

The Pennsylvania State University
The Graduate School

**A METHOD FOR THE EXPLOITATION OF MICROLIFT FOR
UNMANNED AERIAL VEHICLES**

A Thesis in
Aerospace Engineering
by
Nathan A. Kimmel

© 2019 Nathan A. Kimmel

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

December 2019

The thesis of Nathan A. Kimmel was reviewed and approved* by the following:

Jacob W. Langelaan

Associate Professor of Aerospace Engineering

Thesis Advisor

Mark Maughmer

Professor of Aerospace Engineering

Amy Pritchett

Professor of Aerospace Engineering

Head of the Department of Aerospace Engineering

*Signatures are on file in the Graduate School.

Abstract

Disorganized convective structures in the lower part of the atmospheric boundary layer (sometimes called microlift) are recognized as a viable source of energy for manned ultralight sailplanes and high performance hang gliders, and species of soaring birds have been observed to make use of microlift. However, this source of energy remains largely underutilized by autonomous soaring aircraft, and the techniques used to (often incidentally) exploit it are not specifically designed for microlift soaring. Analyzing surveys of the atmosphere and large eddy simulations of the convective boundary layer indicate the presence of long, linear regions of lift that are a potential significant source of energy and present the opportunity for energy extraction without the need to loiter.

This thesis investigates techniques for energy extraction from these structures - referred to as spoke-like structures, thermal walls, thermal streams, or thermal stands - that can be implemented onboard an autonomous microlift soaring aircraft. The primary focus of this thesis is a Kalman filter designed to estimate the location, orientation, and characteristics of a thermal strand. This estimator was then integrated into a flight controller and tested in a realistic wind field obtained from a large eddy simulation of the convective boundary layer.

The choice of a Kalman filter is supported in this thesis by showing that several basic techniques for following a thermal strand are unstable. Furthermore, an idealized model of a thermal strand is developed for estimator and controller design, and various methods of estimating the initial thermal strand state for the Kalman filter are developed and tested in simulation.

Table of Contents

List of Figures	vii
List of Tables	xiv
List of Symbols	xv
Acknowledgments	xvii
Chapter 1	
Introduction	1
1.1 Overcoming the Limitations of Small Unmanned Aerial Vehicles	1
1.2 Atmospheric Energy Extraction	3
1.2.1 Static Soaring	4
1.2.1.1 Thermal Soaring	5
1.2.1.2 Ridge Soaring	6
1.2.1.3 Cloud Street Soaring	6
1.2.2 Dynamic Soaring	7
1.2.3 Gust Soaring	9
1.3 Microlift Soaring and Thermal Strands	10
1.3.1 Why is it Called Microlift?	12
1.3.2 Observations of Birds	14
1.3.3 Characteristics of Thermal Strands	14
1.3.4 Thermal Strands in Atmospheric Simulation	15
1.4 Characteristics of Microlift Soaring Aircraft	17
1.5 Contributions	20
1.6 Reader's Guide	21

Chapter 2

Problem Statement and

Biological Inspiration 22

2.1	The Basics of Soaring Flight	23
2.1.1	The Sink Polar	23
2.1.2	Maximum Glide	25
2.1.3	MacCready and Speed to Fly	27
2.2	Simulating the Convective Boundary Layer	29
2.2.1	Large Eddy Simulation Environment	29
2.2.2	The Idealized Thermal Strand	31
2.3	Vehicle Modeling	31
2.3.1	Definition of Simulation Domain and Reference Frames	34
2.3.2	Vehicle Sensor Modeling	34
2.3.3	Vehicle Dynamics Modeling	37
2.3.4	Integration of the Equations of Motion	40
2.4	Flight Performance of Soaring Birds	43
2.5	Choosing a UAV for Microlift Soaring	52

Chapter 3

Tracking a Thermal Strand 55

3.1	The Challenges of Tracking a Thermal Strand	55
3.2	Determining Likely Approach Angles	61
3.3	Two Unstable Methods of Thermal Strand Following and a Stabilization of One Method	63
3.3.1	Instability of Direct Feedback Methods	63
3.3.2	Instability of Bank Angle Switching	65
3.3.3	Discrete Variation of Heading and a Stabilization of the Method	70
3.4	Estimating Strand Properties Using an Unscented Kalman Filter	73

Chapter 4

Results 79

4.1	Configuration of the Idealized Thermal Strand Simulation	80
4.2	Tracking Along an Ideal Strand	81
4.2.1	Flying Down the Strand Centerline Versus S-Curves Across the Strand	85
4.3	Solving the Initialization Problem	89
4.3.1	Naïve or Arbitrary Initialization	91
4.3.2	Dynamic Control of Maximum Bank Angle and Approach Angle and the Addition of Variable Airspeed	93
4.3.3	Left or Right	96

4.3.4	Left, Right, or Straight Ahead	97
4.3.5	Fitting Using 4 Averaged Measurements	99
4.3.6	Fitting Using 8 Averaged Measurements	101
4.3.7	Identifying Initial State Estimates Unlikely to Result in a Successful Track of the Strand	104
4.3.8	Addition of Initialization Rejection and Reinitialization	110
4.4	Testing in LES	112
Chapter 5		
	Conclusion	121
5.1	Future Work	122
Bibliography		124

List of Figures

1.1	Gary Osoba flying the prototype Carbon Dragon. Photo from the Carbon Dragon Technical Website (http://www.carbondragon.us/) [1]	2
1.2	The life cycle of a thermal. [2]	6
1.3	A diagram of ridge lift. [2]	7
1.4	Cloud streets as seen from the ground. [3]	8
1.5	Two dynamic soaring maneuvers for the extraction of energy from a wind shear.	8
1.6	A top-down diagram of the structure of downdrafts, thermal strands, and thermals in the atmospheric convective boundary layer.	11
1.7	John Bird flying his Glasflügel H-201 Standard Libelle near State College, Pennsylvania. Photo by Phil Chidekel.	13
1.8	Thermal strands in the large eddy simulation wind field of the atmospheric boundary layer provided by Sullivan and Patton (see reference [4] for the paper which describes the generation of this wind field). $z/z_i = 0.06$	16
1.9	Thermal strands in the large eddy simulation wind field provided by Sullivan and Patton. $z/z_i = 0.12$	17
1.10	Comparison of the minimum sink rate and approximate minimum turn radius in a 40 degree bank between the LET Blanik L-13 [5, 6], Glasflügel H-201 Standard Libelle [5, 7], and DG Flugzeugbau DG-800 (18 meter) [5, 8] sailplanes, Pro-Design Effect 38 paragliding wing [9], ATOS VR hang glider [10], Carbon Dragon ultralight glider [1], Turkey Vulture (see section 2.4), and Vulture UAV. The shaded area indicates approximately where it is practical to exploit microlift.	19
2.1	Sink polar of the Vulture UAV. Calculation discussed in Section 2.3.3.	24
2.2	Vertical grid offset in the large eddy simulation data. dx and dz are grid spacings, and k is the vertical index.	30
2.3	Cross section of a thermal strand in the LES data at an altitude of 60 meters. $\sigma_{str} = 70$ meters, $\dot{h}_{max} = 1.5$ meters per second	32

2.4	Cross section of a thermal strand in the LES data at an altitude of 80 meters. $\sigma_{str} = 42$ meters, $\dot{h}_{max} = 2.25$ meters per second	32
2.5	Cross section of a thermal strand in the LES data at an altitude of 15 meters. $\sigma_{str} = 94$ meters, $\dot{h}_{max} = 1.75$ meters per second	33
2.6	African White-Backed Vulture [11]	45
2.7	Black Vulture [12]	46
2.8	Turkey Vulture [13]	46
2.9	Estimated sink rate and lift to drag ratio of the White-Backed Vulture using Pennycuick's sink polar and the corrected sink polar	48
2.10	Estimated sink rates and lift to drag ratios of White-Backed Vultures, Black Vultures, and Turkey Vultures	51
2.11	The author launching the Vulture UAV. Photo by Jack Langelaan. . .	52
2.12	Comparison of the sink polars of the Vulture UAV, a Turkey Vulture, and the Carbon Dragon ultralight sailplane.	53
3.1	Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying across an ideal thermal strand at 90 degrees to the strand axis. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second. .	57
3.2	Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying across an ideal thermal strand at 45 degrees to the strand axis. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second. .	57
3.3	Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying along the strand axis of an ideal thermal strand. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	58
3.4	Variometer and roll disturbance measurements obtained by the simulated Vulture UAV when flying across an ideal thermal strand at 45 degrees to the strand axis. Compare to Figure 3.2. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second, $\dot{h}_{measured} = \dot{h} + N(0, 0.75^2)$, $C_{l,res,measured} = C_{l,res} + N(0, 0.035^2)$. Note that the added Gaussian noise includes contributions from both turbulence and sensor noise. .	59
3.5	Regions selected for analysis of the probability of approaching a thermal strand from a given angle. Regions are marked in white. Large eddy simulation cross-section taken at $z/z_i = 0.12$	62
3.6	Histogram of the probability of approaching a thermal strand from a given angle.	62
3.7	Simulated flight path of an aircraft with bank angle commanded using roll disturbance feedback along an ideal thermal strand. Sensor noise was not included in this simulation. Dotted lines indicate the 2σ distance from the centerline. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	64

3.8	Illustration of the instability of bank angle switching controllers when attempting to track along a thermal strand. Vehicle flight path is depicted in red, and \overline{OL} is the longitudinal axis of the strand. This depicts the case when $\overline{P_n A_n} > \overline{C_n P_{n+1}}$; however, the geometric relations shown are also valid for $\overline{P_n A_n} < \overline{C_n P_{n+1}}$	66
3.9	Simulated flight path demonstrating divergence of a bank angle switching controller with perfect identification of the strand axis. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the sign of the commanded bank angle was flipped. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	69
3.10	Simulated flight path demonstrating divergence of a bank angle switching controller with perfect identification of the strand axis but a simulated 1.5 second variometer lag. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the sign of the commanded bank angle was flipped. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	69
3.11	Simulated flight path demonstrating instability of a controller that alters heading in discrete increments with perfect identification of the strand axis. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the commanded turn direction was flipped. Red dots indicate where the heading command was incremented. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	71
3.12	Ideal track of a strand using a heading switching controller. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate when the commanded heading was flipped. Red dots indicate when the new heading command was applied. $\psi_0 = 60$ degrees, $\Delta\psi = 60$ degrees, $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	71
3.13	Ideal track of a strand using a heading switching controller. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate when the commanded heading was flipped. Red dots indicate when the new heading command was applied. $\psi_0 = 50$ degrees, $\Delta\psi = 60$ degrees, $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.	73
3.14	Definition of thermal strand parameters describing the location of the strand axis \overline{OL} relative to the aircraft and the updraft cross-section.	75
4.1	Flight path of the basic Kalman filter controller when tracking an ideal strand. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.	83

4.2	State estimate and true state of the basic Kalman filter controller when tracking an ideal strand. The true state is plotted in black and the estimated state is plotted in blue.	83
4.3	State estimate error and estimated $\pm 2\sigma$ bounds of the basic Kalman filter controller when tracking an ideal strand. The state estimate error is plotted in blue, and the black dotted lines are the estimated $\pm 2\sigma$ bounds. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second, $\dot{h}_{measured} = \dot{h} + N(0, 0.75^2)$, $C_{l,res,measured} = C_{l,res} + N(0, 0.035^2)$	84
4.4	Partial flight path of the basic Kalman filter controller tracking an ideal thermal strand by flying S-curves during a 20 minute simulation. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.	85
4.5	Flight path of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.	86
4.6	State estimate and true state of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The true state is plotted in black and the estimated state is plotted in blue.	87
4.7	State estimate error and estimated $\pm 2\sigma$ bounds of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The state estimate error is plotted in blue, and the black dotted lines are the estimated $\pm 2\sigma$ bounds.	87
4.8	Partial flight path of the basic Kalman filter controller tracking an ideal thermal strand by flying down the centerline during a 20 minute simulation. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.	88
4.9	Percent of the time that the arbitrary initialization controller can track the strand.	92
4.10	Percent of the time that the arbitrary initialization controller can track the ideal thermal strand on the first attempt.	94
4.11	Percent of the time that the arbitrary initialization controller can track the ideal thermal strand when using the dynamic strand approach method.	95
4.12	Percent of the time that the controller which initializes assuming the strand is either to the left or right can track the ideal thermal strand.	97
4.13	Percent of the time that the 3 case controller can track the ideal thermal strand.	99
4.14	Percent of the time that the controller fitting to 4 average measurements can track the ideal thermal strand.	102

4.15	Percent of the time that the controller fitting to 8 averaged measurements can track the ideal thermal strand.	103
4.16	Scatter plot of initialization values of d when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.20	105
4.17	Scatter plot of initialization values of ψ_{str} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.	106
4.18	Scatter plot of initialization values of \dot{h}_{max} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.21	106
4.19	Scatter plot of initialization values of σ_{str} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.22	107
4.20	Scatter plot of initialization values of d when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.	108
4.21	Scatter plot of initialization values of \dot{h}_{max} when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.	108
4.22	Scatter plot of initialization values of σ_{str} when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.	109
4.23	Percent of the time that the controller fitting to 8 averaged measurements with initialization rejection can track the ideal thermal strand.	111

4.24	Distribution of local updraft encountered when flying a random heading in the large eddy simulation wind field. The mean updraft velocity is marked in gray. The minimum sink rate and sink rate at maximum lift to drag of the Vulture UAV are shown in blue and red, respectfully, as a reference.	114
4.25	Distribution of updraft encountered by the aircraft in the large eddy simulation wind field while using the thermal strand tracking controller. The mean updraft velocity is marked in gray.	114
4.26	Distribution of net energy rate of the aircraft in the large eddy simulation wind field while using the thermal strand tracking controller. The mean net energy rate is marked in gray, and the average sink rate of the Vulture UAV is shown in red as a reference.	115
4.27	3D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second.	116
4.28	2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second.	117
4.29	2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second. Large eddy simulation cross-section taken at 240 meters.	117
4.30	3D flight path of a successful through the large eddy simulation wind field flight using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.62 meters per second, and the average sink rate of the aircraft was 0.47 meters per second.	118
4.31	2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.62 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. Large eddy simulation cross-section taken at 240 meters.	119

4.32	3D flight path of a flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.18 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. . . .	119
4.33	2D flight path of a flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.18 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. Large eddy simulation cross-section taken at 140 meters.	120

List of Tables

1.1	Comparison of the timescale of various soaring techniques to the phugoid mode of the aircraft involved. (* : simulated vehicle)	4
2.1	Proportional gain values for simulating airspeed and roll response to commands	39
2.2	Characteristics of the Vulture UAV	54
4.1	Results of 2,500 5-minute simulations of the aircraft tracking an ideal thermal strand by flying S-curves.	84
4.2	Results of 2,500 5-minute simulations of the aircraft tracking an ideal thermal strand by flying down the center of the strand.	88
4.3	Configuration of the controller testing arbitrary initialization.	92
4.4	Configuration of the controller testing arbitrary initialization with dynamic approach.	95
4.5	Configuration of the controller testing initialization that assumes the strand is either to the left or the right.	96
4.6	Configuration of the controller which initializes assuming the strand is left, right, or straight ahead.	98
4.7	Configuration of the controller which initializes assuming the strand is either to the left or right.	101
4.8	Configuration of the controller which initializes using 8 averaged measurements.	103
4.9	Criteria used for accepting or rejecting initial state estimates.	110
4.10	Configuration of the controller which is able to reject initialization values	111
4.11	Configuration of the controller which is able to reject initialization values	112

List of Symbols

ϕ	Bank angle
θ	Pitch angle
ψ	Heading angle
\mathcal{R}	Wing aspect ratio
b	Wingspan
e_0	Oswald's efficiency factor
g	Standard gravitational acceleration
p	Rate of rotation about the body 1-axis
q	Rate of rotation about the body 2-axis
r	Rate of rotation about the body 3-axis
r_{turn}	Turn radius
u	Component of airspeed parallel to the body 1-axis
v	Component of airspeed parallel to the body 2-axis
w	Component of airspeed parallel to the body 3-axis
C_l	Rolling moment coefficient
C_{lp}	Roll damping of the aircraft, $C_{lp} = \frac{\partial C_l}{\partial \dot{p}}$

C_L	Lift coefficient
C_D	Drag coefficient
D	Drag force
L	Lift force
$N(\mu, \sigma^2)$	Gaussian random noise with mean μ and standard deviation σ
S	Wing area
V	Airspeed
V_0	Trim airspeed of the aircraft
W	Weight
W_{sink}	Aircraft sink rate
\dot{q}	Time derivative of quantity q
q_k	Value of quantity q at time step k
$\mathbf{q}^{\mathbf{B}}$	Vector quantity \mathbf{q} expressed in frame \mathbf{B}
DCM_{21}	Direction cosine matrix from frame 1 to frame 2
\overleftrightarrow{AB}	Line AB
\overline{AB}	Line segment AB
\widehat{AB}	Arc AB
$\angle ABC$	Angle ABC
\parallel	Is parallel to
\perp	Is perpendicular to

Acknowledgments

I would like to thank my family for all of the support they have provided. I am especially grateful for all that my parents have done, particularly the interest they have expressed in my research, their company during the nights that I worked until early morning, and most of all their unwavering support and confidence in me.

I would like to thank my advisor Jack for introducing me to autonomous soaring and suggesting this avenue of research. His advice, insights, and outside perspective helped me get past several stumbling blocks and I could not have done this without him. Thank you, Jack.

I would also like to thank my labmate and friend John Bird for his help. His knowledge of the atmosphere, soaring, and autonomy have been invaluable and I greatly appreciate his willingness to point me towards resources and to act as a sounding board for some of my ideas.

I would like to thank Dr. Mark Maughmer and Dr. Amy Pritchett for reviewing my thesis and for the insights they provided.

Finally, I would like to thank my labmates and all of my friends for their help and support. Thank you all.

Chapter 1 |

Introduction

1.1 Overcoming the Limitations of Small Unmanned Aerial Vehicles

Small unmanned aerial vehicles (UAVs) are widely used for recreational, commercial, and military operations, and the FAA predicts that over the next decade the use of UAVs will dramatically increase. However, when compared to manned aircraft the majority of small UAVs have severely limited internal energy storage capacity, which in turn limits their capability - either in terms of payload, range, endurance, or some combination of the three. While it is not always possible or practical to increase the internal energy storage of a small UAV, the application of biomimetic behaviors can allow the UAV to extract energy from its environment and extend its range and endurance. Furthermore, it is highly desirable for these behaviors to be executed entirely onboard the aircraft, as this eliminates the limitations imposed by requiring a connection to an off-board flight computer.

When attempting to extend the range and endurance of a UAV through environmental energy extraction, it makes sense to provide the vehicle with as diverse a

toolbox of energy extraction techniques as possible. Soaring birds use an extremely wide range of techniques to maintain flight including thermal soaring, slope soaring, dynamic soaring, and the exploitation of small turbulent structures [14,15]. Most of these techniques have been studied and applied to small UAVs either in simulation or on a physical vehicle; however, the exploitation of convective structures unique to the lower third of the atmospheric convective boundary layer has so far escaped study.

The majority of autonomous soaring research thus far has been focused on thermal, slope, and dynamic soaring - likely because these sources of energy can be utilized by manned sailplanes, and the techniques required to extract energy are well understood. Reichmann's book on soaring describes methods for extracting energy from each of the aforementioned types of lift [16].

Sources of lift that are too small or low for the average manned sailplane to safely exploit are termed "microlift" and are generally ignored by the soaring community. Only recently has the introduction of ultralight sailplanes allowed practical manned exploitation of microlift: most notably Gary Osoba and his Carbon Dragon (see Figure 1.1) [1, 17–19].



Figure 1.1. Gary Osoba flying the prototype Carbon Dragon. Photo from the Carbon Dragon Technical Website (<http://www.carbondragon.us/>) [1]

The area of microlift soaring has not entirely gone unnoticed by those working on

autonomous soaring; however, this research has tended to be focused on gust soaring and extracting energy from wind turbulence [20–22]. The small amount of energy that can be extracted from wind gusts places gust soaring on the "micro" end of microlift soaring.

Manned and unmanned aircraft are not the only things that make use of microlift. Many times have this author and others observed Turkey Vultures and Black Vultures soaring at various altitudes over gently rolling fields, making only slight corrections while flying a linear course, and all with no apparent loss of altitude. Investigation of the structure of the atmospheric boundary layer at low altitudes reveals the presence of long linear regions of lift [23], referred to as thermal walls or thermal strands, and it is likely that these birds are making use of these thermal strands to soar without circling. In addition, Gary Osoba has soared in microlift with birds and has noted that soaring birds "frequently and skillfully make use of these minor ... energy sources" [17].

1.2 Atmospheric Energy Extraction

In considering the various methods of extracting energy from the atmosphere that may be employed by an aerial vehicle, it is convenient to organize the techniques by comparing the timescale of the required maneuver(s) to the period of the vehicle's phugoid mode. As noted by Boslough, Hendricks showed in his doctoral thesis that the optimal frequency of the maneuver sequence required for dynamic soaring is equal to the phugoid frequency of the aircraft [24]. More recent analyses have found various other frequencies to be optimal, but these frequencies are still close to that of the phugoid mode of the aircraft. (Furthermore, Sukumar and Selig suggest that dynamic soaring maneuvers can be synchronized with phugoid oscillations and perform dynamic soaring with minimal input, and that the difference between the longitudinal phugoid

frequency determined in trim and that observed while performing dynamic soaring is the result of various nonlinear effects [25]. These effects reduce the frequency of the phugoid motion when performing dynamic soaring maneuvers.) In contrast, the sequence of maneuvers required for an aircraft to center and soar in a thermal - a full sequence determined by the circling period - tend to occur at a significantly longer timescale than that of the phugoid mode of the vehicle. In this way, thermal soaring is relatively "static" in the sense that the maneuvers occur relatively slowly. Extracting energy from atmospheric turbulence and wind gusts requires extremely quick control inputs and maneuvers that occur at a frequency higher than that of the vehicle's phugoid mode. Examples of maneuver frequency compared to the phugoid mode of the aircraft can be found in Table 1.1.

Table 1.1. Comparison of the timescale of various soaring techniques to the phugoid mode of the aircraft involved. (* : simulated vehicle)

Soaring type	Ratio of phugoid frequency to maneuver frequency	Aircraft	Source
Thermal soaring	~ 0.50	CloudSwift UAV	Allen [26]
Dynamic soaring	0.6991	2.93m RC sailplane*	Sukumar
Dynamic soaring	0.713	L-23 Super Blanik* (prescribed flightpath)	Gordon [27] and Selig [25]
Dynamic soaring	0.7559	2.93m RC sailplane* (piloted)	Sukumar and Selig [25]
Gust soaring	~ 3.9	SB-XC*	Langelaan [22]

1.2.1 Static Soaring

Static soaring encompasses soaring techniques that involve relatively slow maneuvers. There are three primary static soaring techniques: thermal soaring, ridge soaring, and cloud street soaring. While most research is focused on only one of these techniques,

Powers, Silverberg, and Gopalarathnam developed a method that can be applied to any form of static soaring. This method involves estimating the energy expected to be available in the forward arc of the aircraft and providing that information to a path planning algorithm [28].

Thermal Soaring

Thermal soaring techniques extract energy from rising columns of air, called thermals, that form in the convective atmospheric boundary layer. The life cycle of a thermal can be seen in Figure 1.2. Thermal soaring is used by soaring birds to gain energy, and has been well studied in manned and unmanned flight, as well as in the flight of soaring birds.

Thermal soaring can be divided into two main phases: seeking a thermal and centering on the thermal. Glider pilots will seek out thermals using meteorological knowledge - visually searching for developing cumulus clouds which form over thermals, and geographic features that are likely to start a thermal - and historical knowledge of the local environment - flying over areas that are known to frequently produce thermals [16]. While using computer vision to find developing cumulus clouds is a very difficult problem, research has been done on thermal search pattern optimization and identifying regions that are likely to develop thermals [29–32].

Once a thermal has been located, it is necessary to have some idea of where its center is to enable the continued extraction of energy. This can be accomplished through various techniques in manned aircraft [16] or using any number of the various algorithms developed for UAVs [33–37]. Rather than programming an algorithm or estimator, Reddy et al. used reinforcement learning to teach a small UAV to soar in thermals [38]. In addition, flocks of soaring birds and groups of glider pilots will work together to find and soar in thermals [16,39], and cooperative soaring techniques have

been successfully applied to UAVs [30].




