

The Pennsylvania State University
The Graduate School

ATMOSPHERIC ENERGY HARVESTING FOR SMALL UNINHABITED
AIRCRAFT BY GUST SOARING

A Thesis in
Aerospace Engineering
by
Nathan Thomas Depenbusch

© 2011 Nathan Thomas Depenbusch

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

August 2011

The thesis of Nathan Thomas Depenbusch was reviewed and approved* by the following:

Jacob W. Langelaan
Assistant Professor of Aerospace Engineering
Thesis Advisor

Mark M. Maughmer
Professor of Aerospace Engineering

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

*Signatures are on file in the Graduate School.

Abstract

Applying bird-inspired flight planning and control techniques to small robotic aircraft can greatly improve flight performance. This paper discusses a method for improving cruise performance of an uninhabited glider by harvesting energy from short period stochastic phenomena (i.e. gusts). Receding horizon control is used to plan a sequence of control inputs that maximizes an energy-based reward function over a time horizon, using only local knowledge of atmospheric conditions. Parameters of the receding horizon controller and parameters in the reward function are tuned using an evolutionary algorithm. The resulting controller is tested using Monte Carlo simulations of flight through Dryden gust fields: results show significant improvement over constant speed flight. Robustness of the receding horizon control approach to changes in aircraft parameters (parasite drag) is also evaluated using Monte Carlo simulations: increasing parasite drag results in gracefully degrading performance over the nominal condition while decreasing parasite drag results in increasing performance.

Table of Contents

List of Figures	vii
List of Tables	viii
Acknowledgments	ix
Chapter 1	
Motivating Energy Efficient Flight Strategies for Small UAVs	1
1.1 Introduction	1
1.2 Motivation	3
1.3 System Overview	5
1.4 Problem Description	6
1.4.1 Unknown Wind	6
1.4.2 Optimal Behavior Determination	7
1.4.3 Uncertain Vehicle Model	7
1.5 Review of Related Work	7
1.5.1 Forms of Atmospheric Energy	8
1.5.2 Avian Soaring Flight	8
1.5.3 Modeling of the Aircraft and Environment	10
1.5.4 Thermal Soaring	10
1.5.5 Orographic Lift	11
1.5.6 Dynamic Soaring	11
1.5.7 Gust Soaring	11
1.5.8 Receding Horizon Control Strategy	12
1.6 Contributions	13
1.7 Reader's Guide	13
Chapter 2	
Receding Horizon Control for Energy Harvesting	14
2.1 Problem Statement	14
2.2 Vehicle Model in a Dynamic Wind Field	16
2.2.1 Relevant Coordinate Frames	17
2.2.2 Derivation of Force Equations	18

2.2.3	Force Equations	21
2.2.4	Derivation of Moment Equations	23
2.2.5	Moment Equations	24
2.2.6	Euler Kinematics Equations	25
2.2.7	Navigation Equations	26
2.2.8	Model Outputs	27
2.2.9	Aerodynamic Forces	27
2.2.10	Thrust Force	28
2.2.11	Longitudinal Aircraft Model	29
2.2.12	Aircraft Models as Applied	30
2.3	Wind Field Prediction	31
2.4	Total Energy	32
2.5	Energy Maximization	34
2.5.1	The Cost Function	34
2.5.2	Aircraft Control Policy	35
2.6	Computation of Problem Solution	36
2.7	Controller Complexity	38
2.8	Wind Fields	38
2.8.1	Discrete gusts	39
2.8.2	Turbulence	39
2.9	Summary	42

Chapter 3

Application of an Evolutionary Algorithm to Optimize Control Variables	43
3.1 Summary of Evolutionary Algorithms	43
3.1.1 Fundamentals	44
3.1.2 Application in Control	45
3.2 Justification for the Use of Evolutionary Computation Methods	45
3.2.1 Downsides	46
3.2.2 Use in Related Literature	47
3.3 Algorithm Choice	48
3.4 Determining the Planning Horizon and Reward Parameters	49
3.4.1 Evolutionary Algorithm Control Variables	49
3.4.2 Algorithm Convergence and Interpretation	50
3.5 Evolved Gust Soaring Controllers	51
3.6 Summary	52

Chapter 4

Simulation Results	53
4.1 Performance Evaluation	53
4.1.1 Setup	53
4.1.2 Discrete gusts	54
4.1.3 Dryden Wind Fields	55
4.2 Discussion	59
4.2.1 Computation Cost of the Controller	61
4.3 Model Uncertainty	62
4.3.1 Robustness of the Controller	62
4.4 Summary	63

Chapter 5	
Conclusions	65
5.1 Summary of Contributions	66
5.1.1 Method for Dealing with Stochastic Nature of Turbulence	66
5.1.2 Development of Energy-Based Reward Function	66
5.1.3 Optimization of Controller Through Evolutionary Methods	66
5.1.4 Performance Verification through Simulation	67
5.1.5 Demonstration of Controller Robustness	67
5.2 Recommendations for Future Research	67
5.2.1 Examination of Controller Variability	67
5.2.2 Control Through Flap Actuation	67
5.2.3 Hardware Implementation	68
5.2.3.1 Improved Gust Data	68
5.2.3.2 Unsteady Aerodynamic Considerations	68
5.2.4 Three-Dimensional Gust Soaring Control	68
Appendix A	
Vehicle Properties	70
References	72

List of Figures

1.1	Dry mass vs. endurance	2
1.2	Small UAVs employed in research.	4
1.3	Several military UAVs of various sizes.	5
1.4	Sink rates of representative aircraft plotted versus airspeed.	6
2.1	Receding Horizon Control premise	16
2.2	Longitudinal reference frames	29
2.3	RHC for gust soaring	36
2.4	State barrier function	37
2.5	Spline point dependence	39
3.1	Evolving control variables	50
3.2	Evolutionary algorithm convergence	50
4.1	Discrete gust performance	54
4.2	Dryden turbulence: performance bars	57
4.3	Dryden turbulence: mean flightpaths	58
4.4	Gust soaring states	59
4.5	Controller performance summary	60
4.6	Computation time	61
4.7	Controller robustness	63
4.8	Robust controller flight paths	64
A.1	Vehicle properties as modeled	71

List of Tables

2.1	Turbulence conditions simulated.	41
3.1	Converged plan parameters.	52
4.1	Dryden simulations	56
A.1	Parameters for Omega II 2M glider.	70
A.2	State limits and control saturation for Omega II 2M glider.	71

Acknowledgments

My parents fostered my desire to engineer cool things from the earliest ages I can remember. They purchased countless Lego kits and put up with the blocks all over the house. As I am now too old to have time to play with Legos (though I never really lost the desire), my parents continue to make incredible sacrifices for my education. They have never stopped supporting me, and for that I thank them.

My advisor Jack Langelaan is ultimately responsible for the successful completion of this thesis. Had it not been for his offer of a position as I completed my undergraduate tenure, I would not have been introduced to this interesting research. Had it not been for his ability to explain the concepts in a manner that made them seem manageable, I would not have been able to comprehend this research. And had it not been for his patience reviewing the results and editing this work, I would not have been able to put this research into a presentable form. Thanks Jack.

I would also like to attribute much of the motivation to complete this work to my labmates over the last two years. Sana Sarfraz, Anjan Chakrabarty, Sean Quinn Marlow, and Jack Quindlen all helped in their own way. Whether it was by offering support, giving advice, answering questions, or just allowing me to vent, each was critical to the completion of my thesis. Sean has been especially helpful, his willingness to put up with me as an office mate never ceases to amaze me. Because I have not expressed the extent of my gratitude in person, this will have to suffice.

Finally, I'd like to thank all of my friends who have put up with me while I've been absorbed by my school work. I understand that I have been less than the friend I should have been during this whole process, but they stuck with me for the duration. Here's to great friends!

Dedication

This one is for my mother and father, the smartest people I know.

Motivating Energy Efficient Flight Strategies for Small UAVs

1.1 Introduction

This thesis presents a method by which useful energy can be gained from stochastic atmospheric turbulence by small uninhabited aerial vehicles (UAVs). Simulation results are provided, and suggestions for implementation on hardware are given. The motivation is to extend the duration over which a small UAV may remain airborne with the limited energy stores inherent to this type of vehicle. In principle no modifications to the vehicle are required (though improvements will be more pronounced on aircraft with more efficient planforms, e.g. gliders), making the research detailed in this thesis and the conclusions drawn immediately applicable and easily verifiable beyond computer simulation. On board inertial measurement sensors and GPS units are common on small UAVs and are the primary sensors used by the developed controller.

Recent developments in UAVs show two trends; at one end of the spectrum, large uninhabited aircraft are gaining capabilities for high altitude flight, long endurance, and are increasingly able to cover great distances. At the other end, UAVs are getting smaller while retaining many of the abilities previously held by significantly larger airplanes. These small aircraft offer advantages such as portability, difficulty for adversaries on the ground to detect or destroy, and comparatively low cost, while also eliminating the need for established ground facilities (airstrips, etc.). Smaller vehicles, however, pay a price in smaller available payloads, severely decreasing their utility for many missions. Additionally, any energy stored on-board, in the form of batteries or chemical fuel, decreases the mission payload capabilities further. Ultimately a trade-off must be made between increasing the energy storage on a small UAV and thus its range, or carrying a larger mission payload.

The relationship between the size a UAV (here measured by the aircraft's dry mass) and the endurance of that aircraft is shown in Figure 1.1. The aircraft that fall into the small UAV

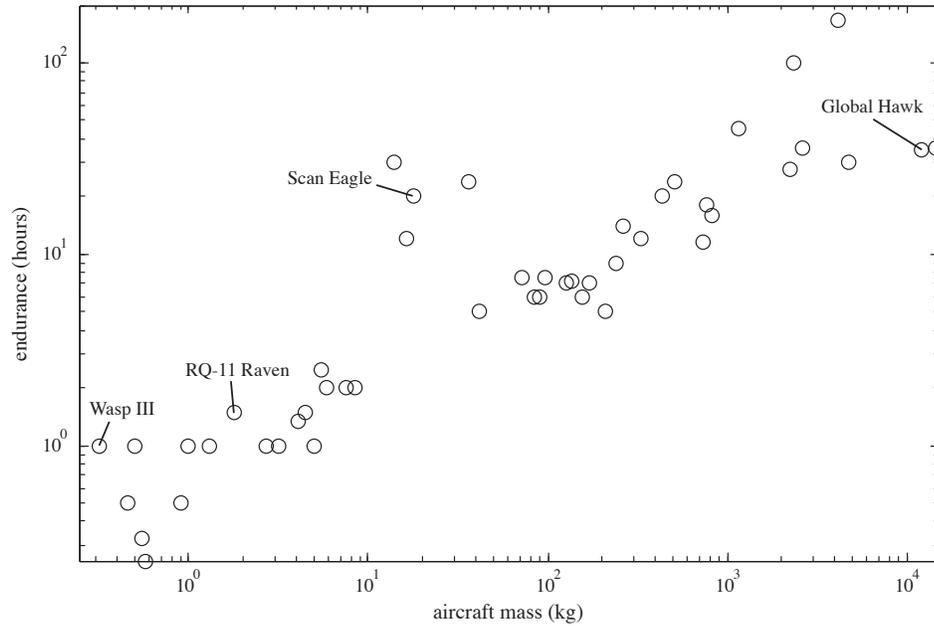


Figure 1.1. A plot of selected UAVs and the relationship between dry mass and endurance.

category can be seen clustered in the lower left corner of the plot. Outliers exist due to differing mission objectives and applications; however, there is a clear trend showing the degradation in endurance with decreasing size [1].

In order to maximize an aircraft's usefulness in light of severe weight restrictions, there is a need for methods to extract as much energy from the surrounding air as possible. An aircraft that is capable of gaining altitude or airspeed from the atmosphere alone is able to realize great gains in key specifications such as loiter time, endurance, or operational range.

This thesis will:

- Provide the framework for a receding-horizon controller to best determine the control inputs required to gain energy from the predicted wind. In this application, the horizon refers to the limited time over which control inputs are planned.
- Provide a simple yet effective method for dealing with the stochastic nature of atmospheric turbulence. From previous wind measurements, future wind conditions along the aircraft's flight path are predicted. Such a method is required in lieu of impossible to obtain data on real wind conditions.
- Develop an energy-based objective function that seeks to maximize energy gained in the short term while maintaining the aircraft's ability to react to future wind estimates in the long term. The objective function will also be shown to drive the aircraft towards stable flight in steady wind conditions.

- Apply an Evolutionary Algorithm to the tuning of relevant control parameters.
- Present simulation results for flight through Dryden turbulence of varying severity. Also presented are simulation results for the case in which the aircraft drag is not perfectly known.

1.2 Motivation

With the advent of increasingly smaller sensor packages such as video cameras, acoustic sensors and chemical sniffers, small uninhabited aerial vehicles are becoming increasingly attractive for many missions. However, these aircraft have significant restrictions in the amount of fuel that can be carried (either chemical fuel or electrical energy stored in batteries). Further, aerodynamic performance of small UAVs is typically worse than full sized aircraft because of the lower operating Reynolds numbers. Large birds such as eagles, vultures and albatrosses operate at similar Reynolds numbers, and they have evolved flight techniques that enable long duration and long distance flight with minimal energy expenditure. Applying bird-like planning and control techniques has the potential to greatly improve performance of small robotic aircraft.

In this thesis a small UAV is defined as an aircraft with a wingspan of up to 5 meters and mass less than 10 kilograms. These aircraft only require a single operator to deploy by hand, and do not rely on a launcher or rockets to launch. Additionally, a UAV must determine much of the control commands on board, excluding traditional radio-controlled aircraft from this definition. Small UAVs have recently become popular due to the relative low cost of developing such systems and their utility in a wide range of applications. Equipped with an autopilot and basic flight hardware, most model aircraft can be flown autonomously by relative amateurs [2].

An example of such a model aircraft often converted to fly as a UAV is the SBXC motor glider with a 4.27 meter wingspan [3] (Figure 1.2(a)) and an average weight on the order of 10 kg. Such aircraft are reasonable substitutes for similarly sized military aircraft and are the type typically examined in the literature, especially that specifically addressing autonomous soaring flight. The specific aircraft modeled in this thesis is the Omega II 2m motor glider shown in Figure 1.2(b). This aircraft features a 2 meter wingspan and a mass of 1.3 kg. Additional aircraft parameters can be found in the Appendix.

Numerous small UAVs are presently deployed in professional and military applications. A primary similarity between all aircraft in this class is that they are able to be launched by hand or with limited equipment while deployed *in situ*. This spans a wide range of aircraft, from the Wasp III [4] (Figure 1.3(a)), to the RQ-11B Raven (Figure 1.3(b)), both designed by Aerovironment. The primary military applications of small UAVs are remote reconnaissance and surveillance as well as target acquisition [5]. Currently these aircraft rely solely on energy stored in on board batteries, and thus their endurance is dependent entirely upon the size of their respective batteries.

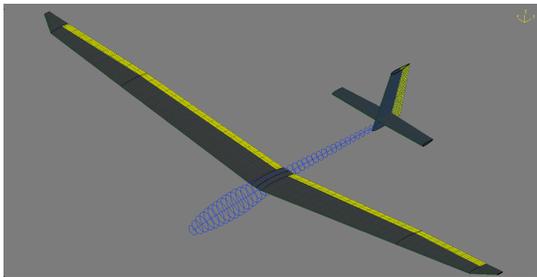
Larger aircraft such as the Scan Eagle [6] (Figure 1.3(c)) use chemical fuel and are similarly



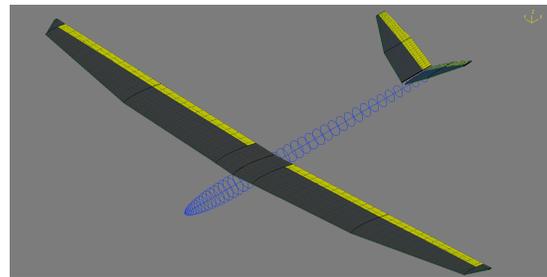
(a) SBXC motor glider (4.27 meter wingspan).



(b) Omega II motor glider (2.0 meter wingspan).



(c) AVL model of SBXC motor glider.



(d) AVL model of Omega II motor glider.

Figure 1.2. Small UAVs employed in research.

limited by the size of their fuel tanks. The need for atmospheric energy extraction methods for such aircraft is less necessary, however, because they are often capable of staying airborne for 20 or more hours on fuel reserves alone (Figure 1.1). These larger aircraft must also be launched and recovered using large, specialized equipment, violating the definition given for a small UAV .

Currently small UAVs are flown with limited regard to local wind conditions, or if considered, effort is put toward mitigating their effect towards an aircraft's performance. This is ignoring a potentially significant source of energy. Rather than limit missions to those that are not prohibited by excessive turbulence, this thesis will show that even in very turbulent conditions, flight is possible, and in many cases, more efficient flights can be made than in still air.

The research detailed in this thesis is made possible by several factors including the progressive miniaturization of computers such that they can fit inside a small aircraft. Additionally, the power consumption of these small computers is low enough that it has become advantageous to do calculations in flight onboard the aircraft. Standard autopilot hardware, such as inertial measurement units and GPS units, has also improved greatly in recent decades, allowing for more precise wind measurements.

This thesis addresses a method by which an aircraft may gain energy from atmospheric turbulence by longitudinal control alone. While further advantage may be gained from lateral control, work by other authors suggests that the lateral deviation in flight path will be extremely



Figure 1.3. Several military UAVs of various sizes.

large, potentially negating the benefits [7]. The method developed (based on receding horizon control, or RHC) is tested in using Monte Carlo simulations in realistic gust fields.

1.3 System Overview

The controller developed and detailed in this thesis is meant to be a part of a complete UAV system to extend range and mission duration. High level commands are typically specified by a human controller or by the mission requirements. A low level controller then carries out these commands. It is as part of this low level controller that the gust energy harvesting methods are employed. All methods here are simulated in two-dimensional flight, but could easily be applied to an aircraft flying in a three dimensional environment. For example, an aircraft that is directed to fly to a destination may make use of the gust soaring controller so that it may arrive having expended as little energy as possible en route.

1.4 Problem Description

The problem at hand is to effectively make use of random wind gusts to extend the range of a small aircraft. This is approached by maximizing the energy that a UAV gains while flying through simulated turbulence. It is assumed that the aircraft is equipped with all of the inertial motion sensors required to determine the wind conditions at its present location. The more general assumption is also made that all of the aircraft states are known and available at each simulation time step.

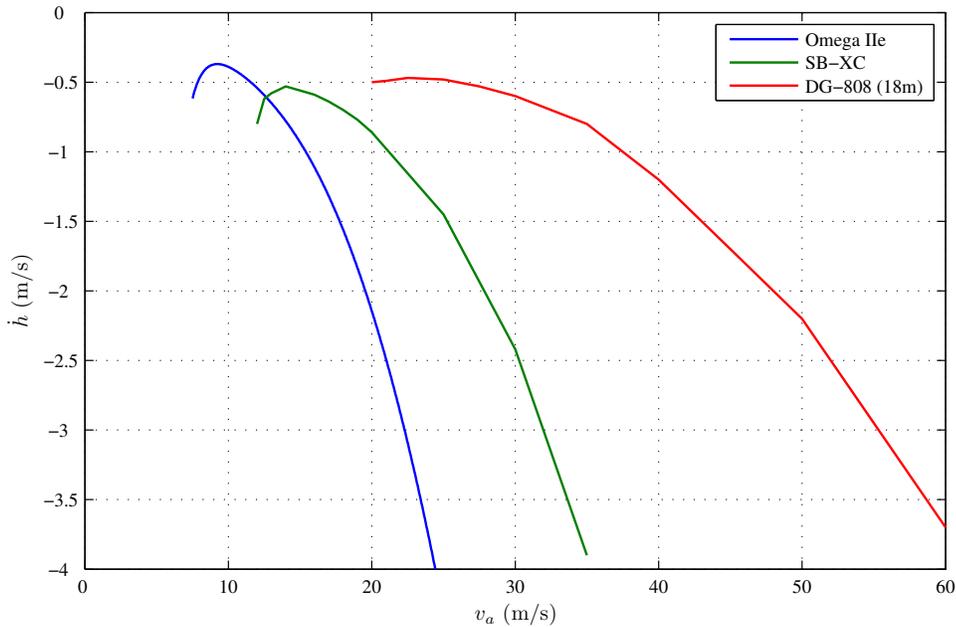


Figure 1.4. Sink rates of representative aircraft plotted versus airspeed.

Small UAVs suffer inherently from performance degradation due to the lower Reynolds number regimes in which they operate. The sink rate curve in Figure 1.4 shows that both best L/D and performance at higher airspeeds gets worse as aircraft size decreases. The DG-808 is a modern high-performance sailplane, representative of the most efficient aircraft available. This further reinforces the need for methods of atmospheric energy extraction to bring the performance of small aircraft closer to that of their larger and more efficient counterparts.

1.4.1 Unknown Wind

Gusts, which are short period spatial and temporal wind gradients, are the form of atmospheric energy examined and exploited here. Because they are random, a gusty wind field cannot be known or accurately predicted ahead of time. This makes the problem of energy extraction particularly difficult.

A strategy for making predictions of wind conditions in the near future is proposed here. A linear estimate based on current wind conditions and those in the recent past, is used to estimate the probable wind in the immediate future. This prediction is then used to determine the optimal set of control inputs for the immediate future. This process is repeated at a high rate (several times per second), giving the receding horizon control architecture that will be discussed further.

1.4.2 Optimal Behavior Determination

Given a prediction of wind in the immediate future, a flight trajectory that maximizes energy gain over a fixed time horizon must be developed. The length of time over which these inputs are computed is known as our planning horizon. The determined control inputs are followed for a portion of the time horizon known as the control horizon. At this point, the process of planning is repeated. Similar receding horizon approaches to path planning for aircraft are present in the literature [8]; however, none deal directly with stochastic wind conditions such as atmospheric turbulence.

In this thesis, a sequence of pitch rate inputs, parameterized as a spline, between optimized spline points is specified. These spline control points are equally spaced in time for each planning horizon. This serves to reduce the dimension of the input space, and smooths the control inputs. The spline points are determined such that they describe an input sequence that maximizes a developed reward function over the course of the planning horizon.

1.4.3 Uncertain Vehicle Model

Even a controller that performs flawlessly in simulation is subject to failure when implemented if the modeled system does not match the actual system. Similarly, it is possible that an aircraft's geometry may change through the course of a mission (e.g. icing, insect accumulation on the wing surfaces). It is important that an effective controller be able to cope with such situations with limited ramifications on the effectiveness of the controller.

In this thesis, the controller developed was used on aircraft models that differed from the baseline to demonstrate robustness to imperfect vehicle modeling.

1.5 Review of Related Work

Large birds and sailplane pilots commonly use the significant energy available in the atmosphere to their advantage by soaring. Soaring is the flight technique by which an aircraft is sustained solely by atmospheric motion. Atmospheric energy, in the form of moving air, is used to prolong a flight and/or cover great distances without the use of engines or flapping wings. It is this energy that UAVs have not previously been adept at utilizing and is the focus of much current research.

1.5.1 Forms of Atmospheric Energy

Energy can be harvested from the atmosphere in three ways: from vertical air motion; from deterministic spatial gradients such as shear layers; or from stochastic spatial and temporal gradients (i.e. gusts). These mechanisms operate on different time scales and differing assumptions are applicable to each. The difficulty and corresponding prevalence of energy harvesting from each form of air motion is dependent upon a number of factors, not the least of which being the energy available and the complexity of a controller designed to take advantage of that energy.

Perhaps the simplest is energy gain from vertical air motion. Vertical air motion is driven by thermal instabilities triggered by uneven heating of the ground, by orographic (i.e. ridge) lift and by long period gravity waves triggered by mountain ranges. The time scale of vertical air motion is long compared with vehicle dynamics, with thermals having lifespans of roughly 15 minutes and orographic and wave lift having lifespan measured in hours or days. Energy gain from vertical air motion is sometimes called *static soaring*, and it is common among raptors and vultures as well as sailplane and hang glider pilots.

Energy can also be harvested from spatial gradients in the wind field. Birds such as albatrosses have been observed gaining energy from wind gradients near the ocean surface [9–11]. This is known as *dynamic soaring*, exemplified by radio-control glider pilots who routinely fly aircraft to extremely high speeds via cyclic flights through the shear layers on the leeward side of ridges. Energy extraction from spatial gradients is often treated as a trajectory optimization problem with *a priori* known wind fields, and typically point mass models are sufficient to define vehicle dynamics.

Gusts are short duration, often very energetic air movements. Because of the stochastic nature of gusts, the wind field cannot be known or accurately predicted, making energy extraction particularly difficult. Further, the short time scales associated with gusts means that higher order dynamics become important. However, significant energy is available. Anecdotes based on observations of birds and radio-controlled (RC) glider flights suggest that on gusty days the performance of RC gliders is greatly reduced relative to birds [12]. The birds are effectively coping with the gusty conditions and, in some cases, gaining energy from them, a feat which RC pilots are unable to reproduce. This thesis focuses on the development of a *gust soaring* controller for small robotic aircraft.

