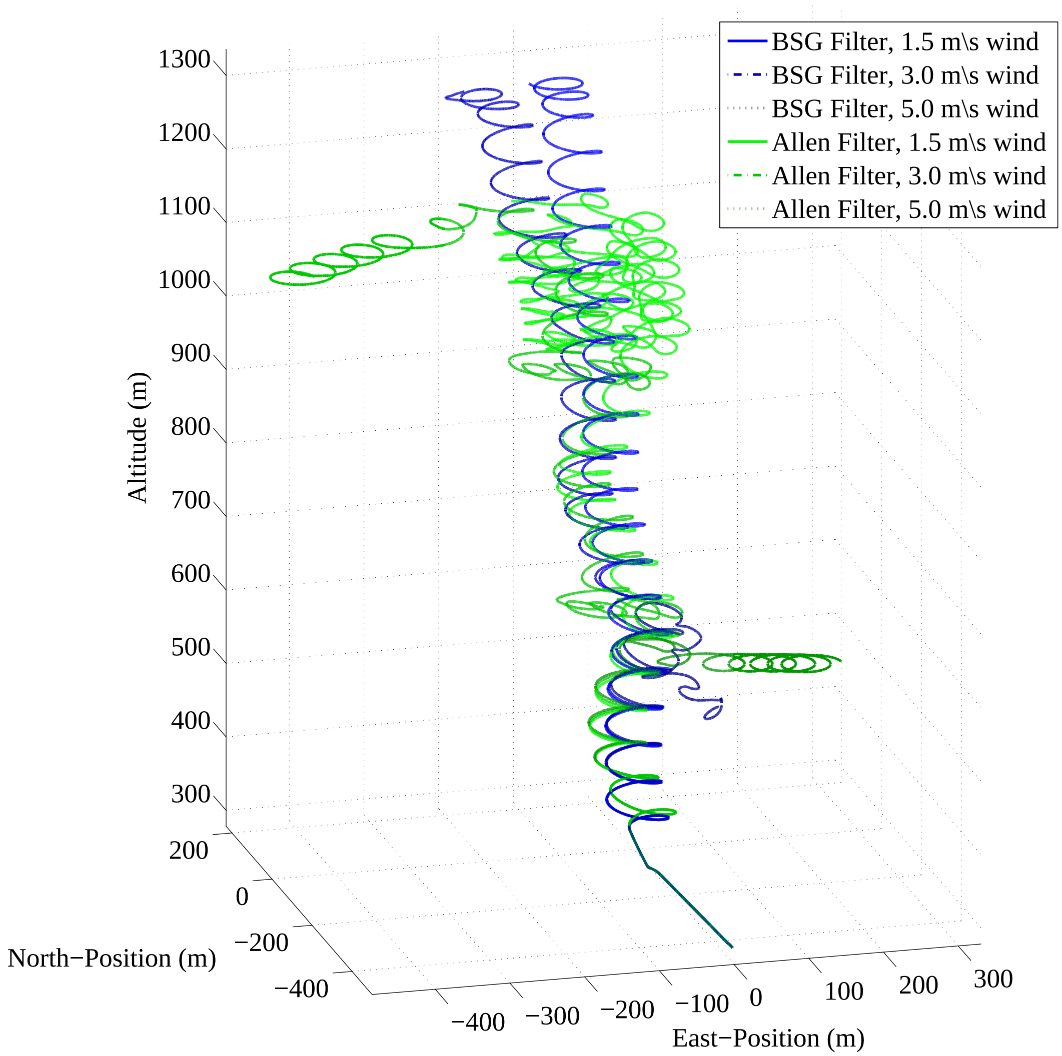

Figure 4.17. Path of the SB-XC in a thermal subjected to various cross-winds.