Figure 1.2. The life cycle of a thermal. [2]

Ridge Soaring

Ridge soaring, also called slope soaring, extracts energy from air that is forced upwards to clear a ridge or similar obstruction as it is pushed along by the prevailing wind [16]. Glider pilots frequently make use of ridge lift by flying parallel to the ridge, staying in the region of rising air, and papers have been published on the optimal maneuver to extract energy from ridge lift [40]. In addition, soaring birds have been noted to make use of ridge lift [41]. Chakrabarty and Langelaan developed a path-planning algorithm for UAVs that made use of ridge soaring [42, 43].

Cloud Street Soaring

Cloud streets are long lines of thermals aligned parallel to the wind, and appear as parallel lines of cumulus clouds, and range from fully continuous to moderately broken. Reichmann notes that cloud streets provide excellent opportunities for long range sailplane flights [16]. Similar to ridge soaring, cloud street soaring allows for the




Figure 1.3. A diagram of ridge lift. [2]

extraction of energy from the atmosphere while maintaining a roughly linear flightpath. Sailplane pilots have developed mathematical equations describing optimal flight both along cloud streets and when the desired flightpath crosses the cloud streets based on MacCready's speed-to-fly theory [16, 44]. A series of cloud streets can be seen in Figure 1.4.

1.2.2 Dynamic Soaring

Dynamic soaring involves using a vertical wind shear to increase the total energy of the vehicle. Various analysis have been conducted in an effort to determine the optimal maneuver for a given set of conditions [45–48]. In general, the maneuver required to extract energy involves repeatedly crossing the vertical wind shear - each time increasing the airspeed of the vehicle. Two dynamic soaring maneuvers are illustrated in Figure 1.5.

Dynamic soaring is used by birds, radio controlled aircraft pilots, and sailplane



Figure 1.4. Cloud streets as seen from the ground. [3]




Figure 1.5. Two dynamic soaring maneuvers for the extraction of energy from a wind shear.

pilots to gain energy from the atmosphere. In addition, dynamic soaring controllers for UAVs have been designed and tested in simulation [47–49]. Albatrosses have been observed to perform dynamic soaring using the vertical shear layer that forms close to the surface of the ocean, which allows them to fly for long distances without ever needing to flap [16, 50]. Pilots of radio controlled airplanes will use the strong vertical wind shear generated on the leeward side of a mountain ridge or peak to accelerate their unpowered aircraft to extremely high speeds - the current record of 545 miles per hour having been set in June 2018 [51].

Manned sailplanes can also make use of dynamic soaring, albeit at higher altitudes, to extract energy and maintain flight [52]. Reichmann relates the experiences of one sailplane pilot who was able to maintain altitude by exploiting a wind gradient of 40 kilometers per hour over 100 meters. However, the pilot and his aircraft were gradually pushed downwind and had to stop after 20 minutes before they drifted too far from the airport [16]. More recently, Gary Osoba has successfully performed dynamic soaring in a manned glider, and has published several articles on the subject [17]. In 2006, Gordon published a thesis on optimal dynamic soaring with manned sailplanes in which he began with a computational analysis, moved onto tests in a flight simulator, and then flight tests using a LET L-23 Super Blanik sailplane [27]. Gordon was able to conclusively demonstrate that a full-size sailplane can extract energy from the atmosphere by performing dynamic soaring maneuvers through a wind shear, although the glide performance of the L-23 Super Blanik was not sufficient to allow a net gain of energy after completing the maneuver. Kiceniuk's paper on manned sailplane energetics [53] describes how to obtain energy wind shear in various directions - both vertical and horizontal - and encompasses both dynamic soaring and exploitation of steady wind shear and gust soaring and the exploitation of dynamic wind shear.

1.2.3 Gust Soaring

Extraction of energy from atmospheric turbulence and gusts is possible through the use of gust soaring techniques. Some birds, such as Black Vultures and Turkey Vultures, use gust soaring to gain energy from the atmospheric turbulence that forms along tree lines [15]. This form of soaring requires very rapid maneuvers to extract energy from the variations in local wind speed and direction.

Research on gust soaring has largely been focused on energy extraction while

maintaining a predetermined flight path. In this type of gust soaring - sometimes called turbulent energy extraction - the maneuver is constrained to the longitudinal plane of the aircraft. Numerous papers have been published on controllers that allow UAVs to extract energy from wind gusts [22]. While it has been shown in simulation that energy extraction is possible, the resulting energy gain tends to be small compared to the sink rate of the vehicle when flying through small and moderate gusts.

1.3 Microlift Soaring and Thermal Strands

In the lower portions of the convective boundary layer there are large atmospheric structures other than thermals that can provide lift. The structures - variously referred to as spoke-like structures, updraft walls, thermal walls, or thermal strands - are formed by air moving from downdrafts back into thermals [23, 54, 55]. As air from downdrafts reaches the lower portions of the convective boundary layer and spreads out, it encounters air from other downdrafts and forms linear regions of converging air encircling each downdraft. The downdraft also suppresses rising air coming from the surface and pushes it towards the converging regions. The warm rising air in the convergence regions is carried along towards their intersections, where thermals begin to form [54].

In 1983, Wallington theorized that pilots of lightweight sailplanes and hang gliders would be able to make use of what he termed "microscale lift structures" [56]. Since then, pilots of manned sailplanes have been able to make use of thermal strands to gain energy, and within the soaring community these smaller-than-thermal convective structures are termed "microlift". Perhaps one of the most notable sailplane pilots to make use of microlift is Gary Osoba, who has written papers and lectured on his experiences with microlift while flying his ultralight sailplane, the Carbon Dragon (see

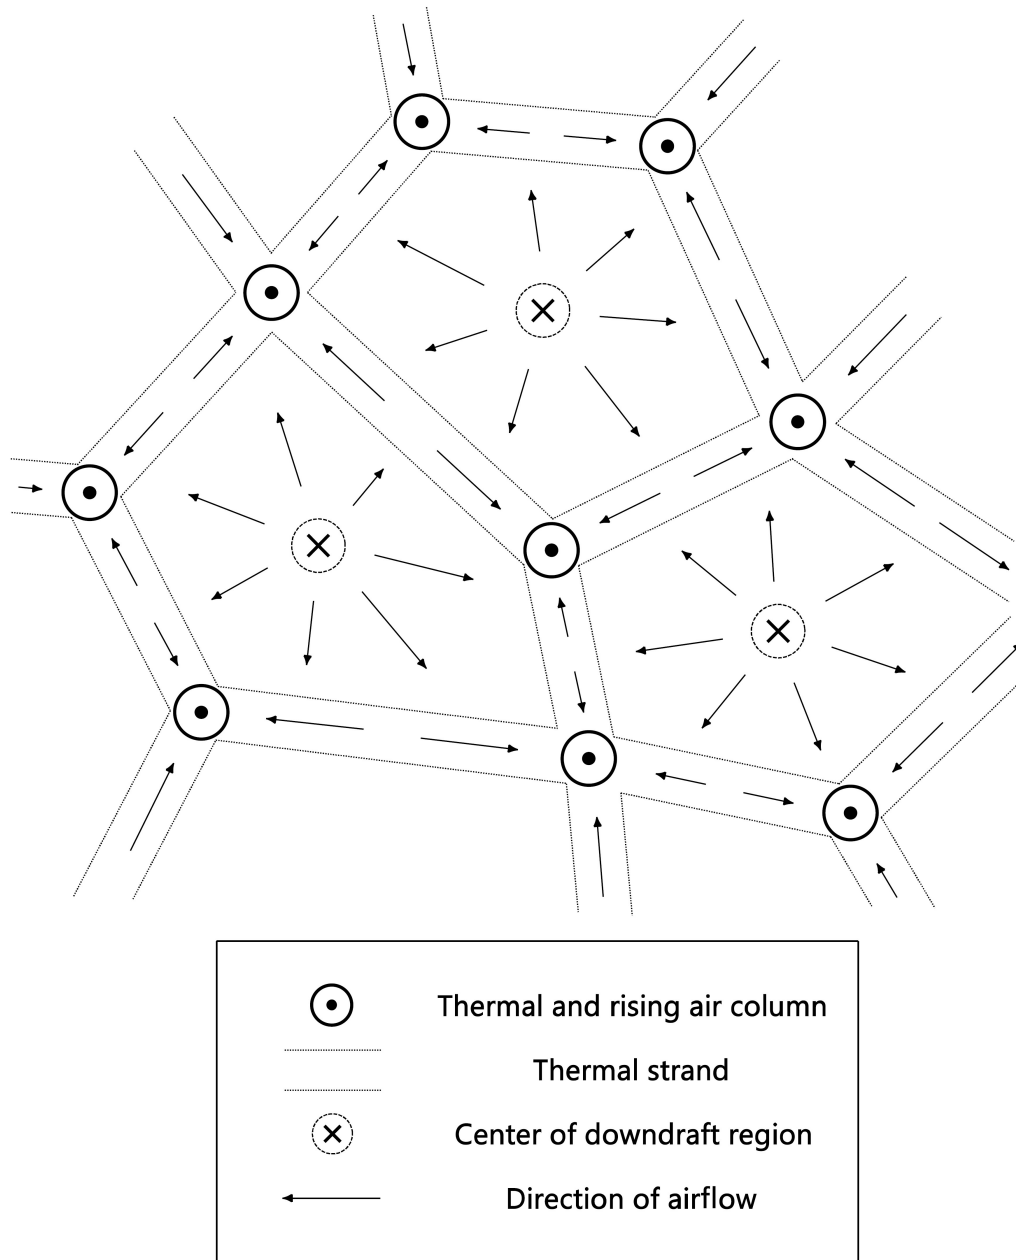


Figure 1.6. A top-down diagram of the structure of downdrafts, thermal strands, and thermals in the atmospheric convective boundary layer.

Figure 1.1) [17–19].

The prospect of soaring along a thermal strand presents two main advantages.

First, from a review of the available data it can be concluded that if the atmospheric boundary layer is being driven by convective activity, then thermal strands are present and would provide a relatively common potential source of energy. Second, unlike thermalling or some forms of dynamic soaring, it is possible to soar along the length of a thermal strand and gain energy without having to stop and loiter. The frequent availability of lift in thermal strands and the potential to be able to continue - at least roughly - in the direction of a goal or waypoint while gaining energy makes microlift soaring and thermal strand tracking potentially a very attractive way for UAVs to gain energy.

1.3.1 Why is it Called Microlift?

Perhaps a good question to ask is why the manned sailplane community refers to thermal strands and other small lift structures as microlift. In general, it is because these forms of lift are too small - hence "micro" - to be of use to the majority of sailplanes. The reasons are rather simple: most sailplanes fly at a speed that makes staying within microlift difficult, microlift occurs at relatively low altitudes, and the lift available in microlift structures is small compared to what is available in a good thermal.

Take, for example, the Glasflügel H-201 Standard Libelle (see Figure 1.7) - a Standard-class competition glider. The speed for best glide angle of a Standard Libelle is roughly 24 meters per second (calculation of speed for best glide angle and longest glide distance is discussed in Section 2.1.2) [7]. This means that a Libelle will fly entirely across the portion of a thermal strand that contains useful lift within 3 to 4 seconds. In contrast, the portion of a thermal that contains useful lift may be a few hundred meters across, and will take more than twice as long to fly through. This

allows the pilot more time to determine where the lift is and to start maneuvering to make use of the lift. In addition, the wing loading and span of most gliders means that their turning radius is too large to allow exploitation of microlift structures [17].



Figure 1.7. John Bird flying his Glasflügel H-201 Standard Libelle near State College, Pennsylvania. Photo by Phil Chidekel.

Thermal strands occur at low altitudes - up to a few hundred meters above ground - at which most glider pilots would (or should) already be looking for potential places to make an off airfield landing [16]. Gary Osoba's ultralight Carbon Dragon is a special case. Its low flight speeds, excellent stability and control, and gentle stall characteristics with a rapid recovery make it able to be safely flown at low altitudes [17].

Finally, the amount of energy within a thermal strand is less than that contained in a thermal, as would be expected given that multiple thermal strands feed energy into a single thermal. This means that in addition to being harder to extract energy from because of aircraft and pilot limitations - many sailplane pilots are unfamiliar with the techniques that are required to extract energy from thermal strands [19] - thermal strands do not offer as much lift as a good thermal. Often it is not worth the time of a larger sailplane pilot to explore the option of microlift, even if it is feasible, because he or she could gain energy much more quickly in a thermal. However, there are cases

when gaining energy from microlift can maintain flight while sailplanes attempting more traditional soaring methods are forced to land [17].

In general, microlift is any source of atmospheric energy that is too small - either physically or energetically - to be of use to the majority of manned sailplanes.

1.3.2 Observations of Birds

It is not just manned ultralight sailplanes exploiting thermal strands. Rarely in soaring is there something done by man that has not first been done by nature.

As previously noted, Black Vultures and Turkey Vultures - which typically extract energy from the atmosphere through thermal soaring - alter their behavior to make uses of other forms of energy when flying at altitudes where thermals are not available. While these behavioral changes include Turkey Vultures gust soaring along tree lines, this researcher and others have also observed Black Vultures and Turkey Vultures overflying farm fields at low to moderate altitudes while flying a roughly linear path without any apparent loss of altitude. Unlike the gust soaring behavior observed by Mallon, Bildstein, and Katzner [15], these flights were linear and occurred over gently rolling farm fields at altitudes estimated to range from 20 to in excess of 100 meters above ground level. It is likely that these birds were performing microlift soaring along thermal strands - hence the apparent constant altitude. In addition to personal observations, Gary Osoba - as quoted by Pierro Morelli - noted that soaring birds often make use of microlift [17].

1.3.3 Characteristics of Thermal Strands

Thermal strands occur in the lower portions of the atmospheric boundary layer and gradually coalesce into thermals as altitude increases. Individually these structures may

be likened to thermal (cloud) streets; however, there are some significant differences. Thermal strands are much smaller than thermal streets and have a length that ranges from an order of magnitude less than the boundary layer depth to the same order of magnitude as the boundary layer depth. In contrast, the convective rolls that form thermal streets may have a length that is 1 to 2 orders of magnitude greater than the boundary layer depth, as seen in thermal streets (cloud streets) that stretch for tens or hundreds of kilometers.

In some ways, the structure of thermal strands mimics that of open cell convection, where regions of descending air are bounded by polygonal lines of rising air. However, while open cell convection in the atmosphere spans the full height of the boundary layer, thermal strands dissipate and are absorbed into thermals well before reaching the top of the atmospheric boundary layer.

Near the ground, thermal strands tend to be small and weak. As altitude increases, thermal strands merge together into larger stronger strands. This can be observed by comparing Figures 1.8 and 1.9.

1.3.4 Thermal Strands in Atmospheric Simulation

Thermal strands have long been observed in simulations of the atmospheric convective boundary layer with low to moderate winds aloft [4, 57–59]. If the atmospheric boundary layer development is dominated by shear forces rather than convective forces, it is unlikely that thermal strands will form. This agrees with the conclusions drawn by Williams and Hacker [23]. The presence of thermal strands in large eddy simulations of the atmospheric convective boundary layer are illustrated in Figures 1.8 and 1.9.

The addition of moderately hilly terrain does not significantly affect the formation




Figure 1.8. Thermal strands in the large eddy simulation wind field of the atmospheric boundary layer provided by Sullivan and Patton (see reference [4] for the paper which describes the generation of this wind field). $z/z_i = 0.06$

of thermal strands, as was seen in Walko et al. [58]. In that paper, the domain included 200 meter tall sinusoidal hills spaced 2 kilometers apart, with an atmospheric boundary layer height of 1 kilometer. In the results, the presence of thermal strands is very clear, and the strands whose axes lay primarily North-South rather than East-West tended to be located near the crests of the hills [58]. This indicates that geographic features can influence the location of thermal strands, as would be expected; however, it should be noted that Walko et al. imposed a spatially and temporally constant surface heat flux to drive the convective boundary layer.

Surface features also do not appear to affect the formation of thermal strands. Park and Baik performed tested the effects of urban-like environments on convective boundary layer development using a lower boundary consisting of a flat plane with a




Figure 1.9. Thermal strands in the large eddy simulation wind field provided by Sullivan and Patton. $z/z_i = 0.12$

grid of 320 meter long by 320 meter wide by 80 meter tall boxes spaced 320 meters apart [60]. Published plots of vertical wind velocity display the linear regions of lift characteristic of thermal strands.

1.4 Characteristics of Microlift Soaring Aircraft

The basic characteristics required to be able to extract energy from microlift are a low sink rate, a decent lift to drag ratio, and a low stall speed and corresponding small turn radius. Any aircraft that soars well needs to have a low sink rate and decent lift to drag ratio; however, to soar in microlift it must also have a low stall speed and tight turning circle to enable tracking along microlift structures. In his 2006 paper, Morelli describes the basic design characteristics required of an ultralight sailplane,

which include a low sink rate and small circling radius [19]. However, exactly what is meant by "low" and "small" was not specified, and proves difficult to identify.

While it is not possible to precisely define exactly what makes an aircraft able to feasibly use microlift, it is possible to identify what is known to work and not work and estimate what combination of characteristics will result in an aircraft that can exploit microlift, as is shown in Figure 1.10.

To begin with, Turkey Vultures have been observed to soar in microlift. The description of how the flight performance of the Turkey Vulture was estimated from published data is detailed in Section 2.4.

As noted by Morelli, high performance hang gliders can soar in microlift; however, older hang gliders have too high a sink rate to be able to gain altitude in microlift [18,19], and the ATOS VR, produced by A-I-R USA, is a very high performance hang glider.

There are relatively few manned sailplanes that are able to fly in microlift; however, the Carbon Dragon ultralight sailplane is one of them. The LET Blanik L-13, Glasflügel Standard Libelle, and DG Flugzeugbau DG-800 (18 meter) are all typical examples of manned gliders and are not able to exploit microlift.

The Effect 38, a paragliding wing manufactured by Pro-Design, is an example of an aircraft with a sink rate too high to effectively be able to make use of microlift. The Effect 38 has a tight turn radius, but its minimum sink rate of 1.2 meters per second [9] means that it would have a hard time getting a net energy gain from microlift. The Effect 38 has comparable performance to older generations of hang glider.

The approximate soaring turn radius for each aircraft was found by assuming that it is flying at the minimum sink lift coefficient for level flight. This airspeed is generally a little bit above the stall speed. The airspeed corresponding to this lift coefficient was calculated for a coordinated level turn with a bank angle of 40 degrees, and the coordinated turn equations were then used to determine the turn radius. The

minimum turn radius of the aircraft would be considerably smaller and would require a much steeper bank.

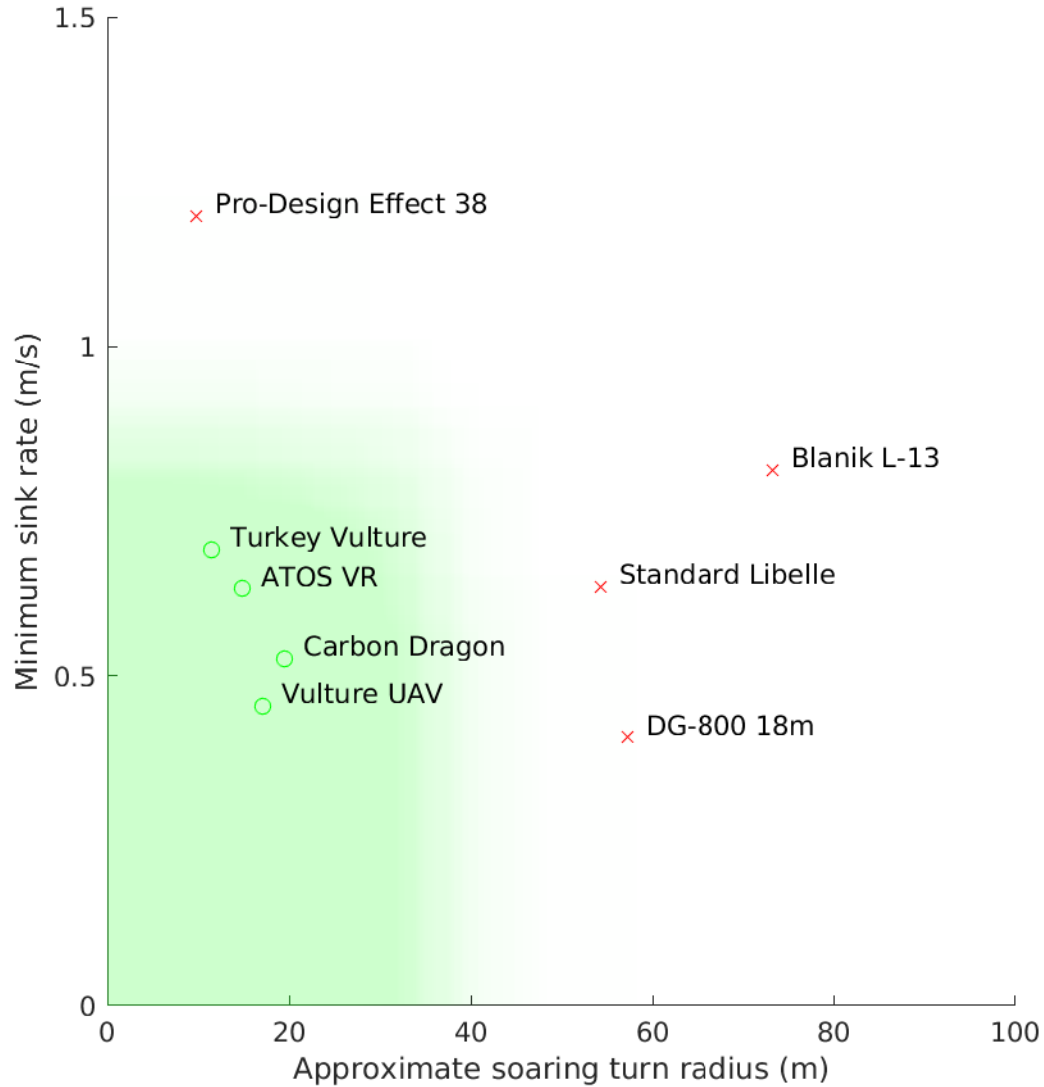


Figure 1.10. Comparison of the minimum sink rate and approximate minimum turn radius in a 40 degree bank between the LET Blanik L-13 [5, 6], Glasflügel H-201 Standard Libelle [5, 7], and DG Flugzeugbau DG-800 (18 meter) [5, 8] sailplanes, Pro-Design Effect 38 paragliding wing [9], ATOS VR hang glider [10], Carbon Dragon ultralight glider [1], Turkey Vulture (see section 2.4), and Vulture UAV. The shaded area indicates approximately where it is practical to exploit microlift.

As will be discussed in detail in Section 2.5, the Vulture UAV is expected to be able to soar in microlift. This is evidenced by the lower minimum sink speed than the Carbon Dragon - indicating that it should be able to obtain a net energy gain from microlift - and it has a turn radius between that of the Carbon Dragon and the ATOS VR - which indicates that it should be able to turn tight enough to track microlift structures.

1.5 Contributions

This thesis presents a method of estimating the location, orientation, and characteristics of a thermal strand using an unscented Kalman filter. This estimator was then integrated with a flight controller designed to track along a thermal strand and extract energy. Finally, the estimator and controller were tested in the wind field obtained from a large eddy simulation of the atmospheric convective boundary layer, and were found to average a positive rate of energy gain from the atmosphere. However, on average this was insufficient to sustain pure soaring flight.

Furthermore, in this thesis it is demonstrated that several methods of attempting to follow a thermal strand cannot result in a stable controller, a model of an idealized thermal strand is developed for use in control system design, and the chance that an aircraft will encounter a thermal strand at a given angle is determined. This thesis also compares two different guidance methods for tracking along a thermal strand, and compares several different initialization methods for the Kalman filter with special consideration of the likelihood of the aircraft's approach from a given angle.

1.6 Reader's Guide

- Chapter 2 contains an overview of soaring flight, a short discussion of maximum glide distance and MacCready's speed to fly theory, a description of the large eddy simulation environment, the development of an idealized thermal strand model to assist in designing a thermal strand tracking controller, the vehicle and sensor models used in simulation, a brief discussion of the flight performance of soaring birds with an emphasis on birds that exploit microlift, and a brief description of the physical Vulture UAV upon which the simulated aircraft is based.
- Chapter 3 discusses the difficulties encountered when attempting to track a thermal strand, demonstrates several methods that are unable to track a thermal strand, and presents the design of a Kalman filter that can estimate the characteristics of a thermal strand.
- Chapter 4 presents the results of testing the ability of the Kalman filter to track an idealized thermal strand, and various methods of generating the initial state estimate for the Kalman filter. It also includes the results of testing the final controller design in a large eddy simulation wind field.
- Chapter 5 is the conclusion and includes a summary of the contributions. It also makes some recommendations of future work.