1.5.2 Avian Soaring Flight

Of course the most prevalent and currently most adept application of flight strategies for taking advantage of atmospheric energy have evolved in birds. Though not every species of bird has been observed taking advantage of each of the three forms of atmospheric energy described above, different bird species can be observed masterfully utilizing different forms of energy. The graceful flight of a condor, or the wave skimming flight of an albatross demonstrates that atmospheric energy harvesting is not just theoretically possible, but is indeed practicable in wide range of locations. This aptitude for sustained flight with little energy expenditure motivates the synthetic

imitation detailed here and elsewhere in the literature.

Though it is true that aerodynamic efficiency may not be the only evolutionary pressure active on soaring birds [13], it is intuitive to believe that such creatures benefit from eons of taking advantage of atmospheric energy to promote their own survival. As such, a review of studies on the soaring flight of birds is pertinent to our understanding of the complexities of atmospheric flight of small aircraft.

Large birds commonly utilize the rising air in atmospheric convection to conserve energy. Numerous studies have been conducted on the thermalling behavior of birds such as pelicans, raptors [14], and frigatebirds [15]. The strategies used by these birds to gain altitude within thermals is understandably an active area of research with applications in UAV endurance optimization [16]. Taking advantage of thermals allows birds to expend very little energy while migrating or foraging for food; however, this strategy necessarily results in low cross-country speeds relative to birds with direct migratory behaviors. It has been determined that strong thermals can be found over land by day, while weaker thermals are available day and night over certain portions of the sea. The abundance of thermals and their frequency restricts the domains of birds that have evolved to rely on them for low-energy flight.

Dynamic soaring in birds is a strategy typically used by seabirds to extract energy from the wind speed gradients along the ocean's surface. Great distances can be covered very efficiently in this manner; albatrosses have been observed traveling thousands of kilometers over several days in this manner [17]. Sachs [11] examines the atmospheric conditions that must be present in order for dynamic soaring by albatrosses to be possible. The author examines cases in which the bird is able to extract enough energy from the wind to stay aloft while using no energy itself. The flight profile of an albatross engaged in dynamic soaring is characterized, and the primary energy gain mechanism is elucidated. The findings show that the minimum required wind speed is commonly found and often exceeded in the regions where albatrosses routinely forage.

The aforementioned methods of atmospheric energy extraction by birds require that specific weather conditions be present to be exploited. Gust soaring is a strategy that is theoretically possible anywhere that there is sufficient turbulence. As should be expected, gust soaring in birds is prevalent, but may require specific adaptations such as advanced sensory organs. Pennycuik asserts that the gust soaring of albatrosses is a complex maneuver that requires the measurement of very small variations in dynamic pressure. As such, he claims that the long beaks and nasal organs of these birds act in the same manner as an aircraft's pitot tube, though with more sensitivity [10].

Barnard observed that several species of African raptors forage for food primarily in gusty conditions [18]. She makes the observation that it is only in gusty conditions that the large birds are able to stay aloft without flapping their wings and therefore without using significant amounts of energy in the pursuit of a meal. The largest of the birds observed by Barnard were Jackal Buzzards (*Buteo rufofuscus*), which with a mean mass of 1.0 kg and wing area of 0.144 m² are very similar in size to the small UAVs to which this thesis proposes to apply such a soaring method. These birds were the more likely to use a gust soaring strategy to conserve energy than

were smaller birds in the same environment.

1.5.3 Modeling of the Aircraft and Environment

Because atmospheric conditions are often very complex, detailed models are required for simulated flight. Controllers that attempt to reproduce the complex behaviors of birds must similarly be very detailed. The trajectory optimization literature generally uses a simplified glider model, which assumes that the pilot has direct control of airspeed. This assumption is certainly appropriate for long duration flights where the glider spends most of its time in a trimmed condition, but this assumption is not valid for periods of transition between trimmed conditions. Some authors have addressed optimal transitions to minimize energy loss [19], and elsewhere Gedeon [20] describes an analysis of dolphin-style flight through thermals.

This simple model is not satisfactory for the rapid transition between trimmed states required for the implementation of a gust soaring controller. In the following chapter, a full six degree of freedom aircraft model will be developed to include the effects of a dynamic atmosphere. The model will then be reduced to model longitudinal motion only while preserving the effects of wind gradients on the dynamics of the aircraft. The model developed and used in this research will be adequate for simulating flight through the wind conditions presented.

The velocity of an atmospheric wind field is typically characterized by quasi-static mean winds and by random deviations. The mean winds are principally of concern for navigation, and thus ignored in this work. This is accomplished by adapting a reference frame that is static relative to these mean winds and is “carried” with the atmosphere. The random deviations from the mean wind comprise the turbulence which is the focus of this research.

Turbulence cannot be defined as an explicit function of time due to its stochastic nature. Therefore, only statistical, probabilistic approaches are appropriate for dealing with this subject. Whereas trajectory optimization for efficient flight is conducted on a time scale long enough that only the mean wind is a factor, a gust controller (and the aircraft model that it is applied to) must account for the instantaneous variations in wind speed that characterize a turbulent atmosphere.

Two assumptions are typically made while modeling turbulence [21]. The first assumption is that turbulence has no time dependence, rather it is a function of position only. This *frozen* wind field assumption is good provided the aircraft’s velocity is large compared to the speed scales of the turbulent deviations, thus this would be a poor assumption to apply to hovering flight. The second assumption is that turbulence is homogeneous at a given altitude. This means that the length and velocity scales encountered by an aircraft will be similar as long as it maintains a constant altitude. Both of these assumptions are made in the Dryden turbulence model described in Military Specification 8785C [22] that is employed in this thesis.

1.5.4 Thermal Soaring

Because thermals have such a prolonged lifespan, it is possible to develop detailed physical models of their structure and behavior. Such models can then be used for the simulation of flight through an environment primarily described by thermal updrafts. Allen [23] presents one such model, including the complexities of atmospheric sink.

Flight tests of a thermal soaring flight system were conducted by Allen and Lin [24] with a small UAV. The aircraft flown in these experiments is representative of the scale at which much UAV work is being conducted, an SBXC motor glider with a wingspan of 4.27m and a mass of 6.8kg (The same aircraft shown in Figure 1.2(a)). The total aircraft energy was used as a measure of thermal location, essentially using the aircraft itself as a thermal sensor. Results showed the ability to gain altitude with a relatively simple thermal centering controller.

1.5.5 Orographic Lift

Langelan [25] shows that utilizing orographic lift to extend the duration of a small UAV's flight becomes a trajectory optimization problem where a point mass model is sufficient to capture the relevant aircraft behavior. An expression for total aircraft energy is also presented and used in the trajectory optimization. Trajectories were discretized into straight line flights of various lengths, allowing for a finer resolution in areas where wind changes dramatically and for coarser resolution where wind speed is relatively consistent. Two solutions are presented for simulations of realistic ridge and valley terrain and wind conditions, one for minimum flight time to a destination, and the other for the maximum energy trajectory.

1.5.6 Dynamic Soaring

Dynamic soaring by both aircraft and birds has become an active area of research with the decreasing size of modern UAVs. Optimal trajectories for energy extraction from wind gradients are described by Zhao [26] and minimum fuel trajectories for power-assisted dynamic soaring are described by Zhao and Qi [27]. In [28] Barate et al. describe a Takagi-Sugeno-Kang (TSK) fuzzy logic controller for albatross-like dynamic soaring in wind gradients.

Dynamic soaring is explored by Barate, Doncieux, and Meyer [28] in the design of a biologically inspired flight controller. The authors use set of fuzzy rules in an attempt to reproduce the dynamic soaring behaviors of albatrosses. The controller was examined in the case of an imperfectly known aircraft system as is probable in a real-world application.

1.5.7 Gust Soaring

While a significant amount of work has been done on exploiting longer-duration atmospheric effects such as thermal and orographic lift as well as the velocity gradients required for dynamic soaring, less work has been performed on exploiting gusts. Phillips describes an approach to compute an equivalent thrust coefficient which occurs due to vertical gusts [29] and concludes that

the effect is too small to be useful in full-sized aircraft; however, extending Phillips approach to small UAVs shows that a significant performance improvement is possible. Interestingly, because full-sized aircraft such as sailplanes are unable to extract significant energy from gusts, small autonomous aircraft (whose aerodynamic efficiency is significantly worse than full-sized aircraft as is shown in Figure 1.4) may be able to outperform full-sized aircraft if they can efficiently extract energy from atmospheric turbulence. Additionally, Phillips ignores the possibilities for energy gain through “dolphin” flight in turbulence, which can lead to increased performance.

A simplistic turbulent atmosphere model was used by Lissaman and Patel [30] for the development of a gust soaring controller. Simulated flight by a small UAV through sinusoidal vertical wind gusts were examined. The authors were able to identify the gust characteristics (period and amplitude) required to sustain neutral energy flight. It was shown that the ability of an aircraft to effectively harvest gust energy was strongly dependent upon the lift-to-drag ratio of the aircraft. An aircraft with a higher lift-to-drag ratio was much better able to gain energy from atmospheric turbulence than was a vehicle in a higher drag configuration.

Work by Patel and Kroo [31] shows that significant energy savings may be achieved by a small aircraft utilizing gust energy harvesting techniques while flying through a Dryden gust field. A point mass model is also used in their research, and control laws are developed that allow for energy gain through the simulated gust environment. Previous work by Langelaan and Bramesfeld uses an environment populated with vertical gusts, but used a full dynamic model of aircraft longitudinal motion to generate control laws which maximized energy gain for flight through sinusoidal gusts [32].

Patel, Lee, and Kroo [33] present a simulation-based optimization procedure for developing a gust soaring controller. To enact energy gain from atmospheric turbulence, Patel et. al. optimize a set of control laws to return the optimal C_L given a measured wind speed. The prescribed C_L is then tracked by using flaps. The controller developed was flight tested on a small “bird-sized” UAV with promising results.

Genetic algorithms were employed for the tuning of controller gains by Langelaan [34]. The controller developed is applied to both vertical sinusoidal gust fields as well as vertical and longitudinal Dryden turbulence. The use of the Dryden turbulence model has been explored by several authors [29, 33, 35] and has become a common, if unverified, way to measure a gust soaring controller’s efficacy.

1.5.8 Receding Horizon Control Strategy

Recent work by Lawrance and Sukkariéh [8] describes a receding horizon control approach that selects an optimal trajectory from a family of possible control inputs. Their work was focused on three-dimensional path planning, but required knowledge of the wind field. An energy-based path planner is developed that uses local wind speeds to determine optimal trajectories. The receding horizon controller presented is used to plan three dimensional paths through an environment populated with thermals. A solely energy based objective function is used to determine which

of a set of available paths will best carry the aircraft through the planning horizon. Significant energy gains were shown with simulated flight through thermal fields and horizontal shear layers. Later work removed the requirement for wind field knowledge by adding a wind field mapping algorithm [36].

The research presented in this thesis has focused on developing a receding horizon control (RHC) strategy for energy harvesting when only local, instantaneous knowledge of the wind field is available. This permits much finer control of the vehicle's flight path. RHC has previously been applied to the problem of obstacle avoidance for small UAVs with limited sensing horizons or limited computational capability since it allows a trade between path quality and computational requirements [37,38].

1.6 Contributions

The primary contributions of this thesis are described below:

- A method for harvesting the often abundant atmospheric energy found in turbulence through the use of a receding horizon controller.
- The application of an evolutionary algorithm as a tool for tuning coefficients in the cost function and parameters of the receding horizon controller.
- Performance verification through numerous simulations of flight through Dryden turbulence, discrete gusts, sinusoidal gusts, and thermal streets.

1.7 Reader's Guide

The remainder of this thesis is organized in the following manner:

- **Chapter 2** defines the control law and the design procedure, as well as defines the dynamics and energetics of flight through gusts. The wind field models employed in simulated flight are also described;
- **Chapter 3** describes the implementation of an Evolutionary Algorithm to tune the control parameters. A brief review of evolutionary methods for controller design is presented, and the chosen evolutionary algorithm is justified;
- **Chapter 4** describes the results of simulated flights through various wind conditions and discussed the implications of the chosen control method. Also presented are results of simulations with an imperfectly known aircraft showing the robustness of the developed controller;
- **Chapter 5** Concludes this work and offers suggestions for future research.

Receding Horizon Control for Energy Harvesting

This chapter provides the necessary equations and methods used in developing the results that are given in this thesis. The structure of a receding horizon controller applied to gust soaring is explained in Section 2.1. Next the derivation of the equations of motion for an aircraft in a dynamic wind field is followed in Section 2.2. The equations of motion for both three and six degree of freedom aircraft models are presented. The method of wind field prediction used here and the equations for aircraft total energy are given in Section 2.3 and Section 2.4 respectively. Methods for maximizing total aircraft energy and the energy based cost function that forms the basis of the gust soaring controller are given in Section 2.5. The constraints on aircraft energy maximization are also explained. The specifics of computing an optimal path are given in Section 2.6 and the computational and physical complexity of the generated path is addressed in Section 2.7. This chapter concludes by presenting in Section 2.8 the equations that describe the turbulent environment in which the simulated aircraft and controllers are directed to fly. The equations for the Dryden turbulence model as given by Military Specification 8785C are described and the steps necessary to incorporate them into the aircraft simulation are detailed.

2.1 Problem Statement

The objective of the current research is to make use of atmospheric turbulence in order to gain as much energy as possible for a given distance traveled. A receding horizon control strategy will be used here. This type of control scheme uses knowledge of the present system, including states, measurements, and the plant model to predict future states and develop a control strategy accordingly. The calculated control strategy is followed over a fixed time horizon.

Thus, the problem at hand is to compute a flight trajectory which maximizes energy gain over a fixed time horizon given knowledge of the present wind conditions. At the same time this

trajectory must not place the vehicle in a state from which departure from controlled flight is likely to occur.

Receding horizon control involves the iterative optimization of a specified cost function over progressive control horizons. A current sample is taken of the system, and the control strategy is computed that best minimizes the cost function. Only the first steps in this strategy are typically followed before the situation is again sampled and a new strategy is computed.

In the context of soaring flight, receding horizon control assumes that atmospheric conditions are known *a priori* over a fixed time horizon (known as the plan horizon). A trajectory optimization algorithm can then be used to maximize energy gain over this plan horizon. The aircraft follows this trajectory for some fraction of the plan horizon (the control horizon), and then the process of planning is repeated (Figure 2.1). In effect RHC allows the robot to “plan over what it knows.”

The length of the plan horizon must be made long enough relative to system dynamics such that the controller can affect meaningful changes in the model. If the plan horizon is too short, the controller will become focused on short-term energy gain to the detriment of a long term performance. Alternatively, too short of a plan horizon may lead the controller to inaction, as any disturbance may yield a negative change in the reward function over the short term. The length of the control horizon is essential as well. Too long of a control horizon means that the predicted measurements of the environment are likely to be greatly in error. Too short of a control horizon will not allow the controller enough time to meaningfully change the system. The determination of these control parameters is introduced now, but will be discussed in more detail in Section 3.4.

In this case, future knowledge of the wind field is unavailable and cannot be predicted, thus receding horizon control is uniquely suited to this problem. Lawrence and Sukkarieh [8, 36, 39] have previously explored receding horizon control as a method for determining optimal paths through complex wind fields. Typically this strategy is used in situations where measurements are only available or reliable in the very near term. For example, robot path planning through environments characterized by many obstacles or with unknown future conditions.

In current physical applications the wind field cannot be known *a priori*, since there is currently no practical sensor that can measure the wind field ahead of the vehicle. This presents a unique problem for the robot. Ostensibly the environment in many other path planning problems may be made clear by more powerful sensors (think higher resolution cameras, longer range LIDAR/RADAR, more accurate geographic knowledge). Contemporary robot motion planning problems typically deal with physical obstacles which must be avoided, but can be detected, where this problem deals with more ephemeral gusts. The rapidly changing environment brought about by stochastic gusts, coupled with the near impossibility of measuring wind speeds remotely necessitate a path planning strategy that does not require clear knowledge of the environment in which the robot must operate.

Here the wind field (magnitude and spatial gradient) is assumed to be measurable only at the vehicle position. This is a valid assumption as wind speed estimation is a relatively simple task with the advanced sensor suite found on most modern UAVs and autopilots. The wind

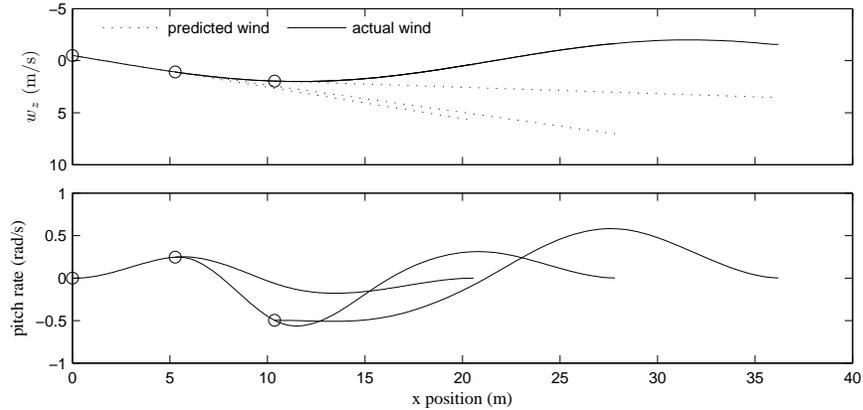


Figure 2.1. The premise of a receding horizon controller as applied to gust soaring. Open circles show the start of a control horizon. The upper figure shows wind speed (solid) and predicted wind (dotted); the lower figure shows the pitch rate input.

measurement along with trends obtained through previous measurements are used to predict future wind speeds. The result is a prediction of wind speeds projected forwards in time over the plan horizon.

A trajectory optimization algorithm is then used to determine a sequence of control inputs that maximizes a given reward function to carry the aircraft to the end of the plan horizon given the projected wind field. The determined sequence of control inputs is followed for the duration of the control horizon. As the aircraft advances in time, it measures the environment again, and the plan horizon advances as well. RHC thus does not require knowledge of the entire wind field and has the potential to be implemented in real time on representative computing hardware.

2.2 Vehicle Model in a Dynamic Wind Field

The equations of motion for a generalized six degree of freedom aircraft model are presented here. The model is then modified to give the dynamics of a three degree of freedom aircraft limited to only longitudinal motion. Both models incorporate the effects of a dynamic atmosphere and are simplified by assuming that the Earth is flat and that the $N E D$ coordinate frame (F_n) is an inertial reference frame. All of the assumptions made in developing these models strive to maintain an accurate representation of the effects of gusts on an aircraft while eliminating needless complexity.

The derivation of many of the following equations follows closely the methodology presented by Stevens and Lewis [40] with a key difference being that wind (specifically time-varying wind) is included in this derivation. Wind and the rate of change of that wind incident on the aircraft is crucial to the current research and, thus, is explicitly expressed in the equations of motion developed for the aircraft model.

2.2.1 Relevant Coordinate Frames

The Earth-centered inertial F_i and Earth-centered Earth-fixed F_e coordinate frames will be presented briefly for completeness, however, with the flat earth assumption, the $N E D$ frame is where relevant aircraft motion occurs.

The aircraft body-frame F_b can be arrived at by a sequence of three right-handed rotations from the $N E D$ coordinate frame F_n . The first rotation is yaw ψ about the original z_n -axis, followed by pitch θ about the new y -axis, and finally roll ϕ about the new x -axis. The transformation matrix to arrive at the body frame is thus $\mathbf{C}_{b/n}$.

$$\mathbf{C}_{b/n} = \begin{bmatrix} \cos \theta \cos \psi & \cos \theta \sin \psi & -\sin \theta \\ (-\cos \theta \sin \psi + \sin \phi \sin \theta \cos \psi) & (\cos \phi \cos \psi + \sin \phi \sin \theta \sin \psi) & \sin \phi \cos \theta \\ (\sin \phi \sin \psi + \cos \phi \sin \theta \cos \psi) & (-\sin \phi \cos \psi + \cos \phi \sin \theta \sin \psi) & \cos \phi \cos \theta \end{bmatrix} \quad (2.1)$$

The stability frame F_s may now be described as a left-handed rotation around the body y_b -axis by the angle of attack α . The relevant transformation matrix $\mathbf{C}_{s/b}$ is:

$$\mathbf{C}_{s/b} = \begin{bmatrix} \cos \alpha & 0 & \sin \alpha \\ 0 & 1 & 0 \\ -\sin \alpha & 0 & \cos \alpha \end{bmatrix} \quad (2.2)$$

The angular velocity vector $\omega_{s/b}$ between F_b and F_s is

$$\omega_{s/b}^s = \omega_{s/b}^b = \begin{bmatrix} 0 \\ -\dot{\alpha} \\ 0 \end{bmatrix} \quad (2.3)$$

and the cross-product matrix of $\omega_{s/b}^s$ is

$$\mathbf{\Omega}_{s/b}^s = \begin{bmatrix} 0 & 0 & -\dot{\alpha} \\ 0 & 0 & 0 \\ \dot{\alpha} & 0 & 0 \end{bmatrix} \quad (2.4)$$

Finally, the wind-axes F_w are the result of a right-handed rotation about the stability z_s -axis by the sideslip angle β . In this frame, the relative wind is pointing directly along the x_w -axis, and the aerodynamic forces L , D , and C are aligned along the z_w , x_w , and y_w axes respectively. It is for this reason that the majority of the development of the equations that follow will be done in F_w .

The rotation matrix from the stability frame to the wind frame can be written as $\mathbf{C}_{w/s}$.

$$\mathbf{C}_{w/s} = \begin{bmatrix} \cos \beta & \sin \beta & 0 \\ -\sin \beta & \cos \beta & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.5)$$

The angular velocity vector $\omega_{w/s}$ between F_s and F_w is

$$\omega_{w/s}^w = \omega_{w/s}^s = \begin{bmatrix} 0 \\ 0 \\ -\dot{\beta} \end{bmatrix} \quad (2.6)$$

The matrix transformation from F_b to F_w can be expressed as

$$\mathbf{C}_{w/b} = \mathbf{C}_{w/s} \mathbf{C}_{s/b} \quad (2.7)$$

2.2.2 Derivation of Force Equations

The most convenient place to start the development of the aircraft equations of motion is with the forces incident on the aircraft. The derivation starts with a statement of the forces that are present on the aircraft due to aerodynamics and propulsion systems and the accelerations incurred in the form of Newton's second law

$${}^i \dot{\mathbf{v}}_{CM/i} = \frac{1}{m} \mathbf{F}_{A,T} + \mathbf{G} \quad (2.8)$$

where ${}^i \dot{\mathbf{v}}_{CM/i}$ is the first derivative of the velocity of the aircraft's center of mass with respect to the inertial frame as taken in the inertial frame. $\mathbf{F}_{A,T}$ includes all of the aerodynamic and thrust forces acting on the aircraft, and \mathbf{G} is the acceleration resultant from the gravitational force on the vehicle.

It is assumed that vehicle mass remains a constant m . For the aircraft that will be simulated using this model, the primary method of energy storage is through batteries, making this assumption perfectly valid. If a chemical fuel source that is depleted over the duration of the mission is present or if the aircraft deploys a payload, this assumption may not be entirely valid.

The velocity of the aircraft in F_i can be expressed as a function of the aircraft's position relative to the origin $\mathbf{p}_{CM/O}$ (which is common to both F_i and F_e) as well as the angular rotation of F_e relative to the inertial frame through the use of the equation of Coriolis

$$\mathbf{v}_{CM/i} = {}^i \dot{\mathbf{p}}_{CM/O} = \mathbf{v}_{CM/e} + \omega_{e/i} \times \mathbf{p}_{CM/O} \quad (2.9)$$

This equation describes the kinematics of the aircraft. The remainder of this derivation involves expressing this equation in a convenient form for aircraft simulation.