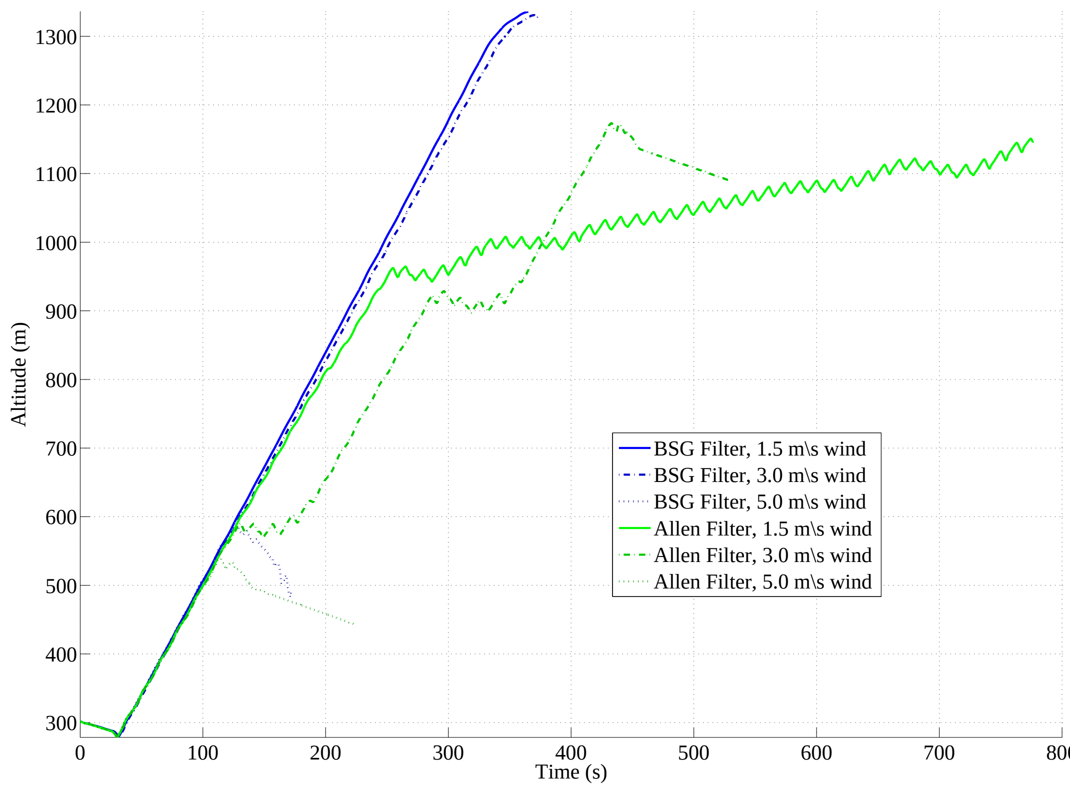

Figure 4.18. Altitude of SB-XC in a thermal subjected to various cross-winds.

Table 4.4. Results of SB-XC subjected to various 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

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.

4.4.2 Extremum Seeking Control

This section shows the results of performance testing of the extremum seeking controller (discussed in Section 2.4). This is a high-level controller that 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. Using the backward Savitzky-Golay energy filter (Section 3.2.3) and the centering controller presented in Section 2.2.4, simulations were performed in Silent Wings soaring simulator.

These simulations exposed the SB-XC aircraft to various thermal sizes ranging from an 80 to 200 m radius. 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 4.19 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 4.5 shows the corresponding average climb rate values.