Chapter 2 |

Problem Statement and Biological Inspiration

This chapter begins with an overview of gliding and soaring flight with a focus on aircraft kinematics in steady flight and calculating the optimal speed to fly (both airspeed for best glide in non-stationary air and MacCready's speed to fly theory). The flight environment - in this case the convective atmospheric boundary layer - is described, as is the large eddy simulation of the convective boundary layer that forms the basis of the simulation environment. An investigation is conducted into the characteristics of the thermal strands found in the large eddy simulation, and an idealized thermal strand model is created to aid in the development of control strategies. The modeling of the sensors and vehicle dynamics are discussed, along with the integration of the vehicle into the simulation environment. The flight characteristics of birds - specifically vultures, as some species are known to exploit microlift - are investigated to help motivate vehicle design. Finally, the physical counterpart to the simulated Vulture UAV is described, and the choice of vehicle is justified through comparison to aircraft and birds that are known to be able to soar in microlift.

2.1 The Basics of Soaring Flight

2.1.1 The Sink Polar

The basis of any theoretical calculation that deals with soaring is the sink polar (also known as the speed polar), which describes the aerodynamic performance of the aircraft in question. The sink polar plots the aircraft's vertical sink rate against horizontal airspeed, and the sink polar of the Vulture UAV can be found in Figure 2.1. Calculation of the sink polar of the Vulture UAV is discussed in Section 2.3.3. It should be noted that a particular sink polar for an aircraft is valid for only one wing loading and geometric configuration (flaps, spoilers, etc.). This is why many sailplanes can carry water ballast - it allows the climb and glide performance of the aircraft to be changed, even "on the fly", without any physical changes to the aircraft.

While soaring birds are unable to change their weight at a whim, they do have various other means at their disposal to alter their glide performance - and hence their sink polar - for optimal flight under any given conditions. As noted by Tucker, gliding birds can adjust their body and wing configuration to achieve a range of sink rates for a given airspeed [61,62]. This includes spreading or retracting their wings to change the aspect ratio, twisting their wings to alter the lift distribution, and changing the position of their wings and/or head to adjust their profile drag [61,63].

In many cases it is convenient to approximate the horizontal airspeed as the airspeed of the vehicle. Provided that the sink rate is relatively small compared to the airspeed, this approximation introduces very little error in sink rate or glide ratio for most airspeeds. For example, this allows the equation for drag coefficient as a function of lift coefficient and zero-lift drag coefficient to be used to derive a theoretical equation for the sink polar, as done by Welch, Welch, and Irving [64]. The quadratic

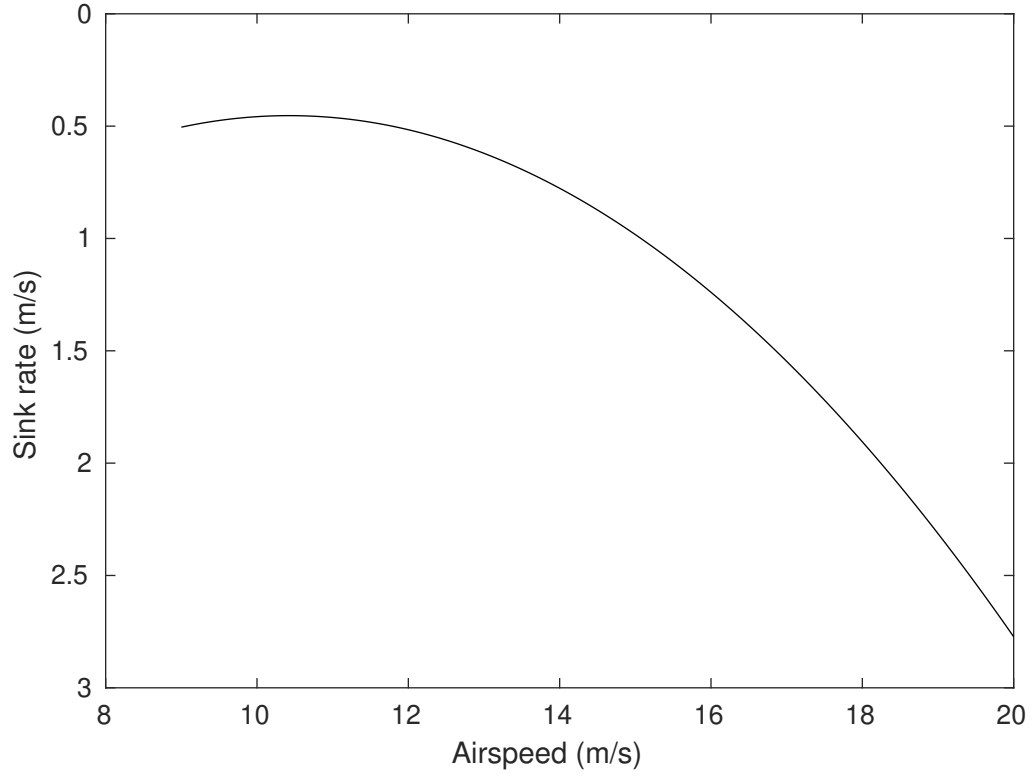


Figure 2.1. Sink polar of the Vulture UAV. Calculation discussed in Section 2.3.3.

expression of total aircraft drag is

$$C_D = \frac{C_L^2}{\pi A Re_0} + C_{D0} \quad (2.1)$$

which can then be manipulated to obtain a sink rate

$$\frac{C_D}{C_L} = \frac{C_L}{\pi A Re_0} + \frac{C_{D0}}{C_L} \quad (2.2)$$

$$\frac{C_D}{C_L} = \frac{2W}{\pi A Re_0 \rho} \frac{1}{V^2} + C_{D0} \frac{\rho S}{2W} V^2 \quad (2.3)$$

$$\frac{W_{sink}}{V_{hor}} = \frac{C_D}{C_L} \quad (2.4)$$

$$W_{sink} = V_{hor} * \left(\frac{2W}{\pi A Re_0 \rho} \frac{1}{V^2} + C_{D0} \frac{\rho S}{2W} V^2 \right) \quad (2.5)$$

$$W_{sink} \approx \frac{2W}{\pi Re_0 \rho} \frac{1}{V_{hor}} + \frac{\rho SC_{D0}}{2W} V_{hor}^3 \quad (2.6)$$

$$W_{sink} \approx \frac{2W}{\pi Re_0 \rho} \frac{1}{V} + \frac{\rho SC_{D0}}{2W} V^3 \quad (2.7)$$

The velocity given by Equations 2.6 and 2.7 is also the weight normalized power required to maintain level flight. These equations are a normalized version of the power required equation for an aircraft. Making the assumption that the horizontal component of airspeed is equal to the airspeed results in relatively little error within the aircraft's flight envelope, and Equation 2.7 is the more common form.

Reichmann presents an alternate and more commonly seen equation for the sink polar. In this form, a parabolic curve is fit to the calculated or experimentally determined sink polar to make it easier to manipulate when performing calculations. The fit is performed by taking three points and finding the equation of the parabola that passes through all three. In order to best capture the performance of the aircraft in the flight regime most critical to efficient soaring calculations - gliding between thermals - it is recommended that one point be located at the maximum lift to drag speed of the aircraft, another at the never exceed speed, and the third halfway between the previous two [16].

$$W_{sink} = aV_{hor}^2 + bV_{hor} + c \quad (2.8)$$

2.1.2 Maximum Glide

One of the primary advantages of environmental energy extraction is that the energy gained can be used to extend the range of the vehicle. When flying through still air, the maximum glide distance is obtained when flying at the airspeed that corresponds to the maximum lift to drag ratio of the aircraft. This speed can easily be obtained from the sink polar by taking a line tangent to the polar curve that passes through

the origin and noting the airspeed at which the line touches the sink polar curve.

However, when the air surrounding the vehicle is moving, flying maximum lift to drag will not yield the maximum glide distance. This is because flying maximum lift to drag ratio will result in the maximum glide distance with respect to the air, it will not result in the maximum glide distance with respect to the ground. Conceptually, if the aircraft is flying with a tailwind, the pilot should slow down to decrease the sink rate and make better use of the assist provided by the wind. Conversely, if the aircraft is flying into a headwind, the pilot should increase speed. Similarly, if the aircraft is flying through rising air the pilot should decrease airspeed to reduce sink rate, and if the aircraft is flying through descending air the pilot should increase airspeed.

The question, of course, is by how much should the airspeed be increased or decreased to account for the motion of the air. If the sink polar is adjusted for the motion of the air, as described by Reichmann [16], then finding the speed that yields the maximum glide distance for any given air motion is as easy as finding the speed for still air. Considering only horizontal air velocities:

$$W_{sink} = f(V_{hor}) \quad (2.9)$$

$$V_{ground} = V_{hor} + U_{wind} \quad (2.10)$$

Then the sink polar with respect to ground speed becomes:

$$W_{sink} = f(V_{ground}) = f(V_{hor} + U_{wind}) \quad (2.11)$$

Graphically, this can be thought of as shifting the sink polar curve horizontally - to the right in the case of a tailwind, and to the left in the case of a headwind. This altered curve then reflects the aircraft's sink rate as a function of ground speed, and

finding the maximum glide distance is as simple as again drawing a line tangent to the curve that passes through the origin [16].

Similarly, for vertical air velocities we are interested in finding the rate of change of altitude as a function of airspeed.

$$\dot{h} = \dot{h}_{wind} - W_{sink} \quad (2.12)$$

And taking the negative to maintain the convention of a positive sink rate or negative altitude rate

$$-\dot{h} = W_{sink} - \dot{h}_{wind} = f(V_{hor}) - \dot{h}_{wind} \quad (2.13)$$

This reflects a vertical shift of the sink polar curve - shifting upwards when the air is rising, and downwards when the air is descending. Again, to find the maximum glide distance a tangent line is taken that passes through the origin [16].

The equation for the rate of change of the altitude of the aircraft as a function of ground speed can then be written as:

$$-\dot{h} = f(V_{hor} + U_{wind}) - \dot{h}_{wind} \quad (2.14)$$

Noting that a positive wind velocity indicates a tailwind, and that a positive \dot{h}_{wind} indicates flying through a rising airmass.

2.1.3 MacCready and Speed to Fly

While flying the speed to obtain maximum glide distance will result in the aircraft traveling as far as possible, it is not the fastest way to traverse a given distance. Rather than optimizing for range, the sink polar equations can be optimized for cruise speed over a given flight segment. MacCready's solution to this optimization problem

considers a flight segment consisting of a glide to some destination (presumably a developing cumulus cloud and associated thermal) followed by a climb back up to the starting altitude [16, 65].

$$t = t_{glide} + t_{climb} = \frac{d_{glide}}{V_{hor}} + \frac{h_{climb}}{\dot{h}_{climb}} \quad (2.15)$$

$$V_{cruise} = \frac{d_{glide}}{t} \quad (2.16)$$

$$h_{climb} = d_{glide} * \frac{W_{sink} - \dot{h}_{wind}}{V_{hor}} \quad (2.17)$$

$$t = \frac{d_{glide}}{V_{hor}} + \frac{d_{glide}}{\dot{h}_{climb}} * \frac{W_{sink} - \dot{h}_{wind}}{V_{hor}} \quad (2.18)$$

In order to find the horizontal airspeed that results in the minimum total time, take the derivative with respect to horizontal airspeed and set equal to zero. Knowing that the glide distance, d_{glide} must be greater than zero yields the MacCready relation: [16, 65]

$$V_{hor} * \left[\frac{dW_{sink}}{dV_{hor}} \right]_{V_{hor}} = (W_{sink} - \dot{h}_{wind}) - \dot{h}_{climb} \quad (2.19)$$

This says that your horizontal airspeed in the glide times the slope of the sink polar at your glide airspeed should be equal to the total sink rate of the aircraft minus the rate at which the aircraft will be climbing when the destination is reached [16].

There are also solutions for variations of the fastest cruise problem. Stojkovic developed a generalized speed to fly theory that expands upon MacCready's work to include all forms of lift [66]. Others have investigated how altitude should impact the optimal airspeed and the minimum thermal strength for which the pilot should stop and circle [67, 68]. DeJong investigated various optimal soaring problems, including the angles at which one should depart a cloud street to return to the desired flightpath,

and the maximum deviation angle from the desired flightpath to intersect a thermal of a given strength - "the optimal zigzagging problem" [69].

2.2 Simulating the Convective Boundary Layer

2.2.1 Large Eddy Simulation Environment

The model of the atmospheric convective boundary layer used to design a thermal strand following controller was generated by a large eddy simulation performed by Sullivan and Patton. In this simulation, the boundary layer was driven by a surface heat flux and the ground was defined to be a flat plane with no features other than surface roughness. The data provided consists of a 5120 meter square by 2048 meter tall domain containing 1024^3 grid points in each of two vertically staggered grids. The grid resolution of both grids is 5 meters horizontally and 2 meters vertically. One grid contains the vertical wind velocity and subfilter scale energy, and the other grid contains the horizontal wind velocities and the virtual potential temperature. The two grids are offset by half of the vertical grid spacing, as can be seen in Figure 2.2 [4].

Sullivan and Patton's simulated convective boundary layer is driven by a constant surface buoyancy flux of $0.24 \text{ K}^*\text{m/s}$, and has a weak geostrophic wind (wind aloft) velocity of 1 m/s heading due North imposed on the upper boundary of the domain. The surface roughness is 0.1 meters, and the Coriolis parameter is defined to be $1.0 * 10^{-4} \text{ seconds}^{-1}$ [4], which corresponds to a latitude of 43.29 degrees. (Large eddy simulations of the atmospheric convective boundary layer have been conducted with higher geostrophic wind - wind aloft - speed and surface roughness, and thermal strands were still present.) The wind field in Sullivan and Patton's LES data shows well developed convection with a convective boundary layer height of approximately 1




Figure 2.2. Vertical grid offset in the large eddy simulation data. dx and dz are grid spacings, and k is the vertical index.

kilometer.

The full LES data file is a 42.9 gigabytes, which is too large to read into MATLAB. The domain was reduced from 2048 meters tall (roughly two boundary layer heights) to 400 meters and alternating layers of the grid were removed to decrease the vertical resolution from 2 meters to 4 meters. This removed from the environment the upper portions of the boundary layer where thermal strands do not exist, and brought the vertical resolution closer to the horizontal resolution. In addition, the subfilter scale energy, virtual potential temperature, and horizontal wind velocities were removed because the most relevant data is the local vertical wind velocity. The resulting vertical wind velocity data set was saved as a MATLAB data file. Removing unnecessary data and downsampling reduced the file size to 808.4 megabytes, which could easily be read into MATLAB.

2.2.2 The Idealized Thermal Strand

To facilitate the design of a control system able to track thermal strands, the LES data was analyzed and used to generate a mathematical model describing a straight, idealized thermal strand. Thermal strands were manually identified in the LES data, and a MATLAB script was used to calculate the vertical velocity along a cross section of the strand. Multiple cross-sections were taken and averaged, and then a Gaussian-based function was fit to the data.

The idealized strand is supposed to capture the basic features of the thermal strands in the LES without overfitting or biasing the controller to only work in a particular wind field. The basic equation of the ideal thermal strand is

$$\dot{h}_{wind} = \frac{-\dot{h}_{max}}{2} + \frac{3\dot{h}_{max}}{2} * \exp\left(\frac{-d_{str}^2}{2\sigma_{str}^2}\right) \quad (2.20)$$

where \dot{h}_{max} is the updraft velocity at the center of the strand, d_{str} is the distance from the central axis of the strand, and σ_{str} is the width parameter of the strand and the standard deviation of the Gaussian component (the second term).

Equation 2.20 is based on the Gaussian model of a thermal, and can be thought of as a Gaussian lift distribution swept along the length of the thermal strand. The term $\frac{-\dot{h}_{max}}{2}$ reflects the background sink surrounding the thermal strand. This model of an ideal thermal strand fits the LES data well, as seen in Figures 2.3, 2.4, and 2.5.

2.3 Vehicle Modeling

The vehicle model used in the simulation reflects the flight performance of the Vulture UAV, a small research UAV that was built from a Magellan XL radio controlled sailplane. The flight characteristics of the Vulture UAV are similar to those of Black

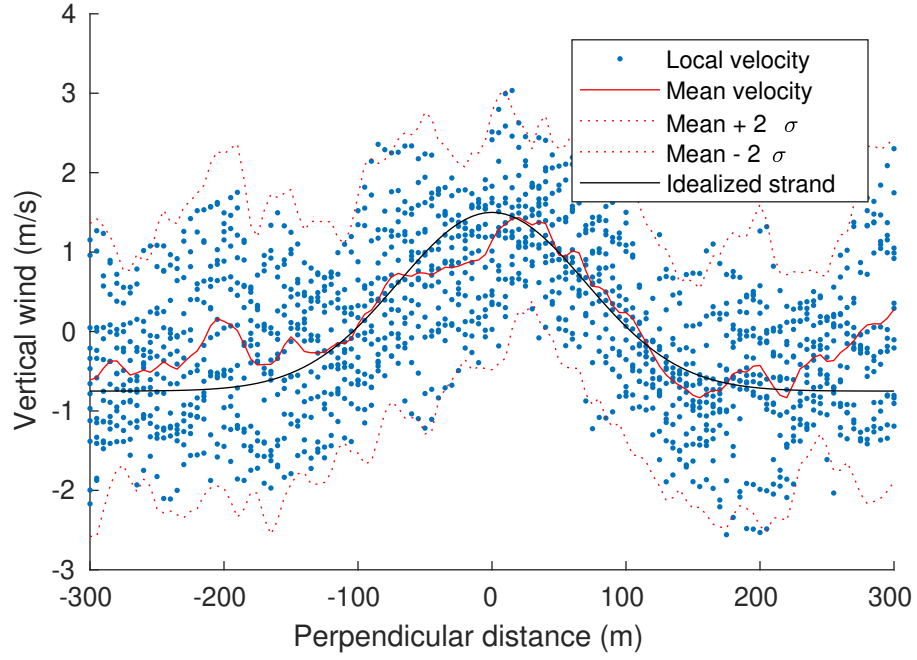


Figure 2.3. Cross section of a thermal strand in the LES data at an altitude of 60 meters. $\sigma_{str} = 70$ meters, $\dot{h}_{max} = 1.5$ meters per second




Figure 2.4. Cross section of a thermal strand in the LES data at an altitude of 80 meters. $\sigma_{str} = 42$ meters, $\dot{h}_{max} = 2.25$ meters per second




Figure 2.5. Cross section of a thermal strand in the LES data at an altitude of 15 meters. $\sigma_{str} = 94$ meters, $\dot{h}_{max} = 1.75$ meters per second

Vultures and Turkey Vultures, as described in Section 2.5, and in general are well-suited to microlift soaring. The low cruising airspeed, low sink rate, and moderately high maximum lift-to-drag ratio make the Vulture UAV well-suited to exploiting microlift.

The model of the aircraft decouples the lateral-directional and vertical dynamics of the aircraft, and simulates the airspeed and pitch dynamics of the aircraft through the use of proportional gains. This considerably simplifies the simulation both structurally and computationally.

For soaring vehicles it is convenient to normalize energy by the vehicle weight because it allows energy to be expressed as height and power to be expressed as a vertical rate. This practice of using normalized energy and power greatly simplifies the calculation of energy gained by the vehicle and vertical rate, as seen in Section

2.3.3. In general, a normalized energy is

$$e = \frac{E}{mg} \quad (2.21)$$

2.3.1 Definition of Simulation Domain and Reference Frames

The simulation was decided to be a 5120 meter square area to make full use of the LES simulation data provided by Sullivan and Patton. The standard aerospace body and flat Earth inertial coordinate systems were used in the simulation of the vehicle. A flat Earth inertial frame was used because the curvature of the Earth would be of negligible impact over the distances involved. The x and y axes of the inertial frame are aligned with geographic North and East, and the z axis points down. In the body frame, the x axis is aligned with the longitudinal axis of the vehicle and the y axis extends out the right wing.

2.3.2 Vehicle Sensor Modeling

The vehicle was assumed to perfectly know its location, airspeed, orientation, and body rates.

The Vulture UAV can simulate various types of variometers using its onboard sensors, including a netto variometer, also known as an airmass variometer. A netto variometer compensates for the total energy of the vehicle and for the vehicle's sink rate, and will always display 0 when the aircraft is flying through still air. The simulated netto variometer onboard the Vulture UAV also compensates for load factor in turns and when thermalling, making it a relative netto variometer, sometimes called a super netto variometer, and will output a measurement of the current estimated vertical wind velocity.

The relative netto variometer in the simulation is simulated by adding zero mean Gaussian noise to the true local vertical wind velocity, and then storing those values in a data buffer to add an optional time delay representative of the variometer's time constant. The measurement at time step k is then given by

$$\dot{h}_{wind,measured,k} = \dot{h}_{wind,k-\kappa} + N(0, \sigma_{vario}) \quad (2.22)$$

where κ is the number of time steps by which the measurement is delayed, and σ_{vario} is the standard deviation of the Gaussian noise.

The roll disturbance detector on the Vulture UAV calculates the residual rolling moment of the aircraft by computing the total rolling moment from known sources using the aircraft's current state, control deflections, and calculated stability and control derivatives. The total rolling moment exerted on the aircraft can be calculated using the measured angular acceleration and the inertial properties of the aircraft, which in turn can be used to calculate the rolling moment resulting from unknown sources, termed the residual rolling moment. This contribution is assumed to be from the motion of the air around the aircraft.

$$C_{l,res} = C_{l,total} - C_{l,state} - C_{l,control} \quad (2.23)$$

The measurements output by the roll disturbance detector are simulated by calculating the rolling moment exerted on the aircraft by local wind gradients and adding Gaussian noise with standard deviation σ_{rdd} . The roll disturbance measurement can be delayed by λ time steps, although the time constant of the real roll disturbance detector is very small to the point of being negligible. The roll disturbance detector measurement

at time step k is then

$$C_{l,res,measured,k} = C_{l,res,k-\lambda} + N(0, \sigma_{rdd}) \quad (2.24)$$

From the residual rolling moment coefficient, it is relatively easy to obtain a measurement of the local wind gradient in the spanwise direction. Using the residual rolling moment coefficient and the roll damping of the aircraft, it is possible to calculate the equivalent normalized roll rate to be

$$\hat{\tilde{p}} = \frac{C_{l,res}}{C_{lp}} \quad (2.25)$$

where the tilde denotes that this is an equivalent roll rate, and the hat denotes that this is a normalized dimensionless quantity. This normalization follows the standard stability and control roll rate normalization, given by

$$\tilde{p} = \frac{b\tilde{p}}{2V_0} \quad (2.26)$$

From this equivalent roll rate, it is possible to calculate the change relative air velocity parallel to the body 3-axis induced by the equivalent roll rate as a function of spanwise position. From there, it is possible to calculate the derivative of the relative air velocity parallel to the body 3-axis with respect to spanwise distance. Beginning with the change in vertical air motion relative to the aircraft at an arbitrary spanwise position

$$\Delta w = \tilde{p}y \quad (2.27)$$

and taking the derivative

$$\frac{\partial \Delta w}{\partial y} = \tilde{p} \quad (2.28)$$

it can be found that

$$\frac{\partial \Delta w}{\partial y} = \frac{C_{l,res}}{C_{lp}} \frac{2V_0}{b} \quad (2.29)$$

If it is assumed that the only contribution to residual rolling moment is a gradient in the local vertical wind velocity, and that the pitch angle is small (as it typically is when soaring), then the derivative of local vertical wind velocity along the horizontal projection of the body 2-axis is found to be

$$\frac{\partial \dot{h}}{\partial y^1} = \frac{C_{l,res}}{C_{lp}} \frac{2V_0}{b} \frac{1}{\cos^2(\phi)} \quad (2.30)$$

where y^1 is the 2-axis in the first intermediate Euler frame. This approximation holds well for small pitch and bank angles, where the gradients of the horizontal components of local wind velocity have relatively little impact on the roll disturbance of the vehicle.

2.3.3 Vehicle Dynamics Modeling

The aircraft is simulated in 3 dimensions using a 5 degree of freedom model. It is assumed that the aircraft performs perfectly coordinated turns (zero sideslip), and that the longitudinal and lateral dynamics of the aircraft are decoupled. In addition, the longitudinal dynamics are simplified by defining the vertical rate of the aircraft to be equal to the normalized excess power. Note that because of the choice of coordinate system, a positive change in energy results in a negative vertical rate. Therefore the vertical rate of the aircraft is

$$-\dot{z} = p_{excess} \quad (2.31)$$

$$p_{excess} = p_{motor} - W_{sink} + \dot{h}_{air} \quad (2.32)$$

where W_{sink} is the sink rate of the simulated Vulture UAV, given by

$$W_{sink} = 0.0253V^2 - 0.5275V + 3.2028 \quad (2.33)$$

The sink rate of the simulated Vulture UAV is calculated using the parabolic form of the sink polar, and is found in Equation 2.33. The sink polar for the Vulture UAV was generated by adjusting the experimentally determined sink polar of a larger autonomous soaring vehicle based on the SB-XC airframe (sold by RnR Products) for the lower wing loading of the Vulture UAV. The method of altering a sink polar for a different wing loading is included in Reichmann's book [16] and is described in detail in Section 2.4. A plot of the sink polar was shown earlier in Figure 2.1.