If (2.9) is differentiated in the inertial frame, it yields

$${}^i \dot{\mathbf{v}}_{CM/i} = {}^i \dot{\mathbf{v}}_{CM/e} + \omega_{e/i} \times {}^i \dot{\mathbf{p}}_{CM/O} \quad (2.10)$$

substituting the derivative of $\mathbf{v}_{CM/e}$ taken in the body frame gives

$${}^i \dot{\mathbf{v}}_{CM/i} = {}^b \dot{\mathbf{v}}_{CM/e} + \omega_{b/i} \times \mathbf{v}_{CM/e} + \omega_{e/i} \times {}^i \dot{\mathbf{p}}_{CM/O} \quad (2.11)$$

(2.9) gives an expression for ${}^i\dot{\mathbf{p}}_{CM/O}$ which is substituted here

$${}^i\dot{\mathbf{v}}_{CM/i} = {}^b\dot{\mathbf{v}}_{CM/e} + \omega_{b/i} \times \mathbf{v}_{CM/e} + \omega_{e/i} \times (\mathbf{v}_{CM/e} + \omega_{e/i} \times \mathbf{p}_{CM/O}) \quad (2.12)$$

which can be rewritten

$${}^i\dot{\mathbf{v}}_{CM/i} = {}^b\dot{\mathbf{v}}_{CM/e} + (\omega_{b/i} + \omega_{e/i}) \times \mathbf{v}_{CM/e} + \omega_{e/i} \times (\omega_{e/i} \times \mathbf{p}_{CM/O}) \quad (2.13)$$

Finally, this expression may be equated to the acceleration presented in (2.8)

$$\frac{1}{m}\mathbf{F}_{A,T} + \mathbf{G} = {}^b\dot{\mathbf{v}}_{CM/e} + (\omega_{b/i} + \omega_{e/i}) \times \mathbf{v}_{CM/e} + \omega_{e/i} \times (\omega_{e/i} \times \mathbf{p}_{CM/O}) \quad (2.14)$$

and the acceleration in F_b solved for

$${}^b\dot{\mathbf{v}}_{CM/e} = \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{G} - (\omega_{b/i} + \omega_{e/i}) \times \mathbf{v}_{CM/e} - \omega_{e/i} \times (\omega_{e/i} \times \mathbf{p}_{CM/O}) \quad (2.15)$$

noting that

$$\mathbf{g} = \mathbf{G} - \omega_{e/i} \times (\omega_{e/i} \times \mathbf{p}_{CM/O}) \quad (2.16)$$

the expression in (2.15) becomes

$${}^b\dot{\mathbf{v}}_{CM/e} = \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - (\omega_{b/i} + \omega_{e/i}) \times \mathbf{v}_{CM/e} \quad (2.17)$$

(2.17) describes the dynamics of the aircraft on a rotating Earth. Note that at this point we can use the flat-Earth assumption to simplify things. The frame F_e is assumed to be an inertial frame, thus

$$\omega_{b/e} = \omega_{b/i} \quad (2.18)$$

and

$$\omega_{e/i} = 0 \quad (2.19)$$

simplifying (2.17) to

$${}^b\dot{\mathbf{v}}_{CM/e} = \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{CM/e} \quad (2.20)$$

The velocity of the vehicle relative to the surrounding air is the cause of the aerodynamic forces and moments. This motivates the development of an expression for aircraft velocity that is not relative to an inertial frame, but rather is relative to the atmosphere. The relative aircraft velocity vector can be written as [40].

$$\mathbf{v}_{rel} = \mathbf{v}_{CM/e} - \mathbf{v}_{W/e} \quad (2.21)$$

The first term on the right-hand side incorporates the velocity of the vehicle's center of mass, while $\mathbf{v}_{W/e}$ is the wind velocity expressed in F_e .

(2.21) can be differentiated in F_b (while differentiating $\mathbf{v}_{W/e}$ in F_e)

$${}^b\dot{\mathbf{v}}_{rel} + \omega_{b/e} \times \mathbf{v}_{rel} = {}^b\dot{\mathbf{v}}_{CM/e} + \omega \times \mathbf{v}_{CM/e} - {}^e\dot{\mathbf{v}}_{W/e} \quad (2.22)$$

\mathbf{v}_{rel} can be eliminated from the left-hand side by substituting (2.21)

$$\begin{aligned} {}^b\dot{\mathbf{v}}_{rel} + \omega_{b/e} \times (\mathbf{v}_{CM/e} - \mathbf{v}_{W/e}) &= {}^b\dot{\mathbf{v}}_{CM/e} + \omega \times \mathbf{v}_{CM/e} - {}^e\dot{\mathbf{v}}_{W/e} \\ {}^b\dot{\mathbf{v}}_{rel} + \omega_{b/e} \times \mathbf{v}_{CM/e} - \omega_{b/e} \times \mathbf{v}_{W/e} &= {}^b\dot{\mathbf{v}}_{CM/e} + \omega \times \mathbf{v}_{CM/e} - {}^e\dot{\mathbf{v}}_{W/e} \\ {}^b\dot{\mathbf{v}}_{rel} - \omega_{b/e} \times \mathbf{v}_{W/e} &= {}^b\dot{\mathbf{v}}_{CM/e} - {}^e\dot{\mathbf{v}}_{W/e} \end{aligned} \quad (2.23)$$

substituting (2.20) for ${}^b\dot{\mathbf{v}}_{CM/e}$ gives

$${}^b\dot{\mathbf{v}}_{rel} - \omega_{b/e} \times \mathbf{v}_{W/e} = \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{CM/e} - {}^e\dot{\mathbf{v}}_{W/e} \quad (2.24)$$

now (2.21) is solved for $\mathbf{v}_{CM/e}$ and substituted in

$$\begin{aligned} {}^b\dot{\mathbf{v}}_{rel} - \omega_{b/e} \times \mathbf{v}_{W/e} &= \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times (\mathbf{v}_{rel} + \mathbf{v}_{W/e}) - {}^e\dot{\mathbf{v}}_{W/e} \\ {}^b\dot{\mathbf{v}}_{rel} - \omega_{b/e} \times \mathbf{v}_{W/e} &= \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{rel} - \omega_{b/e} \times \mathbf{v}_{W/e} - {}^e\dot{\mathbf{v}}_{W/e} \\ {}^b\dot{\mathbf{v}}_{rel} &= \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{rel} - {}^e\dot{\mathbf{v}}_{W/e} \end{aligned} \quad (2.25)$$

The derivative in the body frame presented in (2.25) is now replaced by a derivative in the wind frame F_w

$${}^w\dot{\mathbf{v}}_{rel} + \omega_{w/b} \times \mathbf{v}_{rel} = \frac{1}{m}\mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{rel} - {}^e\dot{\mathbf{v}}_{W/e} \quad (2.26)$$

and the resulting expression is resolved in F_w and cross-products are replaced by their respective cross-product matrices

$${}^w\dot{\mathbf{v}}_{rel} + \mathbf{\Omega}_{w/b}^w \mathbf{v}_{rel} = \frac{1}{m}\mathbf{F}_{A,T}^w + \mathbf{C}_{w/b} \mathbf{C}_{b/n} \mathbf{g}^n - \mathbf{\Omega}_{b/e}^w \mathbf{v}_{rel} - {}^e\dot{\mathbf{v}}_{W/e} \quad (2.27)$$

At this point some attention should be given to the term ${}^e\dot{\mathbf{v}}_{W/e}$ which incorporates the time-varying components of the wind vector. This value is particularly important in an environment characterized by gusts, as the derivative of the wind speed may be very large. In its current form, the time-derivative of wind is taken in F_e . Because wind components are typically specified in F_n and will be in this work, it would be convenient to take the time derivative of wind in F_n .

$${}^e\dot{\mathbf{v}}_{W/e} = {}^n\dot{\mathbf{v}}_{W/e} + \omega_{n/e} \times \mathbf{v}_{W/e} \quad (2.28)$$

Because of the flat-Earth assumption, this is simplified greatly, as $\omega_{n/e} = 0$. Now we can take the wind derivative in a familiar reference frame, and through the use of a transformation matrix,

express that derivative in F_n as well

$${}^n \dot{\mathbf{v}}_{W/n} = \mathbf{C}_{n/e} {}^e \dot{\mathbf{v}}_{W/e} \quad (2.29)$$

From here, ${}^n \dot{\mathbf{v}}_{W/n}$ will be used in place of ${}^e \dot{\mathbf{v}}_{W/e}$. (2.27) is now written entirely in F_w

$${}^w \dot{\mathbf{v}}_{rel}^w + \boldsymbol{\Omega}_{w/b}^w \mathbf{v}_{rel}^w = \frac{1}{m} \mathbf{F}_{A,T}^w + \mathbf{C}_{w/b} \mathbf{C}_{b/n} \mathbf{g}^n - \boldsymbol{\Omega}_{b/e}^w \mathbf{v}_{rel}^w - \mathbf{C}_{w/b} \mathbf{C}_{b/n} {}^n \dot{\mathbf{v}}_{W/n} \quad (2.30)$$

2.2.3 Force Equations

To obtain equations appropriate for modeling the dynamics of the aircraft, (2.30) will be expressed as three equations for \dot{V}_a , $\dot{\alpha}$, and $\dot{\beta}$. As stated previously, the wind frame is advantageous because the relative wind V_a is directly along the x -axis.

$$\mathbf{v}_{rel}^w = \begin{bmatrix} V_a \\ 0 \\ 0 \end{bmatrix}, \quad {}^w \dot{\mathbf{v}}_{rel}^w = \begin{bmatrix} \dot{V}_a \\ 0 \\ 0 \end{bmatrix} \quad (2.31)$$

The rest of the left-hand side of (2.30) can be expanded by writing the angular velocity vector and obtaining the corresponding cross-product matrix

$$\begin{aligned} \omega_{w/b}^w &= \omega_{w/s}^w + \mathbf{C}_{w/s} \omega_{s/b}^s \\ \omega_{w/b}^w &= \begin{bmatrix} -\dot{\alpha} \sin \beta \\ -\dot{\alpha} \cos \beta \\ \dot{\beta} \end{bmatrix} \\ \boldsymbol{\Omega}_{w/b}^w &= \begin{bmatrix} 0 & -\dot{\beta} & -\dot{\alpha} \cos \beta \\ \dot{\beta} & 0 & \dot{\alpha} \sin \beta \\ \dot{\alpha} \cos \beta & -\dot{\alpha} \sin \beta & 0 \end{bmatrix} \end{aligned} \quad (2.32)$$

yielding

$${}^w \dot{\mathbf{v}}_{rel}^w + \boldsymbol{\Omega}_{w/b}^w \mathbf{v}_{rel}^w = \begin{bmatrix} \dot{V}_a \\ \dot{\beta} V_a \\ \dot{\alpha} V_a \cos \beta \end{bmatrix} \quad (2.33)$$

The three primary aerodynamic forces, drag D , side-force C , and lift L comprise \mathbf{F}_A^w

$$\mathbf{F}_A^w = \begin{bmatrix} D \\ C \\ L \end{bmatrix} \quad (2.34)$$

Thrust forces are conveniently defined in the vehicle body frame (as motors are usually fixed to

the airframe), and are included with a transformation applied.

$$\mathbf{F}_{A,T}^w = \mathbf{F}_T^w - \mathbf{F}_A^w \quad (2.35)$$

$$\mathbf{F}_{A,T}^w = \mathbf{C}_{w/s} \mathbf{C}'_{s/b} \begin{bmatrix} F_T \\ 0 \\ 0 \end{bmatrix} - \begin{bmatrix} D \\ C \\ L \end{bmatrix} \quad (2.36)$$

where $\mathbf{C}'_{s/b}$ is equivalent to $\mathbf{C}_{s/b}$ where α is replaced with $\alpha + \alpha_T$. The value α_T represents the inclination of the line on which F_T acts relative to the body x -axis.

$$\mathbf{F}_{A,T}^w = \begin{bmatrix} F_T \cos(\alpha + \alpha_T) \cos \beta - D \\ -F_T \cos(\alpha + \alpha_T) \sin \beta - C \\ -F_T \sin(\alpha + \alpha_T) - L \end{bmatrix} \quad (2.37)$$

The gravity term \mathbf{g}^n is expressed in F_n as $[0 \ 0 \ g_D]^T$

$$\mathbf{g}^w = \mathbf{C}_{w/b} \mathbf{C}_{b/n} \mathbf{g}^n = \begin{bmatrix} g_1 \\ g_2 \\ g_3 \end{bmatrix} \quad (2.38)$$

where

$$\begin{aligned} g_1 &= g_D (-\cos \alpha \cos \beta \sin \theta + \sin \beta \sin \phi \cos \theta + \sin \alpha \cos \beta \cos \phi \cos \theta) \\ g_2 &= g_D (\cos \alpha \sin \beta \sin \theta + \cos \beta \sin \phi \cos \theta - \sin \alpha \sin \beta \cos \phi \cos \theta) \\ g_3 &= g_D (\sin \alpha \sin \theta + \cos \alpha \cos \phi \cos \theta) \end{aligned}$$

The gust term is now expanded

$$\mathbf{C}_{w/b} \mathbf{C}_{b/n} {}^n \dot{\mathbf{v}}_{W/n} = \begin{bmatrix} d_1 \frac{dw_x}{dt} + d_2 \frac{dw_y}{dt} + d_3 \frac{dw_z}{dt} \\ d_4 \frac{dw_x}{dt} + d_5 \frac{dw_y}{dt} + d_6 \frac{dw_z}{dt} \\ d_7 \frac{dw_x}{dt} + d_8 \frac{dw_y}{dt} + d_9 \frac{dw_z}{dt} \end{bmatrix} \quad (2.39)$$

where the coefficients of the above matrix are given by:

$$\begin{aligned} d_1 &= (\cos \beta \sin \alpha (\sin \phi \sin \psi + \cos \phi \cos \psi \sin \theta) - \sin \beta (\cos \theta \sin \psi - \cos \psi \sin \phi \sin \theta) \\ &\quad + \cos \alpha \cos \beta \cos \psi \cos \theta) \\ d_2 &= (\sin \beta (\cos \phi \cos \psi + \sin \phi \sin \psi \sin \theta) - \cos \beta \sin \alpha (\cos \psi \sin \phi - \cos \phi \sin \psi \sin \theta) \\ &\quad + \cos \alpha \cos \beta \cos \theta \sin \psi) \\ d_3 &= (\sin \beta \cos \theta \sin \phi - \cos \alpha \cos \beta \sin \theta + \cos \beta \cos \phi \sin \alpha \cos \theta) \\ d_4 &= -(\cos \beta (\cos \theta \sin \psi - \cos \psi \sin \phi \sin \theta) + \sin \alpha \sin \beta (\sin \phi \sin \psi + \cos \phi \cos \psi \sin \theta) \\ &\quad + \cos \alpha \cos \psi \sin \beta \cos \theta) \end{aligned}$$

$$\begin{aligned}
d_5 &= (\cos \beta (\cos \phi \cos \psi + \sin \phi \sin \psi \sin \theta) + \sin \alpha \sin \beta (\cos \psi \sin \phi - \cos \phi \sin \psi \sin \theta) \\
&\quad - \cos \alpha \sin \beta \cos \theta \sin \psi) \\
d_6 &= (\cos \beta \cos \theta \sin \phi + \cos \alpha \sin \beta \sin \theta - \cos \phi \sin \alpha \sin \beta \cos \theta) \\
d_7 &= (\cos \alpha (\sin \phi \sin \psi + \cos \phi \cos \psi \sin \theta) - \cos \psi \sin \alpha \cos \theta) \\
d_8 &= -(\cos \alpha (\cos \psi \sin \phi - \cos \phi \sin \psi \sin \theta) + \sin \alpha \cos \theta \sin \psi) \\
d_9 &= (\sin \alpha \sin \theta + \cos \alpha \cos \phi \cos \theta)
\end{aligned}$$

The seemingly complex nature of the gust coefficients is a result of their convenient definition in F_n and the necessity of applying them to the vehicle as expressed in F_w .

The equations for vehicle dynamics that result from expanding (2.30) are now given

$$\dot{V}_a = \frac{(F_T \cos(\alpha + \alpha_T) \cos \beta - D)}{m} + g_1 - d_1 \frac{dw_x}{dt} - d_2 \frac{dw_y}{dt} - d_3 \frac{dw_z}{dt} \quad (2.40)$$

$$\begin{aligned}
\dot{\beta} &= -R_s - \frac{(C + F_T \cos(\alpha + \alpha_T) \sin \beta)}{V_a m} + \frac{g_2}{V_a} \dots \\
&\quad - \frac{1}{V_a} \left(d_4 \frac{dw_x}{dt} + d_5 \frac{dw_y}{dt} + d_6 \frac{dw_z}{dt} \right) \quad (2.41)
\end{aligned}$$

$$\begin{aligned}
\dot{\alpha} &= Q - P_s \tan \beta - \frac{(L - F_T \sin(\alpha + \alpha_T))}{V_a m \cos \beta} + \frac{g_3}{V_a \cos \beta} \dots \\
&\quad - \frac{1}{V_a \cos \beta} \left(d_7 \frac{dw_x}{dt} + d_8 \frac{dw_y}{dt} + d_9 \frac{dw_z}{dt} \right) \quad (2.42)
\end{aligned}$$

2.2.4 Derivation of Moment Equations

The moment equations will be derived in the stability frame F_s rather than the wind frame F_w as the aircraft inertia matrix is constant in F_s . The use of any other reference frame would cause unnecessary complications. The moments arrived at can be utilized in aircraft simulations through the simple application of the transformation matrices.

Similar to (2.8), the derivation is started with a basic statement of the moment equation in F_i

$$\mathbf{M}_{A,T} = \dot{\mathbf{h}} \quad (2.43)$$

where $\mathbf{M}_{A,T}$ are the aerodynamic and thrust moments affecting the aircraft. The derivative of the aircraft's angular momentum \mathbf{h} is taken in an inertial reference frame. This derivative is replaced by one taken in F_b

$$\mathbf{M}_{A,T} = {}^b \dot{\mathbf{h}} + \omega_{b/i} \times \mathbf{h} \quad (2.44)$$

Angular momentum can be expressed in F_b as a function of the vehicle inertia matrix J^b and the vehicle's angular velocity with respect to F_i , ($\omega_{b/i}^b$)

$$\mathbf{h}^b = J^b \omega_{b/i}^b \quad (2.45)$$

This can be differentiated in the body-frame

$${}^b \dot{\mathbf{h}}^b = J^b {}^b \dot{\omega}_{b/i}^b \quad (2.46)$$

The result can be combined with (2.52) when the latter is expressed in F_b

$$\mathbf{M}_{A,T}^b = J^b {}^b \dot{\omega}_{b/i}^b + \omega_{b/i}^b \times \mathbf{h}^b \quad (2.47)$$

then substituted (2.45) for \mathbf{h}^b

$$\mathbf{M}_{A,T}^b = J^b {}^b \dot{\omega}_{b/i}^b + \omega_{b/i}^b \times J^b \omega_{b/i}^b \quad (2.48)$$

Because of the flat-Earth assumption, F_e is considered an inertial reference frame, so all angular velocities taken with respect to F_i may equivalently be taken with respect to F_e . The expression is now solved for ${}^b \dot{\omega}_{b/e}^b$

$${}^b \dot{\omega}_{b/e}^b = (J^b)^{-1} \left(\mathbf{M}_{A,T}^b - \omega_{b/e}^b \times J^b \omega_{b/e}^b \right) \quad (2.49)$$

The derivative on the left-hand side in F_b is now replaced with the equivalent derivative in F_s

$${}^s \dot{\omega}_{b/e}^b + (\omega_{s/b} \times \omega_{b/e})^b = (J^b)^{-1} \left(\mathbf{M}_{A,T}^b - \omega_{b/e}^b \times J^b \omega_{b/e}^b \right) \quad (2.50)$$

The entire expression is now given in F_s , and the cross products are replaced by the equivalent cross-product matrices

$${}^s \dot{\omega}_{b/e}^s + \boldsymbol{\Omega}_{s/b}^s \omega_{b/e}^s = (\mathbf{C}_{s/b} J^b \mathbf{C}_{b/s})^{-1} \left(\mathbf{M}_{A,T}^s - \boldsymbol{\Omega}_{b/e}^s (\mathbf{C}_{s/b} J^b \mathbf{C}_{b/s}) \omega_{b/e}^s \right) \quad (2.51)$$

We can call the value $\mathbf{C}_{s/b} J^b \mathbf{C}_{b/s}$ the vehicle inertia matrix in F_s and denote it J^s for simplicity of notation. What results is an expression for the angular acceleration of the vehicle in F_s that can be used to solve for the aircraft angular rates.

$${}^s \dot{\omega}_{b/e}^s = -\boldsymbol{\Omega}_{s/b}^s \omega_{b/e}^s + (J^s)^{-1} \left(\mathbf{M}_{A,T}^s - \boldsymbol{\Omega}_{b/e}^s J^s \omega_{b/e}^s \right) \quad (2.52)$$

2.2.5 Moment Equations

At this point (2.52) will be expanded to explicitly show equations for each of the three aircraft angular rates (P_s , Q , and R_s) in F_s . The left side of (2.52) is equivalent to:

$${}^s \dot{\omega}_{b/e}^s = \begin{bmatrix} \dot{P}_s \\ \dot{Q} \\ \dot{R}_s \end{bmatrix} \quad (2.53)$$

and the first term on the right side can be expanded as:

$$-\mathbf{\Omega}_{s/b}^s \omega_{b/e}^s = \begin{bmatrix} \dot{\alpha} R_s \\ 0 \\ -\dot{\alpha} P_s \end{bmatrix} \quad (2.54)$$

The moments induced by aerodynamic forces must be expressed in F_s as:

$$\mathbf{M}_{A,T}^s = \mathbf{C}_{s/w} \mathbf{M}_{A,T}^w = \begin{bmatrix} l_w \cos \beta - m_w \sin \beta \\ m_w \cos \beta + l_w \sin \beta \\ n_w \end{bmatrix} = \begin{bmatrix} l_s \\ m \\ n_s \end{bmatrix} \quad (2.55)$$

The resulting moment equations of motion are

$$\dot{P}_s = R_s \dot{\alpha} + \frac{J'_{xz}(n_s + J'_x P_s Q - J'_y P_s Q - J'_{xz} Q R_s)}{\Gamma} + \frac{J'_z(l_s + J'_{xz} P_s Q + J'_y Q R_s - J'_z Q R_s)}{\Gamma} \quad (2.56)$$

$$\dot{Q} = \frac{m - P_s(J'_{xz} P_s + J'_x R_s) + R_s(J'_z P_s + J'_{xz} R_s)}{J'_y} \quad (2.57)$$

$$\dot{R}_s = -P_s \dot{\alpha} + \frac{J'_x(n_s + J'_x P_s Q - J'_y P_s Q - J'_{xz} Q R_s)}{\Gamma} + \frac{J'_{xz}(l_s + J'_{xz} P_s Q + J'_y Q R_s - J'_z Q R_s)}{\Gamma} \quad (2.58)$$

where Γ is part of the inverse of the J^s matrix:

$$\Gamma = J'_x J'_z - J'_{xz} \quad (2.59)$$

2.2.6 Euler Kinematics Equations

The approach to obtaining the aircraft Euler-angle derivatives is functionally very similar to a matrix transformation but conceptually different. It should be noted that each Euler angle exists in a separate coordinate system from each other Euler angle. This is due to the series of rotations required to arrive at each angle. The equation for determining the Euler-angle derivatives is

$$\dot{\Phi} = H(\Phi) \omega_{b/e}^b \quad (2.60)$$

where $H(\Phi)$ is

$$H(\Phi) = \begin{bmatrix} 1 & \tan \theta \sin \phi & \tan \theta \cos \phi \\ 0 & \cos \phi & -\sin \phi \\ 0 & \frac{\sin \phi}{\cos \theta} & \frac{\cos \phi}{\cos \theta} \end{bmatrix} \quad (2.61)$$

The equations of motion become:

$$\begin{aligned} \dot{\phi} &= P_s \cos \alpha (1 + \cos \phi \tan \theta) + Q \sin \phi \cos \phi \tan^2 \theta \dots \\ &\quad + R_s (\cos \phi \tan \theta \cos \alpha - \sin \alpha) \end{aligned} \quad (2.62)$$

$$\dot{\theta} = Q \cos \phi - \sin \phi (R_s \cos \alpha + P_s \sin \alpha) \quad (2.63)$$

$$\dot{\psi} = \frac{\cos \phi (R_s \cos \alpha + P_s \sin \alpha)}{\cos \theta} + \frac{Q \sin \phi}{\cos \theta} \quad (2.64)$$

2.2.7 Navigation Equations

The navigation equations are solved in F_n as this is the most useful reference frame for simulation. The motion of the aircraft relative to the ground or set waypoints is far more meaningful than descriptions of air relative motion.

(2.21) can be rewritten due to the flat-Earth assumption and realized in F_n as

$$\mathbf{v}_{rel}^n = \mathbf{v}_{CM/n}^n - \mathbf{v}_{W/n}^n \quad (2.65)$$

where $\mathbf{v}_{CM/n}^n$ represents the change in the aircraft's position in the NED coordinate frame or ${}^n\dot{\mathbf{P}}_{CM/n}^n$.

$$\mathbf{v}_{CM/n}^n = {}^n\dot{\mathbf{P}}_{CM/n}^n = \mathbf{v}_{rel}^n + \mathbf{v}_{W/n}^n \quad (2.66)$$

A matrix transformation is used to realize airspeed in F_w for which we have a simple expression given in (2.31).

$${}^n\dot{\mathbf{P}}_{CM/n}^n = \mathbf{C}_{n/b} \mathbf{C}_{b/w} \mathbf{v}_{rel}^w + \mathbf{v}_{W/n}^n \quad (2.67)$$

Expanding the above expression yields three separate equations for the rate of change of the aircraft's position in F_n .

$$\begin{aligned} \dot{P}_N &= w_N + V_a \left(\frac{\cos \beta \sin \alpha (\sin \phi \sin \psi + \cos \phi \cos \psi \sin \theta)}{G} \right) \\ &\quad - V_a \left(\frac{\sin \beta (\cos \phi \sin \psi - \cos \psi \sin \phi \sin \theta)}{G} + \frac{\cos \alpha \cos \beta \cos \psi \cos \theta}{G} \right) \end{aligned} \quad (2.68)$$

$$\begin{aligned} \dot{P}_E &= w_E + V_a \left(\frac{\sin \beta (\cos \phi \cos \psi + \sin \phi \sin \psi \sin \theta)}{G} \right) \\ &\quad - V_a \left(\frac{\cos \beta \sin \alpha (\cos \psi \sin \phi - \cos \theta \sin \psi \sin \theta)}{G} \right) \\ &\quad + V_a \left(\frac{\cos \alpha \cos \beta \cos \theta \sin \psi (-\cos \phi^2 + \cos \theta \cos \phi + 1)}{G} \right) \end{aligned} \quad (2.69)$$

$$\begin{aligned} \dot{P}_D &= w_D + V_a \left(\frac{\cos \beta \sin \alpha (-\cos^2 \psi \cos^2 \theta + \cos \phi \cos^2 \psi \cos \theta + \cos^2 \theta)}{G} \right) \\ &\quad + V_a \left(\frac{\sin \beta \cos \theta \sin \phi}{G} \right) \\ &\quad - V_a \left(\frac{\cos \alpha \cos \beta (\sin \theta \cos^2 \phi \cos^2 \psi - \sin \theta \cos^2 \phi)}{G} \right) \\ &\quad - V_a \left(\frac{\cos \theta \sin \theta \cos \phi \cos^2 \psi + \sin \phi \sin \psi \cos \phi \cos \psi + \cos \theta \sin \theta \cos \phi}{G} \right) \\ &\quad - V_a \left(\frac{\cos \theta \sin \phi \sin \psi \cos \psi + \sin \theta}{G} \right) \end{aligned} \quad (2.70)$$

Note that \dot{h} is equal to $-\dot{P}_D$. In the above equations, the value G is equivalent to the expression:

$$\begin{aligned} G = & \cos^2 \phi \cos^2 \psi - \cos^2 \phi - \cos \theta \cos \phi \cos^2 \psi + \sin \phi \sin \psi \sin \theta \cos \phi \cos \psi \\ & + \cos \theta \cos \phi - \cos \theta \sin \phi \sin \psi \sin \theta \cos \psi + 1 \end{aligned} \quad (2.71)$$

2.2.8 Model Outputs

In order to properly control an aircraft, some knowledge must be available about the accelerations present in the body-axis of the aircraft F_b . A pilot can feel these accelerations and compensate accordingly; however, an autopilot must be able to derive them. For actions such as coordinating a turn, these body axis accelerations (a_y^b in particular) are important.