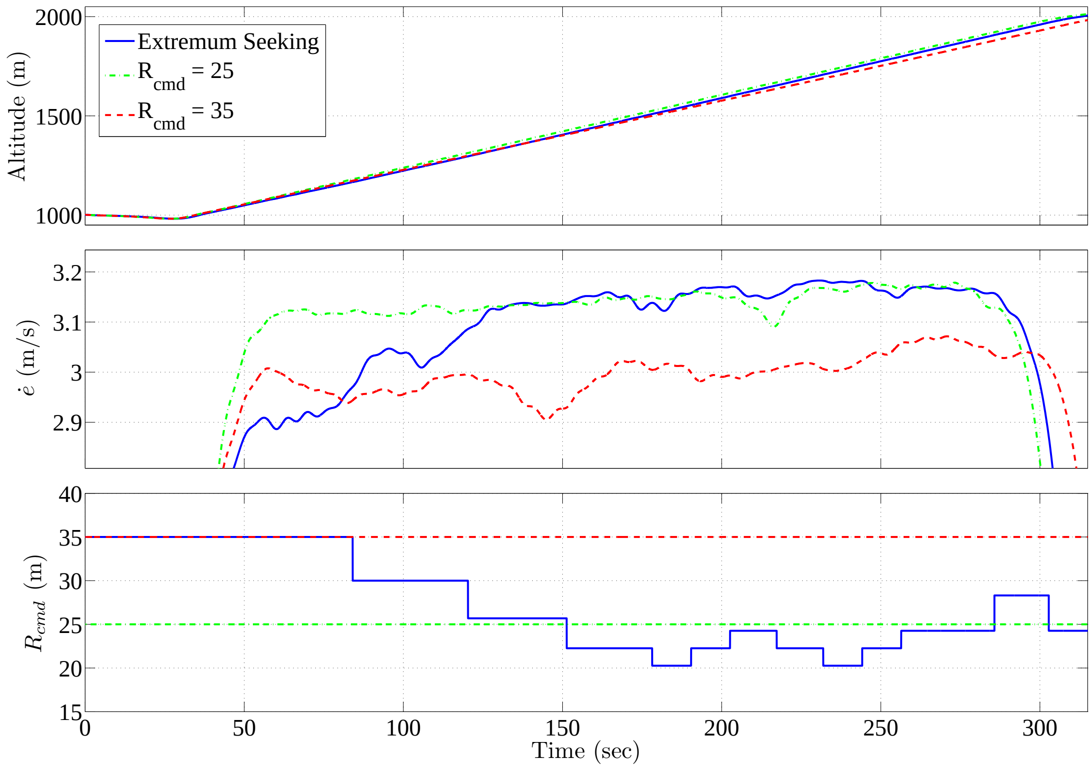

Figure 4.19. Extremum seeking control in a thermal ($w_{core} = 4$ m/s, $R_{th} = 120$ m).

Table 4.5. 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 m	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's $\dot{e}$ is larger than that of the 25 m case. 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, without previous knowledge of the thermal there is no way to know what the optimal value is. Despite the extremum seeking controller starting with an initial turn radius of 35 m, after the (~1,000 m) climb it outperformed the 35 m case by almost 4%. Figure 4.20 and Table 4.6 show results for flight within a wide thermal with a slightly weaker core (w_{core} = 3 m/s, R_{th} =200 m).