When the aircraft is turning the sink rate of the aircraft increases. As derived by Haubenhofer, the sink polar for a turn with bank angle ϕ can be obtained by adjusting the level flight polar as follows [70]:

$$n = \frac{1}{\cos(\phi)} \quad (2.34)$$

$$V_{turn} = \sqrt{n}V \quad (2.35)$$

$$W_{turn} = (\sqrt{n})^3 W \quad (2.36)$$

Therefore, the sink rate of the Vulture UAV can be expressed as a function of airspeed and bank angle, and is given by

$$W_{sink}(V, \phi) = \left(\frac{1}{\cos(\phi)} \right)^{\frac{3}{2}} \left(0.0253V^2 \cos(\phi) - 0.5275V \sqrt{\cos(\phi)} + 3.2028 \right) \quad (2.37)$$

The airspeed and roll dynamics of the vehicle are simulated using proportional gains that have been tuned to reflect the flight characteristics of the real vehicle. The

airspeed gain reflects using pitch angle to control airspeed, and the bank angle gain reflects the significant aileron authority and high roll rate of the aircraft. The values were tuned to accurately reflect the response to commanded changes in airspeed of 1 to 2 meters per second, and commanded changes in bank angle up to $\pi/2$ radians. The gain values are provided in Table 2.1. The airspeed rate and roll rate of the simulated Vulture UAV at an arbitrary time step k are then

$$\dot{u}_{IAS,k} = k_u(u_{IAS,cmd,k} - u_{IAS,k}) \quad (2.38)$$

and

$$\dot{\phi}_k = k_\phi(\phi_{cmd,k} - \phi_k) \quad (2.39)$$

Table 2.1. Proportional gain values for simulating airspeed and roll response to commands

Gain	Value
k_u	1.00
k_ϕ	2.50

Using the decoupled longitudinal dynamics that are based on normalized energy, the pitch angle, θ is no longer needed. Therefore θ and $\dot{\theta}$ can be set to zero as follows

$$\dot{\theta} = \theta = 0 \quad (2.40)$$

The state vector for the simulation can be found in Equation 2.41. Note that because sideslip is assumed to be zero, the component of wind parallel to the body y-axis must be zero, but is included in the state vector for convenience.

$$\mathbf{x} = \begin{bmatrix} \mathbf{V}^B & \boldsymbol{\omega}^B & \boldsymbol{\Phi} & \mathbf{r}^N \end{bmatrix}^T \quad (2.41)$$

$$\mathbf{V}^{\mathbf{B}} = \begin{bmatrix} u & v & w \end{bmatrix}^T \quad (2.42)$$

$$\boldsymbol{\omega}^{\mathbf{B}} = \begin{bmatrix} p & q & r \end{bmatrix}^T \quad (2.43)$$

$$\boldsymbol{\Phi} = \begin{bmatrix} \phi & \theta & \psi \end{bmatrix}^T \quad (2.44)$$

$$\mathbf{r}^{\mathbf{N}} = \begin{bmatrix} x & y & z \end{bmatrix}^T \quad (2.45)$$

2.3.4 Integration of the Equations of Motion

Forward Euler was determined to be sufficient because simulating the exact state of the vehicle is not critical to being able to locate follow a thermal strand. For example, a position error of a few centimeters is negligible when navigation is occurring on a length scale of meters to tens of meters. Furthermore, if the controller uses the same state propagation methods as the simulation, then the error between the estimated state and the true simulated state generated by the state propagation is zero.

Using a first order integration method allows the control system that provides control inputs at a fixed rate to be easily synchronized with the discrete time physical simulation. This reduces the complexity of the simulation structure and the number of computations required per time step. Using an integration method that requires computation of the state derivative between time steps, such as 4th-order Runge-Kutta, would require the local wing velocity to be interpolated twice instead of once - one of the more computationally intensive portions of the simulation.

The time propagation of aircraft airspeed, bank angle, and altitude can be found in Equations 2.47, 2.48, and 2.49. The simulation of these quantities is first order

accurate. The conversion between true airspeed and indicated airspeed is given by

$$u_{TAS} = u_{IAS} \sqrt{\frac{\rho_{SSL}}{\rho}} \quad (2.46)$$

and the airspeed, roll rate, and altitude at the next time step are given by

$$u_{TAS,k+1} = u_{TAS,k} + \dot{u}_{TAS,k} * \Delta t \quad (2.47)$$

$$\phi_{k+1} = \phi_k + \dot{\phi}_k * \Delta t \quad (2.48)$$

$$z_{k+1} = z_k + \dot{z}_k * \Delta t \quad (2.49)$$

To simplify equations, the change in horizontal position is computed in the first intermediate Euler frame before being transformed back into the inertial frame. Note that the first intermediate Euler frame is held fixed at time k . In the first intermediate Euler frame, the 1-axis is aligned with the projection of the vehicle's longitudinal axis onto the horizontal plane. This change in horizontal position can then be applied to the horizontal position components of the state vector. The horizontal position of the aircraft at the next time step can be found using

$$DCM_{N1} = \begin{bmatrix} \cos(-\psi) & \sin(-\psi) & 0 \\ -\sin(-\psi) & \cos(-\psi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.50)$$

$$\Delta \mathbf{r}^N = DCM_{N1} * \Delta \mathbf{r}^1 \quad (2.51)$$

$$\Delta \mathbf{r}^1 = DCM_{N1}^T * \Delta \mathbf{r}^N \quad (2.52)$$

$$\mathbf{r}_{k+1}^N = \mathbf{r}_k^N + DCM_{N1,k} * \Delta \mathbf{r}_k^1 \quad (2.53)$$

When the vehicle is flying wings level, calculating the change in horizontal position in the first intermediate Euler frame is trivial, and is given by

$$\begin{aligned} \dot{x}^1 &= u \\ \phi = 0, \quad \dot{y}^1 &= 0 \\ \dot{z}^1 &= 0 \end{aligned} \tag{2.54}$$

$$\Delta \mathbf{r}^1 = \begin{bmatrix} u * \Delta t \\ 0 \\ 0 \end{bmatrix} \tag{2.55}$$

It is possible to achieve increased accuracy in horizontal position with little additional computational cost by using the equations for a coordinated level turn. It is assumed that the aircraft's bank angle and airspeed are constant between time steps, and hence the aircraft travels along a perfect circular arc. The radius of the turn is given by

$$r_{turn} = \frac{1}{\tan(\phi) * g/u^2} \tag{2.56}$$

with turn rate

$$\dot{\psi} = \frac{u}{r_{turn}} \tag{2.57}$$

and the change in position in the first intermediate Euler frame is then

$$\Delta \mathbf{r}^1 = \begin{bmatrix} \Delta x^1 \\ \Delta y^1 \\ \Delta z^1 \end{bmatrix} = u \Delta t \begin{bmatrix} r_{turn} \sin(\dot{\psi} \Delta t) \\ r_{turn} (1 - \cos(\dot{\psi} \Delta t)) \\ 0 \end{bmatrix} \tag{2.58}$$

To avoid the singularity that occurs in turn radius when the bank angle equals 0, a minimum bank angle of 10^{-5} radians is required to use the turn radius equations, and

therefore bank angles less than 10^{-5} radians are assumed to equal 0. Under typical cruise conditions ($u = 14$ m/s), a bank angle of 10^{-5} radians corresponds to a turn radius of approximately 2000 km. After flying for 5 km at a bank of 10^{-5} radians, the error in position resulting from assuming that the aircraft is in level flight is a mere 6.25 m.

After applying the equations for a steady level turn and solving for the position and orientation, it is necessary to calculate the rotational rates in the body reference frame using the previously determined Euler angle rates. The body rates are given by

$$\begin{bmatrix} p \\ q \\ r \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & 0 & -\sin(\theta) \\ 0 & \cos(\phi) & \sin(\phi)\cos(\theta) \\ 0 & -\sin(\phi) & \cos(\phi)\cos(\theta) \end{bmatrix}_{k+1} \begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix}_{k+1} \quad (2.59)$$

2.4 Flight Performance of Soaring Birds

The equations of aircraft soaring flight can be applied to the flight performance of soaring birds. There have been many papers published on the flight performance of soaring birds; however, relatively few deal with birds that have been observed to make use of microlift and thermal strands. In particular, papers published on Black Vultures and Turkey Vultures have significant data scattering, suspect methods, or both that call into question the accuracy of the results.

Raspet estimated the soaring performance of Black Vultures by flying a glider behind the bird and recording the relative glide performance. He then corrected the measurements to account for the glide performance of the glider [71, 72]. However, Tucker and Parrott later noted that the glide performance calculated by Raspet requires that the parasite drag coefficient of a Black Vulture be less than that of

laminar flow over a flat plate [73]. Furthermore, Raspel's estimated sink polar indicates that a Black Vulture has a maximum lift to drag ratio in excess of 22 [71,72], which is an unreasonably high value when compared to data on other species of soaring birds.

The overly optimistic sink polar estimated by Raspel is likely the result of the extremely noisy method which Raspel used to obtain the relative altitude and distance of the birds. Raspel would visually estimate the relative altitude and distance and transmit that estimate to his ground crew via radio [71]. Twenty years later, Pennycuick would use a more refined version of this method - using a rangefinder and camera to measure the distance between the aircraft and the bird rather than relying on visual estimates made by the pilot - and even with the improved method the data collected by Pennycuick contained a large amount of random noise [74].

Parrott performed wind tunnel tests on a live Black Vulture, and would slowly decrease the inclination of the tunnel until the vulture started flapping in order to estimate the maximum glide angle for a given airspeed [73,75]. While a good method in general, the results of the wind tunnel tests were questionable due to the size of a Black Vulture relative to the wind tunnel used. The average wing span of a Black Vulture is 1.44 meters [76], and the wind tunnel used by Tucker and Parrot had a 1.4 meter wide by 1.1 meter high test section. The data presented by Parrott showed that the Black Vulture used in the tests had extended its wings to a maximum of 1.37 meters when flying in the wind tunnel, but the unconstrained wingspan of the bird was not reported [75]. Even if the Black Vulture was able to fully extend its wings during the experiment, it was still much too large for the wind tunnel, as it is recommended that the model - or bird - being tested have a span no larger than 70 to 80 percent of the width of the test section - much lower than the 97.9 percent recorded in Parrott's experiment [75].

The results of Parrott's experiment indicates that a Black Vulture has a maximum

lift to drag ratio of 11.6 which is more reasonable than the ratio of 22 calculated by Raspet, but is low compared to the maximum lift to drag ratios calculated for various other species of vulture [63,74], and is close to that of soaring birds that are also required to make high-speed dives, such as hawks and falcons [73,77].

The previously mentioned papers are the only ones that contain experimental analysis of the glide characteristics of actual Black Vultures and Turkey Vultures. However, Hoey built a radio controlled models of a Raven, Seagull, Pelican, and Turkey Vulture to investigate the stability and control of birds [78,79]. The majority of his published work deals with the model of a Raven, and the models of the Seagull, Pelican, and Turkey Vulture are barely mentioned in his published works. The Raven model was noted to have an estimated maximum lift to drag ratio of 8. If the glide characteristics of the Turkey Vulture model were analyzed, they were not published.

The papers published that contain the glide performance of Black Vultures and Turkey Vultures contain data that was collected through suspect methods, which makes the resulting sink polars of questionable accuracy. However, it is possible to use data collected for a geometrically similar bird and adjust its sink polar to account for the different wing loading. It so happens that the wingspan, aspect ratio, and body geometry of the White-Backed Vulture is similar to that of the Black Vulture and the Turkey Vulture. This geometric similarity can be seen in Figures 2.6, 2.7, 2.8.



Figure 2.6. African White-Backed Vulture [11]




Figure 2.7. Black Vulture [12]



Figure 2.8. Turkey Vulture [13]

Pennycuick used in-flight measurements taken from a Schleicher ASK-14 motor glider to calculate the sink polar of the White-Backed Vulture. Measurements of the relative position of the gliding vultures were taken using a camera and optical range finder, and were subsequently corrected to account for the glide characteristics of the glider [39, 74]. Pennicuick then fit the adjusted glide measurements to the equation for an aircraft's sink polar derived in Welch, Welch, and Irving [64]. Fitting the data to a theoretical model minimized the impact of measurement error.

A careful inspection of Pennicuick's calculations reveals that he applies the assumption that White-Backed Vultures have an Oswald's efficiency factor of 1, rather than his stated assumption of a span efficiency factor of 1 [74]. This error resulted in his calculated sink polar for the White-Backed Vulture indicating higher lift to drag

ratios than there should be.

Starting with the equation found in Welch, Welch, and Irving's [64] and repeating Pennycuick's method [74] using a more reasonable Oswald's efficiency factor of 0.80, a corrected sink polar equation for the White-Backed Vulture may be found as follows

$$W_{sink} = \frac{2kW}{\pi AR\rho_0 S} \frac{1}{V} + \frac{C_{D0}\rho_0 S}{2W} V^3 \quad (2.60)$$

$$k = \frac{1}{e_0} \quad (2.61)$$

$$W_{sink} = \frac{7.248}{V} + (1.710 * 10^{-4}) V^3 \quad (2.62)$$

Pennycuick's sink polar and the revised sink polar of the African White-Backed Vulture are shown in Figure 2.9. Pennycuick's model predicts a minimum sink rate of 0.76 meters per second and a maximum lift to drag ratio of 15.3, while the model that includes a more reasonable estimate of Oswald's efficiency factor predicts a minimum sink rate of 0.89 meters per second and a maximum lift to drag ratio of 14.2.

The corrected sink polar can then be adjusted to compensate for the lower wing loading of Black Vultures and Turkey Vultures. Reichmann provides equations for adjusting a quadratic sink polar for wing loading [16]. Starting with the basic parabolic form of the sink polar

$$W_{sink,old} = aV^2 + bV + c \quad (2.63)$$

and using the scaling factor

$$A = \sqrt{\frac{\left(\frac{W}{S}\right)_{new}}{\left(\frac{W}{S}\right)_{old}}} \quad (2.64)$$

the adjusted sink polar is then given by

$$W_{sink,new} = \frac{a}{A} V^2 + bV + Ac \quad (2.65)$$

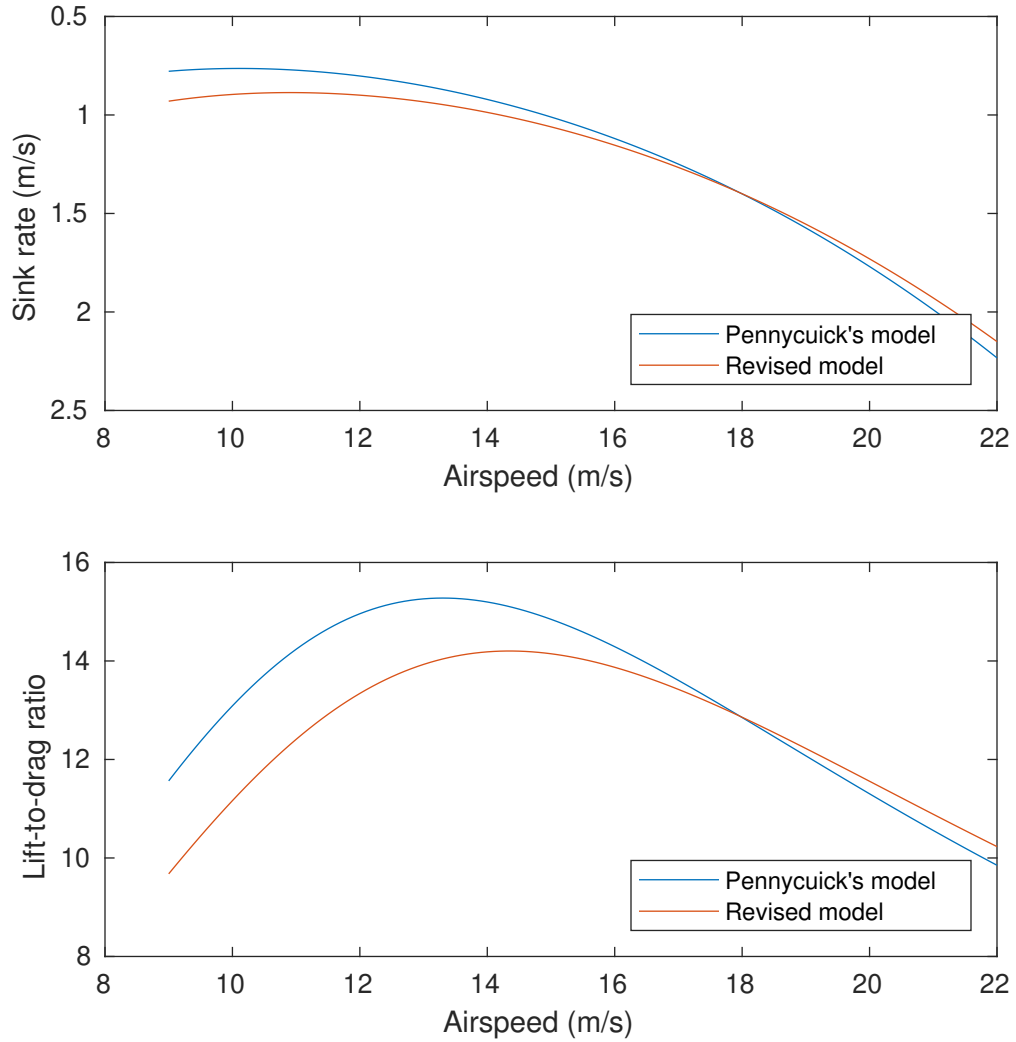


Figure 2.9. Estimated sink rate and lift to drag ratio of the White-Backed Vulture using Pennycuick's sink polar and the corrected sink polar

The above equations result from first recognizing that the drag coefficient of the vehicle for a given lift coefficient is almost entirely a function of the vehicle's geometry, and then assuming that the drag coefficient for a given lift coefficient is constant across airspeed. This assumption ignores Reynolds number effects, which tend to be small when dealing with changes in airspeed and sink rate of 10 to 15 percent. It also assumes that the local air density is constant - Reichmann also provides similar equations that adjust the sink polar for air density and density altitude [16].

Given a desired lift coefficient, the lift generated by the aircraft is

$$L = W = \frac{1}{2}\rho V^2 S C_L \quad (2.66)$$

which allows the conclusion that

$$\frac{W}{S} \propto V^2 \quad (2.67)$$

and therefore

$$V_{new} \sqrt{\left(\frac{W}{S}\right)_{new}} = V_{old} \sqrt{\left(\frac{W}{S}\right)_{old}} \quad (2.68)$$

$$V_{new} = V_{old} \sqrt{\frac{\left(\frac{W}{S}\right)_{old}}{\left(\frac{W}{S}\right)_{new}}} = \frac{V_{old}}{A} \quad (2.69)$$

which gives the new airspeed required to fly at a given lift coefficient when the wing loading of the vehicle is changed. Now starting with the relationship between airspeed and sink rate and lift to drag ratio, and using the relationship between new and old airspeed, it is possible to find the new sink rate. The lift to drag ratio is known to be

$$\frac{V}{W_{sink}} = \frac{L}{D} = \frac{C_L}{C_D} \quad (2.70)$$

and therefore

$$W_{sink,new} = V_{new} \frac{C_D}{C_L} \quad (2.71)$$

$$W_{sink,new} = \frac{V_{old}}{A} \frac{C_D}{C_L} \quad (2.72)$$

which can be simplified to

$$W_{sink,new} = \frac{W_{sink,old}}{A} \quad (2.73)$$

Using these relationships between old and new airspeeds and sink rates, it is

possible to adjust any sink polar of the form

$$W_{sink} = f(V) \quad (2.74)$$

to account for a change in wing loading. Specifically, taking the form of the sink polar equation found in Welch, Welch, and Irving [64] and used by Pennycuick [74] it can be found that:

$$W_{sink} = \frac{a}{V} + bV^3 \quad (2.75)$$

$$W_{sink,new} = \frac{A^2 a}{V_{new}} + \frac{b}{A^2} V_{new}^3 \quad (2.76)$$

This, combined with measurements of the weight and wing area of Black Vultures and Turkey Vultures taken by Graves [76] can be used to adjust the corrected sink polar of the White-Backed Vulture for the wing loading of the geometrically similar Black Vulture

$$W_{sink} = \frac{5.724}{V^2} + (2.165 * 10^{-4})V^3 \quad (2.77)$$

and Turkey Vulture

$$W_{sink} = \frac{4.410}{V^2} + (2.810 * 10^{-4})V^3 \quad (2.78)$$

The resulting sink polars and lift to drag ratios are plotted in Figure 2.10. It should be noted that the maximum lift to drag ratios are equal because the birds were assumed to be perfectly geometrically similar and Reynolds number effects were ignored.




Figure 2.10. Estimated sink rates and lift to drag ratios of White-Backed Vultures, Black Vultures, and Turkey Vultures

2.5 Choosing a UAV for Microlift Soaring

The Vulture UAV (shown in Figure 2.11) was chosen for microlift soaring because its performance is similar to the Carbon Dragon and to that of a Turkey Vulture. A comparison of their sink polars can be found in Figure 2.12. The Vulture UAV has the lowest sink rate of the three, and a maximum lift to drag ratio only slightly lower than that of the Carbon Dragon, and considerably better than that of a Turkey Vulture. The safe airspeed range of the Vulture UAV lies within the range spanned by a Turkey Vulture and the Carbon Dragon. Because both the Carbon Dragon and Turkey Vultures are known to soar in microlift, it is reasonable to assume that the Vulture UAV - with similar flight performance - will also be able to soar in microlift. In addition, the Vulture UAV has already performed several successful autonomous soaring flight. In addition to comparable sink performance, the turn radii of the Carbon Dragon, a Turkey Vulture, and the Vulture UAV are similar, as was shown in Figure 1.10.



Figure 2.11. The author launching the Vulture UAV. Photo by Jack Langelaan.

The Vulture UAV is based on a Magellan-E-XL radio controlled sailplane available from ICARE. It is equipped with a AXi 2220/16 V2 brushless electric motor (910 kV),

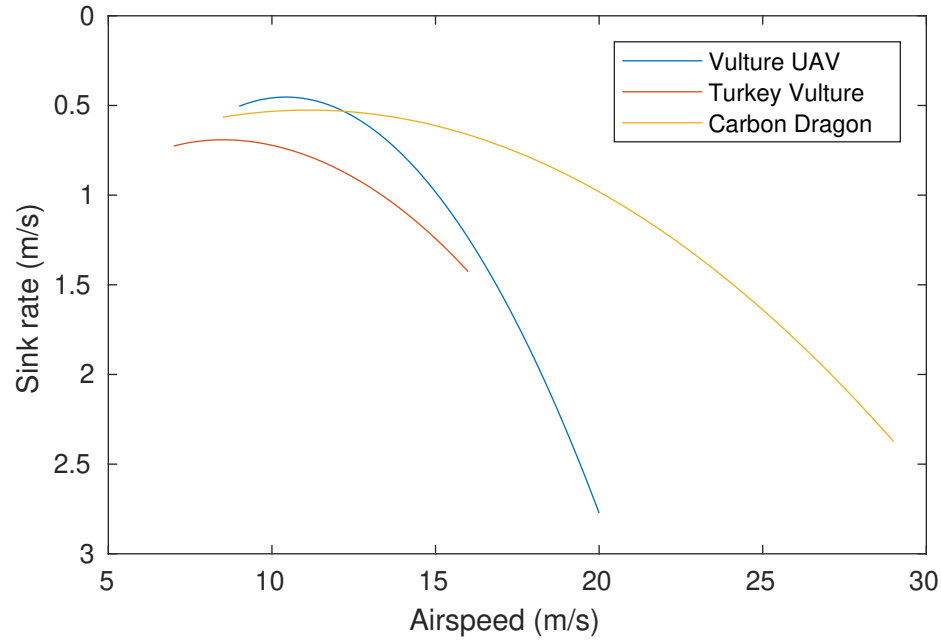


Figure 2.12. Comparison of the sink polars of the Vulture UAV, a Turkey Vulture, and the Carbon Dragon ultralight sailplane.

a two-bladed folding CAMcarbon aero-naut 11x8 propeller, and a 3,000 milliamp-hour 3S lithium polymer motor battery to allow for hand launching and cruising when not soaring. The endurance in a powered level cruise is approximately 20 minutes. The Vulture UAV uses a 3DR Pixhawk autopilot with modified firmware for autonomous flight control, and an ODROID C-1+ handles the higher level autonomy functions such as navigation and soaring. It has a pair of 10 Amp-hour 3.7 Volt lithium polymer avionics batteries in the wings which allow for the servos and onboard computers to run for more than 8 hours. Some basic characteristics of the Vulture UAV can be found in Table 2.2.