The derivation of these body-axis accelerations is begun by restating (2.25):

$${}^b \dot{\mathbf{v}}_{rel} = \frac{1}{m} \mathbf{F}_{A,T} + \mathbf{g} - \omega_{b/e} \times \mathbf{v}_{rel} - {}^e \dot{\mathbf{v}}_{W/e}$$

This equation is realized in F_b , the value ${}^e \dot{\mathbf{v}}_{W/e}$ is replaced as in (2.29), and the cross product is replaced by its respective cross-product matrix.

$${}^b \dot{\mathbf{v}}_{rel}^b = \frac{1}{m} \mathbf{C}_{b/w} \mathbf{F}_{A,T}^w + \mathbf{C}_{b/n} \mathbf{g}^n - \mathbf{\Omega}_{b/e}^b \mathbf{v}_{rel}^b - \mathbf{C}_{b/n} {}^n \dot{\mathbf{v}}_{W/n} \quad (2.72)$$

This expression can be easily expanded by substituting the values that have already been presented to obtain three equations for the aircraft acceleration in F_b . Note that ${}^b \dot{\mathbf{v}}_{rel}^b = [a_x^b \ a_y^b \ a_z^b]^T$.

2.2.9 Aerodynamic Forces

The aerodynamic coefficients that describe the forces on an aircraft are largely non-linear with aircraft states. Within reasonable perturbations of a trimmed condition, they can be approximated as linear functions (with the exception of C_D which is non-linear in the region of interest). The non-dimensional longitudinal aerodynamic coefficients are approximated here by

$$C_D = f_{LD} (C_{L0} + C_{L\alpha} \alpha) + C_{D\delta_e} \delta_e + C_{D\delta_f} \delta_f \quad (2.73)$$

$$C_L = C_{L0} + C_{L\alpha} \alpha + \frac{c}{2V_a} (C_{LQ} Q + C_{L\dot{\alpha}} \dot{\alpha}) + C_{L\delta_e} \delta_e + C_{L\delta_f} \delta_f \quad (2.74)$$

$$C_m = C_{m0} + C_{m\alpha} \alpha + \frac{c}{2V_a} C_{mQ} Q + C_{m\delta_e} \delta_e + C_{m\delta_f} \delta_f \quad (2.75)$$

where $f_{LD} (C_{L0} + C_{L\alpha} \alpha)$ is a polynomial function which relates the aircraft drag coefficient to the aircraft lift coefficient.

The non-dimensional lateral-directional aerodynamic coefficients are approximated by

$$C_C = C_{C\beta} \beta + \frac{b}{2V_a} (C_{C_p} P_s + C_{C_r} R_s) + C_{C\delta_r} \delta_r + C_{C\delta_a} \delta_a \quad (2.76)$$

$$C_l = C_{l\beta} \beta + \frac{b}{2V_a} (C_{l_p} P_s + C_{l_r} R_s) + C_{l\delta_a} \delta_a + C_{l\delta_r} \delta_r \quad (2.77)$$

$$C_n = C_{n_\beta}\beta + \frac{b}{2V_a} (C_{n_p}P_s + C_{n_r}R_s) + C_{n_{\delta_a}}\delta_a + C_{n_{\delta_r}}\delta_r \quad (2.78)$$

The aerodynamic forces affecting the aircraft can now be written as dimensional qualities scaled to the dynamic pressure and the aircraft's wing area

$$D = \bar{q}SC_D \quad (2.79)$$

$$L = \bar{q}SC_L \quad (2.80)$$

$$C = \bar{q}SC_C \quad (2.81)$$

where \bar{q} is the dynamic pressure and is equivalent to $\frac{1}{2}\rho V_a^2$. It is assumed that the aerodynamic forces are expressed in F_w and act in the directions indicated in (2.34). These three equations make up \mathbf{F}_A^w

The moments in F_w can be written:

$$l_w = \bar{q}SbC_l \quad (2.82)$$

$$m_w = \bar{q}SbC_m \quad (2.83)$$

$$n_w = \bar{q}SbC_n \quad (2.84)$$

The three moments make up $\mathbf{M}_{A,T}^w$ as is used in (2.55)

2.2.10 Thrust Force

The thrust force provided by the aircraft's propulsion system is assumed to act at an angle of incidence α_T with respect to the body x -axis. In this model, asymmetric thrust, or thrust acting anywhere but along the aircraft's centerline is not addressed.

The thrust force is expressed in a similar form to the aerodynamic forces on the aircraft as

$$F_T = \bar{q}SC_T \quad (2.85)$$

where C_T is a function of the throttle input δ_t and the propulsion model for the vehicle.

In this work, gliding flight is assumed. From the perspective of the model, this simply means that the thrust force is equal to zero. The assumption of gliding flight does impose several restrictions on the model:

- The flightpath angle with respect to the surrounding air γ_a must always be negative in steady state flight. This implies that the aircraft cannot climb with respect to the atmosphere without giving up airspeed.
- The control inputs available to the controller are limited, and the resulting aircraft motion is restricted. For example, airspeed and flightpath angle may not be controlled simultaneously.

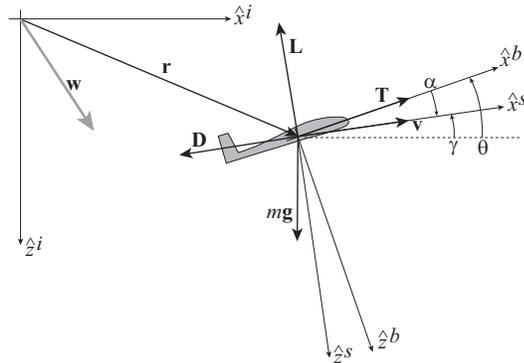


Figure 2.2. Reference frames for the longitudinal aircraft model. Positive rotations are indicated, so positive glideslope is upwards and angle of attack is positive in the conventional sense.

2.2.11 Longitudinal Aircraft Model

Longitudinal vehicle dynamics for motion in a time-varying wind field are derived in earlier work [34]. The equations of motion developed here coincide with the previously developed equations. All longitudinal equations are defined with reference to coordinate frames defined in Figure 2.2.

Using a common longitudinal definition of wind axes, \hat{x}^s is defined as a unit vector in the direction of airspeed (so that $\mathbf{v} = v_a \hat{x}^s$) and \hat{z}^s opposite to lift. The angle γ_a defines the rotation between the wind axes and the inertial axes, and it is the flight path angle with respect to the surrounding airmass. When $\mathbf{w} = 0$ it is also the flight path angle with respect to the inertial frame. In this application γ_a is defined as positive upwards, so for a steady glide the glideslope is negative.

The flat-Earth six degree of freedom aircraft model presented previously can be restricted to yield a three degree of freedom longitudinal aircraft model. All of the lateral variables are set equal to zero, and the notation is changed such that the $N E D$ coordinate frame becomes the $x - z$ plane. Motion previously described in the N direction corresponds directly to motion in the x direction, with similar analogies made between D and z .

Aircraft kinematics are developed by applying (2.68) and (2.70) to longitudinal motion to arrive at expressions for \dot{x} and \dot{z} . Note that w_x and w_z are wind speeds along the x and z axes.

$$\begin{aligned}\dot{x} &= w_x + V_a (\sin \alpha \sin \theta + \cos \alpha \cos \theta) \\ \dot{z} &= w_z + V_a (\sin \alpha \cos \theta - \cos \alpha \sin \theta)\end{aligned}$$

Trigonometric identities are substituted to arrive at a form recognizable in other work [41]. The expression for $\dot{\theta}$ from (2.63) is also included in the longitudinal kinematics equations with the same assumption of exclusively longitudinal motion. It can be seen that the kinematics of the aircraft can now be defined in terms of the airspeed, flight path angle and wind speed only. It is generally more convenient to work in terms of pitch angle and angle of attack, Figure 2.2

shows that $\gamma = \theta - \alpha$:

$$\dot{x} = \bar{w}_x + V_a \cos(\theta - \alpha) \quad (2.86)$$

$$\dot{z} = \bar{w}_z - V_a \sin(\theta - \alpha) \quad (2.87)$$

$$\dot{\theta} = Q \quad (2.88)$$

Vehicle dynamics are copied from (2.40) and (2.42) and written with all lateral and directional terms equal to zero.

$$\begin{aligned} \dot{V}_a &= \frac{(F_T \cos(\alpha + \alpha_T) - D)}{m} - g_D (\cos \alpha \sin \theta - \sin \alpha \cos \theta) \dots \\ &\quad - (\sin \alpha (\sin \theta) + \cos \alpha \cos \theta) \frac{dw_x}{dt} - (\cos \alpha \sin \theta + \sin \alpha \cos \theta) \frac{dw_z}{dt} \\ \dot{\alpha} &= Q - \frac{(L - F_T \sin(\alpha + \alpha_T))}{V_a m} + \frac{g_D (\sin \alpha \sin \theta + \cos \alpha \cos \theta)}{V_a} \dots \\ &\quad - \frac{1}{V_a} \left((\cos \alpha \sin \theta - \sin \alpha \cos \theta) \frac{dw_x}{dt} + (\sin \alpha \sin \theta + \cos \alpha \cos \theta) \frac{dw_z}{dt} \right) \end{aligned}$$

The use of simple trigonometric identities simplifies the equations to a form used in other work [34] and (2.57) for \dot{Q} is included after removing all lateral dependent terms.

$$\dot{V}_a = \frac{(F_T \cos(\alpha + \alpha_T) - D)}{m} - \frac{dw_x}{dt} \cos(\theta - \alpha) + \left(\frac{dw_z}{dt} - g_D \right) \sin(\theta - \alpha) \quad (2.89)$$

$$\dot{\alpha} = Q - \frac{(L - F_T \sin(\alpha + \alpha_T))}{V_a m} - \frac{1}{V_a} \left(\frac{dw_x}{dt} \sin(\theta - \alpha) + \left(\frac{dw_z}{dt} - g_D \right) \cos(\theta - \alpha) \right) \quad (2.90)$$

$$\dot{Q} = \frac{m}{J_y} \quad (2.91)$$

2.2.12 Aircraft Models as Applied

In this work, the longitudinal aircraft model developed here is used for the simulations presented in Chapter 4. In those cases, the non-dimensional aerodynamic coefficients C_L and C_D as presented in (2.74) and (2.73) are simplified to remove any dependence on control-surface deflections. The effects of control surface deflections on vehicle drag is ignored.

$$C_L = C_{L0} + C_{L\alpha} \alpha + \frac{c}{2V_a} (C_{LQ} Q + C_{L\dot{\alpha}} \dot{\alpha}) \quad (2.92)$$

$$C_D = f_{LD} (C_{L0} + C_{L\alpha}) \quad (2.93)$$

The aircraft is controlled with Q , the pitch rate input, effected through (2.90). This makes (2.91) irrelevant for our simulation purposes. For the simulations conducted with the longitudinal aircraft model, it is assumed that an on-board flight controller can follow pitch rate commands.

Additionally, C_T will be set to zero and thus F_T is equal to zero due to the gliding flight assumption as discussed previously.

When the effects of gusts on the aircraft model are to be simulated, a new reference frame

must be introduced. This is because the Dryden turbulence model [22] is not defined as fixed in F_n as wind typically is, but rather in a vehicle carried reference frame. The definitions of the gust components are given by MIL-STD-1797A [42] in the gust frame F_g . The three components of the gusts u_g , v_g , and w_g are defined in F_g . The vertical gust component w_g is considered parallel to the vertical (D) axis in F_n . The longitudinal gust component u_g is to be aligned along the "horizontal mean wind vector," indicating that it lies along the projection of V_a on the $N - E$ plane of F_n . The lateral gust component v_g is orthogonal to the other two.

We start by defining the rotation matrix from F_g to F_n as:

$$\mathbf{C}_{n/g} = \begin{bmatrix} \mathbf{C}_{n/w \ 1,1} & \mathbf{C}_{n/w \ 1,2} & 0 \\ \mathbf{C}_{n/w \ 2,1} & \mathbf{C}_{n/w \ 2,2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.94)$$

where $\mathbf{C}_{n/g}$ is the projection of F_w onto the horizontal plane of F_n . This expression allows us to easily obtain expressions for ${}^n\dot{\mathbf{p}}_{CM/n}^n$ given wind define in F_g .

The rotation matrix from F_g to F_w is now simply

$$\mathbf{C}_{w/g} = \mathbf{C}_{w/n}\mathbf{C}_{n/g} \quad (2.95)$$

The gust vector can now be introduced in F_n as

$$\mathbf{v}_{W/n}^n = \mathbf{C}_{n/g} \begin{bmatrix} u_g \\ v_g \\ w_g \end{bmatrix} \quad (2.96)$$

which may be used in (2.67), and

$${}^n\dot{\mathbf{v}}_{W/n} = \mathbf{C}_{n/g} \begin{bmatrix} \frac{du_g}{dt} \\ \frac{dv_g}{dt} \\ \frac{dw_g}{dt} \end{bmatrix} \quad (2.97)$$

which may be used in place of (2.29).

2.3 Wind Field Prediction

Due to physical and practical limitations, wind conditions cannot be known over the plan horizon. Wind over the plan horizon is approximated by projecting wind at the vehicle's position forward in space over the length of the plan horizon given knowledge of the wind field gradient obtained through previous measurements. The wind prediction becomes

$$\tilde{\mathbf{w}}_P(x_i, z_i) = \mathbf{w}_0 + \nabla \mathbf{w} \begin{bmatrix} x_i \\ z_i \end{bmatrix} \quad (2.98)$$

where x_i and z_i define the position of the aircraft relative to its position at the start of the plan horizon, \mathbf{w}_0 and $\nabla\mathbf{w}$ are wind velocity and gradient at the beginning of each plan horizon (the point when the wind field was last measured). These can be computed using a measurements of airspeed, inertial speed (measured using GPS) and acceleration.

The resulting wind prediction $\mathbf{w}_P(x_i, z_i)$ is only valid over relatively short plan horizons. When the length of time over which the prediction is applied is extended, the prediction will likely differ from the actual wind greatly. Thus, we see that short plan horizons are favored in order to maintain the accuracy of the wind predictions.

In reality, because the turbulent wind field is measured as a sum of sinusoids, as will be shown in Section 2.8.2, an analytical derivative of the wind speed is available at all times. Because it is assumed that no knowledge of the wind field is available however, this analytical derivative is left unused in favor of the numerical derivative described above.

2.4 Total Energy

With the stated goal of energy maximization, it is imperative that we have a relevant expression for vehicle energy. In this application, the vehicle's specific total energy (i.e. total energy divided by weight) is used. This allows for a degree of generalization in the results obtained. The resulting energy is expressed in units of distance which is convenient.

The vehicle's specific total energy is

$$e_{tot} = h + \frac{1}{2g_D} V_a^2 \quad (2.99)$$

where h is height above a datum. Note that in aircraft applications total energy is expressed using airspeed, not inertial velocity. This is because airspeed is a much more meaningful measure of the aircraft's present condition and ability to maintain flight than is inertial or ground-relative speed.

The expression for vehicle total energy is differentiated with respect to time to obtain the rate of change of total energy (i.e. total power)

$$\dot{e}_{tot} = \dot{h} + \frac{V_a \dot{V}_a}{g_D} \quad (2.100)$$

Due to the fact that meaningful energy gain can most easily be accomplished through longitudinal motion, the following derivation is done considering the longitudinal aircraft model developed in Section 2.2.11. The statements regarding an effective strategy for energy gain can be applied to a full six degree of freedom aircraft model as well. Substituting dynamics from (2.89) gives

$$\dot{e}_{tot} = \dot{h} + \frac{qS}{mg_D} (-C_D + C_T \cos \alpha) V_a$$

$$-\frac{\dot{w}_{i,x}}{g_D} V_a \cos \gamma + \frac{\dot{w}_{i,z}}{g_D} V_a \sin \gamma - v_a \sin \gamma \quad (2.101)$$

Recognizing that $\dot{h} = -\dot{z} = V_a \sin \gamma - w_{i,z}$ from (2.87),

$$\dot{e}_{tot} = -w_{i,z} + \frac{qS}{mg_D} (-C_D + C_T \cos \alpha) V_a - \frac{\dot{w}_{i,x}}{g_D} V_a \cos \gamma + \frac{\dot{w}_{i,z}}{g_D} V_a \sin \gamma \quad (2.102)$$

The change in energy with ground distance flown is $\frac{de}{dx}$, and $\frac{de}{dx} = \frac{\dot{e}}{\dot{x}}$. For flight in still air this will always be a negative quantity: gliding flight in still air implies loss of altitude or loss of airspeed, powered flight in still air implies fuel burned. For gliding flight $C_T = 0$, hence

$$\frac{\dot{e}_{tot}}{\dot{x}} = \frac{1}{V_a \cos \gamma + w_x} \left[-w_{i,z} - \frac{qS}{mg_D} C_D v_a - \frac{\dot{w}_{i,x}}{g_D} V_a \cos \gamma + \frac{\dot{w}_{i,z}}{g_D} V_a \sin \gamma \right] \quad (2.103)$$

In steady state $\dot{V}_a = 0$, and

$$\dot{e}_{tot}|_{ss} = \dot{h} \quad (2.104)$$

In steady state flight in still air, the optimal flight efficiency (in the sense of maximizing range) occurs when $\frac{\dot{e}}{\dot{x}}$ is maximized. This occurs when the aircraft flies at best L/D, i.e. when the ratio of lift to drag is maximized. In a glide L/D defines the forward distance flown for each unit of altitude lost, and this occurs at a specific air speed which is vehicle and configuration dependent. When the air is not still the speed for best range changes depending on the wind speed and direction.

This allows for a simple way of validating the performance of any controller attempting to maximize energy gain in a calm environment. Under these conditions, such a controller should drive the aircraft towards the airspeed for best L/D.

Examination of (2.103) reveals that energy gain may be exaggerated when flying through moving air. The first term within the brackets in (2.103) is used advantageously in the traditional thermalling strategies of glider pilots. Regions of lift, where $w_{i,z}$ is negative, lead to positive values of $\frac{\dot{e}_{tot}}{\dot{x}}$. Thus, an obvious strategy for energy gain is maximizing time spent in atmospheric lift. The problem faced by those seeking to develop an optimal path planning algorithm utilizing this atmospheric lift is locating it.

The second terms represents energy loss due to drag on an aircraft. Thus, a low-drag aircraft limits energy loss and facilitates energy gain, validating our choice of a glider model for the application of this controller.

The last two bracketed terms in (2.103) show how energy gain may be maximized by taking advantage of the wind gradients through which an aircraft flies. Because turbulent air often has steep wind gradients resulting from rapid changes in wind speed, these terms become very useful in the development of a gust soaring controller. Note that the vehicle's airspeed is multiplied by both of these values, indicating that if wind gradients are approached at higher speeds, more energy may be harvested.

Because z is defined as positive downwards, negative wind gradients (i.e. increasing upwards

wind, and increasing headwind) allow for positive energy extraction from the wind field.

2.5 Energy Maximization

Merely computing a control input which maximizes the instantaneous value of \dot{e} will result in maximizing \dot{V}_a , i.e. placing the aircraft in a dive. This will soon cause rapid altitude loss and ultimately results in loss of total energy. An onboard pilot with foresight will immediately realize that this is not an appropriate strategy. In order to avoid this problem of greedy gain a cost function which explicitly accounts for future energy gain can be implemented to guide the receding horizon controller.

In receding horizon control a sequence of control inputs which maximizes a reward over a planning horizon T_P is computed. This sequence of inputs is followed over a control horizon T_C , with $T_C < T_P$. Once the control horizon has been reached a new sequence of control inputs is computed.

2.5.1 The Cost Function

The choice of cost function can have a tremendous impact on both mission performance and the final trajectory or control policy. The purpose of the cost function is to guide the behavior of the receding horizon controller. With that in mind, a large amount of attention should be given to the formulation of a cost function that accomplishes the stated goals of the controller.

In order to effectively harvest atmospheric energy, the reward function must accomplish two goals. The first is to maximize energy gain over the entire plan horizon, and the second is to ensure that the aircraft is in a state to maximize energy gain over the following plan horizon. If the second objective is not met, the reward function will tend to cause the aircraft to pitch downwards at the end of a plan horizon (to maximize the value \dot{v}_a), gaining kinetic energy but giving up altitude and the prospect of energy gain over the next horizon. The second goal can be accomplished by placing the aircraft in steady-state flight at the end of each plan horizon.

The goals of maximizing energy gain over the plan horizon, enabling future energy gain, and placing the vehicle in steady state flight at the end of the plan horizon are combined in a reward function of the form

$$R = \kappa_1 \frac{1}{x_i(T_P) - x_i(0)} \int_0^{T_P} \dot{e} dt + (1 - \kappa_1) \left. \frac{\dot{h}}{\dot{x}} \right|_{T_p} + \kappa_2 \dot{v}^2 \Big|_{T_p} \quad (2.105)$$

where $x_i(T_P) - x_i(0)$ is the horizontal distance flown over the planning horizon and κ_1 is a weight used to ensure a balance between energy gain over the planning horizon and the possibility of future gains. The weight κ_2 is used to penalize vehicle acceleration at the final time step, placing the vehicle close to a trimmed flight condition at the end of the planning horizon. The energy maximization term is the first, while the following two terms are in place to ensure that the aircraft does not compromise future energy gains in the current plan horizon.

2.5.2 Aircraft Control Policy

A simplifying assumption is made that an aircraft controller is present that has immediate control over the pitch rate, Q . The problem is now to compute the sequence of pitch rate inputs which maximizes the reward function in (2.105) (roughly corresponding to energy gain) over the plan horizon assuming wind as in (2.98).

To ease computational requirements by reducing the dimension of the input space, the pitch-rate sequence is parameterized as a spline with control points \mathbf{q} evenly spaced in time across the plan horizon. The use of a spline also ensures smoothness of the inputs and gives control over the both pitch rate Q and the pitch acceleration \dot{Q} at the beginning and end of the input sequence. A smooth control input is preferable as steep changes in commanded pitch rate can lead to decreased efficiency of the flight path, and more stress on the vehicle's airframe. A pilot will recognize that smooth control inputs are preferable to sudden changes in pitch attitude. Similarly, it is rare to observe a bird making sudden changes in pitch attitude when smooth maneuvers are preferable.

To incorporate realistic control surface deflection limits, a maximum pitch rate constraint is imposed. The spline generated is then used as an input to the aircraft equations of motion.

The problem of energy maximization can now be cast as a parameter optimization problem:

$$\text{maximize} \quad R(\mathbf{x}, \mathbf{q}) \quad (2.106)$$

$$\text{subject to} \quad \dot{\mathbf{x}} = f(\mathbf{x}, Q, \mathbf{w}, \dot{\mathbf{w}}) \quad (2.107)$$

$$Q = \text{spline}(\mathbf{q}) \quad (2.108)$$

$$Q(0) = Q(T_{C,prev}) \quad (2.109)$$

$$Q_{min} \leq Q \leq Q_{max} \quad (2.110)$$

$$Q(T_P) = 0 \quad (2.111)$$

$$\dot{Q}(T_P) = 0 \quad (2.112)$$

$$\mathbf{x}_{min} \leq \mathbf{x} \leq \mathbf{x}_{max} \quad (2.113)$$

The first constraint (2.107) on the maximization of the reward function is in place as a result of the architecture of the gust soaring controller. The controller seeks to maximize the reward function given the estimated wind field and a model of the aircraft presented in previous sections. This motivates a simple yet accurate aircraft model so that computations of aircraft states may be done quickly yet retain their applicability to the real-world aircraft.

The second constraint (2.108) merely states that the control policy that the aircraft follows Q must be that specified by the spline, itself dictated by the control points \mathbf{q} . This constraint links the optimized control point positions in \mathbf{q} to the enacted control policy Q .