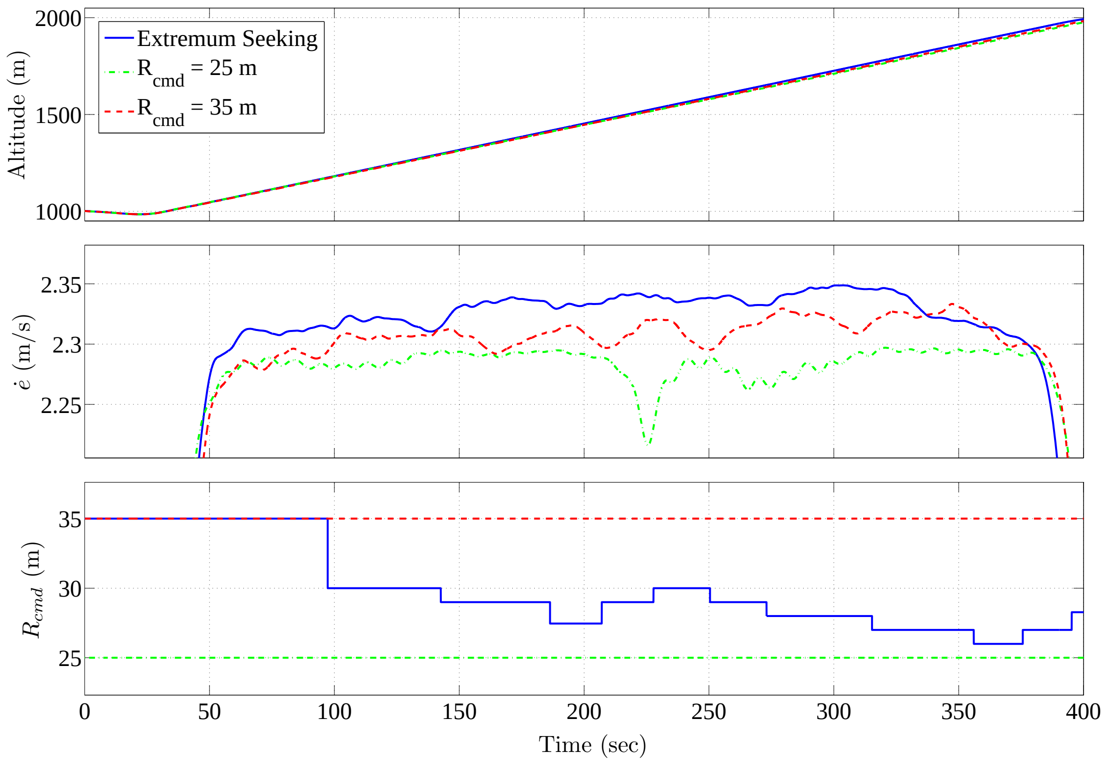

Figure 4.20. Extremum seeking control in a thermal (w_{core} = 3 m/s, R_{th} =200 m).

Table 4.6. Average climb rates (w_{core} = 3 m/s, R_{th} = 200 m).

Radius Cmd (m)	**Avg ROC (m/s)**
Extremum Seeking	2.72
R = 25 m	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 climb. Performance gains 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 m/s, R_{th} =80 m). These results are shown in Figure 4.21 and Table 4.7.

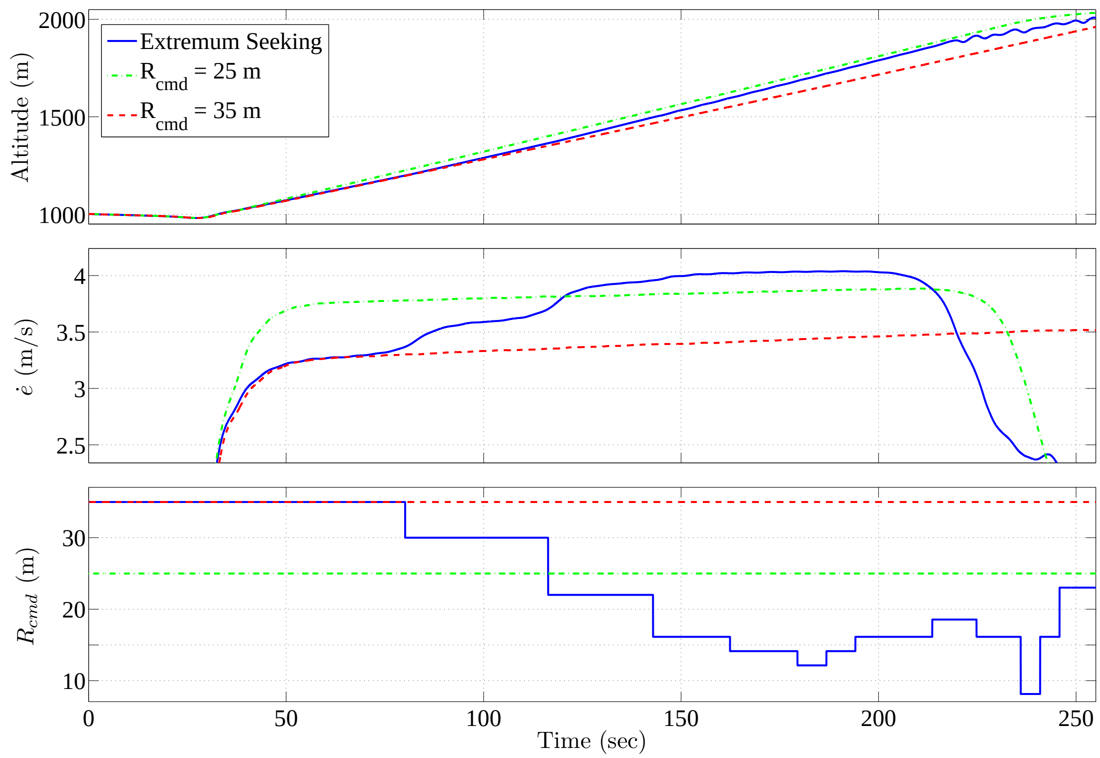

Figure 4.21. Extremum seeking control in a thermal (w_{core} = 5 m/s, R_{th} =80 m).