Onboard sensors include a pitot-static probe, air temperature and humidity sensors, rate gyroscopes, accelerometers, a GPS receiver, and motor telemetry obtained through the Castle Creations electronic speed control. In addition, the Vulture UAV has a program that calculates the roll disturbance experienced by the aircraft, and the

Table 2.2. Characteristics of the Vulture UAV

Parameter	Value
m	2.2 kg
S	0.456 m ²
b	2.51 m
V_{stall}	~ 8 m/s

Vulture UAV can simulate several types of variometers using the onboard sensors.

The most often used of the simulated variometers is the relative netto variometer.

The Vulture UAV has performance that is expected to be adequate for soaring in microlift, is equipped with a comprehensive set of sensors, and has proven its ability to soar autonomously. All of this makes the Vulture UAV an ideal platform for researching microlift soaring and thermal strand tracking, and is why it was decided to use a simulated version of the Vulture UAV to test various methods of tracking thermal strands.

Chapter 3 |

Tracking a Thermal Strand

This chapter starts by discussing some of the challenges faced when attempting to track a thermal strand, with a particular focus on observability of strand characteristics and sensor noise. From there, the likelihood of approaching a thermal strand at a given angle is investigated. Two potential methods of tracking an ideal thermal strand are shown to be unstable, and a third method is investigated and stabilized although it is ultimately shown to have significant limitations. Finally, a case is made to solve the instability problem by directly commanding the angle at which the aircraft crosses the thermal strand, and a method is presented for estimating the location and characteristics of a thermal strand using an unscented Kalman filter.

3.1 The Challenges of Tracking a Thermal Strand

Following a thermal strand presents unique challenges. This includes a lack of existing methods that can be applied, problems with observability, atmospheric turbulence, sensor noise, and a lack of any visual references indicating the presence of a thermal strand.

Part of the difficulty in tracking a thermal strand is that merely adapting traditional

thermal soaring techniques (Reichmann’s technique [16], etc.) does not work when trying to follow a thermal strand, and so new techniques need to be developed. Most thermal centering approaches require the radial symmetry of a thermal for stability, and fail when presented with the axial symmetry of a thermal strand. Furthermore, some basic thermal strand following techniques that are simple extensions of thermal soaring can be proven to be unstable when applied to an idealized thermal strand. For example, the thermal centering technique of rolling towards the lift and turning tighter when lift decreases does not work for tracking thermal strands, as discussed in Section 3.3.1. These issues preclude the use of several previously developed autonomous soaring controllers that were based on these traditional soaring techniques, and require that new specialized techniques for tracking thermal strands be developed. Furthermore, these techniques must not be particularly computationally intensive, otherwise the resulting controller will not be able to be implemented onboard the aircraft.

Observability can be a major issue when tracking a thermal strand. Determining whether a thermal strand is to the left or right of the aircraft, and when the aircraft crosses the center of the strand is critical to being able to follow along the strand. The variometer is good at determining where the center of the strand is when flying perpendicular to the strand axis, but loses observability when approaching flying parallel to the strand. The roll disturbance detector is good at determining whether the strand is to the left or right of the aircraft, but loses observability both when the aircraft is flying perpendicular to the strand and when it is flying near the strand axis. The difficulty in maintaining observability of strand strength and width is illustrated in Figures 3.1, 3.2, and 3.3.

These losses of observability do not occur in thermalling flight unless the aircraft is perfectly centered around the thermal. Daugherty and Langelaan developed a thermal centering controller that seeks the maximum climb rate perturbing the turn radius




Figure 3.1. Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying across an ideal thermal strand at 90 degrees to the strand axis. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.




Figure 3.2. Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying across an ideal thermal strand at 45 degrees to the strand axis. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.




Figure 3.3. Vertical wind velocity and roll disturbance encountered by the simulated Vulture UAV when flying along the strand axis of an ideal thermal strand. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

and hence constantly varying distance from the center of the thermal, which also solves the observability problem [36]. A similar approach could be used in tracking a thermal strand by varying the aircraft’s distance from the center of the strand and orientation relative to the strand.

Atmospheric turbulence and sensor noise present a significant challenge that must be overcome in order to track a thermal strand. As was seen in Figures 2.3, 2.4, and 2.5, there is significant variation about the mean vertical wind velocity profile of a thermal strand that is caused by small-scale turbulence in the atmosphere. In the data provided by Sullivan and Patton, these variations have a standard deviation of 0.5 meters per second or more, depending on the particular strand. Add to this the sensor noise from the variometer, and the variometer signal caused by crossing a thermal strand becomes very noisy. The roll disturbance detector is even more sensitive to small-scale atmospheric turbulence, and simulated flights through the LES wind field

revealed that the signal generated by crossing a strand has a much smaller amplitude than the noise generated by turbulence. The results from simulated flights through the LES data were compared with saved roll disturbance detector measurements from old test flights, and the variations in roll disturbance detector measurements were of similar magnitude. The sensor outputs generated by the simulated Vulture UAV crossing an ideal thermal strand while flying at an airspeed of 11.25 meters per second can be seen in Figure 3.4, where Gaussian random noise of an appropriate magnitude to simulate both turbulence and sensor noise was added to the true vertical wind velocity and roll disturbance of the idealized strand model.




Figure 3.4. Variometer and roll disturbance measurements obtained by the simulated Vulture UAV when flying across an ideal thermal strand at 45 degrees to the strand axis. Compare to Figure 3.2. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second, $\dot{h}_{measured} = \dot{h} + N(0, 0.75^2)$, $C_{l,res,measured} = C_{l,res} + N(0, 0.035^2)$. Note that the added Gaussian noise includes contributions from both turbulence and sensor noise.

It should be noted that the roll disturbance detector is reacting to a spanwise change in local angle of attack - caused by vertical air motion from thermal strands or turbulence - and therefore the ratio of signal generated by the thermal strand to

noise generated by turbulence cannot be altered by changing wing loading or airspeed. However, the ratio of signal to total noise can be slightly by changing wing loading or airspeed. If the airspeed increases, the spanwise change in angle of attack from vertical air motion is decreased, and therefore the magnitude of the signal generated by a thermal strand is decreased, as is the noise generated by turbulence. Similarly, increasing the wing loading will increase the moment of rotational inertial about the longitudinal axis, will make the aircraft have a lower roll acceleration for a roll moment disturbance of a given magnitude, and will decrease both the signal generated by the thermal strand and the noise generated by turbulence. (Increasing the wing loading will also increase the airspeed required to fly at a given lift coefficient.) The noise contribution from sensors is only a function of the sensors used. Therefore increasing airspeed or wing loading will slightly decrease the ratio of signal to total noise, and decreasing airspeed or wing loading will slightly increase the ratio of signal to total noise.

Unlike thermals and cloud streets, there are typically no clear visual indications of the presence of a thermal strand. In order to track a thermal strand, the pilot is forced to rely on feel and instruments rather than visual methods. While there has been some research done on computer vision and recognition of cumulus clouds as a way to guide autonomous sailplanes to lift [80, 81], the majority of UAV sailplanes already rely on instruments alone for their navigation. Even so, the highly chaotic nature of microlift and thermal strands [17] makes it difficult to design a controller that is able to reliably track a thermal strand and extract energy.

3.2 Determining Likely Approach Angles

While it might be assumed that there is an equal likelihood of approaching a thermal strand from any angle, this is not the case. It is significantly more likely that the aircraft will approach the thermal strand from an angle close to 90 degrees than from an angle close to 0 degrees.

Two regions bounded by thermal strands were chosen from the LES data, and the bounding strand centerlines marked. The chosen regions can be seen in Figure 3.5. Twenty thousand starting locations were picked within the two regions, and each location was assigned a heading. The locations and headings were determined by generating uniformly distributed random numbers. The aircraft's flight path was projected from the starting location along the corresponding heading until it intersected one of the thermal strands, and the angle at which the aircraft would cross the strand was calculated. Using the set of 20,000 crossing angles, the probability of crossing the strand at a given angle was calculated, and the results can be seen in Figure 3.6.

The results of the analysis need to be taken into account when designing a controller to track thermal strands. The chance of the aircraft encountering a thermal strand at an initial crossing angle between 80 and 90 degrees is more than 4 times as likely as encountering a thermal strand at an initial crossing angle between 0 and 20 degrees, and there is a greater than 55 percent chance of encountering the strand at an initial crossing angle between 70 and 90 degrees. With this in mind, it is clear that for a thermal following controller to be successful it must be able to reliably track a thermal strand when approaching from a steep angle.




Figure 3.5. Regions selected for analysis of the probability of approaching a thermal strand from a given angle. Regions are marked in white. Large eddy simulation cross-section taken at $z/z_i = 0.12$.

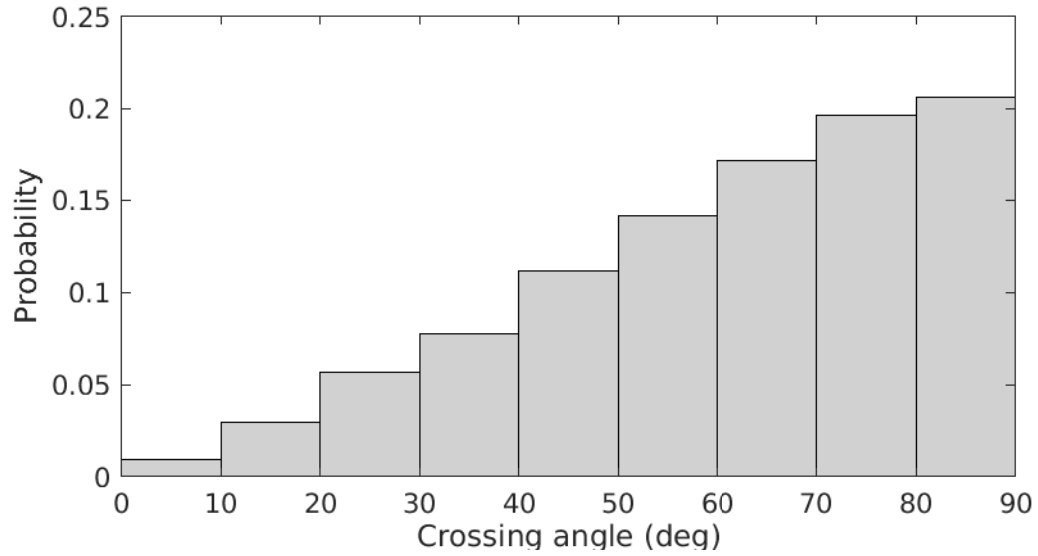


Figure 3.6. Histogram of the probability of approaching a thermal strand from a given angle.

3.3 Two Unstable Methods of Thermal Strand Following and a Stabilization of One Method

3.3.1 Instability of Direct Feedback Methods

One of the basic principles when performing thermal centering in a manned sailplane is to "roll into the lift". As the aircraft flies into the thermal, one wing is closer to the center of the thermal and is pushed upwards more than the other, which rolls the aircraft away from the center of the thermal. The pilot feels this, realizes that there is more lift on one side of the aircraft than the other, and rolls the aircraft towards the region of higher lift and towards the center of the thermal. When rolling into the lift, the pilot is essentially placing a gain on the felt roll disturbance and is using the resulting signal as a roll input.

If the principle of rolling into the lift is applied to a thermal strand; however, the result is an unstable controller that is unable to track a thermal strand. This is because the magnitude of the roll disturbance decreases both with an increase in crossing angle and when near the center of the thermal strand, and because of lag in the response to the controller's commands (instability resulting from lead or lag is discussed in Section 3.3.2).

Let's consider an aircraft equipped with a simple proportional controller that takes in roll disturbance and outputting a bank angle command. When the aircraft flies further from the strand, the magnitude of the roll disturbance decreases and so the bank angle and turn rate will decrease. In addition, when the angle at which the aircraft is approaching the strand increases, the roll disturbance, bank angle, and turn rate all decrease again. The effect of these decreases in turn rate is to gradually

increase how far the aircraft deviates from the strand with each crossing until it reaches a point at which the roll disturbance is decreasing faster than the controller can bring the aircraft back towards the strand. This is illustrated in Figure 3.7, which shows a simulation of an aircraft with perfect sensors and a roll disturbance feedback controller attempting to track an ideal thermal strand.




Figure 3.7. Simulated flight path of an aircraft with bank angle commanded using roll disturbance feedback along an ideal thermal strand. Sensor noise was not included in this simulation. Dotted lines indicate the 2σ distance from the centerline. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

Multiple variations of this controller were tested to see if there was a combination of linear and or nonlinear gains that would result in the aircraft tracking the strand indefinitely. Applying the thermal centering principle of banking less when in higher lift (near the center of the thermal) and banking harder when in low lift (now further away from the center) did slightly improve performance, but it was still unable to perform a stable track along the ideal thermal strand.

3.3.2 Instability of Bank Angle Switching

One of the simplest conceptual designs for a thermal strand tracking controller is one that always attempts to turn back towards the thermal strand. This controller in its simplest form will switch the commanded bank angle from positive to negative (right bank to left bank) based on whether the strand is to the right or to the left. The stability of the controller can be determined by looking at the angle at which the aircraft crosses over the longitudinal axis of the strand and taking the limit as the crossing number goes to infinity.

There are various ways to determine whether the strand is to the left or the right using flight path information and variometer and roll disturbance measurements; however, for any real system the controller will always be unstable and result in the aircraft losing its track along the strand. The proof of this is as follows, with associated diagram found in Figure 3.8.

Assume that the aircraft is already tracking along the strand by performing a series of turns, alternating direction each time the central axis of the strand \overleftrightarrow{OL} is crossed. The process of sensing the peak crossing at P_n - where n denotes the n th crossing of the strand axis - rolling into the turn, and performing the turn is represented by a line segment $\overline{P_n A_n}$, an arc $A_n \widehat{B_n} C_n$, and another line segment $\overline{C_n P_{n+1}}$ at which point the aircraft crosses over the strand axis again at point P_{n+1} .

Proof. Let:

D_n be the center of arc $A_n \widehat{B_n} C_n$

B_n be located such that $\overline{B_n D_n} \perp \overleftrightarrow{OL}$

E_n on $\overleftrightarrow{B_n D_n}$ such that $\overline{A_n E_n} \perp \overline{B_n D_n}$

F_n on $\overleftrightarrow{B_n D_n}$ such that $\overline{C_n F_n} \perp \overline{B_n D_n}$

Assume:




Figure 3.8. Illustration of the instability of bank angle switching controllers when attempting to track along a thermal strand. Vehicle flight path is depicted in red, and \overline{OL} is the longitudinal axis of the strand. This depicts the case when $\overline{P_n A_n} > \overline{C_n P_{n+1}}$; however, the geometric relations shown are also valid for $\overline{P_n A_n} < \overline{C_n P_{n+1}}$.

$$\overline{P_n A_n} > \overline{C_n P_{n+1}}$$

Then:

$$\widehat{A_n B_n} < \widehat{B_n C_n} \text{ because } \overline{P_n A_n} > \overline{C_n P_{n+1}}$$

$$\angle A_n D_n B_n < \angle B_n D_n C_n \text{ because } \widehat{A_n B_n} < \widehat{B_n C_n}$$

$$A_n D_n \perp P_n A_n \text{ because } A_n D_n \text{ is a radius of } \widehat{A_n B_n C_n}$$

$$\angle LP_n A_n = \angle A_n D_n B_n \text{ because } \overleftrightarrow{A_n D_n} \text{ intersects } \overleftrightarrow{OL}, \overline{A_n D_n} \perp \overline{P_n A_n}, \text{ and } \overline{B_n D_n} \perp \overleftrightarrow{OL}$$

$$\angle OP_{n+1} C_n = \angle B_n D_n C_n \text{ because } \overleftrightarrow{C_n D_n} \text{ intersects } \overleftrightarrow{OL}, \overline{C_n D_n} \perp \overline{P_{n+1} C_n}, \overline{B_n D_n} \perp \overleftrightarrow{OL}$$

$$\angle LP_{n+1} A_{n+1} > \angle LP_n A_n$$

□

This proof demonstrates that if $\overline{P_n A_n} > \overline{C_n P_{n+1}}$, then $\angle LP_{n+2} A_{n+2} > \angle LP_{n+1} A_{n+1}$, and it can be concluded that each successive crossing angle will be larger than the last. This progressive increase in crossing angle will continue until the aircraft no longer crosses over the strand axis, as can be seen in the right portion of Figure 3.8.

A similar proof can be constructed by instead assuming that $\overline{P_n A_n} < \overline{C_n P_{n+1}}$, and it can be shown that $\angle LP_{n+1} A_{n+1} < \angle LP_n A_n$. From this it can be concluded then that if $\overline{P_n A_n} < \overline{C_n P_{n+1}}$ each successive crossing angle will be smaller than the last.

In the case when $\overline{P_n A_n} = \overline{C_n P_{n+1}}$ and the flight path is symmetric about the strand axis, the proof can be repeated to show that $\angle LP_{n+1} A_{n+1} = \angle LP_n A_n$, and $\angle LP_{n+2} A_{n+2} = \angle LP_{n+1} A_{n+1}$. Therefore when the flight path is perfectly symmetric about the strand axis, each successive crossing angle will be equal to the last.

From all of this it can be concluded that a bank angle switching controller will be stable if and only if the flight path of the vehicle is symmetric about the strand axis, which is not possible for a real controller of this type. The time it takes to

identify the strand peak and for the vehicle to respond to commands tends to make $\overline{P_n A_n} > \overline{C_n P_{n+1}}$. In theory, a perfect predictor could generate the correct command lead such that $\overline{P_n A_n} = \overline{C_n P_{n+1}}$; however, with anything except a perfect predictor $\overline{P_n A_n} \neq \overline{C_n P_{n+1}}$ and the aircraft would gradually lose its track of the strand centerline.

The change in the angle at which the aircraft crosses the centerline can be found by: $\overline{B_n D_n} \perp \overleftrightarrow{OL}$ therefore $\angle P_n D_n B_n = \angle B_n D_n P_{n+1}$. $\angle P_n D_n A_n = \arctan(P_n A_n / A_n D_n)$ and $\angle C_n D_n P_{n+1} = \arctan(P_{n+1} C_n / D_n C_n)$. $D_n C_n = A_n D_n$, therefore $\angle O P_{n+1} C_n = \angle L P_{n+1} A_{n+1} = \angle L P_n A_n + (\arctan(P_n A_n / A_n D_n) - \arctan(P_{n+1} C_n / A_n D_n))$

The instability of this type of controller was demonstrated in simulations of the Vulture UAV tracking an ideal thermal strand with zero sensor noise, as can be seen in Figure 3.9. In the simulation it is assumed that the controller can perfectly detect when the aircraft crosses over the strand centerline, and the only delay in the aircraft's response is the roughly 1 second that it takes to roll from a 15 degree left bank to a 15 degree right bank, or vice-versa, as simulated by the proportional gain roll dynamics (see Section 2.3.3). The simulation was then repeated using the same initial conditions but with a 1.5 second lag in centerline detection representing the time constant of the variometer, as can be seen in Figure 3.10. In both cases, the simulated flight path of the aircraft matches the geometric prediction, with the simulation that includes additional lag diverging in fewer crossings of the strand centerline, as would be expected.

It should further be noted that in Figure 3.9 after the second crossing of the strand centerline the aircraft is no longer obtaining a net energy gain (including aircraft sink rate) from tracking the strand. This is because the aircraft is spending too much time flying in the sink surrounding the strand and not enough time flying in the lift close to the strand centerline.

Multiple versions of this controller were tested, including several different methods




Figure 3.9. Simulated flight path demonstrating divergence of a bank angle switching controller with perfect identification of the strand axis. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the sign of the commanded bank angle was flipped. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.




Figure 3.10. Simulated flight path demonstrating divergence of a bank angle switching controller with perfect identification of the strand axis but a simulated 1.5 second variometer lag. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the sign of the commanded bank angle was flipped. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

of peak detection. Furthermore, a controller was tested which estimated where the peak will occur in an attempt to eliminate the lag caused by response time to roll commands, as was one that determined which way to bank based on the slope and concavity of recent variometer and roll disturbance detector measurements. Ultimately, each controller exhibited the instability described above.

3.3.3 Discrete Variation of Heading and a Stabilization of the Method

It has been shown that a controller based on reversing the aircraft's bank angle will cause the angle at which the aircraft crosses the strand axis to approach either 0 or 180 degrees for any real controller. One readily apparent solution is to command discrete heading angles instead, and therefore guarantee that the angle at which the aircraft crosses the strand axis will not change. This of course will only work for a perfectly straight thermal strand - if the thermal strand curves, then the crossing angle will change.

It should be noted that if the commanded heading is altered incrementally until the aircraft encounters the strand axis again, the same heading instability found in the bank angle controller will occur. This has been demonstrated in simulations of the Vulture UAS encountering an ideal strand with zero sensor noise, as can be seen in Figure 3.11. The controller was configured to command headings in increments of 22.5 degrees, and increment the heading every time the aircraft came within 1 degree of the commanded heading, and would reverse the sign of the increment every time the central axis of the strand was crossed. The usage of discrete heading increments increases the rate at which the controller diverges, as can be seen by comparing Figures 3.9 and 3.11. The two simulations have identical configurations with the exception of the different controllers. In both simulations it is assumed that the aircraft can perfectly detect the strand centerline.

Rather than continuously incrementing the heading command, the heading command can be flipped between two discrete values depending on which side of the aircraft the strand is on. This can be determined using the variometer, the roll disturbance detector, or a combination of the two. Figure 3.12 shows the ideal case of

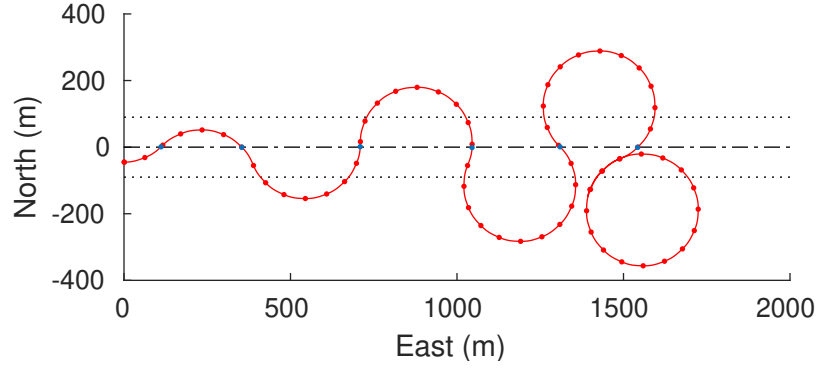


Figure 3.11. Simulated flight path demonstrating instability of a controller that alters heading in discrete increments with perfect identification of the strand axis. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate where the commanded turn direction was flipped. Red dots indicate where the heading command was incremented. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

a heading switching controller with perfect sensing of the strand centerline tracking along an ideal strand. This approach has some significant limitations, although it is stable for a perfectly linear thermal strand.

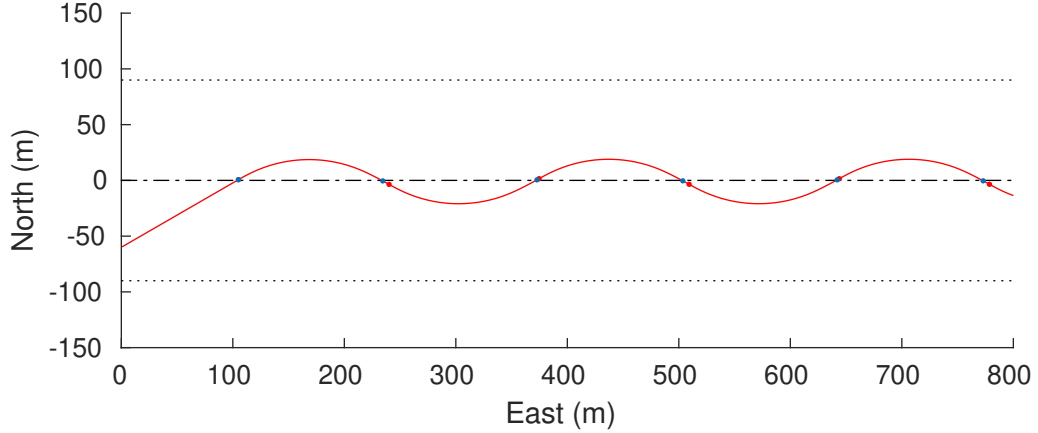


Figure 3.12. Ideal track of a strand using a heading switching controller. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate when the commanded heading was flipped. Red dots indicate when the new heading command was applied. $\psi_0 = 60$ degrees, $\Delta\psi = 60$ degrees, $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

First, strand tracking is limited to strands whose longitudinal axes lie within the heading differential of the aircraft's initial course. This is because when the aircraft makes the initial discrete heading change, that new heading must cause the aircraft

to be headed back towards the strand axis.