The constraint on the initial value of Q in (2.109) ensures continuity of pitch rate inputs between consecutive control horizons. This ensures that no impulsive changes in pitch rate are commanded.

Pitch rates are limited by Q_{min} and Q_{max} in (2.110), values which are properties of the aircraft. This ensures that the sequence of pitch rate commands is realistic and trackable by the aircraft. Similarly, this constraint help to ensure that the aircraft is not driven unstable by the pitch rate commands.

In (2.111) the pitch rate $Q(T_P)$ at the end of the plan horizon is made equal to zero. This ensures that the aircraft is near steady trimmed flight at the end of each plan horizon. For the same reason, (2.112) forces the pitch rate acceleration $\dot{Q}(T_P)$ at the end of the plan horizon to zero.

The final constraint in (2.113) is a limit on the aircraft states \mathbf{x} . This simply requires that state limits cannot be exceeded when maximizing the value of the reward function.

The pitch-rate spline that is found to best maximize the reward function over the plan horizon (T_P) given the simulated aircraft and with predicted wind is now identified as u , the adopted control policy. The control policy is applied to the aircraft flown in the actual wind field and is followed through the control horizon (T_C). At the end of the control horizon, the wind is measured again, the future wind is predicted, and a new path is calculated. At this point the aircraft leaves the previous control input spline and begins to follow the new one. The components of the receding horizon controller are shown in Figure 2.3.

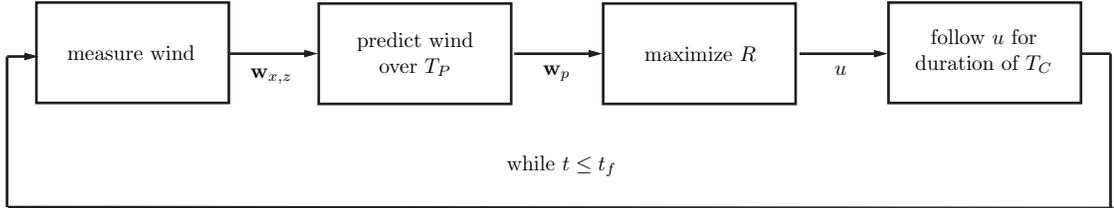


Figure 2.3. Control loop for receding horizon control applied to gust soaring. t_f is the end of the flight.

Note that the length of the planning horizon has not yet been specified, nor have the reward weights κ_1 and κ_2 . Too short a planning horizon will result in input sequences which are not long enough compared with the vehicle's time constant (resulting in diving to maximize the greedy energy gain at the cost of future gain); too long a planning horizon will result in wind predictions which are too much in error from the actual wind for effective path-planning. The control horizon will be fixed to end at the second spline point to prevent short plan horizons. The choice of planning horizon and weights is discussed in Section 3.4.

2.6 Computation of Problem Solution

The parameter optimization problem is cast as an interior point problem by re-defining state constraints as barrier functions. Barrier functions are put in place to limit those states that will result in instability or failure of the mission if the limits are exceeded. For example, airspeed V_a must be limited, as too slow of an airspeed will lead to stall and to high will exceed the structural limits of the aircraft. Angle of attack must be limited to prevent the aircraft from

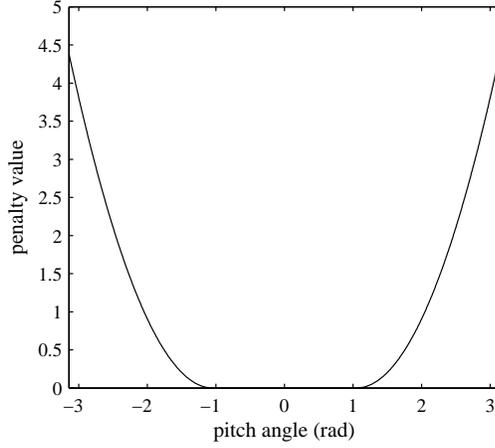


Figure 2.4. Example barrier function for pitch angle θ .

stalling and pitch angle must also be constrained such that manageable flightpaths result from the optimization.

The overall cost function to be minimized is a sum of the negative of the reward function (2.105) and the state barrier functions

$$Cost = -(R) + C_{barrier} \quad (2.114)$$

where

$$C_{barrier} = \int_0^{T_P} (C_{V_a} + C_{\alpha} + C_{\theta}) dt \quad (2.115)$$

The functions $C_{(\cdot)}$ represent barrier functions associated with the cost of exceeding a state limit during the planning horizon. These limits are determined by the vehicle's performance capabilities and are enumerated in Table A.2.

Each barrier function takes the form:

$$C_{x_i} = \begin{cases} 0 & \text{if } x_{min} \leq x_i \leq x_{max} \\ W_x (x_{min} - x_i)^\beta & \text{if } x_i < x_{min} \\ W_x (x_i - x_{max})^\beta & \text{if } x_i > x_{max} \end{cases} \quad (2.116)$$

with x_i representing a particular state (α , V_a , or θ) at each time step, W_x is a weight assigned to the violation of that state and β is an assigned distance metric exponent (here $\beta = 2$), and x_{max} and x_{min} correspond to the limits found in Table A.2. The barrier function for pitch angle θ is shown in Figure 2.4.

At this point the cost function is fully defined. A generic function minimizer (in this case MatLab's `fmincon`) can then be used to compute the optimal pitch rate input sequence.

Interior point problems require a feasible initial guess. For the first plan horizon in a trajectory, a constant pitch-rate input of 0 is used. The optimization function then finds the best

solution with this initial guess. For every plan horizon after the first, the initial guess for pitch-rate inputs is based on the remaining portion of the control input spline not followed in the preceding plan horizon. This method works well provided the wind conditions do not change drastically between plan horizons.

A baseline for performance evaluation is required for meaningful analysis of the developed receding horizon controller. A standard PID controller is used to track airspeed using pitch rate inputs in a similar manner to the receding horizon controller. For all cases it tracks the speed to fly for best L/D in still air. The resulting path generated by this controller is used as a baseline for comparison for all wind conditions simulated.

2.7 Controller Complexity

The complexity of the path developed as a result of (2.106) is an important variable in the gust soaring controller. A measure of complexity is the number of spline points used to generate the control input. Increasing the number of spline points will allow more complex (in the sense of greater variation over the time horizon) control inputs, but this comes at the cost of increasing length of time required to compute the optimal input sequence. Further, because of vehicle dynamics the aircraft will not be able to follow overly complex input sequences.

Intuition would suggest that too few spline points will not allow the vehicle to exploit the energy available in the atmosphere, or would not enable the vehicle to return to a neutral flight attitude quickly enough after attempting to exploit energy available. A three-point spline representing pitch rate input, for example, will allow the vehicle to pitch up but not pitch back down again.

For the simulations that follow, a spline described by five spline points was used. Shown in Figure 2.5, a four-point spline appears to yield a very similar control input to the five-point spline. In more complex wind fields, however, the four-point spline is unable to characterize a complex enough control input, and in simulation typically performs worse than the control input spline described by five points. The six-point spline can be seen over fitting the relatively simple wind field with no appreciable advantage in energy gain (note that the altitude and airspeed changes are nearly identical to the five-point spline).

2.8 Wind Fields

Two wind fields are used to evaluate performance of the designed controller. A discrete gust demonstrates (but does not prove) stability and shows that the receding horizon controller behaves as expected in steady winds. A turbulent wind field is used for determining the RHC control parameters, and demonstrates overall performance.

Both wind conditions are defined [22] to ascertain the effects of dynamic environmental conditions on the flying qualities of aircraft. Similarly, the wind fields were to be used to determine the ability of a pilot to recover an aircraft given upsets caused by atmospheric disturbances. In

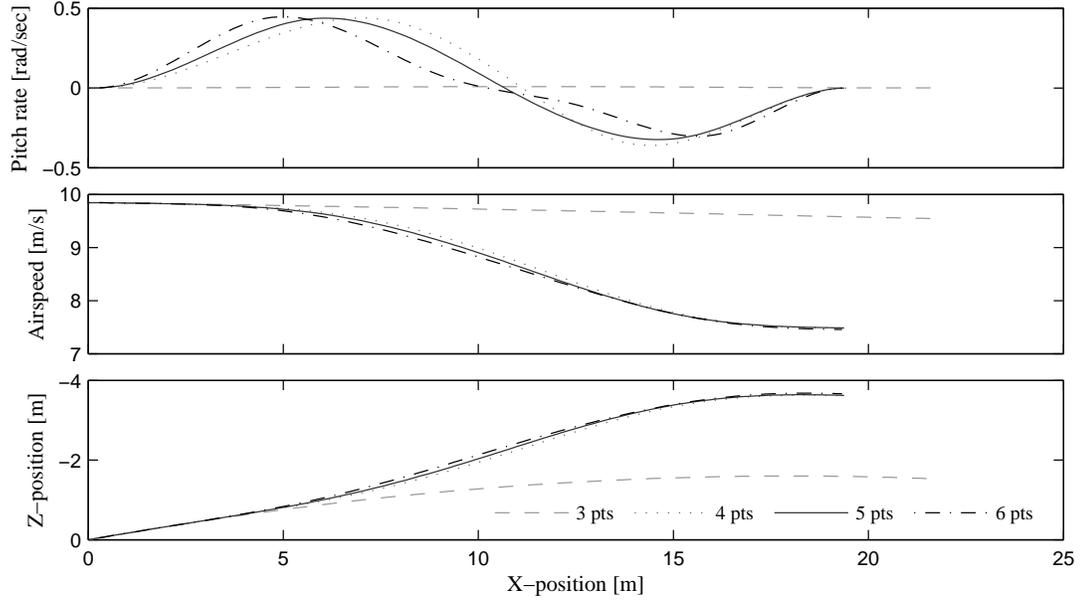


Figure 2.5. The above figure shows the effect of varying the number of spline control points on the optimal control policy. The longitudinal wind field traversed here is $\mathbf{w}_x 0 = 0$, $\mathbf{w}_z 0 = 0$ and $\frac{dw_z}{dx} = 0.1$.

this research, the gust soaring controller takes the place of the pilot, and rather than merely recover, energy gains are sought from atmospheric disturbances.

2.8.1 Discrete gusts

A discrete gust is given as [22]

$$w(x) = \begin{cases} 0 & x < 0 \\ \frac{w_m}{2} \left(1 - \cos \frac{\pi x}{d_m} \right) & 0 \leq x \leq d_m \\ w_m & x > d_m \end{cases} \quad (2.117)$$

where w_m is the magnitude of the wind speed and d_m is a length scale that defines the spatial gradient. In this work only longitudinal and vertical discrete gusts are simulated, as the gust energy harvesting controller is not formulated to take advantage of lateral gusts.

Flight through a discrete gust should show a transient as the aircraft takes advantage of the wind gradient followed by steady state (i.e. constant airspeed) flight as the condition for minimum energy loss is tracked.

2.8.2 Turbulence

The Dryden wind turbulence model has been used in gust energy harvesting simulations for small UAVs in previous research [31, 33]. Turbulence is modeled as a stochastic process based on the

Dryden spectral density function. The resulting wind field provides a realistic test case for the developed gust harvesting controller.

Wind velocities are determined a sum of sinusoids of differing wavelengths and phase angles [43]:

$$w_{(\cdot)} = w_{(\cdot),0} + \sum_{n=1}^N a_{(\cdot),n} \sin(\Omega_{(\cdot),n}s + \phi_{(\cdot),n}) \quad (2.118)$$

where (\cdot) represents a component (either w_g , v_g , or u_g), s is distance along the aircraft's flight path, $\Omega_{(\cdot)}$ is the spatial frequency and $\phi_{(\cdot),n}$ is a random phase angle for each sinusoid. The term $a_{(\cdot),n}$ is the amplitude of a particular sinusoid, and is computed based upon the desired power spectral density of the wind field:

$$a_{(\cdot),n} = \sqrt{2\Phi_{(\cdot)}\Delta\Omega} \quad (2.119)$$

Because s is the distance measured along the aircraft's flight path, when the aircraft six degree of freedom model is used, a separate state accounting for the flightpath distance must be tracked. When only longitudinal motion is considered, it is usually sufficient to substitute x distance traveled in place of s provided that the glide path angle γ is relatively shallow.

If the number of sinusoids summed, N , is made sufficiently large, the resulting wind field will be a continuous mathematical representation of turbulence. In this research, N was taken to be 41, with the spatial frequencies $\Omega_{(\cdot)}$ distributed logarithmically, favoring lower frequency sine waves.

The power spectral density for each of the components of a Dryden gust field is [22]:

$$\Phi_{u_g}(\Omega) = \sigma_u^2 \frac{2L_u}{\pi} \frac{1}{1 + (L_u\Omega)^2} \quad (2.120)$$

$$\Phi_{v_g}(\Omega) = \sigma_v^2 \frac{L_v}{\pi} \frac{1 + 3(L_v\Omega)^2}{(1 + (L_v\Omega)^2)^2} \quad (2.121)$$

$$\Phi_{w_g}(\Omega) = \sigma_w^2 \frac{L_w}{\pi} \frac{1 + 3(L_w\Omega)^2}{(1 + (L_w\Omega)^2)^2} \quad (2.122)$$

where L_u , L_v , and L_w represent turbulence length scales, and σ_u , σ_v , and σ_w represent turbulence intensities in the respective directions.

The Dryden turbulence model also specifies angular velocity disturbances due to turbulence. The spectra are given as [22]:

$$\Phi_{p_g}(\Omega) = \frac{\sigma_w^2}{L_w} \frac{0.8 \left(\frac{\pi L_w}{4b}\right)^2}{1 + \left(\frac{4b}{\pi}\Omega\right)^2} \quad (2.123)$$

$$\Phi_{q_g}(\Omega) = \frac{\Omega^2}{1 + \left(\frac{4b}{\pi}\Omega\right)^2} \Phi_{w_g}(\Omega) \quad (2.124)$$

$$\Phi_{r_g}(\Omega) = \frac{\Omega^2}{1 + \left(\frac{3b}{\pi}\Omega\right)^2} \Phi_{v_g}(\Omega) \quad (2.125)$$

where b is the aircraft wingspan. The angular velocity disturbances are computed as a sum of sinusoids through a function similar to (2.126)

$$\omega_{(\cdot)} = \sum_{n=1}^N a_{(\cdot),n} \sin(\Omega_{(\cdot),n}s + \phi_{(\cdot),n}) \quad (2.126)$$

where (\cdot) represents one of the three angular rate disturbances (either p_g , q_g , or r_g), and all other variables are the same.

The gust components u_g , v_g , and w_g are considered mutually independent and thus do not rely on each other. Likewise p_g is uncorrelated with any other gust component. However, it is easily seen from (2.124) and (2.125) that w_g and q_g are correlated as are v_g and r_g .

The gusts are assumed to act at the vehicle's center of gravity, with angular rate disturbances defined about the center of gravity. The structural modes of the aircraft are assumed to be small and are thus not accounted for. For practical purposes, the aircraft is assumed to be a rigid body.

Table 2.1. Summary of the turbulence parameters for each gust condition used in this thesis.

condition	description	intensity	length
1	low / light	$\sigma_u = 1.106$ m/s	$L_u = 200$ m
		$\sigma_v = 1.106$ m/s	$L_v = 200$ m
		$\sigma_w = 0.7$ m/s	$L_w = 50$ m
2	low / moderate	$\sigma_u = 2.212$ m/s	$L_u = 200$ m
		$\sigma_v = 2.212$ m/s	$L_v = 200$ m
		$\sigma_w = 1.4$ m/s	$L_w = 50$ m
3	medium / light	$\sigma_u = 1.5$ m/s	$L_u = 533$ m
		$\sigma_v = 1.5$ m/s	$L_v = 533$ m
		$\sigma_w = 1.5$ m/s	$L_w = 533$ m
4	medium / moderate	$\sigma_u = 3.0$ m/s	$L_u = 533$ m
		$\sigma_v = 3.0$ m/s	$L_v = 533$ m
		$\sigma_w = 3.0$ m/s	$L_w = 533$ m

The Military Specification also defines length scales and intensities depending on the altitude at which the turbulence is located. At low altitudes, the atmospheric boundary layer leads to heavily anisotropic gusts. While horizontal gusts are of approximately the same magnitude, gusts in the vertical direction are generally smaller.

At low altitude (below 1000 feet Above Ground Level, or AGL) the vertical component length scale is defined as $L_w = h$ and intensity is $\sigma_w = 0.1w_{20}$, where w_{20} is wind speed at 20 feet (6.1 m) AGL. Horizontal gust length and intensity is given by

$$\frac{L_u}{L_w} = \frac{L_v}{L_w} = \frac{1}{(0.177 + 0.000823h)^{1.2}} \quad (2.127)$$

$$\frac{\sigma_u}{\sigma_w} = \frac{\sigma_v}{\sigma_w} = \frac{1}{(0.177 + 0.000823h)^{1.2}} \quad (2.128)$$

with h altitude in feet.

In light turbulence, the value of w_{20} is taken to be 15 knots (7.7 m/s), and in medium turbulence $w_{20} = 30$ knots (15.4 m/s).

Above 1000' AGL gust fields are considered isotropic (i.e. $L_u = L_v = L_w$ and $\sigma_u = \sigma_v = \sigma_w$). For the purposes of this thesis, this is defined as medium altitude. Gust length is defined as 533 meters and intensity is determined based on atmospheric conditions: light turbulence has $\sigma_u = \sigma_v = \sigma_w = 1.5$ m/s and moderate turbulence has $\sigma_u = \sigma_v = \sigma_w = 3.0$ m/s. These gust parameters are summarized in Table 2.1 where each set of parameters corresponds to a gust condition that will be further explored in Chapter 4.

2.9 Summary

This chapter established the primary equations, models, and methods that will be employed elsewhere in this thesis. The aircraft models that will be used are defined and the equations relating aircraft motion to aircraft energy are given. Comments are made regarding the energy-optimal strategies for traversing wind fields that will be followed in Chapter 4 where simulation results are presented.

Application of an Evolutionary Algorithm to Optimize Control Variables

This chapter will describe the tools and procedures used to tune important control parameters. An evolutionary algorithm is used, and a brief background of evolutionary computation is presented in Section 3.1. Previous applications of evolutionary algorithms in a similar manner in control systems is shown, providing an initial justification for the use here. The rationale for using an evolutionary algorithm is continued in Section 3.2 with the specifics of the problem at hand addressed. The particular algorithm used, CMA-ES, is explained in Section 3.3. With the algorithm determined, the details of its implementation in the tuning of controller parameters is covered in Section 3.4. Results for each of the turbulent wind fields are given at the conclusion of this chapter in Section 3.5.

3.1 Summary of Evolutionary Algorithms

The selection of planning and control horizon lengths is essentially arbitrary, but can have a dramatic effect on the capability of the receding horizon controller to gain energy when faced with a complex wind field. Moreover, the choice of cost function weights can also have a tremendous impact on overall energy harvesting performance. With consideration for the importance of finding optimal values for the receding horizon controller's control variables, an evolutionary algorithm is used to find values of T_P , κ_1 and κ_2 which result in good flight performance over a wide range of wind conditions.

First evolutionary algorithms will be introduced in Section 3.1.1 followed by some justification for their use in this investigation presented in Section 3.2.

3.1.1 Fundamentals

Evolutionary computation, the broad field which encompasses tools such as the evolutionary algorithms used here, is a relatively recent technology. The foundations for this form of problem solving were introduced in the late 1950's, and by the mid 1960's, the primary structure of evolutionary computation as it is used today were established [44]. Though the technology was inspired from the start by imitating biological evolution, the term "evolutionary computation" was not coined until 1991 [45]. It was only during the mid-1990's that evolutionary algorithms became popular. At this point, computers were able to handle large population sizes and evaluated generations quickly enough for the evolution of a solution to become practical.

An evolutionary algorithm is a population-based problem solving method. The population is made of individuals that represent the candidate solutions to the problem. A population comprised of numerous individuals contains significant diversity in the solutions tested. A fitness function is used to compare each individual objectively and determine the extent to which it solves the problem posed.

The four mechanisms that characterize all forms of evolutionary computation are reproduction, random variation, competition, and selection based on the fitness of an individual. These four functions imitate their biological analogues closely. It has also been stated that any system featuring all four of these functions will inevitably result in evolution [46].

Reproduction refers to combining individual candidate solutions to produce offspring candidates that share traits from with the parent solutions. Reproduction typically leads to an internal search, where children are often substantially different from their parents, yet lie within the search space bounded by their parents.

As offspring are produced, random variation is introduced. In a natural setting, this is referred to as genetic mutation. Mutation is a form of local search, where the random variation may result in a child only slightly different from a parent, yet it introduces new information into the population. Mutation is highly valued in evolutionary processes for this ability to introduce fresh genetic material even as a solution is converging.

A method for competition must be specified whereby one individual may be compared to others in such a way that winners and losers are evident. This may take the form of a head-to-head competition where one individual prevails and the other is rejected, or by some other method by which two or more individuals are compared with one clearly winning. This is often the defining feature between different evolutionary algorithms.

Selection must occur based on a specified objective function such that the most fit individuals are propagated to the next generation and the population approaches a solution. Fitness functions are input by the user of the algorithm and are what is being solved as evolution progresses. It is the fitness function's complexity and the computational expense of evaluating the fitness function that exposes the biggest drawback to the application of evolutionary algorithms.

With the four mechanisms that act on an evolutionary algorithm in place, it is up to the algorithm's designer to determine how each is affected and to what extent each has an effect on

the problem solution. These are typically tailored for the intended application. For a more in depth look at the specifics of all facets of evolutionary computation, the reader is directed to *The Handbook of Evolutionary Computation* [47] as a comprehensive resource.

3.1.2 Application in Control

Evolutionary algorithms applied to control problems may be utilized in two ways. The algorithm may be used in an off-line capacity to design the controller with the system correspondingly controlled using the now static controller. Alternatively, evolutionary algorithms may be used actively as part of the control system; in this case, they act in an on-line capacity. The later method requires that the evolutionary algorithm return useful results rapidly compared to the time scale of the controller. On-line applications of evolutionary algorithms have the advantage of being able to adapt to systems that change with time [48].

Though on-line path planning is intriguing, it is impractical for the current application. As will be shown in Section 3.3, the convergence times for the evolutionary algorithm as employed are far too long to be functional for real-time use. Additionally, it is assumed in Section 2.8.2 that the properties of the wind field are static with time, negating the need for a controller that is able to adapt to a changing system. In this case, an evolutionary algorithm is used in an off-line manner. This means that no part of the controller itself shares features with evolutionary algorithms, rather that variables in the controller are optimally tuned through the use of an evolutionary algorithm. Off-line optimization has been the most popular and successful area for the application of evolutionary algorithms in the field of controls [49].

3.2 Justification for the Use of Evolutionary Computation Methods

Evolution is by nature an optimization process. The fittest individuals in a given environment are favored by selection processes. An evolutionary “dead-end” can be considered analogous to a local minima encountered during optimization. In natural systems, evolution may be credited with the development of many incredibly complex biological structures. These systems are not necessarily perfect, but they often represent the best solution given the environment in which that biological system resides with the given rules for selection [45].

Evolutionary algorithms can be naturally applied to optimization problems in a similar manner. The objective function specified is maximized (or minimized) as the individuals compete, resulting in the best solution. In the synthetic systems in which evolutionary algorithms operate the objective function fulfills the role of the environment. The competition between individuals results in very functional, unique solutions. At this point, a caveat should be injected, that the solution arrived at by evolution, be it in a natural setting or synthetic, is not necessarily the simplest. This is because it is only the output of the individual that matters (the *phenotype*), not the structure of the individual (the *genotype*), to the objective function.

It should be evident that evolutionary methods excel at solving certain types of problems when compared to classical optimization techniques. This is due to their unique approach to optimization, and their ability to combine both local and global search. In particular, evolutionary algorithms may be superior to classical optimization when the search space is characterized by any of the following features [48]:

- discontinuous
- non-differentiable
- multi-modal
- noisy

In the current investigation, all of these features may be present in the search-space. The complexity of the problem can be attributed to the complex and random wind fields over which the optimization takes place.

The search space in this investigation is far from simple. To make matters more difficult, the controller must not only be optimal for flight through a single wind field but must also be capable of generating a good control policy for any random gust field. It is because of the complexity of the optimization of the objective function that an evolutionary algorithm was settled upon as the best way to determine the control variables.

A final argument in favor of the use of evolutionary algorithms is that they work. That is they are often able to find very good solutions to complex problems. Evolutionary algorithms may be applied to real-world problems with no simplification and garner good results. This is in contrast to simplifying or linearizing a complex model such that a globally optimal solution can be found to the simple representation. In practice, a globally optimal solution to this less complex model may be of less value than an approximate (yet still good) solution to a complex real-world problem.

3.2.1 Downsides

With the advantages offered by evolutionary methods described in the previous section, some caveats to their use are regularly offered in the literature. If evolutionary algorithms are to be used, the disadvantages of such an optimization method should be understood and caution used to avoid them. The use of evolutionary algorithms is unnecessary, and results in inefficient use of resources when the problems to be solved are not mathematically complex. Convex optimization problems in particular are better solved using gradient descent approaches. This includes optimizing objective functions that are linear or quadratic, unimodal, or easily separable.

Additionally, evolutionary algorithms are not particularly efficient; this method of problem solving trades efficiency for an increase in generality of application. Evolutionary algorithms add a significant amount of complexity to the problem solving process, and as such, are to be avoided if a simpler solution is available [48].