Table 4.7. 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 m	4.88
R = 35	4.33

This simulation initially shows rapid convergence of the turn radius. 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.

Chapter 5

Conclusion

This thesis is motivated by the potential that small, unmanned soaring aircraft have to efficiently accomplish various civil and scientific missions. The utility of these vehicles can be significantly improved by increasing their range and endurance. Autonomous soaring in thermal updrafts has been demonstrated as an effective method of atmospheric energy harvesting. However, the drawbacks of these thermal soaring algorithms have prevented small, soaring UAVs to reach their full potential. This thesis focused on improving autonomous soaring performance by the reduction of latency in energy estimates and the development of an extremum seeking controller that optimizes the climb rate in a given thermal.

Silent Wings, a high-fidelity, commercially available flight simulator enabled the testing of the algorithms proposed in this thesis. An aerodynamic model of the SB-XC glider was constructed in Silent Wings and validated using published flight test data. Control calculations were performed on an Arduino-based autopilot that was placed in loop with the simulator via UDP connections. This environment was used as the primary testing unit and allowed the comparison of proposed methods with current capabilities.

Thermal centering performance dictates the range of thermal sizes that can be used as well as what level of disturbances can be rejected without losing the thermal. The centering controller relies on the second derivative of energy as a feedback term. This is problematic because numerically differentiating sensor data amplifies noise significantly. Current estimation methods handle this using a series of low-pass filters. Consequentially, significant latency is introduced into the energy estimates; this has been recognized as the primary hindrance to thermal centering capability.

This problem was handled by adapting a Savitzky-Golay filter to produce minimal lag while accurately estimating the signal derivatives. This was accomplished by fitting a polynomial of a desired order to a time window of data up to the current time step. The polynomial acts a smoothing filter while providing smooth derivative estimates based on the coefficients of the polynomial fit. Using a backward version of this filter provides an estimate of energy states at the current time step with significantly less lag. Due to the asymmetry of this filter, a larger time window provides more smoothing at the expense of increased lag. By choosing a good window size and polynomial fit order, it was demonstrated that good estimates can be acquired with significantly less latency.

Significant improvements were observed in centering time and accuracy, the range of thermals that can be tracked, and rejection of disturbances such as wind gusts and sensor noise . In a series of 9 simulations, the aircraft was introduced to a wide range of thermal sizes and shapes. The proposed estimator was able to maintain thermal lock in all of these simulations whereas the filtering method used by Allen enabled only the largest 6 to be utilized for thermal soaring. The proposed estimator also centered thermals much more accurately; usually within a turn and a half from the time the centering algorithm was triggered. Additionally, enhanced disturbance rejection was observed while encountering various wind shear conditions. Of the three cases explored, the proposed estimator was able to maintain thermal lock on two of them whereas Allen's estimator lost the thermal on all three occasions. Reduced latency in energy estimates proved to be a valuable contribution to the performance of autonomous soaring aircraft.

In addition to thermal tracking, optimal trajectory planning also has the potential to increase soaring performance. Path planning methods for optimal thermal soaring via thermal mapping have been recently explored. However, these methods haven't been effective in solving the optimization problem because they require significant exploration time to build an initial map. Commanding a constant steady-state turn radius typically outperforms the thermal mapping algorithm in single-peak thermals because of the it's ability to start climbing immediately upon detection.

This thesis introduced an adaptive extremum seeking control algorithm that promotes the convergence of the steady-state turn rate to the optimal value. This method uses past energy gradient information to predict the optimal turn rate and continuously perturbs this turn rate to track changes in thermal structure thus adapting to the new optimal turn rate. In simulation this controller converged quickly and seemed to track the optimal turn conditions exceedingly well. Up to a 9% increase in climb rate was observed during these simulations. The largest benefit of this algorithm is it's ability to climb well in a wide range of thermal structures. However, slight instability was observed in exceedingly narrow but strong thermals. This instability was detected in thermals that were significantly out of the range of previous thermal centering algorithms and can be easily avoided by redetermining the turn limits for future cases.

5.1 Summary of Contributions

5.1.1 Energy State Estimator

An energy estimation scheme was developed that uses an asymmetric Savitzky-Golay filter to estimate the first and second derivatives of total vehicle energy. The energy estimates are used as a feedback term to initially center and maintain course within a thermal. This estimator significantly reduces latency in energy estimates compared to previous estimation methods. This provides better disturbance rejection and thermal tracking, increased centering stability, and expands the range of thermals that can be utilized.