Second, having an excessively large discrete heading change is impractical because it can result in the aircraft flying far from the center of the strand before approaching again which reduces the average rate of energy gain. If the aircraft deviates too far from the strand, the energy lost to aerodynamic drag and the sink surrounding the thermal strand will exceed the energy gained by flying across the strand.

Third, the first two limitations result in this form of controller only being able to track strands that are approached at shallow angles. It is significantly more likely that the aircraft will approach a thermal strand at a steep angle than at a shallow angle, which limits the utility of this type of controller.

Fourth, as the angle at which the aircraft will first cross the strand approaches the heading differential, one side of the resulting S-curve across the strand becomes smaller, as seen in Figure 3.13. In this case, the approach angle has been changed by only 10 degrees, but it has a massive effect on the flight path of the aircraft. Note that the aircraft is required to reach the last commanded heading before it is given a new heading command. Removing this constraint would result in both halves of the S-curve becoming smaller as the initial crossing angle approaches the heading differential. While in the simulation it was assumed that the aircraft would be able to perfectly detect the strand centerline, flying close to the strand centerline at a very shallow angle will make it very difficult to detect a peak in the variometer measurements or a sign change in the roll disturbance measurements.

As previously noted, the heading switching controller is stable for a perfectly straight ideal thermal strand. However, the aforementioned limitations and the extreme sensitivity to the angle at which the aircraft approaches the thermal strand make this type of controller unsuitable for real applications.

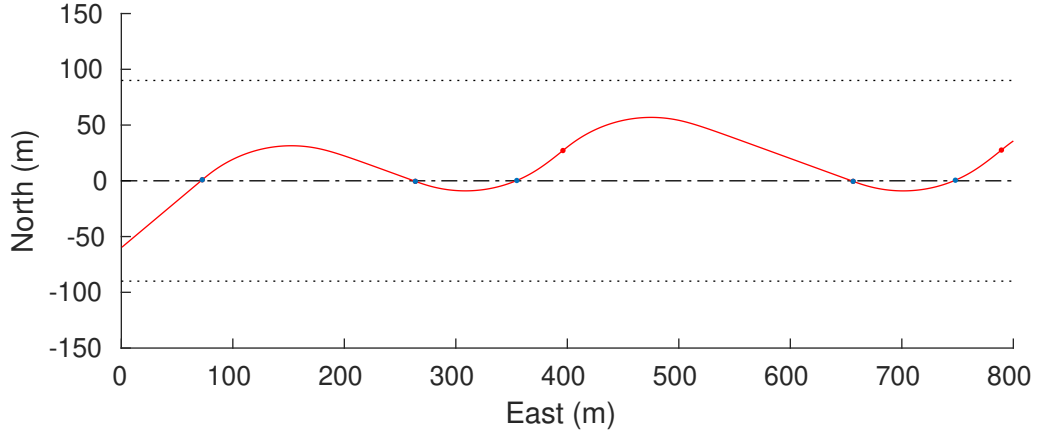


Figure 3.13. Ideal track of a strand using a heading switching controller. Dotted lines indicate the 2σ distance from the centerline. Blue dots indicate when the commanded heading was flipped. Red dots indicate when the new heading command was applied. $\psi_0 = 50$ degrees, $\Delta\psi = 60$ degrees, $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second.

3.4 Estimating Strand Properties Using an Unscented Kalman Filter

A better solution to the instability of the bank angle switching controller is to directly command the angle at which the aircraft crosses the strand. This requires an estimate of the strand's orientation relative to the aircraft but guarantees stability as long as the estimate is accurate. Given the high sensor noise involved and the foreknowledge about the structure of a thermal strand, it seems reasonable to provide the estimate of relative strand location and orientation using a Kalman filter.

To properly make use of knowledge about the structure of thermal strands, the Kalman filter state vector needs to capture both the relative position and orientation of the thermal strand with respect to the aircraft and information about the vertical wind profile of the strand. Using the idealized strand model, the vertical wind profile can be represented by two parameters: σ_{str} and \dot{h}_{max} , which are the standard deviation of the Gaussian model and the maximum vertical air velocity of the strand. The location

of the thermal strand with respect to the aircraft can conveniently be expressed by drawing a vector from the aircraft to the centerline of the thermal strand at such an angle that the vector is perpendicular to the centerline. The magnitude of the vector is d , the distance from the aircraft to the strand, and the heading of the vector relative to the aircraft's current heading is ψ_{str} . The distance to the strand, d , is required to always be positive, and the relative heading to the strand, ψ_{str} , is given the range $-pi \leq \psi_{str} \leq pi$ for convenience when designing the heading controller. The Kalman filter state vector is

$$\mathbf{x} = \begin{bmatrix} d \\ \psi_{str} \\ \dot{h}_{max} \\ \sigma_{str} \end{bmatrix} \quad (3.1)$$

and the definition of the variables is illustrated in Figure 3.14.

It was decided to use an unscented Kalman filter to estimate the thermal strand state because of the nonlinearities in the equations. A particle filter would also be an acceptable choice to capture the nonlinearities; however, a particle filter requires significantly more computations per time step than an unscented Kalman filter. In this case, it was decided that the improvements to performance when dealing with the nonlinear system would likely not be sufficient to justify the additional computational cost.

The unscented Kalman filter presented here is based on the unscented Kalman filter presented in a paper by van der Merwe and Wan [82].

Given a Kalman filter state and state covariance at time step $k - 1$, the time update on the Kalman filter state was performed as follows. Given n states, sigma




Figure 3.14. Definition of thermal strand parameters describing the location of the strand axis \overline{OL} relative to the aircraft and the updraft cross-section.

points are distributed using

$$\boldsymbol{\chi}_{k-1}^+ = \begin{bmatrix} \hat{\mathbf{x}}_{k-1}^+ + \eta\sqrt{\mathbf{P}_{k-1}} & \hat{\mathbf{x}}_{k-1}^+ - \eta\sqrt{\mathbf{P}_{k-1}} \end{bmatrix} \quad (3.2)$$

$$\eta = \sqrt{n} = 2 \quad (3.3)$$

Then state propagation was performed on each of the sigma points. It is assumed that the aircraft's state $(DCM_{1B}, \mathbf{V}^B, \dot{\psi}, \dot{\phi})$ is known perfectly. This is a reasonable assumption because the covariance of the aircraft's internal state estimate is much, much smaller than the covariance of the thermal strand state, as will be seen later. If

the time propagation of the filter state is given by

$$\boldsymbol{\chi}_k^- = \mathbf{F}(\boldsymbol{\chi}_{k-1}^+) \quad (3.4)$$

then the state propagation of the i th sigma point is as follows:

$$\boldsymbol{\chi}_{k-1}^{(i)+} = \begin{bmatrix} d_{k-1}^{(i)+} \\ \psi_{str,k-1}^{(i)+} \\ \dot{h}_{max,k-1}^{(i)+} \\ \sigma_{str,k-1}^{(i)+} \end{bmatrix} \quad (3.5)$$

$$d_k^{(i)-} = d_k^{(i)-} - \left((DCM_{1B} * \mathbf{V}^{\mathbf{B}})^T * \begin{bmatrix} \cos(\psi_{str,k-1}^{(i)+}) \\ \sin(\psi_{str,k-1}^{(i)+}) \\ 0 \end{bmatrix} \right) \quad (3.6)$$

$$\psi_{str,k}^{(i)-} = \psi_{str,k-1}^{(i)+} + (\dot{\psi} * \Delta t) \quad (3.7)$$

$$\dot{h}_{max,k}^{(i)-} = \dot{h}_{max,k-1}^{(i)+} \quad (3.8)$$

$$\sigma_{str,k}^{(i)-} = \sigma_{str,k-1}^{(i)+} \quad (3.9)$$

$$\boldsymbol{\chi}_k^{(i)-} = \begin{bmatrix} d_k^{(i)-} \\ \psi_{str,k}^{(i)-} \\ \dot{h}_{max,k}^{(i)-} \\ \sigma_{str,k}^{(i)-} \end{bmatrix} \quad (3.10)$$

The above equations assume that the center of the thermal strand drifts with horizontal winds. This is not actually the case, as thermal strands lie along convergence regions (see Section 1.3); however, horizontal winds are neglected in this simulation, so this assumption has no impact on the results. Were horizontal winds included in the

simulation, it would be necessary to estimate the background horizontal wind velocity and make the appropriate changes to Kalman filter state propagation.

Finally, the a priori state estimate and covariance is calculated by

$$\hat{\mathbf{x}}_k^- = \sum_{i=1}^{2n} \frac{1}{2n} \boldsymbol{\chi}_k^{(i)-} \quad (3.11)$$

$$\mathbf{P}_k^- = \mathbf{Q} + \sum_{i=1}^{2n} \frac{1}{2n} \left(\boldsymbol{\chi}_k^{(i)-} - \hat{\mathbf{x}}_k^- \right) \left(\boldsymbol{\chi}_k^{(i)-} - \hat{\mathbf{x}}_k^- \right)^T \quad (3.12)$$

Now the measurement update was performed. First, a new set of sigma points (denoted by a tilde) were generated using

$$\tilde{\boldsymbol{\chi}}_k^- = \begin{bmatrix} \hat{\mathbf{x}}_k^- + \eta \sqrt{\mathbf{P}_k^-} & \hat{\mathbf{x}}_k^- - \eta \sqrt{\mathbf{P}_k^-} \end{bmatrix} \quad (3.13)$$

and an estimated measurement was calculated for each sigma point by

$$\boldsymbol{\gamma}_k = \mathbf{H}(\tilde{\boldsymbol{\chi}}_k^-) \quad (3.14)$$

For the i th sigma point, the estimated measurement is as follows:

$$\boldsymbol{\gamma}_k^{(i)} = \begin{bmatrix} \hat{h}^{(i)} \\ \hat{C}_{l,res}^{(i)} \end{bmatrix} \quad (3.15)$$

$$\tilde{\boldsymbol{\chi}}_k^{(i)-} = \begin{bmatrix} \tilde{d}_k^{(i)-} \\ \tilde{\psi}_{str,k}^{(i)-} \\ \tilde{h}_{max,k}^{(i)-} \\ \tilde{\sigma}_{str,k}^{(i)-} \end{bmatrix} \quad (3.16)$$

$$\hat{h}^{(i)} = -0.5 \left(\tilde{h}_{max,k}^{(i)-} \right) + 1.5 \left(\tilde{h}_{max,k}^{(i)-} \right) * \exp \left(\frac{- \left(\tilde{d}_k^{(i)-} \right)^2}{2 \left(\tilde{\sigma}_{str,k}^{(i)-} \right)^2} \right) \quad (3.17)$$

$$C_{l,res} = \left(C_{lp} \frac{b}{2V_0} \cos^2(\phi) \right) * -1.5 \left(\tilde{h}_{max,k}^{(i)-} \frac{\tilde{d}_k^{(i)-}}{\left(\tilde{\sigma}_{str,k}^{(i)-} \right)^2} \right) \exp \left(\frac{- \left(\tilde{d}_k^{(i)-} \right)^2}{2 \left(\tilde{\sigma}_{str,k}^{(i)-} \right)^2} \right) * \sin(\tilde{\psi}_{str,k}^{(i)-}) \quad (3.18)$$

The estimated measurement for the a priori Kalman filter state, the measurement covariance, and the cross covariance are given by

$$\hat{\mathbf{y}}_k = \sum_{i=1}^{2n} \frac{1}{2n} \gamma_k^{(i)} \quad (3.19)$$

$$\mathbf{P}_{y,k} = \mathbf{R} + \sum_{i=1}^{2n} \frac{1}{2n} \left(\chi_k^{(i)-} - \hat{\mathbf{x}}_k^- \right) \left(\chi_k^{(i)-} - \hat{\mathbf{x}}_k^- \right)^T \quad (3.20)$$

$$\mathbf{P}_{xy,k} = \sum_{i=1}^{2n} \frac{1}{2n} \left(\tilde{\chi}_k^{(i)-} - \hat{\mathbf{x}}_k^- \right) \left(\gamma_k^{(i)} - \hat{\mathbf{y}}_k \right)^T \quad (3.21)$$

and the Kalman filter gain is defined to be

$$\mathbf{K}_k = \mathbf{P}_{xy,k} \mathbf{P}_{y,k}^{-1} \quad (3.22)$$

Finally, the a posteriori state estimate and covariance matrix were given by

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + \mathbf{K}_k (\mathbf{y}_k - \hat{\mathbf{y}}_k) \quad (3.23)$$

$$\mathbf{P}_k^+ = \mathbf{P}_k^- - \mathbf{K}_k \mathbf{P}_{y,k} \mathbf{K}_k^T \quad (3.24)$$

Chapter 4 |

Results

The beginning of this chapter deals with the configuration of the idealized thermal strand simulation and some parameters that were common to all of the tested controller designs. The unscented Kalman filter design presented in Section 3.4 is then shown to be able to produce a track of the ideal thermal strand that is stable for all practical purposes. A comparison is then made between the competing command strategies of flying a series of S-curves across the thermal strand and flying straight down the center of the strand. The problem of generating an initial state for the Kalman filter is then discussed, and several different methods are tested and compared. A method was developed for identifying and rejecting initial state estimates that are unlikely to result in the aircraft tracking the strand. Finally, an initialization method is chosen, the controller is tested in the large eddy simulation environment, and the results are presented.

4.1 Configuration of the Idealized Thermal Strand Simulation

After analyzing various other methods of tracking a thermal strand, the Kalman filter described in Section 3.4 was used to estimate strand distance, orientation, strength, and width. This estimate was then provided to a controller that would either track along the center of the strand or repeatedly cross over the centerline at a prescribed angle, making a series of large S-curves across the strand. Various initialization methods were tested and analyzed in an attempt to find one that could reliably provide an initialization of sufficient accuracy to enable the filter to track the strand.

No variometer or roll disturbance lag was simulated in any of the Kalman filter based controllers because with a known variometer or roll disturbance time constant the measurements could be applied at a previous time step so as to largely eliminate the effects of the lag. This would result in the state estimate at time t having the last variometer measurement update at time $t - \tau_{vario}$ and the least roll disturbance measurement update at time $t - \tau_{rdd}$, where τ_{vario} and τ_{rdd} are the variometer and roll disturbance detector lags, respectfully. The primary effect would be to increase the estimated covariance at the current time, but the estimated state and covariance $t - \tau_{vario}$ seconds previously (assuming that the variometer lag is larger than the roll disturbance detector lag) would be identical.

The Kalman filter uses a time covariance given by Equations 4.1 and 4.2. The values shown were generated by starting with a reasonable guess and then tuning in simulation to maintain the state estimate within two standard deviations of the true value 95 percent or more of the time..

$$\mathbf{Q}_{continuous} = \begin{bmatrix} (0.25)^2 & 0 & 0 & 0 \\ 0 & (5\frac{\pi}{180})^2 & 0 & 0 \\ 0 & 0 & (0.1)^2 & 0 \\ 0 & 0 & 0 & (2)^2 \end{bmatrix} \quad (4.1)$$

$$\mathbf{Q}_{discrete} = \mathbf{Q}_{continuous} * (\Delta t)^2 \quad (4.2)$$

For all tests on the ideal thermal strand, the strand was given a width of $\sigma_{str} = 45$ meters and a maximum vertical air velocity of $\dot{h}_{max} = 1.5$ meters per second. This represents a typical strand found in the large eddy simulation data at altitudes between roughly 60 meters ($z/z_i = 0.06$) and 120 meters ($z/z_i = 0.12$).

4.2 Tracking Along an Ideal Strand

The Kalman filter described in Section 3.4 was used to estimate the thermal strand state while a basic controller commanded a heading and bank angle to achieve the desired strand crossing angle. The commanded heading would remain constant (with variations to maintain the same heading relative to the estimated orientation of the strand axis) until aircraft exceeded a distance of $0.5\sigma_{str}$ from the strand axis, at which point it would be commanded to turn back to cross the strand again. The commanded angle at which the aircraft would cross the strand was 20 degrees, and the bank angle was limited to 15 degrees to minimize the sink rate of the aircraft while maintaining an adequate turn radius. The aircraft was commanded to maintain a constant airspeed of 11.25 meters per second which is the best glide airspeed for still air.

This controller design was tested extensively in simulation using the idealized strand model with appropriate measurement noise to simulate the effects of turbulence

and sensor noise on the measurements. The measurements are given by Equations 4.3 and 4.4, and the Kalman filter measurement covariance can be found in Equation 4.5.

$$\dot{h}_{measured} = \dot{h} + N(0, (0.75)^2) \quad (4.3)$$

$$C_{l,res,measured} = C_{l,res} + N(0, (0.035)^2) \quad (4.4)$$

$$\mathbf{R} = \begin{bmatrix} (0.75)^2 & 0 \\ 0 & (0.035)^2 \end{bmatrix} \quad (4.5)$$

When given a reasonable initialization, the controller was able to reliably track the ideal thermal strand. The flight path of the aircraft while tracking an ideal strand can be seen in Figure 4.1. The corresponding state estimates are plotted in Figure 4.2, and the state estimate error and estimated covariance are plotted in Figure 4.3. The initialization values for this simulation were given a normally distributed random error consistent with the initialization covariance given by Equation 4.6.

$$\mathbf{P}_0 = \begin{bmatrix} \left(\frac{45}{4}\right)^2 & 0 & 0 & 0 \\ 0 & \left(15\frac{\pi}{180}\right)^2 & 0 & 0 \\ 0 & 0 & (0.5)^2 & 0 \\ 0 & 0 & 0 & (10)^2 \end{bmatrix} \quad (4.6)$$

The maximum distance from the strand reached by the aircraft after each crossing of the strand axis depends on the Kalman filter state estimate at the time. Estimate errors in the distance do the strand centerline directly impact how far the aircraft actually is from the strand before it starts to turn, but the estimate of strand width also has an impact because the turn point is commanded as a function of strand width.

A Monte Carlo of 2,500 simulated 5 minute flights was conducted to confirm the ability of the controller to intercept and track a strand given random initialization

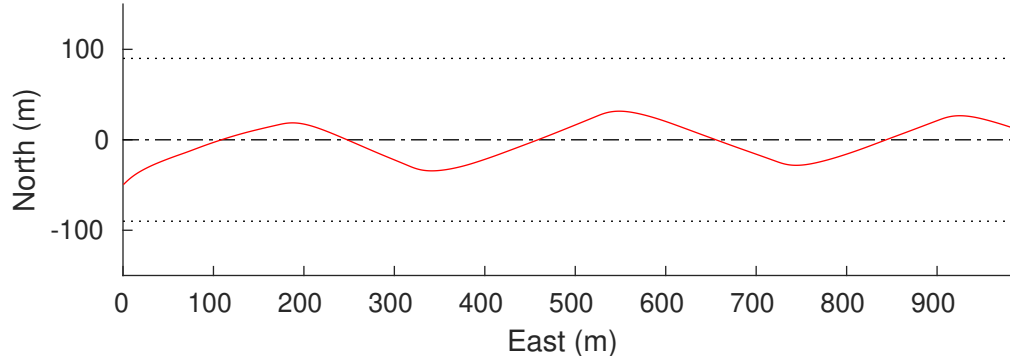


Figure 4.1. Flight path of the basic Kalman filter controller when tracking an ideal strand. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.




Figure 4.2. State estimate and true state of the basic Kalman filter controller when tracking an ideal strand. The true state is plotted in black and the estimated state is plotted in blue.

values as previously described. In each simulation, the aircraft was given a random position and heading so that it would cross the center of the strand in 15 seconds, and would presumably travel along the strand for the remaining 4 minutes and 45 seconds.

The results were filtered for simulations where the final position was more than $0.8\sigma_{str}$. The 47 simulations that met this criteria were manually sorted into cases




Figure 4.3. State estimate error and estimated $\pm 2\sigma$ bounds of the basic Kalman filter controller when tracking an ideal strand. The state estimate error is plotted in blue, and the black dotted lines are the estimated $\pm 2\sigma$ bounds. $\sigma_{str} = 45$ meters, $\dot{h}_{max} = 1.5$ meters per second, $\dot{h}_{measured} = \dot{h} + N(0, 0.75^2)$, $C_{l,res,measured} = C_{l,res} + N(0, 0.035^2)$.

where the aircraft was still tracking the strand but had swung unusually wide, cases where the aircraft had been unable to start tracking the strand, and cases where the aircraft had started tracking the strand but had subsequently diverged. The results are presented in Table 4.1.

Table 4.1. Results of 2,500 5-minute simulations of the aircraft tracking an ideal thermal strand by flying S-curves.

Result	Simulations
Tracked the strand	2487
Unable to start tracking	12
Tracked then diverged	1

In more than 99 percent of the simulations, the aircraft successfully tracked the strand and had a final location within $1 \sigma_{str}$ of the strand centerline. In 12 of the remaining 13 simulations, the aircraft received a Kalman filter initialization that

included one or more states significantly beyond 2 standard deviations of the true value, as determined by the initial covariance. In 1 simulation the aircraft begin tracking the strand and subsequently lost the strand. This represents a total of 196.9 hours spent tracking the ideal thermal strand with the aircraft only losing the strand once.

Multiple long-duration simulations were also conducted in which the aircraft would continuously track the ideal strand for 20 minutes, traveling approximately 13.5 kilometers along the ideal strand. The first 5 kilometers of the flight path from one of these simulations can be seen in Figure 4.4. During this flight, the aircraft had an average rate of energy gain from the atmosphere of 1.33 meters per second - 89 percent of the maximum rate of energy gain rate available. The aircraft had an average net energy rate of 0.83 meters per second, which indicates an average sink rate of 0.50 meters per second.




Figure 4.4. Partial flight path of the basic Kalman filter controller tracking an ideal thermal strand by flying S-curves during a 20 minute simulation. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.

4.2.1 Flying Down the Strand Centerline Versus S-Curves Across the Strand

There are two basic approaches to following along a thermal strand: flying down the center of the strand, or flying a path that repeatedly crosses the strand axis. Flying down the centerline of the strand will maximize the rate of energy gain; however, it

decreases observability of the thermal strand. Intentionally flying a path that causes the aircraft to repeatedly cross the strand axis - while resulting in a lower rate of energy gain - has increased observability of the thermal strand by sweeping across the width of the strand.

Flying along the center of the strand reduces observability of the thermal strand state. All that can be determined is that the aircraft is on or near the centerline of the strand which allows for excellent knowledge of the maximum strength of the strand, but provides no information about the width of the strand or the relative heading angle of the strand. It also provides little about the distance to the strand centerline because of the lack of information about strand heading and width. This lack of observability results in the covariance of the Kalman filter estimate increasing and the quality of the estimate degrading. This in turn causes the aircraft to deviate from its path down the centerline of the strand, and results in a series of unintentional deviations from the strand centerline, as shown in Figures 4.5, 4.6, 4.7.




Figure 4.5. Flight path of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.

The Monte Carlo of 2,500 simulations was run again using a flight controller that would attempt to track down the center of the strand. Again, in more than 99 percent of the simulations the aircraft was able to track the strand for the entire 5 minutes.




Figure 4.6. State estimate and true state of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The true state is plotted in black and the estimated state is plotted in blue.




Figure 4.7. State estimate error and estimated $\pm 2\sigma$ bounds of the basic Kalman filter controller when attempting to fly down the center of an ideal strand. The state estimate error is plotted in blue, and the black dotted lines are the estimated $\pm 2\sigma$ bounds.

However, using the centerline tracking controller there were 4 cases where the aircraft began tracking the ideal thermal strand before wandering off. During 197.5 hours of tracking the ideal strand, the centerline controller's estimate diverged 4 times. This is approximately 4 times as often as the S-curve controller. The results are presented in Table 4.2.

Table 4.2. Results of 2,500 5-minute simulations of the aircraft tracking an ideal thermal strand by flying down the center of the strand.

Result	Simulations
Tracked the strand	2492
Unable to start tracking	4
Tracked then diverged	4

As previously mentioned, flying a series of S-curves across the strand out to a maximum of $0.5\sigma_{str}$ and crossing the strand axis at an angle of 20 degrees resulted in an average rate of energy gain of 1.33 meters per second. The long durations simulations were repeated with the aircraft being commanded to fly down the center of the strand. A partial flight path of the simulation is shown in Figure 4.8, and during the flight the aircraft had an average rate of energy gain of 97 percent of the maximum rate available, or 1.46 meters per second. The average net rate of energy gain was 0.96 meters per second, again indicating an average sink rate of 0.50 meters per second.