Evolutionary algorithms provide no insight into the theory behind a problem's solution, but rather only promote the individual that offers the best results. Thus, the application of an evolutionary algorithm should only be used when the theory behind a system is either fully understood or when the theory is not pertinent to the investigation at hand. In the latter case, the evolutionary algorithm acts like a black-box optimizer [48].

The most pronounced complication encountered over the course of this research was that evolutionary algorithms have a way of exploiting flaws in the model or objective function. From the author's experience, the results of any solution arrived at by the use of an evolutionary algorithm must be examined carefully for flaws in the initial setup, and used with caution.

3.2.2 Use in Related Literature

Evolutionary computation has been recognized as a powerful tool in many engineering disciplines. In the specific areas of UAV systems and controls many researchers have used evolutionary algorithms to further the field. A selection of the literature is presented here to acquaint the reader with the pervasiveness of evolutionary algorithm applications.

Fleming and Purshouse [49] give numerous examples of evolutionary algorithms used in control systems engineering. In particular, evolutionary algorithms that optimize control parameters for such control schemes as neural networks, fuzzy logic controllers, and in system identification are surveyed. These systems are difficult to tune with conventional optimization techniques. This is relevant for the current research because it shows that evolutionary algorithms are being used for similar purposes, the off-line design of complex controllers, elsewhere in the literature.

A significant amount of work has been done on optimal path planning utilizing evolutionary approaches for many different types of autonomous vehicles. Witt and Dunbabin [50] and Fogel and Fogel [51] as well as others [52] have discussed path planning for Autonomous Underwater Vehicles. Path planning for underwater vehicles can be similar to aircraft path planning due to similar spatial degrees of freedom. This is in contrast to the plethora of research on traditional robot path planning in two dimensions.

Off-line control tuning for UAVs was done by Nikolos et al. [53] for the purpose of navigating an obstacle filled environment. Rathburn and Capozzi [54] published work on path planning using evolutionary algorithms for obstacle avoidance as well. Their work focused on obstacles with uncertain locations and the development of a cost function to best avoid those obstacles. In a follow-up work, Rathburn et al. [55], the previous method is applied to time-variant obstacle locations and the corresponding avoidance of those impediments.

Recent work by Mittal and Deb [56] presented a three dimensional, off-line path planner for UAVs. Two objective functions are used, the first with the goal of minimizing total path length, and the second concerned with maximizing the margin of safety when navigating through known terrain. This multi-objective method recognizes the often conflicting objectives involved in piloting a UAV such as the desire to avoid obstacles in the terrain yet follow a minimum distance path to the target.

Obviously, evolutionary algorithms have been met with acceptance in the controls and UAV systems fields. The current research will attempt to apply an evolutionary algorithm as a tool to optimize the control parameters for a receding horizon controller for gust energy harvesting.

3.3 Algorithm Choice

Now that evolutionary algorithms have been established and found to be widely used in the literature and as powerful tools for optimization as well as for the solution of difficult problems, the algorithm used in this investigation is specified. The reader is reminded that in this investigation, an evolutionary algorithm is used as an off-line tool to optimize the controller's parameters. After the controller is decided upon, the control strategy for the aircraft is no longer dependent on the evolutionary algorithm. This allows more freedom in the choice of algorithm, as the speed at which the algorithm converges to a solution is not a factor in the choice.

Because a realistic wind field may be (and often is) very complex, local optimal can abound in the search space. This, coupled with non-linear state and path constraints, necessitates an optimization strategy that can effectively cope with the rugged search landscape. A modern evolutionary algorithm, the Covariance Matrix Adaptation - Evolution Strategy (CMA-ES) [57, 58], which was developed for difficult, non-linear optimization problems, is used here to determine globally optimal values of the controller parameters.

In the optimization of the controller parameters, the variables including both time scales and coefficient weights, are difficult (if not impossible) to separate. The CMA-ES algorithm chosen is able to internally adapt the search distribution for such non-separable objective functions. The covariance matrix used gives a second order model of the system that measures and uses the interaction between parameters in the reward function (κ_1 and κ_2) and the planning horizon T_P .

The CMA-ES algorithm uses small population sizes for faster convergence times. Even with this precaution, however, it is infeasible to compute optimal values in real time. It was found that for the decision variables to converge to within acceptable limits, nearly 36 hours of computation time were required. The optimizations performed in this research are done with MATLAB R2010a on a single core of a four core Intel Xeon (E5506) processor running at 2.13 GHz.

Though a penalty is paid in computation time over more traditional optimization functions, the use of an Evolutionary Algorithm, which initially utilizes a broad global search before optimizing promising local regions, ensures that a global maximum is found within the objective function. Alternatively, the use of an Evolutionary Algorithm is far more efficient than a brute-force method by which all possible combinations of decision variables are tested before finding the ideal combination. The determination of plan and control horizon lengths as well as the weight coefficient must only be done once for each turbulence condition, and are held constant afterwards.

CMA-ES has been applied to numerous problems in widely varied academic and engineering fields. Among other applications in aerospace engineering, CMA-ES has been used in the optimization of the nose shape of hypersonic vehicles [59], drag reduction in airship bodies [60],

and in several studies designing airfoils for specific applications [61, 62]. The generality of these applications speaks to the robustness of the algorithm.

3.4 Determining the Planning Horizon and Reward Parameters

At this point the chosen algorithm, CMA-ES is applied to the optimization of the control parameters. The specifics of the application will be detailed along with the results of evolutionary optimization.

Stochastic turbulence is a challenging flight environment. Here, plan parameters, T_P , κ_1 , and κ_2 were optimized for flight through Dryden gust fields. A population of candidate controllers is flown through a gust field and the fitness of each member is evaluated based on flight performance. To avoid the problem of over-specialization (i.e. a controller designed for excellent performance in one particular instance of a gust field) fitness is determined based on the average of flights through N instances of the random gust field. The same N gust fields were used at each generation, and the fitness function is based on the improvement of flight efficiency over a constant airspeed flight:

$$\Delta \left(\frac{\Delta e}{\Delta x} \right)_{avg} = \frac{\sum_{n=1}^N \left[\left(\frac{\Delta e}{\Delta x} \right)_{GS} - \left(\frac{\Delta e}{\Delta x} \right)_{Cv_a} \right]}{N} \quad (3.1)$$

This value was maximized, and considered converged when the value varied by less than 1%. Figure 3.1 shows graphically the process by which the Evolutionary Algorithm is used to find the optimal control variables for the developed gust energy harvesting controller.

3.4.1 Evolutionary Algorithm Control Variables

Evolutionary algorithms are advantageously robust to poor control parameter choices. The random nature of the search and the learning that is inherent in the evolutionary methods applied, means that even if an optimization is set up poorly, it does not preclude the generation of valuable results [48]. That said, it is important to have some grasp of neighborhood in which parameters should lie.

Perhaps the most important advantage of the CMA-ES algorithm is that it is essentially parameter free. Population size, parent population, and standard deviations are all internally determined. The algorithms author states that the specification of control parameters should be part of the design, and thus little is left to choice for those who implement it [58]. This makes the algorithm easily applicable to a wide range of simulated wind conditions and decision variable setups. In this case, the initial population was seeded with variables that appeared to be optimal by hand-tuning.

The convergence criteria were determined manually such that the algorithm ceased attempting to improve on solutions when further improvement was too small.

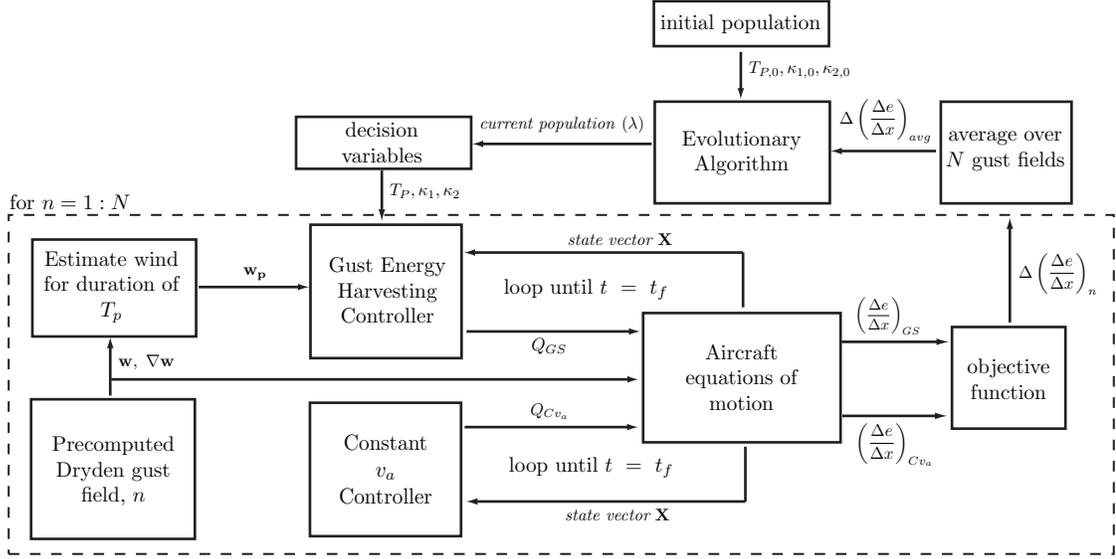


Figure 3.1. Design process for finding the control variables for the Gust Energy Harvesting controller by using an Evolutionary Algorithm. The section of the chart in the dashed box indicates that each population member (a unique set of control variables) is tested in N different wind fields.

As the evolutionary algorithm is not the primary focus of the current research, this hands-off approach to applying CMA-ES to the current problem is justified. The algorithm is used primarily as a black-box optimizer to obtain the most effective controller given the current setup. The evolutionary algorithm is a powerful tool but not a primary research interest.

3.4.2 Algorithm Convergence and Interpretation

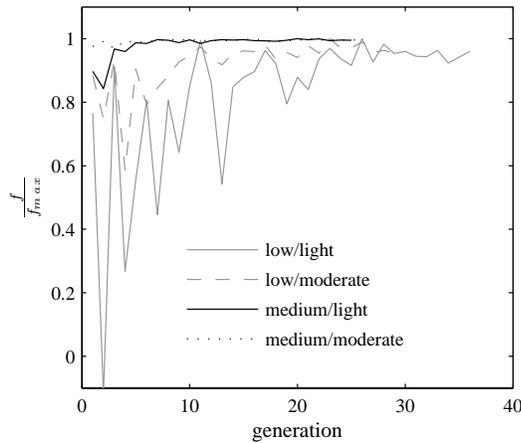


Figure 3.2. A convergence study of the CMA-ES algorithm in each of the four wind cases. Optimized are the values T_P, κ_1 and κ_2 , such that the value $f = \Delta\left(\frac{\Delta e}{\Delta x}\right)_{avg}$ is maximized.

Figure 3.2 shows a convergence study of CMA-ES algorithm as it attempts to maximize the reward function. The value of the objective function (3.1) normalized by the eventual converged fitness value is plotted versus generation. Only the best population member in each generation is shown. The plots do not monotonically converge because the CMA-ES algorithm does not account for elitism when generating the population members of the subsequent generations. Elitism is a property of some evolutionary algorithms where the most fit individual proceeds to the next generation unaltered.

The controllers optimized for the low altitude wind cases took longer to converge than did those in the medium altitude turbulence fields. This is likely because the low altitude wind fields are more complex, with shorter gust wavelengths, and are thus harder optimization problems to solve. The medium altitude turbulence cases feature longer wavelength gusts which are more easily handled by the controller, and thus converge more quickly.

Also shown in Figure 3.2 is that at both altitudes, the moderate turbulence controller converges more quickly than does the controller optimized for light turbulence. The moderate turbulence wind fields are more energetic, and thus there is more differentiation between a poor controller and an adept controller in the same wind field. This gives the algorithm a strong signal in the optimization, and results a solution that converges quickly.

The speed at which a parameter converges when using an evolutionary algorithm to find an optimal solution depends strongly on its marginal contribution to the fitness function [63]. A control variable that, over a wide range of possible values, affects the outcome of a fitness evaluation little, will converge slowly. Likewise, a problem where the solution space is flat with regards to the input parameters will converge more slowly than one that has clearly preferable control variables.

3.5 Evolved Gust Soaring Controllers

The gust soaring controller parameters optimized for each of the four turbulence conditions are used to specify four different gust soaring controllers to be examined in this research. Controllers GS1 through GS4 correspond in ascending order to low altitude with low turbulence, low altitude with moderate turbulence, medium altitude with low turbulence, and medium altitude with moderate turbulence respectively. The performance of each controller will be explored in depth in Chapter 4.

Table 3.1 shows the final plan parameters for each of the four turbulence conditions. At low altitudes a shorter plan horizon is favored: this is likely due to the shorter turbulence length scale at low altitudes, which cause the predicted wind to diverge from the actual wind as the planning horizon increases. Because the length of the control horizon and the time between wind estimates are linked to the length of the planning horizon, a more complex wind field would for this reason be better traversed by a controller with a shorter planning horizon.

In all four wind conditions, κ_2 converged to a negative value. This verifies the conjecture in Section 2.5 where it was stated that greedy controller seeking gains in airspeed at the end of a

Table 3.1. Converged plan parameters.

gust conditions	gust parameters	T_P	κ_1	κ_2	controller
low/light	$L_w = 50$ m, $\sigma_w = 1.0$ m/s	1.77	0.61	-0.15	GS1
low/moderate	$L_w = 50$ m, $\sigma_w = 1.4$ m/s	1.13	0.31	-0.06	GS2
medium/light	$L_w = 533$ m, $\sigma_w = 1.5$ m/s	2.23	0.76	-0.12	GS3
medium/moderate	$L_w = 533$ m, $\sigma_w = 3.0$ m/s	1.85	0.71	-0.11	GS4

planning horizon would result in a poor long term trajectory. A negative value for κ_2 penalizes any rate of change in airspeed at the end of a planning horizon and thus drives the aircraft towards steady-state flight.

3.6 Summary

This chapter justified the use of an evolutionary algorithm, CMA-ES, in the tuning of control parameters. Also presented were the method of optimization used and results of the optimization when applied to the four different wind fields. The different results for each of the four wind fields resulted in four separate controllers that will be examined in the following chapter. Each of the four controllers has the same structure (that given in (2.105)), but differ in the three control parameters optimized (T_P , κ_1 and κ_2) for optimal performance in each of the four wind fields. A detailed performance analysis and comparison of the four gust soaring controllers is given in Chapter 4.

Simulation Results

The previously developed longitudinal aircraft model (derived in Section 2.2 and presented in Section 2.2.11) with the designed controller (described in Section 2.5 and tuned Section 3.5) is modeled in the specified turbulent wind fields (Section 2.8). The results of several such simulations are presented here. Turbulence intensity and altitude are varied to gain a full picture of the capabilities of the gust soaring controller. Generalized results garnered from Monte Carlo simulations are then presented to compensate for the random environment. An attempt is made to clearly show the ways in which the controller successfully navigates the turbulence and to elucidate ways in which the controller may fail. Similar results are presented for the case in which the aircraft model is imperfectly known.

4.1 Performance Evaluation

Using the gust soaring controllers developed in Section 3.5, simulated flight is conducted in discrete gusts and in Dryden gust fields with the aircraft's motion modeled by the longitudinal aircraft model.

4.1.1 Setup

The equations for vehicle dynamics are propagated in time using fourth order Runge-Kutta integration with a time step of 0.02 seconds.

The aircraft model used in these simulations is based on an Omega IIe radio controlled glider. Vehicle properties are given in Table A.1, and the configuration as simulated is shown in Figure 1.2(b). For this vehicle, best L/D is 25 at an airspeed of 9.81 m/s in still air. The aircraft's minimum sink rate is 0.37 m/s when flying at 9.21 m/s, and the aircraft's stall speed is 7.5 m/s. The aircraft employed can be considered typical of a small, high-performance UAV. All flights were begun with the aircraft in a trimmed flight condition at the airspeed for best L/D in still air.

4.1.2 Discrete gusts

Initial validation for each controller is done by flight through a discrete gust. The equations describing a discrete gust are given in Section 2.8.1. Discrete gusts are used in the classical literature to determine aircraft and pilot response to large atmospheric disturbances [42]. Here, they will be used to gauge the gust soaring controller's ability to compensate for large changes in wind condition.

Results are shown in Figure 4.1 for horizontal gusts (left side) and vertical gusts (right side) for the GS3 controller.

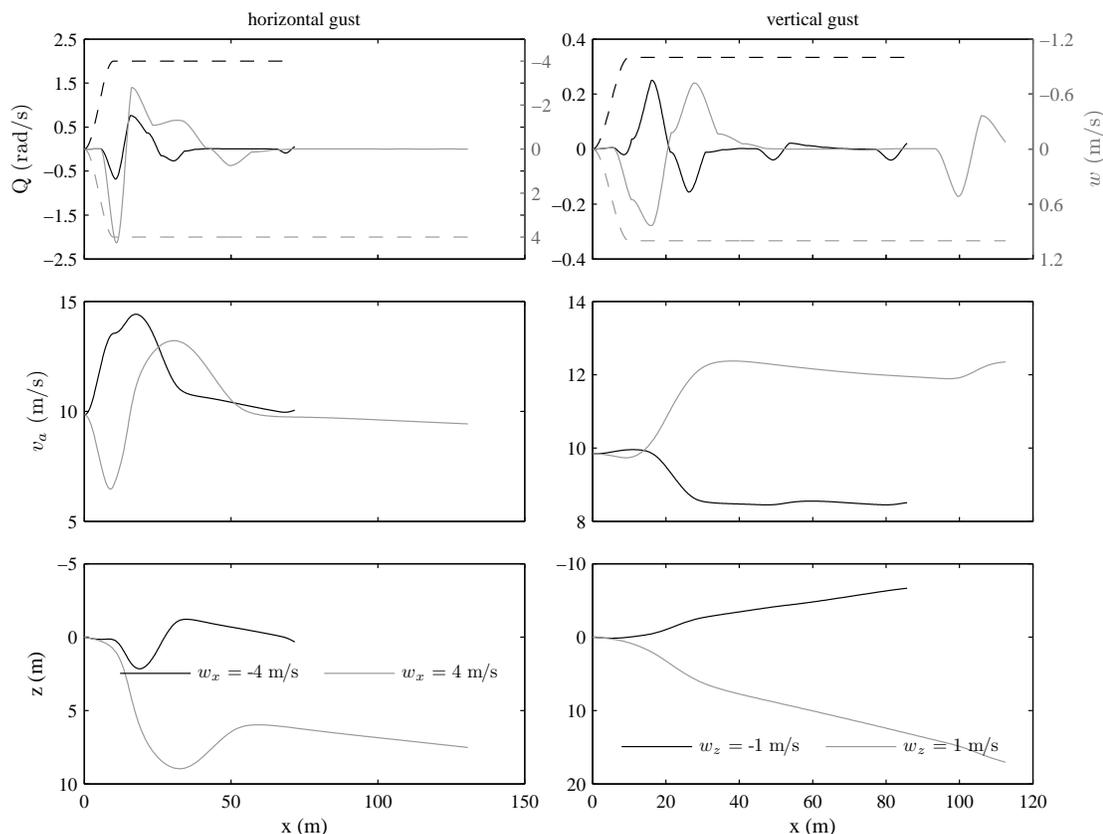


Figure 4.1. Controller performance in discrete gusts. The top plots show wind speed as dotted lines and the commanded control inputs in solid lines. Negative w_x denotes a headwind, negative w_z denotes upwards wind.

Figure 4.1 shows that the RHC strategy can result in a delay before appropriate control action is initiated. In the worst case, this delay is equal to one control horizon. The worst case delay is depicted here. At the beginning of the flight wind and wind gradient are zero, but quickly grow. This means that the controller incorrectly estimates the wind over the first planning horizon to be steady. Pitch rate input remains at zero until the beginning of the second plan horizon (at approximately $x = 8$ m). This delay means that more efficient energy harvesting should be

possible; however, the controller was not designed for discrete gusts.

Figure 4.1 also shows that the controller brings the aircraft near a steady state condition in the steady state part of the wind field. In cases in which the aircraft is flying with a tailwind or in an upwards wind (regions of atmospheric lift) the airspeed commanded is somewhat slower than in still air: this occurs to take advantage of these favorable conditions and minimize the energy lost with respect to distance traveled. In fact, energy can be gained from the vertical wind in this case (visible in the steady increase in altitude: recall that z is positive down).

Because an aircraft is able to take advantage of upwards moving air for energy gain, it is to be expected that the optimal speed in constant upwards wind will be lower than that in no wind or downwards moving air. Conversely, it is to be expected than in a downwards wind field, the aircraft will find it advantageous to fly faster to minimize the energy loss in this region. For similar reasons, to maximize flight path angle with respect to the ground γ_g the aircraft should fly faster in a headwind than in a tailwind.

As stated previously, the performance of the aircraft as the controller responds to a discrete gust disturbance is interesting but not the focus of this research. The manner in which the aircraft reacts can tell us if the controller is performing as expected; however, sub-optimal reactions should not be considered an indicator of a poor controller. Because continuous Dryden turbulence is the wind model of choice in this research, the remainder of this results chapter will focus on the ability of the gust soaring controllers to fly minimum energy trajectories through these wind fields.

4.1.3 Dryden Wind Fields

To assess the effectiveness of the gust soaring controllers (and to determine the effect of using a particular controller for a turbulence condition different from its design condition) Monte Carlo simulations of flights through Dryden gust fields were conducted. As with controller design (Section 3.5), four turbulence conditions were evaluated: low altitude/light intensity, low altitude/moderate intensity, medium altitude/light intensity and medium altitude/moderate intensity. Gust parameters are given in Table 3.1.

The method for optimizing the gust soaring controller parameters should mean that a controller designed for a specific wind condition will out perform the other controllers in that condition. This result will both validate our choice of controller parameters and show that a tailored controller for each turbulence condition is preferable to a general controller to be applied regardless of the wind field. If alternatively, one of the controllers is shown to perform well regardless of the wind condition, this will make the application of such a controller much simpler.

Fifty random gust fields were generated for each condition, and simulations of flights using each of the four gust soaring controllers were conducted for each of the gust conditions. In all cases the simulated flights lasted 8 minutes (for average flight distances of about 4800m in the low/light case and 4100m in the medium/moderate case), long enough for the mean wind speed experienced by the aircraft in each condition to approach zero. Because the simulated Dryden

turbulence fields are calculated as a sum of sinusoids, the mean wind speed after an indefinite amount of time should be equal to zero.

Results are summarized in Table 4.1 and shown graphically in Figure 4.2. Larger values of $\Delta e/\Delta x$ denote better performance, positive values of $\Delta e/\Delta x$ show that the aircraft actually gains energy during flight. Table 4.1 shows that the performance of the constant airspeed controller gets worse as the turbulence intensity increases: in still air $\frac{\Delta e}{\Delta x} = -0.0385$ is the optimal value. Conversely the gust soaring controllers show steadily increasing performance as turbulence intensity increases.

Table 4.1. Summary of results of Monte Carlo simulation of flight through Dryden gust fields.

Altitude / Intensity	Controller	mean	maximum	minimum	σ
Low / Low	GS1	-0.0375	-0.0321	-0.0440	0.0026
	Constant v_a	-0.0397	-0.0345	-0.0473	0.0029
Low / Moderate	GS2	-0.0303	-0.0121	-0.0423	0.0065
	Constant v_a	-0.0416	-0.0262	-0.0556	0.0071
Medium / Low	GS3	-0.0139	0.1262	-0.0954	0.0477
	Constant v_a	-0.0374	0.0736	-0.1220	0.0476
Medium / Moderate	GS4	0.0368	0.4046	-0.1557	0.1215
	Constant v_a	-0.0521	0.2164	-0.4444	0.1354

The degradation of the constant airspeed controller’s performance with increasing turbulence makes theoretical sense. The airspeed tracked by the constant airspeed controller is that for best L/D in still air. As turbulence intensity increases, the instantaneous deviation of the wind speed from zero (still air) becomes more pronounced. Thus, with more intense turbulence, the constant airspeed controller is tracking an increasingly off-optimal operation point.

Figure 4.2(a) shows the difference between the performance of the designed gust soaring controllers and the constant airspeed controller. In almost all cases the gust soaring controller outperforms the constant speed controller, indeed at medium altitude controllers GS3 and GS4 outperform the constant speed controller for all gust fields simulated. In the low altitude turbulence fields, the improvement is relatively small compared to what is possible at medium altitudes.

Figure 4.2(b) shows the absolute performance of each of the gust soaring controllers and the constant airspeed controller. Positive values of $\Delta e/\Delta x$ are those that lie above the solid horizontal line, and values of $\Delta e/\Delta x$ that outperform a glide at the airspeed for best L/D in still air lie above the horizontal dotted line. In the medium/moderate gust condition all of the gust soaring controllers are able to gain total energy during flight on average. At low altitudes, energy gain was not possible in any of the simulated wind fields; however, there is some evidence for improvement over a steady glide. This implies that the aircraft is gaining energy from atmospheric turbulence in most of the simulations run.