5.1.2 Extremum Seeking Controller

A high-level controller was developed that generates steady-state turn commands that seek the optimal energy conditions in a given thermal. The thermal centering controller uses these commands as a reference to track while maintaining centered about the thermal. Using the extremum seeking controller, as opposed to commanding a constant steady-state turn rate, increased robustness was observed as well as improved climb rate without prior knowledge or mapping of the thermal.

5.1.3 Performance Testing and Comparison

A high-fidelity simulator was used to analyze the performance of the system and verify initial predictions. Additionally, these simulations allowed the comparison of previous methods with the methods described in this thesis. This provided performance quantification of the proposed methods.

5.2 Recommendations for Future Work

5.2.1 Stability Analysis for the Extremum Seeking Controller

While the extremum seeking controller demonstrated promising convergence results, it showed stability issues for narrow thermals with exceptionally strong core updraft velocities. An example of this occurrence can be seen in Figure 4.7. As mentioned in Section 4.4.2, this instability was 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).

The simple remedy would be to impose more restrictive limits on bank angle commands. However, a stability analysis would likely reveal the root cause of the issue as well as provide insight into a enhancements that could be made to improve stability and overall performance of the controller. A direct impact of this analysis would allow a more calculated determination of the controller gain (shown in Equation 2.30) which affects the size of the radius perturbation.

5.2.2 Flight Testing and Verification

Although the methods proposed in this thesis were validated in simulation, actual flight testing will provide conclusive insight into their validity. Flight testing will expose the controllers to actual thermal conditions and define their true utility. Additionally, implementation on a real aircraft will certainly provoke any problems that might have been overlooked, principally any hardware related issues.

Flight testing infrastructure was set up to run these tests including: set-up of the Piccolo autopilot for use with the SB-XC; development of an on-board computer system to perform the energy estimation and thermal centering algorithms in real time; creating an interface to communicate with the on-board payload; and running test simulations in Piccolo Command Center . However, acquiring the necessary approval to test autonomously controlled aircraft combined with time limitations prevented actual flight testing to be conducted.

Appendix A

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 A.1. Parameters for SB-XC glider.

variable	**value**	**description**
m	5.7 kg	mass
S	0.996 m^2	wing area
$f(C_L)$	$0.0166{C_L}^3 + 0.0535{C_L}^2 - 0.0437C_L + 0.0276$	
$v_{a,min}$	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 $\frac{m/s}{v_a}$

Bibliography

[1] (2013), "Red-tailed Hawk," Online.
URL http://www.allaboutbirds.org/guide/Red-tailed_Hawk/id

[2] (2013), "Soaring," Online.
URL http://soaring.eu/wp-content/uploads/2011/10/Uys-1.jpg

[3] (2013), "Iran claims it captured another U.S. Raven drone," Online.
URL http://www.worldtribune.com/2013/02/01/iran-claims-it-captured-another-u-s-raven-drone/

[4] Administration, F. A. (2007) *Glider Flying Handbook*, Skyhorse Publishing, Inc.

[5] Reichmann, H. (1978) *Cross-Country Soaring*, Thomson Publications, Santa Monica, California.

[6] Andersson, K., I. Kaminer, K. D. Jones, V. Dobrokhodov, and D.-J. Lee (2009) "Cooperative UAVs Using Thermal Lift to Extend Endurance," in *AIAA Infotech@Aerospace Conference*, AIAA Paper 2009-2043, American Institute of Aeronautics and Astronautics, Seattle, Washington.

[7] Chan, C. D. N. (2012), "File:Savitzky-golay pic gaussien bruite.svg," .
URL http://en.m.wikipedia.org/wiki/File:Savitzky-golay_pic_gaussien_bruite.svg

[8] Edwards, D. (2008) "Performance Testing of RNR's SBXC Using a Piccolo Autopilot," *RC Cross Country Soaring*.