Figure 4.8. Partial flight path of the basic Kalman filter controller tracking an ideal thermal strand by flying down the centerline during a 20 minute simulation. The aircraft is traveling from West to East. Black dotted lines indicate $\pm 2\sigma_{str}$ from the strand centerline.

Flying down the centerline of the strand results in a 9.8 percent higher rate of energy gain when compared to flying a series of S-curves across the strand. However, this comes at the cost of knowledge about the thermal strand, as the state estimate covariances when flying down the center of the strand are higher than when flying S-curves across it - particularly for the estimates of d and σ_{str} . This is particularly apparent when comparing the number of times that each controller started to track the strand and then diverged, although the centerline flying method did result in fewer cases where the aircraft was unable to track the strand at all - likely because it will head directly for where it thinks the strand is rather than attempting to flying at a prescribed angle to the strand. For both controllers, the rate of failure given a good initialization is extremely low.

4.3 Solving the Initialization Problem

For all practical purposes the Kalman filter based controller can track an ideal thermal strand indefinitely; however, to do so it requires a reasonably accurate initialization. Therefore an initialization method must be developed that provides a good estimate of the strand state which is not computationally intensive. Furthermore, this initialization method must reliably function when the angle at which the aircraft approaches the strand is near 90 degrees because - as was found in Section 3.2 - it is much more likely that the aircraft will first approach a thermal strand from a steep angle than from a shallow angle. Several different initialization methods were tested to see how often the Kalman filter and aircraft would be able to track the thermal strand.

It was decided to use the S-curve method of tracking a thermal strand because it is the slightly more difficult method to get to initially track the strand, and because once the aircraft is tracking the ideal strand it is less likely to subsequently lose track

of the strand.

The Kalman filter was initialized when the average vertical air motion sensed by the variometer over a given review period of time exceeds a set value. Similarly, if the average variometer measurement during that same period drops below a predetermined value, the Kalman filter was be turned off and the aircraft returned to seeking a thermal strand to follow. Both values were set to 0.5 meters per second. There was also a period of time after starting to attempt initialization during which the "return to seeking" behavior was locked out. This was designed to give the aircraft a chance to latch onto the thermal strand, and was termed the "latching time". The search behavior of the controller was to level the wings and fly in a straight line until the next thermal strand was encountered. While not ideal when attempting to re-find the ideal thermal strand, this method is acceptable when there is more than one strand available to track, as was the case in the large eddy simulation environment and would be the case in the real atmosphere.

For each initialization method, a Monte Carlo analysis was run consisting of 1000 simulations at each of 10 approach angles evenly spaced from 0 to 90 degrees (parallel to the strand to perpendicular to the strand). The aircraft was positioned so that it would initialize the Kalman filter 15 seconds after the start of the simulation, except when the initial position of the aircraft had it flying parallel to the strand. In that case, the Kalman filter initialized as soon as the buffer of saved measurements was full. Flight was simulated for 90 seconds, at which point the controller was checked to see if it was currently tracking a thermal strand or searching for a thermal strand. The number of times that the Kalman filter was initialized was also recorded.

4.3.1 Naïve or Arbitrary Initialization

The simplest way to initialize the Kalman filter is to naïvely assume that the strand is directly in front of the aircraft with a covariance large enough that it hopefully encompasses the true location of the strand. This method of assigning an arbitrary initialization value does have the drawback that it may take a while for the estimate to converge to the true state of the thermal strand. The chosen initialization values represent a fairly typical thermal strand, although in this case care was taken to prevent the initialization value from being too close to the parameters of the ideal strand upon which the initialization method is being tested.

The configuration of the controller with arbitrary initialization is given in Table 4.3. How often the aircraft was able to track the strand as a function of approach angle is shown in Figure 4.9. It initially appears that the naïve initialization method does decently for how simple it is; however, the results are deceptive.


$$\mathbf{x} = \begin{bmatrix} 55 \\ 0 \\ 1.75 \\ 55 \end{bmatrix} \quad (4.7)$$

$$\mathbf{P}_0 = \begin{bmatrix} (1.5\frac{55}{4})^2 & 0 & 0 & 0 \\ 0 & (30\frac{\pi}{180})^2 & 0 & 0 \\ 0 & 0 & (0.75)^2 & 0 \\ 0 & 0 & 0 & (55)^2 \end{bmatrix} \quad (4.8)$$

In instances when it takes a while for the estimate to converge to the true state of the strand, the aircraft will overshoot the strand but start to turn around. By this

Table 4.3. Configuration of the controller testing arbitrary initialization.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	6 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 degrees
Bank angle limit	15 degrees
Airspeed	11.25 meters per second

**Figure 4.9.** Percent of the time that the arbitrary initialization controller can track the strand.

point the aircraft is often beyond $2\sigma_{str}$ from the strand, and the aircraft is unable to complete the turn and return to the strand before 15 seconds has elapsed and the aircraft returns to seeking. While the aircraft is unable to return to the strand in time, it sometimes turns just far enough that at some point in the future it will cross the strand again. If the aircraft has turned far enough that the second crossing occurs within the 90 second simulation, then the Kalman filter will reinitialize and

start trying to track the strand again. This also occurs if the 15 degree bank limit is insufficient to get the aircraft turned around in time. During the Monte Carlo there were many cases where the Kalman filter initialized twice or more during the 90 second simulation.

When attempting to find a reliable initialization method, it is not sufficient for the controller to have eventually tracked the strand - this includes the possibility that the first initialization was of poor quality and was only sufficient to nudge the aircraft back towards the strand after overshooting. For the initialization method to be reliable, it needs to return a state estimate sufficient to track the thermal strand on the first attempt. The naïve initialization is a particularly dramatic example of this, as can be seen in Figure 4.10. When approaching the thermal strand at steep angles - as will most often be the case - this initialization method enables the aircraft to track the strand between 20 and 40 percent of the time. However, in a large number of cases it is taking two or more initializations before the aircraft can track the strand. For example, during the 1,000 simulations where the aircraft approached perpendicular to the strand it never once tracked the strand on the first attempt.

4.3.2 Dynamic Control of Maximum Bank Angle and Approach Angle and the Addition of Variable Airspeed

The decrease in the ability of the controller to track the thermal strand when approaching from nearly perpendicular is the result of the bank angle limit and commanded crossing angle being too restrictive. By looking at the estimated strand state and dynamically varying the bank angle limit commanded crossing angle, the performance of the controller can be significantly improved.

Rather than always approaching such that the aircraft will cross the strand at




Figure 4.10. Percent of the time that the arbitrary initialization controller can track the ideal thermal strand on the first attempt.

a 20 degree angle, the aircraft was commanded to approach at a 60 degree angle if the strand was estimated to be greater than $1.5\sigma_{str}$ from the aircraft. Furthermore, if the aircraft believed it needed to make a turn of 50 degrees or more to reach the commanded heading, the roll limit was increased from 15 degrees to 30 degrees. In addition, the airspeed controller was turned on, which will decrease the airspeed of the Vulture UAV - and subsequently the turning radius - when flying through lift.

The Monte Carlo simulation of the Kalman filter controller with arbitrary initialization was repeated using the dynamic approach method. The configuration of the controller is given in Table 4.4, and the results are shown in Figure 4.11. The initial state and covariance are identical to the arbitrary initialization values.

The addition of variable desired approach angle and bank angle limit, as well as the addition of variable airspeed, has significantly improved how often the aircraft

Table 4.4. Configuration of the controller testing arbitrary initialization with dynamic approach.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	6 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 / 60 degrees
Bank angle limit	15 / 30 degrees
Airspeed	Best glide
Approach method	Dynamic




Figure 4.11. Percent of the time that the arbitrary initialization controller can track the ideal thermal strand when using the dynamic strand approach method.

can track the strand when approaching from close to 90 degrees. However, with an average rate of tracking of just less than 60 percent, there is considerable room for improvement.

4.3.3 Left or Right

Rather than having a single arbitrary guess of the thermal strand state, the average roll disturbance recorded in the measurement buffer could be used in an attempt to determine whether the thermal strand lies to the left or the right of the aircraft. Then a choice can be made between two initialization values found in Equation 4.9.

$$\mathbf{x} = \begin{cases} \begin{bmatrix} 55 & 45\frac{\pi}{180} & 1.75 & 55 \end{bmatrix}^T & \bar{C}_{res} < 0 \\ \begin{bmatrix} 55 & -45\frac{\pi}{180} & 1.75 & 55 \end{bmatrix}^T & \bar{C}_{res} > 0 \end{cases} \quad (4.9)$$

$$\mathbf{P}_0 = \begin{bmatrix} (1.5\frac{55}{4})^2 & 0 & 0 & 0 \\ 0 & (30\frac{\pi}{180})^2 & 0 & 0 \\ 0 & 0 & (0.75)^2 & 0 \\ 0 & 0 & 0 & (55)^2 \end{bmatrix} \quad (4.10)$$

Table 4.5. Configuration of the controller testing initialization that assumes the strand is either to the left or the right.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	6 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 / 60 degrees
Bank angle limit	15 / 30 degrees
Airspeed	Best glide
Approach method	Dynamic

When compared to arbitrarily initializing the strand in front of the aircraft, this method produces a substantial improvement in tracking the strand at approach angles between 0 and 60 degrees. However, there is a large decrease in performance when the aircraft is approaching nearly perpendicular to the strand axis. This can be seen

in Figure 4.12.




Figure 4.12. Percent of the time that the controller which initializes assuming the strand is either to the left or right can track the ideal thermal strand.

4.3.4 Left, Right, or Straight Ahead

In an attempt to improve the ability of the 2 case controller to track a strand when approaching from nearly perpendicular, a third initialization case was added. If the average of the roll disturbances recorded in the measurement buffer has a magnitude less than 0.006, then the strand is assumed to be directly ahead. Otherwise, the sign of the average of the roll disturbance measurements is used to determine if the strand is to the left or the right. The initialization states and associated average roll disturbance criteria can be found in Equation 4.11, and the associated initial covariance can be found in Equation 4.12.

$$\mathbf{x} = \begin{cases} \begin{bmatrix} 55 & 45\frac{\pi}{180} & 1.75 & 55 \end{bmatrix}^T & \bar{C}_{res} < -0.006 \\ \begin{bmatrix} 55 & 0 & 1.75 & 55 \end{bmatrix}^T & 0.006 < \bar{C}_{res} < 0.006 \\ \begin{bmatrix} 55 & -45\frac{\pi}{180} & 1.75 & 55 \end{bmatrix}^T & \bar{C}_{res} > 0.006 \end{cases} \quad (4.11)$$

$$\mathbf{P}_0 = \begin{bmatrix} (1.5\frac{55}{4})^2 & 0 & 0 & 0 \\ 0 & (20\frac{\pi}{180})^2 & 0 & 0 \\ 0 & 0 & (0.75)^2 & 0 \\ 0 & 0 & 0 & (55)^2 \end{bmatrix} \quad (4.12)$$

Table 4.6. Configuration of the controller which initializes assuming the strand is left, right, or straight ahead.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	6 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 / 60 degrees
Bank angle limit	15 / 30 degrees
Airspeed	Best glide
Approach method	Dynamic

The performance of the 3 case initialization method at approach angles between 0 and 60 degrees is similar to that of the left-or-right initialization method. The cutoff criteria was adjusted several times in an attempt to maximize the performance. As expected, there is a large improvement between approach angles of 60 and 90 degrees in how often the aircraft is able to track on the first attempt.




Figure 4.13. Percent of the time that the 3 case controller can track the ideal thermal strand.

4.3.5 Fitting Using 4 Averaged Measurements

Rather than arbitrarily picking a location (or a number of locations), setting a high yet reasonable covariance, and relying on the Kalman filter to refine the guess, it is possible to use the variometer and roll disturbance measurements - along with the aircraft's state when they were taken - to estimate a strand size and location that could produce the measurements. However, fitting an ideal thermal strand to a 6 second buffer of variometer measurements, roll disturbance detector measurements, and aircraft state vectors all recorded at 20 Hertz will require significant computational resources and time. Note that keeping track of the aircraft's position and orientation for each measurement is critical, as this dramatically affects the measurements.

If instead of attempting to fit the 4 parameters of the ideal thermal model to 242 measurements, those parameters were fit to a small number of measurements then the

computational time required would be decreased. In this controller, the 6 second buffer of saved data is divided into 2 equal length segments. The measurements and state vectors within each of the segments are averaged to generate 1 average roll disturbance measurement and 1 average variometer measurement for each of 2 average aircraft states. The resulting 4 measurements and 2 positions are then used to solve for the 4 ideal thermal strand parameters using MATLAB's nonlinear equation solver "fsolve". The equations input to "fsolve" are structured in such a way that the returned state vector for the ideal thermal strand is relative to the aircraft's current position.

It should be noted that this method of averaging aircraft state works best when the change in the state of the aircraft with respect to time is roughly linear. For example, if the aircraft is flying in a straight line, the average position is a point on that line and is consistent with where the measurements were taken. If instead the aircraft is flying a tight turn, the average position is a point halfway along the turn, offset inside the turn - and from where the measurements were taken - by some distance.

The length of a single segment also has an effect on the accuracy of the result. In addition to exacerbating nonlinear changes in the position and orientation of the vehicle, the measurements of a thermal strand are themselves nonlinear. This means that averaging over an excessive amount of time degrades the result through both the averaged position and the averaged measurements. On the other hand, while averaging for a short period of time decreases the error introduced from averaging position, it increases the noise in the resulting averaged measurement.

The final factor to consider is the distance spanned by the averaged positions and measurements. A larger distance will increase the portion of the strand that is measured and will increase the visibility of the strand parameters, but has the potential to increase the estimate error by increasing the time over which the average is taken.

The ideal thermal strand state obtained from "fsolve" is directly used as the initial state estimate for the Kalman filter. The initial covariance estimate is given by Equation 4.13, and how often the aircraft is able to track the strand using this initialization method can be found in Figure 4.14.

$$\mathbf{P}_0 = \begin{bmatrix} (1.5\frac{55}{4})^2 & 0 & 0 & 0 \\ 0 & (30\frac{\pi}{180})^2 & 0 & 0 \\ 0 & 0 & (0.75)^2 & 0 \\ 0 & 0 & 0 & (55)^2 \end{bmatrix} \quad (4.13)$$

Table 4.7. Configuration of the controller which initializes assuming the strand is either to the left or right.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	6 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 / 60 degrees
Bank angle limit	15 / 30 degrees
Airspeed	Best glide
Approach method	Dynamic

This initialization method performs worse than most of the previous methods tested. However, it does not rely upon the Kalman filter to converge the initial state estimate, and it can react to different thermal strand geometries - something which entirely relied upon the initial covariance in previous methods.

4.3.6 Fitting Using 8 Averaged Measurements

Given that breaking the saved data into 2 segments and performing a nonlinear fit to the resulting 4 averaged measurements did not work very well, the next logical step




Figure 4.14. Percent of the time that the controller fitting to 4 average measurements can track the ideal thermal strand.

was to attempt using more segments. This allows for a larger portion of the thermal strand to be measured, and as previously discussed this will increase the visibility of the strand parameters.

In this controller the saved measurement and aircraft state data was broken into 4 segments rather than 2. This produced 8 averaged measurements associated with 4 averaged aircraft states. The length of saved data was increased to 12 seconds to maintain the length of each segment. The averaged positions and measurements are input into a nonlinear least squares solver and used to find an ideal thermal strand state vector that theoretically results in the smallest least squared error in the measurements. Due to the very small magnitude of the roll disturbance measurements, it was necessary to add a weighting factor to prevent them from being effectively ignored (the unweighted least squared error of an roll disturbance measurement is minuscule when compared to the least squared error of variometer measurements).

The configuration of the controller can be found in Table 4.8, and how often it was able to track the ideal thermal strand is shown in Figure 4.15.

Table 4.8. Configuration of the controller which initializes using 8 averaged measurements.

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	12 seconds
Variometer review period	6 seconds
Latching time	15 seconds
Desired crossing angle	20 / 60 degrees
Bank angle limit	15 / 30 degrees
Airspeed	Best glide
Approach method	Dynamic




Figure 4.15. Percent of the time that the controller fitting to 8 averaged measurements can track the ideal thermal strand.

Increasing the number of segments improved performance over most approach angles, particularly for approach angles between 40 and 90 degrees. Furthermore, this method performs better than all of the previously tested initialization methods

at approach angles between 60 and 90 degrees - which will occur during roughly 60 percent of encounters with thermal strands.

However, this initialization method still performs worse than most of the naïve initialization methods tested at approach angles less than 60 degrees. This is because a naïve initialization method is guaranteed to generate an initial state estimate with values that are reasonable for a real thermal strand, but the nonlinear fitting methods can and sometimes do generate initial state estimates with completely unrealistic values. When this happens, the initial state estimate is so far from the truth that even with a conservative initial covariance estimate the Kalman filter is unable to converge the state estimate.

4.3.7 Identifying Initial State Estimates Unlikely to Result in a Successful Track of the Strand

The initialization values generated by the nonlinear solvers vary wildly with sensor noise. It is not uncommon for MATLAB's "fsolve" or "lsqnonlin" to return state estimates that have the center of the strand kilometers away, the maximum updraft velocity of the thermal strand hundreds or thousands of meters per second, or a σ_{str} that is measured in tens of kilometers. These results are nonsensical, and almost always result in the aircraft failing to track the strand. If these bad initializations can be filtered out, then the performance of the nonlinear fitting initialization methods can be improved.

The first step in dealing with initializations that are unlikely to result in a track of the thermal strand is to identify them. Using the 8 averaged measurements nonlinear fit controller, the initial state estimate from 500 simulations at each approach angle was plotted and color coded to indicate whether the aircraft was able to track the ideal

strand on the first attempt, was able to track eventually, or was completely unable to track the ideal thermal strand. The results can be found in Figures 4.16, 4.17, 4.18, and 4.19.




Figure 4.16. Scatter plot of initialization values of d when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.20

The results for d , \dot{h}_{max} , and σ_{str} are particularly useful, as almost all of the instances in which the aircraft was able to track the strand on the first attempt are grouped tightly together. Figures 4.20, 4.21, and 4.22 show an enlarged view of the previous plots of d , \dot{h}_{max} , and σ_{str} with a focus on the region that contains the cases where the aircraft was able to track the strand.

Using this knowledge about the distribution of strand state variables in cases for which the aircraft was able to track on the first attempt, it was possible to define criteria that will reject a large number of initializations that do not converge while keeping those that do.




Figure 4.17. Scatter plot of initialization values of ψ_{str} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.




Figure 4.18. Scatter plot of initialization values of \dot{h}_{max} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.21




Figure 4.19. Scatter plot of initialization values of σ_{str} when encountering the ideal thermal strand. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track. An enlarged view can be found in Figure 4.22

Almost all of the initialization values of σ_{str} for which the aircraft was able to track the strand were less than 150 meters, and based on published data it was determined to be highly unlikely that a real thermal strand would have a σ_{str} greater than 150 meters. Therefore, a criteria was be set that would filter out all returned initial states for which $\sigma_{str} > 150$. In addition, it was decided to require that $\sigma_{str} > 10$ meters. Thermal strands smaller than this were found to occur in the large eddy simulation data only at altitudes less than 40 meters ($z/z_i = 0.04$), and commonly occur below 20 meters ($z/z_i = 0.02$). The Vulture UAV is flown autonomously only at altitudes greater than 30 meters above ground level, so it is unlikely that the real aircraft would encounter a strand of this size. Having a minimum strand size was necessary because sometimes the nonlinear fit would return an initialization guess that consisted of a very narrow strand ($\sigma_{str} = 0$ to 10 meters) that was right on top of the aircraft. These




Figure 4.20. Scatter plot of initialization values of d when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.




Figure 4.21. Scatter plot of initialization values of \dot{h}_{max} when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.




Figure 4.22. Scatter plot of initialization values of σ_{str} when encountering the ideal thermal strand with a focus on values for which the aircraft tracked on the first attempt. Blue indicates that the controller tracked on the first attempt, yellow that the controller was eventually able to track, and red that the controller failed to track.

initializations were more likely to occur when the aircraft was flying nearly parallel to the strand (see Figure 4.22); however, they did occur for every approach angle. In these cases, the aircraft was almost never able to track the strand, and so it was necessary to determine a minimum strand width to allow for these initializations to be filtered out.

Another criteria was set that required $\dot{h}_{max} \leq 6$ meters per second because this filtered out the majority of initializations resulting in a lost track, and it is highly unlikely that a thermal strand will have a strength greater than 6 meters per second. Another criteria was set that required $\dot{h}_{max} \geq 0$ meters per second. While it is possible to have a linear region of sinking air in the atmospheric convective boundary layer, such a feature is not intended to be modeled by the ideal thermal strand, and is not a feature that the aircraft should be tracking along.

Finally, it was also determined that the distance to the strand should be limited to

no more than 150 meters. This was determined using the scatter plots, as before, and the knowledge that if the aircraft has started attempting initialization, it is unlikely to be more than one σ_{str} away from the strand centerline.

A summary of the criteria used to accept or reject initial state estimate can be found in Table 4.9.

Table 4.9. Criteria used for accepting or rejecting initial state estimates.

State	Criteria
d	less than 150 meters
ψ_{str}	none
\dot{h}_{max}	greater than 0 and less than 6 meters per second
σ_{str}	greater than 10 and less than 150 meters

4.3.8 Addition of Initialization Rejection and Reinitialization

The 8 averaged measurements initialization method was chosen for implementation of initialization rejection and reinitialization because its performance was superior to the 4 averaged measurements method. The initialization rejection and reinitialization was configured such that if the initial state estimate did not meet the given criteria, the estimate would be rejected, the aircraft would continue flying wings level, and another initialization would be attempted after a set period of time - allowing enough change in the saved measurements so that hopefully the next initialization would have a better result. The configuration of this controller is presented in Table 4.10, and the results of testing on the ideal thermal strand are shown in Figure 4.23.

Adding initial state estimate rejection and reinitialization significantly improved the performance of the 8 average measurements initialization. The chance of the aircraft tracking the strand on the first attempt was improved for all approach angles. Furthermore the 8 average measurement initialization with initial state estimate

Table 4.10. Configuration of the controller which is able to reject initialization values

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	12 seconds
Variometer review period	6 seconds
Initialization attempt spacing	1 second
Latching time	15 seconds
Desired crossing angle	20 degrees
Approach method	Dynamic




Figure 4.23. Percent of the time that the controller fitting to 8 averaged measurements with initialization rejection can track the ideal thermal strand.

rejection was able to eventually track the ideal strand more than 80 percent no matter at what angle the aircraft approached the strand. When it is considered that the aircraft is more likely to approach the strand at a steep angle than a shallow angle, this initialization method performed better overall than any other initialization method tested.

4.4 Testing in LES

The controller with 8 averaged measurement initialization and rejection of poor initial state estimates was subjected to a Monte Carlo analysis in the large eddy simulation wind field. The configuration of the controller can be found in Table 4.11. The large eddy simulation data was tested to confirm that the magnitude of the mean updraft velocity at any altitude was less than 10^{-3} meters per second, and that continuity was maintained. The actual magnitude of the mean updraft velocity was much smaller than 10^{-3} at every altitude checked. From this it was concluded that the average updraft velocity encountered by an aircraft flying through the wind field would be almost exactly 0. This was confirmed in simulation.

Table 4.11. Configuration of the controller which is able to reject initialization values

Parameter	Configuration
Tracking method	S-curve
Measurement buffer length	12 seconds
Variometer review period	6 seconds
Initialization attempt spacing	1 second
Latching time	15 seconds
Desired crossing angle	20 degrees
Approach method	Dynamic

The aircraft was started at 120 meters above ground level in every simulation and was configured with an altitude controller that would attempt to make sure that it remained at or above 120 meters at all times. When the aircraft descended below 120 meters, the motor would be turned on and the aircraft would climb back up to 120 meters. If an updraft caused the aircraft to soar above 120 meters, the motor would be turned off.

The aircraft was given a random heading and horizontal position for each of 600 simulations. Each simulation lasted 10 minutes. If the aircraft ran into the horizontal

boundaries of the simulation environment, it would be turned around and the Kalman filter state vector appropriately adjusted. If the aircraft exceeded the upper bound of the simulation, it would descend back down on its own because all vertical wind velocities above 400 meters were set to 0. Violating the lower bound of the simulation was not an issue because of the altitude controller.