As one can expect each controller performs best relative to constant airspeed flight at the gust condition for which it was designed. For all four turbulence intensities, Figure 4.2(a) shows that the mean value of $\Delta(\Delta e/\Delta x)$ is greatest for the gust soaring controller that was designed

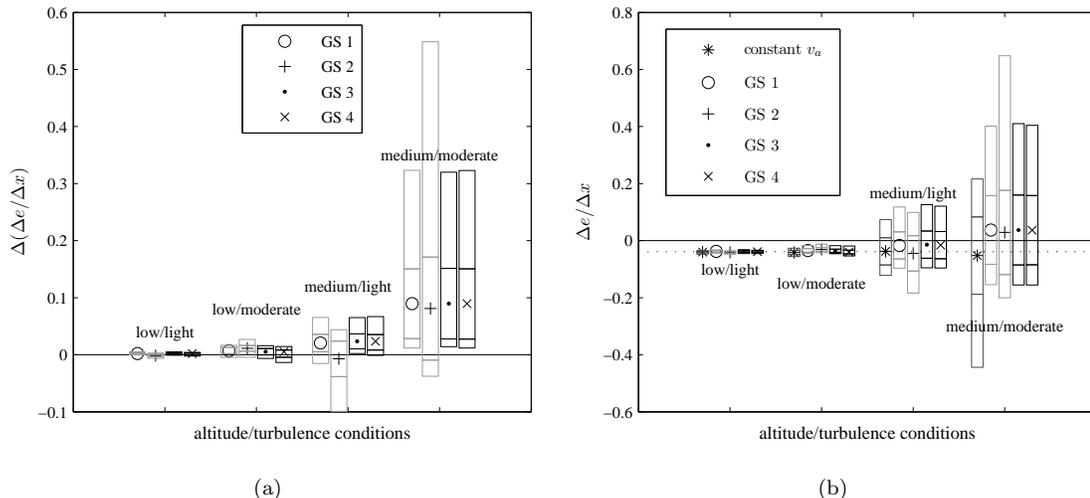


Figure 4.2. Performance of gust soaring controllers. Left: comparison with constant airspeed; Right: absolute performance. Symbols show mean energy change over 50 runs, inner bars show $\pm 1\sigma$, outer bars show minimum/maximum energy change.

in that environment. This is the expected result because Figure 4.2(a) is a plot of an identical metric to the function that was optimized in Section 3.5 to obtain the gust soaring controllers.

Controller GS3 (designed for medium altitude/light intensity turbulence) shows the best overall behavior, with good performance at all gust conditions. For this reason, GS3 will be used specifically for comparison and to clarify some of the features of the gust soaring controllers in later sections.

A two-sample t -test with significance level of 5% was used to assess the statistical significance of the observed difference in performance of each gust controller at the four gust conditions. The performance of each controller was compared with all others at each gust condition. In the medium/moderate gust condition the differences between controllers were not statistically significant, indicating that all four gust soaring controllers showed equivalent (good) performance. Note that Figure 4.2(a) shows that all four gust controllers showed better performance than the constant airspeed controller. In the other three gust conditions the difference between GS1 and GS3 is statistically insignificant. The difference between GS2 and the others is statistically significant in low/light, low/moderate and medium/light, but it is only better in low/moderate (its design condition).

Mean change in total energy is shown in Figure 4.3 for the gust soaring controllers and the constant speed controller. For comparison the energy change in still air at best L/D (the condition which maximizes range) is shown as a dotted line. Mean performance of the gust soaring controller is clearly superior to both constant speed and the still air flight conditions, and the performance improvement with increasing turbulence intensity is clear. The performance improvement associated with flight at higher altitude is likely due to the longer wavelength associated with turbulence, allowing the vehicle to better exploit the energy available.

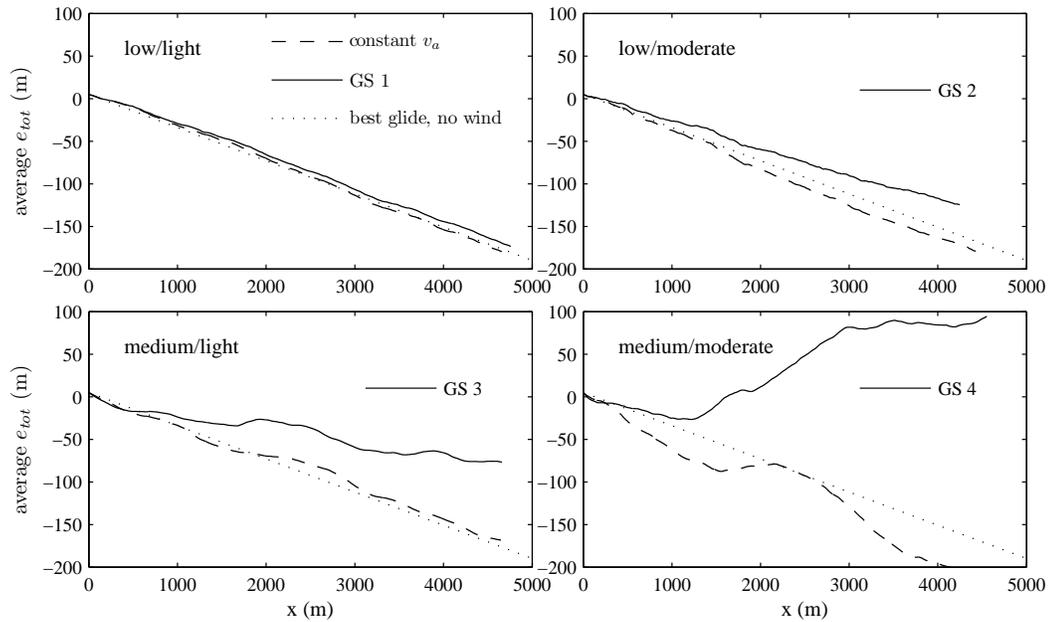


Figure 4.3. The average Δe_{tot} of all 50 runs of the Monte Carlo Simulation are plotted versus x distance in the above figures. The controller optimized for each wind condition is shown, as well as the performance of the constant airspeed controller. The change in total energy of the same aircraft flying at the airspeed for best L/D in still air is shown by the dotted line.

Detailed results of a single representative run (for controller GS3 in a medium altitude, light intensity gust field) are shown in Figure 4.4. For comparison the constant airspeed controller is shown as a dashed line and a steady glide in still air is shown as a dotted line. GS3 shows a steady gain over the constant airspeed controller, even in regions of unfavorable wind. The gust soaring controller reduces energy lost during flight in unfavorable winds (headwind and downwards gusts) by increasing airspeed (this reduces the time spent in an unfavorable region) and increases energy gain during favorable wind (tailwind and upwards gusts) by reducing airspeed. The cost of improved flight performance lies in the increased control actuation required: significantly higher pitch rate inputs are seen for the gust soaring controller.

In Figure 4.5 a comparison is given between the gust soaring controller and the baseline in each of the 50 gust conditions. In 49 of the 50 simulations run, the GS 3 controller performs better than the constant airspeed baseline. The exception is in the 30th simulation in which the constant airspeed controller fairs better, however the difference in energy at the end of the flight is very similar. In simulations 14 and 26 both controllers perform comparably. Regardless, in the vast majority of simulated flights, the gust soaring controller outperforms the baseline by a significant amount.

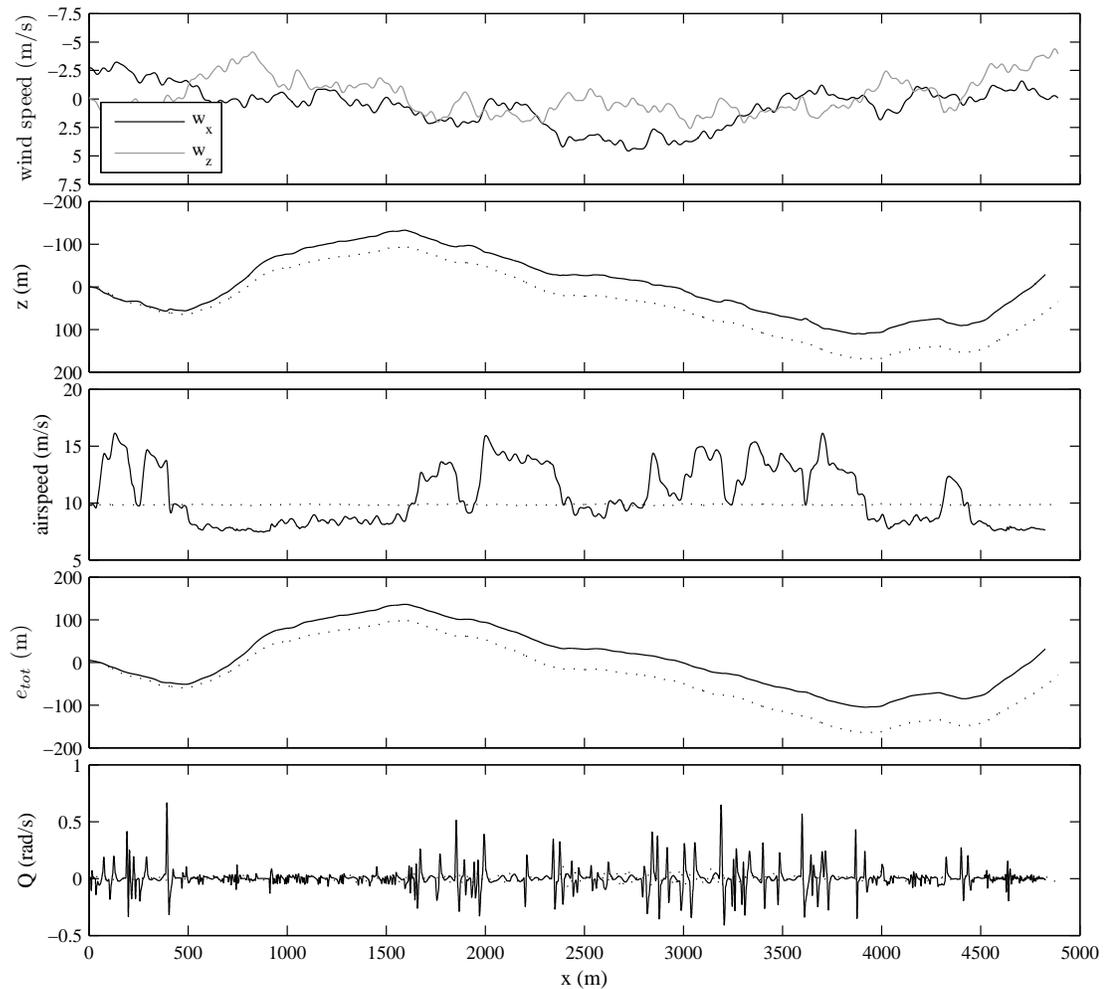


Figure 4.4. Flight history for a single run at medium altitude in light turbulence using the GS3 controller. The top-most figure shows wind speed (negative w_x is a headwind, negative w_z is upwards). Subsequent figures show flight path, air speed, total energy, and commanded pitch rate for the GS3 controller (solid) and constant airspeed controller (dotted).

4.2 Discussion

There is a significant amount of energy available in turbulence. Rather than viewing turbulence as a disturbance to be mitigated, a gust soaring controller seeks to exploit the energy that is available.

The main obstacle to gust soaring is the difficulty in obtaining an accurate wind prediction. The RHC approach avoids this problem by only computing control inputs over a fixed time horizon, which allows the use of a simple wind prediction. Basing the wind prediction on measured wind and wind gradient at the start of the planning horizon means that only local information is required, but this naturally assumes that wind and wind gradient can be measured using sensors

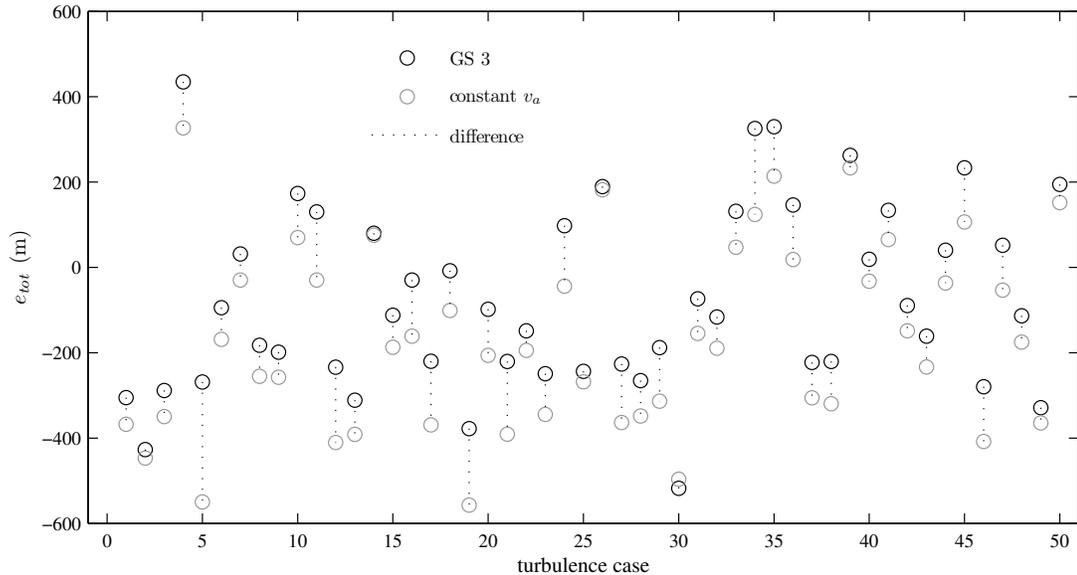


Figure 4.5. A summary plot showing the total specific energy at the end of each flight obtained by the GS 3 controller compared to the baseline in each of the 50 turbulence cases. The dotted line shows the difference between the the two values.

on board the aircraft. A combination of air-relative velocity (air speed and angle of attack) and inertial velocity (obtained using GPS) can be used to compute the three components of wind at the aircraft's current position; a combination of inertial measurements (accelerometers and rate gyros) and numerical differentiation of airspeed measurements can be used to compute wind gradients.

Clearly long horizons (both planning and control) will result in significant errors in the wind prediction, potentially reducing performance. Another trade-off inherent to RHC lies in the length of the planning horizon: if perfect *a priori* knowledge is available then a longer planning horizon results in better paths. However, the cost is increased computation time: too long a plan horizon will make real time computation of plans impossible. In the case of gust soaring via RHC shorter plan horizons give better performance as gust intensity increases: this is likely due to the increase in prediction error over longer plan horizons.

Reducing the length of the control horizon will result in improved ability to account for disturbances or changes in information. Preliminary investigations for this research also attempted to include control horizon as part of the plan parameters to be optimized, but it quickly became clear that the control horizon was converging to its minimum allowable value. Setting T_C to the second spline point (at $0.25T_P$) was done as a balance between re-planning ability and computational cost. Improving the computational speed associated with the trajectory optimization will allow shorter control horizons.

4.2.1 Computation Cost of the Controller

The computations performed in this research are done with MATLAB R2009a on a single core of an Intel Core 2 Duo processor running at 2.20 GHz. For each of the 50 runs in a Monte Carlo simulation, a 480 second flight through a gust field is simulated. With a planning horizon of roughly 1.75s (the average of the plan horizons of GS1 through GS4) and a control horizon of $0.25T_P$, this results in roughly 1,100 plan computations in each run. The computation time statistics presented in Figure 4.6 are thus based on 55,000 plans for each controller in each gust condition.

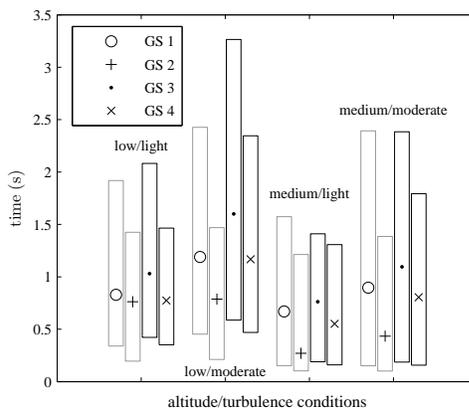


Figure 4.6. Computation time for the gust soaring controllers. Symbols show mean computation time; bars span the 5th and 95th percentile computation times.

The results show that for each altitude, it takes longer to compute a solution in the moderate turbulence fields than in the low intensity turbulence. The cause of this is likely that the more energetic wind fields prove to be more difficult problems to solve for the function minimizer `fmincon`. Additionally, for all wind conditions, the same order of mean computation times is preserved: a solution for controller GS2 is computed most rapidly, followed by GS4, then by GS1, and finally GS3 is the slowest. In the context of the controller parameters presented in Table 3.1, GS2 has the shortest planning horizon T_P and thus the function minimizer must simulate less aircraft motion than in GS3 where the T_P is the longest. From a computation standpoint, this seems to favor controller GS2, however, the longer planning horizon in GS3 allows the function minimizer more time to compute a solution.

Real time operation implies that a plan can be computed in less than one control horizon. While the time required to generate a plan is less than the planning horizon in more than 90% of cases for all controllers, it is rarely less than the control horizon. Significantly better performance should be possible if a solver specific to this problem is used, and this will of course be necessary for any hardware implementation.

In the simulations presented, the optimal path is always available before the current planning horizon begins. This is an impossible assumption to satisfy. However, if the time it takes to

arrive at a solution is sufficiently short, a set of control inputs will be available rapidly for the aircraft to follow. The best function minimizer in this case may be one that is able to present a good solution at any time and refines that solution with the time available. The function minimizer chosen and the time required to return a good solution will have a definite impact on the performance of the receding horizon gust soaring controller.

4.3 Model Uncertainty

Receding horizon control is an example of model predictive control, where a model of the system is used to test a proposed input. The effect of model errors on performance is thus a critical question. If the controller developed is extremely sensitive to model errors, it may be of no use in any real world application. The best-case scenario would see little degradation in controller performance as the aircraft model is varied.

4.3.1 Robustness of the Controller

The receding horizon controller discussed previously uses an assumed nominal aircraft model to compute the optimal sequence of pitch rate inputs. To assess the effect of model errors, the GS3 controller is used to control an off-nominal aircraft and flight performance is determined using Monte Carlo simulations. The simulation setup is identical to that described in Section 4.1.1, the only difference being the changed aircraft model.

An off-nominal aircraft is one whose actual parameters vary from the nominal parameters by some amount. Note that the controllers in Table 3.1 were optimized with a given aircraft model and any change from that model, whether it improves performance or degrades it, is considered an off-nominal aircraft. The GS3 controller was not altered for these additional experiments.

Typically the most difficult parameters to estimate are those associated with drag. Drag is a highly nonlinear force and there are numerous effects (for example, interference drag at the intersection of the wings and fuselage) which are extremely difficult to predict with current technology. At the same time drag is a critical parameter affecting overall aircraft performance. The effect of off-nominal drag on the performance of the controller is therefore assessed.

The term that is primarily composed of skin friction (which is independent of lift coefficient, and is represented by the constant term in (2.93)) is varied by 20%. Simulations are conducted for the nominal aircraft model is modified by increasing parasite drag by 20% and decreasing parasite drag by 20%; both are compared to the nominal aircraft. The variation in drag has the effect of changing best L/D from a nominal value of roughly 26 (with air speed for best L/D 9.84 m/s) to 22 (with air speed 9.6 m/s) for the increased drag case and 31 (with air speed 10.2 m/s) for the reduced drag case. Monte Carlo simulations similar to those in Section 4.1.3 are run (although a new set of random gusts were used). Results are presented for controller GS3.

Results of the Monte Carlo simulations are summarized in Figure 4.7. The effect of changes in parasite drag on mean performance is slight, with increasing drag resulting in slightly decreased

performance relative to the constant airspeed controller and decreasing drag resulting in slightly increased performance relative to the constant airspeed controller. The trends are to be expected;

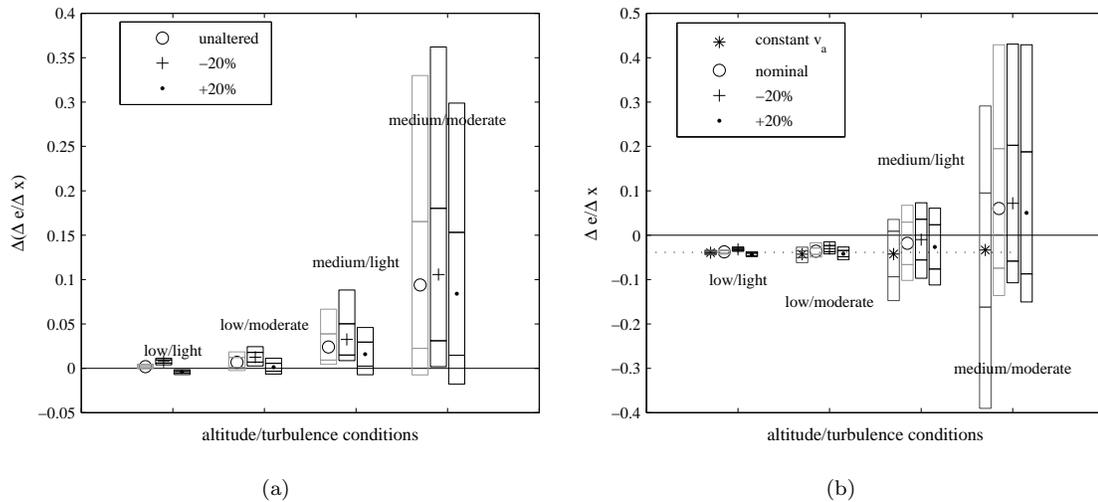


Figure 4.7. Robustness of the GS3 controller to changes in parasite drag. Left: comparison with constant airspeed; Right: absolute performance. Symbols show mean energy change over 50 runs, inner bars show $\pm 1\sigma$, outer bars show minimum/maximum energy change.

the lower drag aircraft should exhibit improved performance regardless of the environmental conditions and a higher drag aircraft will show decreased performance. The important result here is that the GS3 controller is capable of controlling the off-nominal aircraft and similar energy harvesting to that displayed in Figure 4.2 is shown. In the medium altitude turbulence conditions, the gust soaring controller is able to outperform the constant airspeed controller regardless of aircraft configuration.

Figure 4.8 shows the change in total energy with distance traveled for the nominal and off-nominal aircraft. These results may be compared to those in Figure 4.3, however, note that a different set of wind fields are used. The graceful change in performance indicates good robustness to uncertainty in parasite drag: the change in performance is most likely entirely due to the change in aircraft properties, not to the controller operating in an off-nominal condition. The aircraft travels further in the lower drag scenarios because the airspeed at which best L/D is achieved is higher in these cases. Obviously, if an aircraft designer is able to reduce total aircraft drag, it will result in performance gains for the aircraft. Similarly, the airspeed for best L/D is slower for the low drag scenarios and the corresponding aircraft travel shorter distances during the same simulation time.

4.4 Summary

The results of the investigation detailed in this thesis were presented in this chapter. Simulated flights through differing intensity levels of Dryden turbulence were conducted and analyzed.

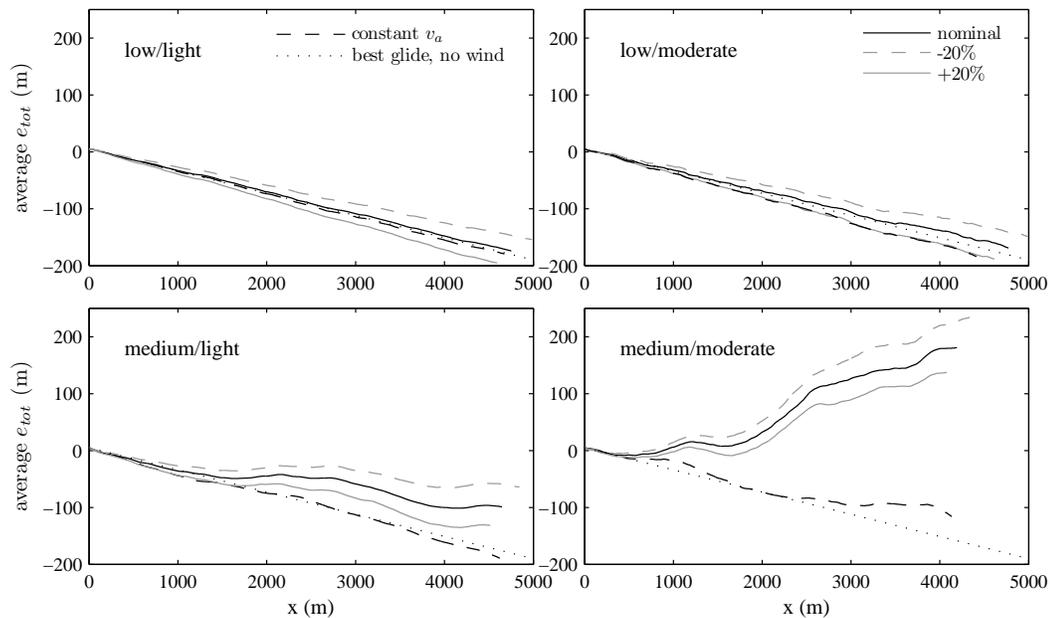


Figure 4.8. The average Δe_{tot} of all 50 runs of the Monte Carlo Simulation are plotted versus x distance in the above figures. The GS3 controller is shown with the nominal aircraft as well as the two off-nominal cases, and the constant airspeed controller for the nominal aircraft is shown. The flight path of the same aircraft flying at the airspeed for best L/D in still air is shown by the dotted line.

The effects of an imperfectly known aircraft model on the performance of one of the controllers was also simulated and analyzed. Finally, the computational requirements of the gust soaring controller were explored.