[9] (2008), "Feedback loop with descriptions.svg," .
URL http://en.wikipedia.org/wiki/File:Feedback_loop_with_descriptions.svg

[10] Allen, M. J. and V. Lin (2007) "Guidance and Control of an Autonomous Soaring Vehicle with Flight Test Results," in *AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada.

[11] Utley, B. (2006) "Steve Fossett Does it Again - Glider World Altitude Record," *Technology Evangelist*.

[12] Edwards, D. J. (2008) "Implementation Details and Flight Test Results of an Autonomous Soaring Controller," in *AIAA Guidance, Navigation and Control Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia.

[13] Andersson, K., I. Kaminer, and K. D. Jones (2010) "Autonomous Soaring; Flight Test Results of a Thermal Centering Controller," in *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2010-8034, American Institute of Aeronautics and Astronautics, Toronto, Canada.

[14] Arho, R. (1974) "Optimal Dolphin Soaring as a Variational Problem," in *OSTIV Publication XIII*, Organisation Scientifique et Technique Internationale du Vol à Voile.

[15] de Jong, J. L. (1981) "The Convex Combination Approach: A Geometric Approach to the Optimization of Sailplane Trajectories," in *OSTIV Publication XVI*, Organisation Scientifique et Technique Internationale du Vol à Voile, pp. 182–201.

[16] Cochrane, J. H. (1999) "MacCready Theory with Uncertain Lift and Limited Altitude," *Technical Soaring*, **23**(3), pp. 88–96.
URL `http://faculty.chicagogsb.edu/john.cochrane/research/Papers/newmcred.pdf`

[17] MacCready Jr., P. B. (1958) "Optimum Airspeed Selector," *Soaring*, pp. 10–11.

[18] Sander, G. and F. X. Litt (1979) "On Global Optimal Sailplane Flight Strategy," in *NASA Conference Publication 2085, Science and Technology of Low Speed and Motorless Flight* (P. W. Hanson, ed.), NASA, Hampton, Virginia, pp. 355–376.

[19] Wharington, J. (1998) *Autonomous Control of Soaring Aircraft by Reinforcement Learning*, Ph.D. thesis, Royal Melbourne Institute of Technologu, Melbourne, Australia.

[20] Wharington, J. and I. Herszberg (1998) *Optimal Semi-Dynamic Soaring*, *Tech. rep.*, Royal Melbourne Institute of Technology.

[21] Allen, M. J. (2005) "Autonomous Soaring for Improved Endurance of a Small Uninhabited Air Vehicle," in *43rd AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada.

[22] Edwards, D. J. and L. M. Silverberg (2010) "Autonomous Soaring: The Montaque Cross-Country Challenge," *Journal of Aircraft*, **47**(5), pp. 1763–1769.

[23] Andersson, K. and I. Kaminer (2009) "On Stability of a Thermal Centering Controller," in *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2009-6114, American Institute of Aeronautics and Astronautics, Chicago, Illinois.

[24] Piggott, D. (2002) *Gliding: A Handbook on Soaring Flight (Flying and Gliding)*, 8 ed., A & C Black.

[25] Andersson, K., I. Kaminer, V. Dobrokhodov, and V. Cichella (2012) "Thermal Centering Control for Autonomous Soaring: Stability Analysis and Flight Test Results," *Journal of Guidance, Control, and Dynamics*, **35**(3), pp. 963–974.

[26] Bird, J. and J. Langelaan (2012) "Spline Mapping to Maximize Energy Exploitation of Non-Uniform Thermals," in *Proceedings of the XXXI OSTIV Congress*, Organization Scientifique et Technique du Vol a Voile.

[27] Chakrabarty, A. and J. W. Langelaan (2011) "Energy-Based Long-Range Path Planning for Soaring-Capable Unmanned Aerial Vehicles," *Journal of Guidance, Control, and Dynamics*, **34**.

[28] Savitzky, A. and M. J. E. Golay (1964) "Smoothing and Differentiation of Data by Simplified Least Squares Procedures," *Analytical Chemistry*, **36**, pp. 1627–1639.

[29] Simon, D. J. (2006) *Optimal State Estimation*, 1st ed., John Wiley & Sons, Inc., Hoboken, New Jersey.

[30] Beard, R. W. and T. W. McLain (2012) *Small Unmanned Aircraft: Theory and Practice*, Princeton University Press.