As a baseline, a set of simulations were run in which the aircraft was given a random starting position and heading, and was commanded to fly a constant airspeed with the wings level. Because the aircraft dynamics in the simulation do not include the effects of disturbances, this had the same effect as commanding the aircraft to fly a constant airspeed and heading. As was expected, the average updraft velocity encountered by the aircraft was nearly 0 meters per second. A histogram of the updraft velocity encountered during these simulations is shown in Figure 4.24.

The results of the 600 simulations and 100 hours of simulated flight are shown in Figures 4.25 and 4.26. The average updraft velocity encountered over the 100 hours of flight was 0.32 meters per second, and the average net energy rate was -0.20 meters per second. This matches the calculated average sink rate of 0.52 meters per second. On average, the average updraft velocity was insufficient to maintain flight, but is considerably better than the average updraft of 0 meters per second that was found when the aircraft was commanded to fly random headings through the large eddy simulation. Furthermore, the average net energy rate of -0.20 meters per second with the controller was considerably better than the -0.47 meters per second obtained from flying the airspeed for maximum lift to drag.

The Monte Carlo analysis was repeated using the centerline tracking method instead of S-curves. It was found that the average updraft experienced by the centerline tracking controller was also 0.32 meters per second, with an average net energy rate of 0.21 meters per second. These values are close enough to those of the




Figure 4.24. Distribution of local updraft encountered when flying a random heading in the large eddy simulation wind field. The mean updraft velocity is marked in gray. The minimum sink rate and sink rate at maximum lift to drag of the Vulture UAV are shown in blue and red, respectfully, as a reference.




Figure 4.25. Distribution of updraft encountered by the aircraft in the large eddy simulation wind field while using the thermal strand tracking controller. The mean updraft velocity is marked in gray.




Figure 4.26. Distribution of net energy rate of the aircraft in the large eddy simulation wind field while using the thermal strand tracking controller. The mean net energy rate is marked in gray, and the average sink rate of the Vulture UAV is shown in red as a reference.

S-curve tracking method that it was concluded that when applied to the large eddy simulation wind field, neither method was superior.

Slightly less than half of the time both the S-curve controller and the centerline controller were able to navigate the aircraft such that the average updraft velocity experienced exceeded the maximum lift to drag sink rate. The flight path of one such successful flight using the centerline tracking controller is shown in Figures 4.27 and 4.28, with the color of the line indicating the total energy of the vehicle - including onboard energy storage in the form of the 3000 milliamp-hour 2S lithium polymer battery. Figure ?? shows the flight path overlaid on top of a cross-section of the large eddy simulation wind field taken at 240 meters. The aircraft is able to track along the strands for short periods of time; however, vertical updrafts from large turbulent eddies tend to throw off the Kalman filter estimate and cause the aircraft to lose track of the strand.

During the simulation, that the aircraft encountered the boundaries of the computational domain three times: once at the Northern boundary, once at the upper boundary of 400 meters, and once at the Western boundary. Both times the aircraft encountered the horizontal boundaries, it was simply turned around with appropriate changes made to the Kalman filter state. When the aircraft encountered the upper boundary of the simulation, the lift experienced by the aircraft decreased - preventing it from soaring higher. Updraft strength at altitudes above 400 meters was set to 0, as described in section 2.2.1.




Figure 4.27. 3D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second.

The aircraft starts with a normalized total energy of 1.23×10^5 meters, and reaches a maximum normalized total energy of 1.25×10^5 meters. Sixty-three seconds after the start of the simulation, the aircraft the minimum normalized total energy experienced during the simulation of 1.19×10^5 meters. This is because the lithium polymer motor battery contains a lot of energy, but a substantial portion of that energy is lost in the




Figure 4.28. 2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second.




Figure 4.29. 2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.56 meters per second, and the average sink rate of the aircraft was 0.52 meters per second. Large eddy simulation cross-section taken at 240 meters.

electronic speed controller, motor, and propeller.

Another successful flight through the large eddy simulation environment is presented in Figures 4.30 and 4.30. In this case, the aircraft was able to track along a thermal strand for 1400 meters before it encountered a large turbulent eddy that caused it to lose track of the strand.




Figure 4.30. 3D flight path of a successful through the large eddy simulation wind field flight using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.62 meters per second, and the average sink rate of the aircraft was 0.47 meters per second.

Figures 4.32 and 4.33 show a flight through the large eddy simulation where the average updraft experienced was 0.18 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. In this simulation, the Vulture UAV made multiple attempts to track various thermal strands; however, it was unable to do so through a combination of inaccurate initializations and large turbulent eddies causing the state estimate to diverge from the truth.




Figure 4.31. 2D flight path of a successful flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.62 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. Large eddy simulation cross-section taken at 240 meters.




Figure 4.32. 3D flight path of a flight through the large eddy simulation wind field using the centerline tracking controller. The flight path is colored to reflect the instantaneous total energy of the vehicle. The average updraft experienced during the flight was 0.18 meters per second, and the average sink rate of the aircraft was 0.47 meters per second.




Figure 4.33. 2D flight path of a flight through the large eddy simulation wind field using the centerline tracking controller. The average updraft experienced during the flight was 0.18 meters per second, and the average sink rate of the aircraft was 0.47 meters per second. Large eddy simulation cross-section taken at 140 meters.

Chapter 5 |

Conclusion

The challenges of sensor noise and turbulence make it difficult to accurately track a thermal strand in the atmospheric boundary layer; however, through the use of a Kalman filter it is possible to reliably estimate the location of an idealized thermal strand. It was shown that when the filter design presented in this thesis is given a reasonable initialization, it is, for all practical purposes, able to track an ideal thermal strand indefinitely.

Several different initialization methods were tested and compared, ranging from a naïve assumption that the thermal strand was directly ahead to a nonlinear estimation of the strand's location and characteristics given a number of averaged measurements. Of the methods tested, using a nonlinear fit to 8 averaged measurements and being able to reject unrealistic initial state estimates and reinitialize was found to be the best of all initialization methods tested. The aforementioned method was able to provide an initialization good enough for the aircraft to track the ideal thermal strand more than 80 percent of the time regardless of from which direction the aircraft approached the ideal strand.

Finally, the Kalman filter and initialization method were tested in a wind field generated from a large eddy simulation of the atmospheric convective boundary layer.

Monte Carlo simulations indicated that the controller was able to obtain an average rate of normalized energy extraction from the atmosphere of 0.32 meters per second. However, this is lower than the average sink rate during of 0.52 meters per second that was found during the analysis, but is considerably better than the 0 meters per second found when the aircraft was commended to fly random headings at constant airspeed. From this it can be concluded that the thermal strand tracking system presented in this thesis works and provides a significant benefit over doing nothing, but it needs improvement to be able to reliably maintain soaring flight.

5.1 Future Work

There is great potential for future work in this area, as there has been very little research thus far into autonomous microlift soaring. One possible method to improve estimation of the strand is to use cooperative soaring to enable better sampling.

The initialization method could be improved by determining the thermal strand state limits as function of atmospheric conditions and the altitude at which the vehicle is flying. This would require further analysis of large eddy simulations of the atmosphere to determine how likely a strand with given characteristics is to form for a known set of atmospheric conditions. Further improvement could come from an investigation into strand acquisition maneuvers designed to increase observability of the thermal strand's location and characteristics.

The model of the ideal thermal strand could be expanded by including the local horizontal wind generated by the strand. This horizontal component should at the very least satisfy the equations for continuity of an incompressible fluid.

The sensors used could be expanded to directly include accelerometers (which are already indirectly used by the Kalman filter through the roll disturbance detector

and simulated variometer). Experienced sailplane pilots use vertical acceleration to help determine when and how steeply to bank while thermalling, and can feel lateral accelerations indicating the presence of lift even before a vertical acceleration is felt - and seconds before the variometer responds. Particularly when combined with the inclusion of horizontal components of wind in the model of the ideal thermal strand, direct accelerometer measurements may improve the ability of the controller to track a thermal strand.

Once a reliable thermal strand following controller has been developed, it could be integrated with a thermal centering controller in an effort to determine whether the aircraft has encountered a thermal strand or a thermal. One possible way to accomplish this would be through the use of a multiple model adaptive estimator running thermal centering and thermal strand tracking Kalman filters in parallel.

References

- [1] ADKINS, S. S. (2005), “Carbon Dragon Technical Website,” .
URL <http://www.carbondragon.us/>
- [2] (2013) *Glider Flying Handbook*, FAA Flight Standards Service, Web.
- [3] (2007), “File:Cloud street.jpg - Wikimedia Commons,” .
URL https://commons.wikimedia.org/wiki/File:Cloud{_}street.jpg
- [4] SULLIVAN, P. P. and E. G. PATTON (2011) “The Effect of Mesh Resolution on Convective Boundary Layer Statistics and Structures Generated by Large-Eddy Simulation,” *Journal of the Atmospheric Sciences*, **68**(10), pp. 2395–2415.
URL <http://journals.ametsoc.org/doi/abs/10.1175/JAS-D-10-05010.1>
- [5] REMDE, P. E., “Cumulus Soaring, Inc. - Glide Navigator II - Polars,” .
URL <https://www.cumulus-soaring.com/polars.htm>
- [6] ALL-AERO (2019), “LET L-13 Blanik / LAK LAK-6,” .
URL <http://all-aero.com/index.php/component/content/article/60-gliders/5973-let-l-13-blanik>
- [7] ——— (2019), “Glasflugel H.201 Standard Libelle,” .
URL <http://all-aero.com/index.php/contactus/34-planes/4217-glasflugel-h201-standard-libelle>
- [8] ——— (2019), “Glaser-Dirks DG-800,” .
URL <http://all-aero.com/index.php/component/content/article/34-planes/4213-glaser-dirks-dg-800>
- [9] ARNDT, U. (2008), “Best Glide and Speed to Fly for Paraglider Pilots,” .
URL https://www.expandingknowledge.com/Jerome/PG/Gear/Wing/Perf/Polar/Compare/Main{_}English.htm
- [10] A-I-R USA (2017), “A-I-R USA | ATOS Rigid Wing Hang Gliders | Resources,” .
URL <https://www.a-i-r-usa.com/resources>

- [11] BIRDS IN FLIGHT (2018), “White-backed Vulture *Gyps africanus* - Birds in Flight - ID guide,” .
URL <http://birds-in-flight.net/?p=9817>
- [12] BUCKLEY, N. J. (1999), “Black Vulture - Introduction | Birds of North America Online,” .
URL <https://birdsna.org/Species-Account/bna/species/blkvul/introduction>
- [13] THE PEREGRINE FUND, “Turkey Vulture (*Cathartes aura*) in Explore Raptors: Facts, Habitat, Diet | The Peregrine Fund,” .
URL https://www.peregrinefund.org/explore-raptors-species/Turkey{_}Vulture
- [14] KHOSRAVIFARD, S., V. VENUS, A. K. SKIDMORE, W. BOUTEN, A. R. MUÑOZ, and A. G. TOXOPEUS (2018) “Identification of Griffon Vulture’s Flight Types Using High-Resolution Tracking Data,” *International Journal of Environmental Research*, **12**(3), pp. 313–325.
- [15] MALLON, J. M., K. L. BILDSTEIN, and T. E. KATZNER (2016) “In-Flight Turbulence Benefits Soaring Birds,” *The Auk*, **133**(1), pp. 79–85.
- [16] REICHMANN, H. (1978) *Cross-Country Soaring*, Thomas Publications, Santa Monica.
- [17] MORELLI, P. (2003) “Why Microlift Soaring ?” in *OSTIV Congress*, Leszno, Poland.
- [18] ——— (2002) “OSTIV Definitions of Light and Ultralight Sailplanes: A Proposal to FAI-IGC,” *Technical Soaring*, **26**(3), pp. 67–69.
- [19] ——— (2006) “Development of the Microlift Glider,” *Technical Soaring*, **30**(1-2), pp. 13–19.
URL <http://journals.sfu.ca/ts/index.php/ts/article/viewFile/195/180>
- [20] LANGELAAN, J. W. (2009) “Gust Energy Extraction for Mini and Micro Uninhabited Aerial Vehicles,” *Journal of Guidance, Control, and Dynamics*, **32**(2), pp. 463–472.
- [21] DEPENBUSCH, N. and J. LANGELAAN (2010) “Receding Horizon Control for Atmospheric Energy Harvesting by Small UAVs,” in *AIAA Guidance, Navigation and Control Conference*, AIAA 2010-8180, August, Toronto, Ontario.
- [22] LANGELAAN, J. W. (2011) “A Gust Soaring Controller for Small Uninhabited Gliders,” *Technical Soaring*, **35**(2), pp. 48–60.

- [23] WILLIAMS, A. G. and J. M. HACKER (2004) “Inside Thermals,” *Technical Soaring*, **28**(4), pp. 24–35.
- [24] BOSLOUGH, M. B. E. (2002) *Autonomous Dynamic Soaring Platform for Distributed Mobile Sensor Arrays*, Tech. Rep. SAND2002-1896, Sandia National Laboratories, Albuquerque, New Mexico.
URL https://cfwebprod.sandia.gov/cfdocs/CompResearch/docs/02-1896{_}MobileSensorArrays.pdf
- [25] SUKUMAR, P. P. and M. S. SELIG (2013) “Dynamic Soaring of Sailplanes Over Open Fields,” *Journal of Aircraft*, **50**(5), pp. 1420–1430.
- [26] ALLEN, M. J. (2006) *Autonomous Soaring 2005 Flight Data Summary*, Tech. Rep. 20070022832, NASA.
- [27] GORDON, R. J. (2006) *Optimal Dynamic Soaring for Full Size Sailplanes*, Master of science, Air Force Institute of Technology.
- [28] POWERS, T. C., L. M. SILVERBERG, and A. GOPALARATHNAM (2019) “Artificial Lumbered Flight for Autonomous Soaring,” *Accepted to Journal of Guidance, Control, and Dynamics*, **November**.
- [29] BOWER, G., T. FLANZER, A. NAIMAN, and S. SARIPALLI (2010) “Dynamic Environment Mapping for Autonomous Thermal Soaring,” in *AIAA Guidance, Navigation and Control Conference*, AIAA 2010-8031, August, Toronto, Ontario.
- [30] DEPENBUSCH, N. T. and J. W. LANGELAAN (2011) “Coordinated Mapping and Exploration for Autonomous Soaring,” in *AIAA Infotech at Aerospace Conference and Exhibit 2011*, AIAA 2011-1436, March, St. Louis, Missouri.
- [31] DEPENBUSCH, N. T., J. J. BIRD, and J. W. LANGELAAN (2018) “The AutoSOAR Autonomous Soaring Aircraft, Part 1: Autonomy Algorithms,” *Journal of Field Robotics*, **35**(6), pp. 868–889.
- [32] ——— (2018) “The AutoSOAR Autonomous Soaring Aircraft Part 2: Hardware Implementation and Flight Results,” *Journal of Field Robotics*, **35**(4), pp. 435–458.
- [33] 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. 963–975.
- [34] ANDERSSON, K. and I. KAMINER (2009) “On Stability of a Thermal Centering Controller,” in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, AIAA 2009-6114, August, Chicago, Illinois.

- [35] EDWARDS, D. J. (2008) “Implementation Details and Flight Test Results of an Autonomous Soaring Controller,” in *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA 2008-7244, August, Honolulu, Hawaii.
- [36] DAUGHERTY, S. C. and J. W. LANGELAANN (2014) “Improving Autonomous Soaring via Energy State Estimation and Extremum Seeking Control,” in *AIAA Guidance, Navigation, and Control Conference*, AIAA 2014-0260, January, National Harbor, Maryland.
- [37] WALTON, C., I. KAMINER, V. DOBROKHODOV, and K. D. JONES (2018) “Alternate Strategies for Optimal Unmanned Aerial Vehicle Thermaling,” *Journal of Aircraft*, **55**(6), pp. 2347–2356.
- [38] REDDY, G., J. WONG-NG, A. CELANI, T. J. SEJNOWSKI, and M. VERGAS-SOLA (2018) “Glider Soaring via Reinforcement Learning in the Field,” *Nature*, **562**(7726), pp. 236–239.
URL <http://dx.doi.org/10.1038/s41586-018-0533-0>
- [39] PENNYCUICK, C. J. (1972) “Soaring Behaviour and Performance of Some East African Birds, Observed From a Motor Glider,” *Ibis*, **114**(2), pp. 178–218.
- [40] WAUGH, W. O. (1997) “Optimizing Figure-8 Soaring for Hill and Wave Lift,” *Technical Soaring*, **21**(1), pp. 3–10.
- [41] PENNYCUICK, C. J. (1982) “The Flight of Petrels and Albatrosses (Procellariiformes), Observed in South Georgia and its Vicinity,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, **300**(1098), pp. 75–106.
- [42] 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**(4), pp. 1002–1015.
URL <http://arc.aiaa.org/doi/10.2514/1.52738>
- [43] CHAKRABARTY, A. and J. LANGELAAN (2013) “UAV Flight Path Planning in Time Varying Complex Wind-fields,” in *2013 American Control Conference*, IEEE, Washington, DC, pp. 2568–2574.
- [44] IRVING, F. G. (1973) “Cloud-street Flying,” *Technical Soaring*, **3**(1), pp. 1–8.
- [45] SACHS, G. (1994) *Optimal Wind Energy Extraction for Dynamic Soaring*, vol. 44.
- [46] LISSAMAN, P. (2005) “Wind Energy Extraction by Birds and Flight Vehicles,” in *43rd AIAA Aerospace Sciences Meeting and Exhibit*, AIAA 2005-241, January, Reno, Nevada.
- [47] ZHAO, Y. J. (2004) “Optimal Patterns of Glider Dynamic Soaring,” *Optimal Control Applications and Methods*, **25**(2), pp. 67–89.

- [48] DEITTERT, M., A. RICHARDS, C. A. TOOMER, and A. PIPE (2009) “Engineless Unmanned Aerial Vehicle Propulsion by Dynamic Soaring,” *Journal of Guidance, Control, and Dynamics*, **32**(5), pp. 1446–1457.
- [49] BIRD, J. (2013) *Wind Estimation and Closed-Loop Control of a Soaring Vehicle*, Master of science, The Pennsylvania State University.
- [50] SACHS, G. (2005) “Minimum Shear Wind Strength Required for Dynamic Soaring of Albatrosses,” *Ibis*, **147**(1), pp. 1–10.
- [51] RCSPEEDS (2019), “RCSpeeds,” .
URL <http://www.rcspeeds.com/>
- [52] SLATER, A. E. (1977) “To Soar Like an Albatross,” *Flight International*, p. 1182.
- [53] KICENIUK, T. (2001) “Dynamic Soaring and Sailplane Energetics,” *Technical Soaring*, **25**(4), pp. 221–227.
- [54] WILLIAMS, A. G. and J. M. HACKER (1993) “Interactions Between Coherent Eddies in the Lower Convective Boundary Layer,” *Boundary Layer Meteorology*, **64**(1-2), pp. 55–74.
- [55] WILLIS, G. E. and J. W. DEARDORFF (1979) “Laboratory Observations of Turbulent Penetrative-convection Planforms,” *Journal of Geophysical Research*, **84**(C1), pp. 295–302.
- [56] WALLINGTON, C. E. (1983) “Potential Exploration and use of Miniscale Lift Patterns,” in *OSTIV Congress*, Hobbs, New Mexico, pp. 167–172.
- [57] MASON, P. J. (1989) “Large-Eddy Simulation of the Convective Boundary Layer,” *Journal of the Atmospheric Sciences*, **46**(11), pp. 1492–1516.
- [58] WALKO, R. L., W. R. COTTON, and R. A. PIELKE (1992) “Large-eddy Simulations of the Effects of Hilly Terrain on the Convective Boundary Layer,” *Boundary-Layer Meteorology*, **58**(1-2), pp. 133–150.
- [59] WAGGY, S. B., S. BIRINGEN, and P. P. SULLIVAN (2013) “Direct Numerical Simulation of Top-down and Bottom-up Diffusion in the Convective Boundary Layer,” *Journal of Fluid Mechanics*, **724**, pp. 581–606.
- [60] PARK, S. B. and J. J. BAIK (2014) “Large-eddy Simulations of Convective Boundary Layers Over Flat and Urbanlike Surfaces,” *Journal of the Atmospheric Sciences*, **71**(5), pp. 1880–1892.
- [61] TUCKER, V. A. (1987) “Gliding Birds: The Effect of Variable Wing Span,” *Journal of Experimental Biology*, **133**(1), pp. 33–58.

- [62] ——— (1988) “Gliding Birds: Descending Flight of the Whitebacked Vulture, *Gyps Africanus*,” *Journal of Experimental Biology*, **140**(1), pp. 325–344.
- [63] PENNYCUICK, C. J. (1971) “Control of Gliding Angle in Rüppell’s Griffon Vulture *Gyps Rüppellii*,” *Journal of Experimental Biology*, **55**(1), pp. 39–46.
- [64] WELCH, A., L. WELCH, and F. IRVING (1955) *The Soaring Pilot*, John Murray, London.
- [65] MACCREADY, P. B. J. (1958) “Optimum Airspeed Selector,” *Soaring*, (January-February), pp. 10–11.
- [66] STOJKOVIC, B. (1993) “Generalized Speed-to-fly Theory,” *Technical Soaring*, **17**(3), pp. 77–83.
- [67] COCHRANE, J. H. (1999) “MacCready Theory With Uncertain Lift and Limited Altitude,” *Technical Soaring*, **23**(3), pp. 88–96.
 URL <http://perso.magic.fr/bertazzo/mc/newmcred.pdf>
<http://www.ncsoaring.org/documents/PresentationsandTraining/macready-cochrane.pdf>
- [68] ALMGREN, R. and A. TOURIN (2015) “Optimal Soaring via Hamilton-Jacobi-Bellman Equations,” *Optimal Control Applications and Methods*, **36**(4), pp. 475–495.
- [69] DE JONG, J. L. (1984) “The “Convex-Combination Approach”: a Geometric Approach to the Optimization of Sailplane Trajectories,” *Technical Soaring*, **8**(3), pp. 98–117.
 URL <http://journals.sfu.ca/ts/index.php/ts/article/view/902>
- [70] HAUBENHOFER, M. (1963) “Die Mechanik des Kurvenfluges,” in *OSTIV Congress*, Junin, Argentine.
- [71] RASPET, A. (1950) “Performance Measurements of a Soaring Bird,” *Aeronautical Engineering Review*, **9**(12), pp. 14–17.
- [72] ——— (1960) “Biophysics of Bird Flight,” *Science*, **132**(3421), pp. 191–200.
- [73] TUCKER, V. A. and G. C. PARROTT (1970) “Aerodynamics of Gliding Flight in a Falcon and Other Birds,” *Journal of Experimental Biology*, **52**(2), pp. 345–367.
- [74] PENNYCUICK, C. J. (1971) “Gliding Flight of the White-Backed Vulture *Gyps Africanus*,” *Journal of Experimental Biology*, **55**(1), pp. 13–38.
- [75] PARROTT, G. C. (1970) “Aerodynamics of Gliding Flight of a Black Vulture *Coragyps Atratus*,” *Journal of Experimental Biology*, **53**, pp. 363–374.

- [76] GRAVES, G. R. (2017) “Sexual Monomorphism in Wing Loading and Wing Aspect Ratio in Black Vulture (*Coragyps atratus*) and Turkey Vulture (*Cathartes aura*),” *Proceedings of the Biological Society of Washington*, **130**(1), pp. 240–249.
- [77] TUCKER, V. A. (1991) “The Effect of Molting on the Gliding Performance of a Harris’ Hawk (*Parabuteo unicinctus*),” *The Auk*, **108**(1), pp. 108–113.
- [78] HOEY, R. G. (1992) “Research On the Stability and Control of Soaring Birds,” in *28th National Heat Transfer Conference*, AIAA 1992-4122, August, San Diego, California.
- [79] ——— (2010) “Exploring Bird Aerodynamics Using Radio-controlled Models,” *Bioinspiration and Biomimetics*, **5**(4).
- [80] STOLLE, M., J. BOLTING, C. DOLL, and Y. WATANABE (2015) “A Vision-based Flight Guidance and Navigation System for Autonomous Cross-Country Soaring UAVs,” *2015 International Conference on Unmanned Aircraft Systems, ICUAS 2015*, pp. 109–117.
- [81] STOLLE, M., Y. WATANABE, and C. DÖLL (2015) “A Sigma-Point Kalman Filter for Remote Sensing of Updrafts in Autonomous Soaring,” in *Advances in Aerospace Guidance, Navigation and Control* (J. Bordeneuve-Guibé, A. Drouin, and C. Roos, eds.), Springer International Publishing, Cham, pp. 283–302.
- [82] VAN DER MERWE, R. and E. A. WAN (2002) “The Square-root Unscented Kalman Filter for State and Parameter-Eestimation,” in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 3461–3464.