Conclusions

The limited endurance of small UAVs compared to their larger counterparts is a severely limiting factor in their deployment. These small aircraft have many advantages, the primary ones being low cost, lack of large ground support facilities, and the ease with which they can be deployed. Birds can be considered natural counterparts to small UAVs and are not so limited in their range or flight times. This is because birds are adept at harnessing the energy available in the atmosphere to augment their flights. Atmospheric turbulence is a large source of atmospheric energy that is, for the most part, unharnessed in contemporary applications.

This thesis describes a method for extracting energy from atmospheric turbulence through longitudinal aircraft control. The complex nature of the environment, the randomness of the variations in local wind speed, and the short time scales in which turbulence changes make this an especially difficult problem. The approach proposed is based on receding horizon control (RHC), where a path which maximizes energy gain over a fixed time horizon is computed. The structure of this approach solves many of the encountered problems, and special care is taken so that the aircraft does not place itself in an unrecoverable condition.

A reward function which combines energy gain over the plan horizon while preserving the potential for future energy gain is defined. This removes a tendency for “greedy” control. The reward function is combined with state cost functions that penalize violations of state constraints, yielding an energy cost function. Aircraft control is dictated through pitch-rate inputs that are the result of the minimization of this cost function.

Parameters of the RHC and the reward function were tuned off-line using an evolutionary algorithm. The four controllers that resulted from optimization in each of the four simulated turbulence conditions were tested using simulations of flight through a discrete gust (to demonstrate stability) and a Monte Carlo simulation of flight through Dryden turbulence. Significant improvement over constant airspeed flight was observed; in many cases, it was even possible to gain total energy during flight through turbulence.

Robustness was demonstrated by using the developed controllers with off-nominal aircraft

conditions. Graceful behavior was observed: increasing parasite drag by 20% resulted in reduced performance (but still better than constant speed flight); decreasing parasite drag by 20% resulted in improved performance. This showed that the controller is applicable to more than just the specific aircraft configuration for which it was designed.

5.1 Summary of Contributions

5.1.1 Method for Dealing with Stochastic Nature of Turbulence

The receding horizon control architecture is selected due to the constraints that a turbulent environment puts on the knowledge available to the controller. In standard path planning applications, some knowledge of the atmospheric environment must be available in order to determine the optimal set of control inputs to achieve the desired path. When the atmospheric conditions are not deterministic, *a priori* information is not available and standard path planning techniques fail. Therefore, current and past wind knowledge is used to make a prediction of the wind conditions in the immediate future. The length over which this prediction is trusted is called the planning horizon, and the optimal set of control inputs is determined for this time-frame only.

With the receding horizon control method, the stochastic properties of turbulence no longer negatively impact the path planning process. The controllers developed in this research show good results when determining energy-optimal control strategies through simulated turbulent wind fields.

5.1.2 Development of Energy-Based Reward Function

The stated goal of this thesis is to provide a controller framework by which energy may be gained from atmospheric turbulence. This goal necessitates a reward function that considers energy gain while preventing the aircraft from entering a flight condition from which it cannot recover. The developed reward function successfully accomplishes this goal. When combined with cost functions that penalize the violation of state constraints, the reward function forms the basis of a gust soaring controller that both gains energy from turbulence and keeps the aircraft in stable flight through energetic wind fields.

5.1.3 Optimization of Controller Through Evolutionary Methods

The coefficients in the reward function (κ_1 and κ_2 in (2.105)) as well as the length of the planning horizon (T_P) are left unset in the formulation of the controller architecture. The values of these parameters have a large influence on the effectiveness of the gust soaring controller. In order to obtain optimal values for these parameters, an evolutionary algorithm is used to tune the controller for the best performance in each of four different turbulence conditions.

5.1.4 Performance Verification through Simulation

Four different sets of parameters are developed yielding four different gust soaring controllers, each one optimal for a different turbulence condition. It is shown that each controller out performs the others in the turbulence condition for which it was optimized. To show definitively the ability of the gust soaring controllers to take advantage of the energy found in gusty atmospheric conditions, Monte Carlo simulations were conducted. These simulations featured the controllers tested in numerous wind fields. The results are presented and conclusions regarding the extent to which each controller effectively harvests gust energy are drawn.

5.1.5 Demonstration of Controller Robustness

A gust soaring controller is of limited usefulness if it relies on an exact aircraft model to return adequate performance results. There are many factors that could result in an imperfectly modeled aircraft, not the least of which being the difficulty of developing such a model and the tendency of small changes in configuration during a mission (such as the accumulation of ice or insects on the wings) to degrade performance. For this reason, it is considered pertinent that the performance of the controller be tested on an off-nominal aircraft model. It is shown that even relatively large changes in the parasitic drag on the aircraft do not noticeably affect the controller's ability to harvest gust energy; any differences in performance may be attributed to the differing capabilities of the off-nominal aircraft model.

5.2 Recommendations for Future Research

5.2.1 Examination of Controller Variability

The results presented in Figure 4.5 show a significant amount of variability in both the performance of the controller but also the extent to which the gust soaring controller is an improvement over the baseline. Because of the random turbulence fields, some variability is expected, however, the apparent lack of a pattern in the results is surprising. The characteristics of a wind field in which the gust soaring controller outperforms the baseline controller should be examined and compared to those fields in which the gust soaring controller fails to meaningfully improve over the baseline.

5.2.2 Control Through Flap Actuation

In the longitudinal simulations conducted in this research, it was assumed that a controller had immediate control over the pitch-rate Q of the aircraft. When the simulation was extended to the six degree of freedom model, pitch rate control was enacted through elevator deflections. Etkin [21] argues in the case of tracking a prescribed aircraft load factor through a gust wind field (for the purpose of gust alleviation), that the time-lag associated with the pitching motion may be too great to obtain good results in this situation. The time lag may be a large source of

error in the current research as well. The solution proposed by Etkin is to use wing flaps, which almost immediately effect changes in wing lift, in conjunction with elevator control. This greatly reduces fluctuations in aircraft load factor and could be feasibly applied to pitch-rate tracking as well.

5.2.3 Hardware Implementation

Much work has been done to implement gust alleviation controllers on real-world aircraft. These controllers ideally seek to make the ride and handling qualities of an aircraft the same regardless of whether the air is rough or smooth [21]. The controller described here is similar but with a nearly opposite goal. Though Etkin states that some sort of gust field sensor array located at the wing tips and on the tail is required for the optimal implementation of such a gust alleviation controller, the gust soaring controller presented needs no such complex sensor array. Rather, the assumption was made in this research that the only sensor available for detecting the wind field was located at the aircraft's center of gravity. Tools for determining the aircraft angle of attack and side slip angle would be valuable in either case.

5.2.3.1 Improved Gust Data

The Dryden turbulence model employed here is frequently used because of its ease of use and its good approximation of realistic wind fields. The approximation is not perfect however, and in some conditions and at some frequencies, the Dryden model [43] can diverge greatly from actual measured wind field data. For this reason, it is suggested in MIL-STD-1797A [42] that data from gusts actually encountered during test flights can and should be used in simulation if available.

Gust data collected during test flights of the aircraft may be used to improve the realism of the simulation and in turn improve the aircraft controller as implemented.

5.2.3.2 Unsteady Aerodynamic Considerations

Flight testing of the aircraft through turbulence fields will also allow for the effects of unsteady aerodynamics on the aircraft and its effect on the ability of the controller to gain energy to be examined. It is expected that the inclusion of unsteady aerodynamic effects will have a negative impact on the performance of the aircraft while making rapid changes in pitch attitude; however, the magnitude of these effects on the controller must be quantified.

5.2.4 Three-Dimensional Gust Soaring Control

The controller developed here enables energy gain through the use of longitudinal control inputs only, specifically by commanding pitch rates. There is some possibility that meaningful energy gain can also be accomplished through lateral and direction control inputs as well. This would involve bank and yaw inputs in addition to the explored pitch inputs.

The addition of further degrees of freedom will make the optimization problem significantly more complex, and an improved function minimization strategy will need to be employed. Caution must also be exercised such that large deviations from the aircraft's desired course are not commanded, however the possibility of further energy gain by allowing the aircraft more freedom to maneuver in the face of turbulent wind conditions should be explored.

Vehicle Properties

Simulation results are based on the Northeast Sailplane Products Omega II 2M radio control glider. This aircraft was chosen because it is available to the author and is of similar size to the birds whose performance this work seeks to imitate. The Omega II glider is representative of modern high-performance UAV s.

Parameters in Table A.1 were obtained from a drag buildup computation executed using the Athena Vortex Lattice program developed by Mark Drela and Hal Youngeren [64]. The aerodynamic coefficients were taken for the aircraft in a trimmed flight condition and thus represent a linearized model about this point. Despite this, the aircraft drag is a nonlinear function of lift coefficient. The state limits in Table A.2 were defined to limit states to “reasonable” bounds and to preserve the accuracy of the linear aerodynamic coefficient assumption.

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

Table A.1. Parameters for Omega II 2M glider.

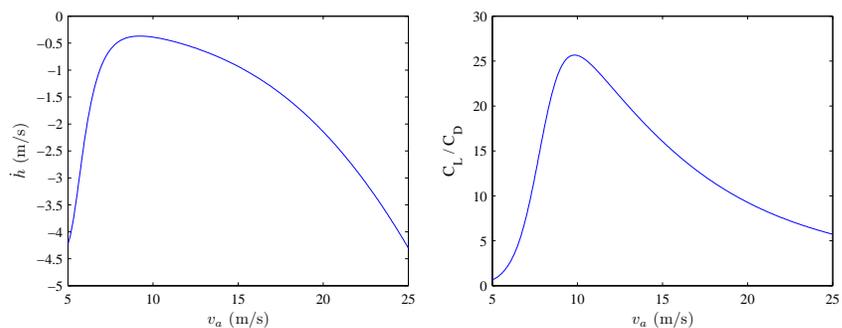
variable	value	description
m	1.31 kg	mass
b	1.99 m	span
c	0.1538 m	MAC
S	.3058 m ²	wing area
I_{yy}	.5483 kg.m ²	pitch moment of inertia
C_{L0}	0.1779	
$C_{L\alpha}$	5.1681 /rad	
C_{LQ}	-2.2189 s/rad	
$f_{LD}(\varphi)$	$0.1488\varphi^4 - 0.2624\varphi^3 + 0.1929\varphi^2 - 0.0511\varphi + 0.0228$	$\varphi = C_{L0} + C_{L\alpha}\alpha$

Some performance plots for the Omega II glider are presented in Figure A.1 to give the reader a sense of the aircraft properties. Note that airspeed for minimum sink is only slightly slower than the airspeed for best L/D. Because all performance metrics are based on energy gain/loss

Table A.2. State limits and control saturation for Omega II 2M glider.

state/control	range	description
θ	$[-60^\circ \ 60^\circ]$	pitch
v_a	$[7.5\text{m/s} \ 20\text{m/s}]$	airspeed
α	$[-5^\circ \ 15^\circ]$	angle of attack
Q	$[-\pi\text{rad/s} \ \pi\text{rad/s}]$	pitch rate

divided by the distance traveled, the airspeed for best L/D is more pertinent to this research.



(a) Sink rate versus airspeed for the Omega II 2M glider (b) L/D versus airspeed for the Omega II 2M glider

Figure A.1. Vehicle properties as modeled

Bibliography

- [1] WILSON, J. (2011) “2011 Worldwide UAV Roundup,” *Aerospace America*, **3**, p. 22.
- [2] ANDERSON, C. (2011), “Do It Yourself Drones,” Website.
URL <http://diydrones.com/>
- [3] EDWARDS, D. (2007) “Performance Testing of RNR’s SBXC Using a Piccolo Autopilot,” .
- [4] (2011), “U.S. Air Force Fact Sheet: Wasp III,” .
URL <http://www.af.mil/information/factsheets/>
- [5] (2011), “U.S. Air Force Fact Sheet: RQ-11B Raven,” .
URL <http://www.af.mil/information/factsheets/>
- [6] (2011), “U.S. Air Force Fact Sheet: Scan Eagle,” .
URL <http://www.af.mil/information/factsheets/>
- [7] PATEL, C. K. (2008) *Energy Extraction from Atmospheric Turbulence to Improve Aircraft Performance*, VDM Verlag Dr. Müller, Saarbrücken.
- [8] LAWRENCE, N. R. J. and S. SUKKARIEH (2009) “Wind Energy Based Path Planning for a Small Gliding Unmanned Aerial Vehicle,” in *AIAA Guidance, Navigation and Controls Conference*, American Institute of Aeronautics and Astronautics, Reston, Virginia.
- [9] RAYLEIGH, J. W. S. (1889) “The Sailing Flight of the Albatross,” *Nature*, **40**, p. 34.
- [10] PENNYCUICK, C. J. (2002) “Gust Soaring as a Basis for the Flight of Petrels and Albatrosses (Procellariiformes),” *Avian Science*, **2**(1), pp. 1–12.
- [11] SACHS, G. (2005) “Minimum Shear Wind Strength Required for Dynamic Soaring of Albatrosses,” *Ibis*, **147**, pp. 1–10.
- [12] BOSLOUGH, M. B. E. (2002) *Autonomous Dynamic Soaring Platform for Distributed Mobile Sensor Arrays*, Tech. Rep. SAND2002-1986, Sandia National Laboratories.
- [13] HEDENSTRÖM, A. (1993) “Migration by Soaring or Flapping Flight in Birds: The Relative Importance of Energy Cost and Speed,” *Philosophical Transactions: Biological Sciences*, **342**(1302), pp. 353–361.
URL <http://www.jstor.org/stable/55771>
- [14] KOCH, G. J. (2006) “Using a Doppler Light Detection and Ranging (Lidar) System to Characterize an Atmospheric Thermal Providing Lift for Soaring Raptors,” *Journal of Field Ornithology*, **77**(3), pp. 315–318.

- [15] WEIMERSKIRCH, H., O. CHASTEL, C. BARBRAUD, and O. TOSTAIN (2003) “Frigatebirds Ride High on Thermals,” *Nature*, **421**, pp. 333–334.
- [16] ÁKOS, Z., M. NAGY, S. LEVEN, and T. VICSEK (2010) “Thermal Soaring Flight of Birds and Unmanned Aerial Vehicles,” *Bioinspiration & Biomimetics*, **5**(4), p. 045003.
URL <http://stacks.iop.org/1748-3190/5/i=4/a=045003>
- [17] WEIMERSKIRCH, H. and G. ROBERTSON (1994) “Satellite Tracking of Light-Mantled Sooty Albatrosses,” *Polar Biology*, **14**, pp. 123–126, 10.1007/BF00234974.
URL <http://dx.doi.org/10.1007/BF00234974>
- [18] BARNARD, P. (1986) “Windhovering Patterns of Three African Raptors in Montane Conditions,” *Ardea*, **74**(2), pp. 151–158.
- [19] IRVING, F. (1978) “The Energy Loss in Pitching Manoeuvres,” in *Proceedings of the XVI OSTIV Congress*, Organisation Scientifique et Technique Internationale du Vol à Voile.
- [20] GEDEON, J. (1974) “Dynamic Analysis of Dolphin Style Thermal Cross Country Flight,” in *Proceedings of the XIV OSTIV Congress*, Organisation Scientifique et Technique Internationale du Vol à Voile.
- [21] ETKIN, B. (2005) *Dynamics of Atmospheric Flight*, Dover Publications, Mineola, N.Y.
URL <http://www.loc.gov/catdir/enhancements/fy0625/2005048479-d.html>
- [22] (1980) *Flying Qualities of Piloted Airplanes, Military Specification MIL-F-8785C*, United States Department of Defense.
- [23] ALLEN, M. J. (2006) “Updraft Model for Development of Autonomous Soaring Uninhabited Air Vehicles,” in *44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada.
- [24] ALLEN, M. J. and V. LIN (2007) “Guidance and Control of an Autonomous Soaring Vehicle with Flight Test Results,” in *AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2007-867, American Institute of Aeronautics and Astronautics, Reno, Nevada.
URL <http://dtrs.dfrc.nasa.gov/archive/00001275/>
- [25] LANGELAAN, J. W. (2007) “Long Distance/Duration Trajectory Optimization for Small UAVs,” in *Guidance, Navigation and Control Conference*, AIAA, Hilton Head, South Carolina.
- [26] ZHAO, Y. J. (2004) “Optimal Patterns of Glider Dynamic Soaring,” *Optimal Control Applications and Methods*, **25**(2), pp. 67–89.
- [27] ZHAO, Y. J. and Y. C. QI (2004) “Minimum Fuel Powered Dynamic Soaring of Unmanned Aerial Vehicles Utilizing Wind Gradients,” *Optimal Control Applications and Methods*, **25**(5), pp. 211–233.
- [28] BARATE, R., S. DONCIEUX, and J.-A. MEYER (2006) “Design of a Bio-Inspired Controller for Dynamic Soaring in a Simulated Unmanned Aerial Vehicle,” *Bioinspiration and Biomimetics*, **1**(3), pp. 76–88.
- [29] PHILLIPS, W. H. (1975) “Propulsive Effects Due to Flight Through Turbulence,” *Journal of Aircraft*, **12**(7), pp. 624–626.
- [30] LISSAMAN, P. B. S. and C. K. PATEL (2007) “Neutral Energy Cycles for a Vehicle in Sinusoidal and Turbulent Vertical Gusts,” in *45th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2007-863, American Institute of Aeronautics and Astronautics, Reno, Nevada.

- [31] PATEL, C. K. and I. KROO (2006) “Control Law Design for Improving UAV Performance using Wind Turbulence,” in *AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2006-0231, American Institute of Aeronautics and Astronautics, Reno, Nevada.
- [32] LANGELAAN, J. W. and G. BRAMESFELD (2008) “Gust Energy Extraction for Mini- and Micro- Uninhabited Aerial Vehicles,” in *46th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2008-0223, American Institute of Aeronautics and Astronautics, Reston, Virginia.
- [33] PATEL, C. K., H.-T. LEE, and I. M. KROO (2008) “Extracting Energy from Atmospheric Turbulence,” in *XXIX OSTIV Congress, Lüsse-Berlin, Germany*.
- [34] LANGELAAN, J. W. (2008) “Biologically Inspired Flight Techniques for Small and Micro Unmanned Aerial Vehicles,” in *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2008-6511, American Institute of Aeronautics and Astronautics, Reston, Virginia.
- [35] LANGELAAN, J. W., N. ALLEY, and J. NEIDERHOEFER (2010) “Wind Field Estimation for Small Unmanned Aerial Vehicles,” in *AIAA Guidance, Navigation and Control Conference*, Toronto, Canada.
- [36] LAWRENCE, N. R. J. and S. SUKKARIEH (2010) “Simultaneous Exploration and Exploitation of a Wind Field for a Small Gliding UAV,” in *AIAA Guidance, Navigation and Control Conference*, AIAA Paper 2010-8032, American Institute of Aeronautics and Astronautics.
- [37] SCHOUWENAARS, T., J. HOW, and E. FERON (2004) “Receding Horizon Path Planning with Implicit Safety Guarantees,” in *Proceedings of the 2004 American Control Conference (ACC)*, vol. 6, Boston, MA, United States, pp. 5576–5581.
- [38] SCHOUWENAARS, T. (2006) *Safe Trajectory Planning of Autonomous Vehicles*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- [39] LAWRENCE, N. R. J. and S. SUKKARIEH (2011) “Path Planning for Autonomous Soaring Flight in Dynamic Wind Fields,” in *IEEE International Conference on Robotics and Automation*, IEEE, Shanghai International Conference Center, Shanghai, China.
- [40] STEVENS, B. L. and F. L. LEWIS (2003) *Aircraft control and simulation*, 2nd ed ed., J. Wiley, Hoboken, N.J.
URL <http://www.loc.gov/catdir/bios/wiley046/2003043250.html>
- [41] CHAKRABARTY, A. and J. W. LANGELAAN (2009) “Energy Maps for Long-Range Path Planning for Small- and Micro- UAVs,” in *AIAA Guidance, Navigation, and Control Conference*, AIAA.
- [42] (1997) *Flying Qualities of Piloted Aircraft, Military Handbook MIL-HDBK-1797A*, United States Department of Defense.
- [43] HOBLIT, F. M. (1988) *Gust Loads on Aircraft: Concepts and Applications*, AIAA Education Series, American Institute of Aeronautics and Astronautics, Washington, D.C.
- [44] KENNETH DE JONG, D. B. F. and H.-P. SCHWEFEL (1997) “A History of Evolutionary Computation,” in *Handbook of Evolutionary Computation* (Z. M. Thomas Bäck, David B. Fogel, ed.), chap. A2.3, IOP Publishing Ltd. and Oxford University Press.
- [45] FOGEL, D. B. (1997) “Why Evolutionary Computation?: Introduction,” in *Handbook of Evolutionary Computation* (Z. M. Thomas Bäck, David B. Fogel, ed.), chap. A1.1, IOP Publishing Ltd. and Oxford University Press.

- [46] ATMAR, W. (1994) “Notes on the Simulation of Evolution,” *IEEE Transactions on Neural Networks*, **5**(1), pp. 130–147.
- [47] BÄCK, T., D. B. FOGEL, and Z. MICHALEWICS (1997) *Handbook of Evolutionary Computation*, Institute of Physics Pub., Bristol.
URL <http://www.loc.gov/catdir/enhancements/fy0638/97004461-d.html>
- [48] BEASLEY, D. (1997) “Possible Applications of Evolutionary Computation,” in *Handbook of Evolutionary Computation* (Z. M. Thomas Bäck, David B. Fogel, ed.), chap. A1.2, IOP Publishing Ltd. and Oxford University Press.
- [49] FLEMING, P. J. and R. C. PURSHOUSE (2002) “Evolutionary Algorithms in Control Systems Engineering: A Survey,” *Control Engineering Practice*, **10**(11), pp. 1223–1241.
- [50] WITT, J. and M. DUNBABIN (2008) “Go with the Flow: Optimal AUV Path Planning in Coastal Environments,” in *Australasian Conference on Robotics and Automation 2008*.
- [51] FOGEL, D. and L. FOGEL (1990) “Optimal Routing of Multiple Autonomous Underwater Vehicles Through Evolutionary Programming,” in *Autonomous Underwater Vehicle Technology, 1990. AUV '90., Proceedings of the (1990) Symposium on*, pp. 44–47.
- [52] SUGIHARA, K. and J. SMITH (1997) “Genetic Algorithms for Adaptive Motion Planning of an Autonomous Mobile Robot,” in *Computational Intelligence in Robotics and Automation, 1997. CIRA '97., Proceedings., 1997 IEEE International Symposium on*, pp. 138–143.
- [53] NIKOLOS, I., K. VALAVANIS, N. TSOURVELOUDIS, and A. KOSTARAS (2003) “Evolutionary Algorithm Based Offline/Online Path Planner for UAV Navigation,” *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, **33**(6), pp. 898–912.
- [54] RATHBUN, D. and B. CAPOZZI (2002) “Evolutionary Approaches to Path Planning Through Uncertain Environments,” in *AIAA Unmanned Unlimited Conference*.
- [55] RATHBURN, D., S. KRAGELUND, and A. PONGPUNWATTANA (2002) “An Evolution Based Path Planning Algorithm for Autonomous Motion of a UAV Through Uncertain Environments,” in *Algebra Universalis*, pp. 551–607.
- [56] MITTAL, S. and K. DEB (2007) “Three-Dimensional Offline Path Planning for UAVs Using Multiobjective Evolutionary Algorithms,” in *Evolutionary Computation, 2007. CEC 2007. IEEE Congress on*, pp. 3195–3202.
- [57] HANSEN, N. (2006) “The CMA-Evolution Strategy: A Comparing Review,” *Studies in Fuzziness and Soft Computing*, **192**, pp. 75–102.
- [58] ——— (2009), “The CMA Evolution Strategy: A Tutorial,” Alogrithm Creator’s Website, <http://www.lri.fr/hansen/cmatutorial.pdf>.
- [59] DAMP, L. and L. GONZALEZ “Optimisation of the Nose of a Hypersonic Vehicle Using DSMC Simulation and Evolutionary Optimisation,” in *5th AIAA ASSC Space Conference*, Melbourne, Australia.
- [60] LUTZ, T. and S. WAGNER (1998) “Drag Reduction and Shape Optimization of Airship Bodies,” *Journal of Aircraft*, **35**(3), pp. 345–351.
- [61] SONODA, T., Y. YAMAGUCHI, T. ARIMA, M. OLHOFFER, B. SENDHOFF, and H.-A. SCHREIBER (2004) “Advanced High Turning Compressor Airfoils for Low Reynolds Number Condition—Part I: Design and Optimization,” *Journal of Turbomachinery*, **126**(3), pp. 350–359.

- [62] NAUJOKS, B., L. WILLMES, W. HAASE, T. BÄCK, and M. SCHÜTZ (2000) “Multi-Point Airfoil Optimization Using Evolution Strategies,” in *European Congress on Computational Methods in Applied Sciences and Engineering*, Barcelona, Spain.
- [63] THIERENS, D., D. E. GOLDBERG, and A. G. PEREIRA (1998) “Domino Convergence, Drift, and the Temporal-Saliency Structure of Problems,” .
- [64] DRELA, M. and H. YOUNGREN (2008), “Athena Vortex Lattice (AVL ver. 3.27),” [Computer Program].